



universität
wien

DISSERTATION / DOCTORAL THESIS

Titel der Dissertation / Title of the Doctoral Thesis

Quantum Information Processing:
Thermodynamics, Complexity, and Multi-Time Phenomena

verfasst von / submitted by

Philip Taranto

angestrebter akademischer Grad / in partial fulfilment of the requirements for the degree of
Doktor der Naturwissenschaften (Dr.rer.nat.)

Wien, 2022 / Vienna, 2022

Studienkennzahl lt. Studienblatt /
degree programme code as it appears on the student
record sheet:

UA 796 605 411

Dissertationsgebiet lt. Studienblatt /
field of study as it appears on the student record sheet:

Physik

Betreut von / Supervisor:

Associate Prof. Mag.rer.nat. Dr.rer.nat. Marcus Huber

*“I am for the birds, not for the cages
in which people sometimes place them.”*

— John Cage, *Pour Les Oiseaux* [1]

Abstract

The dialectic relationship between physics and information is perhaps best exemplified through thermodynamics: A theory that connects our knowledge of the world to our capability to control and thus manipulate it. Active control over physical processes plays a crucial role regarding our ability to implement desired transformations in practice and therefore should be incorporated to define a meaningful notion of complexity. Such complexity often manifests itself in terms of complicated physical behaviour: Intricate multi-partite correlation structures, difficult-to-model evolution, and layered multi-time phenomena. Generally speaking, what one deems to be a difficult task—either from a physical or information-theoretic standpoint—is largely dictated by the required degree of spatio-temporal control, i.e., control over both multiple degrees of freedom as well as memory effects on different timescales. This cumulative thesis aims to provide a holistic picture regarding the intricate interplay between thermodynamics, complexity, and multi-time phenomena, and lay out the ensuing implications for our ability to control and process quantum information. A core thread running throughout is the following question: *What is a complex (quantum) system or process, and how can we describe and exploit control complexity and/or memory effects as a resource for (quantum) information processing?*

We first consider the task of cooling a physical system—the paradigm for a difficult task from a thermodynamic perspective (due to the Third Law of Thermodynamics)—and analyse how the degree of control over the cooling machines impacts performance objectives. We demonstrate that, given *arbitrary* control over a system and a thermal machine, the apparent contradiction between two fundamental statements of thermodynamics—namely Landauer’s famous erasure principle (and protocol) and Nernst’s “unattainability” formulation of the Third Law of Thermodynamics—can be resolved within a unified framework. We show that in order to saturate the Landauer limit for the energy cost of preparing a perfect pure state (corresponding to the lowest temperature) in finite time, one necessarily requires an unbounded level of fine-grained and complex control over spatial degrees of freedom. This result establishes control complexity as a resource that must be accounted for in order to develop a meaningful theory of thermodynamics, in line with the spirit of Nernst’s law. We further demonstrate how three key resources—namely energy, time, and control complexity—can be traded-off amongst one another: For instance, we present an efficient cooling protocol that exploits control over memory effects to dramatically reduce the temperature of the system after a finite cooling time (at finite energy cost). The level of control required in such optimal cooling protocols is reminiscent of Maxwell’s demon, which is in some sense at odds with the true spirit of thermodynamics, i.e., as a theory of *minimal* information and control requirements. Thus, we subsequently develop a formalism

to treat the task of cooling with *solely* thermodynamic resources: Machine states that begin thermal *and* exclusively interactions that can be driven by a heat engine. This setting therefore does not allow for the precise control and work source assumed by Maxwell’s demon and thus embodies the assumption of minimal control over the transformation itself at the very outset. We then derive the ultimate bound for the energy cost of cooling in this fully thermodynamic setting, which we dub the *Carnot-Landauer limit*.

Moving away from the notion of complexity in terms of the difficulty of achieving thermodynamic tasks, we subsequently drop the thermodynamic assumptions and consider the more general setting of open quantum processes, asking similar questions regarding the most general constraints the laws of physics place upon our ability to process information. In the temporal setting, control complexity often manifests itself in the form of memory effects, i.e., *correlations in time*. On the one hand, control over memory effects can lead to significant performance enhancement for many tasks of interest (e.g., as we have seen for cooling) and the simulation of exotic phenomena; on the other, lack of control over the memory leads to correlated noise which makes processes with memory difficult to characterise, simulate, and predict. Starting from a description of the most general type of open quantum dynamics (with memory), we ask the question: *Which physical traits (or processes) are fundamentally quantum and what resources are required to observe such non-classical behaviour?* As we show, the answer depends highly upon whether or not the underlying process is complex (i.e., with memory) or simple (i.e., memoryless). In particular, we connect an operational notion of classicality that is sensible for general processes and probing instruments in terms of measurement non-invasiveness (which can be decided on the basis of observed statistics alone) with various structural notions that are often related to non-classicality, such as coherence, discord, (non-)commutativity, and entanglement. Lastly, we present an entirely novel genuinely quantum multi-time phenomenon: The existence of Markovian (i.e., memoryless) statistics that nonetheless *fundamentally require* memory in the underlying quantum process to be faithfully reproduced; a phenomenon that can be witnessed by the activation of *hidden quantum memory*. Thus, all in all, we provide a well-rounded notion of complexity in the quantum realm: For one thing, by demonstrating the impact of control upon the connection between information and thermodynamics; for another, by laying out the interplay between structural and operational notions of classicality and its dependence on the underlying memory complexity of the process at hand.

We conclude with a brief summary that contextualises our results within the broader research landscape concerning the foundations of both thermodynamics and quantum theory, open quantum dynamics, and optimal control. Finally, we present a number of open challenges that must be overcome if we are to understand, manipulate, and control the complexity that quantum mechanics both affords and challenges us with.

Keywords: Information; Control; Complexity; Memory; Quantum Theory; Thermodynamics; Cooling / Erasure; Open System Dynamics; Temporal Correlations; Classicality.

Kurzfassung

Das dialektische Verhältnis von Physik und Information lässt sich wohl am besten am Beispiel der Thermodynamik verdeutlichen. Diese Theorie verknüpft unser Wissen über die Welt mit unserer Fähigkeit, sie zu kontrollieren und zu beeinflussen. Aktive Kontrolle über physikalische Prozesse spielt eine wichtige Rolle bei der Implementierung gewünschter Transformationen und sollte dementsprechend ein elementarer Baustein für eine sinnvolle Definition des Begriffs der Komplexität sein. Komplexität manifestiert sich häufig anhand von kompliziertem physikalischem Verhalten, wie etwa Korrelationsstrukturen zwischen mehreren Systemen, schwer zu modellierenden Zeitentwicklungen und vielschichtigen, multitemporalen Phänomenen. Was als schwierige Aufgabe angesehen wird—sei es vom physikalischen oder informationstheoretischen Standpunkt her—wird im Allgemeinen weitgehend durch den dafür notwendigen Grad an räumlicher und zeitlicher Kontrolle—sowohl über viele Freiheitsgrade als auch über temporale Erinnerungseffekte—diktiert. Die vorliegende kumulative Dissertation zielt darauf ab, ein ganzheitliches Bild des vielschichtigen Zusammenspiels von Thermodynamik, Komplexität, und multitemporalen Phänomenen zu zeichnen und die sich daraus ergebenden Folgen für unsere Fähigkeit Quanteninformation zu kontrollieren und verarbeiten darzulegen. Der rote Faden, der sich durch die gesamte Arbeit zieht, ist die folgende Frage: *Was ist ein komplexes (Quanten-)System oder ein komplexer (Quanten-)Prozess und wie können Kontrollkomplexität und/oder temporale Erinnerungseffekte als Ressource für die (Quanten-)Informationsverarbeitung beschrieben und genutzt werden?*

Zu diesem Zweck betrachten wir zunächst das Problem der Kühlung physikalischer Systeme, ein aus thermodynamischer Sicht paradigmatisches Beispiel einer anspruchsvollen Aufgabe (auf Grund des dritten Hauptsatzes der Thermodynamik), und analysieren, wie sich der Grad der Kontrolle über die verwendeten Kühlmaschinen auf die Leistungsziele auswirkt. Wir zeigen, dass, bei *unbegrenzter* Kontrolle über ein System und eine zusätzliche thermische Maschine, der scheinbare Widerspruch zwischen zwei grundlegenden Aussagen der Thermodynamik—nämlich das berühmte Landauer-Prinzip und Nernsts „Unerreichbarkeitsformulierung“ des dritten Hauptsatzes der Thermodynamik—aufgelöst werden kann. Wir legen dar, dass man zum Erreichen der Landauer-Grenze der Energiekosten für die Herstellung eines perfekten reinen Zustands in endlicher Zeit notwendigerweise ein unbegrenztes Maß an präziser und komplexer Kontrolle über räumliche Freiheitsgrade benötigt. Dieses Resultat etabliert Kontrollkomplexität als Ressource, die, um eine sinnvolle Theorie der Thermodynamik zu entwickeln—ganz im Sinne Nernsts—notwendigerweise mit in die Betrachtung einbezogen werden muss. Darüber hinaus zeigen wir wie drei Schlüsselressourcen—nämlich Energie, Zeit und Kontrollkomplexität—gegenseitig aus-

getauscht werden können: Zum Beispiel präsentieren wir einen effizienten Kühlalgorithmus, der Kontrolle über temporale Erinnerungseffekte ausnutzt um die Temperatur eines Systems nach endlicher Kühlzeit (und unter endlichem Energieaufwand) dramatisch reduziert. Ein solches Maß an Kontrolle erinnert an Maxwells Dämon, der in gewisser Weise im Widerspruch zu den Grundgedanken der Thermodynamik als einer Theorie steht, die minimale Annahmen über Informations- und Kontrollbedarf macht. Dementsprechend entwickeln wir einen Formalismus, der es erlaubt, Kühlung unter Annahme von ausschließlich thermodynamischen Ressourcen zu behandeln, d.h. Maschinen, die sich anfänglich in einem thermischen Zustand befinden und Wechselwirkungen, die ausschließlich von Wärmekraftmaschinen angetrieben werden. Diese Einschränkungen verhindern die präzise Kontrolle, die Maxwells Dämon benötigen würde, und inkorporieren daher von Beginn an die Annahme minimaler Kontrolle über die verwendeten Transformationen. Im Anschluss leiten wir eine allgemeine untere Schranke für die Energiekosten eines Kühlprozesses, welche wir als *Carnot-Landauer-Grenze* bezeichnen, unter diesen vollständig thermodynamischen Annahmen ab.

Anschließend lassen wir die thermodynamischen Annahmen fallen und betrachten die allgemeinere Situation offener Quantendynamik, indem wir Fragen über die Beschränkungen stellen, die die Gesetze der Physik unserer Fähigkeit zur Informationsverarbeitung auferlegen. Unter diesen Paradigma manifestiert sich Komplexität häufig in der Form temporaler Erinnerungseffekte, d.h. *Korrelationen in der Zeit*. Auf der einen Seite kann Kontrolle über Erinnerungseffekte zu erheblichen Leistungssteigerungen bei vielen praktischen Aufgaben (wie wir zum Beispiel anhand von Kühlprozessen gesehen haben) sowie der Simulation exotischer Phänomene führen. Auf der anderen Seite führt ein Mangel an Kontrolle über die Erinnerungseffekte zu korrelierten Fehlern, was die Charakterisierung, Simulation und Vorhersage von Prozessen mit temporalen Erinnerungseffekten erschwert. Ausgehend von einer allgemeinen Beschreibung offener Quantendynamik (inklusive Erinnerungseffekten) stellen wir die Frage: *Welche Eigenschaften (oder Prozesse) sind grundlegend quantenmechanisch und welche Ressourcen sind erforderlich, damit sich ein solches nichtklassisches Verhalten manifestiert?* Wie wir zeigen, hängt die Antwort dieser Frage stark davon ab ob der zugrunde liegende Prozess komplex (d.h. inklusive Erinnerungseffekte) oder einfach (d.h. ohne Erinnerungseffekte) ist. Insbesondere verbinden wir die operationelle Formulierung klassischer Eigenschaften im Sinne störungsfreier Messungen, ein Kriterium, welches allein auf Grundlage von Beobachtungsstatistiken entschieden werden kann, mit verschiedenen strukturellen Begriffen wie Kohärenz, Discord, (Nicht-)Kommutativität und Verschränkung. Zuletzt stellen wir ein völlig neuartiges, genuines Quanten-Multitemporalphanomen vor: Die Existenz markowscher (d.h. „erinnerungsloser“) Statistik, für deren getreue Erzeugung die zugrunde liegenden Quantenprozesse über temporale Erinnerung verfügen müssen, ein Phänomen, das sich durch die Aktivierung von *verborgenen Quantenerinnerungseffekten* nachweisen lässt. Insgesamt stellen wir ein ausgewogenes Verständnis von Komplexität im Quantenregime bereit: Auf der einen Seite, indem wir den Einfluss von Kontrolle auf die Verbindung zwischen Information und Thermodynamik aufzeigen; auf der anderen Seite, indem wir das Zusammenspiel von strukturellen und operationalen Begriffen von „Klassikalität“, sowie dessen Abhängigkeit von der temporalen Erinnerungskomplexität des zugrunde liegenden Prozesses darstellen.

Anhand einer kurzen Zusammenfassung betten wir unsere Resultate in den größeren Forschungszusammenhang ein, sowohl im Bezug auf die Grundlagen der Quantentheorie und Thermodynamik, offene Quantendynamik, als auch die Theorie optimaler Kontrolle. Zum Abschluss präsentieren wir einige offene Herausforderungen, die überwunden werden müssen um die Komplexität, die Quantenmechanik erlaubt—und mit der sie uns konfrontiert—zu verstehen, zu kontrollieren und zu beeinflussen.

Acknowledgements

The time I spent in Vienna over the past three years has filled me with amazing experiences that go far beyond the work demonstrated explicitly throughout this dissertation, having been incredibly enriching from both a personal as well as scientific standpoint. None of this would have been possible without the support, collaboration, guidance, and kindness of several brilliant people that I had the great fortune of meeting along the journey.

First of all, I would like to express a sincere gratitude towards my supervisor, Marcus Huber. For giving me the opportunity to work on a variety of interesting topics; for helping me feel welcome and at home from day one, some 15,000 km away from where I spent my life before moving to Vienna; for allowing me the freedom to pursue whatever I felt compelling and in whatever way worked for me, all the while knowing that you would have my back if things were to go awry—your continued support means a great deal to me, and for this I am truly grateful.

With the same gravity, I would like to thank my good friend, close collaborator and de-facto supervisor, Simon Milz. Your contagious enthusiasm for life; your ability to create fun and sociable environments in any circumstances; your sense of responsibility, care, and support; and finally your eclectic collection of fun facts, deep love for rom-coms, and exquisitely misaligned Birds on Shirts—I have nothing but the deepest admiration for. I'm out of words, so I'll have to repeat a sentence that I wrote in the acknowledgements for my Masters thesis: It's been a wild ride so far and now we set sail for tomorrow!

Next, I am grateful to the following set of people with whom I made and kept close connections whilst in Vienna. To Jessica Bavaresco, for teaching me how to knit (and a thing or two about physics); for all of the culinary delights; and for the great gossip of course. To Felix Binder, for being a caring and diligent mentor; for always embodying the perfect balance between enthusiasm and critique; and for the weekend coffee adventures. To Marco Túlio Quintino, for always providing the perfect soundtrack; for your (complete) positivity in most situations; and for generally being up for all sorts of shenanigans—"Voilà! Mon Dieu!". To Josh Morris, for keeping me level-headed; for helping me out as a useless flatmate with two broken elbows; and for moderating the vibes whenever necessary. And finally to Kavan Modi, for continuously reinforcing that *the problem is your friend*; for being a great role model; and for always contributing to the biggest day of the calendar year.

To all of the members (and honorary members) of the QUIT group, thank you for creating such a vibrant and stimulating research environment. A special mention to Nicolai Friis for helping me translate my horrifically complicated sentences into German (along with Simon Milz) and always providing a healthy smattering of dad jokes mixed perfectly with an unhealthy level of scepticism. Also many thanks to the YIRG group for letting me be an honorary member from time to time! Thanks again to Marcus, Nico, Simon, Felix, and Jessica, for thoughtful comments on an early draft of this thesis. And to my other collaborators with whom I worked on the papers included in this dissertation. Extra shout-outs to Xiaoqin Gao, Cody Duncan, Francesco Campaioli, Johanna Schäfer, Katarina Baumann, Judith Milz, Top Notoh, and Dacre Smith [2]:

*The willows leaping into leaf
Herald the spring.*

Many of you would have heard my ramblings about open dynamics and how the capabilities of a system are limited by its environment, so thank you all for driving me in interesting directions!

Thanks to the administrative staff at both IQOOI and TU Wien, in particular Lukas Edengruber, Benjamin Skarabela, and Evgeniya Schimpf, for helping me wade through bureaucracy.

From a financial perspective, I would like to thank the Austrian Science Fund (FWF), for funding my PhD through the START project Y879-N27.

Thanks to the local friends I've made in Vienna: In particular, to Katrin Kreiner, for teaching me so many creative techniques and being a constant rock of support; to Sjoerd Sloetjes and Heidi Pein for all of the adventures in der Steiermark; to Tanja Peinsipp for the dancing; and to Lena Kirisits, Martin Gajdos, and Lioba Kuhn for creating a great flat together and putting up with me!

Although there may have been a great deal of physical space between us of late, there will always be a special place in my heart for my loved ones back home in Melbourne. A warm appreciation goes out to Patrick Orlando, Vanessa La Delfa, John Farrugia, Anthony Anastasio, Anthony Moumar, Sam Barone, Jessica and Julia Caissutti, Sam Meulenberg, Nicole Runchey, Pete Srdic, Gorge Bajada, Tony Isaacs, Joe Mayberry, Jenna Van Der Velden, Ashley Bransgrove, Dave MacIndoe, Bradley Hunt, Jack Forbes, and Maddy Muscat. I wouldn't be who, nor where, I am today if it weren't for your constant support and reminder of the lightness of being. To Alanah Guiry: Thanks for the love, magic, and chaos, Bookie! And for teaching me how to play. Keep making waves...

The remainder of my love and appreciation goes out to my family: This journey would have never even started if it were not for your love and dedication, and for that I am forever grateful.

Vienna, 2 May, 2022.

List of Publications

The indicated publications feature throughout this thesis.

Peer Reviewed

- [P1] **Non-Markovian Memory Strength bounds Quantum Process Recoverability**
Philip Taranto, Felix A. Pollock, and Kavan Modi
[npj Quantum Inf. 7, 149 \(2021\)](#)
[arXiv:1907.12583](#)
- [P2] **Experimental Demonstration of Instrument-Specific Quantum Memory Effects and Non-Markovian Process Recovery for Common-Cause Processes**
Yu Guo*, Philip Taranto*, Bi-Heng Liu, Xiao-Min Hu, Yun-Feng Huang, Chuan-Feng Li, and Guang-Can Guo
[Phys. Rev. Lett. 126, 230401 \(2021\)](#)
[arXiv:2003.14045](#)
* Denotes equal contribution
- [P3] **Exponential Improvement for Quantum Cooling through Finite-Memory Effects**
Philip Taranto, Faraj Bakhshinezhad, Philipp Schüttelkopf, Fabien Clivaz, and Marcus Huber
[Phys. Rev. Appl. 14, 054005 \(2020\)](#)
[arXiv:2004.00323](#)
(Chapter 2)
- [P4] **When is a Non-Markovian Quantum Process Classical?**
Simon Milz, Dario Egloff, Philip Taranto, Thomas Theurer, Martin B. Plenio, Andrea Smirne, and Susana F. Huelga
[Phys. Rev. X 10, 041049 \(2020\)](#)
[arXiv:1907.05807](#)
(Chapter 3)
- [P5] **Memory Effects in Quantum Processes**
Philip Taranto
[Int. J. Quantum Inf. 18, 1941002 \(2020\)](#)
[arXiv:1909.05245](#)

- [P6] **The Structure of Quantum Stochastic Processes with Finite Markov Order**
Philip Taranto, Simon Milz, Felix A. Pollock, and Kavan Modi
[Phys. Rev. A **99**, 042108 \(2019\)](#)
[arXiv:1810.10809](#)
- [P7] **Quantum Markov Order**
Philip Taranto, Felix A. Pollock, Simon Milz, Marco Tomamichel, and Kavan Modi
[Phys. Rev. Lett. **122**, 140401 \(2019\)](#)
[arXiv:1805.11341](#)
- [P8] **Emergence of a Fluctuation Relation for Heat in Nonequilibrium Landauer Processes**
Philip Taranto, Kavan Modi, and Felix A. Pollock
[Phys. Rev. E **97**, 052111 \(2018\)](#)
[arXiv:1510.08219](#)

Under Review

- [P9] **Connecting Commutativity and Classicality for Multi-Time Quantum Processes**
Fattah Sakuldee, Philip Taranto, and Simon Milz
[arXiv:2204.11698 \(2022\)](#)
(Chapter 4)
- [P10] **Hidden Quantum Memory: Is Memory There When Somebody Looks?**
Philip Taranto and Simon Milz
[arXiv:2204.08298 \(2022\)](#)
(Chapter 5)
- [P11] **Landauer vs. Nernst: What is the True Cost of Cooling a Quantum System?**
Philip Taranto^{*}, Faraj Bakhshinezhad^{*}, Andreas Bluhm[†], Ralph Silva[†], Nicolai Friis, Maximilian P. E. Lock, Giuseppe Vitagliano, Felix C. Binder, Tiago Debarba, Emanuel Schwarzhans, Fabien Clivaz, and Marcus Huber
[arXiv:2106.05151 \(2021\)](#)
^{*},[†] Denotes equal contribution
(Chapter 1)

Contents

List of Figures	xx
List of Tables	xxii
Preamble	1
Part I	3
Part II	12
Acronyms	25
I Cooling Physical Systems	27
1 Landauer vs. Nernst: What is the True Cost of Cooling a Quantum System?	29
1.1 Introduction	31
1.2 Overview & Summary of Results	32
1.2.1 Perfect Cooling with Coherent Control (Work Source)	34
Energy, Time, and Control Complexity as Resources	34
Role of Control Complexity	35
Imperfect Cooling with Finite Resources	37
1.2.2 Perfect Cooling with Incoherent Control (Heat Engine)	38
1.2.3 The Carnot-Landauer Limit	39
1.3 Main Results	40
1.3.1 Framework	41
1.3.2 Coherent Control	42
1.3.3 Notions of Complexity	45
1.3.4 Incoherent Control	49
1.3.5 Imperfect Cooling with Finite Resources	51
1.4 Discussion	53
1.4.1 Relation to Previous Works	53
1.4.2 Conclusions & Outlook	54

2 Exponential Improvement for Quantum Cooling through Finite-Memory Effects	57
2.1 Introduction	59
2.2 Task: Cooling a Quantum System	60
2.3 Framework: Collision Models with Memory	60
2.4 Memory-Enhanced Cooling	62
2.4.1 Asymptotic Cooling Advantage	63
2.4.2 Step-Wise Optimal Protocol	64
2.5 Relation to Heat-Bath Algorithmic Cooling	65
2.6 Conclusions	66
II Open Quantum Processes & Multi-Time Phenomena	69
3 When is a Non-Markovian Quantum Process Classical?	71
3.1 Introduction	73
3.2 Summary of the Main Results	75
3.3 General Framework	78
3.3.1 Kolmogorov Consistency Conditions and Classicality	78
3.3.2 Measurement Setup	81
3.3.3 Open (Quantum) System Dynamics and Memory Effects	82
3.4 Coherence and Classicality	83
3.4.1 One-to-One Connection in the Markovian Case	84
3.4.2 Coherence in the Non-Markovian Case: Preliminary Analysis	90
3.5 Non-Markovian Classical Processes	92
3.5.1 Classicality and Processes with Memory	92
3.5.2 Choi-Jamiołkowski Representation of General Quantum Processes	96
3.5.3 Structural Properties of Classical Combs	97
3.5.4 Quantifying Non-Classicality	100
3.6 Dynamical Properties of Classical Quantum Processes	103
3.6.1 Discord and Classicality	103
3.6.2 Non-Discord-Generating-and-Detecting (NDGD) Dynamics and Classical Processes	107
3.7 Genuinely Quantum Processes	111
3.8 Conclusions and Outlook	116
3.8.1 Conclusions	116
3.8.2 Outlook	118

4	Connecting Commutativity and Classicality for Multi-Time Quantum Processes	123
4.1	Introduction	125
4.2	Preliminaries	127
4.2.1	Lüders' Theorem: Commutativity and Non-Invasiveness	127
4.2.2	Kolmogorov Consistency and Non-Invasiveness	131
4.2.3	Multi-Time Statistics from (Markovian) Quantum Processes	133
4.3	Multi-Time Dynamics: Kolmogorov Consistency and Commutativity	134
4.3.1	Two-Time vs. Multi-Time Classicality	135
4.3.2	No "Straightforward" Extension of Lüders' Theorem	137
4.3.3	Commutativity as a Notion of Classicality: When is Kolmogorov Consistency Sufficient for Lüders-Type Theorems?	140
4.3.4	Markovian Processes & Non-Coherence-Generating-and-Detecting (NCGD) Dynamics	143
4.4	Discussion and Conclusion	147
5	Hidden Quantum Memory: Is Memory There When Somebody Looks?	149
5.1	Introduction	151
5.2	Framework	152
5.3	Classical Processes	154
5.4	Quantum Processes	155
5.5	Hidden Quantum Memory and Incompatibility	157
5.6	Conclusions	159
	Concluding Discussion	161
	Summary	161
	Outlook	164
	Appendices	169
	Appendix A Supplemental Information for Chapter 1	171
A.1	Equality Forms of the (Carnot-)Landauer Limit	171
A.1.1	Coherent-Control Paradigm: The Landauer Limit	171
A.1.2	Incoherent-Control Paradigm: The Carnot-Landauer Limit	172
A.2	Diverging Energy	175
A.2.1	Sufficiency: Diverging Energy Cooling Protocol	175
A.2.2	Necessity of Diverging Energy for Protocols with Finite Time and Control Complexity	175
A.3	Diverging Time Cooling Protocol for Finite-Dimensional Systems	177
A.3.1	Proof of Theorem 1.2	177
A.3.2	Special Case: Equally Spaced Hamiltonians	179

A.4	Conditions for Structural and Control Complexity	184
A.4.1	Necessary Complexity Conditions	185
A.4.2	Proof of Theorem 1.3, Corollary 1.2, and Theorem A.3	188
A.4.3	Sufficient Complexity Conditions	190
A.4.4	Fine-tuned Control Conditions	191
A.5	Diverging Time / Control Complexity Cooling Protocols for Harmonic Oscillators	203
A.5.1	Preliminaries	204
A.5.2	Diverging Time Cooling Protocol for Harmonic Oscillators	206
A.5.3	Diverging Control Complexity Cooling Protocol for Harmonic Oscillators	213
A.6	Cooling Protocols in the Incoherent-Control Paradigm	215
A.6.1	Diverging Energy: Proof of Theorem 1.5	216
A.6.2	Diverging Time / Control Complexity	221
A.6.3	Saturating the Landauer Limit with an Infinite-Temperature Heat Bath	222
A.6.4	An Analysis of Finite-Temperature Heat Baths	223
A.7	Perfect Cooling at the Carnot-Landauer Limit in the Incoherent-Control Paradigm	225
A.7.1	Qubit Case	225
A.7.2	Qudit Case	232
A.8	Comparison of Cooling Paradigms and Resources for Imperfect Cooling	238
A.8.1	Rates of Resource Divergence for Linear Qubit Machine Sequence	238
A.8.2	Comparison of Coherent and Incoherent Control	239
Appendix B Supplemental Information for Chapter 2		243
B.1	Markovian Embedding of Collision Models with Memory	243
B.2	Proof of Theorem 2.1	247
B.3	Role of System-Memory Carrier Correlations	252
B.4	Step-Wise Optimal Protocol and Finite Time Comparisons	254
B.5	Relation to Heat-Bath Algorithmic Cooling (HBAC) and State-Independent Asymptotically Optimal Protocol	256
Appendix C Supplemental Information for Chapter 3		265
C.1	Connection to Previous Results	265
C.2	Absence of Coherence for a Model System: Qubit Coupled to a Continuous Degree of Freedom	267
C.3	Comb Representation of a Model System: Qubit Coupled to a Continuous Degree of Freedom	269
C.4	Alternative Example for Non-Classical Dynamics that do not Create Coherences	272
C.5	Measure of Non-Classicality	274
C.6	Non-Discord-Creating Maps	278

C.7	Proof that NDGD Dynamics \Rightarrow Classical Process	281
C.8	Classicality $\not\Rightarrow$ NDGD Dynamics	282
C.9	Example of a Genuinely Quantum Process	283
Appendix D Supplemental Information for Chapter 5		289
D.1	Classical Processes	289
D.1.1	Memoryless Classical Processes and Markovian Statistics	289
D.1.2	Sub-Statistics of Memoryless Classical Processes	290
D.2	Quantum Processes	292
D.2.1	Memoryless Quantum Processes and Markovian Statistics	292
D.2.2	Sub-Statistics of Memoryless Quantum Processes	293
D.3	Hidden Quantum Memory and Incompatibility	294
D.3.1	Hidden Quantum Memory	294
D.3.2	Incompatibility	298
Bibliography		301

List of Figures

1.1	Thermodynamic Framework	33
1.2	Notions of Complexity	45
1.3	Imperfect Cooling Comparison	52
2.1	Collision Model with Memory	62
2.2	Cooling Behaviour	64
3.1	Probing a Process with Projective Measurements	82
3.2	General Open Quantum Process	83
3.3	Markovian Process	85
3.4	NCGD Dynamics	87
3.5	Comb of a General Open Quantum Evolution	94
3.6	Consistency Condition for Combs	95
3.7	Nested Set of Processes	99
3.8	NDGD System-Environment Dynamics	108
3.9	Transformation to NDGD Dilation	111
3.10	First Two Times of a Genuinely Quantum Process	113
4.1	Multi-Time Experiments & Kolmogorov Consistency	132
4.2	Multi-Time Probing of a Markovian Quantum Process	134
5.1	Probing a Process	153
5.2	Markovian Statistics that Require Memory	158
A.1	Imperfect Cooling with Coherent and Incoherent Control	241
B.1	Generalised Collision Model with Memory	244
B.2	Operational CP-Divisibility	246
B.3	Markovian Embedding of Generalised Collision Model	247
B.4	Finite-Time Correlations for Global and Local Cooling Protocols	253
B.5	Equivalence Between HBAC and Generalised Collision Model	257
B.6	Finite-Time Advantage with Adaptive Protocols	262
B.7	Finite-Time Behaviour of Step-Wise Optimal Protocol	263
C.1	Labelling of Hilbert Spaces used for the Comb Description of Example 3.1	270

C.2	Dilation for Example 3.1	272
C.3	Non-Classical Process that does not Display Coherences	273
C.4	Non-NDGD Dynamics that leads to Classical Statistics	282
C.5	Genuinely Quantum Process	284
D.1	Circuit with Hidden Quantum Memory / Incompatible Markovian Statistics	295

List of Tables

1.1	Coherent-Control Cooling Protocols for Finite-Dimensional (Qudit) and Harmonic Oscillator Systems	35
1.2	Incoherent-Control Cooling Protocols for Finite-Dimensional Systems	40

Preamble

Information is Physical [3, 4]. This famous realisation of Rolf Landauer is arguably one of the most relevant of modern times: Indeed, it has served as a precursor for the recent paradigmatic shift in the way we understand information processing in the context of our physical world.

Prior to Landauer, the foundations of the theory of information had been laid down largely in terms of an abstract, mathematical formulation, most notably due to Claude Shannon [5]. The theoretic considerations were, nonetheless, deeply rooted in questions of a high degree of practical relevance: *What is the minimal amount of information required to communicate a message? What is the maximum rate at which information can be transmitted over a noisy communication channel? How can one encrypt a message such that it is protected from potential eavesdroppers during transmission?* Early developments of information theory provided the corresponding ultimate limitations by way of deriving optimal compression schemes [5], noisy channel coding theorems [5–7], and secure cryptographic protocols [8], amongst others [9–11].

Given such optimal strategies, the concern then shifted to the possibility of their implementation in realistic scenarios. In retrospect, it is clear that the answer depends upon “the rules of the game”, or, more formally, *control over the resources* at hand: *How exactly is the message being encoded, transmitted, and decoded? What kinds of (and how many) communication channels does one have access to? How powerful can a potential eavesdropper be and what is a sufficient level of security?* In short, information is not simply some abstract entity, but rather must be considered within the context set by the laws of physics, which dictate how it can be encoded, stored, transmitted, and retrieved [12]. Although such concerns were being analysed and overcome within various engineering contexts before Landauer, his principle perhaps best exemplifies the profound and foundational connection between physics and information. What is difficult to achieve in practice determines the realistic performance of implementing any protocol; thus, physics places fundamental limitations on the processing of information.

Importantly, this relationship between physics and information is not one-sided [13]: Reasonable information-theoretic axioms can be posited to derive the laws of (classical, quantum, and even potential post-quantum) physics [14–20], and it seems logical to believe that any task leading to a contradiction with respect to such axioms (e.g., those set by the famous non-signalling Popescu-Rohrlich (PR)-boxes [21], or concerning backwards-in-time signalling [22]) should not be possible to achieve in reality.¹ In this sense, the theory of information sets a boundary upon

¹Of course, just because we have never witnessed a practical implementation of such a task does not mean that we can rigorously exclude such possibilities on a scientific basis.

the laws of physics; or, at least, what we can say about them.²

We therefore have, in the modern scientific approach, a *dialectics* of physics and information.

The mediator of this relationship revolves around the notions of *complexity* and *control*. Physical systems or processes with many intricately connected and interacting “parts” are typically viewed as complex and require a lot of information to describe. On the flip side, given the ability to discern and control such vast amounts of information allows one to simulate such complex phenomena (potentially with overheads; see, e.g., Ref. [23]). Clearly, in practice, one’s ability to extract, process, and utilise such complex information requires a high degree of control over the physical systems at hand. From an operational perspective then, complexity is not only a property of the underlying physical structures (e.g., states, processes) but also hinges upon the envisaged setup; in particular, regarding the kinds of measurements and manipulations of said objects that can be performed. In this sense, complexity is the amalgam of both paradigms, stemming from the interplay between a priori information content, physical restrictions and limitations upon the level of control, and the ensuing difficulty of achieving certain tasks.

Within this context, this dissertation addresses the following question: *What is a complex physical system or process, and what role does control over such complexity play in our ability to process (quantum) information?* Of course, such a complex³ issue cannot be fully distilled and resolved within the confines of a single thesis, and here—bound by natural physical and information-theoretic constraints—we merely “scratch the surface”; nonetheless, we hope that the reader finds the insights presented to be profoundly interesting, captivating, and illuminating.

To this end, we consider two main themes: In **Part I**, we take a paradigmatic difficult task from a thermodynamic perspective, namely the *cooling of a physical system* (or, equivalently, the *erasure of information*), and consider it from both fundamental and practical viewpoints. This analysis will emphasise the necessity of a high degree of control complexity to saturate the ultimate limitations—both regarding how cold a system can be made and at what resource cost—set by the laws of thermodynamics. In this investigation, we will see the pivotal role played by our ability to control the *environment* (or auxiliary *machines*)⁴ in order to implement optimal cooling protocols on a physical system—a crucial information processing task. In **Part II**, we focus our attention on *multi-time phenomena*—in particular, memory effects—as the manifestation of complexity in open processes involving systems interacting with structured environments. We first connect various structural notions of quantumness (or non-classicality), namely coherence, discord, and non-commutativity, to a purely operational one that can be assessed solely via observed experimental data. Leveraging these results, we uncover a novel genuinely quantum multi-time phenomenon—*hidden quantum memory*—which demonstrates a new type of complexity that does not exist in the classical realm, further exemplifying the intricate relationship between complexity, time, and control in quantum information processing.

²Indeed, the broader scientific method relies on a complex social praxis of recording data, repeating experiments, and communicating results—all of which fall within the confines of information processing.

³Pardon the pun.

⁴Strictly speaking, an *environment* is inaccessible by definition; whenever we switch perspectives to permit control over such auxiliary systems, we will use the term *machine*.

Part I

The interplay between physics and information comes to the fore when considering the “village witch” [24] of physical theories: *Thermodynamics*. Indeed, one can reasonably argue that the reason thermodynamics is so unique amongst physical theories is that it appears to be fully robust against the underlying theory, and this might be due to the fact that it is, at its base, a theory of information. As it largely concerns average behaviour, thermodynamics is arguably not a fundamental physical theory per se, but rather seemingly emerges from statistical mechanics whenever one does not have fine-grained information about and control over the system of interest. At the same time, this elevates thermodynamics above other physical theories as perhaps the truly operational one, with questions concerning resources, costs, and practical implementations—all of which depend upon the information and control that an agent assumes—built into its core.

It is within this context that Landauer made his pivotal breakthrough [3], establishing a concrete link between physics and information processing via thermodynamics. Landauer considered the task of erasing information, which is a pivotal sub-routine for various information processing protocols, such as setting up a blank initial register for a computer [12]. Starting from a pragmatic perspective, Landauer realised that in any such procedure, information must first be encoded in some physical system, meaning that one begins with a (generally noisy/mixed) state—such as a bit of information stored in the position of an atom being in one half of a box or the other, unbeknownst to the agent—and the goal is to physically transform said state to a pure one—e.g., such that the atom is in one fiducial half of the box with certainty.

At the outset, said goal can be achieved in a straightforward manner: The agent could simply swap the information-bearing system with a suitable one in a pure state. This solution, however, is at odds with our thermodynamic experience, which stipulates that one is not privy to such detailed information concerning microscopic degrees of freedom of the *environment*, i.e., everything external to the system, but rather is pragmatically assumed to have minimal information, perhaps regarding some macroscopic quantities such as average energy or temperature.

The formal setup of the problem should then demand that the environment cannot begin in a pure state, but instead must be a thermal state of finite temperature, which was historically derived as the state of maximal ignorance that is consistent with certain macroscopic observations [25–27].⁵ Under this assumption, Landauer demonstrated that erasing any information stored in the system *necessitates* a minimum amount of heat being dissipated into the environment [3]. Information processing has physical consequences [12, 28]: Landauer’s bound poses the fundamental minimal energy cost of erasing information.

Current computers erase information all the time, and in doing so necessarily heat up—but the heat dissipated here largely comes from engineering inefficiencies and is typically orders of magnitude above the Landauer limit [29, 30]. Indeed, only recently have some (extremely costly) experiments been performed that come close to saturating the Landauer bound in terms of the dissipated heat [31–34]. To do so, such experiments employ complex interactions between

⁵Indeed, temperature parametrises average energy at maximal degree of ignorance.

the system to be erased and its thermal environment: Although the initial environment state incorporates the assumption of maximal ignorance, the operational control required to achieve the fundamental thermodynamic limit requires precise control over microstate interactions.

This is often the case in implementing the optimal procedure for thermodynamic protocols. Indeed, Landauer’s limit can alternatively be formulated in terms of the *Second Law of Thermodynamics*,⁶ which can famously be saturated by Maxwell’s demon [35]; saturating the Landauer limit requires a similarly omnipotent agent [36–39]. Thus, with such optimal procedures we find ourselves at the boundary of thermodynamics, i.e., concerning the fundamental limitations of what can be achieved with (potentially highly complex) physical transformations upon some system of interest with the aid of an initially thermal environment.

Importantly, Landauer’s limit (and the finite-energy protocol to achieve it) applies to any logically irreversible transformation, e.g., partial erasure of information. From a thermodynamic perspective, the task of erasing information is a special case of *cooling a physical system*, namely bringing a (mixed) initial state closer to its ground state, or more generally, toward any pure state.⁷ Roughly 50 years prior to Landauer’s information-theoretic approach, Walther Nernst considered the task of cooling and demonstrated it to be paradigmatically difficult from a thermodynamic perspective, putting forth the *Third Law of Thermodynamics* [40, 41].

While various formulations of the Third Law abound, the one that we will find to be the most insightful throughout this dissertation is that *infinite resources are required to cool any physical system to absolute zero temperature* [42]; said formulation is also known as the *Unattainability Principle* [43]. Indeed, what is common across experimental demonstrations regarding cooling a system as much as possible (see e.g., Refs. [31–34, 44, 45]) is the enormous amount of resources necessarily invested. Thus, from a thermodynamic perspective, the task of cooling a physical system to absolute zero temperature—i.e., preparing a pure state or completely erasing information—provides an archetype for a fundamentally difficult goal to achieve.

Although both Landauer and Nernst considered this problem in the classical realm, the pervasiveness of thermodynamics applies equally well in the quantum setting, and the core concepts, limitations, and difficulties (and then some) carry over, forming the basis for the emerging field of *quantum thermodynamics* [46]. However, at first glance, their two statements appear at odds with one another: On the one hand, Nernst states that infinite resources are required to prepare a pure state; on the other hand, Landauer explicitly demonstrates an erasure procedure that does just that with finite (indeed, minimal) energy cost. Nonetheless, they cannot be simply “two sides of the same coin”, as that would imply that the Second and Third Laws of Thermodynamics are one and the same. By scrutinising the physical implementation of any perfect cooling procedure, we will come to understand that there is no true contradiction here, but rather that there exist various meaningful resources that can be utilised in different ways to the same end, and that furthermore one such resource must diverge in order to (asymptotically)

⁶The Second Law states that for any system that interacts with some (initially uncorrelated) environment, the sum of the local entropy changes must be non-negative; Landauer’s bound on the dissipated heat arises as a direct consequence of assuming the environment begins in a thermal state at some fixed temperature.

⁷Strictly speaking, there is a finite additional energy cost dictated by the local Hamiltonian of the system with respect to information erasure that must here be considered; however, in practice this is negligible compared to the heat dissipated into the environment, and does not impact any fundamental statements.

attain a pure state (in the spirit of Nernst’s Law): For instance, Landauer’s erasure protocol minimises energy but at the expense of taking infinite time.

Understanding the ultimate limitations set by thermodynamics on the task of cooling under various levels of assumed control, as well as outlining how such fundamental bounds can be saturated, and moreover demonstrating the possible trade-offs between energy, time, control complexity, and memory effects as resources, is the core focus of **Part I**. We begin in **Chapter 1** by developing a cohesive and thermodynamically consistent framework to approach the problem in full generality, before moving to show how control over memory effects can be exploited to drastically enhance cooling performance in **Chapter 2**.

Chapter 1 focuses on the ultimate limitations of cooling and the relationship between the (seemingly contradictory) perspectives put forth by Landauer and Nernst regarding their attainability. Importantly, although computation can be performed using only logically reversible gates in principle [47–49], it also demands the initialisation of the input registers in pure states (as well as resetting states to such “blank” states throughout the computational process for greater efficiency), thus ultimately also requiring logical irreversibility in the form of information erasure and dealing with the ensuing heat that is necessarily dissipated. Heat is, of course, detrimental to reliable information processing as it introduces spurious noise: Such noise can be mitigated in the classical setting where states are stable and error correction schemes scalable and robust, but poses a fundamental issue in the quantum realm, leading to unmanageable errors and short coherence times and rendering any possible “quantum advantages” all but a pipe dream [50, 51]. We must therefore understand the practical limitations of efficiently resetting quantum registers, i.e., the conditions for saturating the Landauer bound and, more importantly, exactly how to do so and at the expense of which resources.

In this chapter, we address these questions and, in doing so, are unavoidably confronted with Nernst’s Third Law of Thermodynamics [40, 41], which states that an infinite amount of resources is required to cool a system to absolute zero temperature [42, 43]. The precise nature of these resources, and potential trade-offs between them, however, were previously left in the dark, leaving one to grasp for isolated examples. Landauer’s erasure protocol, for instance, minimises the energy cost and can be mimicked by a sequence of individually simple interactions, but must be implemented quasi-statically and therefore comes at the expense of diverging time; in this sense, *time* can be seen as a resource that diverges in this protocol to attain a pure state, in line with Nernst’s Third Law. Alternatively, by investing a diverging amount of *energy*, one can create a pure state in a single operation (i.e., minimal time) by swapping the target system with a sufficiently pure (sub)space of an auxiliary thermal machine. Does this mean that one of the two—energy or time—must diverge in order to create a pure state?

As we demonstrate, the answer is, perhaps surprisingly: *No*. For any quantum system, we show how one can prepare a pure state with minimal energy (i.e., Landauer cost) and in finite time. Upon closer inspection, we identify the hidden resource behind this seemingly paradoxical statement: *Control Complexity*. Loosely speaking, control complexity can be quantified in terms of both the number of and manner in which an agent addresses physical degrees of

freedom throughout a protocol. Specifically, we show that with an unconstrained level of control (i.e., diverging control complexity), both finite energy and finite time suffice for perfect cooling; whereas, in the spirit of Nernst’s unattainability principle, finite control complexity remains insufficient. In addition to providing this fundamental conceptual insight, our proofs are constructive in the sense that we present protocols that attain said ultimate limits when any one of the three resources diverges.

By explicitly accounting for the level of assumed control and introducing the notion of control complexity, we emphasise a crucial resource that is oftentimes overlooked in quantum thermodynamics. For example, note that infinite-dimensional thermal systems are, in principle, readily available in many physical contexts [25–27]. The crucial point, though, is that in order to usefully employ these for desired (thermodynamic, or quantum information processing) tasks, one requires a high level of control over them [35, 37]: For instance, such an infinite dimensional thermal machine is *not* sufficient to cool at the Landauer limit if the control complexity of the interaction does not also diverge. Our analysis brings the concept of control complexity to the forefront and demonstrates that it must be considered in order to understand the practical limitations of manipulating quantum systems. In essence, our work takes a complimentary approach to that of standard resource theories by constructing optimal protocols—both for finite-dimensional systems and harmonic oscillators, and with arbitrary Hamiltonians—that saturate ultimate limiting bounds given sufficient resources.

In doing so, we see that protocols saturating the Landauer bound are reminiscent of a Maxwellian demon with perfect control over all degrees of freedom [35, 37]. More concretely, in order to saturate the Landauer bound, the protocols that we derive require fine-tuned microstate control that can only be implemented with the aid of a *coherent* external work source, i.e., a quantum battery [52–55] with time-dependent control and a precise clock [56–58]. From a thermodynamic perspective, this may seem unsatisfactory, as the joint system-plus-machine is not isolated—indeed, it has been argued that Maxwell’s demon “cannot operate” [59]. To remedy this shortcoming, we extend our analysis to explicitly account for a thermodynamic energy source (i.e., a hot and a cold bath) and close the joint system by restricting to global energy-conserving unitary transformations.

This *incoherent* control setting corresponds to a minimal level of overall control, where an agent need only switch on and off an interaction in order to generate a spontaneous transformation; this could be seen as a thermodynamically-driven quantum computer, i.e., with only a heat engine driving the computational process [60–63]. Thus, the assumption of incoherent control provides a fully thermodynamic paradigm, considering both an initial thermal state of the machine *as well as* the system-machine interaction to be implementable with minimal (i.e., thermodynamic) external control. In this chapter, we highlight the distinction between what is possible given the assumption of arbitrary (coherent) control versus thermodynamic (incoherent) control, in particular providing a “no-go” theorem that demonstrates the impossibility of preparing a pure state with incoherent control in finite time and with finite control complexity regardless of the amount of energy drawn from the hot bath; this is in stark contrast with the coherent control setting, where a suitable work source can be used to achieve said goal.

Continuing within this fully thermodynamic paradigm, we then show how one can achieve a pure state with finite energy cost at the expense of either diverging time or control complexity by mimicking the corresponding coherent control protocol with repeated incoherent sub-procedures. However, the price of doing so is that the Landauer bound for said energy cost is not attainable for any realistic (i.e., finite temperature) hot and cold baths. In its stead, we derive a novel ultimate limit for this scenario (in equality form), which we dub the *Carnot-Landauer* bound, being a modification of the Landauer limit with additional factors accounting for the Carnot efficiency of a heat engine connecting the two thermal baths. Importantly, in the limit of an ideal engine connected to an infinitely hot bath, this bound reduces to the Landauer limit, which is consistent with intuition, as an infinite-temperature heat bath can be modelled as a work source that can deliver energy at no entropy change.

This result is novel, fundamental and practically relevant: It is fundamental because it generalises Landauer’s bound to a fully thermodynamic setting and it is practical inasmuch as it provides the ultimate limitations concerning thermodynamically driven quantum information processing. Furthermore, we provide protocols that saturate the Carnot-Landauer bound for baths of any temperature, demonstrating how to achieve said limits for the energetic cost of cooling with *only* thermal resources. This latter point is crucial, as the cost of control is often the “Achilles’ heel” that ruins any potential quantum advantage—said cost is mitigated in the incoherent control paradigm, which embodies the assumption of minimal control at the outset.

Our work here thus both generalises Landauer’s erasure principle and, at the same time, unifies it with the Third Law of Thermodynamics. The unification arises by considering (asymptotic) trade-offs between various resources that are necessary for cooling physical systems or erasing information: As we show, if any two of the resources (energy, time, or control complexity) are restricted to be finite, the third must *necessarily* diverge in order to create a pure state, in line with Nernst’s unattainability principle. The generalisation is provided via the incoherent control paradigm, where only energy-conserving unitaries driven by a heat engine can be used to implement the cooling protocol; thus, not only is the initial state of the machine thermal (as Landauer originally envisaged), but also the level of control of the agent in line with thermodynamic assumptions (i.e., no time-dependent fine-tuned control like a Maxwellian demon, but rather only the ability to switch on and off energy-conserving interactions).

In this first chapter, we mainly consider cooling performance in the asymptotic setting, i.e., with one of the relevant resources allowed to diverge, in order to connect concepts regarding resource usage with Nernst’s unattainability principle to provide fundamental conceptual insights. As a final result, we analyse a pragmatic situation in which all three resources are restricted to be finite. This corresponds to the practically relevant case of having control over only a finite-sized machine as well as both finite time and energy at one’s disposal. The question of what can and cannot be achieved in this setting, as well as the optimal strategies to do so (whenever possible), is a pressing open problem that is difficult to address in general, as the full energy structure and microscopic details of the states and transformations involved come more prominently into play. By focussing on special cases and comparing the finite-resource behaviour of different (coherent) control strategies that all asymptotically achieve a pure state (with either

time or control complexity diverging) and at Landauer energy cost, we demonstrate a number of salient points regarding the interplay between time, energy and control complexity regarding such imperfect cooling (i.e., cooling to finite temperature or creating an approximately pure state). We additionally examine the relationship between the two extremal control paradigms considered by comparing a coherent and a similar incoherent protocol and demonstrate that more resources are required in the latter setting to achieve the same performance as the former.

At this point, we consider the issue of “Landauer vs. Nernst” to be resolved: They are, in a sense, both right, but Landauer’s considerations focussed on the special case of minimising energy resources at the expense of diverging time or control complexity. We now move on to concentrate on such resource trade-offs in the more pragmatic finite-resource setting. Here, the implementation of any cooling protocol to reach a particular desired temperature (when achievable) becomes a complicated expression of various microscopic details which is generally difficult to optimise with respect to, say, the energy cost. Similarly, the coolest achievable temperature itself, given any finite amount of input energy, time, and control complexity, depends upon particularities of the entire configuration. However, disregarding the constraint of minimising over the energy cost, the optimal protocol (regarding the final temperature) is relatively clear and can be deduced. This is the main focus of **Chapter 2**: Understanding the relationship between time and control complexity in terms of the impact on the achievable temperature with finite energy cost (and assuming coherent control).

To return to the broader picture for a moment and set this present work into a relevant context, recall that efficient and reliable quantum technologies rely on using cold states to cleanly encode information [12, 64]. As we have discussed, in any physical implementation of an information processing task, there is a trade-off concerning how resources should be allocated. With respect to the cost of cold state preparations in particular, the Third Law of Thermodynamics stipulates that one requires infinite energy, time, or control complexity to prepare perfectly cold states—none of which is feasible. Thus, perfectly pure states are impossible to prepare, and so a more pragmatic aim concerns preparing states as cold as possible. Our work in this chapter focuses on understanding the (finite) resources needed, and how to best use them, in order to prepare a state of desired (finite) temperature.

With this in mind, we fix the energy cost to be finite (and sufficient to achieve a desired final temperature) and analyse the trade-off between time and control complexity as resources. In the previous chapter, we mainly considered the resource of complexity in terms of control over spatial degrees of freedom; here, we shift perspective to consider complexity in the form of temporal control over *memory effects*. The corresponding physical assumption is that the agent can not only control complex spatial (i.e., many-party or high-dimensional) operations, but can additionally do so on timescales shorter than the natural rethermalisation timescales of the machine. With the ever-increasing precision and control with which one can manipulate quantum systems, such an assumption is not out of near-term reach in natural (i.e., noisy) settings, and even closer in proof-of-principle experiments. With this consideration, we effectively map a spatial notion of control complexity—i.e., in terms of the number machine systems involved in

each interaction—onto a temporal one, namely the ability to control and recycle used machines in time. Such a picture implies that the machine serves as a *memory* mechanism, establishing memory as a particularly pertinent type of complexity in the temporal setting (an idea that we will later return to in **Part II**).

Generally speaking, memory effects indeed play a crucial role in complex phenomena occurring at the nanoscale, including physical and chemical processes underpinning cutting-edge technologies [65, 66]. Depending on the level of control enjoyed by an agent over the system and its interactions with the memory carriers, such effects can lead to either an advantageous or detrimental impact on the figure of merit at hand [67–80].⁸ In this chapter, we highlight how control over such memory effects—concretely, the speed at which the environment can be addressed—can provide significant advantage for the task of cooling quantum systems; in particular, we develop a protocol that displays rapid (i.e., short time resource) cooling at the expense of requiring a higher degree of experimental control over the memory.

To do so, we put forth a general framework for dealing with memory effects in refrigeration schemes over multiple cycles. This is the relevant paradigm for realistic studies of the efficiency of cyclic [24, 81] and autonomous [82, 83] quantum machines as well as algorithmic cooling techniques [44, 84–88]. In the general case (i.e., with arbitrary memory), fairly comparing the performance of the former concrete approaches is difficult due to particular variations and assumptions in each setting; on the other hand, the overarching resource theory of quantum thermodynamics (see, e.g., Refs. [89, 90]) proves too general to isolate the true source of any advantage.

The biggest conceptual challenge that our work in this chapter overcomes thus concerns the inclusion of memory effects that do not immediately trivialise the problem or introduce unwanted artificial advantages from arbitrarily increased control or other resources. We have solved this problem by introducing a generic microscopic mechanism for memory transfer through a physically motivated generalised collision model [91–95]; while not completely general regarding memory per se, this approach permits a tractable amount of memory into the dynamics and allows for meaningful comparison between various memory structures. Specifically, our framework does not offer increased (with respect to the memoryless setting) control over the accessible set of system-machine interactions, coherence, or any other resource, except for the ability to access machine subsystems that have been interacted with in previous cycles.⁹

Our main result shows—perhaps surprisingly—that with memory, one can achieve an exponential improvement in cooling performance in the number of memory carriers. This improvement is reminiscent of a similar enhancement demonstrated by HBAC, where the experimenter is offered control over auxiliary “compression” qubits [44, 84–88]; indeed, not only is our result here equivalent to this setting for qubit systems, it generalises the HBAC scenario to arbitrary target systems, machines, and interaction Hamiltonians. Moreover, our approach offers further generalisation by allowing for *adaptive* cooling strategies: We derive the

⁸Recall that whenever the memory is uncontrollable, such a memory-carrying system is typically referred to as an *environment*; whenever it is controllable, we rather use the term *machine*.

⁹In other words, the memoryless scenario (which has been analysed in Refs. [82, 83]) is a straightforward special case of the situation with memory, namely when the number of memory carriers is zero.

adaptive stepwise optimal protocol (with respect to the ground-state population at each time) and demonstrate a finite-time advantage over the asymptotically optimal non-adaptive scheme. Lastly, our results reduce to all known results in the memoryless setting [82, 83], and, being based upon majorisation theory, apply to all reasonable notions of cooling, such as increasing the ground state population, driving the system to a thermal state at lower temperature or reducing the entropy (which are generally inequivalent for high-dimensional systems) [96].

In summary, our work here both generalises and unifies a number of seemingly disconnected approaches to quantum cooling in the presence of memory. Viewed from the context outlined in **Chapter 1**, we can read these results as demonstrating the trade-off behaviour between time and control complexity (in the particular form of control over memory effects) to achieve a desired cooling procedure with practical resources, providing an important initial foray into the realm of developing optimal strategies for cooling with finite resources.

We now take a step back to consider the “bigger picture” and set said works in a contextual relationship with what is to come in the second part of this dissertation. At its core, the first part of this thesis focuses on the notion of *complexity*. Abstractly, this can be (and historically mainly has been) considered from a purely information-theoretic standpoint: Complex correlations in spatial or temporal datasets [11, 97], algorithmic or computational notions of complexity [9, 10, 98, 99], and in terms of the difficulty of winning strategic games and/or estimating mathematical problems [100]. From the more concrete physical perspective, we intuitively think of complex states or processes as those that are either complicated to describe, display exotic behaviour, or can be used to achieve difficult tasks; with regard to the latter in particular, we see the importance of *control*. Thermodynamics—being, in a sense, the physics of information—inextricably links these perspective, leading to the notion of *Control Complexity*.

Concerning physical processes, we have seen how control complexity has both spatial—i.e., access to multiple environmental degrees of freedom in a fine-tuned manner—and temporal—i.e., control over memory effects across multiple timescales—elements to it. So far, we have analysed this link for a prototypical difficult thermodynamic task, namely cooling; however, the connection does not stop there. More broadly, parallels can be drawn between control complexity and memory effects on the one hand with various thermodynamic phenomena (and beyond) on the other. For instance, below we will briefly (and somewhat “hand-wavily”) discuss two types of complex thermodynamic behaviour that can be simulated given control over the memory, before transitioning to the realm of open (quantum) processes more broadly.

Take, for example, the idea of *ergodicity*, which, loosely speaking, states that for most physical processes, a system will tend to explore its entire phase-space uniformly as it evolves (due to unavoidable interactions with the environment and the ensuing chaos). For non-ergodic processes, there are all sorts of complex structures in the evolution that make the dynamics difficult to characterise when the environment is inaccessible: Aperiodic recurrences, erratic behaviours where the system gets “stuck” in particular “pockets of phase-space”, long-term and complicated multi-time correlations, to name but a few. On the other hand, for ergodic processes, all such behaviours “wash out” over time, thereby allowing one to replace (mathematically difficult and

physically complex) time averages with (relatively simple) ensemble averages. Flipping the paradigm, given the ability to engineer and control the environment and its interactions with the system, one can *simulate* such non-ergodic processes and subsequently the aforementioned complex physical phenomena.

A somewhat related—albeit distinct—example concerns *thermalisation*, i.e., the phenomenon that physical systems tend to equilibrate and therefore (at least for macroscopic observables) take on stable values with small fluctuations. Intuitively, the underlying mechanism for thermalisation is that the environment continuously dissipates information away from the system,¹⁰ thereby driving it towards some fixed equilibrium state, which, in the case of a thermodynamic environment is typically a thermal state for which measurements of macroscopic observables lead to sharp distributions. Such dissipation of information embodies the assumption that the process is memoryless: If, on the other hand, a process has memory, then true thermalisation cannot occur, as it must eventually display recurrent behaviour. Again, given control over the memory (i.e., temporal control over system-environment interactions and environmental degrees of freedom), one can realise the ability to do something that is thermodynamically difficult, namely keep a system from equilibrating (which is its natural tendency).

However, concerning the impact of memory control upon physical behaviours, one need not start from the foundations of thermodynamics; similar connections to those above tend to hold in general throughout physics. Although thermodynamics provides a well-justified starting point for analysing *natural* phenomena (i.e., loosely speaking, what we tend to see under minimal assumptions), it is also interesting to see how far these ideas go when considering *engineered* processes (i.e., given control with respect to the laws of physics more broadly). For instance, both of the above examples occur similarly in the classical and quantum settings. However, there are fundamental differences between classical and quantum processes, both with and without memory. In order to satisfactorily identify and quantify these differences, though, it is first necessary to provide a proper delineation between which processes one considers to be classical and which to be quantum. As we will see, what it means for a process to “look” classical is highly dependent upon whether or not there is memory (and how one probes it). Moreover, there is a fundamental distinction between quantum and classical memory itself, leading to intricate and complex temporal correlations that cannot be attained in the classical realm.

In the first part of this thesis, we saw how control complexity fundamentally limits our ability to process (quantum) information in both space and time with respect to thermodynamics. As we turn to the second part of this dissertation, we drop the thermodynamic assumptions and shift our focus more broadly to consider the ultimate limitations set by physics upon spatio-temporal (quantum) information processing; thus, we move to the more general paradigm of *open (quantum) processes* (without thermodynamic constraints), where we will see complicated *multi-time phenomena* arising as a signature of complexity.

¹⁰More generally, even isolated complex systems have a phase space so large that they effectively become their own environments and equilibrate accordingly; additionally, meaningful notions of equilibration consider the fact that one can typically only measure highly degenerate (i.e., macroscopic) observables, making the formulation yet another instance of an interplay between the underlying dynamics and the level of control. See, e.g., Ref. [101].

Part II

The open systems paradigm begins from the assumption that *no system is truly isolated*: Environmental effects impact the evolution of any system of interest, typically in a non-negligible fashion [65]. This idea can be applied to the classical setting, introducing spurious sources of noise by way of subjective ignorance on the part of the agent, leading to the theory of classical stochastic processes. Similarly, accounting for environmental interactions that are not directly accessible in the quantum realm leads to the formulation of quantum stochastic processes. However, here, in addition to the randomness that arises due to subjective ignorance of the environment, one must deal with the fundamental randomness associated to quantum measurement, as we will discuss in detail below.

In contrast to our earlier considerations, in general, the environment and its interactions with the system need not be thermal—which can be considered synonymous with “simple”—but rather can be structured, correlated, and complex, typically leading to exotic dynamical behaviour such as aperiodic recurrences [65, 66], resistance to thermalisation (as exemplified, e.g., in decoherence-free subspaces [102]), and multi-time correlations [103–106], to name but a few (see Refs. [107–109] for in depth reviews). Control over such memory effects—as we have seen in the special case of open quantum processes involving thermal environments—can improve efficiency and the ability to perform certain difficult tasks. To mention some prominent (non-thermal) examples, such tasks include the creation, manipulation and long-term preservation of coherences and correlations [69, 110], reservoir engineering to simulate complicated dynamical evolution [74, 77, 79, 111–119], sophisticated and robust quantum error characterisation and correction in the presence of (temporally) correlated noise [120–122], optimal dynamical decoupling procedures [68, 123, 124], and advanced design for quantum circuit architectures [106, 125–136].

Thus, we see the importance for understanding, characterising, and exploiting potential memory effects in open quantum processes, which is ever-increasing due to technological improvements that allow us to manipulate interactions with enhanced levels of speed, precision, and complexity. Although characterising classical stochastic processes with memory is already a daunting task due to the sheer amount of data that must be recorded (which, in general, grows exponentially in the length of the memory), there is no problem from a theoretic perspective since the joint probability distribution over sequences of measurement outcomes corresponds *precisely* to the description of the stochastic process itself. On the other hand, when one tries to do similar in the quantum setting, one faces an immediate *breakdown of formalism* [137–139]: Here, the invasive nature of quantum measurements tends to blur the line between the agent and the underlying process per se, which has historically lead to a “zoo” of (non-equivalent and even sometimes incompatible) definitions of memory for quantum stochastic processes [140].

To see the fundamental “issue” in quantum mechanics, note that, in order to experimentally probe an open process (classical, quantum, or even beyond), one sequentially probes a system of interest across multiple points in time and records the corresponding family of probability distributions over outcomes; we refer to this as the *operational description* or *observational*

data. In particular, whenever we refer to an “operational” property, we mean something that is deducible from such measurement statistics alone, without requiring any reference to a particular underlying description. In the classical setting, there is no difference between said operational description and the underlying *dynamical model* that engenders said statistics—they are equivalent descriptions of the same process.¹¹ Thus, from the observational data alone, one can deduce the correct behaviour of the underlying classical stochastic process, completely characterise it and determine all of its properties.

In the quantum realm, things become significantly more complicated. Again, from an operational standpoint, the probing procedure of a quantum process is the same as above; however, here the joint probability distribution over a sequence of outcomes no longer coincides with the dynamical description of the underlying process. This distinction can be seen in multiple ways (as we will explore throughout the second part of this thesis), but a particularly enlightening one is the fact that from a single joint probability distribution recorded from probing an open quantum dynamics over time, one cannot deduce the correct behaviour of the process upon any subset of times, leading to the breakdown of the so-called *Kolmogorov consistency conditions* [145, 146], which we now discuss.

To exemplify the breakdown of Kolmogorov consistency and the ensuing difficulty in characterising quantum processes, consider performing three steps of a Stern-Gerlach experiment in sequence, with the apparatus measuring the spin first in the σ_z direction, then σ_x , and finally σ_z again. For the initial state $|+\rangle := \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, the experimenter would record $\mathbb{P}(z_3, x_2, z_1) = \frac{1}{8}$ for all possible outcomes $z_1, z_3 \in \{\uparrow, \downarrow\}$ and $x_2 \in \{\leftarrow, \rightarrow\}$. However, suppose now that the experimenter does *not* make a measurement at time t_2 , which we denote by the identity instrument \mathcal{I}_2 . For any classical stochastic process, the latter experiment would yield probabilities that can be deduced from those of the former via *marginalisation*; in other words, there is no difference between “not measuring” and “measuring and averaging” over all possible outcomes. However, in the quantum realm, the experimenter would record the probabilities $\mathbb{P}(\uparrow_3, \mathcal{I}_2, \uparrow_1) = \mathbb{P}(\downarrow_3, \mathcal{I}_2, \downarrow_1) = \frac{1}{2}$ and $\mathbb{P}(\uparrow_3, \mathcal{I}_2, \downarrow_1) = \mathbb{P}(\downarrow_3, \mathcal{I}_2, \uparrow_1) = 0$. Evidently, this is *not* equivalent to $\sum_{x_2} \mathbb{P}(z_3, x_2, z_1) = \frac{1}{4}$, i.e., marginalising over the probabilities from the experiment when all three times are measured, which is how one would deduce such behaviour when considering classical stochastic processes. Throughout, we will refer to statistics for which summing over outcomes at superfluous times leads to the correct “contained” probability distributions as *Kolmogorov consistent*, and those that do not (e.g., as above) as *inconsistent*. Finally, note that this example has nothing to do with memory and only concerns measurement invasiveness; as we will see below, things become even more complicated when memory is involved.

To properly describe the second experimental procedure, one must explicitly “do nothing” at the superfluous time rather than measure at all three times and then marginalise over statistics; the latter only leads to the correct description in the classical realm because measurements are non-invasive, and thus doing nothing is equivalent to measuring and subsequently averaging over statistics. In the quantum setting, this is not the case. This breakdown of Kolmogorov

¹¹We implicitly assume the experimenter to make sharp measurements, in line with the standard assumptions pertaining to classical stochastic processes. If noisy classical measuring devices are used, the two descriptions no longer coincide and one runs into similar trouble as in the quantum setting (see, e.g., Refs. [141–144]).

consistency [145, 146]—i.e., the ability to deduce the sub-statistics that correctly describe the behaviour on fewer times from the full statistics—is a result of the fact that quantum measurements are invasive in general, which in turn provides the origin of the lauded Leggett-Garg inequalities which can be understood as a test for measurement invasiveness [147–149]. Put simply, it implies that the joint probability distribution for any fixed choice of instrument does not fully characterise a quantum stochastic process, in stark contrast to the classical case.

As a consequence, in order to operationally characterise a quantum process, one would need to record the (distinct) set of joint probability distributions for *every possible* choice of probing instruments,¹² and from this observational data infer the correct dynamical description of the underlying process at hand, i.e., perform a full multi-time quantum process tomography [129, 150]. Besides being practically much more resource intensive than its classical counterpart, such techniques have only recently been developed and employed to assess multi-time phenomena in quantum theory, i.e., to describe quantum processes *with memory* [128–136, 150, 151].

As we saw in the example above concerning the three Stern-Gerlach devices, measurement invasiveness makes it tricky to operationally describe processes that are sequentially probed at multiple (i.e., more than two) times, even in the absence of memory. To circumvent the issue of measurement invasiveness, many studies regarding open quantum dynamics have attempted to forego the necessity of a truly *multi-time* description and define important properties such as memorylessness in terms of (sets of) *two-time* statements, for which a clear operational picture can be formulated in terms of quantum channels [152–158]. However, such approaches only provide partial information, since memory is a genuinely *multi-time* phenomenon that is unsatisfactorily captured when only considering two-time correlations. Hence, throughout the second part of this dissertation, we will employ the formalism of quantum combs/process tensors [125, 126, 129] to analyse multi-time behaviour. These are genuinely multi-time approaches to open quantum processes that provide *complete* descriptions of complex processes with memory and are not plagued by the aforementioned problems, as they encode the correct probability distributions over sequential measurement outcomes *for all possible* choices of instruments (including, for instance, “not measuring” at certain times). In other words, they constitute a single object that provides a full operational characterisation of the underlying process at hand.

While this approach makes it clear that there exists an unambiguous description of quantum processes, what one sees of them is quite naturally highly dependent upon how one probes them. For instance, modifying the Stern-Gerlach example above to the case where all three devices measure the spin in the computational basis, one would no longer observe any difference in the recorded statistics between $\sum_{z_2} \mathbb{P}(z_3, z_2, z_1)$ and $\mathbb{P}(z_3, \mathcal{I}_2, z_1)$, since the intermediate measurement is now non-invasive with respect to the state at the second timestep (which contains no coherence with respect to said measurement basis). Thus, we see that making inferences about the underlying process from operational data concerns an intricate relationship between the underlying structural properties of the process and one’s ability to resolve them.

Similarly, the discussion regarding a proper analysis of memory effects hinges largely on the interplay between *structural* properties and *operational* ones, which is a general theme

¹²Or at least a full basis of them, due to the linearity of probabilistic theories such as quantum theory.

throughout quantum theory of which many relevant instances abound: Well-known examples of the former include coherence, discord, commutativity, and entanglement; examples of the latter would be criteria based upon measurement invasiveness (e.g., Leggett-Garg inequalities [147–149], Kolmogorov consistency [145, 146]), contextuality (e.g., Kochen-Specker theorem [159, 160]), non-locality (e.g., Bell inequalities [161–163]). From the perspective of practical information processing, it is arguably the latter operational types of criteria that are more important as they have direct observational consequences “built in”. Although beginning with certain structural properties often leads to quite clear operational repercussions, doing so requires a “leap of faith” in the validity of the assumptions and the models. Beginning with concrete operational data is more scientifically justified (as it makes no model-dependent assumptions); however, inferring structural properties based upon observed data alone is oftentimes rather nuanced, especially given certain operational restrictions such as limited control or restricted measurements.

The relationship between entanglement and non-locality provides a prime example. For simplicity, we consider the Clauser, Horne, Shimony, and Holt (CHSH) game [164]. Assuming an initially separable quantum state (structural), it is clear that the CHSH correlation function cannot exceed the classical bound of 2 (operational/experimentally testable) [161], and similarly assuming an initially entangled quantum state (structural), it cannot exceed the Tsirelson bound of $2\sqrt{2}$ (operational) [165]. Conversely, suppose that, for a choice of measurements and a given state, one collects statistics and computes the CHSH correlation function to be of a certain value (operational). From this, and knowledge of the measurements performed, what can be inferred about the entanglement properties of the measured state (structural)? Some questions are easy to answer: If the CHSH correlation function exceeds 2, one can conclude that the state must be entangled (this is an example of an entanglement witness [166, 167]). Other questions are much more difficult: If the CHSH correlation function is less than 2, it might be that the state is entangled, or it might be just that the measurement combinations were not chosen correctly (in the sense that they can reveal the entanglement of the underlying state). For such characterisations (i.e., conclusive inferences about the underlying structural properties based upon operational data), one typically requires performing more complicated experiments to eventually rule out the validity of certain assumptions.

While this delicate relationship between structural properties and their ultimate expression in observed data is well-known and studied for *spatial* properties of quantum states, it is much less explored in the temporal setting, i.e., for (multi-time) quantum *processes*, where invasiveness of measurements and quantum memory play a pivotal role. One might even say that certain Leggett-Garg type inequalities provide the only such connection in the temporal scenario that is broadly known, and the following results that we present go far beyond these.

The overall focus of **Part II** of this dissertation concerns such relationships between structural and operational properties in the context of multi-time open quantum processes. We begin in **Chapter 3**, where we consider the question of “classicality” for open processes with memory, making a connection between Kolmogorov consistency (as an operational definition of a classical stochastic process) and the ability (or lack thereof) of the underlying (potentially quantum) process to generate and detect structural properties that are often related to notions

of classicality, such as coherence and discord. We characterise the structure of all processes that could meaningfully be considered classical (i.e., that hide any potential underlying quantumness) when probed in a fixed basis, and further show that memory plays a fundamental role by demonstrating that there are processes with memory that *cannot* hide their quantumness, no matter how they are probed (which cannot occur for memoryless processes). Following this, we take a similar approach to study the link between commutativity—a structural property at the very core of quantum theory—and classicality in **Chapter 4**. After shedding light upon a number of complications that make the relationship between these two properties rather layered, we generalise Gerhardt Lüders’ theorem [168] (which answers such questions for two-time processes) to the multi-time setting. Lastly, in **Chapter 5**, we conclusively demonstrate a clear distinction between quantum stochastic processes without memory (structural property) and Markovianity of observed statistics (operational property), uncovering two new genuinely quantum multi-time phenomena—namely, *hidden quantum memory* and *incompatibility of Markovian statistics*, i.e., memory effects that are fundamentally present but cannot be directly observed. Neither of these phenomena can occur in the classical setting and they serve to witness the impossibility of positing any possible underlying memoryless description in quantum mechanics, even though the observed statistics do not display any non-Markovianity.

Chapter 3 concerns the connection between structural and operational notions of classicality for multi-time open quantum processes (both with and without memory). Since its inception—more than 100 years ago—quantum theory has not ceased to baffle scientists with counter-intuitive phenomena. One of the more vexing, and notoriously unanswered questions, is: *Which traits are fundamentally quantum?* Determining what makes a phenomenon genuinely quantum and quantifying the underlying dynamical resources that are necessary to observe a deviation from classical behaviour in real experimental conditions constitute two related fundamental questions that still await a fully satisfactory answer.

While there exist myriad proposals for demarcation between classical and quantum phenomena based upon certain structural features, from an experimental perspective, they appear somewhat unsatisfactory. In reality, we do not categorise the world in abstract mathematical terms, nor do we have direct experimental access to them.¹³ Consequently, an *operationally accessible* notion of genuine quantumness should not be based upon structural concepts either, but solely in terms of experimentally measurable entities. Although the concrete mathematical structure of quantum mechanics possesses many traits that are absent in classical theory (e.g., coherence, discord, and entanglement, to name but a few) and have consequently been postulated to be genuinely quantum resources, to date, the explicit relation between such underlying structural features of quantum theory to concrete, observable phenomena that are classically inexplicable remains unclear and under active debate.

One possible avenue that ensures such properties are clearly classified in terms of directly accessible behaviour is to *begin* with the operational approach: Here, the classicality of a phenomenon is investigated directly from observable quantities (e.g., probability distributions

¹³As per Asher Peres [169]: “Quantum phenomena do not occur in a Hilbert space. They occur in a laboratory.”

over measurement outcomes), without any additional assumptions on the underlying theory. This line of reasoning has led to the famous Kochen-Specker [159, 160], Bell [161–163], and Leggett-Garg [147–149] inequalities which must be satisfied in classical physics but can be violated in quantum mechanics, therefore providing an operational test for demarcation. The question regarding *which kind of underlying quantum resources are required to lead to such non-classical observations* then relies upon connecting said operational violations back to the dynamical and structural properties. Importantly, this latter concern returns us to understanding the ensuing implications that underlying structural properties have upon observed statistics: Experimentally accessible entities provide the clear-cut definition of classicality, while structural properties play a secondary (albeit important) role as necessary requirements for (but not definitions of) a delineation from classical behaviour.

In this chapter, we provide a definition for the classicality of a process that is a priori rooted in experimentally observable phenomena only. This shift in perspective allows us to identify the class of processes that rightfully carry the moniker “quantum”, as well as the dynamical resources required to realise them, thereby bridging the apparent divide between genuinely quantum resources and observable non-classical phenomena for general temporal processes.

Our approach is fully general in the sense that it considers all possible processes where a system of interest (potentially interacting with an inaccessible environment carrying memory) is probed at multiple times. In this setting, satisfaction of the fundamental Kolmogorov consistency conditions between observed statistics—meaning that all “sub-statistics” can be readily deduced by marginalising over the full joint probability distribution, i.e., any potential measurement non-invasiveness is not detectable in the statistics—provides a meaningful operation definition of classicality. Indeed, these conditions form the mathematical foundation of the theory of classical stochastic processes and provides a set of experimentally-testable requirements that every *classical* stochastic process must satisfy [145, 146]. Based on this line of demarcation between the classical and the quantum regime, we characterise the set of all quantum processes that fundamentally *cannot* be understood in classical terms, quantify the amount of non-classicality, and identify the underlying physical mechanisms that lead to their non-classicality. While similar results had previously been derived for the limited class of memoryless processes, our present results—based on the quantum comb/process tensor description of general quantum stochastic processes with memory—relinquish previously necessary restrictions on the considered dynamics and provide a comprehensive picture of non-classicality for all conceivable kinds of temporal processes (within quantum theory).

In particular, we show a clear distinction between the quantum resources necessary to observe non-classicality depending on whether or not the process has memory, and identify the relevant resources in both scenarios. For memoryless processes, it is the ability for the process to generate and detect coherences—an often used structural signifier of quantumness—in the system state that is tantamount to non-classical behaviour [170]. Somewhat surprisingly, when memory effects are present, an open quantum process can display non-classical behaviour even if the state of the measured system never displays any coherences (as we demonstrate), thus underlining a fundamental difference between quantum processes that keep track of their history

and those that do not. In the former case, we show that it is the quantum discord that is built up between the probed system and its environment (and subsequently detected), that determines the quantumness of the process. Additionally, we demonstrate that while processes without memory can always “hide” their quantumness (in the sense that there always exists choices of measurements at different times such that the observed statistics will not display invasiveness), on the other hand, there exist *genuinely quantum processes* with memory that can *never* appear classical (i.e., for any choice of measurements).

On the whole, our results here provide a comprehensive picture of non-classicality for open processes, and firmly connect abstract, mathematical notions of quantumness with a fully operational one. Exploring the full space of conceivable processes, we show that the structure of non-classicality in quantum mechanics is multi-layered, and a fundamental gap arises between processes with and without memory. Besides this theoretical progress, the concepts we introduce are immediately applicable to a wide array of experimental situations, like, e.g., quantum transport phenomena [70, 171–173], where it is becoming increasingly important to systematically and quantitatively assess whether or not reported advantages are actually due to quantum resources, or if such quantum resources are merely unavoidably present, but do not necessarily contribute to observed speed-ups.

In **Chapter 4**, we retain the Kolmogorov consistency conditions as an operational notion of classicality but shift the structural concern to the commutativity of measurement operators—yet another fundamentally non-classical structural feature that is at the origin of the Heisenberg uncertainty principle [174–176]. Intuitively, commutativity of a pair of observables implies that they are jointly measurable; in other words, given an arbitrary quantum state, the order in which the measurements are performed is inconsequential—each experimental configuration will yield the same statistics nonetheless. Clearly, such a concept is related to measurement non-invasiveness, i.e., Kolmogorov consistency.

Indeed, this connection has been considered by Lüders [168] for the case of projective measurements and subsequently extended to more general scenarios concerning two sequential measurements [177–179]. In this two-time setting, it has been shown that commutativity and Kolmogorov consistency coincide in many cases, justifying commutativity as a sensible notion of classicality from a structural *as well as* an operational standpoint. However, as we discuss throughout this chapter, such a direct link between these two a priori distinct concepts can only be meaningfully established in the situation where only two sequential measurements are being performed. In the multi-time setting with non-trivial dynamics taking place between measurements, it is unclear how Lüders’ results carry over and what operators play the role of “observables” whose commutativity should be checked for general (i.e., not necessarily projective) measurement schemes.

In particular, both the underlying dynamics and the effects of the measurement devices must be properly accounted for. For processes with memory, deriving the relevant operators is fraught with difficulty due to the requirement of tracking the entire history of both the process and measurement outcomes. Thus, we limit our considerations to the memoryless

scenario throughout this chapter, which, as we will see, is nonetheless significantly rich. As we demonstrate, in the memoryless setting, one can fully capture the interplay between the dynamics and the measurements by combining them into individual operators; it is then not necessarily the commutativity of the bare measurements (i.e., pertaining to the measurement instruments themselves) per se, but rather the effective measurements (i.e., the measurements *combined* with the underlying dynamics) that render the observed statistics classical or not.

An additional difficulty that becomes immediately apparent in the multi-time setting is the fact that one does not have access to a full basis of quantum states as is assumed in the two-time case. For instance, in Lüders’ original consideration, it is only under the assumption that measurement non-invasiveness holds *for all states* that allows one to derive a state-independent criterion on the observables themselves (namely their commutativity). While one could, in principle, invoke similar assumptions in the multi-time case, these would be far removed from the reality of most experimental schemes and would fail to cover even the simplest of special cases, e.g., for sequences of projective measurements. Thus, we make no such assumptions and consider fully general measurement sequences.

The consequences of doing so are two-fold. Firstly, it is no longer possible to make strong statements on the (effective) measurement operators alone, but rather only some concerning their complicated interplay with the system states at each time. In other words, we can no longer hope to derive state-independent statements such as commutativity of pertinent operators, but rather must consign ourselves to searching for weaker ones, such as commutativity *with respect to* the set of possible system states at each point in time. Secondly, allowing for general instruments means that, even though the process itself has no memory (by assumption), all of the statements that we make concern the entire (multi-time) sequence of history and future measurement outcomes. As we discuss throughout the chapter, this is because general instruments can perpetuate information on the level of the system itself since they do not reset the state of the system to a known one (e.g., consider a weak measurement)—thus, although the process per se carries no memory, the effects of the instruments serve to introduce a memory mechanism. As a consequence, the possible states at any given time depend upon an entire sequence of historic outcomes in general, as do all involved future “measurement” operators.

In this chapter, we take all of these considerations into account and identify the relevant operators that determine the non-invasiveness of measurements (and thus classicality of the observed statistics) for the multi-time setting (with non-trivial, memoryless dynamics between times) and analyse the conditions for which commutativity—or weaker versions thereof—correspond to satisfaction of the Kolmogorov consistency conditions (and vice versa). Our analysis—in the same vein as **Chapter 3**, albeit with a focus on commutativity instead of coherence and discord—therefore connects structural with operational notions of classicality for multi-time processes with general measurement settings.

For the special case of two sequential measurements, our results coincide with those of Lüders; however, as we demonstrate and for the reasons described above, the general situation is much more layered. We first show that commutativity (of the relevant operators) is a far stronger condition than Kolmogorov consistency; thus, just as in **Chapter 3** where we saw that there are

processes that generate coherence/discord but the observed statistics remain classical, here we see a similar result where the relevant operators need not commute to ensure classical statistics. We then show that weakening said commutativity to one with respect to the system states at hand provides too weak a condition, before proposing a condition of “absolute” commutativity (with respect to the system states) which—while being weaker than commutativity per se—still implies classicality.

In the converse direction, we find that satisfaction of Kolmogorov consistency does not necessarily imply any of the previously considered commutator relations, as we show by way of various examples. We then derive additional conditions—namely such that the possible set of states at each time essentially forms a basis with respect to all later measurement Kraus operators (which are assumed Hermitian)—whose satisfaction, when taken together with multi-time statistics that are Kolmogorovian, implies commutativity amongst relevant operators.

We finally relate our considerations to the important special case of sequences of projective measurements in a fixed basis, as was considered in the previous **Chapter 3**. There, we saw that in the memoryless setting, Kolmogorov consistency is equivalent to dynamics that do not generate and detect coherences—a three-time statement that can be judged on the basis of neighbouring dynamical maps alone. This restriction to neighbouring maps is made possible since projective measurements—unlike general measurements—break the flow of information through the system and it is thus not necessary to consider entire sequences of outcomes for the structural analysis of processes that yield classical statistics. When considering the relationship between commutativity and classicality, as we do in **Chapter 4**, we show that, although the aforementioned properties concerning the generation and detection of coherence (or lack thereof) follow directly from those we provide here for more general measurement scenarios, it is difficult to identify generally applicable commutator relations even in this special case where the measurements are restricted to sharp projective ones.

All in all, our work in this chapter offers yet another comprehensive analysis regarding the connection between structural—yet not directly observable—properties of multi-time open quantum processes and operational notions of classicality, further exemplifying the complex interplay between dynamics and the types of possible interrogations that quantum mechanics affords in the multi-time setting.

So far, we have seen that measurement non-invasiveness and the ensuing breakdown of Kolmogorov consistency can be related to various quantum dynamical resources, providing a distinctly non-classical signature. Finally, in **Chapter 5**, we take a deeper look at the differences between classical and quantum stochastic processes without memory and the subsequent impact such a structure entails for observational data, in particular regarding the connection between *memorylessness* and *Markovianity*; here, Markovianity concerns the memory properties of the observed statistics, while memorylessness concerns those of the underlying dynamics that engenders them.¹⁴

¹⁴Following the reasoning of **Chapters 3** and **4**, it will not come as a surprise that these two notions—former operational, the latter structural—generally do not coincide.

In the classical realm, as is well known, these two properties are one and the same: Any memoryless process leads to Markovian statistics; and conversely for any Markovian statistics derived from probing a classical process, one can build a memoryless dynamical description in terms of an independent sequence of stochastic matrices that faithfully reproduces the observed statistics. This is due to the non-invasiveness of classical measurements. More concretely, in the classical case, a *single* multi-time joint probability distribution over measured outcomes fully characterises the process—thereby encoding both all structural properties of the underlying dynamics such as memory effects and operational ones like Markovianity [145]. As we have discussed, this is not true for quantum stochastic processes, first and foremost due to the fact that quantum measurements are generally invasive [146].

As alluded to above, if one probes a memoryless quantum process with general instruments, information can be transmitted through the system itself, leading to temporal correlations and, in general, non-Markovian statistics (even in classical physics) [131]. However, this no longer holds true when restricting to sharp, projective measurements:¹⁵ In this case, as we demonstrate, any observed statistics gathered by probing a memoryless quantum process leads to Markovian statistics, just as in the classical setting.

Thus, at first glance, it seems as if the property of memorylessness manifests itself on the operational level (i.e., the observed statistics) in the same way for both classical and quantum processes. In the converse direction, it is clear that given a set of Markovian statistics deduced from probing a quantum process, one does not have enough information to uniquely construct a memoryless dynamical model since—in contrast to the classical setting—projective measurements do not span the full space of quantum measurement operations. Nonetheless, it is reasonable to assume that given any Markovian statistics, there *exists* some (potentially fictitious) memoryless quantum process that faithfully reproduces said statistics.

As we show in this chapter, such a model always exists whenever the measurements happen to be non-invasive: Indeed, this case boils down to the classical settings, where the observation of Markovian statistics is equivalent to the existence of a unique underlying memoryless dynamics. However, in the general case, i.e., where the measurements are invasive, such an equivalence is no longer guaranteed to exist.

Fundamentally, the potential absence of a memoryless quantum model that describes observed Markovian statistics stems from the following issue: Although one might be able to construct an independent sequence of quantum channels (i.e., a memoryless model) that describes the Markovian statistics on the *full* set of times, this is generally insufficient to deduce the correct statistics on any *subset* of times (which we refer to as “sub-statistics”). Again, the root cause of this problem is that quantum measurements are invasive, and thus a memoryless model constructed from the full measurement data will generally produce incorrect statistics for situations in which the experimenter opts to *not* measure the system at certain points in time.

Thus, it is clear that a naïve construction of a memoryless quantum model to describe observed Markovian statistics does not work. More generally though, the question arises whether

¹⁵Or, more generally, any instrument that breaks the flow of information, e.g., by measuring the system and feeding forward a state that is independent of the measurement outcome. Such instruments have been referred to as “causal breaks” and they play a crucial role in the proper definition of memoryless quantum processes [128, 180].

some memoryless model exists that can faithfully reproduce observed Markovian statistics on the full set of times *and* also all corresponding sub-statistics. In this chapter, perhaps surprisingly, we demonstrate that this is not the case and show that quantum processes can lead to Markovian statistics *that fundamentally require memory* for their creation, uncovering a novel genuinely quantum multi-time phenomenon.

We do so by showing that, in addition to Markovianity of the full statistics for any *memoryless* quantum process probed by sharp projective measurements, all *sub-statistics* must also be Markovian. Moreover, all such Markovian sub-statistics must coincide in the sense that they are independent of “how they are obtained”. Specifically, this means that *any* conditional probabilities can only depend on the most recent outcome (this amounts to Markovianity), and additionally these conditional probabilities are independent of whether measurements prior to the most recent outcome were performed or not; we refer to this latter property as “compatibility”.

We prove our main proposition by way of constructing explicit counter-examples. First, we present a four-step quantum process that leads to Markovian statistics when probed in the computational basis at *all four* points in time; however, if the experimenter chooses to *not* measure at the second time, the sub-statistics deduced are non-Markovian—that is, the conditional probabilities at the fourth time given knowledge of the third measurement outcome are *not* independent of the first measurement outcome. More concretely, we have $\mathbb{P}(x_4|x_3, x_2, x_1) = \mathbb{P}(x_4|x_3)$ but $\mathbb{P}(x_4|x_3, \mathcal{I}_2, x_1) \neq \mathbb{P}(x_4|x_3)$ (with the latter showing dependence on x_1). Such “hidden” memory that is activated for certain sub-statistics is in direct contradiction with a dynamical description of a memoryless quantum process.

Hence, we can understand the detection of hidden memory as an operational witness of memory effects in the underlying process. This naturally begs the question: *What if both the full statistics and all sub-statistics are Markovian—does there then always exist a memoryless quantum process that faithfully reproduces them?* We show that this is not the case by constructing an example that satisfies all of these former requirements but the resulting Markovian conditional probabilities are nonetheless incompatible. This again provides a contradiction with the possibility of finding a memoryless quantum process that faithfully reproduces the observed Markovian statistics. Thus, incompatibility of Markovian statistics serves as a finer witness for memory in the underlying process than the activation of hidden memory (with the latter being an example of incompatible statistics by definition).

Our results in this chapter imply that there exist Markovian statistics for which memory in the underlying quantum process is not just present but indeed *required*, since it is impossible to construct *any* memoryless quantum process that reproduces them. Additionally, both situations for witnessing this—whether the sub-statistics are Markovian and incompatible or non-Markovian (i.e., hidden memory)—*cannot* occur in the classical setting and require measurement invasiveness to be possible. Importantly though, the concept of hidden quantum memory is not merely a different manifestation of measurement invasiveness. Rather, it is a distinct and fundamental multi-time quantum memory effect; for instance, while memoryless quantum processes can violate Leggett-Garg inequalities, they cannot lead to hidden quantum memory. In this sense, we have uncovered a novel type of genuinely quantum multi-time phe-

nomenon. Lastly, our results here are similar in spirit to other quantum traits that require precisely the resource in their implementation that they ultimately hide, such as quantum channels that preserve all separable states but cannot be implemented via local operations and classical communication [167, 181, 182], non-signalling maps that require signalling [183], and maximally incoherent operations that necessitate coherent resources [184–186], to name but a few.

Overall, in the second part of this dissertation, we rigorously analyse and exemplify various connections between structural or dynamical properties on the one hand and entirely operational ones on the other, in particular with regard to the setting of open quantum processes probed at multiple points in time. In the end, we once again see that how complex (i.e., how much memory) or how non-classical a process looks boils down to a question of control—e.g., what types of measurements one can implement and what type of correlations one examines. Thus, **Part II** is an analysis of the interplay between memory and control, and their connection to underlying but a priori inaccessible properties. Taken together with the considerations of **Part I**, we have made a substantial foray into developing a holistic understanding of the intricate interplay between **thermodynamics, complexity, and multi-time phenomena**, and in particular the ensuing implications for **quantum information processing**. We finish the dissertation with a **Concluding Discussion**, where we summarise the broader significance of our results and present an outlook on future research directions.

Note

As indicated in the **List of Publications**, this thesis is a cumulation of a number of completed works; as such, although the chapters are linked thematically through the interplay of thermodynamics, complexity, and multi-time phenomena in quantum information processing, they can be read and understood as stand-alone pieces. Consequently, although we have aimed for consistency in notation, language, and general presentation style throughout this dissertation overall, there are inevitably slight inconsistencies to be found between chapters. For the convenience of the reader, we provide a list of commonly used acronyms below.

Acronyms

CNOT Controlled-NOT

XOR Exclusive-OR

CHSH Clauser, Horne, Shimony, and Holt

CJI Choi-Jamiołkowski Isomorphism

CP Completely Positive

CPTP Completely Positive and Trace Preserving

HBAC Heat-Bath Algorithmic Cooling

LP Linear Program

MIO Maximally Incoherent Operation

NCGD Non-Coherence-Generating-and-Detecting

NDGD Non-Discord-Generating-and-Detecting

NISQ Noisy Intermediate-Scale Quantum

POVM Positive Operator-Valued Measure

PR Popescu-Rohrlich

QRF Quantum Regression Formula

RW Reeb and Wolf

Part I

Cooling Physical Systems

Landauer vs. Nernst: What is the True Cost of Cooling a Quantum System?

Philip Taranto*, Faraj Bakhshinezhad*, Andreas Bluhm†, Ralph Silva†, Nicolais Friis, Maximilian P. E. Lock, Giuseppe Vitagliano, Felix C. Binder, Tiago Debarba, Emanuel Schwarzhans, Fabien Clivaz, and Marcus Huber

Abstract. Thermodynamics connects our knowledge of the world to our capability to manipulate and thus to control it. This crucial role of control is exemplified by the third law of thermodynamics, Nernst’s unattainability principle, which states that infinite resources are required to cool a system to absolute zero temperature. But what are these resources and how should they be utilised? And how does this relate to Landauer’s principle that famously connects information and thermodynamics? We answer these questions by providing a framework for identifying the resources that enable the creation of pure quantum states. We show that perfect cooling is possible with Landauer energy cost given infinite time or control complexity. However, such optimal protocols require complex unitaries generated by an external work source. Restricting to unitaries that can be run solely via a heat engine, we derive a novel Carnot-Landauer limit, along with protocols for its saturation. This generalises Landauer’s principle to a fully thermodynamic setting, leading to a unification with the third law and emphasises the importance of control in quantum thermodynamics.

Under review. Manuscript submitted on 25 Jan 2022.

[arXiv:2106.05151](https://arxiv.org/abs/2106.05151)

*,† denotes equal contributions.

Author Contribution

In this work, the doctoral candidate contributed significantly to the conception and formulation of the theoretical framework and methods, the proofs of the main results, the writing and revising of the manuscript, and the organisation and supervision of the project overall. In particular, the main technical contributions of the doctoral candidate were the development and application of both the coherent and incoherent control paradigms to address the task of cooling; the demonstration of all finite-dimensional asymptotic cooling protocols in both settings; the derivation of the Carnot-Landauer limit; the development of a thermodynamically meaningful notion of control complexity and the derivation of relevant conditions upon it regarding the ability to perfectly cool at minimal energy cost; and the comparison of cooling paradigms and resources for imperfect (i.e., finite resource) cooling.

1.1 Introduction

What is the cost of creating a pure state? Pure states appear as ubiquitous idealisations in quantum information processing and preparing them with high fidelity is essential for quantum technologies such as reliable quantum communication [187, 188], high-precision quantum parameter estimation [189–191], and fault-tolerant quantum computation [51, 64]. Fundamentally, pure states are prerequisites for ideal measurements [192] and precise timekeeping [58]. To answer the above question, one could turn to Landauer’s principle, stating that erasing a bit of information has an energy cost of at least $k_B T \log(2)$ [3]. Alternatively, one could consult Nernst’s unattainability principle (the third law of thermodynamics) [40], stating that cooling a physical system to its ground state requires diverging resources. At the outset, it seems that these statements are at odds with one another. However, Landauer’s protocol requires infinite time, thus identifying time as a resource according to the third law [42, 43, 193–195]. *Does this mean either infinite energy or time are needed to prepare a pure state?*

The perhaps surprising answer we give here is: *No*. We show that finite energy and time suffice to perfectly cool any quantum system and we identify the previously hidden resource—*control complexity*—that must diverge (in the spirit of Nernst’s principle) to do so. Intuitively, a good proxy for control complexity is the effective dimension of the system-machine interaction in a given protocol. The ultimate limit on the energetic cost of cooling is still provided by the Landauer limit, but in order to achieve it, either time or control complexity must diverge.

At the same time, heat fluctuations and short coherence times in quantum technologies [50] demand that both energy and time are not only finite, but minimal. Therefore, in addition to proving the necessity of diverging control complexity for perfect cooling with minimal time and energy, we develop protocols that saturate the ultimate limits. We demonstrate that mitigating dissipation comes at the practical cost of controlling fine-tuned interactions that require a *coherent* external work source, i.e., a quantum battery [54, 55, 196–198]. From a thermodynamic perspective, this may seem unsatisfactory: Nonequilibrium resources imply that the total system is not closed, and the optimal protocol (saturating the Landauer bound) is reminiscent of a Maxwellian demon with perfect control.

Accordingly, we also consider an *incoherent* control setting restricted to global energy-conserving unitaries with a heat bath as thermodynamic energy source. This setting corresponds to minimal control, where interactions need only be switched on and off to generate transformations, i.e., a heat engine drives the dynamics [60–63, 199]. The incoherent-control setting is therefore thermodynamically consistent inasmuch as both the machine state is assumed to be thermal (and to rethermalise between control steps) *and* the permitted control operations are those implementable solely via a heat engine.

In this paradigm, we show that the Landauer bound is not attainable, derive a novel limit—which we dub the *Carnot-Landauer* bound—and construct protocols that saturate it. The Carnot-Landauer bound follows from an equality phrased in terms of entropic and energetic quantities that must hold for any state transformation in the incoherent control paradigm; in this sense, the Carnot-Landauer equality adapts the equality version of Landauer’s principle developed in Ref. [200] to a fully (quantum) thermodynamic setting.

Our work thus both generalises Landauer’s erasure principle and, at the same time, unifies it with the laws of thermodynamics. By accounting for control complexity, we emphasise a crucial resource that is oftentimes overlooked but, as we show, must be taken into account for any operationally meaningful theory of thermodynamics. Here, we focus on the asymptotic setting that allows us to connect this resource with Nernst’s unattainability principle; beyond the asymptotic case, the gained insights also open the door to a better understanding of the intricate relationship between energy, time and control complexity when all resources are finite, which will be crucial for practical applications; we provide a preliminary analysis to this end. Lastly, our protocols saturating the Carnot-Landauer bound pave the way for thermodynamically-driven (i.e., minimal control) quantum technologies.

1.2 Overview & Summary of Results

There are two types of thermodynamic laws: Those, like the second law, that bound (changes of) characteristic quantities during thermodynamic processes, and those, like the third law, which state the impossibility of certain tasks. Landauer’s principle is of the former kind (indeed, it can be rephrased as a version of the second law), associating a minimal heat dissipation to any logically irreversible process, thereby placing a fundamental limit on the energy cost of computation. The paradigmatic logically irreversible process is that of erasing information, i.e., resetting an arbitrary state to a blank register. In other words, *perfectly cooling* a system to the ground state, or more generally, preparing a pure state.

Nernst’s unattainability principle is of the latter kind, stating that perfectly cooling a system requires diverging resources. The resources typically considered are energy and time: On the one hand, perfect cooling can be achieved in finite time at the expense of an energy cost that diverges as the ground state is approached; on the other hand, the energy cost can be minimised by implementing a quasistatic process that saturates the Landauer limit but takes infinitely long.

These two types of thermodynamic laws are intimately related, but details of their interplay have remained elusive: Under which conditions can the Landauer bound be saturated and what are the minimal resources required to do so? Which protocols asymptotically create pure states with given (diverging) resources? What type of control




Figure 1.1: Thermodynamic Framework. We consider the task of cooling a quantum system in two extremal control scenarios, with each step of both paradigms comprising two primitives. The top panel depicts the coherent-control scenario: In the control step (left), an agent can use a work source \mathcal{W} to implement any global unitary on the system \mathcal{S} and machine \mathcal{M} , which both begin thermal at inverse temperature β ; in cooling the target, energy and entropy is transferred to the machine. The machine then rethermalises with its environment (right), thereby dissipating the energy it gained in the control step. The bottom panel depicts the incoherent-control scenario: The machine is bipartitioned into a cold part at inverse temperature β and a hot part at inverse temperature $\beta_H < \beta$. In the control step, the agent switches on an interaction between the three systems, represented by a global energy-conserving unitary U_{EC} . In the rethermalisation step, the interaction is turned off and both subsystems of the machine rethermalise to their respective initial temperatures; the hot part draws energy from the heat bath while the cold part dissipates heat to its environment. In both paradigms, we quantify the control complexity as the effective dimension accessed by the unitary operation in a given control step (i.e., the dimension of the system-machine Hilbert space upon which the unitary acts non-trivially).

do such protocols require and how difficult are they to implement? We address these questions by considering the task of cooling a quantum system in two extremal control paradigms (see Fig. 1.1): One driven by a *coherent* work source and the other by an *incoherent* heat engine. Within this context we establish three main results:

1. Perfect cooling is possible with coherent control provided either energy, time, or control complexity diverge. In particular, it is possible in finite time and at Landauer energy cost with diverging control complexity.
2. Perfect cooling is possible with incoherent control, i.e., with a heat engine, provided either time or control complexity diverge. On the other hand, it is impossible with both finite time and control complexity regardless of the amount of energy drawn from the heat bath.
3. No process driven by a finite-temperature heat engine can (perfectly) cool a quantum system at the Landauer limit. Nonetheless, the Carnot-Landauer limit, which we introduce here (as a consequence of a stronger equality), can be saturated for any heat bath, given either diverging time or control complexity.

In the following, we discuss each of these results in turn in more detail and provide a systematic study concerning the asymptotic interplay of energy, time and control complexity as thermodynamic resources, as well as develop insight into the finite-resource regime for some special cases.

1.2.1 Perfect Cooling with Coherent Control (Work Source)

We begin by considering cooling with coherently-controlled resources (see Fig. 1.1, top panel). We first analyse energy, time and control complexity as resources that can be traded off against one another in order to optimise cooling performance, before focusing on the role of control complexity. Lastly, we consider the task of cooling systems (to finite temperatures) with all resources restricted to be finite.

Energy, Time, and Control Complexity as Resources

Whereas Landauer’s limit sets the minimum heat that must be dissipated—and thereby the minimum energy cost—for cooling any physical system, the third law makes no specification that energy must be the resource minimised (or that time must diverge). One might instead consider using a source of unbounded energy to perfectly cool a system as quickly as possible. Additionally, control complexity plays an important role as a resource, inasmuch as its divergence permits perfect cooling at the Landauer limit in finite time. As summarised in Table 1.1, we present coherently-controlled protocols that perfectly cool an arbitrary finite-dimensional target system using thermal machines when any one of the resources—energy, time or control complexity—diverges; moreover, the resources that are kept finite saturate protocol-independent ultimate bounds.

The cooling protocol using diverging energy is the simplest. Here, one exchanges all populations of the target system with those of a thermal machine with suitably large energy gaps to sufficiently concentrate the initial machine population in the ground state of the target. This exchange requires a single system-machine unitary and is of finite complexity (in a sense discussed below). Nonetheless, the energy drawn from the work source in this protocol diverges; indeed, as we show, any protocol that cools perfectly with both finite time and control complexity requires diverging energy.

Instead, one might wish to minimise the energy cost and saturate the Landauer bound. In the coherent control setting, for any transformation $\varrho_S \mapsto \varrho'_S$ enacted via a unitary interaction with a thermal machine at inverse temperature β , the dissipated heat $\Delta E_{\mathcal{M}}$ is given by [200]:

$$\beta\Delta E_{\mathcal{M}} = \tilde{\Delta}S_S + I(\mathcal{S} : \mathcal{M})_{\varrho'_{S,\mathcal{M}}} + D(\varrho'_{\mathcal{M}}\|\varrho_{\mathcal{M}}). \quad (1.1)$$

Here, the r.h.s. involves the (von Neumann) entropy decrease $\tilde{\Delta}S_S$ of the system, the mutual information $I(\mathcal{S} : \mathcal{M})_{\varrho'_{S,\mathcal{M}}}$ built up between system and machine, and the relative

	Energy	Time	Complexity
Qudit	$\rightarrow \infty$	1	$\frac{1}{2}d(d-1)$
	Landauer	$\rightarrow \infty$	$\frac{1}{2}d(d-1)$
	Landauer	1	$\rightarrow \infty$
H. O.	$\rightarrow \infty$	1	$\rightarrow \infty$ (Gaussian)
	Landauer	$\rightarrow \infty$	$\rightarrow \infty$ (Gaussian)
	Finite ($>$ Landauer)	$\rightarrow \infty$	1 (Non-Gaussian)
	Landauer	1	$\rightarrow \infty$ (Gaussian)

Table 1.1: Coherent-Control Cooling Protocols for Finite-Dimensional (Qudit) and Harmonic Oscillator Systems. Landauer energy cost refers to saturation of Eq. (1.2) and complexity refers to the effective dimension (see Definition 1.1); time is measured as the number of unitary operations with a fixed complexity. In the qudit case, the system and machine dimensions are equal: $d_S = d_M =: d$.

entropy $D(\varrho'_M \| \varrho_M)$ of the final machine state with respect to the initial one; the latter two terms are non-negative and can be dropped to yield the Landauer bound:

$$\beta \Delta E_M \geq \tilde{\Delta} S_S. \quad (1.2)$$

We now focus on cooling protocols that saturate this limit. Given a diverging amount of time, the target system can be sequentially coupled with a machine of finite complexity that rethermalises between control steps in such a way that the final target system state is arbitrarily close to the ground state for any initial temperature. Alternatively, one can compress all the operations applied in the diverging-time protocol into one global unitary that achieves the same final states, thereby achieving perfect cooling at the Landauer limit in a single unit of time but with an infinitely complex interaction. That is, the diverging temporal resource of repeated interactions with a single, finite-size machine is replaced by a single interaction with a larger machine of diverging complexity (as described below). However, this particular way of constructing complex control protocols is not necessarily unique. It is thus natural to wonder if diverging control complexity is a generic feature necessary to achieve perfect cooling at the Landauer limit in unit time and indeed, how to quantify control complexity that is operationally meaningful between the extreme cases of being either very small or divergent, as we now turn to discuss.

Role of Control Complexity

To address this issue, we first provide protocol-independent structural conditions that must be fulfilled by the machine to enable (1) *perfect cooling* and (2) *cooling at Landauer cost*; combined, these independent conditions provide a necessary requirement, namely that the machine must have an unbounded spectrum (from above) and be infinite-dimensional (respectively) for the possibility of (3) *perfect cooling at the Landauer limit*.

Such properties of the machine Hamiltonian define the *structural complexity*, which sets the potential for how cool the target system can be made and at what energy cost. As the name suggests, this is entailed by the structure of the machine, e.g., the number of energy gaps and their arrangement, and as such provides a static notion of complexity. However, given a machine with particular structural complexity, one may not be able to utilise said potential due to constraints on the dynamics that can be implemented. For instance, one may be restricted to only two-body interactions, or operations involving only a few energy levels at a time. Assuming a sufficient structural complexity at hand, such constraints limit one from optimally manipulating the systems. Thus, the extent to which a machine’s potential is utilised depends on properties of the dynamics of a given protocol, i.e., the *control complexity*.

Although it is intuitive that a unitary coupling the system to many degrees of freedom of the machine should be considered complex, it is a priori unclear how to quantify control complexity in a manner that both:

1. Corresponds to our intuitive understanding of the word “complex”, meaning “difficult to implement”; and
2. Is consistent with Nernst’s third law in the sense that its divergence is necessary to reach a pure state (when all other considered resources are restricted to be finite).

Many notions of complexity put forth throughout the literature to capture the first point above do not necessarily satisfy the second, as we will discuss later. Here, we take the opposite approach and seek a *minimal* notion of complexity that is first and foremost consistent with the third law of thermodynamics, which we hope to develop further to incorporate the idea of quantifying how difficult a protocol is to implement.

As a first step in this direction, a good proxy measure of control complexity is the effective dimension of a unitary operation, i.e., the dimension of the subspace of the global Hilbert space upon which the unitary acts nontrivially:

Definition 1.1. The *effective dimension* is the minimum dimension of a subspace \mathcal{A} of the joint Hilbert space $\mathcal{H}_{S\mathcal{M}}$ in terms of which the unitary can be decomposed as $U_{S\mathcal{M}} = U_{\mathcal{A}} \oplus \mathbb{1}_{\mathcal{A}^\perp}$:

$$d^{\text{eff}} := \min \dim(\mathcal{A}) : U_{S\mathcal{M}} = U_{\mathcal{A}} \oplus \mathbb{1}_{\mathcal{A}^\perp}. \quad (1.3)$$

Intuitively, given any (sufficiently large) machine dimension, the effective dimension captures how much of the machine takes part in the controlled interaction. While any dynamics that requires a high amount of control must accordingly have large effective dimension, the converse does not necessarily follow: There exist dynamics with corresponding large (even infinite) effective dimensions (e.g., Gaussian operations on two harmonic oscillators, such as those enacted by a beam splitter) that are easily implementable and

do not require high levels of control, as we discuss further below. Nevertheless, using the definition above we show that any protocol achieving perfect cooling at the Landauer limit necessarily involves interactions between the target and infinitely many energy levels of the machine. In other words, no interaction restricted to a finite-dimensional subspace suffices.

The effective dimension therefore provides a minimal quantifier for control complexity: It is the quantity that must diverge in order to (perfectly) cool at minimal energy cost—thus, it satisfies the above point 2. Moreover, it requires no assumption on the underlying structure of the machine, with the results holding for either collections of finite-dimensional systems or harmonic oscillators. However, its drawback comes from the fact that it alone does not necessarily capture a notion of complexity that corresponds to what is always difficult to achieve in practice. In other words, its divergence is insufficient for optimal cooling, as we highlight by example. Put simply, the effective dimension does not necessarily satisfy the above point 1. Indeed, in all of the protocols that we present, the degrees of freedom of the machine must be individually addressed in a fine-tuned manner to permute populations optimally, demonstrating that an operationally meaningful notion of control complexity must take into account factors beyond the effective dimensionality accessed by an operation. Later in this article, we analyse a number of potential candidates for such a definition of control complexity and highlight various shortcomings. Lastly, we demonstrate ways of achieving these conditions with paradigmatic quantum systems.

Imperfect Cooling with Finite Resources

The above results set the ultimate limitations for cooling inasmuch as the protocols saturate optimal bounds by using diverging resources. In reality, however, any practical implementation is limited to having only finite resources at its disposal. According to the third law, a perfectly pure state cannot be achieved in this scenario. Nonetheless, one can prepare a state of finite temperature by investing said resources appropriately. In this finite-resource setting, the interplay between energy, time and control complexity is rather complicated. First, the cooling performance is stringent upon the chosen figure of merit for the notion of cool—the ground-state population, purity, average energy, or temperature of the nearest thermal state are all reasonable candidates, but they differ in general [82]. Second, the total amount of resources available bounds the reachable temperature in any given protocol. Third, the details of the protocol itself influence the energy cost of achieving a desired temperature. In other words, determining the optimal distribution of resources is an extremely difficult task in general and remains an open question.

We therefore focus here on the paradigmatic special case of cooling a qubit target system by increasing its ground-state population in order to highlight some salient points

regarding cooling to finite temperatures. First, we compare the finite performance of two distinct coherent control protocols that both asymptotically saturate the Landauer limit; nonetheless, at any finite time, their performance varies. The first protocol simply swaps the target qubit with one of a sequence of machine qubits whose energy gaps are distributed linearly; the second involves interacting the target with a high-dimensional machine with a particular degeneracy structure. Although the latter cannot be decomposed easily into a qubit circuit (thereby making it more difficult to implement in practice), one can compare the two protocols fairly by fixing the total (and effective) dimension to be equal, i.e., comparing the performance of the linear sequential qubit machine protocol after $N + 1$ qubits have been accessed with that of the latter protocol with machine dimension 2^{N+1} . In doing so, we see that the simpler former protocol outperforms the more difficult latter one in terms of the energy cost at finite times, emphasising the fact that difficulty in practice does not necessarily correspond to complexity as a thermodynamic resource. Additionally, we analyse the cooling rates at which energy and time can be traded off amongst each other in the linear qubit sequence protocol by deriving an analytic expression. Lastly, we compare the performance of a coherent and an incoherent control protocol that use a similar machine structure to achieve a desired final temperature. We see that the price one must pay for running the protocol via a heat engine is that either more steps or more complex operations are required to match the performance of the coherent control setting. This example serves to elucidate the connection between the two extremal control scenarios relevant for thermodynamics.

We now move to consider asymptotic cooling behaviour in the incoherent-control setting.

1.2.2 Perfect Cooling with Incoherent Control (Heat Engine)

The (asymptotic) results presented so far pertain to cooling with the only restriction being that the machines are initially thermal. In particular, there are no restrictions on the allowed unitaries. In general, the operations required for cooling are not energy-conserving and require an external work source. With respect to standard considerations of thermodynamics, this may seem somewhat unsatisfactory, as the joint system is, in the coherent setting, open to the universe. When quantifying thermodynamic resources, one typically restricts the permitted transformations to be energy-conserving, thereby closing the joint system and yielding a self-contained theory.

We therefore analyse protocols using energy-conserving unitaries. With this restriction, it is not possible in general to cool a target system with machines that are initially thermal at a single temperature, as was considered in the coherent-control paradigm [82]. Instead, cooling can be achieved by partitioning the machine into one cold subsystem \mathcal{C} that begins in equilibrium at inverse temperature β and another hot subsystem \mathcal{H} coupled

to a heat bath at inverse temperature $\beta_H < \beta$ [82, 83] (see Fig. 1.1, bottom panel). In other words, one uses a hot and cold bath to construct a heat engine that cools the target. As we demonstrate, perfect cooling can be achieved in this setting as pertinent resources diverge. However, the structure of the hot bath plays a crucial role regarding the resource requirements. In particular, we present a no-go theorem that states that perfect cooling with a heat engine using a single unitary of finite control complexity is impossible, even given diverging energy drawn from the hot bath. This result is in stark contrast to its counterpart in the coherent-control setting, where diverging energy is sufficient for perfect cooling. This highlights that the incoherent-control setting is a fundamentally distinct paradigm that must be considered independently.

1.2.3 The Carnot-Landauer Limit

In the incoherent-control setting, we derive an equality-form adaptation of Landauer's bound on the minimum heat dissipated (or, the minimum amount of energy drawn from the hot bath), which we dub the *Carnot-Landauer limit*. The generalisation of Eq. (1.1) to the incoherent-control setting is:

$$\Delta F_S^{(\beta)} + \eta \Delta E_{\mathcal{H}} = -\frac{1}{\beta} [\Delta S_S + \Delta S_C + \Delta S_{\mathcal{H}} + D(\varrho'_C || \varrho_C) + D(\varrho'_{\mathcal{H}} || \varrho_{\mathcal{H}})]. \quad (1.4)$$

Here, the l.h.s. involves the free-energy change of the target system $\Delta F_S^{(\beta)}$ and the Carnot factor $\eta := 1 - \frac{\beta_H}{\beta}$ of the heat engine, while the r.h.s. involves changes in entropy ΔS of all subsystems and relative entropy terms $D(\bullet || \bullet)$ pertaining to the changes of the hot and cold parts of the machine. Importantly, the sum of terms in parentheses on the right-hand side are non-negative and can be dropped to yield the Carnot-Landauer inequality:

$$\Delta F_S^{(\beta)} + \eta \Delta E_{\mathcal{H}} \leq 0. \quad (1.5)$$

This inequality poses the ultimate limitation for the energy cost of information processing in a fully thermodynamic setting.

By explicitly accounting for a thermal energy source to drive the protocol, we generalise Landauer's erasure principle and, at the same time, unify it with the laws of thermodynamics. The consideration of cooling vis-a-vis information erasure already adapts Landauer's principle to a more concrete physical setting by specifying a target system Hamiltonian. By further incorporating a heat-bath energy source and requiring overall energy conservation we connect the heat dissipated during a cooling process with the thermodynamic resource driving it. Analysing the resource costs in turn provides a unification of cooling (and, as such, erasure) with the laws of thermodynamics. For instance, the Carnot efficiency between the two baths determines the cooling efficiency, implying that Landauer-cost cooling is impossible within any resource theory of thermodynamics with finite-temperature heat baths. Moreover, considering an infinite-temperature heat bath

Heat Bath	Energy	Time	Complexity
$\beta_H \in [0, \beta]$	$\rightarrow \infty$	\times	\times
$\beta_H \in [0, \beta]$	Carnot-Landauer	$\rightarrow \infty$	1
$\beta_H \in [0, \beta]$	Carnot-Landauer	1	$\rightarrow \infty$

Table 1.2: Incoherent-Control Cooling Protocols for Finite-Dimensional Systems. Carnot-Landauer energy cost refers to saturation of Eq. (1.5), β is the initial inverse temperature of the target system and cold part of the machine, and β_H that of the hot part.

resource (which can be identified as a work source) one recovers the relation between work and free-energy difference that embodies the second law of thermodynamics. Alternatively, if one were to consider the incoherent-control scenario with any finite-temperature hot bath and *not* specify an energetic structure of the target system, i.e., not attribute a Hamiltonian to the target, the Landauer information erasure procedure fuelled by a heat engine is recovered. Taking both of these reductions together yields the regular Landauer bound, i.e., perfectly efficient information erasure where the heat dissipated into the environment is *exactly* that drawn from the energy source.

Due to the restriction to energy-conserving unitaries, it is a priori unclear if the Carnot-Landauer bound is attainable and, if so, how to attain it. This problem persists for the special case where the heat-bath temperature tends to infinity and the bound reduces to the standard Landauer limit. To shed light on this issue, we first present cooling protocols that saturate the Landauer bound with an infinite-temperature heat bath. We do so by fine-tuning the machine structure such that the desired cooling transitions are energy-conserving. Subsequently, we study the more general case of finite-temperature heat baths, where we present protocols that saturate the Carnot-Landauer limit [i.e., Eq. (1.5)] for arbitrary heat baths. As in the coherent-control setting, these protocols require either diverging time or control complexity. These results, summarised in Table 1.2, provide a comprehensive understanding of the resources required to perfectly cool at minimum energy cost in a setting that aligns with the resource theories of thermodynamics.

1.3 Main Results

To arrive at these conclusions, we introduce a framework that encompasses both cooling and erasure, providing a unified approach to Landauer’s principle and the laws of thermodynamics. Crucially, this allows us to track all energetic and entropic exchanges throughout the process and relate the heat dissipated into the surrounding environment (i.e., the central quantity of Landauer’s principle) [3, 12, 37, 200] to the resources that an agent utilises to perform a desired transformation (i.e., the costs relevant to the third law) [42, 43, 193–195].

Importantly, all of the protocols we present are constructive proofs in the sense that they prove that perfect cooling is possible under certain conditions (e.g., infinite energy, time, or control complexity), and not the inverse, i.e., that no protocol exists that achieves the same with finite resources. However, note that the existence of a finite-resource protocol that perfectly cools would violate the unattainability principle, implying that either it is incorrect or there is a hidden resource unaccounted for. By establishing such protocols, we show that the three resources can be traded-off amongst each other to approach perfect cooling as long as any one diverges. Whenever we say “achievable with diverging X ”, we mean that one can construct a procedure that takes the initial state to within arbitrarily small distance ϵ from the ground state, i.e., $\varrho_{\mathcal{S}} \rightarrow_{\epsilon} |0\rangle\langle 0|_{\mathcal{S}}$, for any suitably-chosen distance measure, and that X diverges in the limit $\epsilon \rightarrow 0$. Similarly, “at the (Carnot-)Landauer limit” should be interpreted as asymptotic saturation of said bound. The scenarios where energy is minimised correspond to those in which the (Carnot-)Landauer limit is asymptotically attained as either time or control complexity diverge.

Landauer energy cost refers to saturation of Eq. (1.2); Carnot-Landauer energy cost refers to saturation of Eq. (1.5). Both of these equations are derived in Appendix A.1. Here, we use the proxy of effective dimensionality to represent the control complexity (see Definition 1.1). The control complexity of a complete swap between a target system and a machine of the same dimension $d := d_{\mathcal{S}} = d_{\mathcal{M}}$ is $d^{\text{eff}} = \frac{1}{2}d(d-1)$ and that of any exchange of two energy levels (e.g., $|i, j\rangle \leftrightarrow |j, i\rangle$ or $|i, j+1, k\rangle \leftrightarrow |i+1, j, k+1\rangle$ as considered in the incoherent-control protocols) is $d^{\text{eff}} = 1$. We quantify the time duration of any given protocol by the number of required consecutive unitary operations with fixed complexity. While the physical time passing is also a function of the interaction strength and multipartite nature of the interactions giving rise to the respective unitaries [52, 53], it is clear that each operation takes finite time, and a diverging number of unitaries thus implies diverging physical time. In the incoherent-control setting, β refers to the initial inverse temperature of the target system and cold part of the machine, whereas β_H is the initial temperature of the hot part of the machine.

1.3.1 Framework

Consider a target system \mathcal{S} in an initial state $\varrho_{\mathcal{S}}$ described by a unit-trace, positive semidefinite operator with associated Hamiltonian $H_{\mathcal{S}}$. An auxiliary machine \mathcal{M} , initially uncorrelated with \mathcal{S} and in equilibrium with a reservoir at inverse temperature $\beta := \frac{1}{k_B T}$, is used to cool the target system. The initial state of \mathcal{M} is thus of Gibbs form

$$\varrho_{\mathcal{M}} = \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}}) := \frac{e^{-\beta H_{\mathcal{M}}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})}, \quad (1.6)$$

where $H_{\mathcal{M}}$ is the machine Hamiltonian and $\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}}) := \text{tr} [e^{-\beta H_{\mathcal{M}}}]$ its partition function. Throughout this article we only consider Hamiltonians with discrete spectra, i.e., with an associated separable Hilbert space that has a countable energy eigenbasis. Moreover, for the most part we consider finite-dimensional systems (or sequences thereof) and deal with infinite-dimensional systems separately.

As shown in Fig. 1.1, one step of a cooling process comprises two sub-procedures: First, a joint unitary is implemented during the *control* step; second, the machine *rethermalises* to the ambient temperature. A cooling *protocol* is determined by the initial conditions and any concatenation of such primitives.¹ We consider two extremal control paradigms corresponding to two classes of allowed global transformations. The *coherent control* paradigm permits arbitrary unitaries on \mathcal{SM} ; in general, these change the total energy but leave the global entropy invariant and thus require an external work source \mathcal{W} . At the other extreme is the *incoherent control* paradigm, where the energy source is a heat bath. Here, the machine \mathcal{M} is bipartitioned: One part, \mathcal{C} , is connected to a *cold* bath at inverse temperature β , which serves as a sink for all energy and entropy flows; the other, \mathcal{H} , is connected to a *hot* bath at inverse temperature $\beta_H \leq \beta$, which provides energy. The composite system \mathcal{SCH} is closed and thus global unitary transformations are restricted to be energy conserving. The temperature gradient causes a natural heat flow away from the hot bath, which carries maximal entropic change with it. Cooling protocols in this setting can be run with minimal external control, i.e., they only require switching on and off interactions. We first analyse the coherent-control scenario.

1.3.2 Coherent Control

In the coherent-control setting, a transformation $\varrho_S \rightarrow \varrho'_S$ is enacted via a unitary U on \mathcal{SM} involving a thermal machine $\varrho_{\mathcal{M}} = \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})$, i.e.,

$$\varrho'_S := \text{tr}_{\mathcal{M}} [U(\varrho_S \otimes \varrho_{\mathcal{M}})U^\dagger]. \quad (1.7)$$

For such a transformation, there are two energy costs contributing to the total energy change, which must be drawn from a work source \mathcal{W} . The first is the energy change of the target $\Delta E_S := \text{tr} [H_S(\varrho'_S - \varrho_S)]$; the second is that of the machine $\Delta E_{\mathcal{M}} := \text{tr} [H_{\mathcal{M}}(\varrho'_{\mathcal{M}} - \varrho_{\mathcal{M}})]$, where $\varrho'_{\mathcal{M}} := \text{tr}_S [U(\varrho_S \otimes \varrho_{\mathcal{M}})U^\dagger]$. The latter is associated with the heat dissipated into the environment and is given by [200]

$$\beta \Delta E_{\mathcal{M}} = \tilde{\Delta} S_S + I(\mathcal{S} : \mathcal{M})_{\varrho'_{\mathcal{SM}}} + D(\varrho'_{\mathcal{M}} \| \varrho_{\mathcal{M}}), \quad (1.8)$$

where $S(\varrho) := -\text{tr} [\varrho \log(\varrho)]$ is the von Neumann entropy, $\tilde{\Delta} S_A := S(\varrho_A) - S(\varrho'_A)$,² $I(\mathcal{A} : \mathcal{B})_{\varrho_{\mathcal{AB}}} := S(\varrho_A) + S(\varrho_B) - S(\varrho_{\mathcal{AB}})$ (with marginals $\varrho_{\mathcal{A}/\mathcal{B}} := \text{tr}_{\mathcal{B}/\mathcal{A}} [\varrho_{\mathcal{AB}}]$) is the mutual

¹One could refer to both \mathcal{M} and the transformations applied as the *machine* and call the system \mathcal{M} itself the working *medium* inasmuch as the latter passively facilitates the process, in line with conventional parlance; however, we use the terminology established in the pertinent literature.

²Note the differing sign conventions (denoted by the tilde) that we use for changes in energies, $\Delta E_{\mathcal{X}} := E'_{\mathcal{X}} - E_{\mathcal{X}}$, and in entropies, $\tilde{\Delta} S_{\mathcal{X}} := S_{\mathcal{X}} - S'_{\mathcal{X}}$, such that energy *increases* and entropy *decreases* are positive.

information between \mathcal{A} and \mathcal{B} , and $D(\varrho\|\sigma) := \text{tr}[\varrho \log(\varrho)] - \text{tr}[\varrho \log(\sigma)]$ is the relative entropy of ϱ with respect to σ , with $D(\varrho\|\sigma) := \infty$ if $\text{supp}[\varrho] \not\subseteq \text{supp}[\sigma]$. We derive Eq. (1.8) and its generalisation to the incoherent-control setting in Appendix A.1. The mutual information is non-negative and vanishes iff $\varrho_{\mathcal{AB}} = \varrho_{\mathcal{A}} \otimes \varrho_{\mathcal{B}}$; similarly, the relative entropy is non-negative and vanishes iff $\varrho = \sigma$. Dropping these terms leads to the Landauer bound [3]

$$\beta \Delta E_{\mathcal{M}} \geq \tilde{\Delta} S_{\mathcal{S}}. \quad (1.9)$$

The Landauer limit holds *independently* of the protocol implemented, i.e., it only assumes *some* unitary was applied to the target and thermal machine. For large machines, the dissipated heat is typically much greater than the energy change of the target; nonetheless, the contributions can be comparable at the microscopic scale. We will assume that the target begins in equilibrium with the reservoir at inverse temperature β , i.e., in the initial thermal state $\varrho_{\mathcal{S}} = \tau_{\mathcal{S}}(\beta, H_{\mathcal{S}})$, with no loss of generality since such a relaxation can be achieved for free (by swapping the target with a suitable part of the environment; however, see Ref. [201] for a discussion of initial state dependency of the bound). We will track all energetic and entropic quantities and refer to the asymptotic saturation of Eq. (1.9) with $\varrho'_{\mathcal{S}}$ pure as *perfect cooling at the Landauer limit*.

Equation (1.8) provides insight for understanding the conditions required for saturating the Landauer bound. Although for finite-dimensional machines only trivial processes of the form $U_{\mathcal{S}\mathcal{M}} = U_{\mathcal{S}} \otimes \mathbb{1}_{\mathcal{M}}$ saturate the Landauer limit [200], we show how it can be asymptotically saturated with nontrivial processes by considering diverging machine and interaction properties, as we elaborate on shortly. Any such process must asymptotically exhibit no correlations such that $I(\mathcal{S} : \mathcal{M})_{\varrho'_{\mathcal{S}\mathcal{M}}} \rightarrow 0$ and effectively not disturb the machine, i.e., yield $\varrho'_{\mathcal{M}} \rightarrow \varrho_{\mathcal{M}}$ such that $D(\varrho'_{\mathcal{M}}\|\varrho_{\mathcal{M}}) \rightarrow 0$. Indeed, any correlations created between initially thermal systems would come at the expense of an additional energetic cost [202–204] whose minimisation is a problem that has so far only been partially resolved [205]. However, it has been shown that for any rank non-decreasing process, there exists a thermal machine and joint unitary such that for any $\epsilon > 0$, the heat dissipated satisfies $\beta \Delta E_{\mathcal{M}} \leq \tilde{\Delta} S_{\mathcal{S}} + \epsilon$ [200], thereby saturating the Landauer limit. Here, we present protocols that asymptotically achieve both this and perfect cooling (in particular, effectively decrease the rank), and provide necessary conditions on the underlying resources required to do so.

From an alternate perspective, the third law makes no specification that the energy is the resource minimised (or that time must diverge), and so we also present a protocol that uses an unbounded source of energy to perfectly cool the target in unit time. Additionally, as we discuss, the concept of control complexity plays an important role as a resource, inasmuch as allowing for its divergence permits perfect cooling at the Landauer limit in unit time. The following thus provides a comprehensive analysis of cooling with respect

to a trinity of resources—energy, time, and control complexity—that can be traded off amongst each other.

1. *Diverging Energy.*—We first consider the situation in which time and control complexity are fixed to be finite, while the energy cost is allowed to diverge. Here, we present:

Theorem 1.1. *With diverging energy, any finite-dimensional quantum system can be perfectly cooled using a single interaction of finite complexity.*

In addition to being sufficient for perfect cooling with both finite time and control complexity (i.e., using an effectively finite-dimensional machine), diverging energy is also necessary. See Appendix A.2 for details.

2. *Diverging Time.*—We now present a protocol that uses a diverging number of operations of finite complexity to asymptotically attain perfect cooling at the Landauer limit [197, 200, 206]:

Theorem 1.2. *With diverging time, any finite-dimensional quantum system can be perfectly cooled at the Landauer limit via interactions of finite complexity.*

Sketch of proof.—We first show that any system can be cooled from $\varrho_S = \tau_S(\beta, H_S)$ to $\tau_S(\beta^*, H_S)$, with $\beta^* \geq \beta$, using only $\beta^{-1} \tilde{\Delta} S_S$ units of energy. Our proof is constructive in the sense that we provide a protocol that achieves the Landauer energy cost as the number of operations diverges. The individual interactions in this protocol are of finite control complexity as they simply swap the target system with one of a sequence of thermal machines with increasing energy gaps. In this way, the final state $\tau_S(\beta^*, H_S)$ can be made to be arbitrarily close to $|0\rangle\langle 0|_S$ for any initial temperature. \square

The proof is presented in Appendix A.3, along with a more detailed dimension-dependent energy cost function for the special case of equally spaced Hamiltonians.

3. *Diverging Control Complexity.*—By reconsidering the diverging-time protocol above, a trade-off can be made between time and control complexity. As illustrated in Fig. 1.2, one can consider all of the operations $\{U_k = e^{-iH_k t_k}\}_{k=1, \dots, N}$ required in said protocol to make up one single joint interaction $U_{\text{tot}} := \lim_{N \rightarrow \infty} \prod_{k=1}^N U_k = e^{-iH_{\text{tot}} t_{\text{tot}}}$ acting on a larger machine, thus setting the time required to be unity (in terms of the number of operations before the machine rethermalises). In other words, for any finite number N of unitary transformations U_k , there exists a total Hamiltonian $H_{\text{tot}}^{(N)}$ and a finite time t_N that generates the overall transformation $U_{\text{tot}}^{(N)} := \prod_{k=1}^N U_k$; since t_N is finite, we can set it equal to one without loss of generality by rescaling the Hamiltonian as $\tilde{H}_{\text{tot}}^{(N)} = t_N H_{\text{tot}}^{(N)}$. Here, we refer to the limit $N \rightarrow \infty$ as diverging control complexity. Compressing a diverging number of finite-complexity operations thus yields a protocol of diverging control complexity. The fact that there exists such an operation that minimises both the time and

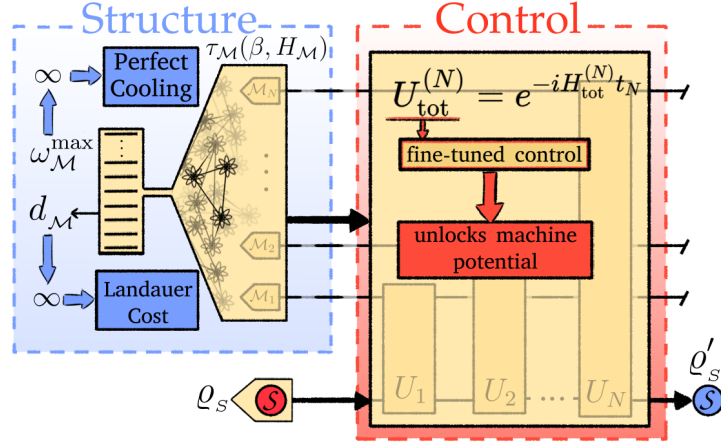


Figure 1.2: Notions of Complexity. We consider structural (left) and control complexity (right). Structural complexity concerns properties of the machine Hamiltonian. For perfect cooling it is necessary that the largest energy gap diverges [see Eq. (1.10)]. Moreover, an infinite-dimensional machine with particular energy-level structure is required for saturation of the Landauer bound. Control complexity refers to properties of the unitary that represents a protocol. The yellow box in the foreground represents a unitary U involving the entire machine, whereas the smaller yellow columns in the background represent a potential decomposition (e.g., of the diverging-time protocol) into unitaries U_i involving certain subspaces of the overall machine. Not only must the target system interact with all levels of an infinite-dimensional machine for Landauer-cost cooling, it must do so in a fine-tuned way.

energy requirements follows from our constructive proof of Theorem 1.2. We therefore have:

Corollary 1.1. *With diverging control complexity, any finite-dimensional quantum system can be perfectly cooled at the Landauer limit in unit time.*

1.3.3 Notions of Complexity

Although the protocol described above has diverging control complexity by construction, one need not construct complex protocols in this way, and so the natural concern becomes understanding the generic features that enable perfect cooling at the Landauer limit in unit time. We provide a detailed study in Appendix A.4, and here summarise the key methods.

Structural Complexity

We split the consideration of complexity into two parts: First, the protocol-independent *structural* conditions that must be fulfilled by the machine and, second, the dynamic *control complexity* properties of the interaction that implements a given protocol. Regarding the former, first note that one can lower-bound the smallest eigenvalue λ_{\min} of the final state ϱ'_S (and hence how cold the system can become) after *any* unitary interaction with a thermal machine by [200]

$$\lambda_{\min}(\varrho'_S) \geq e^{-\beta\omega_{\mathcal{M}}^{\max}} \lambda_{\min}(\varrho_S), \quad (1.10)$$

where $\omega_{\mathcal{M}}^{\max} := \max_{i,j} |\omega_j - \omega_i|$ denotes the largest energy gap of the machine Hamiltonian $H_{\mathcal{M}}$ with eigenvalues ω_i . It follows that perfect cooling is only possible under two conditions: Either the machine begins in a pure state ($\beta \rightarrow \infty$), or $H_{\mathcal{M}}$ is unbounded, i.e., $\omega_{\mathcal{M}}^{\max} \rightarrow \infty$. Requiring $\beta < \infty$, a diverging energy gap in the machine Hamiltonian is thus a necessary structural condition for perfect cooling. Independently, another condition required to saturate the Landauer limit can be derived for any amount of cooling: In Ref. [200], it was shown that for any finite-dimensional machine, there are correction terms to the Landauer bound which imply that it cannot be saturated; these terms only vanish in the limit where the machine dimension diverges.

We thus have two independent necessary conditions on the structure of the machine that must be asymptotically fulfilled to achieve relevant goals for cooling: The former is required for perfect cooling; the latter for cooling at the Landauer limit. Together, these conditions imply the following:

Corollary 1.2. *To perfectly cool a target system with energy cost at the Landauer limit using a thermal machine $\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})$, the machine must be infinite-dimensional and $\omega_{\mathcal{M}}^{\max}$, the maximal eigenvalue of $H_{\mathcal{M}}$, must diverge.*

The unbounded structural properties of the machine support the *possibility* for perfect cooling at the Landauer limit; we now move to focus on the control properties of the interaction that *realise* said potential (see Fig. 1.2). This leads to the distinct notion of control complexity, which differentiates between protocols that access the machine in a more or less complex manner. The structural complexity properties are protocol-independent and related to the energy spectrum and dimensionality of the machine, whereas the control complexity concerns properties of the unitary that represents a particular protocol.

Control Complexity

We begin by demonstrating that the effective dimension (nontrivially) accessed by a unitary (see Definition 1.1) must diverge to achieve perfect cooling at the Landauer limit, thereby providing a good proxy for control complexity in the sense that it aligns with Nernst's third law. Intuitively, the effective dimension of a unitary operation is the dimension of the subspace of the global Hilbert space upon which the unitary acts nontrivially, in other words the part of the joint space that is actually accessed by the control protocol. This quantity can be computed by considering a given cooling protocol and finite unit of time T (which we can set equal to unity without loss of generality) with respect to which the target and total machine transform unitarily by decomposing the Hamiltonian in $U_{S\mathcal{M}} = e^{-iH_{S\mathcal{M}}T}$ in terms of local and interaction terms, i.e., $H_{S\mathcal{M}} = H_S \otimes \mathbb{1}_{\mathcal{M}} + \mathbb{1}_S \otimes H_{\mathcal{M}} + H_{\text{int}}$. The effective dimension then corresponds to $\text{rank}(H_{\text{int}})$. With this definition at hand, we have:

Theorem 1.3. *The unitary representing a cooling protocol that saturates the Landauer limit must act nontrivially on an infinite-dimensional subspace of $\text{supp}(H_{\mathcal{M}})$. This implies $d^{\text{eff}} \rightarrow \infty$.*

Intuitively, we show that if a protocol only accesses a finite-dimensional subspace of the machine, then the machine is effectively finite-dimensional inasmuch as a suitable replacement can be made while keeping all quantities relevant for cooling invariant. Invoking the main result of Ref. [200] then implies that there are finite-dimensional correction terms such that the Landauer limit cannot be saturated.

However, accessing an infinite-dimensional machine subspace is far from sufficient for reaching the Landauer limit. This motivates searching for a more detailed notion of control complexity that takes the energy-level structure of the machine into account. The manifestation of such control seems to be system-dependent, precluding our ability (so far) to present a universal quantifier of control complexity, as we elaborate upon below. Thus, even though further conditions need to be met to achieve perfect cooling at minimal energy cost in unit time (as we discuss below), whenever we talk of an operation with finite control complexity, we mean those represented by a unitary that acts (nontrivially) on a finite-dimensional subspace of the target system and machine. In contrast, by diverging control complexity, we mean a unitary that couples the target (nontrivially) to a full basis of the machine's Hilbert space, whose dimension diverges.

The protocols that we present that achieve perfect cooling at Landauer cost make use of machines and interactions with a far more complicated structure than suggested by the necessary condition of infinite effective dimensionality. In particular, the interactions couple the target system to a diverging number of subspaces of the machine corresponding to distinct energy gaps. Moreover, there are a diverging number of energy levels of the machine both above and below the first excited level of the target. These observations highlight that fine-tuned control plays an important role. Indeed, both the final temperature of the target as well as the energy cost required to achieve this depends upon how the global eigenvalues are permuted via the cooling process. First, how cool the target becomes depends on the sum of the eigenvalues that are placed into the subspace spanned by the ground state. Second, for any fixed amount of cooling, the energy cost depends on the constrained distribution of eigenvalues within the machine. Thus, in general, the optimal permutation of eigenvalues depends upon properties of both the target and machine. To highlight this, we consider the task of cooling a maximally mixed target system with the additional constraint that the operation implemented lowers the temperature as much as possible. This allows us to derive a closed-form expression for the distribution of machine eigenvalues alone that must be asymptotically satisfied as the machine dimension diverges. Although global fine-tuning conditions can be derived similarly for any particular setup considered, a generic set of necessary control complexity conditions (beyond diverging

effective dimensionality) remains an open problem. This concept is even more important in the case where all resources are finite, as particular structures of machines and the types of interactions permitted play a crucial role in both how much time or energy is spent cooling a system and how cold the system can ultimately become (see, e.g., Refs. [83, 207, 208]).

Finally, as a representative for infinite-dimensional systems, we treat harmonic oscillator target systems separately in Appendix A.5. In the infinite-dimensional setting, the difficulty of implementing an operation is often related to the polynomial degree of its generators. Here, we see some friction with respect to Definition 1.1: A generic Gaussian unitary operation (i.e., one generated by a Hamiltonian at most quadratic in the mode operators) between a harmonic oscillator target and machine already implies infinite effective dimensionality. In light of this, we first construct a protocol that achieves perfect cooling at the Landauer limit with diverging time using only sequences of Gaussian operations [i.e., those typically considered to be practically easily implementable (cf. Refs. [54, 209]), but nonetheless with infinite effective dimensionality according to Eq. (1.3)]. In contrast, we then present a protocol that demonstrates that perfect cooling is possible given diverging time and operations acting on only a finite effective dimensionality (i.e., using non-Gaussian operations), with a finite energy cost that is greater than the Landauer limit; whether or not a protocol that saturates the Landauer limit exists in this setting remains an open question. Lastly, we present a unit-time perfect cooling protocol with diverging control complexity that can nonetheless be realised using only Gaussian operations.

The discussion above illustrates some key challenges in defining a measure of control complexity that satisfies natural desiderata: Such a measure should correspond to the difficulty of implementing operations in practice and simultaneously cover all possible physical platforms, including finite-dimensional systems such as, e.g., specific optical transitions of electrons in the shell of trapped ions, and infinite-dimensional systems such as the state-space specific modes of the electromagnetic field. The effective dimension that we use as a proxy manages to cover all such systems and provides a rigorous mathematical criterion that every physical protocol will necessarily have to fulfil in order to cool at minimal energy cost. As we have seen, however, infinite effective dimension is insufficient for cooling at the Landauer limit and it may not be all that difficult to achieve in continuous-variable setups. This begs the question of how this minimal definition of complexity can be extended in order to more faithfully represent what is difficult to achieve in practice. Intuitively, a property that seems to be important in lowering the energy cost of cooling is that the energy levels of the machine are distributed in such a way that they (approximately) densely cover the interval $[\omega_1, \omega^*]$, where ω_1 is the first energy gap of the target system and ω^* is the maximal energy gap, which corresponds to setting the final achievable temperature of the system (for perfect cooling, note that one requires $\omega^* \rightarrow \infty$). Let us denote the number of distinct energy gaps in a (fixed)

interval as the *variety*. Whether or not the variety of energy gaps of the machine used must diverge—in line with the third law—to cool at the Landauer limit remains an open question.

1.3.4 Incoherent Control

We now analyse the energy–time–complexity trinity within the incoherent-control setting (see Fig. 1.1). Here, we focus on finite-dimensional systems and leave the analysis of infinite-dimensional ones to future work. As we will see, the structure of the hot bath plays a crucial role in the ability to cool perfectly, especially regarding the resource requirements.

In the incoherent-control setting, an adaptation of the (equality-form) Landauer bound on the minimum heat dissipated (or, as we phrase it here, the minimum amount of energy drawn from the hot bath) can be derived, which we dub the *Carnot-Landauer limit*:

Theorem 1.4. *Let $F_\beta(\varrho_X) := \text{tr}[H_X \varrho_X] - \beta^{-1}S(\varrho_X)$ be the free energy of a state ϱ_X with respect to a heat bath at inverse temperature β , $\Delta F_S^{(\beta)} := F_\beta(\varrho'_S) - F_\beta(\varrho_S)$, and let $\eta := 1 - \frac{\beta_H}{\beta} \in (0, 1)$ be the Carnot efficiency with respect to the hot and cold baths. In the incoherent-control setting, the quantity*

$$\Delta F_S^{(\beta)} + \eta \Delta E_{\mathcal{H}} = -\frac{1}{\beta} [\Delta S_S + \Delta S_C + \Delta S_{\mathcal{H}} + D(\varrho'_C || \varrho_C) + D(\varrho'_{\mathcal{H}} || \varrho_{\mathcal{H}})] \quad (1.11)$$

satisfies the inequality

$$\Delta F_S^{(\beta)} + \eta \Delta E_{\mathcal{H}} \leq 0. \quad (1.12)$$

Equation (1.12) holds due to the non-negativity of the sum of local entropy changes and the relative-entropy terms. The derivation is provided in Appendix A.1, where we also show that the usual Landauer bound is recovered in the limit of an infinite-temperature heat bath.

The incoherent-control setting is fundamentally distinct from the coherent-control setting in terms of what can (or cannot) be achieved with given resources. For instance, consider the case where one wishes to achieve perfect cooling in unit time and with finite control complexity with diverging energy cost. In the coherent-control setting, this task is possible in principle (see Theorem 1.1). On the other hand, in the incoherent-control setting, we have the following no-go theorem (see Appendix A.6 for a proof):

Theorem 1.5. *In the incoherent control scenario, it is not possible to perfectly cool any quantum system of finite dimension in unit time and with finite control complexity, even given diverging energy drawn from the hot bath, for any non-negative inverse temperature heat bath $\beta_H \in [0, \beta < \infty)$.*

This result follows from the fact that in the incoherent-control setting, the target system can only interact with subspaces of the joint hot-and-cold machine with respect to which it is energy degenerate. For any operation of fixed control complexity, there is always a finite amount of population remaining outside of the accessible subspace, implying that perfect cooling cannot be achieved, independent of the amount of energy drawn from the hot bath.

The above result emphasises the difference between coherent and incoherent controlling, which means that it is a priori unclear if the Carnot-Landauer bound is attainable and, if so, how to attain it. Indeed, the restriction to energy-conserving unitaries generally makes it difficult to tell if the ultimate bounds can be saturated in the incoherent-control setting, and which resources would be required to do so. We present a detailed study of cooling in the incoherent-control setting in Appendix A.6, where we prove the following results. We begin by demonstrating incoherent cooling protocols that saturate the Landauer bound in the regime where the heat-bath temperature goes to infinity. We do so by fine-tuning the machine structure such that the desired cooling transitions between the target system and the cold and hot parts of the machine are rendered energy conserving. In particular, we prove:

Theorem 1.6. *In the incoherent control scenario, for an infinite-temperature hot bath $\beta_H = 0$, any finite-dimensional system can be perfectly cooled at the Landauer limit with diverging time via interactions of finite control complexity. Similarly, the goal can be achieved in unit time with diverging control complexity.*

Following our analysis of infinite-temperature heat baths, we study the more general case of finite-temperature heat baths. In Appendix A.7, we detail cooling protocols that saturate the Carnot-Landauer limit for any finite-temperature heat bath. More precisely, we prove:

Theorem 1.7. *In the incoherent control scenario, for any finite-temperature hot bath $0 < \beta_H < \beta$, any finite-dimensional quantum system can be perfectly cooled at the Carnot-Landauer limit given diverging time via finite control complexity interactions. Similarly, the goal can be achieved in unit time with diverging control complexity.*

As in the coherent-control setting, these protocols use either diverging time or control complexity to asymptotically saturate the Carnot-Landauer bound. The results presented in this section therefore provide a comprehensive understanding of the resources required to perfectly cool at minimum energy cost in a setting that aligns with the resource theories of thermodynamics.

1.3.5 Imperfect Cooling with Finite Resources

Although throughout most of the article we focus on the asymptotic achievability of optimal cooling strategies, the protocols that we construct provide insight into how said asymptotic limits are approached. This facilitates a better understanding of the more practically relevant questions that are constrained when all resources are restricted to be finite: i) *How cold can the target system be made?* and ii) *at what energy cost?* In line with Nernst's third law, the answer to the former question cannot be perfectly cold (i.e., zero temperature). The answer depends upon how said resources are configured and utilised. For instance, given a single unitary interaction of finite complexity in the coherent-control setting, the ground-state population of the output state can be upper bounded in terms of the largest energy gap of the machine, ω_{\max} [see Eq. (1.10)]. On the other hand, supposing that one can reuse a single machine system multiple times, then as the number of operation steps increases, the ground-state population of the output state approaches $(1 + e^{-\beta\omega_{\max}})^{-1}$ from below [82]. There is clearly a trade-off relation here between time and complexity, and a systematic analysis of the rate at which these quantities can be traded off against one another warrants further investigation. Similarly, the energy cost to reach a desired final temperature also depends upon the distribution of resources, as we now examine.

Given access to a machine of a certain size (as measured by its dimension), one could ask: *What is the optimal configuration of machine energy spectrum and global unitary to cool a system as efficiently as possible?* Here, we compare two contrasting constructions for the cooling unitary in the coherent-control setting for a qubit target system (with energy gap ω_S)—both of which asymptotically achieve Landauer cost cooling, but whose finite behaviour differs. The first protocol considers a machine of N qubits whose energy gaps increase linearly from the first excited state energy level of the system $\omega_1 = \omega_S$ to some maximum energy level $\omega_N = \omega_{\max}$, which dictates the final achievable temperature. In this protocol, the target system is swapped sequentially with each of the N qubits in order of increasing energy gaps; we hence refer to it as the *linear qubit machine sequence*. The second protocol we consider is presented in full in Appendix A.4.4 and inspired by one presented in Ref. [200] (see Appendix D therein); we hence refer to it as the Reeb and Wolf (RW) protocol. Here, the global unitary acts on the system and a high-dimensional machine with an equally-spaced Hamiltonian whose degeneracy doubles with each increasing energy level, i.e., it has a singular ground state, a two-fold degenerate first excited state, a four-fold degenerate second excited state, and so on; the final energy level has an extra state so that the total dimension is 2^{N+1} (where N is the number of energy levels). In particular, the unitary performs the permutation that places the maximal amount of population in the ground state of the target system. Due to the structure of both protocols, one can make a fair comparison between them, contrasting the

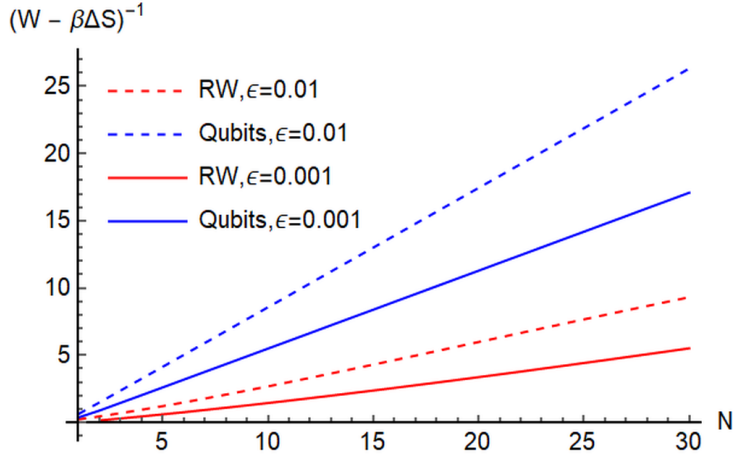


Figure 1.3: Imperfect Cooling Comparison. We compare the cooling performance of a degenerate qubit target system using either N machine qubits of linearly increasing energy accessed sequentially versus a single unitary on a 2^N dimensional machine, the latter being a finite adaptation of a protocol presented Ref. [200]. We set $\beta = 1$ and fix $1 - \epsilon$ to be the desired final ground-state population of the target. We plot the inverse of the excess work cost above the Landauer limit, $W - \beta \tilde{\Delta} S_S$, confirming that the surplus work cost in both cases scales with N^{-1} . Interestingly, we see that the protocol in which the target is sequentially swapped with machine qubits outperforms that which uses a high-dimensional unitary (at equal overall control complexity) in terms of energy cost required to reach a desired temperature.

single unitary on a 2^N -dimensional machine in the RW protocol versus the composition of N two-qubit swap unitaries in the linear machine sequence, i.e., such that both protocols access a machine of the same size overall.

As shown in Fig. 1.3, although both protocols asymptotically tend to the Landauer limit, their finite behaviour differs. Indeed, the work cost of the linear qubit machine sequence protocol outperforms that of the RW protocol. This is somewhat surprising, as the latter is a complex high-dimensional unitary whereas the former a composition of qubit swaps; although both protocols have the same effective dimension in this comparison overall, this highlights that difficulty in the lab setting need not correspond to resourcefulness in a thermodynamic sense. Indeed, developing optimal finite cooling strategies for arbitrary systems and machines is difficult in general and remains an important open question. Nonetheless, in Appendix A.8, we derive the rate of resource divergence of the sequential qubit protocol to further clarify the trade-off between time and energy for this protocol.

Finally, we contrast the two extremal thermodynamic paradigms considered by comparing the energy cost of a coherently-controlled cooling protocol to an incoherently-controlled one that achieves the same final ground-state population. Intuitively, the latter setting requires more resources to achieve the same performance as the former due to the fact that only energy-resonant subspaces can be accessed by the unitary, and hence only a subspace of the full machine is usable. This implies that a greater number of operations (of fixed control complexity) are required to achieve similar results as the coherent setting, as demonstrated in Appendix A.8 explicitly. Indeed, determining the optimal cooling protocols for a range of realistic assumptions remains a major open avenue.

1.4 Discussion

1.4.1 Relation to Previous Works

A vast amount of literature concerning quantum thermodynamics considers resource theories (see Refs. [89, 90] and references therein), whose central question is: *What transformations are possible given particular resources, and how can one quantify the value of a resource?* While this perspective sheds light on what is possible in principle, it does not per se concern itself with the potential implementation of said transformations. Yet, the unitary operations considered in a resource theory will themselves require certain resources to implement in practice. Focusing only on a resource-theoretic perspective would thus overlook the question: *How does one optimally use said resources?* Our results focus on this latter question and highlight the role of complexity in optimising resource use.

Concurrently, by considering arbitrary unitary operations (akin to our coherent-control paradigm without limitations on machine size) Refs. [197, 206] and [200], studied the potential saturation of the second law of thermodynamics and Landauer’s limit, respectively. Refs. [197] and [206] develop a similar protocol to our diverging time protocol in the context of work extraction and demonstrate its optimality for saturating the second law. However, these works do not discuss the practical viewpoint that the goal can be achieved in a smaller number of operations by allowing the latter to be more complex, as we emphasise. On the other hand, Ref. [200] considers the resources required for saturation of the Landauer limit and show an important result regarding structural complexity, namely that the machine must be infinite dimensional to cool at the Landauer limit. Our analysis regarding complexity begins here and continues to elucidate the key complexity properties that enhance the efficiency of a cooling protocol. In particular, we show that an infinite-dimensional machine is not sufficient unless the controlled unitary indeed accesses the entire machine. This leads to the notion of “effective dimension”, which provides a good proxy for control complexity that is consistent with Nernst’s third law for all types of quantum machines—from finite-dimensional systems to harmonic oscillators. Moreover, we highlight that the optimal interactions must be fine-tuned, i.e., they must couple the system to particular energy gaps of the machine in a specific configuration, paving the way for a more nuanced definition of control complexity that takes into account the complicated and precise level of control required. Lastly, we emphasise that the latter discussion concerns the coherent-control scenario, which is only one of the extremal control paradigms that we consider. In addition, we consider the task of cooling in a more thermodynamically consistent setting, namely the incoherent-control paradigm. There we derive the Carnot-Landauer equality and consequent inequality, which are adaptations of the Landauer equality [200] and inequality [3], respectively, where the protocol can only be run via a heat engine.

1.4.2 Conclusions & Outlook

The results of this work have wide-ranging implications. We have both generalised and unified Landauer’s bound with respect to the laws of thermodynamics. In particular, we have posed the ultimate limitations for cooling quantum systems or erasing quantum information in terms of resource costs and presented protocols that asymptotically saturate these limits. Indeed, while it is well-known that heat and time requirements must be minimised to combat the detrimental effects of fluctuation-induced errors and short decoherence times on quantum technologies [50], we have shown that this comes at a practical cost of greater control. In particular, we have demonstrated the necessity of implementing fine-tuned interactions involving a diverging number of energy levels to minimise energy and time costs, which serves to deliver a cautionary message: Control complexity must be accounted for to build operationally meaningful resource theories of quantum thermodynamics. This result posits the effective dimension accessed by a unitary protocol as a minimal quantifier of control complexity that is fully consistent with the third law of thermodynamics. Our analysis of the incoherent-control setting further provides pragmatic ultimate limitations for the scenario where minimal control is required, in the sense that all transformations are driven by thermodynamic energy and entropy flows between two heat baths, which could be viewed as a thermodynamically-driven quantum computer [12]. Nevertheless, the intricate relationship between various resources here will need to be further explored.

Looking forward, we believe it will be crucial to go beyond asymptotic limits. While Landauer erasure and the third law of thermodynamics conventionally deal with the creation of pure states, practical results would need to consider cooling to a finite temperature (i.e., creating approximately pure states) with a finite amount of invested resources [83, 207, 208]. In this context, the trade-off between time and control complexity will gain more practical relevance, as realistic quantum technologies have limited coherence times and interaction Hamiltonians are limited to few-body terms. Here, operational measures of control complexity that fit the envisioned experimental setup present an important challenge that must be overcome to apply our results across various platforms.

Our results strengthen the view that, in contrast to classical thermodynamics, the role of control is one of the most crucial issues to address before a true understanding of the limitations and potential of quantum machines is revealed. On the one hand, in classical systems, control is only ever achieved over few bulk degrees of freedom, whereas addressing and designing particular microstate control is within reach of current quantum technological platforms, offering additional routes towards operations enhanced by fine-tuned control. On the other hand, the cost of such control itself can quickly exceed the energy scale of the system, potentially rendering any perceived advantages a mirage. This is exacerbated by the fact that it is not possible to observe (measure) a quantum machine

without incurring significant additional thermodynamic costs [192, 210] and non-negligible back-action on the operation of the machine itself [211]. A fully developed theory of quantum thermodynamics would need to take these into account and we hope that our study sheds light on the role of control complexity in this endeavour.



Exponential Improvement for Quantum Cooling through Finite-Memory Effects

Philip Taranto, Faraj Bakhshinezhad, Philipp Schüttelkopf, Fabien Clivaz, and Marcus Huber

Abstract. Practical implementations of quantum technologies require preparation of states with a high degree of purity—or, in thermodynamic terms, very low temperatures. Given finite resources, the Third Law of thermodynamics prohibits perfect cooling; nonetheless, attainable upper bounds for the asymptotic ground state population of a system repeatedly interacting with quantum thermal machines have recently been derived. These bounds apply within a memoryless (Markovian) setting, in which each refrigeration step proceeds independently of those previous. Here, we expand this framework to study the effects of memory on quantum cooling. By introducing a memory mechanism through a generalised collision model that permits a Markovian embedding, we derive achievable bounds that provide an exponential advantage over the memoryless case. For qubits, our bound coincides with that of heat-bath algorithmic cooling, which our framework generalises to arbitrary dimensions. We lastly describe the adaptive step-wise optimal protocol that outperforms all standard procedures.

[Phys. Rev. Appl. 14, 054005 \(2020\)](#)

[arXiv:2004.00323](#)

Author Contribution

In this work, the doctoral candidate contributed significantly to the conception and formulation of the theoretical framework and methods, the proofs of the main results, the simulation and analysis of the data, the creation of figures, the writing and revising of the manuscript, and the organisation and supervision of the project overall. In particular, the main technical contributions of the doctoral candidate were the application of collision models with memory to the task of cooling; the demonstration of the exponential cooling advantage; the development of all presented protocols and the comparisons between them; and the discussion relating said work to Heat-Bath Algorithmic Cooling (**HBAC**).

2.1 Introduction

Cooling a physical system is a thermodynamic task of fundamental and practical importance [24, 42, 46, 81, 192, 194]. On the foundational side, the cooling potential is limited by the Third Law of thermodynamics, which posits the necessity of an infinite resource to be able to cool perfectly [43]. This resource is subject to trade-offs: Absolute zero is attainable in finite time given an infinitely-large environment; alternatively, given a finite energy source, one can only perfectly cool asymptotically. Practically, one cannot utilise an infinite resource, so the concern turns to: How cold can a system be prepared given resource constraints?

Formulating a theory with such constraints is typically scenario-dependent; nonetheless, one aims to develop theories that are widely applicable. For example, resource theories of quantum thermodynamics permit energy-conserving unitaries between the system and a thermal environment [89, 90]. Analysing the transformations for various environments and dynamical structures illuminates thermodynamic limitations.

Recent work has examined the task of quantum cooling in such a setting [82, 83]; the main result posits a universal bound for the ground state population of the system in the infinite-cycle limit. However, these results are derived in a memoryless (Markovian) setting, which is often not well-justified in experimental platforms where memory effects can affect the performance. For instance, Landauer’s principle [3] can be violated in the non-Markovian regime [76, 80].

A natural follow-up is to examine the role of memory in quantum cooling. Depending on the task and level of control, memory effects can have a detrimental or advantageous impact [67, 70, 72–75, 77–79, 114]; nonetheless, applications highlight the potential to be unlocked by controlling the memory via reservoir engineering [68, 69, 71]. Attempts to generalise thermodynamics to the non-Markovian setting include trajectory-based dynamical unravellings [212, 213] and those based on the operational process tensor formalism [128, 129, 214–218], among others [219, 220]. However, such general approaches typically obscure insight regarding the crucial resources; it is often unclear whether reported “quantum advantages” are due to genuinely quantum effects (e.g., coherence) or memory.

Here, we propose a mechanism for memory through a generalised collision model [91, 93, 94], which—while not fully general—permits fair comparison between various memory structures. We show that in the asymptotic limit, the memory depth of the protocol plays a critical role and leads to exponential improvement over the Markovian case. Our results coincide with the limits of heat-bath algorithmic cooling protocols [44, 84–88, 221, 222] for qubit targets and our framework both unifies and generalises this setting, applying to all system and environment structures.

2.2 Task: Cooling a Quantum System

A physical system is never isolated, which necessitates working within the theory of open systems, where the joint system-environment are closed, but environmental degrees of freedom are disregarded. Arbitrary environments permit perfect cooling with finite resources, as any physical transformation on a quantum system can be realised unitarily with a sufficiently-large environment; thus, further restrictions are necessary.

We consider a system, S , and environment, E , with Hamiltonians H_S and H_E , respectively. The system and environment begin uncorrelated and in equilibrium at inverse temperature $\beta := \frac{1}{k_B T}$. The joint system-environment evolves unitarily, with the system dynamics between the initial time and a later one t described by the dynamical map, $\varrho_S^{(t)}(\beta) := \Lambda^{(t)}[\tau_S^{(0)}(\beta)]$, defined such that:

$$\varrho_S^{(t)}(\beta) = \text{tr}_E \left[U^{(t)}(\tau_S^{(0)}(\beta) \otimes \tau_E^{(0)}(\beta)) U^{(t)\dagger} \right], \quad (2.1)$$

where $\tau_X(\beta)$ denotes a thermal state of X at inverse temperature β , i.e., $\tau_X(\beta) := \mathcal{Z}_X^{-1}(\beta) \exp(-\beta H_X)$ with partition function $\mathcal{Z}_X(\beta) := \text{tr}[\exp(-\beta H_X)]$.

The aim is to prepare $\varrho_S^{(t)}(\beta)$ as cold as possible. Cooling a system, however, can have several meanings: For one remaining in equilibrium, it could mean driving it to a thermal state of lower temperature; otherwise, one could consider increasing its ground state population or purity, or decreasing its entropy or energy. As such notions are generally nonequivalent, any study of cooling depends on the objective function [82]. We focus on achieving states that majorise all other potential states; this ensures optimisation of all Schur-convex/concave functions of the vector of populations ordered with respect to non-decreasing energy eigenstates, in particular all above notions of temperature.

2.3 Framework: Collision Models with Memory

Above we have described one step of a cooling protocol. In thermodynamic tasks, however, one is oftentimes interested in the multiple-cycle behaviour. Here, one faces a choice in how to proceed: One could implement each operation independently of those previous, i.e., completely refresh the environment between steps, leading to Markovian dynamics; or, one could temporally correlate the cycles, leading to non-Markovian dynamics. The main difficulty in treating the latter is that memory effects can arise in various ways: They can be the manifestation of initial correlations, recurring system-environment or intra-environment interactions; or any combination thereof. In any case, for multiple cycles, the dynamical map in Eq. (2.1) fails to completely describe the system dynamics, since system-environment correlations can influence later evolution, in contradistinction to the Markovian setting, where the environment is entirely forgotten between steps. In general, one must track all system-environment degrees of freedom to describe the system

evolution, which becomes unfeasible. Thus, we seek a framework that permits tractable memory and comparison between different memory structures.

We propose a microscopic model for the environment and its interactions with the system. We consider a d_S -dimensional system with $H_S = \sum_{i=0}^{d_S-1} E_i |i\rangle\langle i|_S$ and assume the environment comprises a number of identical units—which we call *machines*—each being a d_M -dimensional quantum system with associated Hamiltonian $H_M = \sum_{i=0}^{d_M-1} \mathcal{E}_i |i\rangle\langle i|_M$. We order Hamiltonians with respect to non-decreasing energies, and set $E_0 = \mathcal{E}_0 = 0$ and $\mathcal{E}_{\max} = \mathcal{E}_{d_M-1}$. Assuming that the dynamics proceeds via successive unitary “collisions” between the system and subsets of machines yields a collision model with memory.

The memory effects that arise from endowing such models with various dynamical structures have been examined: Considerations include initially correlated machines [223, 224], inter-machine [91, 95, 225–227] or repeated system-machine collisions [112, 115], or hybrid variations [92–94, 228]. In certain cases, the model exhibits finite-length memory [106, 130–132]. In the limit of many machines, the system is expected to interact with only mutually-exclusive subsets of machines; since any used machines never play a subsequent role, one yields a microscopic picture of Markovian evolution that gives a Lindbladian master equation in the continuous-time limit [229–232].

Although not fully general, this setting captures tractable non-Markovian dynamics. In this article, we will analyse the memory effects that arise from repeated system-machine interactions (see Fig. 2.1). More precisely, we consider k machines to interact with the system between timesteps, with some $\ell \leq k$ of these carrying memory forward; this reduces to a Markovian protocol involving k machines for $\ell = 0$. The assumptions are that the system and all machines begin uncorrelated, and there are no interactions between memory-carrying machines and fresh ones other than those involving the system. These are valid whenever the memory-carrying machines relax much slower than those that rethermalise between steps. We can vary the number of machines in each interaction, k ,¹ the number of memory carriers, ℓ , the initial temperature, β , and the Hamiltonians.

This generic framework applies to a wide range of protocols. For instance, one can compare adaptive strategies, where different unitaries are performed between steps, versus non-adaptive ones, where a fixed dynamics is repeated. Additionally, one can restrict the allowed unitaries, such as limiting the set from general “coherent” ones (that require an external energy source) to “incoherent” energy-conserving transformations (where the cooling resource is an additional hot bath) [83, 233]. Lastly, one could allow the memory structure itself to be adaptive, where k and ℓ vary between times; we do not consider this and instead focus on cooling limits for fixed structures. A choice of k and ℓ , along with the system and machine dimensions, determines the control complexity afforded to the

¹One could consider the restricted case of $(k+1)$ -partite system-machine interactions that are decomposable into sequences of p -partite interactions for $p \leq k$. We do not make this restriction and allow any multi-partite interaction to be genuinely so.




Figure 2.1: Collision Model with Memory. At each step, the system S interacts unitarily with k machines, of which ℓ carry forward memory (red lines). Here we illustrate $k = 2, \ell = 1$, with m the total number of machines used by timestep n .

experimenter: Intuitively, k is related to spatial complexity and ℓ to temporal complexity. We now compare the achievable cooling of a system for different memory structures.

2.4 Memory-Enhanced Cooling

The fundamental Markovian cooling bounds have been derived in Refs. [82, 83]. The optimally-cool system state at any finite time depends upon the energy-level structure between the system and machines and the level of control. However, in the asymptotic limit of Markovian operation, the vector of eigenvalues of the asymptotic state (in any aforementioned control paradigm) is majorised by that of

$$\varrho_S^*(\mathcal{E}_{\max}, \beta, k) = \sum_{n=0}^{d_S-1} \frac{e^{-\beta n k \mathcal{E}_{\max}}}{\mathbb{Z}_S(\beta, k \mathcal{E}_{\max})} |n\rangle\langle n|_S, \quad (2.2)$$

whenever the initial state $\tau_S^{(0)}(\beta)$ is majorised by $\varrho_S^*(\mathcal{E}_{\max}, \beta, k)$; here $\mathbb{Z}_X(\beta, \mathcal{E}) := \sum_{n=0}^{d_X-1} e^{-\beta n \mathcal{E}}$ is a quasi-partition function (depending only on the maximum energy gap of each machine, \mathcal{E}_{\max}). The state in Eq. (2.2) is attainable with coherent control, positing the ultimate Markovian cooling limit.

The intuition is that the optimal protocol reorders the eigenvalues of the system and relevant machines at each step such that the maximum population is placed into the ground state subspace of the system, the second largest into the first excited state subspace, and so on. When this cycle is repeated with fresh machines at each timestep, the asymptotic state looks as if it had interacted with only the qubit subspace of each machine with maximum energy difference. However, the result cannot immediately be extended to the non-Markovian regime, as its derivation relies on an inductive argument on the system state at each step; for non-Markovian dynamics, this cannot be expressed in terms of the previous state, posing a logical roadblock.

Whenever $\ell > 0$ the generalised collision model is non-Markovian. Nonetheless, a relevant result states that such non-Markovian collision models can be lifted to a Markovian dynamics on a larger state space [95]. For a system interacting with k machines at each

step, of which ℓ feed forward, the dynamics can be embedded into a Markovian one by considering the system and ℓ memory carriers as a unified system, which interacts at each step with $k - \ell$ fresh machines; such a process has *memory depth* ℓ . In Appendix B.1, we detail the Markovian embedding, which leads to the following results.

2.4.1 Asymptotic Cooling Advantage

We now present the universal cooling bound for the non-Markovian collision model in the infinite-cycle limit:

Theorem 2.1. *For any d_S -dimensional system interacting at each step with k identical d_M -dimensional machines, with ℓ of the machines (labelled L) used at each step carrying the memory forward, in the limit of infinitely many cycles:*

i) The ground state population of S is upper bounded by

$$p^*(\mathcal{E}_{\max}, \beta, k, \ell) = \left(\sum_{n=0}^{d_S-1} e^{-\beta n d_M^\ell (k-\ell) \mathcal{E}_{\max}} \right)^{-1}. \quad (2.3)$$

ii) The vector of eigenvalues of the output system state is majorised by that of the following attainable state

$$\varrho_S^*(\mathcal{E}_{\max}, \beta, k, \ell) = \sum_{n=0}^{d_S-1} \frac{e^{-\beta n d_M^\ell (k-\ell) \mathcal{E}_{\max}}}{\mathbb{Z}_S(\beta, d_M^\ell (k-\ell) \mathcal{E}_{\max})} |n\rangle\langle n|_S, \quad (2.4)$$

whenever the initial state $\tau_S(\beta) \otimes \tau_M(\beta)^{\otimes \ell}$ is majorised by

$$\varrho_{SL}^*(\mathcal{E}_{\max}, \beta, k, \ell) = \sum_{n=0}^{d_{SL}-1} \frac{e^{-\beta n (k-\ell) \mathcal{E}_{\max}}}{\mathbb{Z}_{SL}(\beta, (k-\ell) \mathcal{E}_{\max})} |n\rangle\langle n|_{SL}. \quad (2.5)$$

Sketch of Proof. We use the Markovian embedding to lift the non-Markovian dynamics of the target S to a Markovian process for the target-plus-memory carriers SL system, which interacts with $k - \ell$ fresh machines (which we label R) at each step. This implies that optimally cooling SL is necessary to optimally cool S . From Ref. [82], the asymptotically-optimal state of SL has the same eigenvalue distribution as Eq. (2.5), whenever the initial SL state is majorised by ϱ_{SL}^* , and is thus unitarily equivalent to it. As majorisation concerns all partial sums, given that initial condition, whatever protocol one chooses to cool S , the asymptotic state cannot be colder than ϱ_S^* (which is the coldest S state in the unitary orbit of ϱ_{SL}^*). This implies that the asymptotic ground state population is upper bounded by p^* . See Appendix B.2. \square

There are many noteworthy points: Firstly, the optimal ground state population is enhanced by d_M^ℓ compared to the Markovian case, highlighting the drastic role of memory; in particular, one achieves an exponential improvement in ℓ . Secondly, as the factors in Eq. (2.4) arise independently from various sources (i.e., S , L and R), the bound extends to the case where L is an arbitrary d_L -dimensional system and R an arbitrary d_M -dimensional




Figure 2.2: Cooling behaviour. We simulate the ground state population after m machines have been used in the step-wise optimal cooling protocol (see Theorem 2.2) for a qubit system and machines with fixed $\beta = 0.2$, $E_{\max} = 1$ and $\mathcal{E}_{\max} = 2$. The asymptotic hierarchy agrees with Corollary 2.1, whereas the complex short-term behaviour can exhibit crossovers.

system (with maximum energy gap \mathcal{E}'_{\max}), with $d_M^\ell \rightarrow d_L$ and $(k - \ell)\mathcal{E}_{\max} \rightarrow \mathcal{E}'_{\max}$. This clarifies that the asymptotic bound only depends on the dimension of L , not on its energy structure. Lastly, the asymptotic SL state of Eq. (2.5) is unitarily equivalent to a tensor product state that has Eq. (2.4) as its reduced state on S . Nonetheless, throughout the cooling protocol correlations build up, due to the finite-time dependence on the energy structures of the systems involved, before dying out asymptotically; in Appendix B.3, we explore the role of correlations in more detail.

Returning to Theorem 2.1, Eq. (2.4) allows us to compare limits for various k, ℓ, β and \mathcal{E}_{\max} (see Fig. 2.2):

Corollary 2.1. *The asymptotic hierarchy is determined via:*

$$\begin{aligned} \varrho_S^*(\mathcal{E}_{\max}, \beta, k, \ell) \prec (\succ) \varrho_S^*(\mathcal{E}'_{\max}, \beta', k', \ell') \\ \text{if } \beta(k - \ell)d_M^\ell \mathcal{E}_{\max} \leq (>) \beta'(k' - \ell')d_M^{\ell'} \mathcal{E}'_{\max}. \end{aligned} \quad (2.6)$$

2.4.2 Step-Wise Optimal Protocol

The bound is achievable and one protocol to do so reorders the global eigenspectrum at each step such that they are non-increasing with respect to non-decreasing energy eigenstates of SL . In the last step, the protocol additionally reorders the eigenvalues of the obtained SL state largest-to-smallest with respect to non-decreasing energy eigenstates of S . Precisely, at each step j the system and memory carriers are optimally cooled via a unitary $V_{SLR}^{(j)} : \varrho_{SLR}^{(j)} = V_{SLR}^{(j)} \varrho_{SLR}^{(j-1)} V_{SLR}^{(j)\dagger}$ that acts as²

$$\varrho_{SLR}^{(j)} = \sum_{\xi=0}^{d_S d_M^\ell - 1} \sum_{\omega=0}^{d_M^{k-\ell} - 1} \lambda_{\xi, d_M^{k-\ell} + \omega}^\downarrow |\xi\omega\rangle \langle \xi\omega|, \quad (2.7)$$

²Here, R refers to the fresh machines included at each step, with an implied identity map on all other systems.

where $|\xi\omega\rangle = |\xi\rangle_{SL} \otimes |\omega\rangle_R$ and λ^\downarrow denotes the eigenvalues of $\varrho_{SLR}^{(j-1)}$ in non-increasing order. This unitary dissipates maximal heat into the machines that play no subsequent role, and thus at any finite timestep j , the protocol has achieved the coldest SL state possible given its history, which is crucial for finite-time optimality. By implementing the sequence $\{V_{SLR}^{(j)}\}_{j=1,\dots,n}$, although the final (timestep n) SL state $\varrho_{SL}^{(n)}$ generically exhibits correlations, there always exists a unitary $W_{SL}^{(n)}$ that ensures S is optimally cool by further reordering the eigenvalues under the previous constraint; in the asymptotic limit, i.e., when $n \rightarrow \infty$, said unitary completely decorrelates SL (whereas at finite times, some correlation generically remains). However, while this strategy attains the optimally-cool S for any final timestep n , this protocol is not necessarily *step-wise* optimal.

To derive the step-wise optimal protocol, consider the unitary $W_{SL}^{(j)} : \sigma_{SL}^{(j)} = W_{SL}^{(j)} \varrho_{SL}^{(j)} W_{SL}^{(j)\dagger}$ that acts as

$$\sigma_{SL}^{(j)} = \sum_{\mu=0}^{d_S-1} \sum_{\nu=0}^{d_M^\ell-1} \lambda_{\mu \cdot d_M^\ell + \nu}^\downarrow |\mu\nu\rangle \langle \mu\nu|, \quad (2.8)$$

where $|\mu\nu\rangle = |\mu\rangle_S \otimes |\nu\rangle_L$ and λ^\downarrow here denotes the eigenvalues of $\varrho_{SL}^{(j)}$ in non-increasing order. $W_{SL}^{(j)}$ optimally cools S given any SL state by unitarily transferring maximal entropy towards L ; thus, if we apply $W_{SL}^{(j)}$ at each step j after having optimally cooled SL via $V_{SLR}^{(j)}$ until then, i.e., implement $U_{SLR}^{(j)} = W_{SL}^{(j)} V_{SLR}^{(j)}$, S is guaranteed to be optimally cool. This leads to the following, proven in the Appendix B.4, where we examine finite-time behaviour.

Theorem 2.2 (Step-wise optimal cooling protocol). *By applying $U_{SLR}^{(j)}$ described above at each step, the cooling protocol is step-wise optimal regarding the temperature of the system.*

2.5 Relation to Heat-Bath Algorithmic Cooling

Above we have derived the cooling limit in a controlled non-Markovian setting; through the Markovian embedding, we can further make direct connection with Heat-Bath Algorithmic Cooling (HBAC) [44, 84–88, 221, 222], the limitations of which align with our results for qubit targets. Here, one cools a “target” system by cooling a larger ensemble of “compression/refrigerant” systems³ via interactions with “reset” systems that rethermalise between steps. This permits better cooling than cooling the target alone with only reset systems as a resource; indeed, HBAC protocols are non-Markovian and a special case of our framework, which treats the compression/refrigerant systems as memory-carrying machines and the reset systems as fresh machines, as detailed below and in Appendix B.5.

i) Each of the target, memory carrier (compression/refrigerant), and reset systems can comprise multiple subsystems of arbitrary dimension, with arbitrary energy spectra and

³Together, the target and compression/refrigerant systems constitute what is often called the “computation” system.

initial temperatures, which determines the asymptotic hierarchy for different strategies. In contrast, many HBAC studies focus only on target and reset qubits [84, 87, 221]; although some consider qudit compression [86] and reset systems [85], no HBAC study has shown results pertaining to the general qudit-qudit-qudit case. ii) Our results are based on majorisation (as are those in Refs. [84, 85]), and therefore applicable to more general notions of cooling than the often-considered ground state population (e.g., in Refs. [44, 86, 221]), which crucially differ for high-dimensional systems [82]. This is important for quantum computing—for which cooling is a critical requirement—where high-dimensionality can simplify logical structures [234–236]. iii) Finally, our results extend the partner-pairing algorithm—introduced to maximise the ground state population of a qubit in Ref. [84]—to the most general setting (described above). The partner-pairing algorithm is step-wise optimal with a complexity that scales polynomially; our protocol achieves the same scaling, as the sorting required at each step can be achieved with a single operation. Although these operations depend on the global state at each step, in Appendix B.5 we present a simple robust algorithm (based on one presented in Ref. [87]) that uses only a fixed state-independent two-body interaction to reach the asymptotically-optimal state.

By contextualising HBAC within the framework of collision models with memory, our work provides both a unification and generalisation of HBAC. Moreover, our approach lends itself to modelling realistic HBAC experiments, where reset systems only partially thermalise, as considered in Ref. [222].

2.6 Conclusions

In this article, we have put forward a framework for consistently dealing with memory when cooling quantum systems; indeed, the generalised collision model proposed is versatile enough to analyse the role of memory in various thermodynamic tasks. In doing so, we have revealed the potential for exponential improvement in the reachable ground state population (and more general notions of cooling), yielding drastic enhancement already for modest memory depths. Through a Markovian embedding of our framework, we could connect our framework with HBAC, shedding new light on the latter as a particular class of non-Markovian dynamics. Our results can be read as a generalisation of HBAC applicable to arbitrary target and compression systems and bath spectra; by putting all HBAC protocols on an equal footing, our work opens the door to comparative studies that can now be made fairly. Moreover, we clarify the origin of the advantages that make HBAC so effective. Together with that of Refs. [82, 83], our work unifies HBAC with the resource theory of thermodynamics, as all results can be achieved either via coherent control or energy-conserving unitaries on enlarged systems.

The exponential improvement with respect to the memory carriers stands in contrast to the only linear enhancement in the number of rethermalising systems, highlighting the

importance of controllable memory. Thus, given the ability to perform $(k + 1)$ -partite interactions, having $\ell = k - 1$ of these systems as memory carriers is both the minimum requirement (as the system-and-memory must be open, otherwise any cooling ability rapidly diminishes) and, moreover, the optimal configuration. In particular, this implies that one can achieve the exponential advantage via interactions involving the system and memory carriers and only one additional reset system. Of course, if ℓ is large, the ability to implement complex many-body interactions presents a difficult challenge. To this end, we have developed an explicit protocol (see Appendix B.5) which necessitates only a fixed, two-body interaction to achieve the fundamental bound. In particular, this implies that the attainability of the optimal asymptotic ground state population does not require implementing highly non-local unitaries. In fact, this robust protocol applies to arbitrary dimensional systems for the target, memory and reset parts, which is significant because high-dimensional systems are becoming increasingly relevant for fault-tolerant quantum computing [234–236]—a major motivation for cooling quantum systems in the first place.

Our results consolidate the limits for quantum refrigeration in a setting with perfect control and high-quality isolation of the target and memory carriers. However, in most experimental scenarios, further challenges arise. We have assumed that the uncontrolled system-environment interactions are negligible compared to the controlled ones. For finite times, our results are reliable due to the exponential scaling, which makes it sufficient to run the protocol for a short time to approximate the asymptotics. An immediate concern is the impact of uncontrolled interactions: Either to model imperfect target isolation or to better understand the realistic asymptotics, as for infinite steps, rethermalisation of the target cannot always safely be neglected. Another assumption worth analysing is that of perfect control: To implement a unitary perfectly, one requires both precise clocks [237], which have their own thermodynamic costs [57, 238], and high control over the interaction terms. While this is plausible for quantum computing devices, other systems are more challenging to control, particularly those with multi-partite interactions that are only perturbatively accessible [239]. Lastly, a resource-theoretic approach has derived the minimum amount of energy required to implement transformations such as those considered here, which should also be accounted for [240]. Our framework lends itself to such pragmatic analyses of cooling; deriving similar bounds in realistic settings highlights the potential to elevate our results beyond fundamental limits and towards practical guidelines for quantum experiments.



Part II

Open Quantum Processes & Multi-Time Phenomena

When is a Non-Markovian Quantum Process Classical?

Simon Milz, Dario Egloff, [Philip Taranto](#), Thomas Theurer, Martin B. Plenio, Andrea Smirne, and Susana F. Huelga

Abstract. More than a century after the inception of quantum theory, the question of which traits and phenomena are fundamentally quantum remains under debate. Here we give an answer to this question for temporal processes which are probed sequentially by means of projective measurements of the same observable. Defining classical processes as those that can—in principle—be simulated by means of classical resources only, we fully characterise the set of such processes. Based on this characterisation, we show that for non-Markovian processes (i.e., processes with memory), the absence of coherence does not guarantee the classicality of observed phenomena and furthermore derive an experimentally and computationally accessible measure for non-classicality in the presence of memory. We then provide a direct connection between classicality and the vanishing of quantum discord between the evolving system and its environment. Finally, we demonstrate that—in contrast to the memoryless setting—in the non-Markovian case, there exist processes that are *genuinely quantum*, i.e., they display non-classical statistics independent of the measurement scheme that is employed to probe them.

[Phys. Rev. X **10**, 041049 \(2020\)](#)
[arXiv:1907.05807](#)

Author Contribution

In this work, the doctoral candidate contributed to the conception and formulation of the theoretical framework and methods, the proofs of the main results, and the writing and revising of the manuscript. In particular, the main technical contributions of the doctoral candidate were the extension of the notion of classicality to the setting of processes with memory; the revisiting of the memoryless setting in light of these more general considerations; the rephrasing of relevant conditions on the dynamics in terms of the Choi-Jamiołkowski isomorphism; the derivation of the structure of processes that look classical and the notion of Non-Discord-Generating-and-Detecting (NDGD) dynamics; and the presentation of the “genuinely quantum process”.

3.1 Introduction

Quantum coherence is considered to be one of the fundamental traits that distinguishes quantum from classical mechanics [241–243]. Beyond its mathematical deviation from classical theory, it plays an important role in the enhancement of quantum metrology tasks [186, 244], constitutes a fundamental requirement for many quantum algorithms [245, 246], and has been conjectured to be necessary for the formulation of efficient transport models in biology that are consistent with spectroscopic data [70, 247, 248]. Consequently, the resource theory of coherence [184, 185, 249–255] has been of tremendous interest in recent years, and has seen rapid development both on the theoretical as well as the experimental side [256].

Despite such progress and the growing wealth of accompanying evidence that links coherence to non-classical phenomena, the explicit connection between the two remains unclear and subject to active debate [257–261]. Put differently, the mere presence of coherence does not guarantee the existence of effects that cannot be explained on purely classical grounds, and an unambiguous relationship between coherence and non-classicality has not been established yet.

In order to provide such a connection, an operationally meaningful and clear-cut definition of classicality is crucial. One such possible definition is based on experimentally attainable quantities only, namely the joint probability distributions obtained from sequential measurements of an observable.¹ If these satisfy the Kolmogorov consistency conditions for all considered sets of measurement times—which provide the starting point for the formulation of the theory of classical stochastic processes [145, 146]—then they can, in principle, be explained by a fully classical model and there is therefore nothing inherently quantum about the observed phenomenon. If they do not, then there exists no underlying classical stochastic process that could lead to the observed joint probability distributions, and the corresponding process is considered non-classical. This characterisation of classicality is in the spirit of the derivation of Leggett-Garg inequalities, where, instead of classicality, non-invasiveness and macroscopic realism are put to the test [147, 148]. Indeed, any set of probability distributions that satisfies the Kolmogorov conditions does not violate the corresponding Leggett-Garg inequalities [149, 263].

Following this line of reasoning, and in a sense to be further specified later more precisely, in Ref. [170] a one-to-one connection was derived between the notion of classicality based on the Kolmogorov conditions and the coherence properties of the dynamics of Markovian (i.e., memoryless) quantum processes: Such a process is classical iff the corresponding dynamical propagators can never create coherence that can be detected at any later time. Thus, a direct relation between the mathematical notion of coherence and an operationally

¹For a different demarcation line between classical and quantum physics, based on the memory cost required to simulate a given process, see, e.g., Ref. [262].

well-defined and broadly applicable notion of classicality has been established. In turn, this relation provides a direct interpretation of Markovian processes that violate Leggett-Garg inequalities in terms of the underlying quantum resources. However, this connection only holds in the memoryless case and does not straightforwardly apply to the non-Markovian scenario, where, amongst other issues, such propagators cannot be used to compute multi-time statistics [264].

Here, we go beyond this paradigm of memoryless processes and consider the general case of non-Markovian dynamics. Such general processes can be described in terms of higher-order quantum maps, so-called quantum combs [125, 126]. Recently, this framework has been tailored to the description of open quantum system dynamics [128, 129], and has—amongst others—found direct application in the characterisation of multi-time memory effects [106, 130–132] and within the field of stochastic thermodynamics [216–218]. Here, we employ it to extend the results of Ref. [170] to the non-Markovian case. In particular, we link spatial quantum correlations or, more precisely, the discord between an observed system and an environment to the non-classicality of the observed measurement statistics. Somewhat surprisingly, for the case of general processes—where memory effects play a non-negligible role—the presence of non-classical phenomena is not solely dependent on the ability of the process to create or detect coherence, in stark contrast to the memoryless case. As we will show, the absence of detectable coherence is not necessarily sufficient to enforce classical behaviour in general. Rather, classicality of multi-time statistics is inherently linked to quantum discord—which was originally introduced as a means to distinguish classical spatial correlations from non-classical ones [265–268]—between the evolving system and its environment. We characterise the complete set of classical processes and derive a concrete relation between the presence and detectability of discord and the non-classicality of observed multi-time measurement statistics. This, in turn, allows for the derivation of experimentally accessible quantifiers of non-classicality and the categorisation of the resources required for the implementation of a non-classical, non-Markovian process, paving the way to a clear-cut understanding of non-classicality on operational grounds.

In a similar manner to the analysis of coherences, our results will predominantly be phrased with respect to measurements in an arbitrary, but predetermined basis i.e., with respect to a fixed observable, raising the question if classicality is merely a question of perspective; in principle, for every process, there could exist a sequential measurement scheme, that yields classical statistics. While this always holds true for processes in classical physics, as well as memoryless quantum processes, we show by means of an explicit example, that this is not necessarily the case for quantum processes with memory; in the presence of quantum memory, there exists a fundamentally new class of processes, which we will call *genuinely* quantum processes, that lead to non-classical statistics independent of how they are probed.

Throughout this article, we investigate the question of when a physical process—with or without memory—can be considered classical, and what classicality implies if we assume the underlying theory to be quantum mechanics. Concretely, for the most part, we consider the scenario of a quantum system of interest that is sequentially probed in a fixed basis, that is, interrogated at successive points in time—like, for example, in Leggett-Garg type experiments—and we are interested in characterising when the multi-time measurement statistics resulting from such a scenario can be simulated by a classical stochastic process, and thus be reasonably considered classical.

As we will make no assumption about the underlying dynamics, the system of interest can be coupled to an environment that is out of the experimenter’s control and can thus undergo an open evolution that displays complex classical and quantum memory effects. The classicality of the observed statistics then depends on the interplay of the dynamics of the system of interest, the pertinent memory effects, and the way in which the system is probed. We derive both the structural as well as dynamical properties of general classical non-Markovian processes, providing an answer to the question: *What is a non-classical process, and what are its key features?*

Finally, by dropping the restriction to fixed instruments, we show that an observer-independent notion of non-classicality exists, i.e., that there are processes that, no matter how they are probed, display statistics that cannot be simulated by classical stochastic processes. As such processes cannot exist in the absence of memory, the interplay of quantum memory effects and quantum dynamics leads to a fundamentally new class of processes—genuinely quantum processes—that cannot hide their non-classicality.

3.2 Summary of the Main Results

Before providing detailed derivations in the subsequent sections, here, we give a more concrete overview of the main results of our work. Throughout this article, we define the classicality of a process based on observed multi-time statistics $\mathbb{P}_n(x_n, t_n; \dots; x_1, t_1)$ for measurements at different times $\{t_1, \dots, t_n\}$. The number of possible outcomes is always considered to be finite, and, unless stated otherwise, the measurements are given by measurements in the computational basis $\{|x_k\rangle\langle x_k|\}$. With respect to these statistics, a process is considered classical (on K times), if the made measurements are non-invasive, i.e., they satisfy the Kolmogorov conditions

$$\begin{aligned} & \mathbb{P}_{n-1}(x_n, t_n; \dots; \cancel{x_j, t_j}; \dots; x_1, t_1) \\ &= \sum_{x_j} \mathbb{P}_n(x_n, t_n; \dots; x_j, t_j; \dots; x_1, t_1) \quad \forall n \leq K, \forall j. \end{aligned} \quad (3.1)$$

On the other hand, it is Markovian, i.e., memoryless, if the respective conditional probabilities satisfy

$$\mathbb{P}(x_n | x_{n-1}, \dots, x_1) = \mathbb{P}(x_n | x_{n-1}) \quad \forall n \leq K. \quad (3.2)$$

In quantum mechanics, such a process can be modelled by means of Completely Positive and Trace Preserving (CPTP) $\{\Lambda_{t_j, t_{j-1}}\}$, which act on the probed system and describe the dynamics between measurements, as well as an initial system state ρ_{t_0} .

Going beyond the results of Ref. [170], we show that (see Theorem 3.1) a Markovian process is classical iff it can be modelled by a state ρ_{t_0} that is diagonal in the measurement basis $\{|x_k\rangle\langle x_k|\}$ and Non-Coherence-Generating-and-Detecting (NCGD) maps $\Lambda_{t_k, t_{k-1}}$, i.e., maps that satisfy

$$\Delta \circ \Lambda_{t_{j+1}, t_j} \circ \Delta \circ \Lambda_{t_j, t_{j-1}} \circ \Delta = \Delta \circ \Lambda_{t_{j+1}, t_j} \circ \Lambda_{t_j, t_{j-1}} \circ \Delta \quad \forall j, \quad (3.3)$$

where Δ is the completely dephasing map in the measurement basis, and \circ denotes composition. Intuitively, maps that satisfy the above equation can create coherences, but not in a way that can be detected at a later time by means of the employed measurement basis. Thus, Theorem 3.1 provides a direct connection between coherence and an experimentally testable notion of classicality in the Markovian case.

Extending to the non-Markovian setting, we show that this direct connection between coherence and classicality breaks down when memory is present. We provide an explicit example (Example 3.1) of a dynamics $U_{t_j, t_i} |\ell, p\rangle = e^{i\phi_{\ell p}(t_j - t_i)} |\ell, p\rangle$ acting on a qubit system (represented by ℓ) coupled to a continuous degree of freedom (represented by p) that—for the right choice of initial environment state—never displays coherences in the system state, but exhibits non-classical statistics nonetheless.

When memory plays a non-negligible role, individual CPTP maps that act on the system alone are insufficient for the computation of multi-time probabilities. Rather, probabilities are computed by means of higher order quantum maps, called quantum combs [125, 126]. These maps contain all information about the underlying process at hand, and multi-time joint probabilities can then be expressed as

$$\mathbb{P}_K(x_K, t_K; \dots; x_1, t_1) = \mathcal{C}_K[\mathcal{P}_{x_K}, \dots, \mathcal{P}_{x_1}], \quad (3.4)$$

where \mathcal{C}_K is the quantum comb of the process and $\{\mathcal{P}_{x_j}\}$ are the Completely Positive (CP) maps corresponding to measurements with outcome x_j , i.e., $\mathcal{P}_{x_j}[\rho] = \langle x_j | \rho | x_j \rangle |x_j\rangle\langle x_j|$.

We derive a full characterisation of combs that lead to classical statistics in Theorem 3.2, and make this characterisation more concrete in Theorem 3.2', employing the Choi-Jamiołkowski Isomorphism (CJI) that allows one to map higher order quantum maps \mathcal{C}_n onto multipartite quantum states C_n .

Using this full characterisation, a measure $M(C)$ for the non-classicality of a process C can be derived. We phrase this problem in terms of the operational task of deciding whether or not a given comb C is classical, and show that the corresponding maximum probability to guess correctly is given by [see Eq. (3.54)]

$$\mathbb{P}(C) = \frac{1}{2}[1 + M(C)], \quad (3.5)$$

where $M(C)$ can both be computed efficiently via a linear program [see Eq. (3.56)] and is accessible experimentally—and could be evaluated based on already existing experimental data (e.g., in Ref. [269]). We show that, e.g., in the two-time case

$$M(C) \leq \sum_{x_2} \left| \mathbb{P}(x_2) - \sum_{x_1} \mathbb{P}(x_2, x_1) \right|, \quad (3.6)$$

holds, where the right hand side of the above equation is a natural quantifier of classicality, that is used both theoretically, as well as experimentally (for example in Leggett-Garg type scenarios) to quantify the non-classicality of sequential measurement statistics.

In the same vein as in the Markovian case, the dynamical properties (in contrast to the aforementioned structural ones) of classical processes can be obtained. In the non-Markovian case, a process is fully defined by an initial system-environment state $\eta_{t_0}^{se}$ and intermediate system-environment CPTP maps $\Gamma_{t_j, t_{j-1}}$. We show that in the non-Markovian case, rather than the coherences of the system it is the (basis dependent) system-environment discord [265–268] that determines the classicality of the observed statistics. In particular, we demonstrate (see Thms. 3.4 and 3.5) that a process is classical iff it can be modelled by an initial state $\eta_{t_0}^{se}$ with vanishing (basis dependent) discord, i.e., $\eta_{t_0}^{se} = \sum_m p_m |x_m\rangle\langle x_m| \otimes \xi_m$, and a set of system-environment maps that is Non-Discord-Generating-and-Detecting (NDGD), i.e.,

$$\Delta \circ \Gamma_{t_{j+1}, t_j} \circ \Delta \circ \Gamma_{t_j, t_{j-1}} \circ \Delta = \Delta \circ \Gamma_{t_{j+1}, t_j} \circ \Gamma_{t_j, t_{j-1}} \circ \Delta, \quad (3.7)$$

where the completely dephasing map Δ acts on the system alone. Analogously to the Markovian case, the above equation implies that the maps $\{\Gamma_{t_j, t_{j-1}}\}$ can create discord, but said discord cannot be detected by means of later measurements on the system in the chosen measurement basis. In turn, this result provides a direct connection between quantum discord and the classicality of a quantum process. Additionally, it also gives an a posteriori explanation why the absence of coherence in Example 3.1 did not lead to classical statistics (for an explicit discussion of the discord that leads to of non-classical statistics in Example 3.1, see its continuation Example 3.3).

While, in principle, these aforementioned results do not rely on the fact that we assume measurements in onefixed basis, but could similarly be obtained for different (but fixed) instruments at every time, they still depend on the fact that one specific measurement scheme is chosen beforehand. Classicality (or the absence thereof) of the observed statistics could thus depend on the respective choice of measurement schemes. This holds true in the Markovian case, where there is always a choice of measurement bases that renders the observed statistics classical. However, as we show by explicit example (see Section 3.7), there are processes with memory—dubbed genuinely quantum—that display non-classical statistics independent of the employed measurement scheme.

The paper is structured as follows: In Section 3.3 we introduce the basic concepts that will be employed throughout this article to examine classicality. In Section 3.4,

we reiterate and slightly generalise the results of Ref. [170] linking non-classicality and coherence for the Markovian case, and discuss their breakdown when memory effects are present. This motivates our consideration of the non-Markovian case in Section 3.5, where we fully characterise the set of general classical processes by means of the quantum comb framework. This characterisation then enables us to formulate a quantifier of non-classicality, that is both experimentally accessible and can be computed efficiently. Based on these results, in Section 3.6, we subsequently establish the direct connection between (basis dependent) quantum discord and the classicality of temporal processes. Finally, in Section 3.7, we go beyond the paradigm of measurements in a fixed basis, and provide an example for processes that appear quantum independent of the scheme that is used to probe them. The paper concludes in Section 3.8 with a summary and an outlook on further research directions and open problems.

3.3 General Framework

The overarching aim of this paper is to characterise when a general quantum mechanical process can be considered classical in an operationally consistent manner and identify the structural properties consequently implied on the underlying evolution. Importantly, our investigation will be operational in the sense that it is based solely on experimentally accessible quantities; as such, it applies to situations where the underlying theory is classical mechanics, quantum mechanics, or some more general theory [270].

Ultimately, any physical theory provides predictions about possible observations—only these can be tested by experiments. That is, any theory must (in principle) provide the correct probabilities for measurement outcomes (or sequences thereof) to occur when a system of interest is experimentally probed. The difference between predictions made regarding such observable quantities by classical physics and quantum (or post-quantum) theory can then be used to unambiguously demarcate between the theories on the investigated spatial and temporal scales.

Following Ref. [170], we will thus define our notion of classicality by means of joint probability distributions pertaining to sequences of measurement outcomes, as these are precisely what is obtained when a temporal process is probed.

3.3.1 Kolmogorov Consistency Conditions and Classicality

In classical physics, a stochastic process on a set of K times is fully described by a joint probability distribution

$$\mathbb{P}_K(x_K, t_K; \dots; x_1, t_1) \tag{3.8}$$

which yields the probability to measure the realisations $\{x_K, \dots, x_1\}$ of the random variables $\{X_K, \dots, X_1\}$ at times $\{t_K, \dots, t_1\}$. For example, $\mathbb{P}_2(x_2, t_2; x_1, t_1)$ could describe

the probability to obtain both outcomes $\{x_2, x_1\}$ when measuring the position of a particle undergoing Brownian motion at times t_1 and $t_2 > t_1$. In what follows, we will often omit the explicit time label, with the understanding that x_j denotes an outcome of a measurement at time t_j .

Crucially, in classical physics, joint probability distributions describing a stochastic process for different sets of times satisfy the so-called *Kolmogorov consistency conditions* [145, 271–273]: Given a joint probability distribution \mathbb{P}_K for a set of times, the probability distributions for all subsets of times can be obtained by marginalisation, that is

$$\begin{aligned} & \mathbb{P}_{n-1}(x_n, t_n; \dots; \cancel{x_j, t_j}; \dots; x_1, t_1) \\ &= \sum_{x_j} \mathbb{P}_n(x_n, t_n; \dots; x_j, t_j; \dots; x_1, t_1) \quad \forall n \leq K, \forall j. \end{aligned} \quad (3.9)$$

Just like the Leggett-Garg inequalities [147–149] for temporal correlations, the satisfaction of these requirements is based on the assumptions of realism per se, i.e., the assumption that x_j has a definite value at any time t_j , and the possibility to implement non-invasive measurements [146].

Importantly, an experimenter obtaining a family of joint probability distributions that satisfies the Kolmogorov conditions when probing a temporal process at different sets of times would not be able to distinguish said process from a classical one, as every such finite family can be obtained from a—potentially exotic—underlying *classical stochastic process*. More generally, the *Kolmogorov extension theorem* states that if all joint probability distributions for finite subsets of a time interval $[0, t]$ satisfy the consistency conditions of Eq. (3.9) amongst each other, then there exists an underlying classical stochastic process on said time interval that leads to the observed probability distributions [145, 271–273]. In other words, if the Kolmogorov consistency conditions of Eq. (3.9) are satisfied (for all considered choices of t_j), then there is nothing inherently quantum mechanical about the observed process. We therefore define:

Definition 3.1 (*K-classical process* [170]). Let \mathcal{X} be a finite set. A process defined on a set of times \mathcal{T} , with $|\mathcal{T}| = K$, that is described by the joint probabilities $\mathbb{P}_n(x_n, t_n; \dots; x_1, t_1)$, with $t_n \geq \dots \geq t_1$, $t_i \in \mathcal{T}$, $n \leq K$ and $x_i \in \mathcal{X}$, is said to be *K-classical* if the Kolmogorov consistency conditions of Eq. (3.9) are satisfied up to $n = K$.

Throughout this article, we will call a family of joint probabilities on a set of K times a *K-process* and denote it by $\{\mathbb{P}_n(x_n, \dots, x_1)\}_{n \leq K}$. Here, the label $n \leq K$ is a short-hand notation for all the subsets of \mathcal{T} with n ordered times $t_n \geq \dots \geq t_1$, where $t_i \in \mathcal{T}$, for any $n \leq K$; moreover from here on we will not indicate explicitly the time arguments in the probability distributions, implying that the outcome x_j refers to time t_j .

While the above definition of classicality seems intuitive, some comments are in order. First, we choose to define classicality for a finite set of K times. While this is motivated

on a practical ground, the general definition of a classical stochastic process involves the joint probability distributions associated with any number of ordered time instants $t_K \geq \dots \geq t_1$, with $K \in \mathbb{N}$, and any choice of such instants. In particular, as said, the Kolmogorov extension theorem infers the existence of a stochastic process from the validity of the consistency conditions on all such joint distributions. Here, instead, we fix a finite value of K and the sequence of time instants beforehand, so that, given the K -time joint probability distribution of a K -classical process, the involved hierarchy of probability distributions can be constructed by iteratively applying the consistency conditions, at any intermediate time.

Second, the above definition of classicality is a priori device independent, as it only relies on the inferred statistics without any assumptions on the underlying theory and/or measurement devices; as a consequence, the classicality of a process according to the above definition depends upon the manner in which the system of interest is probed. Although often overlooked, this is also the case in classical physics: Given some underlying classical stochastic process, not every set of measurements that an experimenter might be able to perform will lead to a family of probability distributions that satisfies the above definition of K -classicality. In fact, if performing such measurements might potentially disturb the system (i.e., the measurement is invasive), the Kolmogorov condition fails in general, even if the underlying evolution is classical [146].

For example, suppose that instead of merely measuring the position of a particle at different times when probing a Brownian motion process, an experimenter chooses to displace the particle at each time depending on where it was found. In this case, Eq. (3.9) would generally fail to hold for the joint probability distributions observed. Consequently, the Kolmogorov consistency conditions [Eq. (3.9)] are in fact a statement of the *non-invasiveness* of the performed measurements: If they hold true, then *not* performing a measurement at any given time cannot be distinguished (for the given experimental situation) from averaging over their probabilities (i.e., forgetting the outcomes of the measurements performed).

In classical physics one assumes that, in principle, one could measure the system without disturbing it, and that therefore there exists a family of joint probability distributions that can consistently explain all possible outcome probabilities. Such a non-invasive and complete measurement is often referred to as an “ideal measurement” in the literature [274]. On the other hand, in quantum mechanics any measurement disturbs *some* system state and therefore ideal measurements do not exist in general in the strong sense discussed above. As a consequence, quantum mechanical processes generically do not satisfy Kolmogorov conditions [108, 146], a fact that fundamentally distinguishes them from the classical realm.

More generally, the violation of Bell, Kochen-Specker, or Leggett-Garg inequalities, which can be observed in quantum mechanics, are different manifestations of the impossi-

ility to obtain the measured data by non-invasive measurements. Particularly, in the case of Leggett-Garg inequalities [147, 275], it is precisely the breakdown of Kolmogorov conditions that is being probed [146, 170], and our above definition of classicality is hence in line with the wider program of determining fundamentally quantum traits of nature.

3.3.2 Measurement Setup

As mentioned above, the structural properties of families of joint probability distributions depend on the way in which a system of interest is probed. Consequently, before being able to analyse the set of quantum processes, it is crucial to fix the measurements that are used to probe a process at hand. Although there are no ideal measurements in quantum mechanics, projective measurements share some basic features with the classical ideal measurements discussed above, and are thus a natural choice. In particular, they guarantee repeatability, i.e., that two sequential measurements (without any evolution in between) would give the same value with unit probability, as well as a weaker form of ideality, namely that if an outcome occurs with certainty, then the state of the system before the measurement is not disturbed by the latter [276]. It therefore suggests itself to start our analysis on the classical reproducibility of quantum processes by focusing on projective measurements; moreover, also following Ref. [170], we will further restrict to the case of orthogonal rank-1 (sharp) projectors, like, e.g., projective measurements with respect to the eigenbasis of any non-degenerate self-adjoint operator.

In many experimental situations of interest, there is a preferred basis to select. For instance, if the dynamics is such that the system dephases to a given basis, the latter provides a natural choice. This occurs, e.g., in the case of open quantum systems dynamics that are subject to environmental fluctuations. In other cases it may make sense to choose the basis more arbitrarily (in advance), for instance when analysing a specific protocol, or attempting to optimise it (see Ref. [277] for more details). Finally, the experimental setup might only allow for a measurement of one particular observable, in which case the chosen basis would correspond to the eigenbasis of said observable.

In what follows, we will analyse the classicality of a process based on the joint probability distributions obtained from sequential sharp measurements in a fixed basis $\{|x\rangle\}_{x=1}^d$ —henceforth also called the classical, standard, or computational basis—with the action of a measurement with outcome x on a state ρ given by

$$\rho \mapsto \mathcal{P}_x[\rho] := |x\rangle\langle x|\rho|x\rangle\langle x|. \quad (3.10)$$

See Fig. 3.1 for a graphical depiction.

This freedom in the considered measurements makes the property of classicality fundamentally contingent on the respective choice of measurement basis. However, this basis dependence is unsurprising and mirrored by coherence theory [242]. There, the existence of off-diagonal elements $\langle m|\rho|n\rangle$, i.e., coherences, depends on the choice of the




Figure 3.1: Probing a Process with Projective Measurements. At each time t_j , the process (depicted in blue) is probed by a projective measurement (depicted in green) with outcomes x_j , where each x_j belongs to the same finite set \mathcal{X} . If the resulting family of probability distributions \mathbb{P}_n (depicted are the cases $n \leq 4$) satisfies the Kolmogorov consistency conditions, then not performing a measurement at a time t_j cannot be distinguished from performing a measurement and averaging over the outcomes. In this case, this experiment cannot be distinguished from a classical one, even though the underlying evolution might be quantum mechanical.

basis a quantum state is represented in. As they are considered to be a fundamentally quantum property, it is a natural question to ask how coherences (with respect to the computational basis) and classicality of a process (with respect to the same basis) are interrelated. Importantly, while the existence of coherences cannot be determined by projective measurements in the computational basis alone, the prevalence of non-classical effects can be. Thus, as we shall see below, providing an operationally accessible notion of classicality allows one to link coherence (and, more generally, quantum correlations) in a quantitative manner to experimentally observable deviations from classical physics.

3.3.3 Open (Quantum) System Dynamics and Memory Effects

The definition of classicality we use answers the question of whether or not there exists a classical stochastic process that can explain the multi-time probabilities obtained by measuring a quantum system at given times in the computational basis. To make our analysis as general as possible, we will consider the possibility that the measured system interacts with a surrounding environment, which can influence the resulting statistics. Explicitly, assuming that the system and environment in state η are together closed and described by quantum mechanics, their joint dynamics between measurements is given by unitary evolution $\mathcal{U}_{t_{j+1}, t_j}[\eta] = U_{t_{j+1}, t_j} \eta U_{t_{j+1}, t_j}^\dagger$. The resulting joint probability distributions read

$$\mathbb{P}_n(x_n, \dots, x_1) = \text{tr} \left\{ (\mathcal{P}_{x_n}^s \otimes \mathcal{I}^e) \circ \mathcal{U}_{t_n, t_{n-1}} \circ \dots \circ \mathcal{U}_{t_2, t_1} \circ (\mathcal{P}_{x_1}^s \otimes \mathcal{I}^e) [\eta_{t_1}^{se}] \right\}, \quad (3.11)$$

where $\eta_{t_1}^{se}$ is the system-environment state at time t_1 , \mathcal{I}^e signifies the identity channel on the environment, $\mathcal{P}_{x_j}^s$ corresponds to a measurement on the system in the computational basis at time t_j with outcome x_j and \circ denotes composition (see Fig. 3.2 for a graphical representation). Whenever there is no risk of confusion, we will drop the additional superscripts s and e throughout this paper. Naturally, the classicality of the family of




Figure 3.2: General Open Quantum Process. The state of the system at time t_1 is correlated with the environment (depicted by the yellow triangle representing the joint state). Measurements on the system (green boxes) are performed at times t_1, t_2, \dots . In between, the system and the environment undergo a unitary evolution (blue boxes). The distinction between system and environment is given by the degrees of freedom that the experimenter controls (system) and those that remain inaccessible to experimental control (environment).

joint probability distributions obtained via Eq. (3.11) crucially depends on the properties of the intermediate evolutions $\mathcal{U}_{t_{j+1}, t_j}$ and the initial state $\eta_{t_1}^{se}$.

In general, such a multi-time statistics displays memory effects, i.e., it is non-Markovian: At any point in time t_j , the future statistics does not only depend on the measurement outcome x_j at time t_j , but also on (potentially) all previous outcomes x_{j-1}, \dots, x_1 . Indeed, all information about future statistics at t_j is contained in the joint state of system and environment, which depends upon the previous measurement outcomes. As this total state cannot be accessed by measurements on the system alone, this dependence on past measurements manifests itself as memory effects on the system level (see Section 3.5 for a detailed discussion).

However, under some specific circumstances, the influence of such memory effects on the multi-time statistics can be neglected; this is essentially the case when the Quantum Regression Formula (QRF) can be applied [272, 278–280]. Under this assumption, the observed statistics can be understood in terms of dynamical propagators that act on the system alone, which, in turn, enables one to directly link the classicality of a process to the properties of said propagators in terms of coherence production and detection. The corresponding result has been obtained in Ref. [170], and we will reiterate and expand upon it in the coming section. Subsequently, employing quantum combs—a powerful framework for the description of general, possibly non-Markovian open quantum processes—we characterise the set of quantum processes that can be described classically.

3.4 Coherence and Classicality

In this section, we reiterate the main result of Ref. [170] on the connection between coherence and classicality for the memoryless case, generalising it to the case of a divisible (but not necessarily semigroup [272, 281, 282]) dynamics. As mentioned above, such a direct connection may be established, because memoryless processes can be understood in terms of propagators that are defined on the system alone, while this property fails to hold in the general, non-Markovian, case.

After introducing an operational notion of Markovianity associated with the multi-time

statistics due to sequential measurements of a (non-degenerate) observable, we present a one-to-one connection between the non-classicality of such statistics and the capability of the open system dynamics to generate and detect coherences with respect to the relevant basis. We also clarify the relation between the notion of Markovianity used in this paper and the QRF, which allows us to straightforwardly recover the main result of Ref. [170]. Finally, we lay out the subtleties that arise when generalising the framework to allow for memory effects, motivating the main results of this work.

3.4.1 One-to-One Connection in the Markovian Case

Classically, a process is Markovian (i.e., memoryless), if, for any chosen time t_j , the future statistics only depend upon the outcome at time t_j , but not on any prior outcomes at t_{j-1}, t_{j-2}, \dots ; explicitly, a classical stochastic process is Markovian if its statistics satisfy

$$\mathbb{P}(x_j|x_{j-1}, \dots, x_1) = \mathbb{P}(x_j|x_{j-1}) \quad \forall j, \quad (3.12)$$

where $\mathbb{P}(x_j|x_{j-1}, \dots, x_1)$ is the conditional probability to obtain outcome x_j at time t_j given that outcomes x_{j-1}, x_{j-2}, \dots were measured at earlier times t_{j-1}, t_{j-2}, \dots [272]. Extending this definition to general (i.e., not necessarily classical) statistics and taking into account that, in practice, one only deals with systems probed at a finite number of times, we obtain the following definition of K -Markovianity:

Definition 3.2. Let \mathcal{X} be a finite set. A process defined on a set of times \mathcal{T} , with $|\mathcal{T}| = K$ is called K -Markovian if it satisfies:

$$\mathbb{P}(x_n|x_{n-1}, \dots, x_1) = \mathbb{P}(x_n|x_{n-1}) \quad \forall n \leq K, \quad (3.13)$$

for all ordered tuples of times $t_n \geq \dots \geq t_1$, with $t_i \in \mathcal{T}$, and $x_i \in \mathcal{X}$.

Just like our earlier definition of classicality and coherence, the absence of memory effects as defined in Definition 3.2 is basis dependent: A process that appears Markovian in one basis may appear non-Markovian when probed in a different one. While there exist basis independent notions of Markovianity in the quantum case [103, 104, 128, 129, 140], the basis dependent one introduced here is best suited for the experimental situation we envision; as such, in what follows, we predominantly understand Markovianity with respect to measurements in the computational basis. We will briefly return to the relation between this basis dependence and the basis independent notion of Markovianity in Section 3.5.

To establish a connection between non-classicality of a Markovian process and the coherence properties of the underlying dynamics, we need to introduce the maps that characterise the dynamical evolution of the open system. To this end, assume that at an initial time t_0 (with $t_0 \leq t_1$) the system and the environment are in a product




Figure 3.3: Markovian Process. For a Markovian process, the system dynamics in between intermediate times (depicted as the blue boxes) can be modelled by maps Λ_{t_{j+1}, t_j} that do not depend on previous outcomes (i.e., there is no memory). The measurement statistics are obtained by measuring in the classical basis at times t_1, t_2, t_3, \dots (depicted in green); before the first measurement the system is in the state ρ_{t_1} (depicted in yellow).

state $\eta_{t_0}^{se} = \rho_{t_0} \otimes \sigma_{t_0}$ (for some fixed environment state σ_{t_0}), so that we can define the Completely Positive and Trace Preserving (CPTP) dynamical maps $\{\Lambda_{t_j, t_0}\}$ of the open system evolution between the initial time and the measurement times t_j [272, 283]

$$\rho_{t_j} = \Lambda_{t_j, t_0}[\rho_{t_0}] = \text{tr}_e \left[U_{t_j, t_0} (\rho_{t_0} \otimes \sigma_{t_0}) U_{t_j, t_0}^\dagger \right], \quad (3.14)$$

where tr_e denotes the trace over the environmental degrees of freedom. Additionally, let us also assume that the dynamics is divisible [284], i.e, we can define the corresponding propagators $\{\Lambda_{t_k, t_j}\}$ between any two times via the composition rule

$$\Lambda_{t_k, t_0} = \Lambda_{t_k, t_j} \circ \Lambda_{t_j, t_0} \quad \forall t_k \geq t_j \geq t_0, \quad (3.15)$$

and they satisfy the composition law $\Lambda_{t_\ell, t_j} = \Lambda_{t_\ell, t_k} \circ \Lambda_{t_k, t_j}$ for all times $t_\ell \geq t_k \geq t_j$. Under these assumptions, it is natural to ask, what conditions the propagators $\{\Lambda_{t_k, t_j}\}$ must satisfy in order for the resulting statistics to be classical. However, Eq. (3.15) does not yet tell us how to obtain multi-time statistics [151].

The relation we seek is provided by the QRF, which, for example, holds in the weak coupling and the singular coupling limits [285], and constitutes a relation between the definition of Markovian processes given by Definition 3.2 and the corresponding open system dynamics (see also Ref. [140] for an extensive discussion of the QRF and its generalisations). For the case of rank-1 projective measurements (in the computational basis), the QRF states that the multi-time probability distributions in Eq. (3.11) can be equivalently expressed by

$$\mathbb{P}_n(x_n, \dots, x_1) = \text{tr} \left[\mathcal{P}_{x_n} \circ \Lambda_{t_n, t_{n-1}} \circ \dots \circ \Lambda_{t_2, t_1} \circ \mathcal{P}_{x_1} \circ \Lambda_{t_1, t_0}[\rho_{t_0}] \right]. \quad (3.16)$$

Importantly, this means that the full multi-time statistics can be obtained by means of maps that are independent of the respective previous measurement outcomes and which act on the system alone (see Fig. 3.3 for a graphical representation).

It is straightforward to see that satisfaction of the QRF [see Eq. (3.16)] implies Markovian statistics in the sense of Eq. (3.13) and in particular we have the identities

$$\langle x_k | \Lambda_{t_k, t_j} [|x_j\rangle\langle x_j|] |x_k\rangle = \mathbb{P}(x_k | x_j) \quad \forall j \geq 1, \quad (3.17)$$

$$\text{and } \langle x_1 | \Lambda_{t_1, t_0}[\rho_{t_0}] |x_1\rangle = \mathbb{P}(x_1). \quad (3.18)$$

In other words, the action of the propagators on the populations (i.e., the diagonal terms of ρ_{t_j} , the state of the system at t_j) can be identified with the conditional probabilities between any two times. Crucially, this is not generally the case, and breaks down in situations where the QRF cannot be applied [155].

More generally, even if the QRF applies, the composition rule on the level of propagators does not imply a composition rule on the level of the resulting measurement statistics, i.e., for a divisible process that satisfies the QRF, we generally have

$$\sum_{x_k} \mathbb{P}(x_\ell|x_k)\mathbb{P}(x_k|x_j) \neq \mathbb{P}(x_\ell|x_j), \quad (3.19)$$

which captures the deviation of quantum Markovian processes from classical ones. As mentioned previously, in order for the resulting process to be classical, not performing a measurement must be indistinguishable from performing a measurement and averaging over all possible outcomes. Put differently, for an observer that can only perform measurements in a fixed basis, the process is classical if they cannot detect the invasiveness of measurements in said basis.

A measurement at time t_j in the fixed basis where the measurement outcomes are averaged over can be represented by the *completely dephasing map*

$$\Delta[\rho] = \sum_{x_j} \mathcal{P}_{x_j}[\rho] = \sum_{x_j} \langle x_j|\rho|x_j\rangle |x_j\rangle\langle x_j|. \quad (3.20)$$

The natural property of the propagators to look at in relation to classicality is thus that for all t_j :

$$\begin{aligned} & \Delta_{j+1} \circ \Lambda_{t_{j+1},t_j} \circ \Delta_j \circ \Lambda_{t_j,t_{j-1}} \circ \Delta_{j-1} \\ &= \Delta_{j+1} \circ \Lambda_{t_{j+1},t_j} \circ \mathcal{I}_j \circ \Lambda_{t_j,t_{j-1}} \circ \Delta_{j-1} \\ &= \Delta_{j+1} \circ \Lambda_{t_{j+1},t_{j-1}} \circ \Delta_{j-1}, \end{aligned} \quad (3.21)$$

where \mathcal{I}_j and Λ_j are the identity map and the completely dephasing map at time t_j , respectively (see Fig. 3.4 for a graphical representation). In the last line of Eq. (3.21) we used the composition law $\Lambda_{t_{j+1},t_{j-1}} = \Lambda_{t_{j+1},t_j} \circ \Lambda_{t_j,t_{j-1}}$. Eq. (3.21) is satisfied, e.g., if none of the maps $\{\Lambda_{t_{j+1},t_j}\}$ create coherences. More generally, each of the maps in Eq. (3.21) can in principle create coherences, as long as these coherences cannot be detected at the next time by means of measurements in the classical basis. Therefore, such a collection of maps satisfying Eq. (3.21) has been named Non-Coherence-Generating-and-Detecting (NCGD) [170].




Figure 3.4: NCGD Dynamics. If the process is NCGD, then for a classical observer, “doing nothing” (i.e., performing the identity map \mathcal{I} , bottom) cannot be distinguished from a measurement in the classical basis and averaging over the outcomes (i.e., performing the map Δ , top) at any point in time.

The precise connection between NCGD and classicality is expressed by the following theorem:

Theorem 3.1. *Let $\{\mathbb{P}_n(x_n, \dots, x_1)\}_{n \leq K}$ be a K -Markovian process (Definition 3.2). Then, the process is also K -classical (Definition 3.1) if and only if there exist a system state ρ_{t_0} (at a time $t_0 \leq t_1$) which is diagonal in the computational basis $\{|x\rangle\}_{x \in \mathcal{X}}$ and a set of propagators $\{\Lambda_{t_j, t_{j-1}}\}_{j=1, \dots, K}$ which are NCGD with respect to $\{|x\rangle\}_{x \in \mathcal{X}}$, such that ρ_{t_0} and $\{\Lambda_{t_j, t_{j-1}}\}_{j=1, \dots, K}$ yield $\{\mathbb{P}_n(x_n, \dots, x_1)\}_{n \leq K}$ via Eq. (3.16).*

Proof. We first show that if a Markovian process can be reproduced by means of NCGD propagators $\{\Lambda_{t_{j+1}, t_j}\}$ and an initial diagonal state (both properties with respect to the computational basis), then it yields classical statistics. If the statistics is Markovian, then it follows from Eq. (3.13) that the joint probability distribution on any set of times $t_n \geq \dots \geq t_1$, with $t_i \in \mathcal{T}$, is given by

$$\mathbb{P}_n(x_n, \dots, x_1) = \mathbb{P}(x_n | x_{n-1}) \cdots \mathbb{P}(x_2 | x_1) \mathbb{P}(x_1). \quad (3.22)$$

As the process can, by assumption, be reproduced by the maps $\{\Lambda_{t_j, t_{j-1}}\}$ via Eq. (3.16), then for any time t_j we have

$$\begin{aligned} \sum_{x_j} \mathbb{P}(x_{j+1} | x_j) \mathbb{P}(x_j | x_{j-1}) &= \sum_{x_j} \text{tr}\{\mathcal{P}_{x_{j+1}} \circ \Lambda_{t_{j+1}, t_j}[\Pi_{x_j}]\} \text{tr}\{\mathcal{P}_{x_j} \circ \Lambda_{t_j, t_{j-1}}[\Pi_{x_{j-1}}]\} \\ &= \text{tr}\{\mathcal{P}_{x_{j+1}} \circ \Lambda_{t_{j+1}, t_j} \circ \Delta_j \circ \Lambda_{t_j, t_{j-1}}[\Pi_{x_{j-1}}]\} \\ &= \text{tr}\{\mathcal{P}_{x_{j+1}} \circ \Lambda_{t_{j+1}, t_{j-1}}[\Pi_{x_{j-1}}]\}, \end{aligned} \quad (3.23)$$

where we have set $\Pi_{x_j} = |x_j\rangle\langle x_j|$ and the NCGD property was used in the last line. This equation implies

$$\sum_{x_j} \mathbb{P}(x_{j+1} | x_j) \mathbb{P}(x_j | x_{j-1}) = \mathbb{P}(x_{j+1} | x_{j-1}). \quad (3.24)$$

Moreover, the (initial) diagonal state ρ_{t_0} guarantees that we have

$$\sum_{x_1} \mathbb{P}(x_2, x_1) = \mathbb{P}(x_2). \quad (3.25)$$

As a consequence of these two previous relations, the family of joint probability distributions computed via Eq. (3.22) satisfies Kolmogorov conditions, and is thus classical.

Conversely, if the process is classical and Markovian, Eq. (3.24) holds. We can then define the maps

$$\tilde{\Lambda}_{t_{j+1}, t_j} [|x_j\rangle\langle y_j|] = \delta_{x_j y_j} \sum_{x_{j+1}} \mathbb{P}(x_{j+1} | x_j) \Pi_{x_{j+1}}, \quad (3.26)$$

and the initial diagonal state

$$\tilde{\rho}_{t_0} = \sum_{x_1} \mathbb{P}(x_1) \Pi_{x_1}, \quad (3.27)$$

which also means that we identify the initial time as the time of the first measurement, $t_1 = t_0$. The set of maps $\{\tilde{\Lambda}_{t_{j+1}, t_j}\}$ defined in this way, in conjunction with $\tilde{\rho}_{t_0}$, reproduces the correct statistics via Eq. (3.16). As they are diagonal in the computational basis for any pair of times t_j and t_{j+1} , they form an NCGD set. \square

Crucially, the connection between classicality and NCGD dynamics is one-to-one: If the obtained Markovian statistics cannot be reproduced by a set of maps that are NCGD, then the process is non-classical. Before discussing classicality in the presence of memory effects below, it is worth discussing the intuitive meaning of this theorem, and NCGD dynamics in particular.

If the process at hand is Markovian and classical, the maps $\{\tilde{\Lambda}_{t_{j+1}, t_j}\}$ (as well as the initial state $\tilde{\rho}_{t_0}$) introduced in the proof of Theorem 3.1 define an artificial reduced dynamics of the system, whose propagators correctly reproduce all joint probability distributions for measurements in the (fixed) classical basis via Eq. (3.16). Note that the actual propagators of the dynamics [i.e., those fixed by the unitary evolution in Eq. (3.11) via Eqs. (3.14) and (3.15)] might differ from the maps $\tilde{\Lambda}_{t_{j+1}, t_j}$ above (and $\tilde{\rho}_{t_0}$ might differ from the actual initial state ρ_{t_0}); indeed, the fact that they do not coincide is simply a manifestation of the basis dependence of the (sequential) measurement scheme we are focusing on here.

Crucially, a composition rule on the level of the actual propagators does not imply a composition rule on the level of the propagators of the populations. This implication only holds if the propagators of the dynamics are NCGD and the resulting statistics can be computed via Eq. (3.16), in which case Eq. (3.21) results in

$$\tilde{\Lambda}_{t_{j+1}, t_{j-1}} = \tilde{\Lambda}_{t_{j+1}, t_j} \circ \tilde{\Lambda}_{t_j, t_{j-1}} \quad \forall t_j, \quad (3.28)$$

with

$$\tilde{\Lambda}_{t_k, t_j} [|x_j\rangle\langle y_j|] = \delta_{x_j y_j} \langle x_k | \Lambda_{t_k, t_j} [|x_j\rangle\langle x_j|] |x_k\rangle \Pi_{x_j} \quad (3.29)$$

[see Eqs. (3.17) and (3.26)]. These reduced propagators still produce the correct populations, which is the only relevant part for the considered statistics, and set all coherences

to zero. This composition law is then—as already seen in Eq. (3.24)—equivalent to the well-known classical Chapman-Kolmogorov equations

$$\sum_{x_j} \mathbb{P}(x_{j+1}|x_j)\mathbb{P}(x_j|x_{j-1}) = \mathbb{P}(x_{j+1}|x_{j-1}), \quad (3.30)$$

which hold for classical Markovian processes: If the measurement statistics of a Markovian process can be reproduced by a set of NCGD maps $\{\Lambda_{t_j, t_{j-1}}\}$, then it can also be reproduced by the set of maps $\{\tilde{\Lambda}_{t_j, t_{j-1}}\}$, which act non-trivially on only the populations of the computational basis and satisfies a composition law, thus the process is classical.

Conversely, if the classical composition rule of Eq. (3.30) holds for a Markovian process, then there exists a set $\{\tilde{\Lambda}_{t_{j+1}, t_j}\}$ of propagators [e.g., those defined in Eq. (3.26)] that are NCGD and correctly reproduce all joint probability distributions for measurements in the (fixed) classical basis.

Theorem 3.1 is a generalisation of the main result of Ref. [170] in two ways. First, it does not impose any restriction on the propagators of the underlying quantum evolution, while in Ref. [170] these were required to form a semigroup, i.e., $\Lambda_{t_{j+1}, t_j} = e^{\mathcal{L}(t_{j+1}-t_j)}$, for some Lindbladian \mathcal{L} [281, 282].

Second, the definition of Markovianity used here coincides with the standard definition of classical stochastic processes, whereas in Ref. [170], a definition based on Eq. (3.16) (for semigroups) was used. Consequently, while the maps $\{\Lambda_{t_{j+1}, t_j}\}$ cannot be fully probed by measurements in the computational basis alone, the requirement of Eq. (3.30) can be tested for by simply performing sequences of measurements in the classical basis at the relevant times, thus making our theorem fully operational. However, this comes at the cost of dealing with propagators $\{\tilde{\Lambda}_{t_{j+1}, t_j}\}$ which possibly do not correspond to those of the actual reduced dynamics.

On the other hand, as we show in Appendix C.1, a one-to-one correspondence between the dynamical propagators Λ_{t_{j+1}, t_j} and the non-classicality of the multi-time statistics can be established also in the general (non-semigroup) divisible case, when the QRF applies, provided that one assumes a proper invertibility condition on the restriction of the dynamical maps to the populations of the computational basis. Indeed, this also allows one to recover in a straightforward way the main result of Ref. [170] as a corollary by further imposing the semigroup composition law.

Importantly, Theorem 3.1 characterises the connection between coherences and the classicality of a Markovian process. While it is not necessary that the underlying propagators do not create coherences in order for a Markovian process to be classical, it is necessary and sufficient that coherences—should they be created—cannot be detected at a later point in time by means of measurements in the computational basis. Put differently, the propagators must be such that a classical observer could not decide whether at any point in time an identity map or a completely dephasing map was performed (which is

depicted in Fig. 3.4). This requirement is exactly encapsulated in the NCGD property of the propagators.

3.4.2 Coherence in the Non-Markovian Case: Preliminary Analysis

The above connection between quantum coherence and non-classicality fails to hold in the non-Markovian case. On the one hand, in this case propagators between two times are no longer sufficient to fully characterise the multi-time statistics.² On the other hand, even if the state of the system is diagonal in the computational basis at all times, dephasing can still be invasive due to correlations with the environment, breaking the connection between coherences and the classicality of statistics. We will discuss the former problem in the subsequent sections. Using an open system model from Refs. [104, 288, 289], an explicit *ante litteram* example of the latter case has already been provided in Ref. [170] (note also a similar investigation in Ref. [290]), albeit not with an emphasis on the lack of coherence in the system state at all times (even in between the measurements). Here, we reiterate this example, focusing on the absence of coherences in the state of the system. The details of this discussion can be found in Appendices C.2 and C.3. A simpler, although non-continuous, example for a non-Markovian process that yields non-classical statistics but never displays coherences in the system state is provided in Appendix C.4.

Example 3.1. Let the system of interest s consist of a qubit described by $\rho_s(t)$ which is coupled to a continuous degree of freedom p of the environment. The global dynamics of system and environment is governed by the unitary evolution U_{t_j, t_i} , acting as

$$U_{t_j, t_i}|\ell, p\rangle = e^{i\phi_\ell p(t_j - t_i)}|\ell, p\rangle, \quad (3.31)$$

where $\{|\ell\rangle\}_{\ell=0,1}$ is the eigenbasis of the system Pauli operator $\hat{\sigma}_z$ and $\phi_\ell = (-1)^\ell$. The initial system-environment state is assumed to be of product form $\eta(0) = \rho_s(0) \otimes |\varphi^e\rangle\langle\varphi^e|$, with $|\varphi^e\rangle = \int_{-\infty}^{\infty} dp f(p)|p\rangle$, where $f(p)$ satisfies the normalisation condition $\int_{-\infty}^{\infty} dp |f(p)|^2 = 1$. By defining

$$k(t) := \int_{-\infty}^{\infty} dp |f(p)|^2 e^{2ipt}, \quad (3.32)$$

it is straightforward to show that, expressed in the eigenbasis of $\hat{\sigma}_z$, the free open evolution of the state of the system (i.e., without intermediate measurements) is given by

$$\rho_s(t) = \begin{pmatrix} \rho_{00} & k(t)\rho_{01} \\ k^*(t)\rho_{10} & \rho_{11} \end{pmatrix}, \quad (3.33)$$

where $\rho_{mn} := \langle m|\rho_s(0)|n\rangle$.

²For a characterisation of non-Markovian processes in terms of collections of CPTP maps (or sequences thereof), see Refs. [286, 287]. Notably, the characterisation employed in these references is equivalent to the one provided here.

If $\rho_s(0)$ is initialised in a convex mixture of the eigenvectors $\{|\pm\rangle = \frac{|0\rangle \pm |1\rangle}{\sqrt{2}}\}$ of the $\hat{\sigma}_x$ operator, i.e., $\rho_s(0) = \alpha|+\rangle\langle+| + (1 - \alpha)|-\rangle\langle-|$, then

$$\begin{aligned} \rho_s(t) &= \frac{1}{2} \begin{pmatrix} 1 & k(t)(2\alpha - 1) \\ k^*(t)(2\alpha - 1) & 1 \end{pmatrix} \\ &= \frac{1}{2} \left\{ |+\rangle\langle+| [1 + (2\alpha - 1) \operatorname{Re}(k(t))] - |+\rangle\langle-| (2\alpha - 1) \operatorname{Im}(k(t)) \right. \\ &\quad \left. + |-\rangle\langle+| (2\alpha - 1) \operatorname{Im}(k(t)) + |-\rangle\langle-| [1 - (2\alpha - 1) \operatorname{Re}(k(t))] \right\}, \end{aligned} \quad (3.34)$$

i.e., no coherence w.r.t. $\hat{\sigma}_x$ will be generated if $k(t)$ is a real function of time (as noted in Ref. [170]); this is, e.g., the case if $f(p)$ corresponds to a Lorentzian distribution centred around zero,

$$|f(p)|^2 = \frac{\Gamma}{\pi(\Gamma^2 + p^2)} \mapsto k(t) = e^{-2\Gamma|t|}. \quad (3.35)$$

A priori, the fact that there are no $\hat{\sigma}_x$ -coherences created in the free evolution does not mean that none are created if the system is probed at intermediate times. However, here, no $\hat{\sigma}_x$ -coherence is generated even when we take into account how the measurements modify the system's state. Specifically, immediately after a measurement in the $\hat{\sigma}_x$ -basis is performed at time t_1 (yielding outcome \pm), the total system-environment state is of product form

$$\eta^{(\pm)}(t_1) = |\pm\rangle\langle\pm| \otimes \xi^{(\pm)}(t_1), \quad (3.36)$$

where $\xi^{(\pm)}(t_1)$ is a state of the environment that depends on the measurement outcome. As we show in Appendix C.2, any state of the system evolved from the post-measurement state of Eq. (3.36) according to the described dynamics remains diagonal in the $\{|\pm\rangle\}$ basis; this also holds true for the state of the system after any sequence of such measurements. Together with the fact that the statistics resulting from measurements in the $\{|\pm\rangle\}$ basis is non-classical (i.e., it does not satisfy Kolmogorov conditions, as has been shown in Ref. [170]), this constitutes an example of a non-classical process without any coherence with respect to the measured observable ever being generated. Evidently, this behaviour is only possible since the chosen example is non-Markovian.

Unlike in the Markovian case, where the absence of coherences trivially leads to classical statistics, when memory effects are present, it is the coherences of the system state as well as the non-classical correlations between the system and its environment that can lead to non-classical behaviour—in a way which will be specified in the following. Intuitively, while the completely dephasing map leaves the system unchanged if no coherences are created, it does not necessarily leave the overall system-environment state invariant. In detail, in general we can have $\Delta[\rho_{t_j}^s] = \mathcal{I}[\rho_{t_j}^s] \forall t_j$, without it implying $\Delta \otimes \mathcal{I}^e[\eta_{t_j}^{se}] = \mathcal{I}[\eta_{t_j}^{se}] \forall t_j$. As we will see, the latter property is sufficient, but not necessary, for the satisfaction of the

Kolmogorov conditions. First, though, in order to be able to go beyond the investigation of Markovian processes, and extend the existing connection between classicality and coherences, it is important to introduce *quantum combs*—a suitable framework to describe general quantum processes [126, 129].

3.5 Non-Markovian Classical Processes

The previous example illustrates the subtle relation between coherence and classicality in the case of open quantum processes with memory. There, although no coherence is ever generated on the level of the system with respect to the chosen measurement basis, the system-environment correlations built up throughout the dynamics lead to non-classical statistics. To develop a more in-depth understanding of the interplay between coherences and classical phenomena, we require a suitable operational framework for approaching such scenarios. We can then employ this framework to comprehensively characterise all quantum processes that display classical statistics.

3.5.1 Classicality and Processes with Memory

The necessity of such a novel framework for the description of quantum processes that display memory effects stems from the breakdown of their modelling in terms of propagators that could be used in the Markovian case; this can already be seen for classical stochastic processes. Here, a joint probability distribution $\mathbb{P}_K(x_K, \dots, x_1)$ fully describes a K -process. This probability distribution can equivalently be represented in terms of multi-time conditional probabilities as

$$\mathbb{P}_K(x_K, \dots, x_1) = \mathbb{P}_K(x_K | x_{K-1}, \dots, x_1) \cdots \mathbb{P}_2(x_2 | x_1) \mathbb{P}_1(x_1). \quad (3.37)$$

Importantly, all of the above conditional probabilities generally depend upon all preceding measurement results, in contrast to the Markovian case where they only depend on the most recent outcome. Consequently, two-point transition probabilities of the form $\mathbb{P}(x_j | x_{j-1})$ are not sufficient in general to build up all joint probability distributions and thus do not completely describe the process. Similarly, two-time propagators $\{\Lambda_{t_j, t_{j-1}}\}$ are generally not sufficient to compute multi-time joint probabilities in the quantum case and therefore fail to fully characterise the process [107, 155].

For classical statistics, the joint probability distribution $\mathbb{P}_K(x_K, \dots, x_1)$ contains all information about the K -process, since all distributions for fewer times, as well as all conditional probabilities, can be derived once \mathbb{P}_K is known. In exactly the same way, a general quantum K -process is fully characterised by the joint probabilities for *all possible sequences* of K measurements (at times t_1, \dots, t_K), including non-projective and non-orthogonal ones.

As discussed in the previous section, if the complete system-environment dynamics is known, then all joint probability distributions (on times $\{t_j\}_{j=1}^n$) obtained from sequential measurements of the system can be computed via (recall Fig. 3.2)

$$\mathbb{P}_n(x_n, \dots, x_1) = \text{tr} \left[(\mathcal{P}_{x_n} \otimes \mathcal{I}^e) \circ \mathcal{U}_{t_n, t_{n-1}} \circ \dots \circ (\mathcal{P}_{x_1} \otimes \mathcal{I}^e) [\eta_{t_1}^{se}] \right]. \quad (3.38)$$

Here, $\{\mathcal{P}_{x_j}\}$ correspond to projective measurements in the computational basis, but evidently the same relation can also be used to obtain the correct probabilities when using different probing instruments, e.g., those that measure sharply in a different basis or those that perform generalised measurements. More formally, an *instrument* $\mathcal{J}_k = \{\mathcal{M}_{x_k}\}$ (at time t_k) is a collection of CP maps that add up to a CPTP map [276]. For instance, the instrument corresponding to a measurement in the computational basis is given by $\mathcal{J}_k = \{\mathcal{P}_{x_k}\}$, and all of its elements add up to the CPTP map $\sum_{x_k} \mathcal{P}_{x_k} = \Delta_k$. Intuitively, each outcome of an instrument corresponds to one of its constituent CP maps, which, in turn, describes how the state of the system changes upon the realisation of a specific measurement outcome. With this, the probability to obtain the sequence of outcomes x_1, \dots, x_K , given that the instruments $\mathcal{J}_1, \dots, \mathcal{J}_K$ were used to probe the system, is given by

$$\begin{aligned} & \mathbb{P}_K(x_K, \dots, x_1 | \mathcal{J}_K, \dots, \mathcal{J}_1) \\ &= \text{tr} \left[(\mathcal{M}_{x_K} \otimes \mathcal{I}^e) \circ \mathcal{U}_{t_K, t_{K-1}} \circ \dots \circ (\mathcal{M}_{x_1} \otimes \mathcal{I}^e) [\eta_{t_1}^{se}] \right] \\ &=: \mathcal{C}_K[\mathcal{M}_{x_K}, \dots, \mathcal{M}_{x_1}]. \end{aligned} \quad (3.39)$$

Moreover, the joint probability distribution for any subset of ordered times $t_n \geq \dots \geq t_1$, with $n < K$, can be obtained by replacing in the formula above \mathcal{M}_{x_j} with the identity operator, in correspondence with the times t_j not contained in the subset.

In what follows, whenever we drop the explicit instrument labels, it is understood that the probabilities were the result of a measurement in the computational basis at each time. The multi-linear functional \mathcal{C}_K introduced above is a special case³ of a *quantum comb* [125, 126] and provides a natural generalisation to the concept of quantum channels that by construction allows for the inclusion of memory effects [128, 129, 291, 292] (see Fig. 3.5 for a graphical representation).

It maps any sequence of possible experimental transformations enacted on the system to the corresponding joint probability of their occurrence. In this sense, \mathcal{C}_K plays exactly the same role that the joint probability distribution \mathbb{P}_K plays in the classical setting, and thus allows one to decide on the classicality of the resulting statistics. For example, for the completely memoryless case, i.e, the case of Markovianity with respect to measurements in

³In contrast to the combs discussed in Refs. [125, 126], the combs we consider do not start on an open input line, and do not end on an open output line; or, equivalently, in our case, the Hilbert spaces of this initial input and final output space are trivial. Such combs are also called *testers* in the literature.




Figure 3.5: Comb of a General Open Quantum Evolution. The probabilities characterising a quantum process can be understood as the action of a comb \mathcal{C}_K on the sequence of CP maps $\{\mathcal{M}_{x_j}\}$ that correspond to the respective measurement outcomes.

any basis, the evolution between any two points in time is described solely by a sequence of independent CPTP maps that act on the system alone [128, 180], and we have

$$\mathcal{C}_K^{\text{Markov}}[\mathcal{M}_{x_K}, \dots, \mathcal{M}_{x_1}] = \text{tr} \left[\mathcal{M}_{x_K} \circ \Lambda_{t_K, t_{K-1}} \circ \dots \circ \mathcal{M}_{x_2} \circ \Lambda_{t_2, t_1} \circ \mathcal{M}_{x_1}[\rho_{t_1}] \right]. \quad (3.40)$$

In general, however, the comb of a K -process does not split in the way above into independent portions of evolution between times. Thus, when analysing the relation between coherence and classicality in the presence of memory, instead of investigating the properties of individual CPTP maps, one must consider those of the multi-time comb \mathcal{C}_K .

The comb \mathcal{C}_K is an operationally well-defined object that can—just like the joint probability distribution \mathbb{P}_K —be obtained by means of probing measurements on the system alone through a generalised tomographic scheme [129, 150]. Specifically, for its reconstruction, it is not necessary to explicitly know the system-environment dynamics: The comb does not contain direct information about the environment, but solely that of its influence on the multi-time statistics observed from measurements on the system. As such, it encapsulates all that is out of control of the experimenter and thereby clearly separates the underlying process at hand from what can be controlled (i.e., the experimental interventions). An explicit example of the comb formalism is provided in Appendix C.3, where we rephrase Example 3.1 in terms of the comb description.

Crucially, the comb framework allows us to consider what it means for a stochastic process with memory to be classical, thereby permitting an extension of the results of Ref. [170] to the non-Markovian case: Given the comb \mathcal{C}_K of a process on times in \mathcal{T} , *all* combs correctly describing the process on fewer times $\mathcal{T}' \subseteq \mathcal{T}$ can be deduced by letting \mathcal{C}_K act on the identity map at the appropriate superfluous times [129, 146]. For example, we have (see also Fig. 3.6)

$$\begin{aligned} & \mathcal{C}_{K-1}[\mathcal{M}_{x_K}, \dots, \mathcal{M}_{x_{j+1}}, \mathcal{M}_{x_{j-1}}, \dots, \mathcal{M}_{x_1}] \\ &= \mathcal{C}_K[\mathcal{M}_{x_K}, \dots, \mathcal{M}_{x_{j+1}}, \mathcal{I}_j, \mathcal{M}_{x_{j-1}}, \dots, \mathcal{M}_{x_1}]. \end{aligned} \quad (3.41)$$




Figure 3.6: Consistency Condition for Combs. Letting a comb defined on times \mathcal{T} act on identity maps at a set of times $\mathcal{T} \setminus \mathcal{T}'$ (i.e., the set of times in \mathcal{T} but not in \mathcal{T}') yields the correct comb on \mathcal{T}' . Depicted is the situation for $\mathcal{T} = \{t_1, t_2, t_3\}$ and $\mathcal{T}' = \{t_1, t_3\}$.

As we have discussed in the previous sections, classicality of a process means that the action of the completely dephasing map cannot be distinguished (by means of measurements in the classical basis) from not performing an operation. With the method of “generalised marginalisation” given by Eq. (3.41) at hand, we obtain the following characterisation of classical combs:

Theorem 3.2 (K -classical quantum combs). *A comb \mathcal{C}_K on times \mathcal{T} , with $|\mathcal{T}| = K$, yields a K -classical process via Eq. (3.38) iff it satisfies*

$$\mathcal{C}_K \left[\bigotimes_{t_j \in \mathcal{T}'} \mathcal{I}_j, \bigotimes_{t_k \in \mathcal{T} \setminus \mathcal{T}'} \mathcal{P}_{x_k} \right] = \mathcal{C}_K \left[\bigotimes_{t_j \in \mathcal{T}'} \Delta_j, \bigotimes_{t_k \in \mathcal{T} \setminus \mathcal{T}'} \mathcal{P}_{x_k} \right], \quad (3.42)$$

for all subsets $\mathcal{T}' \subseteq \mathcal{T}$ and all possible sequences of outcomes on $\mathcal{T} \setminus \mathcal{T}'$.

In slight abuse of notation, here, the argument $\bigotimes_{t_j \in \mathcal{T}'} a_j, \bigotimes_{t_k \in \mathcal{T} \setminus \mathcal{T}'} b_{x_k}$ of the comb \mathcal{C}_K signifies that it acts on the maps a_j at times $t_j \in \mathcal{T}'$ and on b_{x_k} at times $t_k \in \mathcal{T} \setminus \mathcal{T}'$.

Theorem 3.2 expresses in a concise way that a general process is K -classical iff measurements in the computational basis cannot distinguish the action of completely dephasing maps from the action of identity maps. Let us emphasize again that the completely dephasing map does not only destroy coherences of the systems reduced state, but also quantum correlations between the system and the environment. Therefore, Theorem 3.2 does not directly link coherence and non-classicality as Theorem 3.1 did for the case without memory.

Proof. The proof of Theorem 3.2 is thus straightforward: If a comb satisfies Eqs. (3.42), then the resulting statistics satisfy Kolmogorov conditions. Conversely, any joint probability distribution on a set of times $\mathcal{T}' \subseteq \mathcal{T}$ can either be obtained by direct measurement, or by marginalisation of the corresponding distribution on \mathcal{T} . The former can be computed via the first line of Eq. (3.42), the latter via the second one. If the statistics of the process appear classical, then both resulting distributions have to coincide, and Eq. (3.42) must hold. \square

In the (basis dependent) Markovian case that we discussed in the previous section, Eq. (3.42) directly reduces to Eq. (3.28), the NCGD property at the level of propagators

of populations. Theorem 3.2 therefore provides the proper generalisation of the results of Ref. [170] to the non-Markovian case. Nonetheless, its consequences for the structural properties of classical combs, and, in particular, the relation of classicality and coherence remain somewhat opaque in the way Theorem 3.2 is phrased. In order to address these questions, we now introduce a representation of quantum combs that is favourable for the purposes of our work.

3.5.2 Choi-Jamiołkowski Representation of General Quantum Processes

Both the quantum comb describing the K -process at hand and the experimental interventions applied at each time are linear maps (the former being a higher-order multi-linear map). Any such map can be represented in a variety of ways, but the most natural for our present purposes makes use of the Choi-Jamiołkowski isomorphism [293, 294] between quantum maps and positive semi-definite Hermitian matrices.

A general quantum map—e.g., one that corresponds to a generalised measurement—at time t_k is a CP transformation $\mathcal{M}_{x_k} : \mathcal{B}(\mathcal{H}_k^{\text{i}}) \rightarrow \mathcal{B}(\mathcal{H}_k^{\text{o}})$ that takes bounded linear operators on the (input) Hilbert space \mathcal{H}_k^{i} onto bounded linear operators on the (output) Hilbert space \mathcal{H}_k^{o} . Throughout this paper, we will consider the input and output spaces of such maps to be isomorphic (and of finite dimension), and the labels i and o , as well as the time label, are merely introduced for better accounting of the involved spaces. Any such quantum map \mathcal{M}_{x_k} can be isomorphically mapped onto a positive semi-definite Hermitian matrix that we will call its *Choi state*, $M_{x_k} \in \mathcal{B}(\mathcal{H}_k^{\text{o}} \otimes \mathcal{H}_k^{\text{i}})$, by letting it act on one half of an unnormalised maximally-entangled state $\Phi^+ = \sum_{x_k, y_k} |x_k x_k\rangle \langle y_k y_k| \in \mathcal{B}(\mathcal{H}_k^{\text{i}} \otimes \mathcal{H}_k^{\text{i}})$, i.e.,

$$M_{x_k} := (\mathcal{M}_{x_k} \otimes \mathcal{I})[\Phi^+] \in \mathcal{B}(\mathcal{H}_k^{\text{o}} \otimes \mathcal{H}_k^{\text{i}}). \quad (3.43)$$

This isomorphism implies, e.g., the following identifications:

$$\text{Identity Map :} \quad \mathcal{I}_k \Leftrightarrow \Phi_k^+, \quad (3.44)$$

$$\text{Proj. Map :} \quad \mathcal{P}_{x_k} \Leftrightarrow |x_k\rangle \langle x_k| \otimes |x_k\rangle \langle x_k|, \quad (3.45)$$

$$\text{C. Deph. Map :} \quad \Delta_k \Leftrightarrow \sum_{x_k} |x_k x_k\rangle \langle x_k x_k| := D_k. \quad (3.46)$$

Here and throughout this article, we typically denote maps with calligraphic upper-case letters (as we have already done above) and their Choi state with the corresponding non-calligraphic variant—with the exception of the identity map [Eq. (3.44)] and the completely dephasing map [Eq. (3.46)]. For better orientation, we will continue to denote the respective time at which the maps act by an additional subscript.

Analogously, as a quantum comb \mathcal{C}_K is a multi-linear map it can—in a similar way to Eq. (3.43)—be mapped onto a positive semi-definite Hermitian matrix C_K [126, 129, 264]. The action of a quantum comb on a sequence of CP maps $\{\mathcal{M}_{x_K}, \dots, \mathcal{M}_{x_1}\}$ is then

equivalently given by [126]

$$\mathcal{C}_K[\mathcal{M}_{x_K}, \dots, \mathcal{M}_{x_1}] = \text{tr} \left[(M_{x_K}^T \otimes \dots \otimes M_{x_1}^T) C_K \right], \quad (3.47)$$

where \bullet^T denotes the transposition with respect to the computational basis. Eq. (3.47) constitutes the Born rule for temporal processes [125, 295], where C_K plays the role of a quantum state over time and the Choi states M_{x_K}, \dots, M_{x_1} play the role that Positive Operator-Valued Measure (POVM) elements play in the standard Born rule.

Concretely, given an instrument sequence $\mathcal{J}_K, \dots, \mathcal{J}_1$, by combining Eqs. (3.39) and (3.47), the joint probability over the sequence of outcomes x_K, \dots, x_1 is given by

$$\mathbb{P}_K(x_K, \dots, x_1 | \mathcal{J}_K, \dots, \mathcal{J}_1) = \text{tr} \left[(M_{x_K}^T \otimes \dots \otimes M_{x_1}^T) C_K \right]. \quad (3.48)$$

Through this isomorphism, memory effects of the temporal process correspond directly to structural properties of its Choi state [106, 130–132, 264]; analogously, the classicality of a process is reflected in the properties of C_K .

Represented in this way, quantum combs and the channels that they generalise have particularly nice properties. Complete positivity and trace preservation for a quantum channel \mathcal{M} correspond respectively to $M \geq 0$ and satisfaction of $\text{tr}_o[M] = \mathbb{1}_i$. Analogously the Choi state of a quantum comb has to satisfy $C_K \geq 0$ as well as a hierarchy of trace conditions that fix the causal ordering of events [126], i.e., they ensure that later events cannot influence the statistics of earlier ones.

It is important to note that all K -processes can be represented through the Choi-Jamiołkowski isomorphism as (unnormalised) quantum states C_K . In the converse direction, any operator satisfying the aforementioned properties admits an underlying open quantum dynamics description [125, 126, 129]. Specifically, this means that for every proper comb, there is a (possibly fictitious) environment and a set of system-environment unitaries such that the action of the comb on any sequence of instruments can be written as in Eq. (3.39). Quantum combs are hence the most general descriptors of open quantum system processes (when the system of interest is probed at fixed times). We will call the respective underlying unitary description that includes the environment the dilation of the comb. As is the case for quantum channels, any such dilation is non-unique. On the other hand, the comb \mathcal{C}_K resulting from some underlying evolution is unique, and—just like the joint probability distribution \mathbb{P}_K in the classical case—constitutes the maximal descriptor of the process on the respective set of times.

3.5.3 Structural Properties of Classical Combs

As a first step to a structural understanding of classical combs, we rephrase Theorem 3.2 in terms of Choi states:

Theorem 3.2' (*K*-classical quantum combs). *A comb C_K on times \mathcal{T} , with $|\mathcal{T}| = K$, yields a *K*-classical process iff its Choi state satisfies*

$$\mathrm{tr} \left[\left(\bigotimes_{t_j \in \mathcal{T}'} \Phi_j^+ \bigotimes_{t_k \in \mathcal{T} \setminus \mathcal{T}'} P_{x_k} \right) C_K \right] = \mathrm{tr} \left[\left(\bigotimes_{t_j \in \mathcal{T}'} D_j \bigotimes_{t_k \in \mathcal{T} \setminus \mathcal{T}'} P_{x_k} \right) C_K \right]. \quad (3.49)$$

for all subsets $\mathcal{T}' \subseteq \mathcal{T}$ and all possible sequences of outcomes on $\mathcal{T} \setminus \mathcal{T}'$.

Using the relations (3.44) – (3.46) as well as Eq. (3.48), it is straightforward to see that this theorem is indeed equivalent to Theorem 3.2. Importantly, as it is stated in terms of Choi states, Theorem 3.2' allows one to derive a direct connection between general correlations and the classicality of a *K*-process.

To see how the requirement in Eq. (3.49) translates to structural constraints on classical combs, first note that any comb that yields the joint probability distribution $\mathbb{P}_K(x_K, \dots, x_1)$ when probed in the classical basis can be written as

$$C_K = \tilde{C}_K^{\mathrm{Cl.}} + \chi, \quad (3.50)$$

where the term

$$\tilde{C}_K^{\mathrm{Cl.}} = \sum_{x_K, \dots, x_1} \mathbb{P}_K(x_K, \dots, x_1) P_{x_K} \otimes \dots \otimes P_{x_1}, \quad (3.51)$$

contains the joint probability distribution \mathbb{P}_K on its diagonal and $\mathrm{tr}[(P_{x_K} \otimes \dots \otimes P_{x_1})\chi] = 0$ for all x_K, \dots, x_1 [150]. Intuitively, $\tilde{C}_K^{\mathrm{Cl.}}$ corresponds to the part of C_K that can be probed by measurements in the classical basis alone, while χ contains all the information about the underlying process that such measurements are blind to. If $\chi = 0$, then C_K clearly satisfies the conditions of Eq. (3.42), as $\mathrm{tr}[P_{x_j} \Phi_j^+] = \mathrm{tr}[P_{x_j} D_j]$ for all x_j .⁴ In words, for $\chi = 0$, the corresponding comb is classical, as it is diagonal in the classical product basis. However, this is not necessary for Eq. (3.42) to hold; rather, it suffices if χ is such that it does not allow one to distinguish between the action of the identity map and the completely dephasing map. We thus arrive at the following lemma:

Lemma 3.1. *Let C_K be the comb of a *K*-process on \mathcal{T} , with $|\mathcal{T}| = K$, and let $A_j := \Phi_j^+ - D_j$. C_K yields a *K*-classical process iff it is of the form*

$$C_K = \tilde{C}_K^{\mathrm{Cl.}} + \chi, \quad (3.52)$$

where $\tilde{C}_K^{\mathrm{Cl.}}$ is obtained from some joint probability distribution \mathbb{P}_K via Eq. (3.51) and χ satisfies

$$\mathrm{tr} \left[\left(\bigotimes_{t_j \in \mathcal{T}'} A_j \bigotimes_{t_k \in \mathcal{T} \setminus \mathcal{T}'} P_{x_k} \right) \chi \right] = 0 \quad (3.53)$$

for all subsets $\mathcal{T}' \subseteq \mathcal{T}$ and $\mathcal{T}' = \emptyset$.

⁴For $\chi = 0$, C_K is actually not a proper comb, as it does not satisfy the hierarchy of trace conditions that ensure causal ordering. Nonetheless, this lack of causality could not be picked up by means of projective measurements in the classical basis alone, and does thus not pose a problem for our discussion.




Figure 3.7: Nested Set of Processes. Processes that cannot produce coherence and destroy any coherence that is fed in (i.e., their Choi states are diagonal in the computational basis) form a strict subset of processes that appear classical when sequentially probed in the computational basis. Both of these sets, as well as the set of all quantum processes, are convex.

Proof. It is straightforward to see that a comb of the form of Eq. (3.52) satisfies Eq. (3.49), whenever χ fulfills Eq. (3.53), and thus yields K -classical statistics. Conversely, *any* comb C_K on K times can be written as $C_K = \tilde{C}_K^{\text{Cl.}} + \chi$, where $\tilde{C}_K^{\text{Cl.}}$ is of the form of Eq. (3.51) for some \mathbb{P}_K and $\text{tr}[(P_{x_K} \otimes \cdots \otimes P_{x_1})\chi] = 0$ [150]. When measuring (in the computational basis) at K times, the resulting joint probability distribution is given by \mathbb{P}_K . As, by assumption, the process is classical, summation over outcomes obtained at any time C_K is defined on must yield the same statistics as letting the comb act on the identity channel at this time. As this has to hold for any collection of times in \mathcal{T} , χ has to satisfy the additional requirements given by Eq. (3.53). \square

Intuitively, Eq. (3.53) ensures that the action of Δ_j cannot be detected at any point in time by means of measurements in the classical basis. Therefore Lemma 3.1 is equivalent to Theorem 3.2. However, the former provides an explicit constraint on the structure of such combs that contain coherences that can be present in the process without making the resulting statistics non-classical.

Indeed, if $\chi = 0$, then the corresponding comb C_K is diagonal in the classical product basis and, as such, cannot create coherences and destroys any kind of coherences that could be fed into the process (e.g., by performing coherence creating operations at some time). On the other hand, if $\chi \neq 0$ and the comb contains off-diagonal terms (with respect to the classical basis), then coherences can be created over the course of the process. However, if χ satisfies Eq. (3.53), then these coherences—or rather the invasiveness of the completely dephasing map—cannot be detected at any later time by measurements in the classical basis. This understanding of classical non-Markovian combs mirrors the intuition we had built in the Markovian setting for the case of NCGD dynamics. Consequently, Lemma 3.1 fully characterises the relation between coherences and the non-classicality of a process (see Fig. 3.7 for a graphical representation of the different sets of processes we consider).

Somewhat unsurprisingly, the above lemma implies that combs leading to classical processes are of measure zero in the set of all combs: While any comb can be written in the form of Eq. (3.52), Eq. (3.53) places further linear constraints on the χ term, which

must be satisfied by combs leading to classical processes, but not by general combs. The set of combs leading to classical processes is thus confined to a lower dimensional subset, implying that it is of zero measure (with respect to any reasonable measure in the set of all non-Markovian combs). This fact falls in line with the intuition built above; for a randomly chosen comb, the action of a completely dephasing map in a given basis will generally be detectable. Furthermore, the vanishing volume of classical combs within the set of all combs mirrors the analogous property in the spatial setting: There, quantum states that display no discord are of measure zero in the set of all bipartite quantum states [296] (the relation between quantum discord and classicality of processes is discussed in detail in Section 3.6).

In the non-Markovian case, the characterisation of classical processes comes at a price. In order to decide on the K -classicality of a given process, it is no longer sufficient to investigate propagators between pairs of times, but rather the full part of the comb C_K that is relevant for sequential projective measurements must be known, due to the importance of multi-time effects. However, this behaviour is to be expected, as can already be seen in the case of classical stochastic processes: The full characterisation of a non-Markovian process only happens on the level of the full joint probability distribution \mathbb{P}_K , and not by way of transition probabilities between adjacent times only. Despite the additional complexity brought in by the presence of memory, as we will see in the following section, measures for classicality that are both experimentally and computationally accessible can be derived based on the characterisation of classical processes we have provided.

3.5.4 Quantifying Non-Classicality

As we have seen above, the set of combs leading to classical processes is of measure zero in the set of all combs. Importantly though, this fact does not render our original definition of classicality meaningless, but rather—in conjunction with Lemma 3.1—allows for the derivation of a meaningful measure of non-classicality that is experimentally accessible and can be formulated by means of a Linear Program (LP).

More specifically, we can exploit the characterisation of classical processes provided by Eqs. (3.52) and (3.53) in order to define a measure of non-classicality with a clear operational meaning. Such a measure not only classifies whether or not a comb is non-classical, but also quantifies the degree to which it is. This is crucial when assessing whether any potential non-classicality arises from inherently quantum features of the experiment or from experimental errors. In order to clarify its operational interpretation, we formulate our measure in the context of a game with two adversaries, Alice and Bob, and one referee, Rudolph. The task of Alice is to construct a classical stochastic process that is a good model for a comb she receives from Rudolph. The task of Bob is to design a test that distinguishes this model from the original comb. Let C be the given comb in its

Choi representation (i.e., a positive operator with some additional causality constraints). The game then proceeds as follows:

0) Rudolph begins with a given comb C and sends its description to both Alice and Bob.

A) Alice prepares a classical process $C^{\text{Cl.}}$ and sends it to Rudolph.

R1) Rudolph sends the description of the classical process $C^{\text{Cl.}}$ prepared by Alice to Bob.

B1) Bob prepares a testing sequence $\{T_i(\vec{x})\}_{\vec{x}}$ and sends it to Rudolph.

R2) Rudolph takes randomly either C or $C^{\text{Cl.}}$ and applies the testing sequence chosen by Bob. He yields an outcome \vec{x} , which he announces.

B2) Bob announces whether the comb is C or $C^{\text{Cl.}}$

R3) Rudolph announces whether Bob is correct or not and hence who wins the game.

Let us recall at this point that our definition of classicality relies exclusively on the statistics obtained by probing the process with projective measurements in fixed, orthonormal bases. Therefore, to only probe what is relevant within our framework, we restrict the testing sequences that Bob is allowed to prepare to only involve such measurements, i.e., the testing sequence must be of the form $T_i(\vec{x}) = \bigotimes_{t_j \in \tau_i} \Phi_j^\dagger \bigotimes_{t_k \in \tau_i^c} P_{x_k}$.

The figure of merit that we are interested in is the probability for Bob to win if both players play optimally. This is an operational quantity describing how well said comb can be distinguished from its best classical approximation, given that one has only access to the aforementioned restricted testing strategies that can be used to probe classicality. Making use of the arguments of Lemma 3.1 to simplify the structure of the classical combs, in Appendix C.5 we derive this quantity; here we simply present the main results.

The probability for Bob winning the game is given by:

$$\mathbb{P}_B(C) = \frac{1}{2} [1 + M(C)], \quad (3.54)$$

with $M(C)$ being one half of the solution of

$$\begin{aligned} \text{minimise: } & \max_i \sum_{\vec{x}} \left| \text{tr}[(C^{\text{Cl.}} - C)T_i(\vec{x})] \right| \\ \text{subject to: } & C^{\text{Cl.}} = \sum_{y_K, \dots, y_1} \mathbb{P}_K(\vec{y}) P_{y_K} \otimes \dots \otimes P_{y_1}, \\ & \mathbb{P}_K(\vec{y}) \text{ joint probability distribution.} \end{aligned} \quad (3.55)$$

This can be transformed into the following linear program (and hence can be solved efficiently numerically; the error can be estimated and one can compute the optimal $C^{\text{Cl.}}$

and $T_i(\vec{x})$ [297]):

$$\begin{aligned}
& \text{minimise: } a \\
& \text{subject to: } \sum_j b_{ij} - a \leq 0, \\
& \quad \sum_k p_k \alpha_{ijk} - \beta_{ij} - b_{ij} \leq 0, \\
& \quad - \sum_k p_k \alpha_{ijk} + \beta_{ij} - b_{ij} \leq 0, \\
& \quad \sum_k p_k - 1 = 0, \\
& \quad p_k \geq 0, a \geq 0, b_{ij} \geq 0,
\end{aligned} \tag{3.56}$$

where we have defined $\alpha_{ijk} := \text{tr}[(P_{y_K(k)} \otimes \cdots \otimes P_{y_1(k)})T_i(\vec{x}_j)]$, $\beta_{ij} := \text{tr}[CT_i(\vec{x}_j)]$ and $p_k := \mathbb{P}_K(\vec{y}(k))$. For completeness we also give the dual program, which by definition turns a minimisation into a maximisation. The dual problem is useful to give bounds on the found solution, to solve the problem, and potentially to find different interpretations of the quantity in question. The dual of the program above can be formulated as:

$$\begin{aligned}
& \text{maximise: } Z \\
& \text{subject to: } Z \leq \sum_{ij} (\alpha_{ijk} - \beta_{ij}) (2Y_{ij} - X_i) \quad \forall k, \\
& \quad \sum_i X_i = 1, \\
& \quad X_i, Y_{ij}, X_i - Y_{ij} \geq 0, \\
& \quad Z \in \mathbb{R}.
\end{aligned} \tag{3.57}$$

It follows directly from the interpretation as the solution of the game defined above that the quantity $M(C)$ is faithful, i.e., its value is zero if the statistics is classical, and that it measures how difficult it is to simulate the given comb by a classical stochastic process. As such, it provides us with a properly motivated quantifier of the degree of non-classicality of quantum processes, which describes how well the obtained statistics can be simulated by a classical process.

The full evaluation of $M(C)$ would, in principle, require testing over every sequence of projective measurements [to compute the maximisation in Eq. (3.55)] and the comparison with every classical multi-time probability distribution [to compute the minimisation in Eq. (3.55)]. Practically, it is then useful to consider bounds to this quantifier of non-classicality, which can be accessed via a limited number of measurements. In particular, lower bounds can be obtained by using a subset of measurement sequences $T_i(\vec{x})$ (in a similar way as to how one can use entanglement witnesses to construct bounds on meaningful entanglement measures [166, 298–300]). If such a lower bound is non-zero, this is already sufficient to conclude that the comb is non-classical. On the other hand,

upper bounds can be attained by restricting our consideration to some classical combs. As a relevant example, for any given comb C one can focus on a single classical comb \bar{C}^{Cl} , realised by applying a dephasing map before and after each measurement. This yields the statistics resulting from the marginals of the joint statistics one would obtain by measuring at every time. Note that, while this specific choice of a classical comb only provides us with an upper bound on our measure defined above, it is nonetheless faithful. In the simplest case where only two times are involved, $K = 2$, one can easily see that by replacing C^{Cl} with \bar{C}^{Cl} in Eq. (3.55), we derive the following upper bound

$$M(C) \leq \sum_{x_2} \left| \mathbb{P}(x_2) - \sum_{x_1} \mathbb{P}(x_2, x_1) \right|. \quad (3.58)$$

Such a natural quantifier of non-classicality has already been used to investigate coherence properties in transport phenomena [301] and, more recently, to control the departure from any classical random walk via the manipulation of quantum coherence in a time-multiplexed quantum walk experiment [269]. Let us note at this point that the experimental data that was used in Ref. [269] to evaluate the right hand side of Eq. (3.58) allows one to calculate $M(C)$ too. Hence, $M(C)$ can be evaluated without further acquisition of experimental data, which demonstrates the applicability of our measure to current experiments. In addition, our measure—or lower bounds thereof—can be employed to investigate more complex experiments with $K > 2$.

3.6 Dynamical Properties of Classical Quantum Processes

Theorem 3.2 and Lemma 3.1 provide a full characterisation of processes that yield classical statistics. Together, they allow for the derivation of classically testable quantifiers of non-classicality. For further clarification, and in order to connect non-classical processes to the respective underlying evolution, we now discuss some concrete cases of underlying non-Markovian dynamics that lead to classical statistics. Moreover, we will connect the classicality of temporal processes to vanishing quantum discord in the joint state of the system and the environment.

3.6.1 Discord and Classicality

Recall that in the Markovian case, the classicality of a process can be decided solely in terms of propagators between pairs of times that are defined on the system of interest alone and it is linked to the ability of those maps to create and detect coherences. In particular, the set of dynamics that does not create coherences on the level of the system is contained in the set of maps that lead to classical statistics [170]. As we have seen above, this fails to hold in the non-Markovian case, where, even if the state of the system is diagonal in the computational basis at all times, i.e., no coherence on the system level is ever generated, the statistics might not satisfy the Kolmogorov conditions.

As soon as memory effects play a non-negligible role, it is both the coherences of the system state and the correlations between the system and its environment that can lead to non-classical behaviour. It is thus desirable to derive a more explicit relation between coherence, correlations and classicality.

To do so, first recall that while the completely dephasing map leaves the system unchanged if the state of the system is classical at all times, it does not necessarily leave the overall system-environment state, which, at every time t_j contains all relevant memory, invariant. Specifically, in this case we have $\Delta_j[\rho_{t_j}^s] = \mathcal{I}[\rho_{t_j}^s] \forall t_j$ but not necessarily $\Delta_j \otimes \mathcal{I}_j^e[\eta_{t_j}^{se}] = \mathcal{I}^{se}[\eta_{t_j}^{se}] = \forall t_j$. While the latter is not necessary for the satisfaction of the Kolmogorov conditions, it is sufficient:

Lemma 3.2. *Let $\{p_{t_i}^m\}$ be sets of probabilities that sum to unity, $\{\Pi_j^m\}$ orthogonal projectors (not necessarily rank-1) on the system that are diagonal in the computational basis, and $\{\xi_j^m\}$ states on the environment. If at all times $t_j \in \mathcal{T}$, with $|\mathcal{T}| = K$, the system-environment state is of the form*

$$\eta_{t_j}^{se} = \sum_m p_{t_j}^m \Pi_j^m \otimes \xi_j^m, \quad (3.59)$$

then the underlying process is K -classical, i.e., it satisfies the Kolmogorov conditions of Eq. (3.9).

Note that we assume the computational basis to be the same at every time, so that the additional subscript of Π_j^m is somewhat superfluous and merely added to clearly signify the respective time at which the state is defined. In principle, one could define classicality with respect to projective measurements in different bases at each time t_j , in which case the additional subscript of Π_j^m would denote projectors in different bases, and the above lemma would still hold. Analogously, all other results of this paper can straightforwardly be adapted to such more general probing schemes, but for simplicity, we understand classicality with respect to a fixed basis that does not change in time (the only exception being Section 3.7, where we will extend the setting to allow for arbitrary measurement schemes in order to examine the nature of genuinely quantum processes.). Naturally, the environment states ξ_j^m in Eq. (3.59) can be diagonal in arbitrary bases, as it is only invasiveness with respect to measurements on the system that we are concerned with.

Before we prove Lemma 3.2, it is insightful to discuss the relation between the concept of classical temporal processes and the classical spatial system-environment correlations it introduces. Firstly, recall the full system-environment state at each time encapsulate all memory effects. Concretely, in contrast to the state of the system alone, they contain all information that is relevant to predict the future statistics. In particular, for states of the form given in Eq. (3.59), at each time t_j , this memory is stored in the probabilities $\{p_{t_j}^m\}$ and the environment states $\{\xi_j^m\}$. States of said form have vanishing quantum discord [265–268, 302], i.e., they do not display any genuinely quantum correlations

between the system and the environment. For a general zero-discord state, the set $\{\Pi_j^m\}$ in Eq. (3.59) could be any set of mutually orthogonal projectors, and the correlations between the system and the environment are considered to be classical, since there exists a measurement on the system with perfectly distinguishable outcomes that overall leaves the total state undisturbed [267, 268] (see also the proof below).

As we only consider measurements on the system in a fixed basis in our setting, here, vanishing discord at all times does not yet force the resulting statistics to be classical; rather, the discord must vanish in the correct basis, i.e., the one in which the experimenter’s measurements act. While discord is often considered as a basis independent quantity—obtained by a minimisation procedure over all possible measurement scenarios [268]—here, and throughout the remainder of this article, we will always consider its basis dependent formulation [265–268, 277, 303] and call states of the form in Eq. (3.59) discord-zero with respect to the classical basis. That is, whenever we consider a state to be of zero discord, we will always implicitly mean that it can be represented as per Eq. (3.59) with the projectors being diagonal in the classical basis of the measurements. Importantly, this basis dependence mirrors the basis dependence of coherence, which is also always defined with respect to a fixed classical basis.

Proof. For states of the form in Eq. (3.59), the completely dephasing map Δ on the system has the same effect as the “do-nothing” identity channel \mathcal{I} , i.e.,

$$\Delta_j \otimes \mathcal{I}_j^e \left[\sum_m p_{t_j}^m \Pi_j^m \otimes \eta_j^m \right] = \mathcal{I}_j^s \otimes \mathcal{I}_j^e \left[\sum_m p_{t_j}^m \Pi_j^m \otimes \eta_j^m \right]. \quad (3.60)$$

Consequently, if the system-environment state is of this form at all times, the resulting statistics satisfy the Kolmogorov conditions. \square

It is insightful to re-examine Example 3.1 in light of Lemma 3.2; there, we provided an example of a process for which the state of the system never displayed coherence, but nonetheless led to non-classical statistics. Consequently, the system-environment state must have non-zero (basis dependent) discord over the course of the dynamics:

Theorem 3.3. (3.1') *As we discuss in Appendix C.2, in Example 3.1, the system-environment state before the first measurement ($t < t_1$) is given by*

$$\begin{aligned} \rho_{se}(t) = \frac{1}{4} \sum_{i,j \in \{-,+\}} |i\rangle\langle j| \otimes \{ & i \cdot j |\varphi^-(t)\rangle\langle\varphi^-(t)| + i(2\alpha - 1) |\varphi^-(t)\rangle\langle\varphi^+(t)| \\ & + j(2\alpha - 1) |\varphi^+(t)\rangle\langle\varphi^-(t)| + |\varphi^+(t)\rangle\langle\varphi^+(t)| \}, \end{aligned} \quad (3.61)$$

where both

$$|\varphi^+(t)\rangle = \int_{-\infty}^{\infty} dp f(p) e^{ipt} |p\rangle \quad (3.62)$$

and

$$|\varphi^-(t)\rangle = \int_{-\infty}^{\infty} dp f(p) e^{-ipt} |p\rangle \quad (3.63)$$

are valid quantum states. This state has zero discord with respect to the eigenbasis of $\hat{\sigma}_x$ iff

$$|\varphi^+(t)\rangle\langle\varphi^+(t)| - |\varphi^-(t)\rangle\langle\varphi^-(t)| = 0 \quad (3.64)$$

and either $\alpha = \frac{1}{2}$ or

$$|\varphi^+(t)\rangle\langle\varphi^-(t)| - |\varphi^-(t)\rangle\langle\varphi^+(t)| = 0. \quad (3.65)$$

In the case of the Lorentzian distribution, it follows from

$$\langle\varphi^-(t)|\varphi^+(t)\rangle = k(t) = e^{-2\Gamma|t|} \quad (3.66)$$

that Eq. (3.64) cannot be satisfied for $t > 0$, i.e., basis dependent discord is created during the evolution (and subsequently destroyed by the measurement at t_1). Since the state of the system itself is not altered by the measurement, but the probabilities to obtain \pm at a later time are (as has been discussed in Ref. [170]), the discord necessarily must be converted into populations by the following portion of evolution. Below, we will examine this connection between the creation and detection of basis dependent discord and non-classicality in a rigorous manner.

If a state is of zero discord, it displays neither coherences on the level of the system nor non-classical correlations between the system and the environment, which is, to reiterate, sufficient for the classicality of the resulting process, but not necessary. In this sense, Lemma 3.2 is a direct extension of the analogous statement in the Markovian case; there, the absence of coherence in the system state at all times is also sufficient but not necessary for the process to be classical. Put differently, if all of the individual maps making up a Markovian dynamics are Maximally Incoherent Operations (MIO) [233, 242], i.e., they map all incoherent states onto incoherent states, then the resulting dynamics satisfies Kolmogorov conditions. However, MIOs are a strict subset of NCGD maps [170].

While somewhat intuitive, the above lemma sheds light on the properties that a general non-Markovian dynamics has to satisfy in order to appear classical. For system-environment states that are discord-zero in the computational basis (with respect to the system), a measurement on the system in the computational basis is non-invasive, i.e., it leaves the full state unchanged (and not just the system state, as it would be the case if the system state is incoherent at all times). For comprehensiveness, in Appendix C.6 we provide a characterisation of non-discord creating processes in terms of their dynamical building blocks.

In general, the absence of discord at all times is not necessary for a process to appear classical. However, what is necessary is that at no time can there be coherences or non-classical system-environment correlations that can be detected by means of measurements

in the computational basis at a later time. This mirrors the requirement for classical processes in the Markovian case, where the individual propagators have to be NCGD, i.e., the propagators must be such that they cannot create coherences whose existence can be picked up at a later time by means of measurements in the classical basis; yet, it is still possible that the individual maps create coherences [170]. NCGD maps are the fundamental building blocks that constitute classical Markovian combs. In what follows, utilising the connection of classicality and discord discussed above, we will provide a characterisation of the building blocks that make up classical non-Markovian processes.

3.6.2 Non-Discord-Generating-and-Detecting (NDGD) Dynamics and Classical Processes

In the Markovian case, classicality of a process can be decided on the level of CPTP maps, since in the absence of memory all higher order probability distributions can be obtained from the system state ρ_{t_1} and the two-time propagators $\{\Lambda_{t_j, t_{j-1}}\}$. It suggests itself to employ this intuition in the non-Markovian case, as every non-Markovian process corresponds to a Markovian one if enough additional degrees of freedom are taken into account.

In detail, as we discussed, every non-Markovian process can be dilated to a concatenation of a (potentially correlated) system-environment state and unitary total dynamics [126, 129], interspersed by the operations of the experimenter on the system alone that are performed at times $\{t_j\}$ (see Fig. 3.5 for reference). If the experimenter had access to all the degrees of freedom necessary for the dilation, then the underlying process would appear Markovian, and the results of Ref. [170] could be applied on the system-environment level for the characterisation of a classical process. Here, using the Markovian case as a guideline, we aim for a similar characterisation of classical processes when only the system degrees of freedom can be accessed.

To compactify notation and simplify later discussions, we can equivalently consider a general open process as a concatenation of CPTP maps that act on both the system and the environment, interspersed by the operations on the system alone. This way of describing general open system dynamics is simply a notational compression of the general case with global unitaries that allows for an easier connection to the Markovian case, but does not lead to a different set of possible combs. In what follows, we will denote these CPTP maps by $\Gamma_{t_j, t_{j-1}}$ to clearly distinguish them from the memoryless scenario (where the respective maps $\Lambda_{t_j, t_{j-1}}$ act only on the system), so that Eq. (3.11) generalises to

$$\mathbb{P}_n(x_n, \dots, x_1) = \text{tr} \left[(\mathcal{P}_{x_n} \otimes \mathcal{I}^e) \circ \Gamma_{t_n, t_{n-1}} \circ \dots \circ (\mathcal{P}_{x_1} \otimes \mathcal{I}^e) [\eta_{t_1}^{se}] \right]. \quad (3.67)$$

Moreover, for the sake of generality and to ease the comparison with the Markovian case, we allow for the state before the first measurement to be evolved from some other state at




Figure 3.8: NDGD System-Environment Dynamics. From the perspective of a classical observer performing projective measurements in a fixed basis, the identity map at any time t_j cannot be distinguished from the completely dephasing map. Any discord (with respect to the classical basis) that is present in the system-environment state, and/or created by the system-environment CPTP maps, cannot be detected by such a classical observer.

an initial reference time $t_0 \leq t_1$, i.e.,

$$\eta_{t_1}^{se} = \Gamma_{t_1, t_0} \eta_{t_0}^{se}; \quad (3.68)$$

of course, if the first measurement occurs at the initial time, then $t_1 = t_0$.

On this dilated level, the dynamics is Markovian—there are no additional external wires that can carry memory forward—and all higher order joint probability distributions could be built up when the individual CPTP maps $\{\Gamma_{t_j, t_{j-1}}\}$ (and the initial system-environment state) are known. With this, we can define NDGD dynamics (we provide a graphical representation in Fig. 3.8):

Definition 3.3 (NDGD dynamics). A global system-environment dynamics with CPTP maps $\{\Gamma_{t_j, t_{j-1}}\}_{j=1}$ is called Non-Discord-Generating-and-Detecting (**NDGD**) if it satisfies

$$\Delta_{j+1} \circ \Gamma_{t_{j+1}, t_j} \circ \Delta_j \circ \Gamma_{t_j, t_{j-1}} \circ \Delta_{j-1} = \Delta_{j+1} \circ \Gamma_{t_{j+1}, t_j} \circ \mathcal{I}_j \circ \Gamma_{t_j, t_{j-1}} \circ \Delta_{j-1} \quad (3.69)$$

for all $\{t_{j-1}, t_j, t_{j+1}\}$, where the maps $\Gamma_{t_k, t_{k-1}}$ act on the system and the environment, while Δ_k act on the system alone.

Formally, Definition 3.3 is equivalent to the definition of NCGD dynamics, with the difference that the involved intermediary maps between times are now the system-environment maps, instead of the maps $\{\Lambda_{t_j, t_{j-1}}\}$ acting on the system alone in the Markovian case.

Analogously to the case of NCGD, a NDGD dynamics cannot create discord (with respect to the classical basis) that can be detected at the next time (and, as such, at any later time) by means of classical measurements. Or, equivalently, an experimenter who can only perform measurements in the classical basis cannot distinguish between a completely dephasing map and an identity map implemented at any time in \mathcal{T} . As such, it provides the natural extension of NCGD to the non-Markovian case. We then have the following theorem:

Theorem 3.4 (NDGD dynamics and classicality). *Consider a general, possibly non-Markovian, process on \mathcal{T} , with $|\mathcal{T}| = K$, obtained from a system-environment dynamics as in Eqs. (3.67) and (3.68); then the process is K -classical if the initial system-environment state $\eta_{t_0}^{se}$ and the set $\{\Gamma_{t_j, t_{j-1}}\}$ of maps that corresponds to it are zero discord and NDGD, respectively.*

The proof of this theorem is provided in Appendix C.7. It relies on the fact that measurements in the classical basis commute with the completely dephasing map and proceeds along the same lines as the analogous proof for NCGD dynamics in the Markovian setting provided in Ref. [170]. Importantly, though, it is not a necessity for classical statistics that the corresponding maps are NDGD, as we will discuss below.

In order to further elucidate the relation of discord and classicality for general quantum stochastic processes, it is insightful to discuss the proximity of Theorem 3.4 to the corresponding results in Ref. [170] for the Markovian case. Theorem 3.4 establishes the importance of the role of quantum discord for the classicality of non-Markovian processes. In the memoryless case, it is coherence—or the impossibility of detection thereof—that makes a process classical. Here, this role is played by discord, with the only difference being that instead of describing the process in terms of maps that are solely defined on the system of interest, we are forced to dilate the process to the system-environment space, where it is rendered Markovian. Consequently, the classicality of a process cannot be decided based on the master equation or dynamical maps that describe the evolution of the system alone (as has already been pointed out in Ref. [170]). However, given, e.g., a Hamiltonian that generates the corresponding system-environment dynamics, whether or not the resulting process can be simulated classically can be decided by checking the validity of Eq. (3.69).

It would be desirable if NDGD dynamics were a sufficient and necessary criterion for the classicality of non-Markovian processes; however, this is not the case. We provide an example of dynamics that is not NDGD, but nevertheless leads to classical dynamics, in Appendix C.8. NDGD dynamics as defined in Eq. (3.69) is a statement about the entire system-environment dynamics, and holds for any possible initial state on the environment. However, by means of projective measurements on the system alone, one only has access to the system part, and the system-environment dynamics cannot be fully probed. Consequently, the criterion of Eq. (3.69) will, in general, be too strong for a given experimental scenario. Crucially, though, Theorem 3.4 allows us to understand the role of the discord generated by the system-environment interaction and subsequently detected via projective measurements on the system in establishing non-classical statistics.

Nonetheless, even though it is not necessary for the underlying dynamics to be NDGD in order for a non-Markovian process to display classical statistics, for any K -classical process, there always exists a dilation that is NDGD. That is, there exists a set $\{\tilde{\Gamma}_{t_j, t_{j-1}}\}$

of system-environment CPTP maps that are NDGD and a zero-discord initial system-environment state $\tilde{\eta}_{t_0}^{se}$ that yield the correct classical family of joint probability distributions when probed in the classical basis. Specifically, we have the following theorem:

Theorem 3.5. *Let $\{\mathbb{P}_n(x_n, \dots, x_1)\}_{n \leq K}$ define a process on \mathcal{T} , with $|\mathcal{T}| = K$, coming from an underlying evolution, fixed by the system-environment maps $\{\Gamma_{t_j, t_{j-1}}\}$ and the state $\eta_{t_0}^{se}$, according to Eqs. (3.67) and (3.68). The resulting statistics $\{\mathbb{P}_n(x_n, \dots, x_1)\}_{n \leq K}$ is K -classical iff there exists a NDGD evolution given by system-environment maps $\{\tilde{\Gamma}_{t_j, t_{j-1}}\}$ defined on times in \mathcal{T} and a zero-discord state $\tilde{\eta}_{t_0}^{se}$ that yield $\mathbb{P}_n(x_n, \dots, x_1)$ when probed in the classical basis.*

Before we prove this statement, it is important to contrast it with Theorem 3.1, the analogous result for Markovian processes. There, NCGD propagators of the system dynamics guarantee that the process associated with sequential projective measurements is classical, and classical Markovian processes can be reproduced by a set of NCGD maps (which do not necessarily identify with the actual dynamical propagators). Analogously, here, the NDGD property of the actual system-environment evolution ensures the classicality of the process; while the converse holds for particular dilations, there can be non-NDGD dilations that nonetheless yield classical statistics.

In both cases the projective measurements in a fixed basis only provide a limited amount of information about the overall evolution underlying the probed statistics. While in the Markovian case the statistics can be traced back to dynamical maps acting on the open system alone, in the more general non-Markovian case it is the whole system-environment evolution that enters into play. As a consequence, only the former case allows one to establish a one-to-one correspondence between classicality and the properties of the actual evolution by enforcing a proper condition on the dynamics, as discussed at the end of Section 3.4.1.

Proof. As we have already seen in the discussion of Theorem 3.4, the joint probability distributions obtained from an NDGD dynamics are always classical. We thus only need to prove the opposite direction. Let the underlying system-environment dynamics of the process between times be given by the maps $\{\Gamma_{t_j, t_{j-1}}\}$. As the process is classical, the set of maps $\{\tilde{\Gamma}_{t_j, t_{j-1}} = \Delta_j \circ \Gamma_{t_j, t_{j-1}} \circ \Delta_{j-1}\}$ together with a state $\tilde{\eta}_{t_1}^{se} = \Delta_1[\eta_{t_1}^{se}]$, where, again, Δ_k only acts on the system degrees of freedom, yields the same joint probability distributions when probed in the classical basis (see Fig. 3.9 for reference). The process given by this set $\{\tilde{\Gamma}_{t_j, t_{j-1}}\}$ is NDGD by construction and $\tilde{\eta}_{t_1}^{se}$ has vanishing discord, which means that for every K -classical process there is an NDGD dilation that reproduces it correctly, where we identify the initial time as the time of the first measurement, $t_0 = t_1$. \square




Figure 3.9: Transformation to NDGD Dilation. Any dilation of an open dynamics can be mapped onto an NDGD one by inserting completely dephasing maps on the level of the system. If the process is classical, then the transformed dilation yields the same statistics as the original one when probed in the classical basis.

Theorems 3.4 and 3.5 complete our results for the non-Markovian setting and provide an intuitive connection between non-classical spatial correlations (i.e., discord) and classical processes.

3.7 Genuinely Quantum Processes

As we have alluded to throughout this article, the classicality of a process depends on the measurement scheme that is employed to probe it; a process that appears classical in one basis—and is thus NDGD with respect to said basis—might display non-classical correlations when probed differently. This raises the question if non-classicality is merely a matter of perspective. In principle, for any process, there could exist a probing scheme that yields classical statistics. More concretely, for an experimenter that can perform arbitrary measurements, it might always be possible to “hide” the quantum nature of a process by choosing their respective measurements at the times $\{t_j\}$ such that the resulting statistics are classical.

Naturally, such schemes with (potentially non-projective) measurements go beyond the discussion of classicality that we have conducted so far. As we will not limit the employed instruments of such schemes to be the same at every time, we will call them *unrestricted* in what follows. However, we still assume that the instrument at each time is fixed in advance and is independent of previous measurements—if the choice of instruments could depend on previous outcomes, then the employed probing scheme would be temporally correlated and marginalisation at a given time would not be well-defined.

In this case, our previous results allow us to show that there exist *genuinely quantum* processes, i.e., processes that display non-classical statistics with respect to *every* unrestricted measurement scheme (in the sense described above) which reveals something about the probed process.

To reiterate, up to this point, our discussion of Markovianity focused on situations, where an experimenter measures in the computational basis only, thus employing the same instrument $\mathcal{J} = \{\mathcal{P}_{x_j}\}$ at each time, where all of the (projective) CP maps \mathcal{P}_{x_j} comprising the instrument added up to the completely dephasing map Δ_j . More generally, an experimenter could use instruments $\mathcal{J}_1 = \{\mathcal{M}_{x_1}\}, \mathcal{J}_2 = \{\mathcal{M}_{x_2}\}, \dots$, each adding up to

the CPTP maps $\mathcal{M}_1, \mathcal{M}_2, \dots$, respectively, to sequentially probe the system of interest. With this, for a process defined on times \mathcal{T} , they could collect the joint probability for all subsets $\mathcal{T}' \subseteq \mathcal{T}$ and check if Kolmogorov consistency holds. For example, in the simplest case of two times, with $\mathcal{T} = \{t_1, t_2\}$ and a given comb \mathcal{C}_2 on \mathcal{T} , an experimenter would consider the process classical, if $\mathbb{P}(x_2|\mathcal{J}_2) = \sum_{x_1} \mathbb{P}(x_2, x_1|\mathcal{J}_2, \mathcal{J}_1)$ holds for all x_2 , i.e., if

$$\mathcal{C}_2[\mathcal{M}_{x_2}, \mathcal{I}_1] = \mathcal{C}_2[\mathcal{M}_{x_2}, \mathcal{M}_1] \quad \forall \mathcal{M}_{x_2} \in \mathcal{J}_2. \quad (3.70)$$

Note that, due to causality, the second condition, i.e., $\mathcal{C}_2[\mathcal{I}_2, \mathcal{M}_{x_1}] = \mathcal{C}_2[\mathcal{M}_2, \mathcal{M}_{x_1}] \forall \mathcal{M}_{x_1} \in \mathcal{J}_1$ holds automatically, independent of whether the process is classical or not.

In principle, there could always exist a set of instruments $\{\mathcal{J}_K, \dots, \mathcal{J}_1\}$ for a given process \mathcal{C}_K on \mathcal{T} , such that the resulting statistics appear classical. Naturally, for this question to make sense, the respective instruments actually have to extract information from the process at hand. In principle, an instrument could consist of a random number generator and a set of CPTP maps that the experimenter implements depending on the respective output of the random number generator. Considering these outputs as outcomes of the instruments, the experimenter could then collect statistics that are independent of the process at hand (they only depend on the statistics of the random number generators), and satisfy Kolmogorov consistency conditions (if the respective random number generators at different times are independent of each other). However, this apparent classicality would not be a statement about the properties of the underlying process, and we thus exclude such pathological instruments. We can do so by demanding that at any time t_j , none of the elements \mathcal{M}_{x_j} of the instrument \mathcal{J}_j is proportional to a CPTP map. Under this reasonable assumption, we now show that there are processes that are genuinely quantum, i.e., they violate Kolmogorov conditions for arbitrary choices of instruments.

To this end, in the first step, we argue that genuinely quantum processes only exist in the non-Markovian setting, while in the memoryless case there always exists a measurement scheme that yields classical statistics. This conclusion follows from the fact that all features of a Markovian process are governed by the dynamical maps acting on the space of the system alone. Suppose then that a Markovian process is deemed to be non-classical with respect to some basis of measurements: This means that the dynamical maps constituting the process generate and detect coherence with respect to said basis. However, at each point in time throughout the process, the system to be measured is diagonal in some basis (namely, its eigenbasis); thus, in principle, if the experimenter were able to choose an unrestricted measurement scheme that is always diagonal in the same basis as the system, no coherence with respect to this basis will ever be generated and detected, implying that the statistics measured will appear classical. Consequently, in our proposed framework, genuinely quantum processes can only exist in the presence of (quantum) memory.




Figure 3.10: First Two Times of a Genuinely Quantum Process. The system-environment begin in a Bell state φ_{se}^+ . Between times t_1 and t_2 , the map Γ_z is implemented, which biases the system in the z -basis if any CPTP map $\mathcal{M}_1 \neq \mathcal{I}_1$ is performed [see Eq. (3.71)]. The label η_2^{se} refers to the joint system-environment state immediately prior to t_2 [see Eqs. (3.72) and (3.73)]. Classicality implies that the POVM Π_2 must be chosen such that it is unable to detect biases in the z -basis. Although this is always possible when only two times are considered, in general, classicality requires satisfaction of a growing number of constraints on the choices of later measurements, which can eventually lead to contradiction, implying the existence of genuinely quantum processes.

A similar argument as in the Markovian case holds for the special case of non-Markovian dynamics where the system-environment state at each time is of zero discord in a basis independent sense, i.e., when there exists a basis with respect to which the joint state at each time is discord-zero. Recall that if the system-environment dynamics is NDGD (with respect to a fixed basis), then the statistics observed are classical. Now, if at each time the system-environment state has zero discord, then an experimenter can (in principle) choose the measurement basis at each time to be the one with respect to which the performed measurement is non-invasive. For such a sequence of measurements, the experimenter would not be able to distinguish between having implemented the identity map or the dephasing map (with respect to the chosen basis) at any time, since the measurement is non-invasive on the joint system-environment state (due to the lack of discord). Thus, in such a scenario, there always exists some choice of bases in which such a process looks classical. It follows then that no non-Markovian process with zero basis independent discord between system and environment at every time is genuinely quantum.

However, the above logic fails in the general setting, which we now show by explicit example. To provide intuition, we first outline the logical implication that is a consequence of the classicality demand for a chosen (two-step) process (depicted in Fig. 3.10 and described below). While for two times it is always possible to find a measurement scheme such that the statistics appear classical (even in the non-Markovian case), when a non-Markovian process extends over multiple times, finding such a measurement scheme is not possible in general. We show this in detail in Appendix C.9 by considering a multi-time variant of the process shown in Fig. 3.10 that is extended over four times, proving the existence of genuinely quantum processes.

The explicit example process we consider begins with a two-qubit system-environment state in the Bell state $\varphi_{se}^+ = \frac{1}{2} \sum_{ij} |ii\rangle\langle jj|$. The experimenter can choose to measure the system (in whichever basis, or, more generally, employing any non-pathological instrument they like) at time t_1 . Following this, the dynamics consists of a system-environment CPTP

map $\Gamma_z : \mathcal{B}(\mathcal{H}^{s^i} \otimes \mathcal{H}^{e^i}) \rightarrow \mathcal{B}(\mathcal{H}^{s^o} \otimes \mathcal{H}^{e^o})$ whose action is to measure its joint inputs in the Bell basis, and output φ^+ if the measurement outcome indeed corresponds to φ^+ , or else output a system-environment state whose system part is a pure state in the z -basis. The action of Γ_z on a system-environment state η^{se} is thus given by

$$\Gamma_z[\eta^{se}] = \text{tr}(\eta^{se} \varphi_{se}^+) \varphi_{se}^+ + \text{tr}[(\mathbb{1}_{se} - \varphi_{se}^+) \eta^{se}] |0\rangle\langle 0|_s \otimes \tau_e, \quad (3.71)$$

where τ_e is some quantum state on the environment. It is straightforward to check that such a map is indeed CPTP. Following this part of the dynamics, the experimenter has access to measure the system at time t_2 .

For a genuinely quantum process, we demand that the statistics are non-classical with respect to any possible measurement choices at times t_1 and t_2 ; if this is not the case, then there exists a POVM at t_2 that cannot distinguish between the experimenter having implemented the identity map \mathcal{I}_1 or an arbitrary CPTP map \mathcal{M}_1 at time t_1 , such that the statistics look classical with respect to said measurement scheme. By tracking the joint system-environment state for either choice of operation at t_1 , we first show that such a POVM always exists. This implies that there is no genuinely quantum process defined on just two times, even in the non-Markovian setting. However, the POVM that does the trick is constrained by the demand of classicality, as we now detail. Extending the considered process to more times then imposes a number of constraints on the employed measurement devices which must be concurrently satisfied, such that finally there is no unrestricted measurement scheme that can yield classical statistics.

Suppose that the experimenter implements \mathcal{I}_1 at time t_1 ; then, the system-environment state at t_2 is given by

$$\eta_2^{se}(\mathcal{I}_1) := \Gamma_z[(\mathcal{I}_1^s \otimes \mathcal{I}^e)(\varphi_{se}^+)] = \varphi_{se}^+, \quad (3.72)$$

where the notation $\eta_2^{se}(\mathcal{I}_1)$ refers to the joint state immediately prior to t_2 given that the experimenter implemented the identity map at t_1 . On the other hand, if the experimenter overall implements some CPTP map $\mathcal{M}_1 \neq \mathcal{I}_1$ (corresponding to their instrument \mathcal{J}_1 at t_1), then the initial Bell pair will be perturbed (as it is only locally invariant under the identity map) and therefore the system-environment state prior to t_2 is

$$\eta_2^{se}(\mathcal{M}_1) := \Gamma_z[(\mathcal{M}_1^s \otimes \mathcal{I}^e)(\varphi_{se}^+)] = p\varphi_{se}^+ + (1-p)|0\rangle\langle 0|_s \otimes \tau_e, \quad (3.73)$$

where $p := \text{tr}[\varphi_{se}^+(\mathcal{M}_1 \otimes \mathcal{I}^e)[\varphi_{se}^+]] < 1$. The statistics observed are gathered by making measurements on only the system, so we are now interested in the reduced system state at t_2 in either case: From Eq. (3.72), we have the maximally mixed state $\eta_2^s(\mathcal{I}_1) = \frac{\mathbb{1}}{2}$, whereas from Eq. (3.73) we yield a state that is biased in the z -basis, $\eta_2^s(\mathcal{M}_1) = \frac{p}{2}\mathbb{1} + (1-p)|0\rangle\langle 0|$. As previously mentioned, classicality dictates that the POVM implemented at t_2 must not be able to distinguish between these two states, which leads to the fact that the chosen

measurement must be blind to any bias in the z -basis. Mathematically, we demand

$$\mathbb{P}_2(x_2|\mathcal{I}_1) \stackrel{!}{=} \mathbb{P}_2(x_2|\mathcal{M}_1), \quad (3.74)$$

which can only be satisfied if the experimenter chooses a POVM $\Pi_2 = \{\Pi_2^{(x_2)}\}$ such that

$$\mathrm{tr}[\Pi_2^{(x_2)} \eta_2^s(\mathcal{I}_1)] = \mathrm{tr}[\Pi_2^{(x_2)} \eta_2^s(\mathcal{M}_1)] \quad \forall x_2. \quad (3.75)$$

A POVM that satisfies the above equation can be readily constructed: The elements $\{\Pi_2^{(a)}, \mathbb{1} - \Pi_2^{(a)}\}$ can always be described by $\Pi_2^{(a)} = r_2^{(0)} \mathbb{1} + \vec{r}_2 \cdot \vec{\sigma}$, where $\vec{r}_2 = (r_2^{(x)}, r_2^{(y)}, r_2^{(z)})$ and $\vec{\sigma} = (\sigma^{(x)}, \sigma^{(y)}, \sigma^{(z)})$ is the vector of Pauli matrices (note that we have changed notation and use the letter “ a ” to label the measurement outcome in order to avoid potential confusion with the x -basis direction). Demanding classicality, i.e., Eq. (3.74), then implies that $r_2^{(z)} = 0$. In other words, any POVM that is not able to detect biases in the z -basis satisfies Eq. (3.75) and thus the statistics measured by such a POVM will appear classical. Importantly, here, and in what follows, we can restrict our analysis to the case of POVMs/instruments with only two elements, as any other POVM/instrument (except for the trivial case of single element ones) can always be coarse-grained to a two-element one. If such a coarse-grained instrument can detect non-classicality of statistics, then so too will the original one be able to, since it necessarily reveals more information about the process upon implementation.

However, although it might always be possible to find a basis/POVM such that the two-time statistics for a non-Markovian process look classical, this is not the case in general. Intuitively, demanding that the experimenter cannot distinguish between implementing the identity map and an arbitrary CPTP map at different times leads to a number of constraints (e.g., above we have the constraint $r_2^{(z)} = 0$) on the later measurement basis. In Appendix C.9, we consider a process defined across four times that is a logical extension of the two-time process considered here: In each of the first three times, depending on whether or not the system has previously been biased in either the x -, y - or z -basis, the process either performs an identity map (in the affirmative case) or else acts to bias the system in one of the bases. In the end, for an arbitrary CPTP map being implemented at each one of the first three times (with identity map being enacted at the others), the system state at the fourth time is biased in one of the three basis directions, and it is completely unbiased (i.e., maximally mixed) only if three consecutive identity maps are implemented. The only possible POVM at the final time that yields classical statistics must not be able to detect biases in any of the basis vector directions; the only POVM that achieves this is the one with elements proportional to the identity matrix, which corresponds to one of the measurements that we excluded because they reveal nothing about the process. Thus, the process is non-classical with respect to every possible non-pathological measurement scheme and is therefore genuinely quantum.

A relevant side-note seems in order here. Suppose that someone claims that a given process is genuinely quantum. To falsify such a statement it is enough to probe the process by whatever (non-trivial) devices one chooses; if the statistics one gets is classical, the statement is wrong. The processes that are not genuinely quantum can therefore be device-independently verified [304–307].⁵ In turn, this makes the genuinely quantum processes quite peculiar, as it is impossible to hide their quantumness, and it might come as a surprise that the set of these processes is non-empty; in fact we even conjecture that almost all many-time processes are genuinely quantum.

3.8 Conclusions and Outlook

3.8.1 Conclusions

In this paper, we have provided an operationally motivated definition of general classical stochastic processes and discussed its structural consequences and relation to quantum coherence in a system’s evolution as well as to the generation and activation of non-classical correlations between the system and the surrounding environment. While we phrased our results predominantly in the language of quantum mechanics, there is—a priori—nothing particularly quantum mechanical about the notion of non-classicality we introduced. Rather, any process for which the potential invasiveness of performed measurements can be detected by means of said measurements is non-classical, independent of the underlying theory; as such an invasiveness is experimentally detectable, this is a fully operational notion. The question of whether or not a process is classical can thus be answered on experimentally accessible grounds and is a priori independent of concepts that the experimenter might not be able to check for, like, e.g., coherences in the system of interest.

Nonetheless, our definition allows for the derivation of a direct connection between the classicality of a process and coherences/non-classical correlations that might be present. While this connection can be formulated in terms of a necessary and sufficient condition for memoryless processes, there are additional subtleties to be considered in the non-Markovian case. In general, it is not sufficient for the state of the system to be diagonal in the classical basis at all times for the resulting multi-time statistics to be classical. Rather, it is the interplay of coherences, non-classical system-environment correlations, and the underlying dynamics that is of importance, as we have highlighted through a number of examples presented throughout. Using the comb framework—which can encapsulate this complex interplay—for the description of general quantum processes with memory,

⁵Note, that the term device independent does not imply that one needs no assumptions on the devices, but just that the assumptions are trivial from the setting. In the case of device independent quantum key distribution, for instance, the assumption that the devices have some kind of independence needs to be assumed. In our case, we need to assume that the measurements are not trivial, because otherwise there is nothing one can say from their statistics at all.

we have provided a characterisation of quantum processes that yield classical statistics, and derived the structural properties of such processes. In principle, analogous structural properties could be derived for processes that display classical statistics when probed by means of different measurements, e.g., non-projective and/or non-orthogonal ones. However, while still enabling the derivation of structural properties, the clear connection between classicality and quantum discord would be lost as soon as sharp measurements in the computational basis are not the probing mechanism of choice anymore. In this paper, orthogonal projections were chosen as the kind of measurements that come closest to the ideal non-invasiveness displayed by classical measurements. More generally, our results could in principle also be extended to post-quantum theories. As the definition of classicality we provided is fully operational, the structure of classical processes in such theories could be derived in the same vein as we presented in this paper, with coherence and discord being replaced by the analogous properties of the respective theory.

Unsurprisingly, the set of classical processes turns out to be of measure zero within the set of all quantum processes. The full characterisation we have provided equips the set of classical processes with an experimentally accessible measure of non-classicality that can be formulated as an linear program, thereby providing an operationally clear-cut quantification of the degree of non-classicality of a given quantum process and a general theoretical framework to define practically useful measures of non-classicality. As an example, we showed how within our approach one can recover and motivate a quantifier of non-classicality which is exploited in different contexts [301] and has been used to analyse the properly quantum features of a given experimental setup [269].

Furthermore, we investigated the relation between the non-classicality of the statistics observed throughout a process and the quantumness of the prevalent spatial system-environment correlations in the underlying dynamics. While the absence of coherence in the state of the system of interest is no longer sufficient in the non-Markovian case to guarantee classicality, the absence of (basis dependent) discord is. This latter fact is somewhat intuitive, as the absence of discord at all times means that there are neither non-classical system-environment correlations nor coherences in the system that could influence the multi-time statistics deduced. Specifically, we have shown that the non-Markovian case to some extent mirrors the memoryless one: If the underlying dynamics is NDGD, i.e., any discord that is created at some point in time cannot be detected at a later time, then the process appears classical. While the converse of this statement does not hold, we have further shown that any classical process admits an NDGD dilation.

Finally, we demonstrated that, even if we extend our notion of classicality to the case of unrestricted measurement schemes, there exist processes that display non-classical statistics independent of how they are probed. This can happen only for non-Markovian processes, thus showing that genuine non-classicality can be seen as a further degree of complexity introduced by the presence of memory effects in the multi-time statistics of

quantum systems.

As our definition of classicality is tantamount to the assumptions of realism and non-invasiveness that underlie the derivation of Leggett-Garg inequalities, our results furnish experiments that test for the aforementioned properties with a clear interpretation: If the observed statistics satisfy a Leggett-Garg inequality, then the underlying process can be assumed to be NDGD. It does not have to be composed of fully classical resources, though. On the other hand, violation of a Leggett-Garg inequality implies that quantum discord must have been created (and later detected) over the course of the experiment.

3.8.2 Outlook

While we have provided a comprehensive picture of the interplay between the non-classical resources that are present in the underlying process and the non-classicality of the resulting non-Markovian multi-time statistics, the mechanisms that lead to the emergence of classical behaviour on macroscopic scales remain unclear. Naïvely, the fact that classical processes only constitute a vanishing fraction of the set of all processes, renders it puzzling that classical processes can be observed at all. This apparent “puzzle” is reminiscent of the superposition principle which restricts the set of states that are diagonal in a fixed basis to be of measure zero in the set of all pure states, yet superpositions are generally not observed in the macroscopic domain, where one fixed basis seems to be singled out.⁶ While for the latter case, decoherence has been identified as the mechanism that fixes a preferred basis—and as such leads to the emergence of classicality in the spatial setting [311, 312]—an analogous investigation for temporal processes remains outstanding. Our results pave the way towards the analysis of the onset of classicality in general quantum processes when system and/or environment size increases.

Beyond this foundational perspective, the characterisation of the set of classical processes, as well as the measure of non-classicality we have provided, lend themselves naturally to the development of a resource theory of non-classicality in which processes defined by Eqs. (3.52) and (3.53) are free. Additionally, our approach yields a definite theoretical background which allows one to deal with different quantifiers of the degree of non-classicality, related to practical situations where different sets of operations are available to investigate the quantumness of physical processes.

On the structural side, we have fully characterised the set of classical processes and have shown that there exist processes that are genuinely quantum. However, the explicit partitioning of the set of quantum processes into classical, non-classical, and genuinely quantum processes remains opaque and requires further investigation. It suggests itself to assume that the set of genuinely quantum processes is of full measure: As the set of discordant states is of full measure in the set of all states [296], for a randomly chosen

⁶For experiments that detect somewhat macroscopic superpositions, see, e.g., Refs. [308–310].

process, at any time t_j there will generally not exist a measurement that leaves the respective system-environment state invariant, and the subsequent dynamics would have to be highly fine-tuned in order to disguise this invasiveness. More specifically, based on the arguments employed in the explicit construction example of a genuinely quantum process we provided, where four measurement times were necessary to prove the genuine quantumness, we conjecture that almost all processes associated with a d -dimensional system are genuinely quantum, if the system is probed d^2 or more times. A rigorous proof of this statement is subject of future research. Moreover, since genuinely non-classical processes lead to non-classical statistics in a device-independent manner, their quantumness cannot be disguised. It then seems natural to explore if these processes can be used for technological applications.

Finally, the full characterisation of general, non-Markovian quantum processes possessing an equivalent classical description, will likely be useful to better understand the different facets of memory effects in the classical and quantum realm. Although the operational framework of quantum combs does not a priori concern any inherent timescales, as the choice of the discrete set of times is arbitrary, from a physical perspective one expects a connection between some relevant timescales of an underlying system-environment Hamiltonian generating a dynamics and the properties of the corresponding comb that arises upon specification of a set of times. Analogously, the timescales—and number of measurements—over which the non-classicality of a process can be deduced experimentally will be related to the pertinent timescales of the dynamics. However, determining the properties of an underlying system-environment Hamiltonian that leads to classicality and how the different timescales relate is an interesting, yet multi-layered and far from trivial, open problem.

The complexity arises due to the various temporal effects that play a significant role in determining the classicality (or absence thereof) of a given process and the relevant timescales over which it can be detected. For instance, we have already seen that the presence of multi-time memory effects is one such property; however, the connection between memory and classicality is a subtle one. One of the key differences between classical and quantum memory effects arises from the generically invasive nature of measurements in quantum mechanics, which leads to an inherent dependence of memory effects on the probing instruments employed [106]. The very notion of relevant memory timescales associated with the evolution of a quantum system therefore crucially depends on whether one wants to infer such timescales via sequential measurements over the course of the evolution or only at some final (possibly varying) time, as is done, e.g., in master equation approaches. In the latter case, the memory of the final statistics on the previous states of the system is dictated by the interplay of different timescales, related with the system of interest, its environment and their mutual interaction [272]. Such a memory ultimately determines the complexity of the description of the system evolution,

as provided, e.g., by memory kernels [171, 172], Green functions [313] or path integrals [314].

In the case where the temporal correlations of the environment rapidly decay, the process can often be approximated as a Markovian one. When the process is indeed Markovian, i.e., described by a sequence of individual channels between times, as we have shown, it is the NCGD property of the evolution that is necessary and sufficient for classicality; however, this is not easy to relate to the relevant timescales. A property that would be sufficient for classicality, and more straightforwardly related to the inherent timescales of a Hamiltonian generating the evolution is forgetfulness of any initial system state.

For instance, suppose one has a Markovian process generated by some Hamiltonian, which has a natural timescale of system forgetfulness, e.g., one that leads to an exponential decay of correlations between any preparations and final measurements. Then, if one probes such a process at sufficiently spaced time instants, one should expect to see classicality: The Markovianity property means that all relevant information can be determined solely on the system level, and forgetfulness ensures that any temporal correlations—in particular the ability to detect a distinction between a complete dephasing and an identity map—between adjacent times vanish. Strictly speaking, in the standard setting of testing for classicality, where a choice of measurement basis is fixed, one only requires forgetfulness with respect to projective measurements in said basis, rather than complete forgetfulness, for this argument to hold; however, besides being too strict a requirement, connecting such an instrument-specific forgetfulness to the relevant timescales is—like in the NCGD case—a difficult task.

In the presence of memory, the connection between classicality and the relevant timescales of the evolution is more involved yet. Here we have a subtle interplay between the question concerning the forgetfulness of the system of any initial non-classicality, as well as how much any non-classical effects can be transmitted through the environment via the memory mechanism. The fact that forgetfulness of the system alone here is insufficient to imply classicality is related to the crucial point that all multi-time effects must be captured in order to properly describe processes with memory. Thus, in the non-Markovian setting, the relevant timescales must typically be determined via sequential measurements over the course of the evolution.

However, different interrogation procedures will lead to the exhibition of different multi-time memory effects. For instance, when the system is left unperturbed, the memory can be solely attributed to properties of the underlying Hamiltonian (e.g., those leading to the decay of environmental correlations), whereas when the system is measured, the effect of conditioning the environment state also plays a role. Similarly to the Markovian setting discussed above, the question of classicality of a non-Markovian process does not necessarily concern all such temporal correlations in the process (both those transmitted

on the level of the system itself and the genuine memory effects due to the environment), but rather only those that can distinguish between the completely-dephasing instrument and the identity map applied to the system. These memory effects are, in turn, a special case of instrument-specific quantum Markov order, which has been recently introduced using the quantum comb formalism [130, 131]. Connecting such memory effects of the process, and their subsequent impact on the classicality of observed statistics, with the timescales associated to the corresponding Hamiltonian that generates a given process poses a promising avenue for future research.

While we anticipate the above open questions to generate much theoretical interest, we also expect our results to find immediate application in a broad range of situations where it is relevant to assess whether experimental outcomes are not amenable to a classical description in order to certify some type of quantum advantage or benchmark some genuinely quantum behaviour. The former include metrological schemes operating beyond the standard quantum limit [191, 315–317] while the latter can refer to the simulation of many body quantum systems [318–322]. Also, the role that the emergence of classicality plays in system thermalisation and homogenisation can be investigated in a systematic and quantitatively tractable manner within our proposed approach.



Connecting Commutativity and Classicality for Multi-Time Quantum Processes

Fattah Sakuldee, [Philip Taranto](#), and Simon Milz

Abstract. Understanding the demarcation line between classical and quantum is an important issue in modern physics. The development of such an understanding requires a clear picture of the various concurrent notions of “classicality” in quantum theory presently in use. Here, we focus on the relationship between *Kolmogorov consistency* of measurement statistics—the foundational footing of classical stochastic processes in standard probability theory—and the *commutativity* (or absence thereof) of measurement operators—a concept at the core of quantum theory. Kolmogorov consistency implies that the statistics of sequential measurements on a (possibly quantum) system could be explained entirely by means of a classical stochastic process, thereby providing an operational notion of classicality. On the other hand, commutativity of measurement operators is a structural property that holds in classical physics and its breakdown is the origin of the uncertainty principle, a fundamentally quantum phenomenon. Here, we formalise the connection between these two a priori independent notions of classicality, demonstrate that they are distinct in general and detail their implications for memoryless multi-time quantum processes.

Under review. Manuscript submitted on 30 Apr 2022.

[arXiv:2204.11698](#)

Author Contribution

In this work, the doctoral candidate contributed to the conception and formulation of the theoretical framework and methods, the proofs of the main results, and the writing and revising of the manuscript. In particular, the main technical contributions of the doctoral candidate were the extension of the considerations of Lüders to the multi-time setting; the derivation of the relevant operators whose commutativity should be checked regarding classicality; the comparative analysis between the two-time and multi-time settings; the development of many of the examples presented; and the connection to Non-Coherence-Generating-and-Detecting (NCGD) dynamics.

4.1 Introduction

Since the inception of quantum theory, various notions of “classicality” for the states of physical systems and measurements thereof have been put forth, including those based on the commutativity of observables [168, 323–327] (the breakdown of which being the origin of Heisenberg’s uncertainty principle [174–176]), the absence of coherence [242, 243, 249–251, 253] or quantum discord [265, 267, 268], the non-negativity of the Wigner function [328–332], the broadcastability of states [333, 334], or the objectivism that emerges through Darwinist arguments [335–339]. Most of these concepts of classicality are *static* in the sense that they focus on properties of quantum states and/or the compatibility of measurements in situations where there is no dynamics taking place between them. When extending such considerations to *processes*, i.e., physical systems that display non-trivial evolution and are measured at several points in time, classicality is often linked to the inability of a process to generate and/or detect states displaying such aforementioned properties, as well as certain properties of the resulting multi-time statistics (see, e.g., Refs. [170, 269, 340, 341]). Additionally for the case of sequential measurements, non-commutativity is the key ingredient in the generalisation of stochastic processes to the quantum realm [104]. However, besides partial results, the links between such—a priori inequivalent— notions of classicality (or non-classicality) remain poorly understood and subject to debate [257–261], both in static and dynamic scenarios.

Broadly speaking, existing notions of classicality fall into one of two categories: *Structural* ones, i.e., criteria for classicality based on mathematical properties like the commutativity (of observables) or the coherence of quantum states; and *operational* ones, i.e., those based only upon experimentally accessible entities, such as the multi-time statistics obtained from probing an evolving quantum state at different points in time. While both types of considerations are well-motivated in their own right, the connection between such—generally inequivalent—structural and operational notions of classicality has not yet been fully established; in the static case, it is only known for special cases [168, 177–179] and in the dynamic scenarios that we will focus on such links are only known when restricting to projective measurements of a fixed observable [170, 269, 340, 341]. Here, we establish more general connections between structural and operational notions of classicality for the case of a quantum system that undergoes non-trivial dynamics and is probed at multiple points in time with general instruments. Specifically, we analyse the connection between the satisfaction of so-called *Kolmogorov consistency conditions*—an operational notion of classicality—and commutativity of the operators that “naturally” assume the role of observables in multi-time processes (we motivate and identify these operators in Section 4.3).

Satisfaction of the former criterion implies the existence of a classical (i.e., described by standard probability theory) stochastic process that leads to the same statistics as

the one observed when the underlying quantum process is measured [145, 146]; in other words, although such a process might actually be quantum in nature, one cannot conclude this from the collected statistical data alone. Importantly, checking the Kolmogorov consistency conditions amounts to a clear operational notion of classicality that can be tested without any additional knowledge of the underlying dynamics or physical theory. For the case of quantum theory restricted to sequential measurements in a fixed basis, this criterion has been connected to the ability of the dynamics to generate and detect coherences, both in the Markovian (memoryless) [170, 269] as well as non-Markovian setting [340, 341]. Extensions to more general measurements have remained elusive.

On the other hand, (non-)commutativity of observables—the structural property that we consider in this article—lies at the core of quantum theory. Intuitively, commutativity of two observables A and B implies that they are jointly measurable; that is, given an arbitrary quantum state ρ , the probability of obtaining an outcome pertaining to observable B is independent of whether A was measured before it or not (and vice versa). This connection between measurement *non-invasiveness*—an operational notion of classicality—and commutativity of observables—a structural notion of classicality—was first considered by Lüders [168] for the case of projective measurements and later extended to more general scenarios [177–179] (see also Section 4.2), where it was shown that commutativity and measurement non-invasiveness coincide in many cases. Importantly, such a direct connection between these two a priori distinct concepts can only be meaningfully established under the assumption that there are only *two* sequential measurements being considered.

In the multi-time setting with non-trivial dynamics between measurements, it is a priori unclear how Lüders’ results carry over and what “observables” are the appropriate ones to consider when checking commutativity. In particular, both the underlying dynamics in between measurements and the effects of general measurement instruments must be accounted for in the temporal setting. While this can be done by combining the chosen measurements with the dynamics, it is then not necessarily the commutativity of the bare measurements (i.e., pertaining to the measurement device itself) per se, but rather the *effective* measurements (i.e., those with the dynamics accounted for) that render the observed statistics “classical” or “non-classical” accordingly.

Here, we identify the operators that determine the non-invasiveness of measurements for the case of multiple sequential measurements with non-trivial intermediate dynamics and analyse the conditions for which the commutativity—or weaker versions thereof—of these operators corresponds to the satisfaction of the Kolmogorov consistency conditions (and vice versa). Our analysis thus connects structural with operational notions of classicality for multi-time dynamics and general measurement settings. For the special case of two sequential measurements without intermediate dynamics, our results coincide with those of Lüders. However, in general, the situation presents itself considerably more

layered, and a “straightforward” extension of Lüders’ results is not possible. We show that commutativity (of the relevant operators) is a stronger condition than Kolmogorov consistency in general; while the former implies the latter, the converse does not hold. Additionally, we derive the conditions under which Kolmogorov consistency implies the vanishing of pertinent commutators in a restricted—but still multi-time—setting and highlight the ensuing physical implications in order to develop intuition concerning the interplay of these two notions of classicality. Finally, we relate our considerations to the well-known case of projective measurements in a fixed basis—where structural properties that are equivalent to Kolmogorov consistency have been identified [170]—and show that, while said structural considerations follow directly from those we provide for more general measurement scenarios, even in this special case, it is difficult to identify generally applicable commutator relations.

Together, our results offer a comprehensive analysis regarding the connection of structural—yet mostly not directly observable—properties of quantum dynamics and operational—i.e., experimentally accessible— notions of classicality, and underline the complicated interplay between dynamics and measurements that arises in the multi-time scenario.

This article is organised as follows. We begin by outlining some preliminary concepts, including the considerations of Lüders [168] that motivate the examination of commutativity, and similarly the Kolmogorov consistency conditions, in Section 4.2. We then explore the link between these two concepts within the setting of multi-time Markovian quantum dynamics throughout Section 4.3, where we first derive a commutator expression whose vanishing is sufficient to imply classical statistics, before deriving a necessary condition for classicality to imply vanishing commutators of the pertinent operators. We subsequently connect our work with the special case of dynamics that do not generate and detect coherences, which constitutes perhaps the most physically relevant special case [170, 269, 340, 341] that our results apply to. Finally, we conclude with a discussion and outlook in Section 4.4.

4.2 Preliminaries

We begin by introducing the relevant concepts for both a structural and an operational definition of classicality in multi-time processes. To this end, first, we recall the connection between commutativity of observables and non-invasiveness in the two-time case, in particular the considerations of Lüders [168].

4.2.1 Lüders’ Theorem: Commutativity and Non-Invasiveness

As a preliminary example concerning the connection between structural and operational notions of classicality, we consider the simplest case: The sequential measurement of two

observables A and B on a state ρ without intermediate evolution. Let $\{\Pi^{(a)}\}$ and $\{\Omega^{(b)}\}$ be projectors onto the eigenspaces of A and B , respectively, with eigenvalues $\{a\}$ and $\{b\}$. The probability to first measure outcome a and then b is given by

$$\mathbb{P}(b, a) = \text{tr}(\Omega^{(b)}\Pi^{(a)}\rho\Pi^{(a)}) = \text{tr}(\rho\Pi^{(a)}\Omega^{(b)}\Pi^{(a)}). \quad (4.1)$$

In classical¹ physics, future statistics are unaffected by whether or not a previous measurement was conducted (when that previous measurement outcome is not recorded, i.e., averaged over). Consequently, if the above situation were classical, then

$$\sum_a \mathbb{P}(b, a) = \mathbb{P}(b, \emptyset) \quad (4.2)$$

would hold, where $\mathbb{P}(b, \emptyset)$ denotes the probability to obtain outcome b if the first measurement was *not* performed. In quantum mechanics, Eq. (4.2) (or generalisations thereof, see Section 4.2.2) fails to hold in general, since quantum measurements are invasive. As a consequence, not performing a measurement is distinguishable from measuring and averaging over outcomes.

One example where Eq. (4.2) can be satisfied in the quantum setting, independently of ρ , is when $[\Pi^{(a)}, \Omega^{(b)}] = 0$ for all a, b (which is equivalent to $[A, B] = 0$), since then $\sum_a \Pi^{(a)}\Omega^{(b)}\Pi^{(a)} = \Omega^{(b)}$ (where we have used $\Pi^{(a)}\Pi^{(a)} = \Pi^{(a)}$ and $\sum_a \Pi^{(a)} = \mathbb{1}$), and thus

$$\sum_a \mathbb{P}(b, a) = \sum_a \text{tr}(\Pi^{(a)}\Omega^{(b)}\Pi^{(a)}\rho) = \text{tr}(\Omega^{(b)}\rho) = \mathbb{P}(b, \emptyset). \quad (4.3)$$

If the above is satisfied, then, just like in classical physics, all information is contained in the two-point probability distribution $\mathbb{P}(b, a)$ in the sense that both single-time distributions $\mathbb{P}(a)$ and $\mathbb{P}(b)$ can be obtained from it by marginalisation and the process is thereby fully characterised. Consequently, throughout this article, we call *classical* those experimental situations that satisfy this property, i.e., that yield probabilities which can all be obtained from one single multi-time probability distribution by means of marginalisation (see Section 4.2.2 for a rigorous discussion). In this case, we also say that said probability distributions for different subsets of times are *compatible* or *consistent*. Importantly, this notion of classicality amounts to measurement non-invasiveness: Whether or not a measurement has been performed at a given time has no bearing on the outcome probabilities at different times if the statistics are classical. In this sense, commutativity of observables A and B in a two-point measurement scenario implies classicality of the observed statistics (the converse is also true, but not obvious, see below), establishing a direct connection between a structural notion (commutativity of operators) of classicality and an operational definition (measurement non-invasiveness) thereof.

More generally, instead of performing sharp measurements of an observable, an experimenter could first probe the state ρ with a general instrument, described by a set of

¹Here, we employ “classical” in a somewhat colloquial sense, as pertaining to the macroscopic world. Below, we properly define what we mean exactly by “classical” throughout this article.

Kraus operators $\{K^{(a)}\}$, where each $K^{(a)}$ corresponds to a measurement outcome a , and subsequently a POVM $\{Q^{(b)}\}$, each element of which corresponds to an outcome b . The two-point statistics are then given by

$$\mathbb{P}(b, a) = \text{tr}(Q^{(b)} K^{(a)} \rho K^{(a)\dagger}) = \text{tr}(\rho K^{(a)\dagger} Q^{(b)} K^{(a)}) =: \text{tr}(\rho \mathcal{K}^{(a)\dagger}[Q^{(b)}]). \quad (4.4)$$

Here, $K^{(a)}$ and $Q^{(b)}$ respectively play analogous roles to $\Pi^{(a)}$ and $\Omega^{(b)}$ in the previous example. Now, setting $\mathcal{K}^\dagger[\bullet] := \sum_a \mathcal{K}^{(a)\dagger}[\bullet]$, we see that non-invasiveness of the first measurement amounts to the satisfaction of

$$\text{tr}(\rho \mathcal{K}^\dagger[Q^{(b)}]) = \text{tr}(\rho Q^{(b)}) \quad \forall b. \quad (4.5)$$

If the above must hold for *all* states ρ , then non-invasiveness is equivalent to

$$\mathcal{K}^\dagger[Q^{(b)}] = Q^{(b)} \quad \forall b. \quad (4.6)$$

This criterion has been connected to commutation relations by Lüders, yielding the following theorem:

Theorem 4.1 (Lüders [168, 177]). *Let \mathcal{K}^\dagger be defined as above and let Q be a positive semi-definite operator. If all Kraus operators $K^{(a)}$ are Hermitian, then $\mathcal{K}^\dagger[Q] = Q$ is equivalent to $[K^{(a)}, Q] = 0 \quad \forall a$.*

Since some of our proofs below follow a similar line of reasoning, we recall from the literature, see e.g. Ref. [177], for the proof of this theorem.

Proof. First, it is easy to see that $[K^{(a)}, Q] = 0 \quad \forall a$ implies $\mathcal{K}^\dagger[Q] = Q$. To see the converse, let $|\varphi\rangle$ be an arbitrary vector in the Hilbert space \mathcal{H} that Q is defined on. Suppose that $\mathcal{K}^\dagger[Q] = Q$, and decompose $Q = \sum_\mu \lambda_\mu P_\mu$ with $\lambda_1 > \lambda_2 > \dots$ with $\{P_\mu\}$ being mutually orthogonal projection operators. It follows that

$$\begin{aligned} \lambda_1 \|P_1|\varphi\rangle\|^2 &= \langle P_1|\varphi|Q P_1|\varphi\rangle \\ &= \langle P_1|\varphi|\mathcal{K}^\dagger[Q] P_1|\varphi\rangle \\ &= \sum_a \langle K^{(a)} P_1|\varphi|Q|K^{(a)} P_1|\varphi\rangle \\ &\leq \lambda_1 \sum_a \langle K^{(a)} P_1|\varphi|K^{(a)} P_1|\varphi\rangle \\ &= \lambda_1 \sum_a \langle K^{(a)\dagger} K^{(a)} P_1|\varphi|P_1|\varphi\rangle \\ &= \lambda_1 \|P_1|\varphi\rangle\|^2, \end{aligned} \quad (4.7)$$

where we have used $\lambda_1 \mathbb{1} - Q \geq 0$ for the first inequality and $\sum_a K^{(a)\dagger} K^{(a)} = \mathbb{1}$ for the last equality. From the above, we see that

$$\langle K^{(a)} P_1|\varphi|(\lambda_1 \mathbb{1} - Q)|K^{(a)} P_1|\varphi\rangle = \|(\lambda_1 \mathbb{1} - Q)^{1/2} K^{(a)} P_1|\varphi\rangle\|^2 = 0 \quad (4.8)$$

holds, and we thus have

$$QK^{(a)}P_1|\varphi\rangle = \lambda_1 K^{(a)}P_1|\varphi\rangle, \quad (4.9)$$

implying that $K^{(a)}$ leaves the λ_1 -eigensubspace invariant. This means that $K^{(a)}P_1 = P_1K^{(a)}P_1$, and since $K^{(a)}$ is assumed to be Hermitian, it follows that $[K^{(a)}, P_1] = 0$ for all a . Now, we set $Q_\mu = Q_{\mu-1} - \lambda_{\mu-1}P_{\mu-1}$ and $Q_0 = Q$, and repeat the same steps as above with Q_2 and P_2 and so on. This iteration then leads to the fact that $[K^{(a)}, P_\mu] = 0$ for all a and μ . Hence $[K^{(a)}, Q] = 0$ as claimed. \square

Since $\mathcal{K}^\dagger[Q^{(b)}] = Q^{(b)}$ is equivalent to non-invasiveness of the first measurement, Theorem 4.1 says that non-invasiveness for arbitrary initial states ρ is equivalent to commutativity of the Kraus operators of the first measurement and the POVM elements of the second one, if *all* Kraus operators are Hermitian (which is, e.g., the case for projective measurements of two observables).

We emphasise that Hermiticity of the Kraus operators is crucial for the derivation of Theorem 4.1, and without this assumption, it no longer holds in general. This can be seen by considering a counterexample provided in Ref. [179]:

Example 4.1. Let the POVM elements $\{Q^{(1)}, Q^{(2)}\}$ be given by

$$Q^{(1)} = \frac{1}{2} \begin{pmatrix} 2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad Q^{(2)} = \frac{1}{2} \begin{pmatrix} 0 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (4.10)$$

and the Kraus operators by

$$\begin{aligned} K^{(1)} &= \frac{1}{2} \begin{pmatrix} \sqrt{2} & 0 & -1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad K^{(2)} = \frac{1}{10} \begin{pmatrix} 0 & 0 & 0 \\ 0 & -\sqrt{10} & 2\sqrt{10} \\ 0 & 0 & 0 \end{pmatrix}, \\ K^{(3)} &= \frac{1}{2} \begin{pmatrix} 0 & 0 & 0 \\ 0 & \sqrt{2} & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad K^{(4)} = \frac{1}{20} \begin{pmatrix} 0 & 0 & 0 \\ 0 & 4\sqrt{10} & 2\sqrt{10} \\ 0 & 0 & 0 \end{pmatrix}, \\ K^{(5)} &= \frac{1}{2} \begin{pmatrix} \sqrt{2} & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \end{aligned} \quad (4.11)$$

For this setup, it is straightforward to see that both $Q^{(1)}$ and $Q^{(2)}$ are invariant under $\sum_{a=1}^5 K^{(a)\dagger} \bullet K^{(a)}$, i.e., the measurement is non-invasive overall and therefore the

resulting statistics are classical, but, for example, $[K^{(5)}, Q^{(1)}] \neq 0$. Consequently, classicality of statistics and commutativity of the Kraus operators are generally inequivalent notions. Nonetheless, the direct connection between commutativity and measurement non-invasiveness in quantum mechanics, i.e., that provided by Lüders' theorem, has subsequently been extended to more general (e.g., non-Hermitian) Kraus operators [177–179].

■

Importantly for our purposes, Theorem 4.1 identifies the relevant operators whose commutation relations are related to non-invasiveness. As a first step, in what follows we will investigate which operators play the roles of $K^{(a)}$ and $Q^{(b)}$ in the multi-time—i.e., more than two consecutive measurements—case.

Before doing so, we first note that Theorem 4.1 (and its extensions) are restricted in their realm of application. Firstly, they are limited to only two sequential measurements with no intermediate evolution.² Additionally, the (potential) equivalence between $\mathcal{K}^\dagger[Q^{(b)}] = Q^{(b)}$ and classicality requires non-invasiveness for *all* states ρ . As we will discuss, in the multi-time scenario, one is not always guaranteed to have access to a full basis of quantum states at each time of interest. Consequently, in the multi-time case, the relation between compatible statistics and the commutativity of some appropriate operators presents itself as a more layered issue than in the static or two-time cases, even when the Kraus operators of performed measurements are limited in a similar way to the assumptions of Theorem 4.1.


We now move to consider the operational notion of non-invasiveness in the multi-time case, namely the *Kolmogorov consistency conditions*, which are naturally suited to analysing the classicality (or not) of general quantum processes. We will see that Lüders' considerations amount to a special case of Kolmogorov consistency, before moving on to develop multi-time “Lüders-type” theorems, in the sense that they connect non-invasiveness of measurements to the vanishing of a set of pertinent commutator expressions.

4.2.2 Kolmogorov Consistency and Non-Invasiveness


In Eq. (4.2), we provided an experimentally accessible notion of non-invasiveness—and thus classicality—for two sequential measurements. The natural way to extend this definition to the multi-time case is as follows. Suppose that an agent probes a physical system at n discrete points in time, recording the corresponding joint probability distribution $\mathbb{P}(m_n, \dots, m_1)$ over possible outcomes $\{m_n, \dots, m_1\}$ observed at the respective times $\{t_n, \dots, t_1\}$ [see Fig. 4.1(a)]. Importantly, analogous to the two-time case discussed above, for any classical stochastic process, the recorded probability distribution is guaranteed to satisfy the *Kolmogorov consistency conditions* [145], illustrated in Figs. 4.1(b) and 4.1(c):

$$\mathbb{P}(m_n, \dots, \cancel{m_i}, \dots, m_1) = \sum_{m_i} \mathbb{P}(m_n, \dots, m_i, \dots, m_1) \quad \forall i. \quad (4.12)$$


²This is, technically, not a strong restriction, since any intermediate dynamics could be absorbed into the Kraus operators/POVM elements.



(a) Overall joint probability distribution $\mathbb{P}(m_4, m_3, m_2, m_1)$ for measurements at each time t_1, \dots, t_4 .



(b) Marginalised probability distribution $\sum_{x_2} \mathbb{P}(x_4, x_3, x_2, x_1)$ for times t_1, t_3 , and t_4 .



(c) Overall joint probability distribution $\mathbb{P}(m_4, m_3, \cancel{m_2}, m_1)$ for measurements at each of the times t_1, t_3, t_4 .

Figure 4.1: Multi-Time Experiments & Kolmogorov Consistency: (a) An unknown process (depicted by the “trajectories”) is probed at different times, here $\{t_1, t_2, t_3, t_4\}$, and the resulting joint probability distribution $\mathbb{P}(m_4, m_3, m_2, m_1)$ is recorded. (b) From the four-time distribution obtained in (a), one can compute a three-time joint distribution by marginalising over the outcomes at a given time (here, t_2). (c) The three-time joint probability $\mathbb{P}(m_4, m_3, \cancel{m_2}, m_1)$ is obtained by *not* performing a measurement at t_2 . In general, this is a different experiment than the one of (b). If the process is classical—and the employed measurements are non-invasive—then the joint probabilities of (b) and (c) coincide [see Eq. (4.12)].

Evidently, Eq. (4.12) is a multi-time generalisation of the two-time scenario considered in Eq. (4.2). The distribution on the l.h.s. corresponds to what an experimenter observes if they do *not* perform any interrogation at time t_i , whereas on the r.h.s. the full statistics is recorded and then marginalised over at time t_i . As in the two-time case, Kolmogorov consistency states that there is no difference between not having performed a measurement and measuring but averaging over all possible outcomes at any time, corresponding to a sensible notion of classicality in terms of measurement non-invasiveness: For classical stochastic processes, measurements simply reveal a pre-existing property of the system, in line with the assumptions of macroscopic realism used, e.g., in the derivation of Leggett-Garg inequalities [147]. This property fails to hold for quantum processes, since quantum measurements generically alter the state of the system being measured [274].

We emphasise that a breakdown of Kolmogorov consistency does not necessarily imply that the probed process at hand is non-classical per se. For instance, in the theory of classical causal modelling [342], where invasive interrogations can be implemented (e.g., by first measuring the value of some property and then setting it to some other value) in order to potentially infer causal influence, the recorded statistics generally do not satisfy the Kolmogorov consistency conditions [145, 146]. Nonetheless, testing the validity of Eq. (4.12) provides a theory-independent, operational procedure to decide on the non-

invasiveness of interrogations. In particular, satisfaction of the Kolmogorov consistency conditions implies that there exists a—potentially exotic—classical stochastic process that can reproduce the observed statistics. To do so, said classical stochastic process merely needs to correctly recreate the full joint probability distribution $\mathbb{P}(m_n, \dots, m_1)$, and, due to satisfaction of the Kolmogorov conditions, it then also correctly recreates all joint probability distributions for any subset of times, thereby fully characterising the process from an operational standpoint. Thus, we will interchangeably use the terms “Kolmogorov consistency”, “measurement non-invasiveness” and “classicality”.

Recently, in Refs. [170, 269, 340, 341], the implications of the satisfaction of the Kolmogorov consistency conditions for general multi-time processes (including those with memory) that are probed by means of pure projective measurements have been characterised, thus connecting the operational, experimentally accessible notion of classicality with certain properties of the underlying quantum dynamics, namely their ability to generate and detect coherence or discord with respect to the chosen measurement basis (we discuss the relationship of these results with our present work in Section 4.3.4). Here, we allow for *general measurements* and phrase our results in terms of commutation relations (i.e., in the spirit of Lüders’ theorem), rather than in terms of the coherence- or discord-related properties of the underlying quantum maps that engender the observed statistics.

4.2.3 Multi-Time Statistics from (Markovian) Quantum Processes

To make the relation between commutation relations and non-invasiveness in quantum theory more concrete and identify the relevant operators, we now examine how observed statistics are related to the underlying dynamics of a quantum process. In order to collect joint statistics at times t_1, \dots, t_n , at each time t_i an experimenter probes the system of interest with an *instrument*, $\mathcal{J}_i = \{\mathcal{K}_i^{(m_i)}\}$, which is a collection of Completely Positive (CP) maps that sum up to a Completely Positive and Trace Preserving (CPTP) map, i.e., $\mathcal{K}_i := \sum_{m_i} \mathcal{K}_i^{(m_i)}$ is a CPTP map [103]. Each CP map $\mathcal{K}_i^{(m_i)}$ corresponds to a possible outcome m_i and captures the state change of the system upon measurement. For simplicity, we assume that every element $\mathcal{K}_i^{(m_i)}$ can be represented by a single Kraus operator, i.e., $\mathcal{K}_i^{(m_i)}[\rho] = K_i^{(m_i)}\rho K_i^{(m_i)\dagger}$. In between these measurements (e.g., between t_i and t_{i+1}), the system of interest undergoes non-trivial dynamics, possibly interacting with an environment, described by CPTP maps $\Lambda_{i+1:i}$. Assuming the dynamics to be *Markovian* (i.e., memoryless), as we do throughout, then these maps are mutually independent and act on the system alone (see Fig. 4.2). Any statistics observed by probing a Markovian process can be computed via the *quantum regression formula* [272, 278, 279]:

$$\mathbb{P}(m_n, \dots, m_1 | \mathcal{J}_n, \dots, \mathcal{J}_1) = \text{tr} \left(\mathcal{K}_n^{(m_n)} \circ \Lambda_{n:n-1} \circ \dots \circ \Lambda_{2:1} \circ \mathcal{K}_1^{(m_1)}[\rho] \right), \quad (4.13)$$

where all maps act on the system alone. It is important to note that the statistics observed depend on *both* the CP maps $\mathcal{K}_i^{(m_i)}$ implemented by the experimenter and the generally




Figure 4.2: Multi-Time Probing of a Markovian Quantum Process. A quantum process without memory on a discrete set of times t_1, \dots, t_n can be described by an initial state ρ of the system and a collection of independent CPTP maps $\{\Lambda_{i:i-1}\}$ that act on the system alone between neighbouring times (blue). At each time t_i , we envisage an agent probing the process and observing a measurement outcome m_i , with the post-measurement state feeding forward (yellow). Such a probing is represented by a CP map $\mathcal{K}_i^{(m_i)}$ at each time.

uncontrollable dynamics of the process given by the CPTP maps $\{\Lambda_{i:i-1}\}$. Whether or not the measured statistics satisfies the Kolmogorov consistency conditions thus depends on the complex interplay between measurements *and* intermediate dynamics. With Eq. (4.13) at hand, we can now identify the relevant operators and commutation relations concerning the satisfaction of Kolmogorov conditions.

4.3 Multi-Time Dynamics: Kolmogorov Consistency and Commutativity

To identify the relevant commutation relations, let us rewrite Eq. (4.13) entirely in terms of Kraus operators:

$$\begin{aligned} & \mathbb{P}(m_n, \dots, m_1 | \mathcal{J}_n, \dots, \mathcal{J}_1) \\ &= \sum_{\ell_2 \dots \ell_n} \text{tr} \left[K_n^{(m_n)} L_{n:n-1}^{(\ell_n)} \dots L_{2:1}^{(\ell_2)} K_1^{(m_1)} \rho K_1^{(m_1)\dagger} L_{2:1}^{(\ell_2)\dagger} \dots L_{n:n-1}^{(\ell_n)\dagger} K_n^{(m_n)\dagger} \right], \end{aligned} \quad (4.14)$$

where we have set $\Lambda_{i:i-1}[\bullet] = \sum_{\ell_i} L_{i:i-1}^{(\ell_i)} \bullet L_{i:i-1}^{(\ell_i)\dagger}$.

In order to connect this equation to non-invasiveness of a measurement at a time t_i , we split the above expression into three parts (using the cyclicity of the trace): One that corresponds to the state immediately prior to the measurement, one corresponding to the measurement itself, and one corresponding to everything that happens after the measurement at said time of interest. Specifically, setting $R_{i:i-1}^{\ell_i, m_{i-1}} := L_{i:i-1}^{(\ell_i)} K_{i-1}^{(m_{i-1})}$, we see that the pre-measurement (subnormalised) state at time t_i is given by

$$\tilde{\rho}_i(\mathbf{m}_{i-1:1}) := \sum_{\ell_2 \dots \ell_i} R_{i:i-1}^{\ell_i, m_{i-1}} \dots R_{2:1}^{\ell_2, m_1} \rho R_{2:1}^{\ell_2, m_1\dagger} \dots R_{i:i-1}^{\ell_i, m_{i-1}\dagger} \quad (4.15)$$

for $i \geq 2$ (with $\tilde{\rho}_1 := \rho$). Note that $\tilde{\rho}_i(\mathbf{m}_{i-1:1})$ depends upon all measurement outcomes $\mathbf{m}_{i-1:1} := (m_{i-1}, \dots, m_1)$ up to t_i , and its trace corresponds to the probability to observe said sequence of outcomes, i.e.,

$$\mathbb{P}(m_{i-1}, \dots, m_1 | \mathcal{J}_{m-1}, \dots, \mathcal{J}_1) = \text{tr}[\tilde{\rho}_i(\mathbf{m}_{i-1:1})]. \quad (4.16)$$

On the other hand, grouping all post-measurement operators together, we can define the positive semi-definite operator

$$Q_i(\mathbf{m}_{n:i+1}) := \sum_{\ell_{i+1} \dots \ell_n} L_{i+1:i}^{(\ell_{i+1})\dagger} R_{i+2:i+1}^{(\ell_{i+2}, m_{i+1})\dagger} \dots R_{n:n-1}^{(\ell_n, m_{n-1})\dagger} K_n^{(m_n)\dagger} K_n^{(m_n)} R_{n:n-1}^{(\ell_n, m_{n-1})} \dots R_{i+2:i+1}^{(\ell_{i+2}, m_{i+1})} L_{i+1:i}^{(\ell_{i+1})}. \quad (4.17)$$

With this, the multi-time statistics of Eq. (4.14) can be expressed succinctly as

$$\begin{aligned} \mathbb{P}(m_n, \dots, m_1 | \mathcal{J}_n, \dots, \mathcal{J}_1) &= \text{tr} \left[\tilde{\rho}_i(\mathbf{m}_{i-1:1}) K_i^{(m_i)\dagger} Q_i(\mathbf{m}_{n:i+1}) K_i^{(m_i)} \right] \\ &=: \text{tr} \left[\tilde{\rho}_i(\mathbf{m}_{i-1:1}) \mathcal{K}_i^{(m_i)\dagger} [Q_i(\mathbf{m}_{n:i+1})] \right]. \end{aligned} \quad (4.18)$$

Intuitively, $\tilde{\rho}_i$ is the time-evolved (subnormalised) state that is to be measured at time t_i , while Q_i corresponds to the effect of each measurement outcome after t_i (all the way up until some fixed final time t_n), with the dynamics of the process in between accounted for, “rolled back” in time to t_i , when the measurement described by $\mathcal{K}_i^{(m_i)\dagger}$ occurs. With this, setting $\mathcal{K}_i^\dagger[\bullet] := \sum_{m_i} \mathcal{K}_i^{(m_i)\dagger}[\bullet]$, Kolmogorov consistency is equivalent to

$$\text{tr} \left[\tilde{\rho}_i(\mathbf{m}_{i-1:1}) \mathcal{K}_i^\dagger [Q_i(\mathbf{m}_{n:i+1})] \right] \text{tr} \left[\tilde{\rho}_i(\mathbf{m}_{i-1:1}) Q_i(\mathbf{m}_{n:i+1}) \right], \quad (4.19)$$

which can be expressed as a commutator expression via

$$\sum_{m_i} \text{tr} \left[\tilde{\rho}_i(\mathbf{m}_{i-1:1}) K_i^{(m_i)\dagger} [K_i^{(m_i)}, Q_i(\mathbf{m}_{n:i+1})] \right] = 0 \quad (4.20)$$

for all t_i and all outcomes $\{m_1, \dots, m_{i-1}, m_{i+1}, \dots, m_n\}$. Formally, apart from the dependence on past and future outcomes, the above equation coincides with Eq. (4.5), which states the Kolmogorov consistency conditions for two sequential measurements without intermediate dynamics. This seemingly implies that there should be a direct relation between commutation relations of the involved operators in Eq. (4.20) and classicality of the observed statistics. However, we now discuss some important differences between classicality for two-time vis-a-vis multi-time processes and subsequently demonstrate that there is no “straightforward” extension of Lüders’ theorem to the multi-time setting, except under rather restrictive assumptions.

4.3.1 Two-Time vs. Multi-Time Classicality

There are various major differences between the two- and multi-time scenario. Firstly, in the two-time case, assuming that measurement non-invasiveness holds for arbitrary initial states ρ , one can conclude that satisfaction of the Kolmogorov consistency conditions is equivalent to $\mathcal{K}^\dagger[Q^{(b)}] = Q^{(b)}$. On the other hand, in the multi-time case, even when assuming that the Kolmogorov consistency conditions hold for arbitrary initial states (i.e., those at time t_1), one is no longer guaranteed that the system states span a full basis at

every later time t_i . To see this, consider the natural case of measurements performed in the computational basis, i.e., $\mathcal{K}_{i-1}^{(m_{i-1})}[\rho] = \langle m_{i-1} | \rho | m_{i-1} \rangle |m_{i-1}\rangle\langle m_{i-1}|$. Then, the state $\tilde{\rho}_i$ immediately prior to the measurement at time t_i —independent of that at the beginning of the experiment—is proportional to $\Lambda_{i:i-1}[|m_{i-1}\rangle\langle m_{i-1}|]$ and so the set of states $\{\tilde{\rho}_i\}$ can at most span a d dimensional space, which cannot coincide with the d^2 dimensional space spanned by all quantum states. Consequently, in the multi-time case, Kolmogorov consistency is—in contrast to the two-time scenario—manifestly *not* equivalent to

$$\mathcal{K}_i^\dagger[Q_i(\mathbf{m}_{n:i+1})] = Q_i(\mathbf{m}_{n:i+1}). \quad (4.21)$$

In particular, satisfaction of the above equation for all i is sufficient for satisfaction of Kolmogorov conditions—as can be seen by direct insertion into Eq. (4.19)—but not necessary (see below). Nonetheless, the formulation of Eq. (4.20) informs us that the commutation relations $[K_i^{(m_i)}, Q_i(\mathbf{m}_{n:i+1})]$, or variants thereof, are the relevant ones to investigate with respect to satisfaction of Kolmogorov consistency conditions.

Note that if *all* CP maps $\{\mathcal{K}_i^{(m_i)}\}$ and *all* intermediate dynamics $\{\Lambda_{i:i-1}\}$ are invertible (the former is a choice, while the latter is generally true for Markovian dynamics [343]), then one is guaranteed a full basis of states at each time t_i provided that the very initial state is arbitrary; this implies that the Kolmogorov consistency conditions are then indeed equivalent to satisfaction of Eq. (4.21), even in the multi-time scenario. In this case, one can, in the spirit of Theorem 4.1, establish a direct connection between measurement non-invasiveness for all input states and the vanishing of the commutator $[K_i^{(m_i)}, Q_i(\mathbf{m}_{n:i+1})]$ (provided some additional assumptions are met, e.g., Hermiticity of the Kraus operators $K_i^{(m_i)}$ —we will revisit these assumptions and the issue of requiring a full basis in Section 4.3.3). However, with respect to the choice of instruments in particular, this restriction is rather strong and would (as mentioned above) fail to cover the most natural scenario of measurements in a fixed basis.

Additionally, independent of the fact that at t_i one does not have access to a full basis of states, when considering general quantum measurements in multi-time processes, information can be transmitted through the system alone, and thus measurement statistics can be correlated over multiple points in time, even for Markovian processes. Thus, in the multi-time scenario, one must deal with entire sequences of outcomes—for example, $Q_i(\mathbf{m}_{n:i+1})$ is an operator that pertains to the entire sequence of future outcomes—instead of just outcomes at single or neighbouring times. As we discuss in detail in Section 4.3.4, this added complexity cannot be circumvented as soon as general measurements are considered, and consequently all of our results will be phrased with respect to operators that generally correspond to measurement outcomes at multiple different points in time.

We now detail how Eq. (4.21) has to be modified in order to yield a direct relation between commutativity of pertinent operators and the classicality of the observed statistics in a multi-time experiment.

4.3.2 No “Straightforward” Extension of Lüders’ Theorem

In Eq. (4.19), we have expressed satisfaction of Kolmogorov conditions at an arbitrary time t_i in terms of the measurement map \mathcal{K}_i^\dagger , the post-measurement operators $Q_i(\mathbf{m}_{n:i+1})$ and the pre-measurement subnormalised states $\tilde{\rho}(\mathbf{m}_{i-1:1})$. The close formal relation of said equation to those that appear in the two-time scenario, and thus in Lüders’ theorem, informs us that the commutators $[K_i^{(m_i)}, Q_i(\mathbf{m}_{n:i+1})]$ play a pivotal role for the classicality of the observed statistics. As mentioned, it is easy to see that $[K_i^{(m_i)}, Q_i(\mathbf{m}_{n:i+1})] = 0$ implies satisfaction of Kolmogorov conditions, but the converse is not true (see Example 4.3 for a concrete counterexample). Indeed, such a commutation relation is, at the outset, far too strict a condition to be necessary for Kolmogorov consistency: It would, for instance, imply that the measurements are non-invasive for *arbitrary* system states at each time, which is not a necessary requirement for classicality, since for many relevant scenarios, the possible states $\tilde{\rho}_i$ before a measurement at time t_i do not span the full space of quantum states. As a result, our aim is to find weaker commutation relations that still guarantee the satisfaction of Kolmogorov conditions, and, conversely, to work out the consequences of Kolmogorov conditions on the commutation relations of the relevant operators. Here, we begin with the former direction.

Recall that satisfaction of Kolmogorov consistency conditions is given by Eq. (4.19). As we have emphasised, this does not necessarily imply that $\mathcal{K}_i^\dagger[Q_i] = Q_i$ and is thus not equivalent to $[K_i^{(m_i)}, Q_i] = 0$. Following the logic of Theorem 4.1, one might then suspect that

$$\mathrm{tr} \left[\tilde{\rho}_i(\mathbf{m}_{i-1:1}) [K_i^{(m_i)}, Q_i(\mathbf{m}_{n:i+1})] \right] = 0, \quad (4.22)$$

i.e., commutativity of the measurement operators *with respect to* $\tilde{\rho}_i(\mathbf{m}_{i-1:1})$ is equivalent to the satisfaction of Kolmogorov conditions at time t_i —at least for the case of Hermitian Kraus operators $K_i^{(m_i)}$. However, this is not the case, as the following example shows:

Example 4.2. Let the pre-measurement state (for some history of outcomes, which we renormalise and suppress for the sake of conciseness) be given by $\tilde{\rho}_i = \frac{1}{2}(\mathbb{1} + \sigma_z)$ followed by a measurement described by Kraus operators $K_i^{(\pm)} = \frac{1}{2}(\mathbb{1} \pm \sigma_x)$ and let the post-measurement part (for some sequence of future outcomes) be encoded in the operator $Q_i(\pm) = \frac{1}{2}(\mathbb{1} \pm \sigma_z)$. This situation can, e.g., arise in a two-step process without intermediate evolution, where the measurement with Kraus operators $\{K_1^{(\pm)}\}$ is made at t_1 , the pre-measurement state is prepared as $\tilde{\rho}_1 = \frac{1}{2}(\mathbb{1} + \sigma_z)$ and at t_2 the observable σ_z is measured with outcomes \pm , corresponding to the post-measurement operators $Q_1(\pm) = \frac{1}{2}(\mathbb{1} \pm \sigma_z)$. We observe that (for future outcome $+$) we have $\mathrm{tr} \left[\tilde{\rho}_1 [K_1^{(\pm)}, Q_1(+)] \right] = 0$ but

$$\begin{aligned} & \mathrm{tr} \left[\tilde{\rho}_1 \left(\sum_{m_1=\pm} K_1^{(m_1)\dagger} Q_1(+) K_1^{(m_1)} - Q_1(+) \right) \right] \\ &= \mathrm{tr} \left[\tilde{\rho}_1 \left(\mathcal{K}_1^\dagger [Q_1(+)] - Q_1(+) \right) \right] = \frac{1}{2} \neq 0, \end{aligned} \quad (4.23)$$

implying that the Kolmogorov consistency conditions are not satisfied even though commutativity with respect to $\tilde{\rho}_1$ [i.e., Eq. (4.22)] holds and all involved Kraus operators are Hermitian. ■

We emphasise that even though we only explicitly consider two measurements here, the considered scenario is indeed a multi-time one; in contrast to the scenario envisioned by Lüders, we do not assume the states before the first measurement at t_1 to span a full basis, which is an implicit assumption of the two-time setting with arbitrary initial preparations. This, in turn, can be understood as the pre-measurement state $\tilde{\rho}_1$ being the result of a previous measurement with a fixed outcome (or sequence thereof), making the scenario of the example a genuine multi-time one of which we only explicitly investigated the two times t_1 and t_2 . The fact that Kolmogorov conditions are not satisfied despite the weak commutativity of Eq. (4.22) holding then signifies that in the multi-time scenario, one requires a stricter commutation relation for the involved operators in order to obtain classical statistics.

Although the weak commutation relation with respect to $\tilde{\rho}_i(\mathbf{m}_{i-1:1})$ is not restrictive enough, the following theorem informs us that *absolute* commutativity with respect to $\tilde{\rho}_i(\mathbf{m}_{i-1:1})$ is indeed sufficient to guarantee the satisfaction of Kolmogorov conditions:

Theorem 4.2. *Let $\mathbb{P}(m_n, \dots, m_1 | \mathcal{J}_n, \dots, \mathcal{J}_1)$ be a joint probability obtained from Eq. (4.13), i.e., by probing a Markovian process. If absolute commutativity*

$$\text{tr} \left[\tilde{\rho}_i(\mathbf{m}_{i-1:1}) \left| [K_i^{(m_i)}, Q_i(\mathbf{m}_{n:i+1})] \right| \right] = 0 \quad (4.24)$$

holds at all times t_i and for all possible $\mathbf{m}_{i-1:1}$, m_i and $\mathbf{m}_{n:i+1}$, where $|X| := \sqrt{X^\dagger X}$, then $\mathbb{P}(m_n, \dots, m_1 | \mathcal{J}_n, \dots, \mathcal{J}_1)$ satisfies the Kolmogorov consistency conditions [given explicitly in Eq. (4.20)].

Proof. For simplicity, we will omit the explicit arguments of the involved operators throughout the proof. We first show that Eq. (4.24) implies $\text{tr} \left[\tilde{\rho}_i K_i^{(m_i)\dagger} [K_i^{(m_i)}, Q_i] \right] = 0$. To this end, we note that Eq. (4.24) implies

$$\left| [K_i^{(m_i)}, Q_i] \right| \tilde{\rho}_i = 0, \quad (4.25)$$

since both $\tilde{\rho}_i$ and $\left| [K_i^{(m_i)}, Q_i] \right|$ are positive semidefinite. Now, let us employ the polar decomposition $[K_i^{(m_i)}, Q_i] = V^{(m_i)} M^{(m_i)}$, where $V^{(m_i)}$ is unitary and $M^{(m_i)} = \left| [K_i^{(m_i)}, Q_i] \right| \geq 0$, i.e., $M^{(m_i)} \tilde{\rho}_i = 0$. With this, we obtain

$$\text{tr} \left[\tilde{\rho}_i K_i^{(m_i)\dagger} [K_i^{(m_i)}, Q_i] \right] = \text{tr} \left[K_i^{(m_i)\dagger} V^{(m_i)} M^{(m_i)} \tilde{\rho}_i \right] = 0. \quad (4.26)$$

By summing this expression over m_i , Eq. (4.20)—and thus satisfaction of the Kolmogorov consistency conditions—is recovered. □

Theorem 4.2 informs us that absolute commutativity with respect to the state of the system at each time is sufficient for classicality of the observed statistics. However, in contrast to the two-time scenario, this requirement is *not* necessary for classicality—even in the case where the involved Kraus operators are Hermitian. To see this, consider the following example:

Example 4.3. We employ Example 4.2 with a change in the (renormalised) pre-measurement part to $\tilde{\rho}_1 = \frac{1}{2}(\mathbb{1} + \sigma_y)$ but still followed by a measurement $K_1^{(\pm)} = \frac{1}{2}(\mathbb{1} \pm \sigma_x)$ and the post-measurement parts are encoded in the operators $Q_1(\pm) = \frac{1}{2}(\mathbb{1} \pm \sigma_z)$. We observe that—for these choices—Kolmogorov consistency holds, i.e.,

$$\mathrm{tr} \left\{ \tilde{\rho}_1 \left[\sum_{m_1=\pm} K_1^{(m_1)\dagger} Q_1(\pm) K_1^{(m_1)} - Q_1(\pm) \right] \right\} = 0, \quad (4.27)$$

since $\sum_{m_1=\pm} K_1^{(m_1)\dagger} Q_1(\pm) K_1^{(m_1)} - Q_1(\pm) = \pm \frac{\sigma_z}{2}$, which is trace orthogonal to $\tilde{\rho}_1$. However we find that $|[K_1^{(\pm)}, Q_1(\pm)]| = \frac{\mathbb{1}}{4}$ leading to

$$\mathrm{tr} \left[\tilde{\rho}_1 |[K_1^{(\pm)}, Q_1(\pm)]| \right] = \frac{1}{4} \neq 0. \quad (4.28)$$

Consequently, this example shows that classical statistics in a multi-time experiment do not imply absolute commutativity with respect to the state of the interrogated system over time, even when all Kraus operators are Hermitian (which is the case here). In turn, since absolute commutation with respect to $\tilde{\rho}_i$ is weaker than commutativity itself, this makes the considered case also an example of a situation where satisfaction of Kolmogorov consistency conditions does not imply $[K_i^{(m_i)}, Q_i(\mathbf{m}_{n:i+1})] = 0$, as mentioned at the beginning of this section. ■

While not being equivalent to satisfaction of the Kolmogorov consistency conditions, absolute commutativity with respect to the state $\tilde{\rho}_i$ guarantees classical statistics and, in contrast to the much stronger standard commutativity condition, does not necessarily imply Kolmogorov consistency independent of the sequentially measured system states, making it a more relevant consideration for the envisaged scenario.

Regarding this connection between commutativity and classicality, two remarks are in order. On the one hand, if $\tilde{\rho}_i$ is *full rank*, then it is easy to see that

$$\mathrm{tr} \left[\tilde{\rho}_i(\mathbf{m}_{i-1:1}) |[K_i^{(m_i)}, Q_i(\mathbf{m}_{n:i+1})]| \right] = 0 \quad (4.29)$$

implies $[K_i^{(m_i)}, Q_i(\mathbf{m}_{n:i+1})] = 0$, thus equating the assumption of absolute commutativity with respect to $\tilde{\rho}_i$ to the (rather strong) assumption of standard commutativity. However the states $\tilde{\rho}_i$ do not necessarily have to be full rank. This holds true, for example, for the case of pure projective measurements in a fixed basis on a qutrit and intermediate dynamics that only map to the (lower dimensional) space that is spanned by $\{|0\rangle, |1\rangle\}$. In

turn, this makes the assumption of Theorem 4.2 strictly weaker than full commutativity, while still being strictly stronger than commutativity with respect to $\tilde{\rho}_i$ [i.e., satisfaction of Eq. (4.22)], which, as we have seen, is *not* sufficient to guarantee classical statistics.

On the other hand, the states $\tilde{\rho}_i$ at time t_i are the result of a state preparation at the initial time, followed by a sequence of measurements and intermediate dynamics. Assuming that a full basis of initial states can be prepared (as is assumed in the two-time scenario envisioned by Lüders), then it is—in principle—possible that, for each sequence of outcomes, the corresponding states $\tilde{\rho}_i$ also span a basis at each time t_i . In this case, satisfaction of Kolmogorov conditions at t_i would exactly coincide with $\mathcal{K}_i^\dagger[Q_i(\mathbf{m}_{n:i+1})] = Q_i(\mathbf{m}_{n:i+1})$ [see Eq. (4.19)] and, following the same reasoning that led to Lüders' Theorem 4.1, we would be able to recover the equivalence between classical statistics and the vanishing of the commutators $[K_i^{(m_i)}, Q_i(\mathbf{m}_{n:i+1})] = 0$. In this sense, it might seem artificial to investigate the case where states at each time do *not* span a full basis, which, as we will see, leads to a more layered relationship between commutativity and classicality. However, this latter case exactly mirrors many physically relevant scenarios (like, e.g., the case of sequential projective measurements).

This inequivalence between commutation relations and classicality naturally raises the question: *What further assumptions, in addition to classicality of statistics, must be satisfied in order to ensure commutation relations of the relevant operators?*

4.3.3 Commutativity as a Notion of Classicality: When is Kolmogorov Consistency Sufficient for Lüders-Type Theorems?

In this section we investigate under which conditions satisfaction of Kolmogorov consistency implies the vanishing of pertinent commutator expressions. Unlike the previous sections, here—just like in the scenario considered by Lüders—we have to restrict the Kraus operators of the probing instruments to be Hermitian in order to establish a clear connection between Kolmogorov consistency and vanishing commutators.

As mentioned previously, a key element that makes the multi-time setting substantially different to the two-time one is that one is no longer guaranteed a full basis of quantum states at each time. Nonetheless, below we outline a condition that ensures that the set of possible states at each time (conditioned on previous outcome sequences) essentially forms a basis *with respect to* any subsequent measurements (additionally accounting for the intermediate dynamics). Analogously to the case considered by Lüders, our argument requires Hermiticity of the measurement Kraus operators; we leave the analysis of sufficient conditions regarding more general measurements in this setting for future work. The conditions that we detail below consequently ensures a connection between Kolmogorov consistency and commutativity in the multi-time setting (for Hermitian Kraus operators). Importantly, just like in the case of Lüders, under the additional assumptions we make,

commutativity and classicality are equivalent.

To establish the connection between Kolmogorov consistency and commutation relations, consider the set of possible pre-measurement states at some time (say, t_i) of interest. As mentioned, we can follow a Lüders type argument, if these states form a basis with respect to the post-measurement operators Q_i . Formally, we can express this by letting \mathbb{S} be a set of initial states ρ and \mathbb{H}_i be the span of the union of the images of all possible pre-measurement sequences up until time t_i :

$$\mathbb{H}_i := \text{span} \bigcup_{\mathbf{m}_{i-1:1}} \sum_{\ell_2 \dots \ell_i} R_{i:i-1}^{(\ell_i, m_{i-1})} \dots R_{2:1}^{(\ell_2, m_1)} \mathbb{S} R_{2:1}^{(\ell_2, m_1)\dagger} \dots R_{i:i-1}^{(\ell_i, m_{i-1})\dagger}, \quad (4.30)$$

for $i \geq 2$, i.e., \mathbb{H}_i is the span of all attainable states $\tilde{\rho}_i(\mathbf{m}_{i-1:1})$ at time t_i . Furthermore, we take the union of all possible projections for the post-measurement operators to define:

$$\mathbb{F}_i := \text{span} \bigcup_{\mathbf{m}_{n:i-1}} \left\{ P_\mu : Q_i(\mathbf{m}_{n:i-1}) = \sum_\mu \lambda_\mu P_\mu \right\}, \quad (4.31)$$

where, for technical reasons, we will assume non-degeneracy of Q_i (see the proof of Theorem 4.3). Demanding that the pre-measurement states form a basis with respect to the post-measurement operators now amounts to the requirement $\mathbb{F}_i \subseteq \mathbb{H}_i$. As it turns out, together with the satisfaction of Kolmogorov consistency, this implies that $\text{tr}[P_\mu \mathcal{K}_i^\dagger[Q_i]] = \text{tr}[P_\mu Q_i] \forall \mu$, which suffices to prove that the pertinent commutation relations hold (under the assumption that all Kraus operators pertaining to the measurement map \mathcal{K}^\dagger are Hermitian):

Theorem 4.3. *Let $\mathbb{P}(m_n, \dots, m_1 | \mathcal{J}_n, \dots, \mathcal{J}_1)$ be a joint probability obtained from Eq. (4.13), i.e., by probing a Markovian process. Assume that $\mathbb{P}(m_n, \dots, m_1 | \mathcal{J}_n, \dots, \mathcal{J}_1)$ satisfies the Kolmogorov consistency conditions [given explicitly in Eq. (4.20)] for all initial state ρ in \mathbb{S} and for every measurement time t_i and that $Q_i(\mathbf{m}_{n:i+1})$ is non-degenerate for all $\mathbf{m}_{n:i+1}$. If all Kraus operators $K_i^{(m_i)}$ are Hermitian for all m_i and*

$$\mathbb{F}_i \subseteq \mathbb{H}_i, \quad (4.32)$$

then the commutation relations hold, i.e.,

$$[K_i^{(m_i)}, Q_i(\mathbf{m}_{n:i+1})] = 0 \quad (4.33)$$

for all post-measurement sequences $\mathbf{m}_{n:i+1}$ and all m_i .

Before providing the proof of Theorem 4.3, we emphasise that the converse trivially holds (even without any assumptions), since commutativity [i.e., Eq. (4.33)] directly implies the satisfaction of Kolmogorov consistency.

Proof. From the assumption Eq. (4.32), one can see that for a given post-measurement sequence $\mathbf{m}_{n:i+1}$ and for any pre-measurement sequence $\mathbf{m}_{i-1:1}$, there exists an initial

state ρ in \mathbb{S} leading to $\tilde{\rho}_i(\mathbf{m}_{i-1:1}) = P_\mu$ for any P_μ defined via $Q_i(\mathbf{m}_{n:i-1}) = \sum_\mu \lambda_\mu P_\mu$. In other words, Kolmogorov consistency in the form of Eq. (4.19) leads to $\text{tr}[P_\mu \mathcal{K}_i^\dagger[Q_i]] = \text{tr}[P_\mu Q_i] \forall \mu$.

Now, using the same arguments as those of the proof of Theorem 4.1, one sees that $\text{tr}[P_\mu \mathcal{K}_i^\dagger[Q_i]] = \text{tr}[P_\mu Q_i]$ for $\mu = 1$ leads to

$$\sum_m \langle K_i^{(m_i)} \varphi_1 | (\lambda_1 \mathbb{1} - Q_i) K_i^{(m_i)} \varphi_1 \rangle = \sum_m \| (\lambda_1 \mathbb{1} - Q_i)^{1/2} K_i^{(m_i)} | \varphi_1 \rangle \|^2 = 0, \quad (4.34)$$

where we set $P_1 =: |\varphi_1\rangle\langle\varphi_1|$ and made use of the fact that $\lambda_1 \mathbb{1} - Q_i \geq 0$. In other words, $(\lambda_1 \mathbb{1} - Q_i)^{1/2} K_i^{(m_i)} P_1 = 0$ or

$$Q_i K_i^{(m_i)} P_1 | \varphi \rangle = \lambda_1 K_i^{(m_i)} P_1 | \varphi \rangle \quad (4.35)$$

for arbitrary states $|\varphi\rangle$, i.e., $K_i^{(m_i)}$ leaves the λ_1 -eigensubspace invariant. Thus, due to non-degeneracy of Q_i , we have $K_i^{(m_i)} P_1 = P_1 K_i^{(m_i)} P_1$ and assuming that the $K_i^{(m_i)}$ are Hermitian, we observe that $[K_i^{(m_i)}, P_1] = 0$. Again, we set $Q_i^{(\mu)} = Q_i^{(\mu-1)} - \lambda_{\mu-1} P_{\mu-1}$ and $Q_i^{(0)} = Q_i$. Since $[K_i^{(m_i)}, P_1] = 0$, the expression $\text{tr}[P_2 \mathcal{K}_i^\dagger[Q_i]] = \text{tr}[P_2 Q_i]$ can be reduced to $\text{tr}[P_2 \mathcal{K}_i^\dagger[Q_i^{(2)}]] = \text{tr}[P_2 Q_i^{(2)}]$ and then it follows that $[K_i^{(m_i)}, P_2] = 0$ (by invoking the same previous argument but replacing Q_i and P_1 with $Q_i^{(2)}$ and P_2 , respectively). Iterating this argument—as in the proof of Theorem 4.1—we obtain that $K_i^{(m_i)} P_\mu = P_\mu K_i^{(m_i)} P_\mu$ for all m and μ , i.e., $K_i^{(m_i)}$ leaves all eigensubspaces of Q invariant. Then $[K_i^{(m_i)}, Q_i(\mathbf{m}_{n:i+1})] = 0$ for all post-measurement sequences $\mathbf{m}_{n:i+1}$ and all m_i as claimed. \square

We emphasise that—as in the case of Lüders' theorem—this logic can fail to hold if the $K_i^{(m_i)}$ are not Hermitian (as can already be explicitly seen by considering Example 4.1).

For illustration of the above theorem, let us consider the following example.

Example 4.4. Recall the scenario of Example 4.2. We modify it to be a three step process with measurements in the σ_z -basis at the first and third time, while at the second time, a measurement in the σ_x -basis is carried out. In addition, let the dynamics between the first and the second measurement, as well as between the second and the third measurement be given by a Hadamard gate H , with $H|0\rangle = |\pm\rangle$ and $|\pm\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$.

We focus on time t_2 as the measurement time of interest (i.e., the time for which we analyse Kolmogorov consistency). For an arbitrary initial state ρ , the measurement at t_1 leads to the set of possible post-measurement states $\tilde{\rho}_1 \in \{|0\rangle\langle 0|, |1\rangle\langle 1|\}$, where we omit potential subnormalisation. These states will then evolve to $H[\tilde{\rho}_1]H = \tilde{\rho}_2 \in \{|+\rangle\langle +|, |-\rangle\langle -|\}$. Thus, at time t_2 , we have $\mathbb{H}_2 = \text{span}\{|+\rangle\langle +|, |-\rangle\langle -|\}$.

Since a measurement in the σ_z -basis is performed at t_3 , the post-measurement part at t_2 amounts to the corresponding projectors, “rolled back” by means of the evolution $\Lambda_{3:2}^\dagger[\bullet] = H[\bullet]H$, i.e., we have $Q_2(0/1) \in \{|+\rangle\langle +|, |-\rangle\langle -|\}$. With this, we observe that

$$\mathbb{F}_2 = \mathbb{H}_2 \quad (4.36)$$

and thus the condition of Eq. (4.32) holds. Likewise, for the measurement $K_2^{(\pm)}$, the Kolmogorov consistency condition reads

$$\text{tr} \left[\tilde{\rho}_2 \left(\sum_m K_2^{(m)\dagger} [K_2^{(m)}, Q_2(0/1)] \right) \right] = 0. \quad (4.37)$$

where $m \in \{+, -\}$. We can calculate the commutativity expression of Eq. (4.33) explicitly:

$$[K_2^{(\pm)}, Q_2(0/1)] = [|\pm\rangle\langle\pm|, |\pm\rangle\langle\pm|] = 0. \quad (4.38)$$

■

The inclusion property of Eq. (4.32) is a rather strong requirement, which—as we have seen in Example 4.4—can be checked for and satisfied in particular cases. However, it can fail to hold for many experimentally relevant situations that yield classical statistics (like, e.g., measurements in a fixed basis, see Example 4.6 below). Additionally, one would ideally like to deduce similar conditions that apply to arbitrary (e.g., non-Hermitian) measurements at the expense of potentially weakening the vanishing commutator expression; so far, such results have proved elusive. As a consequence, in the multi-time setting, the relation between observed classicality and the commutation of pertinent operators presents itself much more layered than in the two-time case considered by Lüders, and must seemingly be decided on a case-by-case basis.

We now finish our discussion of Markovian classical multi-time processes by discussing why, even though the underlying process is memoryless, it is, in general, indeed necessary to consider the *full* history (future) of outcomes $\mathbf{m}_{i-1:1}$ ($\mathbf{m}_{n:i+1}$) at each time t_i , and not just the preceding (subsequent) ones m_{i-1} (m_{i+1}). The latter (i.e., only considering outcomes at t_{i-1} and t_{i+1} for each time t_i) can be done for the special case of projective pure measurements in a fixed basis, which leads to an equivalent formulation of classicality in terms of non-coherence-generating-and-detecting maps. However, as we will see in the following section, a direct connection between such dynamics to pertinent commutation relations is generally not direct.

4.3.4 Markovian Processes & Non-Coherence-Generating-and-Detecting (NCGD) Dynamics

Up to this point, we have investigated the conditions under which Kolmogorov consistency of observed statistics and the vanishing of pertinent commutator expressions—i.e., structural properties pertaining to the dynamics and measurement scheme—are related in the multi-time setting. Crucially, we see that Kolmogorov consistency concerns a deep interplay between the choice (and assumptions) of measurements and the underlying dynamics that depends on entire sequences of measurement outcomes.

On the other hand, for the Markovian case we consider, recent work has demonstrated a one-to-one connection between Kolmogorov consistency and the (in)ability for pairs of

neighbouring dynamical maps $\Lambda_{i:i-1}$ (describing the open evolution of Markovian process) to generate and detect coherence with respect to a fixed basis determined by the measurement scheme [170, 269, 340, 341]. Specifically, these works showed equivalence between classical statistics and the set of so-called Non-Coherence-Generating-and-Detecting (NCGD) dynamics, i.e., maps that can create coherences, but those coherences cannot be “detected” by the subsequent dynamics [see below Eq. (4.42) for a proper definition]. While this criterion can be phrased entirely in terms of dynamical maps pertaining to neighbouring times, our work has required the consideration of *entire measurement sequences* of past $\mathbf{m}_{i-1:1}$ and future $\mathbf{m}_{n:i+1}$ outcomes in general, instead of simply adjacent ones.

At the outset, this necessity seems to be overkill, since, intuitively, the statistics measured at each time of a *Markovian* process should only depend upon the most recent outcome, and not on the entire history. We now return to elucidate why, even though the underlying dynamics that we study are assumed to be Markovian, one must indeed consider commutator expressions of the relevant operators corresponding to entire sequences. The important subtlety to note here is that general quantum measurements *do not* break the flow of information through the measured system, and thus even though there is no *non-Markovian memory* (travelling through an environment), the observed statistics can still be correlated over multiple times. Put differently, after a general measurement, the state of the measured system is unknown, and might depend on earlier measurements, even though the underlying process itself does not exhibit any non-Markovian memory.

Importantly, this point is not critically related to any inherently “quantum” notion regarding the measurement (such as being a POVM comprising non-projective, non-Hermitian, or non-orthogonal elements), but can occur for *any* measurement for which an outcome does not fully determine the post-measurement state of the system. This can happen in classical physics for “fuzzy” measurements that coarse grain over different levels [131], and is generally the case for measurements in quantum mechanics described by CP maps that do *not* necessarily break the information flow through the system. For such measurements, the far past can still have an influence on the future [106, 128–132]. As a result, measurement invasiveness might not be detected at the next step, but possibly only further in the future, and any conditions pertaining to neighbouring dynamical maps alone are insufficient to characterise classicality. To see this explicitly, consider the following example, which concerns noisy (i.e., not rank 1) orthogonal measurements:

Example 4.5. Let ρ be the state of a four level system that is measured at times $\{t_1, t_2, t_3\}$ by means of projective—but not rank-1—measurements, i.e., $\mathcal{J}_i = \{K_i^{(1)} = \Pi_{(12)}, K_i^{(2)} = \Pi_{(34)}\}$, where $\Pi_{(xy)}$ is the projector on the space spanned by $\{|x\rangle, |y\rangle\}$, with $x, y \in \{1, 2, 3, 4\}$. Now, let the dynamics in between measurements be given by the Kraus operators $\{L_{2:1}^{(1)} = |1\rangle\langle 1| + |2\rangle\langle 4|, L_{2:1}^{(2)} = |2\rangle\langle 2| + |4\rangle\langle 3|\}$ and $\{L_{3:2}^{(1)} = \frac{1}{\sqrt{2}}(|3\rangle\langle 1| + |3\rangle\langle 2|), L_{3:2}^{(2)} = \frac{1}{2}(|1\rangle\langle 1| - |1\rangle\langle 2| - |2\rangle\langle 1| + |2\rangle\langle 2|) + |3\rangle\langle 3| + |4\rangle\langle 4|\}$, respectively. These choices of Kraus

operators correspond to CPTP maps, since $\sum_{\ell_{i+1}} L_{i+1:i}^{(\ell_{i+1})\dagger} L_{i+1:i}^{(\ell_{i+1})} = \mathbb{1}$ for $i \in \{1, 2\}$. It is easy to see that in this case, the statistics of the measurement at t_2 is independent of whether or not the measurement at t_1 was performed. Overall, the measurement at t_1 reduces the initial state ρ to a block diagonal structure; however, the statistics at t_2 only depend on the diagonal terms of ρ , such that the invasiveness of the first measurement is not detected. Specifically, we have

$$\begin{aligned}\mathbb{P}(m_2 = 1, \emptyset) &= \rho_{11} + \rho_{22} + \rho_{44} = \sum_{m_1} \mathbb{P}(m_2 = 1, m_1), \\ \mathbb{P}(m_2 = 2, \emptyset) &= \rho_{33} = \sum_{m_1} \mathbb{P}(m_2 = 2, m_1).\end{aligned}\quad (4.39)$$

As a result, the two-time statistics do not reveal the non-classicality of the observed statistics, despite the dynamics being Markovian. However, the invasiveness of the first measurement can be observed via the measurement at time t_3 . Concretely, the dynamics between t_2 and t_3 is such that it maps off-diagonal terms to diagonal ones, and thus the joint probability to measure $m_2 = 1$ and $m_3 = 1$ at times t_2 and t_3 (with *no* measurement at t_1), respectively, is given by

$$\mathbb{P}(m_3 = 1, m_2 = 1, \emptyset) = \frac{1}{2} (\rho_{11} - 2\text{Re}(\rho_{14}) + \rho_{22} + \rho_{44}). \quad (4.40)$$

Since the above probability depends on the entry ρ_{14} of the initial state ρ , it cannot coincide with the corresponding probability for the case where a measurement was performed at t_1 . As mentioned, the overall action of the measurement at t_1 is to force ρ into a block-diagonal structure, implying in particular $\rho_{14} \mapsto 0$ if a measurement at t_1 is performed. Consequently, we have

$$\sum_{m_1} \mathbb{P}(m_3 = 1, m_2 = 1, m_1) = \frac{1}{2} (\rho_{11} + \rho_{22} + \rho_{44}) \neq \mathbb{P}(m_3 = 1, m_2 = 1, \emptyset). \quad (4.41)$$

Accordingly, for the case of general instruments, one indeed must consider the full past and full future statistics for the relevant commutation relations in order to characterise classicality. ■

In the example above, we see that the invasiveness of the first measurement “skips” a time, i.e., it is not detected at time t_2 but rather only by the measurement at time t_3 . Such “skipping” of detectability is not limited to measurement invasiveness and has recently been analysed with respect to the activation of hidden quantum memory [344].

Such behaviour highlights the intricacies involved when considering quantum processes probed sequentially at multiple times by general instruments. However, for particular types of measurements, the flow of information through the system is broken, and one can therefore connect classicality to structural properties of the underlying dynamical maps between only *adjacent* times. This is, for example, the case if all measurements are rank-1 projective measurements in a fixed basis. Then, it is easy to see that $\tilde{\rho}_i(\mathbf{m}_{i-1:i}) \propto$

$\Lambda_{i:i-1}[|m_{i-1}\rangle\langle m_{i-1}|]$ and $Q_i(\mathbf{m}_{n:i+1}) \propto \Lambda_{i+1:i}^\dagger[|m_{i+1}\rangle\langle m_{i+1}|]$, where $\Lambda_{i+1:i}^\dagger[\bullet] = \sum_{\ell_{i+1}} L_{i+1:i}^{(\ell_{i+1})\dagger} \bullet L_{i+1:i}^{(\ell_{i+1})}$. With this, Eq. (4.19) reduces to

$$\begin{aligned} & \langle m_{i+1} | \Lambda_{i+1:i} \circ \Delta_i \circ \Lambda_{i:i-1}[|m_{i-1}\rangle\langle m_{i-1}|] | m_{i+1} \rangle \\ &= \langle m_{i+1} | \Lambda_{i+1:i} \circ \Lambda_{i:i-1}[|m_{i-1}\rangle\langle m_{i-1}|] | m_{i+1} \rangle \end{aligned} \quad (4.42)$$

for all t_i and all $\{m_{i-1}, m_{i+1}\}$, where $\Delta_i[\rho] := \sum_{m_i} \langle m_i | \rho | m_i \rangle |m_i\rangle\langle m_i|$ is the completely dephasing map at time t_i . Notably, Eq. (4.42) rephrases satisfaction of the Kolmogorov consistency conditions in terms of the properties of adjacent dynamical maps $\{\Lambda_{i+1:i}, \Lambda_{i:i-1}\}$ only, thus allowing for a full characterisation of Markovian dynamics that yield classical statistics when probed in a fixed basis, as is provided in Ref. [170]. However, this is only possible since for rank-1 projective measurements, the state of the system after measurement is known (up to normalisation). Any measurement with this property breaks the information flow through the system, in the sense that, upon observing a given outcome, the future outcome statistics of a Markovian process cannot depend on any previous outcomes, since the state of the system has been completely reset [128, 129, 344]. This, in turn, is what allows one to characterise the classicality of Markovian processes in terms of neighbouring dynamical maps only, as per Eq. (4.42). The above example shows why, in the case of *general* measurements, one must consider operators corresponding to the entire future and history when discussing classicality, even for Markovian processes, as we have done throughout this article.

Finally, given that for the special case of (rank-1) projective measurements, NCGD dynamics provides a necessary *and* sufficient condition for the classicality of the observed statistics, and all Kraus operators of the measurements are Hermitian, one might expect that Lüders-type assertions can be made with respect to commutation relations of pertinent operators (like those of Theorems 4.2 and 4.3). However, this is not the case, as the following example demonstrates:

Example 4.6. Consider a three step qubit process on times $\{t_1, t_2, t_3\}$ with measurements in the computational basis and intermediate dynamics given by the CPTP maps

$$\begin{aligned} \Lambda_{2:1}[\bullet] &= \frac{1}{2} \begin{pmatrix} 1 & 1 \\ i & -i \end{pmatrix} \bullet \begin{pmatrix} 1 & -i \\ 1 & i \end{pmatrix} \\ \text{and } \Lambda_{3:2}[\bullet] &= \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \bullet \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \end{aligned} \quad (4.43)$$

i.e., a rotation from the computational basis to the eigenbasis of σ_y between t_1 and t_2 , followed by a Hadamard gate between t_2 and t_3 . It is easy to see that, for measurements in the computational basis

$$\langle m_3 | \Lambda_{3:2} \circ \Delta_2 \circ \Lambda_{2:1}[|m_1\rangle\langle m_1|] | m_3 \rangle = \langle m_3 | \Lambda_{3:2} \circ \Lambda_{2:1}[|m_1\rangle\langle m_1|] | m_3 \rangle = \frac{1}{2} \quad (4.44)$$

for all $m_1, m_3 \in \{0, 1\}$, and thus the dynamics is NCGD [since it satisfies Eq. (4.42)]. However, it neither satisfies the inclusion property (4.32) of Theorem 4.3, nor any of the commutation relations we have discussed throughout this article. With respect to the former, it is easy to see that $\mathbb{H}_2 = \text{span}\{|+i\rangle\langle+i|, |-i\rangle\langle-i|\}$ and $\mathbb{F}_2 = \text{span}\{|+\rangle\langle+|, |-\rangle\langle-|\}$ holds, where $|\pm i\rangle = 1/\sqrt{2}(|0\rangle \pm i|1\rangle)$. Since these are the spaces spanned by the eigenvectors of σ_y and σ_x , respectively, neither of them is included within the other.

With respect to commutation relations, as mentioned above, we have $\tilde{\rho}_2(m_1) \propto \Lambda_{2:1}[|m_1\rangle\langle m_1|]$ and $Q_2(m_3) = \Lambda_{3:2}^\dagger[|m_3\rangle\langle m_3|]$, which implies (up to normalisation)

$$\begin{aligned}\tilde{\rho}_2(0/1) &= \frac{1}{2}(|0\rangle \pm i|1\rangle)(\langle 0| \mp i\langle 1|) =: |i\pm\rangle\langle i\pm|, \\ Q_2(0/1) &= \frac{1}{2}(|0\rangle \pm |1\rangle)(\langle 0| \pm \langle 1|) =: |\pm\rangle\langle \pm|.\end{aligned}\quad (4.45)$$

Together with $K_2^{(m_2)} = |m_2\rangle\langle m_2|$, we then obtain, for example,

$$[K_2^{(0)}, Q_2(0)] = \frac{1}{2}(|0\rangle\langle+| - |+\rangle\langle 0|) \neq 0, \quad (4.46)$$

i.e., commutativity à la Lüders (and Theorem 4.3) does not hold. Furthermore, as a consequence of Eq. (4.46), we have $[[K_2^{(m_2)}, Q_2(m_3)]] \propto \mathbb{1}$, such that

$$\text{tr}[\tilde{\rho}_2(m_1)[[K_2^{(m_2)}, Q_2(m_3)]]] = \text{tr}[\tilde{\rho}_2(m_1)] \neq 0. \quad (4.47)$$

Thus, for this example, neither commutativity nor absolute commutativity with respect to $\tilde{\rho}_2$ hold, and no inclusion property of the relevant spaces is satisfied. ■

Overall, we thus see that, even for the simple case of Markovian dynamics and projective measurements—where necessary and sufficient conditions for classicality are known in terms of NCGD dynamics—no commutation relations between the relevant operators are implied, at least none of the ones discussed in this paper.

4.4 Discussion and Conclusion

Throughout this article, we have analysed the connection between classicality and commutativity for Markovian processes probed at multiple points in time. In the two-time setting, it is straight forward to identify the pertinent operators whose commutativity should be assessed. Using the the availability of a full basis of input states, one can then proof an equivalence between commutativity and non-invasiveness, providing a connection between operational and structural notions of classicality. In the multi-time setting, Kolmogorov consistency conditions provide an operationally meaningful notion of classicality; however, it is, a priori, unclear what the relevant operators are to check for commutativity. Here, we have identified the relevant operators, and our work can be seen as a multi-time extension of Lüders' theorem. As discussed, many crucial assumptions of

Lüders' theorem immediately break down (or become too restrictive) in the multi-time setting, e.g., the guarantee of a full basis of system states at each time. Nonetheless, we have detailed the relevant operators and commutator expressions that imply a connection between operational and structural notions of classicality, putting these distinct notions on a comparable mathematical footing. We have thus overcome a number of complications that arise naturally in physically meaningful scenarios, including probing open system dynamics over multiple times with general quantum measurements.

In particular, in Section 4.3.2, we first exemplified how Kolmogorov consistency does not guarantee the vanishing of the analogous commutator expression to the one Lüders originally considered, and subsequently derived a novel relevant “absolute” commutator expression that indeed implies that satisfaction of Kolmogorov consistency conditions (see Theorem 4.2). Following this, in Section 4.3.3, we derived additional assumptions such that Kolmogorov consistency implies commutativity (see Theorem 4.3). Lastly, in Section 4.3.4, we connected our results with existing literature to demonstrate the connection between commutativity, classicality, and the ability of the dynamics to generate and detect coherence with respect to sharp measurements in a fixed (but otherwise arbitrary) basis.

Our results provide a connection between commutation relations and the classicality of the observed statistics. However, the absence of necessary *and* sufficient conditions, highlighted via the examples of processes that satisfy none, or just some, of the commutator relations we identified demonstrate that in the multi-time case, a direct connection between mathematical and operational notions of classicality is far more elusive than in the two-time case (even in the simplest case of projective measurements in a fixed basis). Looking forward, our work opens the door to a number of interesting avenues for exploration. Following our general exposition regarding the structural implications of operational classicality, it would firstly be interesting to identify necessary *and* sufficient conditions for the classicality of observed statistics. While this is a daunting task in general, starting from our considerations, such results might be readily derivable for dynamics that are particularly relevant to certain physical situations: Just as NCGD dynamics equates structural properties to Kolmogorov consistency, we expect it to be possible to derive similarly strong correspondences between particular types of dynamics (e.g., dephasing, depolarising, thermalising, etc.) and the classicality of statistics observed for certain types of instruments (e.g., measure-and-prepare, unital instruments, etc.). Since Kolmogorov consistency ensures the existence of an underlying classical stochastic process that reproduces the statistics correctly, this would in turn shed light on the types of noise that can be effectively replaced by classical environments [345], which would have profound impact on the fields of optimal quantum control, reservoir engineering, and the simulation of complex open dynamics.

Hidden Quantum Memory: Is Memory There When Somebody Looks?

Philip Taranto and Simon Milz

Abstract. In classical physics, memoryless processes and Markovian statistics are one and the same. This is not true for quantum processes, first and foremost due to the fact that quantum measurements are invasive. Independently of measurement invasiveness, here we derive a novel distinction between classical and quantum processes, namely the possibility of *hidden quantum memory*. While Markovian statistics of classical processes can always be reproduced by a memoryless dynamics, our main result establishes that this is not the case in quantum mechanics: We first provide an example of quantum non-Markovianity that depends on whether or not a previous measurement is performed—a phenomenon that is impossible for memoryless processes; we then strengthen this result by demonstrating statistics that are Markovian independent of how they are probed, but are nonetheless *still* incompatible with memoryless quantum dynamics. Thus, we establish the existence of Markovian statistics that *fundamentally* require quantum memory for their creation.

Under review. Manuscript submitted on 27 Apr 2022.

[arXiv:2204.08298](https://arxiv.org/abs/2204.08298)

Author Contribution

In this work, the doctoral candidate contributed to the conception and formulation of the theoretical framework and methods, the proofs of the main results, the writing and revising of the manuscript, and the organisation and supervision of the project overall. In particular, the main technical contributions of the doctoral candidate were the formulation regarding the distinction between operational and structural notions of memory; the subsequent analysis of said properties in both the classical and quantum settings; the proofs of all lemmata, corollaries and theorems; and the demonstration of hidden quantum memory and incompatibility via the two examples.

5.1 Introduction

Our ability to understand and control memory effects in the evolution of open quantum systems is becoming increasingly important as technology allows us to manipulate interactions with increasing levels of speed, precision and complexity [50, 51]. Control over memory can be advantageous in various tasks, such as the creation, manipulation and preservation of coherences and correlations [69, 110], reservoir engineering to simulate complex dynamics [74, 77, 79, 111, 112, 114–119, 346], sophisticated randomised benchmarking and quantum error correction [120–122], optimal dynamical decoupling [68, 123, 347], designing quantum circuit architectures [125–127, 133–136], and improving the efficiency of thermodynamic machines [73, 75, 207].

One has no choice but to account for complex noise and memory effects when modelling realistic dynamical systems, as no system is truly isolated; in general, the environment stores information about the past and perpetuates it in time, leading to memory effects that manifest themselves as complex multi-time correlations [106–109, 140]. A special case of open dynamics are memoryless processes, for which the environment retains *no* memory of previous interactions with the system.

Memoryless dynamics have been studied extensively due to their accuracy in many practically relevant situations and their exponentially reduced complexity from the general scenario. Both in the classical and quantum setting, such efficient descriptions arise by way of (time-local) master equations that efficiently simulate the system dynamics [281, 282, 348]; in practice, the assumption of memorylessness is often made for simplicity and describes many “real-world” scenarios with high accuracy [65, 66, 272, 279].

However, experimentally determining that a quantum process is indeed memoryless requires full process tomography, which necessitates a myriad of complex sequential measurements and has consequently only been done for low-dimensional cases [133, 134, 136, 349]. A more experimentally tractable situation is the sequential probing of a fixed observable via sharp, projective measurements. In this case, memoryless quantum processes—just like their classical counterparts—lead to Markovian statistics, i.e., statistics where the future is conditionally independent of the past. Thus, at first glance, memorylessness seems to manifest on the experimental level in the same way for both classical and quantum processes. However, this is not the case; for one, quantum measurements are generally invasive, leading to inconsistent statistics [146] and the violation of Leggett-Garg type inequalities [147–149]. Moreover—independent of measurement invasiveness—as we demonstrate in this Letter, quantum processes can yield Markovian statistics *that fundamentally require memory* for their creation.

More concretely, in classical physics, *any* Markovian statistics can be described by a memoryless dynamical model (i.e., as emerging from a set of independent stochastic matrices). In the quantum case, measurements in a fixed basis are not tomographically

complete; consequently, the existence of processes with memory that nonetheless lead to Markovian statistics when probed in said basis is not surprising *per se* and has been demonstrated [130–132]. This phenomenon notwithstanding, for any experiment that yields Markovian statistics, it is reasonable to believe that there *always exists* some memoryless quantum process that faithfully reproduces the observed statistics. Such a description is known as the Quantum Regression Formula (QRF) [272, 279, 350] and is a widely used assumption that links operational quantities—namely, recorded statistics—to dynamical ones—namely, a model of the underlying process.

Here, we ask the question: *Can Markovian statistics always be faithfully reproduced by a memoryless dynamical model?* In other words, can the QRF always be employed to describe Markovian statistics? Our main result, perhaps surprisingly, answers this in the negative. Since this is not possible in classical physics, we thus uncover a new type of genuinely quantum phenomenon: *Hidden quantum memory*. This observation makes quantum memory an emergent phenomenon: Observing Markovianity with respect to a fixed measurement basis is not sufficient to guarantee the existence of a memoryless dynamical descriptor. Such hidden quantum memory is similar in spirit to other quantum traits that require precisely the resource in their implementation that they ultimately hide, such as quantum channels that preserve all separable states but cannot be implemented via local operations and classical communication [167, 181, 182], non-signalling maps that require signalling [183], and maximally incoherent operations that necessitate coherent resources [184–186], to name but a few. We begin by outlining the envisaged setup before detailing key properties of memoryless processes (both classical and quantum).

5.2 Framework

In any experimental scheme concerning temporal processes, an experimenter probes a system of interest at (any subset of) times $\mathcal{T}_n := \{t_1, \dots, t_n\}$ (with $t_n > \dots > t_1$) and records the corresponding probability distributions $\{\mathbb{P}(\mathbf{x}_\Gamma)\}$, where $\Gamma \subseteq \mathcal{T}$ and $\mathbf{x}_\Gamma := \{x_j | t_j \in \Gamma\}$ (see Fig. 5.1). These capture, for instance, the probability that x_1 is observed at time t_1 and x_2 at t_2 , and so on, with all possible combinations of measurement times. Note that the experimenter can also *not* make a measurement at any intermediate time, e.g., $\mathbb{P}(x_3, x_1)$.

Independent of the physical scenario—it could be classical, quantum, or even post-quantum—one can define the concept of *Markovianity* based on the observed statistics alone, as conditional independence of any current outcome from all but the most recent. Concretely, we have the following working definition:




Figure 5.1: Probing a Process. By probing a process—for example Brownian motion, or the evolution of a spin degree of freedom—sequentially (here, at times $\mathcal{T}_4 = \{t_1, t_2, t_3, t_4\}$, an experimenter can deduce the probability distribution $\mathbb{P}(x_4, x_3, x_2, x_1)$. In the classical case, this also includes all “contained” distributions, for example $\mathbb{P}(x_4, x_2, x_1)$ via marginalisation [see Eq. (5.4)]. In the quantum case, due to invasiveness, deducing said distributions would require a new experiment where *no* measurement is performed at t_3 (depicted by the shaded box).

Definition 5.1. A *Markovian statistics* on a set of times \mathcal{T}_n is a collection of conditional probability distributions $\{\mathbb{P}(x_j|x_{j-1}, \dots, x_1)\}_{t_j \in \mathcal{T}_n}$ for which

$$\mathbb{P}(x_j|x_{j-1}, \dots, x_1) = \mathbb{P}(x_j|x_{j-1}) \quad (5.1)$$

for all $t_j \in \mathcal{T}_n$.

Defined as such, the question of Markovianity is, a priori, theory-independent and concerns only the observed statistics. As we shall see, though, the concept of conditional probabilities is a subtle issue that depends on the envisaged scenario. Throughout this Letter, we distinguish Markovianity from the notion of *memoryless processes* which corresponds to the memory properties of the *underlying* dynamics that engenders the observed statistics, making the latter a *theory-dependent* concept.

Specifically, the question of memorylessness concerns whether, in the evolution of a system that is coupled to some inaccessible environment, said environment perpetuates past information about the system forward in time or irretrievably dissipates it. The description of such open evolution differs across physical theories: In the classical setting, the most general state transformations are stochastic matrices, whereas in the quantum realm, these are quantum channels. Probability distributions arising from interrogating either classical or quantum processes therefore have different properties since they are calculated via different rules. Our main result shows that this distinction holds for the relationship between Markovianity and memorylessness: Although equivalent in the classical case, in the quantum realm the observation of Markovian statistics does *not* guarantee even the existence of a memoryless dynamics that engenders them.

5.3 Classical Processes

We begin with a discussion of memoryless *classical* processes:

Definition 5.2. A *memoryless classical process* on \mathcal{T}_n is a set of mutually independent stochastic matrices $\{S_{j:j-1}\}_{j=2,\dots,n}$ and an initial state (i.e., probability vector) \mathbf{p}_1 such that the probability distribution over any sequence of outcomes x_1, \dots, x_n is given by

$$\mathbb{P}(x_n, \dots, x_1) = \langle x_n | S_{n:n-1} | x_{n-1} \rangle \langle x_{n-1} | \dots | x_2 \rangle \langle x_2 | S_{2:1} | x_1 \rangle \langle x_1 | \mathbf{p}_1, \quad (5.2)$$

where $|x_j\rangle\langle x_j|$ are projectors corresponding to measurement outcomes x_j .

Although in classical physics the environment plays a role in dictating the state transitions between any times t_{j-1} and t_j —namely via the stochastic matrices $S_{j:j-1}$, which are matrices with non-negative elements whose columns sum to unity—in *memoryless* processes, the environment does not perpetuate information, i.e., the stochastic matrices in Eq. (5.2) are mutually *independent*. On the other hand, Markovianity (see Definition 5.1) concerns only the observed statistics [l.h.s. of Eq. (5.2)]. In classical physics, we can make the following simple observation (see Appendix D.1):

Observation 1. *In the classical setting, memoryless processes are equivalent to Markovian statistics.*

Specifically, this equivalence is given by setting $\langle x_j | S_{j:j-1} | x_{j-1} \rangle = \mathbb{P}(x_j | x_{j-1})$, and it follows from Eq. (5.2) that for any Markovian statistics/memoryless classical process we have

$$\mathbb{P}(x_n, \dots, x_1) = \mathbb{P}(x_n | x_{n-1}) \mathbb{P}(x_{n-1} | x_{n-2}) \dots \mathbb{P}(x_2 | x_1) \mathbb{P}(x_1). \quad (5.3)$$

In one direction, Observation 1 states that for any memoryless process, the observed statistics are Markovian—this is also true in the quantum setting (see below). Conversely, if one records Markovian statistics, then one can always construct a unique, memoryless process that faithfully reproduces them—as we will see, this is not true for statistics gathered from quantum processes.

A major distinction between classical and quantum processes (memoryless or not) is that in the classical realm, the single n -time probability distribution $\mathbb{P}(x_n, \dots, x_1)$ contains the entire set of statistics on all subsets of times $\Gamma \subseteq \mathcal{T}_n$. That is, the probability $\mathbb{P}(\mathbf{x}_\Gamma, \mathcal{I}_{\bar{\Gamma}})$ to observe a sequence of outcomes \mathbf{x}_Γ when probing the process at times Γ and *not* measuring (denoted by the “do-nothing instrument” $\mathcal{I}_{\bar{\Gamma}}$) at the remaining times $\bar{\Gamma} := \mathcal{T}_n \setminus \Gamma$ can be deduced via marginalisation

$$\mathbb{P}(\mathbf{x}_\Gamma, \mathcal{I}_{\bar{\Gamma}}) = \sum_{\mathbf{x}_{\bar{\Gamma}}} \mathbb{P}(x_n, \dots, x_1), \quad (5.4)$$

This *non-invasiveness* of measurements in classical physics underlies Observation 1 and similarly fails to hold in quantum mechanics. As a direct consequence of measurement non-invasiveness, the properties of a memoryless classical process on \mathcal{T}_n translate to all “sub-processes” that are probed only at times $\Gamma \subseteq \mathcal{T}_n$ (see Appendix D.1):

Corollary 5.1. *All sub-statistics of a memoryless classical process are Markovian and the corresponding conditional probabilities are mutually compatible.*

By *compatible*, we mean that all conditional probabilities are independent of how they are obtained, i.e.,

$$\frac{\mathbb{P}(x_j, \mathbf{x}_{\Gamma^{(i)}})}{\mathbb{P}(\mathbf{x}_{\Gamma^{(i)}})} = \frac{\mathbb{P}(x_j, \mathbf{x}_{\Gamma^{(i)'}})}{\mathbb{P}(\mathbf{x}_{\Gamma^{(i)'}})} =: \mathbb{P}(x_j|x_i), \quad (5.5)$$

for all $t_j, t_i \in \mathcal{T}_n$ (with $t_j > t_i$) and all subsets $\Gamma^{(i)}, \Gamma^{(i)'} \subseteq \mathcal{T}_n$ that contain t_i as their largest time; for a classical memoryless process, knowledge of any outcome x_i suffices to erase all historic information (including whether or not a previous measurement was made) and is therefore the only relevant parameter for predicting future outcomes. Such compatibility between Markovian sub-statistics of a memoryless quantum process also holds, although it is less obvious, and we will later employ the breakdown of compatibility as a witness for memory.

5.4 Quantum Processes

In contrast to classical physics, in quantum mechanics, measurements are generally *invasive* such that there is a difference between averaging over outcomes and not performing a measurement. This makes (conditional) probabilities protocol-dependent entities that require specification; in what follows, whenever we consider a probability distribution $\mathbb{P}(\mathbf{x}_\Gamma)$ (or corresponding conditional probabilities), we always mean the statistics obtained from *only* performing measurements at times in the set $\Gamma \subseteq \mathcal{T}_n$, and doing nothing (denoted by $\mathcal{I}_{\bar{\Gamma}}$ at the remaining times $\bar{\Gamma} = \mathcal{T}_n \setminus \Gamma$). Importantly, in quantum mechanics—in contrast to Eq. (5.4)—generally $\mathbb{P}(\mathbf{x}_\Gamma) := \mathbb{P}(\mathbf{x}_\Gamma, \mathcal{I}_{\bar{\Gamma}}) \neq \sum_{\mathbf{x}_{\bar{\Gamma}}} \mathbb{P}(x_n, \dots, x_1)$. Such measurement invasiveness is well-studied and has recently been used as a witness for the non-classicality of an underlying process [146, 170, 340, 341]. Despite these added subtleties in the definition of (conditional) probabilities, memoryless quantum processes lead—just like in the classical case—to well-defined, compatible Markovian statistics and sub-statistics. To see this, let us first provide the generalisation of Definition 5.2 to the quantum case:

Definition 5.3. *A memoryless quantum process on \mathcal{T}_n is a set of mutually independent Completely Positive and Trace Preserving (CPTP) maps $\{\Lambda_{j:j-1}\}_{j=2,\dots,n}$ and an initial state (density operator) ρ_1 such that the probability distribution over any sequence of outcomes x_1, \dots, x_n is given by*

$$\mathbb{P}(x_n, \dots, x_1) = \text{tr} \left[\mathcal{P}_n^{(x_n)} \Lambda_{n:n-1} \dots \Lambda_{2:1} \mathcal{P}_1^{(x_1)} [\rho_1] \right], \quad (5.6)$$

where $\mathcal{P}_j^{(x_j)}[\bullet] := |x_j\rangle\langle x_j|\bullet|x_j\rangle\langle x_j|$ are maps corresponding to sharp (i.e., rank-1) projective measurements.

Analogous to the classical case, CPTP maps are the most general state transformations in the presence of noise, and the absence of memory corresponds to the mutual independence of the maps $\Lambda_{j:j-1}$ in the definition. The above equation to compute probabilities is commonly known as the Quantum Regression Formula (QRF) [272, 279, 350]. Importantly, it allows for the computation of sub-statistics on any $\Gamma \subseteq \mathcal{T}_n$, not via marginalisation, but by replacing the projection operators corresponding to times in $\bar{\Gamma}$ in Eq. (5.6) by identity maps. Of course, one need not perform projective measurements, and the above formula can be used to calculate the probability distribution over *any* sequence of outcomes for arbitrary instruments. In contrast to classical measurements, such general quantum measurements do not necessarily reset the state of the system, which means that memoryless quantum processes can lead to non-Markovian statistics for general instruments [128–131]. However, when restricted to sharp, projective measurements, then—just as in the classical setting—memorylessness in the quantum realm manifests itself on the observational level as Markovianity (see Appendix D.2):

Lemma 5.1. *Any memoryless quantum process leads to Markovian statistics (for sharp, projective measurements).*

We saw that memoryless classical processes also lead to (compatible) Markovian sub-statistics (see Corollary 5.1), where compatibility is given by Eq. (5.5). This is also true for memoryless quantum processes, with the important difference that sub-statistics are not obtained by marginalisation, but by “doing nothing” at the excessive times, i.e., performing the experiment in a different way. Probing sub-statistics in this way then yields meaningful conditional probabilities and we have the following (see Appendix D.2):

Lemma 5.2. *Any memoryless quantum process leads to Markovian sub-statistics (for sharp, projective measurements) that are mutually compatible.*

In both quantum mechanics and classical physics, memoryless processes—when probed in a fixed basis—*always* lead to Markovian statistics and Markovian, compatible sub-statistics. In the classical setting, the converse is also true. From the observation of Markovian statistics one can construct a memoryless process describing the situation at hand. As discussed, measuring a quantum process in a fixed basis cannot provide enough information to fully determine the underlying process. Nonetheless, it is reasonable to assume that whenever one observes Markovian statistics, there should exist some memoryless description that correctly reproduces them (indeed, this is the assumption of employing the QRF to describe Markovian statistics). Thus, we now ask the question:

Given Markovian statistics (deduced via sharp, projective measurements), does there always exist a memoryless quantum process that faithfully reproduces them?

5.5 Hidden Quantum Memory and Incompatibility

We answer the above question in the negative, first by showing the existence of Markovian statistics with non-Markovian sub-statistics, and then by constructing a process with Markovian statistics and sub-statistics that are however incompatible.

Observation 2. *Given Markovian statistics on \mathcal{T}_n (deduced via sharp, projective measurements), there does not always exist a memoryless quantum process that faithfully reproduces them.*

Proof. Our proof is by way of constructing an explicit example, depicted in Fig. 5.2. The process is over four times and the experimenter always measures in the computational σ_z -basis. An initial state $\rho_1 = \frac{1}{2}$ is sent to the experimenter, who measures it. The process between times t_1 and t_2 is a Hadamard gate. Following the measurement at t_2 , the system is swapped with a fiducial environment state $\tau = |0\rangle$, which is what the experimenter measures at time t_3 . Meanwhile, the dynamics of the environment consists of a measurement in the σ_x -basis, followed by a preparation of $|0\rangle(|1\rangle)$ whenever $+(-)$ is recorded. Between times t_3 and t_4 , the process comprises a Controlled-NOT (CNOT) gate, controlled on the environment. Due to the gates that act on the system and environment, this circuit can, in principle, display memory effects for the system dynamics. In Appendix D.3, we calculate the full statistics $\mathbb{P}(x_4, x_3, x_2, x_1)$ and show them to be Markovian, i.e., $\mathbb{P}(x_4|x_3, x_2, x_1) = \mathbb{P}(x_4|x_3)$ and $\mathbb{P}(x_3|x_2, x_1) = \mathbb{P}(x_3|x_2)$. This is because the measurement in the σ_z -basis at t_2 yields an output state that is unbiased with respect to the σ_x -basis measurement in the environment and therefore all memory of x_1 is lost. However, by calculating the sub-statistics where the experimenter does *not* measure at time t_2 , i.e., $\mathbb{P}(x_4, x_3, \mathcal{I}_2, x_1)$, we see that they are *non*-Markovian since information about x_1 is now *not* fully scrambled by the “intervention” (or rather lack thereof) at t_2 , and we have $\mathbb{P}(x_4|x_3, \mathcal{I}_2, x_1) \neq \mathbb{P}(x_4|x_3)$ with dependence on x_1 . As we proved in Lemma 5.2, this cannot happen for *any* memoryless quantum process. Thus, even though the statistics on \mathcal{T}_n is Markovian, there is *no* memoryless quantum process that faithfully reproduces the statistics on all four times, since the sub-statistics fail to be Markovian. \square

Here, we have uncovered a new type of genuinely quantum phenomenon: *Hidden quantum memory*. The fact that full statistics can be Markovian but sub-statistics can be non-Markovian is impossible in the classical realm. Moreover, this property cannot occur for memoryless quantum processes either. Thus, we have shown the existence of Markovian statistics that, not only potentially come from a quantum process with memory




Figure 5.2: Markovian Statistics that Require Memory. When probed by projective measurements in the σ_z -basis at *all* times, the circuit yields Markovian statistics. Without the Exclusive-OR (XOR) gate, memory becomes apparent when no measurement is performed at t_2 (Observation 2). With the XOR gate, all sub-statistics are Markovian, but mutually incompatible (Observation 3). Both situations are *only* possible in the presence of memory.

(which can happen, as is well known, when measured in a fixed basis), but *fundamentally require* memory to reproduce them.

Another way of viewing this result is that non-Markovian sub-statistics serves as a witness for the necessity of memory in the underlying quantum process. This naturally begs the question: *If the full statistics and all sub-statistics are Markovian, does there always exist a memoryless quantum process that faithfully reproduces them?* In other words, is the ability to detect non-Markovian sub-statistics a requirement for ruling out a memoryless description of the quantum process? Here, we also answer this in the negative, providing a stronger result than above:

Observation 3. *Given Markovian statistics and sub-statistics on \mathcal{T}_n and all subsets thereof (deduced via sharp, projective measurements), there does not always exist a memoryless quantum process that faithfully reproduces them.*

Proof. The proof is again by constructing an explicit example. The process is similar to above, with the only modification being that after the final CNOT gate, there is an additional Exclusive-OR (XOR) gate performed on both system and environment (see Fig. 5.2). In Appendix D.3, we explicitly calculate the full statistics $\mathbb{P}(x_4, x_3, x_2, x_1)$ and all sub-statistics [e.g., $\mathbb{P}(x_4, x_3, \mathcal{I}_2, x_1)$, etc.] and show them to be Markovian. However, crucially, the conditional probabilities $\mathbb{P}(x_4|x_3, x_2, x_1)$ and $\mathbb{P}(x_4|x_3, \mathcal{I}_2, x_1)$ differ, i.e., we have $\mathbb{P}(x_4|x_3, x_2, x_1) =: \mathbb{P}(x_4|x_3) \neq \tilde{\mathbb{P}}(x_4|x_3) := \mathbb{P}(x_4|x_3, \mathcal{I}_2, x_1)$. As stated in the second part of Lemma 5.2, such an incompatibility cannot occur for memoryless quantum processes, and so we again come to a contradiction with the existence of a memoryless process that engenders the observed statistics. \square

5.6 Conclusions

In this Letter, we have presented the concept of hidden quantum memory, i.e., the existence of processes that yield Markovian statistics that *cannot* be explained without underlying memory effects. In a similar vein to the violation of Leggett-Garg type inequalities, this phenomenon can only occur when the performed measurements are invasive, since otherwise the observed statistics are fully classical and Markovianity and memorylessness coincide. However, hidden quantum memory is not merely a different manifestation of measurement invasiveness, but a distinct and fundamental quantum memory effect; while memoryless quantum processes can violate Leggett-Garg type inequalities, they cannot lead to hidden quantum memory.

Importantly, our results differ from the (known) fact that probing a quantum processes with memory in a fixed basis can yield Markovian statistics. For the Markovian statistics that we reported, there exists *no* memoryless model that reproduces them, either because they become non-Markovian when measurements are not performed at some times, or because all observed statistics and sub-statistics are Markovian but mutually incompatible. In turn, this implies that even if one observes Markovian statistics in a given basis, one cannot confidently employ a QRF to compute the statistics in said basis. As a consequence, even detecting the possibility of a memoryless description of a process is an experimentally complex undertaking that not only requires one to deduce joint probabilities on \mathcal{T}_n , but also on all subsets thereof. Naturally, one might expect that simultaneously demanding Markovianity *and* compatibility of all observed sub-statistics should suffice to guarantee a memoryless description. However, even under such strong requirements, the existence of a memoryless model is a priori not clear, and investigations into this question are subject to future work.

Together, our results expose a fundamentally new quantum effect in time and demonstrate the richness of effects that arise from the intricate interplay of invasiveness, memory, and the freedom to choose different instruments that quantum mechanics affords.



Concluding Discussion

Summary

Throughout this dissertation, we have demonstrated the connection between thermodynamics, complexity, and multi-time phenomena from a variety of angles, highlighting their complicated relationship and the ensuing implications for quantum information processing.

We began in **Part I**, where we studied the archetype of a difficult task from a thermodynamic perspective, namely cooling a physical system (or, equivalently, erasing information). In **Chapter 1**, we developed a unifying framework to encompass the—seemingly contradictory—statements of both Landauer and Nernst by explicitly accounting for the level of control, in full thermodynamic spirit. From this vantage point, we considered two extremal levels of overall control—namely the coherent control setting (where one has access to either a quantum battery or classical control with a precise clock) and the incoherent control setting (where one can only use a heat engine to drive the transformations)—and provided a number of fundamental insights regarding the asymptotic attainability of what is technically forbidden, i.e., reaching a pure state.

First, given coherent resources, we constructed protocols that saturate the Landauer limit when any one of three relevant resources—namely energy, time, or control complexity—diverges. Our protocols apply to both finite and infinite dimensional systems and machines with arbitrary Hamiltonians, thereby conclusively exploring the ultimate limitations of quantum information processing with thermal machines across a variety of platforms. Second, in the incoherent control paradigm, we began by showing a stark contrast with that of coherent control by proving a “no-go” theorem that implies the impossibility of preparing a pure state with either finite time or control complexity (i.e., a finite sized machine) regardless of how much energy is drawn from a hot bath. We then derived the ultimate limitation for the energy cost of cooling in this fully thermodynamic setting, namely the Carnot-Landauer bound, and provided protocols for its asymptotic saturation.

Following this, in **Chapter 2**, we shifted our focus away from the asymptotic setting and considered how finite amounts of resources can be traded off against one another to cool a system. Disregarding the energy cost, we showed that control over memory effects—a relevant proxy for control complexity in the temporal setting—leads to an

exponential enhancement of the ground state population (and thus in cooling) at any given time. By considering the task of cooling with memory within the framework of generalised collision models, we demonstrated a powerful tool for treating memory effects in quantum thermodynamics and showed that dramatic improvements in refrigeration performance are possible with realistic (i.e., finite) control. As the corresponding results here recover the well-studied memoryless setting as a special case, we have unified a number of seemingly disparate approaches to cooling with practical resources.

Focussing more concretely on the properties of memory itself, we then dropped the thermodynamic assumptions and considered the intricacies regarding quantum information processing within the more general open dynamics framework in **Part II**. Whenever an environment interacts non-trivially with a system of interest, it generically perpetuates memory effects, which lead to complicated multi-time correlations as a manifestation of the underlying complexity. Continuing with our previously developed perspective, on the one hand, if one can control such memory effects, they can be incredibly useful (as we showed regarding cooling with memory); on the other hand, if they are uncontrollable, then they are observed as correlated noise, making processes with memory incredibly difficult to characterise and manipulate advantageously (as is well known).

In **Chapters 3** and **4**, we asked the question: *Which traits are fundamentally quantum and what resources are required to observe non-classical behaviour?* Of course, there are a plethora of potential candidates (coherence, discord, non-commutativity, to name but a few), all of which being well-justified as markers of non-classicality in their own right. As we showed throughout the second part of this dissertation, their relationship to a fully operational notion of classicality—namely measurement non-invasiveness (or Kolmogorov consistency), which implies the existence of a potentially exotic underlying classical stochastic process—plays out over a rich and interesting tapestry.

In **Chapter 3**, we first demonstrated that a quantum process need not be able to generate or detect coherence in the system state to yield non-classical statistics. We then analysed the connection between non-classicality and the ability to generate and detect quantum discord between the system and environment, highlighting that while the inability regarding the latter implies classicality, the converse is not true. Both of these results demonstrate a clear distinction between processes with memory and those without. Perhaps the most emblematic and insightful take-away from these results is the realisation that, while all memoryless quantum processes can hide their non-classicality in principle (via measurements performed with respect to the eigenbasis of the system), there exist quantum processes with memory that are *genuinely* non-classical, in the sense that their non-classicality can never be hidden, no matter the choice of instruments.

We subsequently analysed the connection between measurement non-invasiveness and commutativity in **Chapter 4**. The results of this chapter can be read as a multi-time

extension of Lüders’ theorem, which provides a one-to-one relationship in the two-time setting. However, as we demonstrated by way of various examples, this link is not so clear in the multi-time setting, even for memoryless processes. Nonetheless, we uncovered various Lüders-type theorems in the sense that they relate the vanishing of pertinent commutator expressions—typically involving a complicated interplay of the dynamics, the system state, and the measuring devices themselves at each time—with said operational notion of classicality.

Finally, having detailed a number of insightful distinctions between classical and quantum processes with respect to the criterion of measurement non-invasiveness, in **Chapter 5** we considered the connection between the underlying dynamical property of *memorylessness* and the related operationally observed notion of *Markovianity*. In the classical setting, these two concepts are one and the same; however, this connection is not immediately clear for quantum processes due to measurement invasiveness. We asked the question: *Given Markovian statistics deduced from probing a quantum process with sharp, projective measurements, does there always exist an underlying memoryless dynamical model that faithfully reproduces said statistics?* Perhaps surprisingly, we answered this question in the negative by constructing processes that lead to Markovian statistics which can either have “hidden” non-Markovian, or Markovian but nonetheless incompatible, sub-statistics. In addition to being impossible in the classical setting, such multi-time phenomena stand in direct contradiction with the possibility of an underlying memoryless quantum description. Thus, we have here uncovered a novel type of genuinely quantum effect, namely the existence of Markovian statistics that *fundamentally require* memory for their physical implementation.

Overall, our work throughout this dissertation makes significant progress towards a more holistic understanding of the crucial role of control regarding quantum information processing. We highlighted this in the first part by demonstrating the impact that control has on thermodynamic tasks, showing—amongst other insights—that control complexity *must* be considered as a meaningful thermodynamic resource. In the second part, we considered the control over physical systems in time, i.e., over memory effects, and demonstrated the ensuing implications in terms of complex multi-time phenomena; yet again, we saw that complexity (either in the form of memory effects or non-classicality) depends upon the interplay between the underlying process as well as the control one has to probe it. Coming full circle to highlight the intricacies involved regarding a suitable understanding of complexity, we finally showed that simple observations—e.g., Markovian statistics—can not only arise from complex processes (as is well known), but can in some cases indeed *necessitate* underlying complexity (in our case, dynamics *with* memory). We now move to consider the broader implications of our work, in particular the rich landscape for future investigation.

Outlook

We anticipate that our introduction of control complexity as a resource will initialise a shift in the way that resources are perceived, especially in quantum thermodynamics, but also more broadly in all areas of quantum information processing, from quantum computing to quantum sensing and beyond.

In particular, our results lay the foundations for a plethora of practically relevant follow-up opportunities concerning finite-resource trade-offs in quantum technologies and the intricate relationship between energy, time and control complexity. For instance, given a finite maximum amount of energy, time, and control (which could concern limitations on interaction Hamiltonians, machine sizes, gate sets, etc.) that one is willing or able to invest, *what is the optimal allocation of such resources to reach a desired final temperature?* As we discussed, such an optimisation is complicated and perhaps no “one-size-fits-all” solution exists; nonetheless, we expect there to be assumptions that can be reasonably justified with respect to certain platforms that might lead to further insights. Additionally, directly connecting the notion of control complexity to what is experimentally feasible in near-term quantum platforms will provide another important step to eventually transcend the era of Noisy Intermediate-Scale Quantum (NISQ) devices.

More generally, properly accounting for the overall cost of control in quantum information processing is critical, as it can quickly outweigh any potential benefits to be had. In other words, any reported “quantum advantage” over a classical protocol might not be so desirable if the quantum algorithm requires substantially more resources to implement. This is in stark contrast with the classical setting, where states encoding information are generally stable, errors easily correctable, and read-out procedures accurate—all of which allows one to more-or-less neglect the cost of control and impact of heat dissipation in classical computers. From this perspective, it would seem desirable to understand the limitations of quantum information processing in a setting that embodies the assumption of minimal control at the outset, as we have put forth regarding the thermodynamic paradigm of incoherent control, where one only requires the ability to switch on and off interactions. This, of course, comes with its own resource cost—namely that of precise timing [57, 58]—which should also be accounted for a more rounded treatment moving forwards. Nonetheless, it will be interesting to understand the ultimate possibilities of such a thermodynamically driven quantum computer: *What kind of algorithms can be run efficiently? To what extent is it scalable? Which experimental platforms could be used in the near term to build such devices?*

On a more foundational level, we have shown that a diverging level of either control or time is required to saturate the Landauer limit and asymptotically attain the “unattainable”, namely a perfectly pure state. The Landauer bound is closely related to the Second Law of Thermodynamics, which states that the sum of local entropy changes in

any (initially uncorrelated) closed system-environment dynamics cannot decrease; on the other hand, the unattainability principle is essentially the Third Law of Thermodynamics. Unlike the Third Law, which explicitly focuses on the ability to prepare a perfectly pure state, the Second Law makes no assumptions on the purity of the final state. Although one would require a high degree of control to cool a system to *any* finite temperature at the Landauer limit, such a task is, in principle, possible with finite time and control resources; it is only when one wants to cool the system to absolute zero that this resource requirement necessarily diverges. On the other hand, the Third Law makes no mention of the efficiency or energy cost—given any finite amount of energy (which could be much larger than the Landauer cost), one would still require either diverging time or control complexity to prepare a pure state asymptotically. In this sense, we see the intricate relationship between the Second and Third Laws of Thermodynamics (in the asymptotic setting) in terms of *control as a resource*: The former dictates the minimum energy requirements for any entropy-decreasing (i.e., irreversible) procedure, which requires ever-increasing levels of control to saturate as the system gets cooler; whereas the latter concerns itself with how cold a system can (or cannot) be made with the investment of any finite amount of energy in terms of the level of control complexity and/or time requirements. With this in mind, developing general energy-optimal cooling protocols given finite time and control complexity to achieve any desired finite temperature would further strengthen the connection between these two fundamental laws.

The last future research direction related to thermodynamics that we will mention in this thesis concerns its connection to multi-time phenomena. Most thermodynamic processes are, in some sense, simple: Systems rapidly equilibrate with their environments; dynamics are typically ergodic; and environments do not carry memory. It is clear that such behaviour follows from the standard assumptions of thermodynamics. However, operationally testing the validity of such assumptions often requires one to probe the evolving system at multiple points in time and to try to make inferences on the underlying dynamical behaviour from the multi-time observations. We saw that the connection between Markovianity of statistics and the memory structure of the underlying process is—in contrast to the classical setting—rather complicated in the quantum realm. Our results to this end strengthen the notion of memory as an emergent phenomenon, i.e., as a fundamentally multi-time property that cannot be readily determined from sub-statistics.

On the other hand, although almost all processes carry memory, in our typical experience, we mostly do not see it. This raises a natural question: *To what extent is underlying complexity necessary for other emergent concepts, for example, those in quantum thermodynamics?* Many dynamical thermodynamic phenomena such as equilibration, thermalisation, and ergodicity, look simple at the macroscopic level in spite of all the possible underlying complexity. Even more so, they seem to *require* a sufficient level of complexity: For instance, simple—e.g., non-interacting or integrable—systems do not

tend to thermalise (loosely speaking, because they have too many constants of motion). However, such connections between complexity on the one hand and the simplicity of observed phenomena on the other are yet to be fully and rigorously understood, especially in the multi-time setting with general measurements. For instance, amongst the various notions of thermalisation that abound [101], some make explicit assumptions on the kinds of observables being measured whereas others do not; the former fall more in line with what we have called operational considerations and the latter with structural properties. As we have seen, in quantum theory such approaches are generally not equivalent, and clearly demonstrating a fundamental incompatibility between them for such thermodynamic concepts would drive a definitive wedge between classical and quantum thermodynamics. In this context, note the work in Refs. [135, 214–218, 351–353] that set a preliminary foundation for addressing multi-time phenomena in quantum thermodynamics from an operational standpoint.

Continuing with understanding the distinction between classical and quantum processes, our work in the second part of this dissertation naturally begs the question: *How does operational classicality (in terms of measurement non-invasiveness) relate to other structural notions of classicality that are prevalent throughout the community?* For instance, how is it related to non-negativity of phase-space distributions? Or uncertainty relations? Or broadcastability criteria? Furthermore, having derived the existence of genuinely quantum processes (i.e., those that cannot hide their non-classicality), a natural follow-up study would characterise such processes and detail their properties. A good starting point would be to develop generic witnesses for such genuine non-classicality.

Lastly, our work on the relationship between Markovianity and memorylessness in the quantum realm poses the following concrete open question: *Given a Markovian statistics, deduced from probing a quantum process with sharp projective measurements, such that all sub-statistics are Markovian and compatible, does there always exist a dynamical description in terms of a memoryless quantum process?* An answer in the negative would provide extremely strong evidence for not just a subjective but rather a fundamental difference regarding complexity (in this case, memory effects) as measured on an observational level and that concerning the underlying physical resources. Furthermore, we have seen throughout this thesis both the contrasts between classical and quantum processes on the one hand, and between processes with and without memory (in both settings) on the other; an interesting research direction would focus on the overlap: *What can be achieved by quantum processes that only have classical memory?* To this end, note the recent work of Refs. [354, 355] that begin such an exploration. Understanding such processes would have profound implications for near-term quantum devices, as—in contrast to quantum memory—classical memory can be reliably stored, processes, and read-out without any measurement “back-action” and is therefore more practical for controlling quantum dynamics in “real-world” applications. Thus, aside from being of

theoretical interest in its own right, the development of a framework for describing such processes is imperative for building reliable and efficient quantum devices and gauging their information processing power.

In summary, the above questions all pose promising routes for future research that apply across a wide array of areas: From practical implementations of quantum computers with minimal control, to the very foundations of thermodynamics; from mathematical concepts at the heart of quantum theory, to their manifestation on an operational level (and back again); from a theoretic description of open quantum processes, to sophisticated quantum information processing across both space and time. Accordingly, we believe that the fundamental insights and results presented in this dissertation will serve as an important springboard for future endeavours.



Appendices

Supplemental Information for Chapter 1

A.1 Equality Forms of the (Carnot-)Landauer Limit

In this section, we present lower bounds on the energy change of the machine (or heat dissipated into its environment) in terms of the entropy change of the target system, both in the coherent and incoherent-control settings outlined in the main text. In the coherent setting, this amounts to the well-known Landauer’s principle [3], whereas the incoherent setting requires an extension of this derivation. These lower bounds are important, because they put limits on the optimal energetic performance of the machines for cooling. Note, finally, that the initial state of the machine is diagonal in its energy eigenbasis and must remain so for any process saturating the (Carnot-)Landauer limit; moreover, the target begins similarly and ends up in the pure state $|0\rangle\langle 0|$ when perfect cooling is achieved. As a result, all quantities relevant to perfect cooling at the (Carnot-)Landauer limit can be computed in terms of their “classical” counterparts, i.e., $\varrho_{\mathcal{X}} \rightarrow p_{\mathcal{X}} := (p_0, \dots, p_d)$ with $p_n = e^{-\beta E_n}$, $\text{tr}[H\varrho_{\mathcal{X}}] \rightarrow \langle E \rangle_{p_{\mathcal{X}}} := \sum_n p_n E_n$, $S(\varrho_{\mathcal{X}}) \rightarrow S(p_{\mathcal{X}}) := -\sum_n p_n \log(p_n)$, $\mathcal{Z}(\beta, H_{\mathcal{X}}) = \sum_n e^{-\beta E_n}$, and so on. Nonetheless, all of the results presented hold for the more general “quantum” properties.

A.1.1 Coherent-Control Paradigm: The Landauer Limit

The coherent setting was already studied in detail in Ref. [200], where the authors derived an equality version of Landauer’s principle. We restate the results here for convenience, since we will also use them in the incoherent paradigm. Recall that the setting we consider consists of two parts, the target system \mathcal{S} and the machine \mathcal{M} . In the beginning, the joint state is $\varrho_{\mathcal{SM}} = \varrho_{\mathcal{S}} \otimes \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})$ for some arbitrary (but fixed) Hamiltonian $H_{\mathcal{M}}$ and $\beta \in \mathbb{R}$. Note that any full-rank state ϱ can be associated to some

chosen temperature β , which sets the energy scale, and a Hamiltonian $H = -\frac{1}{\beta} \log(\varrho)$; as we consider arbitrary Hamiltonians, we only write the state dependence on these parameters when necessary. If the state is not full rank, the rank can be used to re-define the dimension. We assume that both systems are finite-dimensional. Let U be a global unitary on \mathcal{SM} . We write $\varrho'_{\mathcal{SM}} := U[\varrho_{\mathcal{S}} \otimes \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})]U^\dagger$ and denote by ϱ'_S and ϱ'_M the respective reduced states. The quantity $I(\mathcal{S} : \mathcal{M})_{\varrho'_{\mathcal{SM}}} = S(\varrho'_S) + S(\varrho'_M) - S(\varrho'_{\mathcal{SM}})$ is the final mutual information between \mathcal{S} and \mathcal{M} and $D(\varrho'_M || \varrho_M) = \text{tr}[\varrho'_M \log(\varrho'_M)] - \text{tr}[\varrho'_M \log(\varrho_M)]$ is the relative entropy of the final machine state with respect to its initial state.

Lemma A.1 ([200, Lemma 2]). *Let the setting be as above. Then*

$$[S(\varrho'_S) - S(\varrho_S)] + [S(\varrho'_M) - S(\varrho_M)] = I(\mathcal{S} : \mathcal{M})_{\varrho'_{\mathcal{SM}}} \geq 0. \quad (\text{A.1})$$

Proof. We note that

$$[S(\varrho'_S) - S(\varrho_S)] + [S(\varrho'_M) - S(\varrho_M)] = S(\varrho'_S) + S(\varrho'_M) - S(\varrho'_{\mathcal{SM}}), \quad (\text{A.2})$$

since the von Neumann entropy is additive for product states and invariant under unitary evolution. The assertion follows from the definition of the mutual information and the fact that it is non-negative. \square

Theorem A.1 (Equality form of Landauer's principle, [200, Theorem 3]). *Let the setting be as above. Then*

$$\beta \text{tr}[H_{\mathcal{M}}(\varrho'_M - \varrho_M)] - [S(\varrho_S) - S(\varrho'_S)] = I(\mathcal{S} : \mathcal{M})_{\varrho'_{\mathcal{SM}}} + D(\varrho'_M || \varrho_M) \geq 0. \quad (\text{A.3})$$

Proof. From Lemma A.1, it follows that

$$[S(\varrho_S) - S(\varrho'_S)] + I(\mathcal{S} : \mathcal{M})_{\varrho'_{\mathcal{SM}}} = S(\varrho'_M) - S(\varrho_M). \quad (\text{A.4})$$

Using the fact that $\varrho_M = \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})$, we infer that $D(\varrho'_M || \varrho_M) = -S(\varrho'_M) + \beta \text{tr}[H_{\mathcal{M}} \varrho'_M] + \log[\text{tr}(e^{-\beta H_{\mathcal{M}}})]$ and $S(\varrho_M) = \beta \text{tr}[H_{\mathcal{M}} \varrho_M] + \log[\text{tr}(e^{-\beta H_{\mathcal{M}}})]$. Re-expressing the first of these for $S(\varrho'_M)$ and inserting both yields Eq. (A.4). The inequality results from non-negativity of relative entropy and mutual information. This completes the proof. \square

A.1.2 Incoherent-Control Paradigm: The Carnot-Landauer Limit

Landauer's principle provides a relationship between how much heat must necessarily be dissipated into the thermal background environment upon manipulating the entropy of a given quantum system. Until now, we have assumed that the system of interest can interact arbitrarily with its environment (i.e., the machine); in other words, we have considered general joint unitary interactions between system and machine, without restriction. In doing so, we have tacitly assumed the ability to draw energy from some external resource (i.e., a work source) in order to implement said unitaries, which are

in general not energy preserving. The particularities of such a resource are left as an abstraction. However, from a thermodynamicists' perspective, this setting may seem somewhat unsatisfactory, as the joint target-machine system is not energetically closed. In order to provide a more self-contained picture of the cooling procedure, one can explicitly include the energy resource, modelled as a quantum system itself, into the setting.

To this end, note first that said resource must be out of thermal equilibrium with respect to the target and machine in order to perform any meaningful thermodynamic transformation. Furthermore, it is sensible to assume that the energy resource system is in thermal equilibrium with its own environment to begin with. The joint target-machine-resource system is then considered to be energetically closed; as such, global unitaries in this setting are restricted to be energy conserving. In order to act as a resource for cooling the target in this picture, the energy source here must begin in equilibrium with a heat bath that is hotter than the initial temperature of the machine (assuming that the machine and resource both begin in thermal states), such that a natural heat flow is induced that leads the environment of the machine to act as a final heat sink. This setting is what we call the incoherent-control scenario. In this context, Landauer's principle translates to studying the relationship between the heat that is necessarily dissipated into the machine's environment upon manipulating the entropy of the target system. Finally, note that the relationship between the coherent and the incoherent-control paradigms is interesting in itself: While on the one hand the incoherent setting includes an additional system and therefore increases the dimensionality of the overall joint system, on the other hand by restricting the transformations on this larger space to be energy conserving, one limits the orbit of attainable states.

Now let us consider the incoherent-control setting. Here, we have the target system \mathcal{S} and the machine comprises of one part \mathcal{C} coupled to the cold bath and another \mathcal{H} coupled to the hot bath. We assume that all systems are finite-dimensional. Every subsystem \mathcal{A} is associated to a Hamiltonian $H_{\mathcal{A}}$ and \mathcal{C} , \mathcal{H} are initially in a thermal state; the cold bath has inverse temperature β and the hot bath has inverse temperature $\beta_H < \beta$. We assume β, β_H . Thus, the initial joint state is $\varrho_{\mathcal{SCH}} = \varrho_{\mathcal{S}} \otimes \tau_{\mathcal{C}}(\beta, H_{\mathcal{C}}) \otimes \tau_{\mathcal{H}}(\beta_H, H_{\mathcal{H}})$. The global evolution on \mathcal{SCH} is implemented via a unitary U , leading to $\varrho'_{\mathcal{SCH}} = U(\varrho_{\mathcal{SCH}})U^\dagger$. We further assume that the unitary evolution on the joint system is energy conserving, i.e., $[U, H_{\mathcal{S}} + H_{\mathcal{C}} + H_{\mathcal{H}}] = 0$. We write $\Delta S_{\mathcal{A}} := S(\varrho'_{\mathcal{A}}) - S(\varrho_{\mathcal{A}})$ for the entropy change on subsystem \mathcal{A} and $\Delta E_{\mathcal{A}} := \text{tr}[H_{\mathcal{A}}(\varrho'_{\mathcal{A}} - \varrho_{\mathcal{A}})]$ for the average energy change. Moreover, the free energy of a state $\varrho_{\mathcal{A}}$ with respect to the inverse temperature β is $F_{\beta}(\varrho_{\mathcal{A}}) = \text{tr}[H_{\mathcal{A}}\varrho_{\mathcal{A}}] - \beta^{-1}S(\varrho_{\mathcal{A}})$.

In the incoherent setting, it makes sense to look at the energy decrease in the hot bath \mathcal{H} , since the hot bath can be seen as the energetic resource one must to expend in order to cool the system \mathcal{S} (alternatively, as we present after the following theorem, one can consider the energy dissipated into the cold bath \mathcal{C} , which serves as the heat sink).

Theorem A.2. *In the above setting, it holds that*

$$\Delta F_S^{(\beta)} + \eta \Delta E_{\mathcal{H}} = -\frac{1}{\beta} [\Delta S_S + \Delta S_C + \Delta S_{\mathcal{H}} + D(\varrho'_C || \varrho_C) + D(\varrho'_{\mathcal{H}} || \varrho_{\mathcal{H}})] \leq 0, \quad (\text{A.5})$$

where $(0, 1) \ni \eta := 1 - \frac{\beta_H}{\beta}$ is the Carnot efficiency and $\Delta F_S^{(\beta)} = F_{\beta}(\varrho'_S) - F_{\beta}(\varrho_S)$.

Proof. Let us consider

$$I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{SC\mathcal{H}}} := S(\varrho'_S) + S(\varrho'_C) + S(\varrho'_{\mathcal{H}}) - S(\varrho'_{SC\mathcal{H}}) \geq 0. \quad (\text{A.6})$$

Note that the quantity $I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{SC\mathcal{H}}}$, which quantifies the tripartite mutual information of the state $\varrho'_{SC\mathcal{H}}$, is non-negative via subadditivity $S(\varrho_A) + S(\varrho_B) \geq S(\varrho_{AB})$ for any state ϱ_{AB} . Furthermore, since the von Neumann entropy is invariant under unitary transformations and additive for tensor product states, we have

$$I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{SC\mathcal{H}}} = \Delta S_S + \Delta S_C + \Delta S_{\mathcal{H}}. \quad (\text{A.7})$$

We also have that

$$\Delta S_C = \beta \Delta E_C - D(\varrho'_C || \varrho_C) \quad (\text{A.8})$$

and

$$\Delta S_{\mathcal{H}} = \beta_H \Delta E_{\mathcal{H}} - D(\varrho'_{\mathcal{H}} || \varrho_{\mathcal{H}}). \quad (\text{A.9})$$

Thus,

$$I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{SC\mathcal{H}}} = \Delta S_S + \beta \Delta E_C - D(\varrho'_C || \varrho_C) + \beta_H \Delta E_{\mathcal{H}} - D(\varrho'_{\mathcal{H}} || \varrho_{\mathcal{H}}). \quad (\text{A.10})$$

Since the unitary is energy conserving, we infer that $\Delta E_S + \Delta E_C + \Delta E_{\mathcal{H}} = 0$. Hence, we have

$$\Delta S_S - \beta \Delta E_S + (\beta_H - \beta) \Delta E_{\mathcal{H}} = I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{SC\mathcal{H}}} + D(\varrho'_C || \varrho_C) + D(\varrho'_{\mathcal{H}} || \varrho_{\mathcal{H}}). \quad (\text{A.11})$$

Using the free energy, we can rewrite this as

$$-\beta [F_{\beta}(\varrho'_S) - F_{\beta}(\varrho_S)] - (\beta - \beta_H) \Delta E_{\mathcal{H}} = I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{SC\mathcal{H}}} + D(\varrho'_C || \varrho_C) + D(\varrho'_{\mathcal{H}} || \varrho_{\mathcal{H}}). \quad (\text{A.12})$$

Dividing by $-\beta$, we obtain the assertion, since, in particular, $I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{SC\mathcal{H}}} + D(\varrho'_C || \varrho_C) + D(\varrho'_{\mathcal{H}} || \varrho_{\mathcal{H}}) \geq 0$ by the non-negativity of each term. \square

In particular, we have shown that the energy extracted from the hot bath is lower-bounded by the increase in free energy, weighted by the inverse Carnot efficiency:

$$\text{tr}[H_{\mathcal{H}}(\varrho_{\mathcal{H}} - \varrho'_{\mathcal{H}})] \geq \frac{1}{\eta} [F_{\beta}(\varrho'_S) - F_{\beta}(\varrho_S)]. \quad (\text{A.13})$$

Note that if $\varrho_S = \tau_S(\beta, H_S)$, the r.h.s. is non-negative for any nontrivial thermodynamic process, i.e., any for which the target system is heated or—of particular relevance for

us—cooled. This follows by the Gibbs variational principle which states that the free energy of ϱ is minimal if and only if ϱ is the Gibbs state.

Finally, in order to make a more concrete connection to the spirit of Landauer’s original derivation, note that one can consider bounding the heat dissipated into the cold bath, rather than that drawn from the hot bath. Substituting $\Delta E_{\mathcal{H}} = -(\Delta E_S + \Delta E_C)$ into Eq. (A.10) leads to

$$-\tilde{\Delta}S_S - \beta_H \Delta E_S + (\beta - \beta_H) \Delta E_C \geq 0, \quad (\text{A.14})$$

which recovers the standard Landauer bound for the dissipated heat in the limit of an infinitely-hot heat bath, i.e., $\beta_H \rightarrow 0$.

A.2 Diverging Energy

A.2.1 Sufficiency: Diverging Energy Cooling Protocol

This cooling protocol is arguably the simplest of those presented. The thermal populations of any target system can be exchanged with a machine system of the same dimension, in the thermal state of $H_{\mathcal{M}} = \omega_{\mathcal{M}} \sum_{n=0}^{d-1} n |n\rangle\langle n|$. As $\omega_{\mathcal{M}} \rightarrow \infty$, the machine state $\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})$ approaches $|0\rangle\langle 0|_{\mathcal{M}}$ independently of β (as long as $\beta \neq 0$). Such a population-exchange operation is a single interaction (i.e., the protocol occurs in unit time) which is of finite complexity (in a sense that we discuss below). However, the energy drawn from the resource \mathcal{W} upon performing said swap is at least $E = (p_S^{(1)} - p_{\mathcal{M}}^{(1)})(\omega_{\mathcal{M}} - \omega_S^{(1)})$, where $p_{\mathcal{X}}^{(1)}$ is the initial population of the first excited level of system \mathcal{X} and $\omega_S^{(1)}$ is the first energy eigenvalue of the target system. Denoting by $\omega_S^{(k)}$ the energy eigenvalue of the k^{th} excited level of the target system, we have above assumed that $\omega_S^{(0)} = 0$ (which we do for all Hamiltonians without loss of generality) and $\omega_{\mathcal{M}} > \omega_S^{(d-1)}$. As such, perfect cooling will incur diverging energy cost.

A.2.2 Necessity of Diverging Energy for Protocols with Finite Time and Control Complexity

Consider the following Hamiltonians for the target system and machine with finite but otherwise arbitrary energy levels, $H_S = \sum_{n=0}^{d_S-1} \omega_S^{(n)} |n\rangle\langle n|_S$ and $H_{\mathcal{M}} = \sum_{n=0}^{d_{\mathcal{M}}-1} \omega_{\mathcal{M}}^{(n)} |n\rangle\langle n|_{\mathcal{M}}$, respectively. For any finite inverse temperature β , the initial thermal states $\tau_S(\beta, H_S)$ and $\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})$ are of full rank. Suppose now that one can implement a single unitary transformation (i.e., a unit time protocol) of finite control complexity on the joint target and machine, yielding the joint output state $\varrho'_{S\mathcal{M}} = \text{tr}_{\mathcal{M}} [U(\tau_S(\beta, H_S) \otimes \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}}))U^\dagger]$, and wishes to attain perfect cooling of the target in doing so. By invariance of the rank under unitary transformations and the fact that the system and machine begin

uncorrelated, we have

$$\text{rank}[\tau_{\mathcal{S}}(\beta, H_{\mathcal{S}})] \text{rank}[\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})] = \text{rank}[\varrho'_{\mathcal{SM}}] \leq \text{rank}[\varrho'_{\mathcal{S}}] \text{rank}[\varrho'_{\mathcal{M}}], \quad (\text{A.15})$$

where the inequality follows from the subadditivity of the Rényi-zero entropy [356], which is the logarithm of the rank. To achieve perfect cooling of the target, one must (at least asymptotically) attain $\text{rank}[\varrho'_{\mathcal{S}}] < \text{rank}[\tau_{\mathcal{S}}(\beta, H_{\mathcal{S}})]$, which implies that $\text{rank}[\varrho'_{\mathcal{M}}] > \text{rank}[\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})]$. However, if this condition is achieved, then $D[\varrho'_{\mathcal{M}} \parallel \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})]$ diverges, implying a diverging energy cost by Eq. (1.8). The above argument already appears in Ref. [200].

The other situation that one must consider is the case where one attains a $\varrho'_{\mathcal{S}}$ such that $\text{rank}[\varrho'_{\mathcal{S}}] = \text{rank}[\tau_{\mathcal{S}}(\beta, H_{\mathcal{S}})]$ but nonetheless $\varrho'_{\mathcal{S}}$ is arbitrarily close to a pure state, as is the case, for instance, in the protocols that we present. Consider a sequence of machines $\varrho_{\mathcal{M}}^{(i)}$ and unitaries $U^{(i)}$ such that $\varrho_{\mathcal{M}}^{(i)} \rightarrow \varrho_{\mathcal{M}}$ and $U^{(i)} \rightarrow U$. Note that since we fixed the dimensions of \mathcal{S} and \mathcal{M} , any sequence of machines has a converging subsequence by the Bolzano-Weierstrass theorem and the fact that the set of quantum states is compact. Here, $\varrho_{\mathcal{M}}$ and U achieve perfect cooling. If we fix $\varrho_{\mathcal{S}}$, we obtain a corresponding sequence $(\varrho'_{\mathcal{M}})^{(i)}$ such that $(\varrho'_{\mathcal{M}})^{(i)} \rightarrow \varrho'_{\mathcal{M}}$. Crucially, here, since we restrict the unitary transformation to be of finite control complexity, the states $\varrho_{\mathcal{M}}$ and $\varrho'_{\mathcal{M}}$ are effectively finite-dimensional, in the sense that whatever their true dimension, they can be replaced by finite-dimensional versions without changing any of the relevant quantities (see Appendix A.4). Since the relative entropy $(\varrho, \sigma) \mapsto D(\varrho \parallel \sigma)$ is lower semicontinuous [357, 358] and since $D(\varrho'_{\mathcal{M}} \parallel \varrho_{\mathcal{M}}) \rightarrow \infty$ by the arguments above, we infer that $D[(\varrho'_{\mathcal{M}})^{(i)} \parallel \varrho_{\mathcal{M}}^{(i)}] \rightarrow \infty$ as $i \rightarrow \infty$. This argument holds independently of $\text{rank}[\varrho'_{\mathcal{S}}]$; in particular, for the special case $\text{rank}[\varrho'_{\mathcal{S}}] = \text{rank}[\tau_{\mathcal{S}}(\beta, H_{\mathcal{S}})]$ that we are considering here. Thus, to approach perfect cooling in finite time and with finite control complexity, one would need a diverging energy cost. Thus, we see that within the resource trinity of energy, time and control complexity, if the latter two are finite, then energy must diverge to asymptotically achieve a pure state. Whether or not there exist other (unaccounted for) resources that allow one to achieve this with all three of the aforementioned resources being finite remains an open question.

Importantly, the above argument no longer holds if the time or control complexity is allowed to diverge. In such cases, both $\varrho_{\mathcal{M}}$ and $\varrho'_{\mathcal{M}}$ can be infinite-dimensional, and because of this the rank argument no longer applies and the relative entropy does not necessarily diverge in the limit of perfect cooling. On the contrary, as we show, it is even possible to saturate the Landauer bound.

A.3 Diverging Time Cooling Protocol for Finite-Dimensional Systems

A.3.1 Proof of Theorem 1.2

Proof. Consider a target system \mathcal{S} of dimension d with associated Hamiltonian

$$H_{\mathcal{S}} = \sum_{k=0}^{d-1} \omega_k |k\rangle\langle k|_{\mathcal{S}}, \quad (\text{A.16})$$

where we also set $\omega_0 = 0$ without loss of generality. Consider also the machine \mathcal{M} to be composed of N subsystems, $\{\mathcal{M}_n\}_{n=1,\dots,N}$, each of the same dimension d as the target, whose local Hamiltonians are

$$H_{\mathcal{M}}^{(n)} = (1 + n\epsilon)H_{\mathcal{S}}, \quad (\text{A.17})$$

where $\epsilon = \frac{\beta_{\max} - \beta}{N\beta}$. We first cool the system initially at non-zero β to some fixed, finite β_{\max} , which we will eventually take $\beta_{\max} \rightarrow \infty$ in order to asymptotically achieve perfect cooling. We treat the case $\beta = 0$ as a limiting case of $\beta \rightarrow 0$: Here, as $\beta \rightarrow 0$, we let $N \rightarrow \infty$ such that $N\beta \rightarrow \infty$, e.g., we specify a suitable function $N(\beta)$ such that $N(\beta) \rightarrow \infty$ “faster” than $\beta \rightarrow 0$.

We now show that, given the ability to perform a diverging number of operations on such a configuration, one can reach the target state $\tau_{\mathcal{S}}(\beta_{\max}, H_{\mathcal{S}})$. In particular, we show that the protocol presented uses the minimal amount of energy to do so, and explicitly calculate this to be $\beta^{-1}\tilde{\Delta}S$ units of energy, where $\tilde{\Delta}S := S[\tau_{\mathcal{S}}(\beta, H_{\mathcal{S}})] - S[\tau_{\mathcal{S}}(\beta_{\max}, H_{\mathcal{S}})]$. In other words, as the number of operations in the protocol diverges, we approach perfect cooling at the Landauer limit, thereby saturating the ultimate bound.

The diverging time cooling protocol is as follows. At each step, the target system interacts with a single machine labelled by n via the swap operator

$$\mathbb{S}_{\mathcal{S}\mathcal{M}_n}^d := \sum_{i,j=0}^{d-1} |i, j\rangle\langle j, i|_{\mathcal{S}\mathcal{M}_n}. \quad (\text{A.18})$$

As the target and machine subsystems considered here are of the same dimension, we will drop the subscript on the states associated to each subsystem, for ease of notation. Such a transformation is, in general, not energy conserving, but one can calculate the energy change for both the target system and the machine due to the n^{th} interaction as

$$\Delta E_{\mathcal{S}}^{(n)} = \text{tr} \left[H_{\mathcal{S}} \tau(\beta, H_{\mathcal{M}}^{(n)}) \right] - \text{tr} \left[H_{\mathcal{S}} \tau(\beta, H_{\mathcal{M}}^{(n-1)}) \right], \quad (\text{A.19})$$

and so the total energy change of the system over the entire N -step protocol is given by

$$\Delta E_{\mathcal{S}} = \sum_{n=1}^N \Delta E_{\mathcal{S}}^{(n)} = \text{tr} \left[H_{\mathcal{S}} \tau(\beta, H_{\mathcal{M}}^{(N)}) \right] - \text{tr} \left[H_{\mathcal{S}} \tau(\beta, H_{\mathcal{M}}^{(0)}) \right]. \quad (\text{A.20})$$

The energy change of the machine subsystem that is swapped with the target system at each step is given by

$$\begin{aligned}\Delta E_{\mathcal{M}}^{(n)} &= \text{tr} \left[H_{\mathcal{M}}^{(n)} \tau(\beta, H_{\mathcal{M}}^{(n-1)}) \right] - \text{tr} \left[H_{\mathcal{M}}^{(n)} \tau(\beta, H_{\mathcal{M}}^{(n)}) \right] \\ &= \sum_{k=0}^{d-1} (1 + n\epsilon) \omega_k \left[p_k(\beta, H_{\mathcal{M}}^{(n-1)}) - p_k(\beta, H_{\mathcal{M}}^{(n)}) \right],\end{aligned}\quad (\text{A.21})$$

where $p_k(\beta, H_{\mathcal{M}}^{(n)}) = e^{-\beta(1+n\epsilon)\omega_k} / \mathcal{Z}_{\mathcal{M}_n}(\beta, H_{\mathcal{M}}^{(n)})$ is the population in the k^{th} energy level of the thermal state of the n^{th} machine subsystem at inverse temperature β , with $\mathcal{Z}_{\mathcal{M}_n}(\beta, H_{\mathcal{M}}^{(n)}) = \text{tr} \left[e^{-\beta H_{\mathcal{M}}^{(n)}} \right]$ being the partition function.

By summing the contributions of the energy changes in each step, one can obtain the total energy change for the overall machine throughout the entire process:

$$\Delta E_{\mathcal{M}}^{(N)} = \sum_{n=1}^N \Delta E_{\mathcal{M}}^{(n)} = \sum_{n=1}^N \sum_{k=0}^{d-1} (1 + n\epsilon) \omega_k \left[p_k(\beta, H_{\mathcal{M}}^{(n-1)}) - p_k(\beta, H_{\mathcal{M}}^{(n)}) \right], \quad (\text{A.22})$$

In general, it is complicated to calculate the energy cost for the protocol up until a finite time step N , since this depends on the full energy structure of the target system and machine subsystems involved (we will return to resolve this problem for the special case of equally-spaced system and machine Hamiltonians in the coming section). Here, we focus on a special case in which $N \rightarrow \infty$, i.e., there is a diverging number of machine subsystems that the target system interacts with throughout the protocol. This limit physically corresponds to that of requiring a diverging amount of time (in terms of the number of steps). Furthermore, we take the limit $\epsilon \rightarrow 0$ for any fixed $\beta, \beta_{\text{max}}$. Considering the differentials

$$\Delta p_k^{(n)} := p_k(\beta, H_{\mathcal{M}}^{(n)}) - p_k(\beta, H_{\mathcal{M}}^{(n-1)}), \quad (\text{A.23})$$

and

$$\Delta x_n := x_n - x_{n-1} \quad \text{with} \quad x_n := 1 + n\epsilon. \quad (\text{A.24})$$

In order for x_n to become infinitesimal, and noting the explicit form of the machine subsystem Hamiltonians $H_{\mathcal{M}}^{(n)} = (1 + n\epsilon)H_S$, we can make the replacement

$$-\frac{\Delta p_k^{(n)}}{\Delta x_n} \Delta x_n \rightarrow -\frac{\partial p_k(\beta, xH_S)}{\partial x} dx, \quad (\text{A.25})$$

where $x := 1 + n\epsilon$ has become a continuous parameter. This way we can express the limit $N \rightarrow \infty$ of Eq. (A.22) as a Riemann integral in the following form

$$\lim_{N \rightarrow \infty} \Delta E_{\mathcal{M}}^{(N)} = - \int_1^{x_{\text{max}}} \sum_{k=0}^{d-1} x \omega_k \frac{\partial p_k(\beta, xH_S)}{\partial x} dx, \quad (\text{A.26})$$

where $x_{\max} := \frac{\beta_{\max}}{\beta}$. Both the summation and the integral converge, so one can swap the order of their evaluation. Integrating by parts then gives

$$\begin{aligned}
 \lim_{N \rightarrow \infty} \Delta E_{\mathcal{M}}^{(N)} &= \sum_{k=0}^{d-1} \left[-x \omega_k p_k(\beta, x H_S) \Big|_1^{x_{\max}} + \int_1^{x_{\max}} \omega_k p_k(\beta, x H_S) dx \right] \\
 &= \sum_{k=0}^{d-1} \left[-x \omega_k p_k(\beta, x H_S) \Big|_1^{x_{\max}} \right] - \int_1^{x_{\max}} \frac{1}{\beta} \frac{\partial}{\partial x} \left[\log \mathcal{Z}(\beta, x H_S) \right] dx \\
 &= E[\tau(\beta, H_S)] - E[\tau(\beta, x_{\max} H_S)] - \frac{1}{\beta} \log \mathcal{Z}(\beta, x_{\max} H_S) + \frac{1}{\beta} \log \mathcal{Z}(\beta, H_S),
 \end{aligned} \tag{A.27}$$

where in the second line we again swapped the order of the integral and the sum to write $\sum_{k=0}^{d-1} \omega_k p_k(\beta, x H_S) = -\frac{1}{\beta} \frac{\partial}{\partial x} [\log \mathcal{Z}(\beta, x H_S)]$ and in the last line we invoke $E[\tau(\beta, x H)] = \text{tr}[x H \tau(\beta, x H)]$. Finally, writing the partition function in terms of the average energy and entropy, i.e., $\log[\mathcal{Z}(\beta, x H)] = -\beta E[\tau(\beta, x H)] + S[\tau(\beta, x H)]$, the total energy change of the machine is given by

$$\begin{aligned}
 \lim_{N \rightarrow \infty} \Delta E_{\mathcal{M}}^{(N)} &= E[\tau(\beta, H_S)] - E[\tau(\beta, x_{\max} H_S)] + E[\tau(\beta, x_{\max} H_S)] \\
 &\quad - \frac{1}{\beta} S[\tau(\beta, x_{\max} H)] - E[\tau(\beta, H_S)] + \frac{1}{\beta} S[\tau(\beta, H_S)] \\
 &= \frac{1}{\beta} \left\{ S[\tau(\beta, H_S)] - S[\tau(\beta_{\max}, H_S)] \right\} = \frac{1}{\beta} \tilde{\Delta} S_S,
 \end{aligned} \tag{A.28}$$

where we have made use of the property $\tau_S(\beta, x_{\max} H_S) = \tau_S(\beta_{\max}, H_S)$ and the entropy decrease of the target system corresponds to that associated with the transformation $\tau(\beta, H_S) \rightarrow \tau(\beta_{\max}, H_S)$. Thus, as the number of timesteps diverges, this cooling process saturates the Landauer limit for the heat dissipated by the machine. In order to achieve perfect cooling at the Landauer limit, i.e., the final target state to approach $|0\rangle\langle 0|$ and thus prove Theorem 1.2, we can now take the limit $\beta_{\max} \rightarrow \infty$. \square

The above proof holds for systems and machines of arbitrary (but equal) dimension, either finite or infinite, with arbitrary Hamiltonians. We now present some more detailed analysis regarding the special case where the Hamiltonians of the target system and all machine subsystems are equally spaced; this provides an opportunity both to derive a more detailed formula for the energy costs involved and to build intuition regarding some of the important differences between the finite- and infinite-dimensional settings.

A.3.2 Special Case: Equally Spaced Hamiltonians

Consider a finite d -dimensional target system beginning at inverse temperature β with an equally spaced Hamiltonian $H_S(\omega_S) = \omega_S \sum_{n=0}^{d-1} n |n\rangle\langle n|_S$. In this case, we can derive a more precise dimension-dependant function for the energy cost dissipated by the machines throughout the optimal cooling protocol presented above.

Consider an initial target system $\tau_S(\beta, H_S)$ and a diverging number N of machines $\{\mathcal{M}_\alpha\}_{\alpha=0,\dots,N}$ of the same dimension d as the target, which all begin in a thermal state at inverse temperature β with respect to an equally spaced Hamiltonian whose gaps between neighbouring energy levels $\omega_{\mathcal{M}_\alpha}$ are ordered non-decreasingly. Each machine is used once and then discarded; the particular interaction is the aforementioned swap between the target system and the n^{th} qudit machine, i.e., that represented by the unitary $\mathbb{S}_{S\mathcal{M}_\alpha}^d := \sum_{i,j=0}^{d-1} |i, j\rangle\langle j, i|_{S\mathcal{M}_\alpha}$. After applying such an operation, the state of the target system is given by

$$\tau_S(\beta, \omega_\alpha) := \frac{e^{-\beta H_S(\omega_\alpha)}}{\mathcal{Z}_S(\beta, \omega_\alpha)}, \quad (\text{A.29})$$

where $H_S(\omega_\alpha) := \omega_\alpha \sum_{n=0}^{d-1} n |n\rangle\langle n|_S$ and $\mathcal{Z}_S(\beta, \omega_\alpha) := \text{tr} [e^{-\beta H_S(\omega_\alpha)}]$.

We now calculate the energy cost explicitly for the diverging time cooling protocol, which saturates the Landauer bound in the asymptotic limit. In order to minimise the energy cost of cooling, the target system must be cooled by the qudit system in the machines with the smallest gap between neighbouring energy levels (that permits cooling) as much as possible at each stage. In order to optimally use the given machine structure at hand, we thus order the set of energy gaps ω_α in non-decreasing order. In addition, the protocol to reach the Landauer erasure bound, i.e., minimal energy cost, dictates that one must infinitesimally increase ω_α of the machines in order to dissipate as little heat as possible throughout the interactions. Since we are here considering a diverging time limit, we have access to a diverging number of qudit machine with distinct energy gap ω_α at our disposal; the task is then to use these in an energy-optimal manner.

It is straightforward to see that to minimise the total energy cost, one must apply the sequence of unitaries $\mathbb{S}_{S\mathcal{M}_\alpha}^d$ such that $\mathbb{S}_{S\mathcal{M}_0}^d$ is first applied to reach the optimally cool $\tau_S(\beta, \omega_0)$, then $\mathbb{S}_{S\mathcal{M}_1}^d$ to reach $\tau_S(\beta, \omega_1)$, and so on. The heat dissipated by the reset machines in each stage of such a cooling protocol (i.e., for each value of α) can thus be calculated as

$$\begin{aligned} \Delta E_{\mathcal{M}_\alpha}(\omega_\alpha) &= - \left\{ \text{tr} [H_{\mathcal{M}_\alpha}(\omega_\alpha) \tau_{\mathcal{M}_\alpha}(\beta, \omega_\alpha)] + \text{tr} [H_{\mathcal{M}_\alpha}(\omega_\alpha) \tau_{\mathcal{M}_\alpha}(\beta, \omega_{\alpha-1})] \right\} \\ &= -\text{tr} [H_S(\omega_\alpha) [\tau_S(\beta, \omega_\alpha) - \tau_S(\beta, \omega_{\alpha-1})]]. \end{aligned} \quad (\text{A.30})$$

In the second line, we have made use of the fact that the Hamiltonians of both the target system and each of machine are d -dimensional and equally spaced. So far, we have obtained the energy dissipated by the reset machines. To investigate the total energy cost of cooling in such a process, we also must consider the contribution of energy transferred to the target system \mathcal{S} , which is characterised via its local Hamiltonian H_S and calculated via

$$\Delta E_S(\omega_\alpha) = \text{tr} [H_S(\omega_S) \tau_S(\beta, \omega_\alpha)] - \text{tr} [H_S(\omega_S) \tau_S(\beta, \omega_{\alpha-1})], \quad (\text{A.31})$$

in which we set $\omega_0 = \omega_S$. Using Eqs. (A.30, A.31), the total energy cost for each stage of cooling is given by

$$\begin{aligned}\Delta E_{S\mathcal{M}}(\omega_\alpha) &= \Delta E_S(\omega_\alpha) + \Delta E_{\mathcal{M}}(\omega_\alpha) \\ &= \text{tr} \left\{ \left[H_S(\omega_S) - H_S(\omega_\alpha) \right] \left[\tau_S(\beta, \omega_\alpha) - \tau_S(\beta, \omega_{\alpha-1}) \right] \right\},\end{aligned}\quad (\text{A.32})$$

which leads to the overall energy cost after N stages, where N is the number of non-zero distinct energy gaps of the reset machines, as

$$\Delta E_{S\mathcal{M}}^{(N)} = \sum_{\alpha=1}^N \Delta E_{S\mathcal{M}}(\omega_\alpha) = \sum_{\alpha=1}^N \text{tr} \left\{ \left[H_S(\omega_S) - H_S(\omega_\alpha) \right] \left[\tau_S(\beta, \omega_\alpha) - \tau_S(\beta, \omega_{\alpha-1}) \right] \right\}.\quad (\text{A.33})$$

Now, we can obtain the total energy cost for each stage of the protocol (i.e., each value of α considered) in terms of the transformation of the target system alone. Note that in this protocol, each stage corresponding to each of the N distinct energy gaps $\{\omega_\alpha\}$ in itself requires only one operation to perfectly reach $\tau_S(\beta, \omega_\alpha)$. The end result of this protocol is that the target system is cooled from the initial thermal state $\tau_S(\beta, \omega_S)$, where ω_S is the energy gap between each pair of adjacent energy levels in the system, to $\tau_S(\beta, \omega_{\max})$ in the energy-optimal manner.

Starting from Eq. (A.33), we have

$$\begin{aligned}\Delta E_{S\mathcal{M}}^{(N)} &= \sum_{\alpha=1}^N \text{tr} \left\{ \left[H_S(\omega_S) - H_S(\omega_\alpha) \right] \left[\tau_S(\beta, \omega_\alpha) - \tau_S(\beta, \omega_{\alpha-1}) \right] \right\} \\ &= \sum_{\alpha=1}^N (\omega_S - \omega_\alpha) \left[\left(\frac{e^{-\beta\omega_\alpha}}{1 - e^{-\beta\omega_\alpha}} - \frac{e^{-\beta\omega_{\alpha-1}}}{1 - e^{-\beta\omega_{\alpha-1}}} \right) - \left(\frac{d e^{-\beta d\omega_\alpha}}{1 - e^{-\beta d\omega_\alpha}} - \frac{d e^{-\beta d\omega_{\alpha-1}}}{1 - e^{-\beta d\omega_{\alpha-1}}} \right) \right] \\ &= \sum_{\alpha=1}^N (\omega_S - \omega_\alpha) \times \\ &\quad \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\left(e^{-\beta(k+1)\omega_\alpha} - e^{-\beta(k+1)\omega_{\alpha-1}} - d \left(e^{-\beta(k+1)d\omega_\alpha} - e^{-\beta(k+1)d\omega_{\alpha-1}} \right) \right) \right] \\ &= \sum_{\alpha=1}^N (\omega_S - \omega_\alpha) \times \\ &\quad \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[e^{-\beta(k+1)\omega_\alpha} \left(1 - e^{-\beta(k+1)(\omega_{\alpha-1} - \omega_\alpha)} \right) - d e^{-\beta d(k+1)\omega_\alpha} \left(1 - e^{-\beta d(k+1)(\omega_{\alpha-1} - \omega_\alpha)} \right) \right].\end{aligned}\quad (\text{A.34})$$

Here, since both $H_{\mathcal{M}_\alpha}$ and H_S are equally spaced Hamiltonians, the average energy can be written as

$$\begin{aligned}E(\omega_x, \omega_y) &= \text{tr} [H_S(\omega_x) \tau_S(\beta, \omega_y)] = \frac{\sum_{n=0}^{d-1} n \omega_x e^{-n\beta\omega_y}}{\sum_{n=0}^{d-1} e^{-n\beta\omega_y}} \\ &= \omega_x \left(\frac{e^{-\beta\omega_y}}{1 - e^{-\beta\omega_y}} - \frac{d e^{-\beta d\omega_y}}{1 - e^{-\beta d\omega_y}} \right)\end{aligned}\quad (\text{A.35})$$

by evaluating the geometric series

$$\mathcal{Z}(\beta, \omega_y) = \sum_{n=0}^{d-1} e^{-\beta n \omega_y} = \frac{1 - e^{-\beta d \omega_y}}{1 - e^{-\beta \omega_y}} \quad (\text{A.36})$$

and writing

$$\begin{aligned} E(\omega_x, \omega_y) &= \sum_{n=0}^{d-1} n \omega_x \frac{e^{-\beta n \omega_y}}{\mathcal{Z}(\beta, \omega_y)} \\ &= \frac{\omega_x}{\omega_y} \left\{ -\frac{\partial}{\partial \beta} \log [\mathcal{Z}(\beta, \omega_y)] \right\} \\ &= -\frac{\omega_x}{\omega_y} \frac{\partial}{\partial \beta} \left[\log (1 - e^{-\beta d \omega_y}) - \log (1 - e^{-\beta \omega_y}) \right] \end{aligned} \quad (\text{A.37})$$

as we do in the second line of Eq. (A.34) and then using the infinite series expression $(1-x)^{-1} = \lim_{K \rightarrow \infty} \sum_{k=0}^K x^k$ for any $|x| < 1$ as per the third line.

As we will see in Appendix A.5.2, the energy cost for cooling an infinite-dimensional system when both target and machines have equally spaced Hamiltonians (i.e., harmonic oscillators) is similar to the form of Eq. (A.34). Importantly, the second term in square parenthesis vanishes as $d \rightarrow \infty$, simplifying the expression even further.

We now assume that the energy gaps of the machine are given by $\omega_\alpha = \omega_S + \epsilon \alpha$ and so the total energy cost can be written as follows:

$$\begin{aligned} \Delta E_{S\mathcal{M}}^{(N)} &= \\ &- \lim_{K \rightarrow \infty} \sum_{\alpha=1}^N \alpha \epsilon \sum_{k=0}^K e^{-\beta k (\omega_S + \alpha \epsilon)} (1 - e^{\beta k \epsilon}) + \lim_{K \rightarrow \infty} \sum_{\alpha=1}^N \alpha d \epsilon \sum_{k=0}^K e^{-\beta k d (\omega_S + \alpha \epsilon)} (1 - e^{\beta k d \epsilon}) \\ &= \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[e^{-\beta k \omega_S} (e^{\beta k \epsilon} - 1) \left(\sum_{\alpha=1}^N \alpha \epsilon e^{-\beta k \alpha \epsilon} \right) \right] \\ &- \lim_{K \rightarrow \infty} \sum_{k=0}^K e^{-\beta k d \omega_S} \left[(e^{\beta k d \epsilon} - 1) \left(\sum_{\alpha=1}^N d \alpha \epsilon e^{-\beta k d \alpha \epsilon} \right) \right], \end{aligned} \quad (\text{A.38})$$

where we can swap the order of summation since both sums converge and the summands are non-positive. This can be seen from the first line above, using the fact that $e^{-\alpha x} (1 - e^x) \in [-1, 0]$ for all $\alpha \geq 1$ and $x \geq 0$. We now calculate the sum over α .

$$\begin{aligned} \sum_{\alpha=1}^N \alpha \epsilon e^{-\beta \alpha \epsilon} &= -\frac{\partial}{\partial \beta} \sum_{\alpha=0}^N e^{-\beta \alpha \epsilon} = -\frac{\partial}{\partial \beta} \left(\frac{1 - e^{-\beta (N+1) \epsilon}}{1 - e^{-\beta \epsilon}} \right) \\ &= - \left(\frac{(N+1) \epsilon e^{-\beta (N+1) \epsilon} - (N+1) \epsilon e^{-\beta (N+2) \epsilon} - \epsilon e^{-\beta \epsilon} + \epsilon e^{-\beta (N+2) \epsilon}}{(1 - e^{-\beta \epsilon})^2} \right) \\ &= \frac{\epsilon e^{-\beta \epsilon}}{(1 - e^{-\beta \epsilon})^2} \left(1 - (N+1) e^{-\beta N \epsilon} + N e^{-\beta (N+1) \epsilon} \right) \\ &= \frac{\epsilon e^{-\beta \epsilon}}{(1 - e^{-\beta \epsilon})^2} \left(1 - e^{-\beta N \epsilon} - N e^{-\beta N \epsilon} (1 - e^{-\beta \epsilon}) \right). \end{aligned} \quad (\text{A.39})$$

Combining Eqs. (A.38) and (A.39), we arrive at

$$\begin{aligned} \Delta E_{S\mathcal{M}}^{(N)} &= \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k \omega_S}}{k} \frac{k \epsilon (1 - e^{-\beta N k \epsilon})}{(1 - e^{-\beta k \epsilon})} - N \epsilon e^{-\beta k (\omega_S + N \epsilon)} \right] \\ &\quad - \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k d \omega_S}}{k} \frac{k d \epsilon (1 - e^{-\beta N k d \epsilon})}{(1 - e^{-\beta k d \epsilon})} - N d \epsilon e^{-\beta k d (\omega_S + N \epsilon)} \right]. \end{aligned} \quad (\text{A.40})$$

In order to optimise the energy cost, we now assume that the energy gaps of the machines can be chosen to be smoothly increasing in such way that $\epsilon = \Delta \frac{\omega}{N} := \frac{\omega_{\max} - \omega_S}{N}$. Substituting this expression for ϵ into the above equation yields

$$\begin{aligned} \Delta E_{S\mathcal{M}}^{(N)} &= \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k \omega_S}}{k} \frac{k \Delta \omega (1 - e^{-\beta k \Delta \omega})}{N (1 - e^{-\beta k \frac{\Delta \omega}{N}})} - \Delta \omega e^{-\beta k (\omega_S + \Delta \omega)} \right] \\ &\quad - \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k d \omega_S}}{k} \frac{k d \Delta \omega (1 - e^{-\beta k d \Delta \omega})}{N (1 - e^{-\beta k d \frac{\Delta \omega}{N}})} - d \Delta \omega e^{-\beta k d (\omega_S + \Delta \omega)} \right]. \end{aligned} \quad (\text{A.41})$$

We now wish to take the limit of $N \gg K \rightarrow \infty$. This assumption means that energy change of the system is approximately equal to its free energy change; in other words, the process occurs quasi-adiabatically. The ability to switch the order of taking the limits of K and N going to ∞ follows from the monotonic convergence of the sum over k . In particular, note that the term inside square parentheses in each summand converges and the first term in each summation (which is the only part that depends on N) is positive and bounded.

Under this assumption, we can use the approximation $\lim_{\beta x \rightarrow 0} \frac{x}{1 - e^{-\beta x}} = \frac{1}{\beta}$; since $0 < e^{-\beta x} < 1$ for any positive x , the sum over k converges to a finite value. In general, this approximation introduces a correction term for the energy change, however under said assumption the error incurred becomes negligible. Then, the total energy change $\Delta E_{S\mathcal{M}}^{\text{tot}}$ for the transformation $\tau_S(\beta, \omega_S) \rightarrow \tau_S(\beta, \omega_{\max})$ throughout the overall process is

$$\begin{aligned} \Delta E_{S\mathcal{M}}^{\text{tot}} &= \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k \omega_S}}{\beta k} - \frac{e^{-\beta k \omega_{\max}}}{\beta k} - (\omega_{\max} - \omega_S) e^{-\beta k \omega_{\max}} \right] \\ &\quad - \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k d \omega_S}}{\beta k} - \frac{e^{-\beta k d \omega_{\max}}}{\beta k} - d (\omega_{\max} - \omega_S) e^{-\beta k d \omega_{\max}} \right]. \end{aligned} \quad (\text{A.42})$$

As a side remark, note that here one can see that in the special case of equally spaced Hamiltonians, one indeed requires a diverging number of machine subsystems to attain perfect cooling at the Landauer limit, as this is the only way to fulfil the condition of Theorem 1.3. This follows from the fact that the approximation $\frac{x}{1 - e^{-\beta x}} \approx \frac{1}{\beta}$ only holds for small βx and in general one would need to include higher-order terms that lead to an increase in energy cost.

We then have, using the expression for $E(\omega_x, \omega_y)$ derived earlier:

$$\begin{aligned}
 \Delta E_{\mathcal{SM}}^{\text{tot}} &= -\frac{1}{\beta} \log(1 - e^{-\beta\omega_S}) + \frac{1}{\beta} \log(1 - e^{-\beta\omega_{\max}}) - \frac{(\omega_{\max} - \omega_S) e^{-\beta\omega_{\max}}}{1 - e^{-\beta\omega_{\max}}} \\
 &\quad + \frac{1}{\beta} \log(1 - e^{-\beta d\omega_S}) - \frac{1}{\beta} \log(1 - e^{-\beta d\omega_{\max}}) + \frac{d(\omega_{\max} - \omega_S) e^{-\beta d\omega_{\max}}}{1 - e^{-\beta d\omega_{\max}}} \\
 &= \frac{1}{\beta} \log\left(\frac{1 - e^{-\beta d\omega_S}}{1 - e^{-\beta\omega_S}}\right) - \frac{1}{\beta} \log\left(\frac{1 - e^{-\beta d\omega_{\max}}}{1 - e^{-\beta\omega_{\max}}}\right) \\
 &\quad - (\omega_{\max} - \omega_S) \left(\frac{e^{-\beta\omega_{\max}}}{1 - e^{-\beta\omega_{\max}}} - \frac{d e^{-\beta d\omega_{\max}}}{1 - e^{-\beta d\omega_{\max}}} \right) \\
 &= \frac{1}{\beta} \log[\mathcal{Z}_S(\beta, \omega_S)] - \frac{1}{\beta} \log[\mathcal{Z}_S(\beta, \omega_{\max})] \\
 &\quad - \text{tr}[H_S(\omega_{\max}) \tau_S(\beta, \omega_{\max})] + \text{tr}[H_S(\omega_S) \tau_S(\beta, \omega_{\max})] \\
 &= \frac{1}{\beta} \log[\mathcal{Z}_S(\beta, \omega_S)] - \frac{1}{\beta} \log[\mathcal{Z}_S(\beta, \omega_{\max})] \\
 &\quad - \text{tr}[H_S(\omega_{\max}) \tau_S(\beta, \omega_{\max})] + \text{tr}[H_S(\omega_S) \tau_S(\beta, \omega_S)] \\
 &\quad - \text{tr}[H_S(\omega_S) \tau_S(\beta, \omega_S)] + \text{tr}[H_S(\omega_S) \tau_S(\beta, \omega_{\max})] \\
 &= \frac{1}{\beta} \Delta S_S + \Delta E_S, \tag{A.43}
 \end{aligned}$$

where we have explicitly written the von Neumann entropy $S(\rho) = -\text{tr}[\rho \log(\rho)]$ of a thermal state at inverse temperature β as $S[\tau_S(\beta, \omega)] = \log[\mathcal{Z}_S(\beta, \omega)] + \beta E[\tau_S(\beta, \omega)]$. Since the energy change of the target system only concerns its local Hamiltonian, we immediately see that the heat dissipated by the resetting of machines in such a cooling process, i.e., $\Delta E_{\mathcal{M}}$, saturates the Landauer bound as it is equal to $\beta^{-1} \Delta S_S$. The process described is thus energy-optimal.

A.4 Conditions for Structural and Control Complexity

Here we begin by considering the protocol-independent structural conditions that must be fulfilled by the machine Hamiltonian to enable (1) *perfect cooling* and (2) *cooling at Landauer cost*; combined, these independent conditions provide a necessary requirement, namely that the machine must be infinite-dimensional with a spectrum that is unbounded (from above) for the possibility of (3) *perfect cooling at the Landauer limit*. We then turn to analyse the control complexity, which concerns the properties of the interaction that implements a given protocol. The properties of the machine Hamiltonian define the *structural complexity*, which set the potential for how cool the target system can be made and at what energy cost; the extent to which a machine's potential is utilised in a particular protocol then depends on the properties of the joint unitary, i.e., the *control complexity*. Here, we show that it is necessary that any protocol achieving perfect cooling at the Landauer limit involves interactions between the target and infinitely-many levels of the machine to realise the full cooling potential. We then analyse some sufficient

conditions that arise as observations from our diverging control complexity protocols. This then leads us to demonstrate that individual degrees of freedom of the machine must be addressed in a fine-tuned manner to permute populations, highlighting that an operationally meaningful notion of control complexity must take into account factors beyond the effective dimensionality.

A.4.1 Necessary Complexity Conditions

Necessary Structural Conditions

1. *Perfect Cooling.*—Let us consider the task of perfect cooling, independently from protocol-specific constraints, in the envisaged setting. One can lower-bound the smallest eigenvalue λ_{\min} of the final state ϱ'_S (and hence how cold the system can become) after any unitary interaction with a thermal machine by [200]

$$\lambda_{\min}(\varrho'_S) \geq e^{-\beta\omega_{\mathcal{M}}^{\max}} \lambda_{\min}(\varrho_S), \quad (\text{A.44})$$

where $\omega_{\mathcal{M}}^{\max} := \max_{i,j} |\omega_j - \omega_i|$ denotes the largest energy gap of the machine Hamiltonian $H_{\mathcal{M}}$ with eigenvalues ω_i . Without loss of generality, throughout this article we set the ground state energy of any system to be zero, i.e., $\omega_0 = 0$, such that the largest energy gap coincides with the largest energy eigenvalue. As we have made no restrictions on the size or structure of the target or machine, the above inequality pertains to cooling protocols that could, for instance, be realised via sequences of unitaries on the target and parts of the machine. It follows that perfect cooling is only possible under two conditions: Either the machine begins in a pure state ($\beta \rightarrow \infty$), or $H_{\mathcal{M}}$ is unbounded, i.e., $\omega_{\mathcal{M}}^{\max} \rightarrow \infty$. Requiring $\beta < \infty$, a diverging energy gap in the machine Hamiltonian is thus a necessary structural condition for perfect cooling. Indeed, the largest energy gap of the machine plays a crucial role in limiting how cool the target system can be made (see also, e.g., Refs. [82, 359]). We now detail an independent property that is required for cooling with minimal energetic cost.

2. *Cooling at the Landauer Limit.*—Suppose now that one wishes to cool an initial target state $\tau_S(\beta, H_S)$ to any thermal state $\tau'_S(\beta^*, H_S)$ with $\beta^* > \beta$ (not necessarily close to a pure state), at an energy cost saturating the Landauer limit. In Ref. [200], it was shown that for any finite-dimensional machine, there are correction terms to the Landauer bound which imply that it cannot be saturated; these terms only vanish in the limit where the machine dimension diverges. Thus, a necessary condition for achieving cooling with energy cost at the Landauer limit is provided by the following:

Theorem A.3. *To cool a target system $\tau_S(\beta, H_S)$ to $\tau_S(\beta^*, H_S)$, with $\beta^* > \beta$, using a machine in the initial state $\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})$ with energy cost at the Landauer limit, the machine must be infinite-dimensional.*

As we will discuss below, this minimal requirement for the notion of complexity is far from sufficient to achieve cooling at Landauer cost.

3. *Perfect Cooling at the Landauer Limit.*—We have two independent necessary conditions on the structure of the machine that must be asymptotically achieved to enable relevant goals for cooling: The former is required to achieve perfect cooling; the latter for cooling at the Landauer limit. Together, these conditions imply that in order to achieve perfect cooling at the Landauer limit, one must have an infinite-dimensional machine with a spectrum that is unbounded (from above), as stated in Corollary 1.2.

Henceforth, we will assume that these conditions are satisfied by the machine. The question then becomes: *How does one engineer an interaction between the target system and machine to achieve perfect cooling at Landauer cost?*

Necessary Control Complexity Conditions

The unbounded structural properties of the machine support the possibility for perfect cooling at the Landauer limit; however, we now focus on the control properties of the interaction that realise said potential (see Fig. 1.2). This leads to the distinct notion of *control complexity*, which aims to differentiate between protocols that access the machine in a more or less complex manner. The structural complexity properties are protocol-independent and related to the energy spectrum and dimensionality of the machine, whereas the control complexity concerns properties of the unitary that represents a particular protocol. For instance, the diverging-time protocol previously outlined comprises a sequence of interactions, each of which is individually not very complex; at the same time, the unconstrained control complexity protocol accesses the total (overall infinite-dimensional) machine “at once”, and thus the number of (nontrivial) terms in the interaction Hamiltonian, or the effective dimensionality of the machine accessed by the unitary, becomes unbounded. Nonetheless, the net energy cost of this protocol with unconstrained control complexity remains in accordance with the Landauer limit, as the initial and final states of both the system and machine are identical to those in the diverging-time protocol.

Effective Dimensionality.—We begin by considering the effective dimensionality accessed (nontrivially) by a unitary, whose divergence is necessary but insufficient for achieving perfect cooling at the Landauer limit, as we show in the next section. This in turn motivates the desire for a more detailed notion of control complexity that takes into account the energy-level structure of the machine.

We define the effective dimension as the dimension of the subspace of the global Hilbert space upon which the unitary acts nontrivially, which can be quantified via the minimum dimension of a subspace \mathcal{A} of the joint Hilbert space \mathcal{H}_{SM} in terms of which the unitary

can be decomposed as $U_{S\mathcal{M}} = U_{\mathcal{A}} \oplus \mathbb{1}_{\mathcal{A}^\perp}$, i.e.,

$$d^{\text{eff}} := \min \dim(\mathcal{A}) : U_{S\mathcal{M}} = U_{\mathcal{A}} \oplus \mathbb{1}_{\mathcal{A}^\perp}. \quad (\text{A.45})$$

One can relate this quantity to properties of the Hamiltonian that generates the evolution in a finite unit of time T (which we can set equal to unity without loss of generality) by considering the interaction picture. In general, any global unitary $U_{S\mathcal{M}} = e^{-iH_{S\mathcal{M}}T}$ is generated by a Hamiltonian of the form $H_{S\mathcal{M}} = H_S \otimes \mathbb{1}_{\mathcal{M}} + \mathbb{1}_S \otimes H_{\mathcal{M}} + H_{\text{int}}$. However, all protocols considered in this work have vanishing local terms, i.e., $H_S = H_{\mathcal{M}} = 0$. More generally, one can argue that the local terms play no role in how the machine is used to cool the target. As such, one can consider unitaries generated by only the non-trivial term H_{int} to be those representing a particular protocol of interest. That is, we can restrict our attention to $U_{S\mathcal{M}} = e^{-iH_{\text{int}}T}$, where H_{int} is a Hermitian operator on $\mathcal{H}_{S\mathcal{M}}$ of the form $\sum_i A_S^i \otimes B_{\mathcal{M}}^i$ such that none of the $A_S^i, B_{\mathcal{M}}^i$ are proportional to the identity operator. In doing so, it follows that the effective dimension corresponds to $\text{rank}(H_{\text{int}})$. Lastly, note that the above definition in terms of a direct sum decomposition provides an upper bound on any similar quantification of effective dimensionality based on other tensor factorisations of the joint Hilbert space considered and makes no assumption about the underlying structure. On the other hand, knowledge of said structure would permit a more meaningful notion of complexity to be defined. For instance, the effective dimensionality of a unitary acting on a many qubit system is better captured by considering its decomposition into a tensor product factorisation rather than the direct sum. We leave the exploration of such considerations to future work.

The effective dimensionality provides a minimal quantifier for a notion of control complexity, insofar as its divergence is necessary for saturating the Landauer bound, as we prove in the next section. In fact, we prove a slightly stronger statement, namely that the dimension of the machine Hilbert space to which the unitary (nontrivially) couples the target system to must diverge. However, as we will discuss below, $d^{\text{eff}} \rightarrow \infty$ is generally insufficient to achieve said goal, and fine-tuned control is required. Nonetheless, the manifestation of such control seems to be system-dependent, precluding our ability (so far) to present a universal quantifier of control complexity. Thus, even though further conditions need to be met to achieve perfect cooling at minimal energy cost in unit time (see Theorem A.4), whenever we talk of an operation with finite control complexity, we mean those represented by a unitary that acts (nontrivially) only on a finite-dimensional subspace of the target system and machine. In contrast, by diverging control complexity, we mean a unitary that couples the target (nontrivially) to a full basis of the machine's Hilbert space, whose dimension diverges. With this notion at hand, we have Theorem 1.3, which is proven below. Intuitively, we show that if a protocol only accesses a finite-dimensional subspace of the machine, then the machine is effectively finite-dimensional inasmuch as a suitable replacement can be made while keeping all quantities relevant for

cooling invariant. Invoking then the main result of Ref. [200], there are finite-dimensional correction terms that then imply that the Landauer limit cannot be saturated.

Note finally that in Theorem 1.3 no particular structure of the systems is presupposed and the effective dimensionality relates to various notions of complexity put forth throughout the literature (see, e.g., Refs. [360, 361]). For instance, for a finite-dimensional target system with equally spaced energy levels ω_S , suppose that the machine structure is decomposed as N qubits with energy gaps $\omega_{M_n} \in \{\omega_S + n\epsilon\}_{n=1,\dots,N}$, with arbitrarily small $\epsilon > 0$ and $N \rightarrow \infty$. Then the overall unitary that approaches perfect cooling at the Landauer limit has circuit complexity equal to the diverging N .

A.4.2 Proof of Theorem 1.3, Corollary 1.2, and Theorem A.3

Here we prove Theorem 1.3, which implies Theorem A.3 and leads to Corollary 1.2.

Proof. Let $\mathcal{H}_{\mathcal{X}}$ be a separable Hilbert space associated with the system \mathcal{X} . Consider

$$H_{\mathcal{M}} = \sum_{n=0}^{\infty} \omega_n |n\rangle\langle n| \quad \text{and} \quad H_{\mathcal{M}'} = \text{span}_{n \leq m} \{|n\rangle\}, \quad (\text{A.46})$$

for some finite m . In other words, $H_{\mathcal{M}'}$ is a finite-dimensional restriction of $H_{\mathcal{M}}$. We will show that any unitary that (nontrivially) interacts the target system with only a subspace spanned by finitely many eigenstates of $H_{\mathcal{M}}$ cannot attain Landauer's bound. Consider a general unitary U . Suppose that U only couples \mathcal{H}_S with $\mathcal{H}_{\mathcal{M}'}$; whenever we talk of an operation with finite control complexity in this article, we mean specifically such a U , and by diverging control complexity we mean a unitary that couples the target to any subspace of $H_{\mathcal{M}}$ whose dimension diverges. Since

$$\mathcal{H}_S \otimes \mathcal{H}_{\mathcal{M}} = \mathcal{H}_S \otimes (\mathcal{H}_{\mathcal{M}'} \oplus \mathcal{H}_{\mathcal{M}'}) \simeq (\mathcal{H}_S \otimes \mathcal{H}_{\mathcal{M}'}) \oplus (\mathcal{H}_S \otimes \mathcal{H}_{\mathcal{M}'})^{\perp}, \quad (\text{A.47})$$

we can associate the subspace $\mathcal{H}_S \otimes \mathcal{H}_{\mathcal{M}'}$ with the label \mathcal{A} and $\mathcal{H}_S \otimes \mathcal{H}_{\mathcal{M}'})^{\perp}$ with \mathcal{B} and write $U = U_{\mathcal{A}} \oplus \mathbb{1}_{\mathcal{B}}$. Then the initial configuration can be expressed as:

$$\varrho_S \otimes \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}}) = \begin{bmatrix} \varrho_S \otimes \varrho_{\mathcal{M}'} & 0 \\ 0 & \varrho_S \otimes \varrho_{\mathcal{M}'}^{\perp} \end{bmatrix}, \quad (\text{A.48})$$

where

$$\varrho_{\mathcal{M}'} := \frac{1}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \sum_{n \leq m} e^{-\beta \omega_n} |n\rangle\langle n| \quad \text{and} \quad \varrho_{\mathcal{M}'}^{\perp} := \frac{1}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \sum_{n > m} e^{-\beta \omega_n} |n\rangle\langle n| \quad (\text{A.49})$$

add up to a (normalised) thermal state. Now consider the state

$$\tilde{\varrho}_{\mathcal{M}} = \begin{bmatrix} \varrho_{\mathcal{M}'} & 0 \\ 0 & \text{tr}[\varrho_{\mathcal{M}'}^{\perp}] \end{bmatrix}. \quad (\text{A.50})$$

It is straightforward to check that is indeed a quantum state; moreover, it is the Gibbs state (at inverse temperature β) associated with the Hamiltonian

$$\widetilde{H}_{\mathcal{M}} = \sum_{n \leq m} \omega_n |n\rangle\langle n| - \frac{1}{\beta} \log \left(\sum_{n > m} e^{-\beta \omega_n} \right) |m+1\rangle\langle m+1|. \quad (\text{A.51})$$

To see this, note that $\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}}) = \mathcal{Z}_{\mathcal{M}}(\beta, \widetilde{H}_{\mathcal{M}})$ and that

$$\exp \left\{ -\beta \left[-\frac{1}{\beta} \log \left(\sum_{n > m} e^{-\beta \omega_n} \right) \right] \right\} = \sum_{n > m} e^{-\beta \omega_n}. \quad (\text{A.52})$$

Thus $\tilde{\varrho}_{\mathcal{M}} = \tau_{\mathcal{M}}(\beta, \widetilde{H}_{\mathcal{M}})$. To ease notation in what follows, we will write $\tilde{\omega}_{m+1} := -\frac{1}{\beta} \log \left(\sum_{n > m} e^{-\beta \omega_n} \right)$. In the rest of the proof, we will show that the unitary U and the Hamiltonian $H_{\mathcal{M}}$ can be replaced by finite-dimensional versions without changing the quantities relevant for Landauer's principle.

Let $\tilde{U} = U_{\mathcal{A}} \oplus (\mathbb{1}_{\mathcal{S}} \otimes |m+1\rangle\langle m+1|)$. We then have

$$\tilde{U} (\varrho_{\mathcal{S}} \otimes \tilde{\varrho}_{\mathcal{M}}) \tilde{U}^{\dagger} = \begin{bmatrix} U_{\mathcal{A}}(\varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^{\dagger} & 0 \\ 0 & \frac{e^{-\beta \tilde{\omega}_{m+1}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \varrho_{\mathcal{S}} \end{bmatrix} \quad (\text{A.53})$$

and

$$\text{tr}_{\mathcal{M}} \left[\tilde{U} (\varrho_{\mathcal{S}} \otimes \tilde{\varrho}_{\mathcal{M}}) \tilde{U}^{\dagger} \right] = \text{tr}_{\mathcal{M}'} \left[U_{\mathcal{A}} (\varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'}) U_{\mathcal{A}}^{\dagger} \right] + \frac{e^{-\beta \tilde{\omega}_{m+1}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \varrho_{\mathcal{S}}. \quad (\text{A.54})$$

Compare this to the expression

$$\begin{aligned} \text{tr}_{\mathcal{M}} \left[U (\varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}}) U^{\dagger} \right] &= \text{tr}_{\mathcal{M}'} \begin{bmatrix} U_{\mathcal{A}}(\varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^{\dagger} & 0 \\ 0 & \varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'}^{\perp} \end{bmatrix} \\ &= \text{tr}_{\mathcal{M}'} \left[U_{\mathcal{A}}(\varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^{\dagger} \right] + \text{tr} \left[\varrho_{\mathcal{M}'}^{\perp} \right] \varrho_{\mathcal{S}} \\ &= \text{tr}_{\mathcal{M}'} \left[U_{\mathcal{A}}(\varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^{\dagger} \right] + \frac{e^{-\beta \tilde{\omega}_{m+1}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \varrho_{\mathcal{S}}, \end{aligned} \quad (\text{A.55})$$

since $\text{tr} \left[\varrho_{\mathcal{M}'}^{\perp} \right] = \frac{1}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \sum_{n > m} e^{-\beta \omega_n}$. Thus, the final system state is the same as it would be if we replaced the full initial machine state with $\tilde{\varrho}_{\mathcal{M}}$; in particular, the entropy decrease of the system for any unitary that cools it is also unchanged.

The last thing we need to check is that the energy change of the machine similarly remains invariant. To that end, we have that

$$\begin{aligned} \tilde{\mathcal{E}}_{\mathcal{M}} &= \text{tr}_{\mathcal{S}} \left[\tilde{U} (\varrho_{\mathcal{S}} \otimes \tilde{\varrho}_{\mathcal{M}}) \tilde{U}^{\dagger} \right] = \text{tr}_{\mathcal{S}} \left[U_{\mathcal{A}} (\varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'}) U_{\mathcal{A}}^{\dagger} \right] + \frac{e^{-\beta \tilde{\omega}_{m+1}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} |m+1\rangle\langle m+1|, \\ \tilde{\varrho}_{\mathcal{M}} &= \varrho_{\mathcal{M}'} + \frac{e^{-\beta \tilde{\omega}_{m+1}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} |m+1\rangle\langle m+1|. \end{aligned} \quad (\text{A.56})$$

Thus, we have

$$\text{tr} \left[\widetilde{H}_{\mathcal{M}} (\tilde{\varrho}'_{\mathcal{M}} - \tilde{\varrho}_{\mathcal{M}}) \right] = \text{tr} \left\{ H_{\mathcal{M}} \left[\text{tr}_{\mathcal{S}} \left[U_{\mathcal{A}} (\varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'}) U_{\mathcal{A}}^{\dagger} \right] - \varrho_{\mathcal{M}'} \right] \right\}, \quad (\text{A.57})$$

since U_S only acts on $\mathcal{H}_S \otimes \mathcal{H}_{\mathcal{M}'}$ and $\widetilde{H}_{\mathcal{M}|\mathcal{M}'} = H_{\mathcal{M}|\mathcal{M}'}$. In the same way, we have

$$\begin{aligned} \text{tr}_S [U(\varrho_S \otimes \varrho_{\mathcal{M}})U^\dagger] &= \text{tr}_S [U_{\mathcal{A}}(\varrho_S \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^\dagger] + \varrho_{\mathcal{M}}^\perp \\ \varrho_{\mathcal{M}} &= \varrho_{\mathcal{M}'} + \varrho_{\mathcal{M}'}^\perp. \end{aligned} \quad (\text{A.58})$$

Thus, the energy difference is also

$$\text{tr} \left\{ H_{\mathcal{M}} \left[\text{tr}_S [U_{\mathcal{A}}(\varrho_S \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^\dagger] - \varrho_{\mathcal{M}'} \right] \right\}. \quad (\text{A.59})$$

Hence, we have shown that we can replace \mathcal{M} by some $m + 1$ dimensional machine $\widetilde{\mathcal{M}}$ if the joint unitary U only acts on m levels of $H_{\mathcal{M}}$. By Theorem 6 of Ref. [200], there are finite-dimensional corrections to the Landauer bound, which then imply that it cannot be reached for finite m . Thus, the effective machine dimension, i.e., that which is actually (nontrivially) accessed throughout the interaction, must diverge in order for cooling to be possible at the Landauer limit. This proves Theorem 1.3, which implies Theorem A.3. \square

A.4.3 Sufficient Complexity Conditions

Having shown the necessary requirements for cooling at Landauer cost, namely a control interaction that acts nontrivially on an infinite-dimensional (sub)space of the machine's Hilbert space, let us now return to emphasise the properties of the machine and cooling protocol that are sufficient to achieve perfect cooling at Landauer cost. For simplicity, we consider the case of a qubit, which exemplifies the discussion of finite-dimensional systems. The case of infinite-dimensional systems shall be treated independently in the next appendix.

We first consider the structural properties of the machine. The diverging-time protocol discussed in Appendix A.3 makes use of a diverging number N of machines. Thus, the machine begins in the thermal state $\tau(\beta, H_{\mathcal{M}}^{\text{tot}})$ of a (2^N) -dimensional system (with N eventually diverging), with energy-level structure given by the sum of the Hamiltonians in Eq. (A.17), i.e.,

$$H_{\mathcal{M}}^{\text{tot}} = \sum_{n=1}^N H_{\mathcal{M}_n}^{(n)} = \sum_n (1 + n\epsilon) H_S^{(n)}, \quad (\text{A.60})$$

that acts on the full Hilbert space (we use the usual convention that it acts as identity on unlabelled subspaces, e.g., $H_{\mathcal{M}}^{(1)} \equiv H_{\mathcal{M}}^{(1)} \otimes \mathbb{1}^{(2)} \otimes \dots \otimes \mathbb{1}^{(N)}$). Let us analyse in detail the properties of this Hamiltonian. The ground state is $|0\rangle^{\otimes N}$, which is set at zero energy. More generally, the energy eigenvalue corresponding to an eigenstate $|i_0, i_1, \dots, i_N\rangle$ is given by ω_1 multiplied by the number of indices i_k that are equal to 1, plus a sum of terms $k\epsilon$ where k is the label of each index equal to 1. Thus, the energy eigenvalue of the eigenstate $|1, \dots, 1\rangle$ diverges as the number of subsystems diverges. At the same time, letting the factor ϵ go to zero renders all eigenstates with the same (constant) number of indices such that $i_k = 1$ approach the same energy. Thus, in the limit $\epsilon \rightarrow 0$, one

obtains subspaces of energy $E_{\mathcal{M}}^{(k)} = k\omega_1$ with degeneracy given by $D_k = \binom{N}{k}$, which also diverges for each constant k and diverging N . Therefore, in addition to satisfying the structural conditions that are necessary for perfect cooling, as stated in Theorem A.3, the machine used here features additional properties, which are crucially important for this particular protocol, in particular because they are sufficient for perfect cooling at Landauer cost. As a remark, we also emphasise that for fixed (large) N and (small) ϵ , the machine is finite-dimensional and has a non-degenerate Hamiltonian without any energy levels formally at infinity.

Concerning the control complexity properties of the unitary that achieves perfect cooling in unit time, note that it is a cyclic shift operator, which can be written as

$$\begin{aligned} U_{S\mathcal{M}} &= \Pi_{n=1}^N \mathbb{S}_{S\mathcal{M}_n}^2 = \Pi_n \left(\sum_{i,j_n=0}^1 |i, j_1, \dots, j_n, \dots, j_N\rangle\langle j_n, j_1, \dots, i, \dots, j_N|_{S\mathcal{M}} \right) \\ &= \sum_{i,j_1 \dots j_N=0}^1 |i, j_1, \dots, j_N\rangle\langle j_N, i, j_1, \dots, j_{N-1}|_{S\mathcal{M}}. \end{aligned} \quad (\text{A.61})$$

As it is evident from its form, this unitary acts nontrivially on all of the (divergingly many) energy levels of the machine. The only basis vectors of the system-plus-machine Hilbert space that are left invariant are $|i = 0, j_1 = 0, \dots, j_N = 0\rangle$ and $|i = 1, j_1 = 1, \dots, j_N = 1\rangle$.

A.4.4 Fine-tuned Control Conditions

Theorem 1.3 captures a notion of control complexity as a resource in a thermodynamically consistent manner, i.e., in line with Nernst's unattainability principle. However, following the discussion around Theorem A.3 and that above, the protocols that we present that achieve perfect cooling at Landauer cost make use of machines and interactions with a far more complicated structure than suggested by the necessary condition of infinite effective dimensionality. In particular, the interactions couple the target system to a diverging number of subspaces of the machine corresponding to distinct energy gaps in a fine-tuned manner. Moreover, there are a diverging number of energy levels of the machine both above and below the first excited level of the target.

This suggests that an operationally meaningful quantifier of control complexity must take into account the energy-level structure of the machine that is accessed throughout any given protocol; additionally that of the target system plays a role. Indeed, both the final temperature of the target as well as the energy cost required to achieve this depends upon how the global eigenvalues are permuted via the cooling process. First, how cool the target becomes depends on the sum of the eigenvalues that are placed into the subspace spanned by the ground state. Second, for any fixed cooling amount, the energy cost depends on the constrained distribution of eigenvalues within the machine. Thus, in general, the optimal permutation of eigenvalues depends upon properties of the target and machine.

For instance, consider an arbitrary initially thermal target qubit, whose state is given by $\text{diag}(p, 1 - p)$ and a thermal machine of dimension $d_{\mathcal{M}}$ with spectrum $\{\lambda_{\mathcal{M}}^i\}_{i=0, \dots, d_{\mathcal{M}}-1}$. Now consider the decomposition of the joint Hilbert space into two orthogonal subspaces, \mathcal{B}_0 and \mathcal{B}_1 , corresponding to the ground and excited eigenspaces of the target. The initial joint state is $p \text{diag}(\lambda_{\mathcal{B}_0}^i) \oplus (1 - p) \text{diag}(\lambda_{\mathcal{B}_1}^i)$, where we have written $\lambda_{\mathcal{B}_j}^i$ to denote the i^{th} machine eigenvalue in the subspace \mathcal{B}_j . The total population in the subspaces \mathcal{B}_0 and \mathcal{B}_1 are p and $(1 - p)$ respectively. To achieve perfect cooling one must permute the eigenvalues such that approximately a net transfer of population $(1 - p)$ is moved from \mathcal{B}_1 to \mathcal{B}_0 . To do this, one can take any subset K of $\{\lambda_{\mathcal{B}_1}^i\}$ such that as $d_{\mathcal{M}} \rightarrow \infty$, $\sum_{i \in K} \lambda_{\mathcal{B}_1}^i \rightarrow (1 - p)$ and a subset K' (with $|K| = |K'|$) from $\{\lambda_{\mathcal{B}_0}^i\}$ such that $\sum_{i \in K'} \lambda_{\mathcal{B}_0}^i \rightarrow 0$ and exchange them. Although the choice of eigenvalues permuted is non-unique, the requirement must be fulfilled for some sets to perfectly cool the target. For any pair of eigenvalues exchanged between the subspaces, demanding that the exchange costs minimal energy amounts to a fine-tuning condition of the form $\lambda_{\mathcal{M}}^i \rightarrow p \lambda_{\mathcal{B}_0}^i + (1 - p) \lambda_{\mathcal{B}_1}^i$ that must be satisfied. In general, the fine-tuned eigenvalue conditions that must be asymptotically attained depend upon target and machine eigenvalues, making it difficult to derive a closed-form expression. However, in the restricted scenario in which the target qubit begins maximally mixed (i.e., at infinite temperature), the machine begins thermal at some $\beta > 0$ and of dimension $d_{\mathcal{M}}$, and that the unitary implemented is such that the target is cooled as much as possible, one can derive precise conditions in terms of the machine structure alone, as we demonstrate below. The case for higher-dimensional target systems is similar.

This discussion highlights the importance of capturing properties beyond the effective dimensionality, e.g., those regarding the distribution of machine (and, more generally, target system) eigenvalues, in order to meaningfully quantify control complexity in thermodynamics. Our protocols display similar behaviour to that discussed above asymptotically. Moreover, the machines exhibit an energy-level structure such that every possible energy gap is present, i.e., the set of machine energy gaps $\{\omega_{ij} = \omega_i - \omega_j\}$ densely covers the interval $[\omega_S, \infty)$, where ω_S is the energy of the first excited level of the target. Whether or not such a condition is necessary in general remains an open question.

Here we analyse the fine-tuned control conditions that are asymptotically required for cooling at the Landauer limit. We begin with some general considerations before focusing on a special case for which an analytic expression can be derived. Furthermore, we demand that the unitary implemented is such that the target is cooled as much as possible: This does not preclude the possibility for cooling the target system less (albeit still close to a pure state) at a cost closer to the Landauer bound without satisfying all of the fine-tuning conditions. Nonetheless, in general there are a number of such conditions to be satisfied, and the special case serves as a pertinent example that demonstrates how the particular set of fine-tuning conditions for any considered scenario can be derived.

Consider an arbitrary thermal target system and machine of finite dimensions, with re-

spective spectra $\boldsymbol{\lambda}_S := \{\lambda_S^0, \dots, \lambda_S^{d_S-1}\}$ and $\boldsymbol{\lambda}_M := \{\lambda_M^0, \dots, \lambda_M^{d_M-1}\}$. The states begin uncorrelated, so the global spectrum of the initial joint state is $\boldsymbol{\lambda}_{S,M} := \{\lambda_{S,M}^0, \dots, \lambda_{S,M}^{d_S d_M-1}\} = \{\lambda_S^0 \lambda_M^0, \lambda_S^0 \lambda_M^1, \dots, \lambda_S^{d_S-1} \lambda_M^{d_M-1}\}$. Consider now a global unitary transformation; such a transformation cannot change the values of the spectrum, but merely permute them. In other words, the spectrum of the final global state after any such unitary is invariant and we have equivalence of the (unordered) sets $\boldsymbol{\lambda}'_{S,M}$ and $\boldsymbol{\lambda}_{S,M}$.

The transformation that cools the target system as much as possible¹ is the one that places the d_M largest of the global eigenvalues into the subspace spanned by the ground state of the target, the second d_M largest into that spanned by the first excited state of the target, and so forth, with the smallest d_M global eigenvalues placed in the subspace corresponding to the highest energy eigenstate of the target system (we prove this statement shortly). More precisely, we denote by $\boldsymbol{\lambda}^\downarrow$ the non-increasing ordering of the set $\boldsymbol{\lambda}$. Since the target and machine begin thermal, the local spectra $\boldsymbol{\lambda}_S$ and $\boldsymbol{\lambda}_M$ are already ordered in this way with respect to their energy eigenbases, which we consider to be labelled in non-decreasing order. Cooling the target system as much as possible amounts to achieving the final reduced state of the target

$$\rho'_S = \sum_{i=0}^{d_S-1} \left(\sum_{j=0}^{d_M-1} \lambda_{S,M}^{\downarrow i d_M + j} \right) |i\rangle\langle i|_S. \quad (\text{A.62})$$

As a side remark, note that since each of the global eigenvalues are a product of the initial local eigenvalues (due to the initial tensor product structure), which are in turn related to the energy-level structure of the target system and machine (as they begin as thermal states), one can already see here that in order to approach perfect cooling, the machine must have some diverging energy gaps, such that the (finite) sum of the global eigenvalues contributing to the ground-state population of the target approaches 1.

Of course, there is an equivalence class of unitaries that can achieve the same amount of cooling; in particular, any permutation of the set of the d_M global eigenvalues within each energy eigenspace of the target system achieves the same amount of cooling, since it is the sum of these values that contribute to the total population in each subspace. Importantly, although such unitaries cool the target system to the same extent, their effect on the machine differs, and therefore so too does the energy cost. However, demanding that such cooling is achieved at minimal energy cost amounts to a unique constraint on the global post-transformation state, namely that it must be

$$\rho'_{S,M} = \sum_{i=0}^{d_S-1} \sum_{j=0}^{d_M-1} \lambda_{S,M}^{\downarrow i d_M + j} |ij\rangle\langle ij|_{S,M}. \quad (\text{A.63})$$

¹We take majorisation among passive states to be the measure of cooling; this implies the highest possible ground state population and purity, and lowest possible entropy and average energy via Schur convexity.

We can derive the above form of the final joint state as follows. Consider the following ordering for the energy eigenbasis of \mathcal{SM} chosen to match the above form

$$\begin{aligned} & \{|00\rangle_{\mathcal{SM}}, |01\rangle_{\mathcal{SM}}, \dots, |0, d_{\mathcal{M}} - 1\rangle_{\mathcal{SM}}, \\ & |10\rangle_{\mathcal{SM}}, \dots, |1, d_{\mathcal{M}} - 1\rangle_{\mathcal{SM}}, \\ & \vdots \\ & |d_{\mathcal{S}} - 1, 0\rangle_{\mathcal{SM}}, \dots, |d_{\mathcal{S}} - 1, d_{\mathcal{M}} - 1\rangle_{\mathcal{SM}} \}. \end{aligned} \quad (\text{A.64})$$

This ordering is monotonically non-decreasing primarily with respect to the energy of \mathcal{S} , and secondarily w.r.t. \mathcal{M} . We take the final state $\varrho'_{\mathcal{SM}}$ to be expressed in this basis. To maximise the cooling in a single unitary operation, we maximise the sum of the first $k \cdot d_{\mathcal{M}}$ diagonal elements, for each $k \in \{1, 2, \dots, d_{\mathcal{S}}\}$, as each sum corresponds to the total population in the k^{th} lowest energy eigenstate of \mathcal{S} . The initial state $\varrho_{\mathcal{SM}}$ is diagonal in this basis, so the vector of initial diagonal elements, which we label $\boldsymbol{\theta} := \text{diag}(\varrho_{\mathcal{SM}})$, is also the vector of eigenvalues, $\boldsymbol{\lambda}_{\mathcal{SM}}$, i.e., $\boldsymbol{\theta} = \boldsymbol{\lambda}_{\mathcal{SM}}$. Furthermore, since the unitary operation leaves the set of eigenvalues invariant, we have via the Schur-Horn lemma [362] that the vector of final diagonal elements, which we label $\boldsymbol{\theta}' := \text{diag}(\varrho'_{\mathcal{SM}})$, is majorised by the vector of initial ones, i.e., $\boldsymbol{\theta}' \prec \boldsymbol{\theta}$. It follows that the partial sums we wish to maximise are upper bounded by the corresponding partial sums of the $k \cdot d_{\mathcal{M}}$ largest diagonal elements of the initial state. We claim that the unitary that cools this maximal cooling amount at minimum energy cost is the one that permutes the diagonal elements to be ordered w.r.t. the basis ordering in Eq. (A.64).

More precisely, via the Schur-Horn lemma, one can always write $\boldsymbol{\theta}' = D\boldsymbol{\theta}$, with D a doubly stochastic matrix. The partial sums of the $k \cdot d_{\mathcal{M}}$ first elements are linear functions of the elements of $\boldsymbol{\theta}$. Thus the maximum values are obtained at the extremal points of the convex set of doubly stochastic matrices, which are the permutation matrices, via the Birkhoff-von Neumann theorem [362]. One can see by inspection that the optimal permutation matrices are the ones that place the largest $d_{\mathcal{M}}$ diagonal elements in the first block (i.e., the ground state eigenspace of \mathcal{S}), the next largest $d_{\mathcal{M}}$ elements in the second block (i.e., the first excited state eigenspace of \mathcal{S}), and so on. Within each block, the ordering does not affect the cooling of the target, so there is an equivalence class of permutations that satisfy the maximal cooling criterion. However, adding the optimisation over the energy cost eliminates this freedom. We may consider the reduced set of stochastic matrices that satisfy maximal cooling, generated by the permutations described above. Since the average energy of the final state is again a linear function of the diagonal elements, here too the minimum corresponds to a permutation matrix. Clearly the permutation that minimises the average energy is the one that orders the elements within each block to be decreasing w.r.t. the energies of \mathcal{M} . Thus, the unique²

²Note that degeneracies in energy eigenvalues would lead to sets of equal diagonal elements, and prevent one

stochastic matrix D which leads to maximal cooling at the least energy cost possible is the one that permutes the energy eigenvalues to be ordered decreasing primarily w.r.t. the system energies, and secondarily w.r.t. the machine energies. The action of the stochastic matrix on diagonal elements of the state is related to the unitary operation on the entire quantum state by $|U_{ij}|^2 = D_{ij}$, so that the unitary operation is also a permutation (up to an energy dependent phase, which is irrelevant since the initial and final states are diagonal).

We may understand this optimal operation through the notion of passivity, by noting that it cools at minimal energy cost by rendering the machine into the most passive reduced state in the joint unitary orbit with respect to the cooling constraint on the target. Intuitively, one has cooled the target system maximally at the expense of heating the machine as little as possible. The final reduced state of the machine corresponding to this energetically-optimal cooling transformation is

$$\varrho'_{\mathcal{M}} = \sum_{j=0}^{d_{\mathcal{M}}-1} \left(\sum_{i=0}^{d_{\mathcal{S}}-1} \lambda_{\mathcal{S}\mathcal{M}}^{\downarrow id_{\mathcal{M}}+j} \right) |j\rangle\langle j|_{\mathcal{M}}. \quad (\text{A.65})$$

In general, any unitary that achieves these desired conditions simultaneously depends upon the energy-level structure of both the target system and machine, precluding a closed-form set of conditions that can be expressed only in terms of the machine. However, for the special case of a maximally-mixed initial target state (i.e., cooling a thermal state at infinite temperature or erasing quantum information from its most entropic state), one can deduce this ordering precisely and moreover relate it directly to properties of the machine Hamiltonian, as we now demonstrate. In the following, we assume that $d_{\mathcal{M}}$ is even; the case for odd $d_{\mathcal{M}}$ can be derived similarly.

Theorem A.4. *Consider the target system to begin in the maximally mixed state and a thermal machine at temperature $\beta > 0$, whose eigenvalues are labelled in non-increasing order, $\{\lambda_{\mathcal{M}}^{\downarrow i}\}_{i=0,\dots,d_{\mathcal{M}}-1}$. In order to cool the target perfectly, with the restriction that the target must be cooled as much as possible, at an energy cost that saturates the Landauer limit, the machine eigenvalues must satisfy*

$$\sum_{i=0}^{\frac{d_{\mathcal{M}}}{2}-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 1, \quad \sum_{i=\frac{d_{\mathcal{M}}}{2}}^{d_{\mathcal{M}}-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 0, \quad (\text{A.66})$$

and

$$\frac{\frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2} + \lfloor \frac{i}{2} \rfloor} \right)}{\lambda_{\mathcal{M}}^{\downarrow i}} \rightarrow 1 \quad (\text{A.67})$$

from choosing a unique permutation. However, as the state in such degenerate subspaces is proportional to the identity matrix, we may take any unitary that is block diagonal w.r.t. the degeneracies without affecting the state, and hence the final cooling or average energy change.

for all $i \in \{0, \dots, d_{\mathcal{M}} - 1\}$, where $\lfloor \cdot \rfloor$ denotes the floor function and \rightarrow denotes that the condition is satisfied asymptotically, i.e., as $d_{\mathcal{M}} \rightarrow \infty$.³

Proof. We consider a qubit for simplicity, but the generalisation to cooling an arbitrary-dimensional maximally-mixed state is straightforward. The initial joint spectrum of the system and machine is

$$\boldsymbol{\lambda}_{S\mathcal{M}} = \frac{1}{2}\{\boldsymbol{\lambda}_{\mathcal{M}}^{\downarrow}, \boldsymbol{\lambda}_{\mathcal{M}}^{\downarrow}\} = \frac{1}{2}\{\lambda_{\mathcal{M}}^{\downarrow 0}, \lambda_{\mathcal{M}}^{\downarrow 1}, \dots, \lambda_{\mathcal{M}}^{\downarrow d_{\mathcal{M}}-1}, \lambda_{\mathcal{M}}^{\downarrow 0}, \lambda_{\mathcal{M}}^{\downarrow 1}, \dots, \lambda_{\mathcal{M}}^{\downarrow d_{\mathcal{M}}-1}\}. \quad (\text{A.68})$$

As each $\lambda_{\mathcal{M}}^{\downarrow i} = \frac{1}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} e^{-\beta \omega_i}$ for any thermal state with Hamiltonian $H_{\mathcal{M}} = \sum_i \omega_i |i\rangle\langle i|_{\mathcal{M}}$ written with respect to non-decreasing energy eigenvalues, it follows that the globally ordered spectrum is

$$\boldsymbol{\lambda}_{S\mathcal{M}}^{\downarrow} = \frac{1}{2}\{\lambda_{\mathcal{M}}^{\downarrow 0}, \lambda_{\mathcal{M}}^{\downarrow 0}, \lambda_{\mathcal{M}}^{\downarrow 1}, \lambda_{\mathcal{M}}^{\downarrow 1}, \dots, \lambda_{\mathcal{M}}^{\downarrow d_{\mathcal{M}}-1}, \lambda_{\mathcal{M}}^{\downarrow d_{\mathcal{M}}-1}\}. \quad (\text{A.69})$$

Expressing the global states with respect to the product of local energy eigenbases, we have that the initial joint state is $\frac{1}{2} \otimes \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}}) = \text{diag}(\boldsymbol{\lambda}_{S\mathcal{M}})$ [see Eq. (A.68)] and the unitary that cools the target as much as possible at minimum energy cost is the one achieving the globally passive final joint state $\varrho'_{S\mathcal{M}} = \text{diag}(\boldsymbol{\lambda}_{S\mathcal{M}}^{\downarrow})$. This leads to the following reduced states

$$\varrho'_S = \left(\sum_{i=0}^{\frac{d_{\mathcal{M}}-1}{2}} \lambda_{\mathcal{M}}^{\downarrow i} \right) |0\rangle\langle 0|_S + \left(\sum_{i=\frac{d_{\mathcal{M}}}{2}}^{d_{\mathcal{M}}-1} \lambda_{\mathcal{M}}^{\downarrow i} \right) |1\rangle\langle 1|_S, \quad (\text{A.70})$$

$$\begin{aligned} \varrho'_{\mathcal{M}} &= \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow 0} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2}} \right) |0\rangle\langle 0|_{\mathcal{M}} + \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow 0} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2}} \right) |1\rangle\langle 1|_{\mathcal{M}} + \\ &+ \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow 1} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2}+1} \right) |2\rangle\langle 2|_{\mathcal{M}} + \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow 1} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2}+1} \right) |3\rangle\langle 3|_{\mathcal{M}} + \dots \end{aligned} \quad (\text{A.71})$$

Intuitively, the reduced target state has the larger half of the initial machine eigenvalues in the ground state and the smaller half in the excited state; the reduced machine state has the sum of the largest elements from each of these halves in its ground state, the next largest element from each half (which, in this case, is equal to the first) in its first excited state, and so forth. Let us denote the spectrum of the final state of the machine by $\boldsymbol{\lambda}_{\mathcal{M}}^{\downarrow} := \{\lambda_{\mathcal{M}}^{\downarrow 0}, \lambda_{\mathcal{M}}^{\downarrow 1}, \dots, \lambda_{\mathcal{M}}^{\downarrow d_{\mathcal{M}}-1}\} = \frac{1}{2}\{\lambda_{\mathcal{M}}^{\downarrow 0} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2}}, \lambda_{\mathcal{M}}^{\downarrow 0} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2}}, \dots, \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2}+1} + \lambda_{\mathcal{M}}^{\downarrow d_{\mathcal{M}}-1}\}$. Importantly, by construction, the reduced state of the final machine has its local eigenvalues in non-increasing order, i.e., it is energetically passive.

We therefore have the final reduced states of the protocol that cools the initially maximally-mixed target as much as possible at minimal energy cost, in particular with minimal heat dissipation by the machine, given the structural resources at hand. We can now analyse the properties that are required to saturate the Landauer limit by considering

³Strictly speaking, in the limit $d_{\mathcal{M}} \rightarrow \infty$ the conditions in Eq. (A.67) must only be satisfied for almost all i , i.e., for all but a small subset that contributes negligibly to the relative entropy, as we discuss below.

the terms on the r.h.s. of Eq. (1.8) for any fixed initial inverse temperature of the machine $\beta \geq 0$.

First note that cooling the target system by any amount fixes the change in entropy of the target system, so the first term is irrelevant. The second term concerns the mutual information built up between the target system and machine. In general, this is non-vanishing, although one can achieve any desired amount of cooling without generating such correlations (as per our constructions). Furthermore, in the case where one wants to consider attaining a perfectly cool final state, as we do here, the final reduced state of the target is approximately pure and so $I(S : M)_{\varrho'_{S,M}} \rightarrow 0$. In terms of the reduced states above, this means that $\sum_{i=0}^{\frac{d_M}{2}-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 1$ and $\sum_{i=\frac{d_M}{2}}^{d_M-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 0$, which can only occur if the largest half of energy eigenvalues of the machine, i.e., ω_i for all $i \geq \frac{d_M}{2}$, diverge (since the summation contains only non-negative summands).

The final term that must be minimised to saturate the Landauer limit is the relative entropy of the final with respect to the initial machine state, $D(\varrho'_{\mathcal{M}} \parallel \varrho_{\mathcal{M}})$. Here one can already see that an infinite-dimensional machine is required to saturate the Landauer bound: From Ref. [200], $D(\varrho'_{\mathcal{M}} \parallel \varrho_{\mathcal{M}}) \geq f(\Delta S_{\mathcal{M}}, d_{\mathcal{M}})$, where f is a dimension-dependant function of the entropy difference of the machine that exhibits non-negative correction terms that only vanish in the limit $d_{\mathcal{M}} \rightarrow \infty$. The relative entropy vanishes iff $\varrho_{\mathcal{M}} = \varrho'_{\mathcal{M}}$; moreover, by Pinsker's inequality one has $\frac{1}{2} \|\varrho_{\mathcal{M}} - \varrho'_{\mathcal{M}}\|_1^2 \leq D(\varrho_{\mathcal{M}} \parallel \varrho'_{\mathcal{M}})$, so one can bound the trace distance between the initial and final state of the machine for any desired value of the relative entropy. Although $\varrho_{\mathcal{M}} = \varrho'_{\mathcal{M}}$ implies a trivial process that cannot cool the (initially thermal) target system, as our protocols that saturate the Landauer limit demonstrate, there are processes that asymptotically display the behaviour $\varrho'_{\mathcal{M}} \rightarrow \varrho_{\mathcal{M}}$ and cool the target system. For the asymptotic machine states to converge, in particular their eigenvalues must become approximately equal asymptotically. Demanding this on the spectrum in Eq. (A.71) leads to a generic term that must be asymptotically satisfied of the form:

$$\frac{\frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_M}{2} + \lfloor \frac{i}{2} \rfloor} \right)}{\lambda_{\mathcal{M}}^{\downarrow i}} \rightarrow 1 \quad \forall i \in \{0, \dots, d_M - 1\}. \quad (\text{A.72})$$

□

In order to achieve perfect cooling at the Landauer limit, one thus must simultaneously satisfy the conditions outlined in Theorem A.4. In other words, to minimise the relative entropy term with the additional constraints $\sum_{i=0}^{\frac{d_M}{2}-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 1$ and $\sum_{i=\frac{d_M}{2}}^{d_M-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 0$. The first thing to note is that since the eigenvalues $\lambda_{\mathcal{M}}^{\downarrow i}$ contribute to different sums depending on whether i is in the larger half $\{0, \dots, \frac{d_M}{2} - 1\}$ or smaller half $\{\frac{d_M}{2}, \dots, d_M\}$, one cannot have $\lambda_{\mathcal{M}}^{\downarrow \frac{d_M}{2} + \lfloor \frac{i}{2} \rfloor} = \lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor} \forall i$ (i.e., a completely degenerate machine), since then

both summations would be over identical values and there is no way for them to converge to distinct values. This precludes the trivial solution that satisfies the constraints of Eq. (A.67) alone, namely the maximally mixed machine state which cannot be used to perform any cooling [as, in particular, it does not satisfy the constraints of Eq. (A.66)]. For the conditions to be simultaneously satisfied, we intuitively require that, although they must be distinct, for each i both $\lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor}$ and $\lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2} + \lfloor \frac{i}{2} \rfloor}$ become “close” to each other, but with a difference that decays rapidly as $d_{\mathcal{M}} \rightarrow \infty$, such that in the infinite-dimensional limit the larger “half” of the eigenvalues sum to one and the smaller “half” sum to zero. A subtle point to note is that because the relative entropy involves the ratio of final to original eigenvalues it is not enough that the absolute difference $|\lambda_{\mathcal{M}}^{\downarrow i} - \lambda_{\mathcal{M}}^{\downarrow i}|$ goes to zero, as in the infinite $d_{\mathcal{M}}$ limit, it is possible for this to happen for all of the eigenvalues approaching zero without the ratios of final to initial eigenvalues approaching unity (and hence the relative entropy not vanishing). One manner of satisfying such a constraint, as evidenced by the construction we proceed with next, is for the ratios of final to initial eigenvalues go to unity for all but a small number energy levels, with the population in this exceptional subspace going to zero in the infinite $d_{\mathcal{M}}$ limit (along with the ratios not diverging within said subspace).

The natural question that arises here is whether or not it is possible to satisfy these constraints concurrently. (Note that none of the cooling protocols provided throughout this article use the max-cooling operation, so do not serve as examples.) To this end, we now construct a family of machine Hamiltonians $H_{\mathcal{M}}$ of increasing dimension that in the limit $d_{\mathcal{M}} \rightarrow \infty$ manages to attain both perfect cooling of a maximally-mixed qubit and the Landauer limit for the energy cost using the maximal cooling operation discussed above. The form of the Hamiltonian is instructive regarding the complexity requirements for perfect cooling at the Landauer limit. The construction is inspired by the infinite-dimensional Hamiltonian found in Ref. [200] (Appendix D), therein used to perfectly cool a qubit with energy cost arbitrarily close to the Landauer limit. Their construction already begins with infinitely many machine eigenvalues, as well as infinitely many of them corresponding to diverging energy levels. In the following, we demonstrate that one can arbitrarily closely attain perfect cooling and the Landauer limit with finite-dimensional Hamiltonians, and by taking the limit $d_{\mathcal{M}} \rightarrow \infty$, recover the result of Ref. [200].

The Hamiltonian of the machine is $d_{\mathcal{M}} := 2^{N+1}$ dimensional,

$$H_{\mathcal{M}} = \sum_{n=0}^N \sum_{j=1}^{2^n} \left(n\Delta |n; j\rangle \langle n; j|_{\mathcal{M}} \right) + N\Delta |N; 2^N+1\rangle \langle N; 2^N+1|_{\mathcal{M}}. \quad (\text{A.73})$$

Here, each energy eigenvalue labelled by n is 2^n -fold degenerate. Thus the ground state is unique, the first excited state is two-fold degenerate, the second excited state four-fold degenerate, and so on, with the degeneracy doubling every energy level. In order to make the Hamiltonian of even dimensionality for convenience, we have added an extra

degenerate state to the final level (which makes this level $(2^N + 1)$ -fold degenerate). Also note that the Hamiltonian is equally spaced with energy gap Δ . In the following, we use the index n to denote any one of the degenerate states in the n^{th} energy level from $n = 0$ to $n = N$, and the index i to denote individual energy eigenstates from $i = 1$ to $i = 2^{N+1}$ (note that in contrast to the previous section, we are here beginning with $i = 1$ in order to simplify some future notation). With these indices, the eigenvalues are related by

$$\lambda_{\mathcal{M}}^{\downarrow i} = e^{-\beta\Delta} \lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor} \quad \forall i \in \{2, \dots, d_{\mathcal{M}} - 1\}, \quad (\text{A.74})$$

$$\lambda_{\mathcal{M}}^{\downarrow n} = e^{-\beta\Delta} \lambda_{\mathcal{M}}^{\downarrow n-1} \quad \forall n \in \{1, \dots, N\}. \quad (\text{A.75})$$

We introduce a parameter ϵ to express the Gibbs ratio as

$$e^{-\beta\Delta} = \frac{1 - \epsilon}{2}, \quad (\text{A.76})$$

where $0 < \epsilon < 1$, and we will eventually take the limit $\epsilon \rightarrow 0$ appropriately as the dimension diverges. Note that this constrains the Gibbs ratio to be smaller than $\frac{1}{2}$, which in turn ensures that the total population over all of the degenerate eigenstates in the n^{th} level is smaller than that in the $(n - 1)^{\text{th}}$ level (as it has twice the number of eigenstates, but less than half the population in each). If this constraint failed to hold, then in the asymptotic limit, all of the population would lie in energy levels that diverge.

We now consider using this machine to cool a maximally-mixed qubit target. The final ground-state population of the qubit under the maximal cooling operation is the sum over the larger half of the eigenvalues of the machine, corresponding to the eigenvalues from $i = 1$ to $i = 2^N$ (equivalently, from $n = 0$ to $n = N - 1$ plus a single eigenvalue from the $n = N$ energy level), and is thus given by

$$p'_0 = \frac{1}{\mathcal{Z}_{\mathcal{M}}} \left(\sum_{n=0}^{N-1} 2^n \left(\frac{1 - \epsilon}{2} \right)^n + \left(\frac{1 - \epsilon}{2} \right)^N \right), \quad (\text{A.77})$$

$$\text{where } \mathcal{Z}_{\mathcal{M}} = \sum_{n=0}^N 2^n \left(\frac{1 - \epsilon}{2} \right)^n + \left(\frac{1 - \epsilon}{2} \right)^N \quad (\text{A.78})$$

is the partition function of the machine. The geometric series above evaluates to

$$p'_0 = \left(1 + \frac{\epsilon(1 - \epsilon)^N}{1 - (1 - \epsilon)^N + \epsilon(1 - \epsilon)^N 2^{-N}} \right)^{-1}. \quad (\text{A.79})$$

As an ansatz, supposing that ϵ scales inversely with N as $\epsilon := \frac{\theta}{N}$ leads to the simplification $(1 - \epsilon)^N \rightarrow e^{-\theta}$ as $d_{\mathcal{M}}$ (and hence N) diverges. The asymptotic behaviour of the ground-state population is thus

$$p'_0 = 1 - \frac{1}{N} \left(\frac{\theta}{e^{\theta} - 1} \right) + O\left(\frac{1}{N^2}\right), \quad (\text{A.80})$$

and so $p'_0 \rightarrow 1$ in the $N \rightarrow \infty$ limit.

We now move to calculate the energy cost. Rather than considering the optimal max-cooling operation described above, we consider a slight modification in order to make the connection to the construction in Ref. [200] clear as well as to simplify notation. Nonetheless, the energy cost of this modified protocol upper-bounds that of the max-cooling operation (for the same achieved ground-state population), and so showing that the Landauer limit is attained for the modified protocol implies that it would be too for the max-cooling protocol. The modification is simply to relabel the smallest eigenvalue of the machine $\lambda_{\mathcal{M}}^{2^{N+1}}$ as $\lambda_{\mathcal{M}}^0$, and treat it as the ground-state eigenvalue in the max-cooling operation. For general machine states, such a switch would lead to less cooling (if the same unitary were applied), but in this case it does not because the sum of the first half of the machine eigenvalues, from $i = 0$ to $i = 2^N - 1$, is the same as the original sum from $i = 1$ to $i = 2^N$, due to the relabelling $\lambda_0 = \lambda_{2^N}$, since they are both eigenvalues of states corresponding to the maximum excited energy level of the machine spectrum. The spectrum of the final state of the machine is then given by

$$\lambda_{\mathcal{M}}^{\uparrow i} = \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor} + \lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor + \frac{d_{\mathcal{M}}}{2}} \right) \quad \forall i \in \{0, \dots, d_{\mathcal{M}} - 1\}, \quad (\text{A.81})$$

which leads to

$$\begin{aligned} \lambda_{\mathcal{M}}^{\uparrow 0} &= \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow 0} + \lambda_{\mathcal{M}}^{\downarrow 2^N} \right) = \lambda_{\mathcal{M}}^{\downarrow 0}, & \lambda_{\mathcal{M}}^{\uparrow 1} &= \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow 0} + \lambda_{\mathcal{M}}^{\downarrow 2^N} \right) = \lambda_{\mathcal{M}}^{\downarrow 0}, \\ \lambda_{\mathcal{M}}^{\uparrow i} &= \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor} + \lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor + \frac{d_{\mathcal{M}}}{2}} \right) & \forall i \in \{2, \dots, d_{\mathcal{M}} - 1\} \\ &= \frac{1}{2} \left(\frac{2}{1 - \epsilon} \lambda_{\mathcal{M}}^{\downarrow i} + \lambda_{\mathcal{M}}^{n=N} \right) = \frac{1}{2} \frac{1}{\mathcal{Z}_{\mathcal{M}}} \left[\left(\frac{1 - \epsilon}{2} \right)^{n-1} + \left(\frac{1 - \epsilon}{2} \right)^N \right], \end{aligned} \quad (\text{A.82})$$

where we observe that the index $\lfloor \frac{i}{2} \rfloor + \frac{d_{\mathcal{M}}}{2}$ corresponds to the largest energy level of the machine for all i , and we have used Eq. (A.74) for the spectrum of initial eigenvalues. Using the index n instead to denote a generic eigenvalue of the n^{th} energy level, we have the simpler expression

$$\lambda_{\mathcal{M}}^{\uparrow n} = \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow n-1} + \lambda_{\mathcal{M}}^{\downarrow N} \right) \quad \forall n \in \{1, 2, \dots, N\}. \quad (\text{A.83})$$

The energy cost can now be simply calculated from the difference in the average energy of the machine state,

$$\Delta E_{\mathcal{M}} = \sum_{i=0}^{d_{\mathcal{M}}-1} \left(\lambda_{\mathcal{M}}^{\uparrow i} - \lambda_{\mathcal{M}}^{\downarrow i} \right) \omega_i, \quad (\text{A.84})$$

where we denote the i^{th} energy eigenvalue by ω_i . $\lambda_{\mathcal{M}}^{\downarrow 0}$ is unchanged, and although $\lambda_{\mathcal{M}}^{\downarrow 1}$ does change, $\omega_1 = 0$ corresponds to the ground state and thus this eigenvalue change does not affect the energy cost. We can thus express the energy cost in terms of the index n

instead, starting from $n = 1$ (corresponding to $i = 2$ onward), as

$$\begin{aligned}\Delta E_{\mathcal{M}} &= \sum_{n=1}^N (\lambda_{\mathcal{M}}^{\downarrow n} - \lambda_{\mathcal{M}}^{\uparrow n}) \omega_n \\ &= \frac{1}{\beta} \left[1 - \frac{2(1-\epsilon)^N}{1 - (1-\epsilon)^N + (1-2^{-N})(1-\epsilon)^N \epsilon} \right] \log\left(\frac{2}{1-\epsilon}\right).\end{aligned}\quad (\text{A.85})$$

As we did above, we parameterise $\epsilon = \frac{\theta}{N}$. The asymptotic behaviour of the energy cost is then

$$\beta \Delta E_{\mathcal{M}} = \log(2) + \frac{1}{N} \left(1 - \frac{2 \log(2)}{e^{\theta} - 1} \right) \theta + O\left(\frac{1}{N^2}\right), \quad (\text{A.86})$$

or in terms of the decrease in entropy of the system,

$$\beta \Delta E_{\mathcal{M}} = \tilde{\Delta} S_{\mathcal{M}} + \frac{\log N}{N} \left(\frac{\theta}{e^{\theta} - 1} \right) + O\left(\frac{1}{N}\right). \quad (\text{A.87})$$

Combining (A.80) and (A.86), we have that in the limit $N \rightarrow \infty$, which is also $d_{\mathcal{M}} \rightarrow \infty$, the ground state population approaches 1—corresponding to perfect cooling—and the energy cost approaches $\beta^{-1} \log(2)$, which is the Landauer limit for the perfect erasure of a maximally-mixed qubit.

To connect this construction to the constraints of Eq. (A.67), note that in the limit $N \rightarrow \infty$ (recalling that $\epsilon = \frac{\theta}{N}$),

$$\frac{\lambda_{\mathcal{M}}^{\downarrow n}}{\lambda_{\mathcal{M}}^{\uparrow n}} = \lim_{N \rightarrow \infty} \frac{\frac{1}{2} \left[\left(\frac{1-\epsilon}{2}\right)^{n-1} + \left(\frac{1-\epsilon}{2}\right)^N \right]}{\left(\frac{1-\epsilon}{2}\right)^n} = \lim_{N \rightarrow \infty} \left[\frac{1}{1-\epsilon} + \frac{1}{2^{N-n+1}} \left(\frac{e^{-\theta}}{(1-\epsilon)^n} \right) \right] = 1, \quad (\text{A.88})$$

for all $n \geq 1$, leaving only the ground state eigenvalue (corresponding to $n = 0$ and $i = 1$) not satisfying the condition. However this term is actually a negative contribution to the relative entropy as this eigenvalue decreases, and in any case can be verified independently to approach zero.

To see this, note that a necessary condition that ensures the contribution of any set of eigenvalues that do not satisfy Eq. (A.67) to the relative entropy to be negligible is that the total population of the relevant subspace is vanishingly small. Writing the relative entropy between two states in terms of their eigenvalues, we have $D(\varrho' \parallel \varrho) = \sum_n \lambda'_n \log\left(\frac{\lambda'_n}{\lambda_n}\right)$, which we split up into two sets: S_0 containing all n for which Eq. (A.67) is satisfied and S_{\pm} containing the all n for which Eq. (A.67) is not satisfied. The contribution of the first term to the relative entropy is asymptotically zero, so we are left with $D(\varrho' \parallel \varrho) = \sum_{n \in S_{\pm}} \lambda'_n \log\left(\frac{\lambda'_n}{\lambda_n}\right)$. For each term in the sum here, one can write $\lambda_n = \lambda'_n (1 + \Delta_n)$ with the condition $|\Delta_n| \geq \theta > 0$ for some θ , i.e., the ratio of eigenvalues is bounded away from unity (on either side) by at least θ . This leads to the expression

$$D(\varrho' \parallel \varrho) = - \sum_{n \in S_{\pm}} \lambda'_n \log(1 + \Delta_n) = -N_{\pm} \sum_{n \in S_{\pm}} p_n \log(1 + \Delta_n), \quad (\text{A.89})$$

where we have renormalised the eigenvalues (which here correspond to a subnormalised probability distribution) by writing $\lambda'_n = N_{\pm} p_n$, with $N_{\pm} := \sum_{n \in S_{\pm}} \lambda'_n$ being the total population of the subspace S_{\pm} and $\{p_n\}$ here forming a probability distribution. Note that the ratio of eigenvalues going to unity in the S_0 subspace implies that the total populations of initial and final eigenvalues in this subspace are equal, i.e., $\sum_{n \in S_0} \lambda_n = \sum_{n \in S_0} \lambda'_n$, which in turn implies that the same is true for the S_{\pm} subspace, leading to $\sum_{n \in S_{\pm}} p_n \Delta_n = 0$.

We argue from the concavity of the logarithm function that

$$\frac{1}{2} \log(1 + \theta) + \frac{1}{2} \log(1 - \theta) \geq \sum_{n \in S_{\pm}} p_n \log(1 + \Delta_n). \quad (\text{A.90})$$

Visualising the graph of the function $y = \log(1 + x)$, the latter expression above must evaluate to a point that lies within the intersection of the convex hull of $(\Delta_n, \log(1 + \Delta_n))$ and the linear equality $\sum_{n \in S_{\pm}} p_n \Delta_n = 0$, the latter of which is the line $x = 0$. By the concavity of the logarithm, the aforementioned convex hull lies entirely below the line segment connecting $(1 - \theta, \log(1 - \theta))$ to $(1 + \theta, \log(1 + \theta))$, and thus the expression is upper bounded by the intersection of this line segment with $x = 0$, which is precisely the l.h.s. of the inequality above. Thus we have the inequality

$$D(\varrho' \parallel \varrho) \geq -N_{\pm} \left[\frac{1}{2} \log(1 + \theta) + \frac{1}{2} \log(1 - \theta) \right] = -\frac{N_{\pm}}{2} \log(1 - \theta^2) \geq \frac{N_{\pm}}{2} \theta^2, \quad (\text{A.91})$$

where we used $\log(1 - \theta^2) \leq -\theta^2$ for all $\theta \in [-1, 1]$. As $\theta > 0$, the only way that this contribution to the relative entropy by the eigenvalues that do not satisfy Eq. (A.67) can be asymptotically negligible is if the total population of their associated subspace N_{\pm} goes to zero.

Finally note that, as mentioned in the main text, the above result pertains to the restricted setting where the target system is cooled as much as possible. However, this is not the only way to approach perfect cooling at the Landauer cost: Instead of the largest half of global eigenvalues being placed into the ground state subspace of the target system, any amount of them such that their sum is sufficiently close to one would suffice. Although it is complicated to derive an exact set of conditions that would need to be satisfied in such cases (since it depends upon exactly which eigenvalues are permuted to which subspaces), the fact that fine-tuned control over particular degrees of freedom is required remains. Lastly, note that even in the restricted setting of cooling the target as much as possible, the situation becomes even more complicated when considering target systems that begin at a finite temperature. Here, the choice of which global eigenvalues should be permuted to which subspaces to cool the system as much as possible at minimal energy cost depends on the microscopic structure of both the system and machine. This means that one can no longer determine the final eigenvalue distributions of the reduced states in terms of the initial machine eigenvalues alone, as we were able to do for the maximally mixed state. In turn, one can no longer derive a condition on properties

of the machine itself, independently of the target system. Nonetheless, again, the key message that cooling at minimal energy cost requires fine-tuned control to access precisely distributed populations still holds true. We leave the further exploration of such scenarios, for instance constructing optimal machines for particular initial target systems, to future work.

A.5 Diverging Time / Control Complexity Cooling Protocols for Harmonic Oscillators

We now analyse the case of cooling infinite-dimensional quantum systems in detail. More specifically, we consider ensembles of harmonic oscillators. For the sake of completeness, we first briefly present some key concepts that will become relevant throughout this analysis. Following this, in Appendix A.5.2, we construct a protocol that achieves perfect cooling at the Landauer limit using a diverging number of Gaussian operations. Although such operations are typically considered to be relatively “simple” both when it comes to experimental implementation and theoretical description, according to the effective dimension notion of control complexity that we have shown must necessarily diverge to cool at the Landauer limit [see Eq. (1.3)], such Gaussian operations have infinite control complexity. Subsequently, in Appendix A.5.2, we consider the task of perfect cooling with diverging time but restricting the individual operations to be of finite control complexity. In particular, note that such operations are non-Gaussian in general. Here, we present a protocol that approaches perfect cooling of the target system as the number of operations diverges, with finite energy cost—albeit not at the Landauer limit. Whether or not a similar protocol exists that also saturates the Landauer bound remains an open question. Finally, in Appendix A.5.3, we reconsider the protocol from Appendix A.5.2 in terms of a single transformation, i.e., unit time. By explicitly constructing the joint unitary transformation that is applied throughout the entire protocol, we show this to be a multi-mode Gaussian operation acting on a diverging number of harmonic oscillators. The key message to be taken away from these protocols is that, while the distinction between Gaussian and non-Gaussian operations is a significant one in terms of experimental feasibility, and it certainly plays a role regarding the task of cooling—in particular, the energy cost incurred—these concepts alone cannot be used to characterise a notion of control complexity that must diverge to approach perfect cooling at the Landauer limit. On the other hand, the effective dimension of the machine used does precisely that; however, in a manner that is far from sufficient (for the case of harmonic oscillators), as even a single two-mode swap, which cannot cool perfectly at Landauer cost, would have infinite control complexity. Indeed, a more nuanced characterisation of control complexity in the infinite-dimensional setting, which takes more structure regarding the operations and energy levels into account, remains an open problem to be addressed.

A.5.1 Preliminaries

We consider ensembles of N harmonic oscillators (i.e., infinite-dimensional systems consisting of N bosonic modes), which are associated to a tensor product Hilbert space $\mathcal{H}_{\text{tot}} = \bigotimes_{j=1}^N \mathcal{H}_j$ and (lowering, raising [respectively]) mode operators $\{a_k, a_k^\dagger\}$ satisfying the bosonic commutation relations:

$$[a_k, a_{k'}^\dagger] = \delta_{kk'}, \quad [a_k, a_{k'}] = 0, \quad \forall k, k' = 1, 2, \dots, N. \quad (\text{A.92})$$

The free Hamiltonian of any such system can be written as $H_{\text{tot}} = \sum_{k=1}^N \omega_k a_k^\dagger a_k$, where ω_k represents the energy gap of the k -th mode (in units where $\hbar = 1$). Position- and momentum-like operators for each mode can be defined as follows (for simplicity, we use the rescaled version below where the ω_k are omitted from the pre-factors)

$$q_k := \frac{1}{\sqrt{2}}(a_k + a_k^\dagger), \quad p_k := \frac{1}{i\sqrt{2}}(a_k - a_k^\dagger). \quad (\text{A.93})$$

As a consequence of the commutation relations in Eq. (A.92), the generalised position and momentum operators satisfy the canonical commutation relations

$$[q_k, p_l] = i\delta_{kl}. \quad (\text{A.94})$$

To simplify notation, one may further introduce the vector of quadrature operators $\mathbb{X} := (q_1, p_1, \dots, q_N, p_N)$; then, the commutation relations can be expressed succinctly as

$$[\mathbb{X}_k, \mathbb{X}_l] = i\Omega_{kl}, \quad (\text{A.95})$$

where the Ω_{kl} are the components of the symplectic form

$$\Omega = \bigoplus_{j=1}^N \Omega_j, \quad \Omega_j = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}. \quad (\text{A.96})$$

The density operator associated to N harmonic oscillators can be written in the so-called *phase-space representation* as

$$\varrho = \frac{1}{(2\pi)^N} \int \chi(\Omega\xi) \mathcal{W}(-\Omega\xi) d^{2N}\xi, \quad (\text{A.97})$$

where $\mathcal{W}(\xi) := e^{i\xi^T \mathbb{X}}$ is the Weyl operator and $\chi(\xi) := \text{tr}[\varrho \mathcal{W}(\xi)]$ is called the characteristic function.

Throughout our analysis, we will see that a particular class of states and operations, namely those that are known as *Gaussian*, are of particular importance. A Gaussian state is one for which the characteristic function is Gaussian

$$\chi(\xi) = e^{-\frac{1}{4}\xi^T \Gamma \xi + i\bar{\mathbb{X}}^T \xi}. \quad (\text{A.98})$$

Here, $\bar{\mathbb{X}} := \langle \mathbb{X} \rangle_{\varrho}$ is the *displacement vector* or *vector of first moments*, and Γ is a real symmetric matrix that collects the *second statistical moments* of the quadratures, which is known as the *covariance matrix*. Its entries are given by

$$\Gamma_{mn} := \langle \mathbb{X}_m \mathbb{X}_n + \mathbb{X}_n \mathbb{X}_m \rangle_{\varrho} - 2 \langle \mathbb{X}_n \rangle_{\varrho} \langle \mathbb{X}_m \rangle_{\varrho}. \quad (\text{A.99})$$

We see that any Gaussian state is thus uniquely determined by its first and second moments. As an example of specific interest here, we recall that any thermal state τ of a harmonic oscillator with frequency ω is a Gaussian state and has vanishing first moments, $\bar{\mathbb{X}} = 0$. Here and throughout this article, we are assuming that the infinite-dimensional thermal state is well-defined (see, e.g., Ref. [363] for discussion). The covariance matrix of a thermal state is proportional to the 2×2 identity, and given by $\Gamma[\tau(\beta, H)] = \coth\left(\frac{\beta\omega}{2}\right) \mathbb{1}_2$.

Gaussian operations are transformations that map the set of Gaussian states onto itself. Such operations, which include, e.g., beam-splitting and phase-space displacement, are generally considered to be relatively easily implementable in the laboratory. Although non-unitary Gaussian operations exist as well, all of the examples mentioned above are Gaussian unitaries. Such Gaussian unitaries are generated by Hamiltonians that are at most quadratic in the raising and lowering operators. Conversely, any Hamiltonian that can be expressed as a polynomial of at most second order in the mode operators generates a Gaussian unitary. Any unitary Gaussian transformation can be represented by an affine map (M, κ) ,

$$\mathbb{X} \mapsto M\mathbb{X} + \kappa, \quad (\text{A.100})$$

where $\kappa \in \mathbb{R}^{2N}$ is a displacement vector in the phase-space representation and M is a symplectic $2N \times 2N$ matrix that leaves the symplectic form Ω invariant, i.e.,

$$M \Omega M^T = \Omega. \quad (\text{A.101})$$

Under such a mapping, the first and second moments transform according to

$$\bar{\mathbb{X}} \mapsto M\bar{\mathbb{X}} + \kappa, \quad \Gamma \mapsto M\Gamma M^T. \quad (\text{A.102})$$

Lastly, note that the energy of a Gaussian state ϱ_G with respect to its free Hamiltonian $H = \sum_k \omega_k a_k^\dagger a_k$ can be calculated in terms of the first and second moments as follows [54]

$$E(\varrho_G) = \sum_k \omega_k \left(\frac{1}{4} \text{tr} [\Gamma^{(k)} - 2] + \frac{1}{2} \|\bar{\mathbb{X}}^{(k)}\|^2 \right), \quad (\text{A.103})$$

where $\|\cdot\|$ denotes the Euclidean norm. Here, $\Gamma^{(k)}$ indicates the (2×2) -submatrix of the full covariance matrix Γ corresponding to the reduced state of the k^{th} mode. Similarly $\bar{\mathbb{X}}^{(k)}$ denotes the two-component subvector of first moments for the k^{th} mode of the displacement vector $\bar{\mathbb{X}}$.

A.5.2 Diverging Time Cooling Protocol for Harmonic Oscillators

Diverging Time Protocol using Gaussian Operations (with Diverging Control Complexity)

We now consider a simple protocol for lowering the temperature of a single-mode system within the coherent-control paradigm using a single harmonic oscillator machine. This protocol will form the basic step of a protocol for achieving perfect cooling at the Landauer limit using diverging time, which we subsequently present.

In the situation we consider here, the target system \mathcal{S} to be cooled is a harmonic oscillator with frequency $\omega_{\mathcal{S}}$ interacting with a harmonic oscillator machine \mathcal{M} at frequency $\omega_{\mathcal{M}} \geq \omega_{\mathcal{S}}$ via a (non-energy-conserving) unitary acting on the joint system \mathcal{SM} initialised as a tensor product of thermal states $\tau_{\mathcal{S}}(\beta, H_{\mathcal{S}}) \otimes \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})$ at inverse temperature β with respect to their local Hamiltonians $H_{\mathcal{S}}$ and $H_{\mathcal{M}}$, respectively. The joint covariance matrix of the system and machine modes is block-diagonal since the initial state is of product form, i.e.,

$$\Gamma[\tau_{\mathcal{S}}(\beta, H_{\mathcal{S}}) \otimes \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})] = \Gamma[\tau_{\mathcal{S}}(\beta, H_{\mathcal{S}})] \oplus \Gamma[\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})], \quad (\text{A.104})$$

and the 2×2 blocks of the individual modes are also diagonal, with the explicit expression $\Gamma[\tau_X(\beta, H_X)] = \coth\left(\frac{\beta\omega_X}{2}\right) \mathbb{1}_2$.

In this setting, it has been shown that the minimum reachable temperature of the target system is given by $T_{\min} = \frac{\omega_{\mathcal{S}}}{\omega_{\mathcal{M}}} T$ (for the case $\omega_{\mathcal{M}} \geq \omega_{\mathcal{S}}$) [83]. The non-energy-conserving unitary transformation that achieves this is of the form

$$U = e^{-i\frac{\pi}{2}(a^\dagger b + ab^\dagger)}, \quad (\text{A.105})$$

where the operators a (a^\dagger) and b (b^\dagger) denote the annihilation (creation) operators of the target system and machine, respectively. This beam-splitter-like unitary acts as a swap with a relative phase factor imparted on the resultant state; nonetheless, this phase is irrelevant at the level of the covariance matrix, which fully characterises the (Gaussian) thermal states considered, and transforms it according to a standard swapping of the systems. After acting with such a swap operator, which is a Gaussian operation, the first moment remains vanishing and the covariance matrix transforms as [see Eq. (A.102)]

$$\begin{bmatrix} \coth\left(\frac{\beta\omega_{\mathcal{S}}}{2}\right) \mathbb{1}_2 & 0 \\ 0 & \coth\left(\frac{\beta\omega_{\mathcal{M}}}{2}\right) \mathbb{1}_2 \end{bmatrix} \xrightarrow{\text{SWAP}} \begin{bmatrix} \coth\left(\frac{\beta\omega_{\mathcal{M}}}{2}\right) \mathbb{1}_2 & 0 \\ 0 & \coth\left(\frac{\beta\omega_{\mathcal{S}}}{2}\right) \mathbb{1}_2 \end{bmatrix}. \quad (\text{A.106})$$

This means that both the output target system and machine are thermal states at different temperatures $T'_{\mathcal{S}} = \frac{\omega_{\mathcal{S}}}{\omega_{\mathcal{M}}} T$ and $T'_{\mathcal{M}} = \frac{\omega_{\mathcal{M}}}{\omega_{\mathcal{S}}} T$. Making use of Eq. (A.103), we can calculate

the energy change for the system and machine as

$$\begin{aligned}
 \Delta E_S &= E \left[\tau_S \left(\frac{\omega_M}{\omega_S} \beta, H_S \right) \right] - E \left[\tau_S(\beta, H_S) \right] \\
 &= \frac{\omega_S}{2} \left[\coth \left(\frac{\beta \omega_M}{2} \right) - \coth \left(\frac{\beta \omega_S}{2} \right) \right], \\
 \Delta E_M &= E \left[\tau_M \left(\frac{\omega_S}{\omega_M} \beta, H_M \right) \right] - E \left[\tau_M(\beta, H_M) \right] \\
 &= \frac{\omega_M}{2} \left[\coth \left(\frac{\beta \omega_S}{2} \right) - \coth \left(\frac{\beta \omega_M}{2} \right) \right].
 \end{aligned} \tag{A.107}$$

The total energy cost associated to such a swap operation is thus

$$\begin{aligned}
 \Delta E_{SM} &= \Delta E_S + \Delta E_M = \frac{(\omega_M - \omega_S)}{2} \left[\coth \left(\frac{\beta \omega_S}{2} \right) - \coth \left(\frac{\beta \omega_M}{2} \right) \right] \\
 &= (\omega_M - \omega_S) \frac{e^{-\beta \omega_S} (1 - e^{-\beta(\omega_M - \omega_S)})}{(1 - e^{-\beta \omega_S})(1 - e^{-\beta \omega_M})}.
 \end{aligned} \tag{A.108}$$

Note that this form is similar to that for finite-dimensional systems with equally spaced Hamiltonian [cf., Eq. (A.34)]; the dimension-dependent term vanishes as $d \rightarrow \infty$, simplifying the expression in the infinite-dimensional case.

With this simple protocol for lowering the temperature of a harmonic oscillator target using a single harmonic oscillator machine at hand, we are now in a position to describe an energy-optimal (in the sense of saturating the Landauer bound) cooling protocol when a diverging number of operations, i.e., diverging time, is permitted. In other words, we now show how to achieve perfect cooling with minimal energy at the expense of requiring diverging time, i.e., infinitely many steps of finite duration. As mentioned above, in this specific protocol, the control complexity as per Eq. (1.3) is infinite in each of these infinitely many steps. As we will argue after having presented the protocol, this is an artefact of the simple structure of the Gaussian operations used. Indeed, we will later present a non-Gaussian diverging-time protocol for cooling a single harmonic oscillator to the ground state using finite control complexity in each of the infinitely many steps, and at an overall finite (albeit not minimal, i.e., not at the Landauer limit) energy cost. Before presenting this non-Gaussian protocol, let us now discuss the details of the Gaussian diverging-time protocol for cooling at the Landauer limit.

We consider a harmonic oscillator with the frequency ω_S as the target system and the machine to comprise N harmonic oscillators, where the n^{th} oscillator has frequency $\omega_{M_n} = \omega_S + n\epsilon$. In addition, we assume that all modes are initially uncorrelated and in thermal states at the same inverse temperature β with respect to their free Hamiltonians, i.e., the target system is $\tau_S(\beta, H_S)$ and the multi-mode thermal machine is $\tau_M(\beta, H_M) = \bigotimes_{n=1}^N \tau_{M_n}(\beta, H_{M_n})$.

In this case, the cooling process is divided into N time steps. During each step, there is an interaction between the target system and one of the harmonic oscillators in the

machine. Here, we assume that at the n^{th} time step, the target system interacts only with the n^{th} harmonic oscillator, which has frequency $\omega_S + n\epsilon$. To obtain the minimum temperature for the target system, we perform the previously outlined cooling process at each step, which is given by swapping the corresponding two modes. Using Eq. (A.106), the covariance matrix transformation of the two-mode process at the first time step takes the form

$$\begin{aligned} \Gamma^{(1)}(\tau_S(\beta) \otimes \tau_{\mathcal{M}_1}(\beta)) &= \begin{bmatrix} \coth\left(\frac{\beta\omega_S}{2}\right) \mathbb{1}_2 & 0 \\ 0 & \coth\left(\frac{\beta(\omega_S+\epsilon)}{2}\right) \mathbb{1}_2 \end{bmatrix} \\ \xrightarrow{\text{SWAP}} \Gamma_{\text{opt}}^{(1)} &= \begin{bmatrix} \coth\left(\frac{\beta(\omega_S+\epsilon)}{2}\right) \mathbb{1}_2 & 0 \\ 0 & \coth\left(\frac{\beta\omega_S}{2}\right) \mathbb{1}_2 \end{bmatrix}. \end{aligned} \quad (\text{A.109})$$

By repeating this process on each of the harmonic oscillators in the machine, after the $(n-1)^{\text{th}}$ step, the 2×2 -block corresponding to the target system \mathcal{S} in the covariance matrix is given by $\coth\left(\frac{\beta(\omega_S+(n-1)\epsilon)}{2}\right) \mathbb{1}_2$. Therefore, one can show inductively that the covariance matrix transformation associated to the n^{th} interaction is given by

$$\begin{aligned} \Gamma^{(n)}(\tau_S(\beta) \otimes \tau_{\mathcal{M}_n}(\beta)) &= \begin{bmatrix} \coth\left(\frac{\beta(\omega_S+(n-1)\epsilon)}{2}\right) \mathbb{1}_2 & 0 \\ 0 & \coth\left(\frac{\beta(\omega_S+n\epsilon)}{2}\right) \mathbb{1}_2 \end{bmatrix} \\ \xrightarrow{\text{SWAP}} \Gamma_{\text{opt}}^{(n)} &= \begin{bmatrix} \coth\left(\frac{\beta(\omega_S+n\epsilon)}{2}\right) \mathbb{1}_2 & 0 \\ 0 & \coth\left(\frac{\beta(\omega_S+(n-1)\epsilon)}{2}\right) \mathbb{1}_2 \end{bmatrix}. \end{aligned} \quad (\text{A.110})$$

Based on this process, after N steps (i.e., after the system has interacted with all N harmonic oscillators), the minimal achievable temperature of the target system is $T_{\text{min}}^{(N)} = \frac{\omega_S}{\omega_S+N\epsilon}T$. Moreover, by using Eq. (A.107), one can calculate the energy changes of the target system and the machine at each time step as

$$\begin{aligned} \Delta E_S^{(n)} &= \frac{\omega_S}{2} \left[\coth\left(\frac{\beta(\omega_S+n\epsilon)}{2}\right) - \coth\left(\frac{\beta(\omega_S+(n-1)\epsilon)}{2}\right) \right], \\ \Delta E_{\mathcal{M}_n}^{(n)} &= \frac{(\omega_S+n\epsilon)}{2} \left[\coth\left(\frac{\beta(\omega_S+(n-1)\epsilon)}{2}\right) - \coth\left(\frac{\beta(\omega_S+n\epsilon)}{2}\right) \right]. \end{aligned} \quad (\text{A.111})$$

The total energy change for the target system during the overall process (i.e., throughout the N steps) is thus given by

$$\begin{aligned} \Delta E_S &= \sum_{n=1}^N \Delta E_S^{(n)} = \sum_{n=1}^N \frac{\omega_S}{2} \left[\coth\left(\frac{\beta(\omega_S+n\epsilon)}{2}\right) - \coth\left(\frac{\beta(\omega_S+(n-1)\epsilon)}{2}\right) \right] \\ &= \frac{\omega_S}{2} \left[\coth\left(\frac{\beta(\omega_S+N\epsilon)}{2}\right) - \coth\left(\frac{\beta\omega_S}{2}\right) \right] \\ &= \omega_S \left[\frac{e^{-\beta(\omega_S+N\epsilon)}}{1-e^{-\beta(\omega_S+N\epsilon)}} - \frac{e^{-\beta\omega_S}}{1-e^{-\beta\omega_S}} \right]. \end{aligned} \quad (\text{A.112})$$

Here, we have written $\coth(x) = 1 + \frac{2e^{-2x}}{1-e^{-2x}}$. Similarly, one can obtain the total energy change of the overall machine

$$\begin{aligned}\Delta E_{\mathcal{M}} &= \sum_{n=1}^N \Delta E_{\mathcal{M}_n}^{(n)} = \sum_{n=1}^N \frac{\omega_S + n\epsilon}{2} \left[\coth\left(\frac{\beta(\omega_S + (n-1)\epsilon)}{2}\right) - \coth\left(\frac{\beta(\omega_S + n\epsilon)}{2}\right) \right] \\ &= \sum_{n=1}^N (\omega_S + n\epsilon) \left[\frac{e^{-\beta(\omega_S + (n-1)\epsilon)}}{1 - e^{-\beta(\omega_S + (n-1)\epsilon)}} - \frac{e^{-\beta(\omega_S + n\epsilon)}}{1 - e^{-\beta(\omega_S + n\epsilon)}} \right].\end{aligned}\quad (\text{A.113})$$

It is straightforward to check that the total energy change, i.e., the sum of Eqs. (A.112) and (A.113), is equal to the energy cost obtained in Eq. (A.34) with $d \rightarrow \infty$. In particular, this can be seen by considering the second line of Eq. (A.34), where the second term in round parenthesis vanishes as $d \rightarrow \infty$ for any value of N . Thus, when the number of operations diverges $N \rightarrow \infty$ and $\epsilon = \frac{\omega_{\max} - \omega_S}{N} \rightarrow 0$, where $\omega_{\max} := \frac{\beta_{\max}}{\beta} \omega_S$ is the maximum frequency of the machines, the heat dissipated by the machines throughout the process saturates the Landauer bound and is therefore energetically optimal. Moreover, by taking $\omega_{\max} \rightarrow \infty$ one approaches perfect cooling.

At this point, a comment on the notion of control complexity is in order. According to Eq. (1.3), the effective dimension of the machine in the protocol we consider here diverges in addition to time. Indeed, the notion of control complexity thusly defined diverges for any Gaussian operation acting on the machine, in particular, it diverges for any single one of the infinitely many steps of the protocol, as each operation is a two-mode Gaussian operation. At first glance, this appears to be in contrast to the common conception that Gaussian operations are typically easily implementable (cf. Refs. [54, 209]). However, an alternative way of interpreting this protocol is that, exactly because of the simple structure of Gaussian operations, reaching the ground state at finite energy cost requires a diverging number of two-mode Gaussian unitaries, and thus divergingly many modes on which to act (see also Appendix A.5.3). In fact, if non-Gaussian unitaries are employed, then the ground state can be approached at finite energy cost using just a single harmonic oscillator machine, as we now show.

Diverging Time Protocol using Non-Gaussian Operations (with Finite Control Complexity)

We now consider a protocol for cooling a single harmonic oscillator at frequency ω_S to the ground state using a diverging amount of time, but requiring only a finite overall energy input as well as finite control complexity in each of the diverging number of steps of the protocol. In this protocol, the machine \mathcal{M} is also represented by a single harmonic oscillator whose frequency matches that of the target oscillator that is to be cooled, $\omega_{\mathcal{M}} = \omega_S =: \omega$. The initial states of both the target system \mathcal{S} and machine \mathcal{M} are assumed to be thermal at the same inverse temperature β , and are hence both described

by thermal states of the form

$$\tau(\beta) = \frac{e^{-\beta H}}{\text{tr}[e^{-\beta H}]} = \sum_{n=0}^{\infty} e^{-\beta\omega 2n} (1 - e^{-\beta\omega}) |n\rangle\langle n| = \sum_{n=0}^{\infty} p_n |n\rangle\langle n|_{\mathcal{SM}}, \quad (\text{A.114})$$

where the Hamiltonian H is given by $H = \sum_{n=0}^{\infty} n\omega |n\rangle\langle n|$ and the $p_n = e^{-\beta\omega 2n} (1 - e^{-\beta\omega})$ are the eigenvalues of τ . The joint initial state is a product state that we can then write as

$$\tau_{\mathcal{S}}(\beta) \otimes \tau_{\mathcal{M}}(\beta) = \sum_{m,n=0}^{\infty} p_m p_n |m\rangle\langle m|_{\mathcal{S}} \otimes |n\rangle\langle n|_{\mathcal{M}} = \sum_{m,n=0}^{\infty} \tilde{p}_{m+n} |m, n\rangle\langle m, n|, \quad (\text{A.115})$$

where we have defined $\tilde{p}_k := e^{-\beta\omega 2k} (1 - e^{-\beta\omega})^2$. We then note that the eigenvalues \tilde{p}_k of the joint initial state have degeneracy $k + 1$. For instance, the largest value $\tilde{p}_0 = p_0 p_0$, corresponding to both the system and machine being in the ground state, is the single largest eigenvalue, but there are two eigenstates, $|0, 1\rangle$ and $|1, 0\rangle$, corresponding to the second largest eigenvalue \tilde{p}_1 , three states, $|0, 2\rangle$, $|1, 1\rangle$, and $|2, 0\rangle$ for the third largest eigenvalue \tilde{p}_2 , and so forth. Obviously, not all of these eigenvalues correspond to eigenstates for which the target system is in the ground state.

In order to increase the ground-state population of the target system oscillator, we can now apply a sequence of ‘two-level’ unitaries, i.e., unitaries that act only on a subspace spanned by two particular eigenstates and exchange their respective populations. The two-dimensional subspaces are chosen such that one of the two eigenstates corresponds to the system \mathcal{S} being in the ground state, $|0, k\rangle$, while the other eigenstate corresponds to \mathcal{S} being in an excited state, $|i \neq 0, j\rangle$. In addition, these pairs of levels are selected such that, at the time the unitary operation is to be performed, the population of $|0, k\rangle$ is smaller than that of $|i \neq 0, j\rangle$, such that the two-level exchange increases the ground-state population of \mathcal{S} at each step.

More specifically, at the k^{th} step of this sequence, the joint system \mathcal{SM} is in the state $\varrho_{\mathcal{SM}}^{(k)}$ and one determines the set Ω_k of index pairs $(i \neq 0, j)$ such that $\tilde{p}_k < \langle i, j | \varrho_{\mathcal{SM}}^{(k)} | i, j \rangle$, i.e., the set of eigenstates for which \mathcal{S} is not in the ground state and which have a larger associated population (at the beginning of the k^{th} step) than $|0, k\rangle$. One then determines an index pair (m_k, n_k) for which this population is maximal, i.e., $\langle m_k, n_k | \varrho_{\mathcal{SM}}^{(k)} | m_k, n_k \rangle = \max\{\langle i, j | \varrho_{\mathcal{SM}}^{(k)} | i, j \rangle | (i, j) \in \Omega_k\}$, and performs the unitary

$$U_{\mathcal{SM}}^{(k)} = \mathbb{1}_{\mathcal{SM}} - |0, k\rangle\langle 0, k| - |m_k, n_k\rangle\langle m_k, n_k| + \left(|0, k\rangle\langle m_k, n_k| + |m_k, n_k\rangle\langle 0, k| \right). \quad (\text{A.116})$$

If there is no larger population that is not already in the subspace of the ground state of the target system, i.e., when $\Omega_k = \emptyset$, which is only the case for the first step ($k = 1$), then no unitary is performed. After the k^{th} step, the joint state $\varrho_{\mathcal{SM}}^{(k+1)}$ is still diagonal in the energy eigenbasis, and the subspace of the joint Hilbert spaces for which \mathcal{S} is in the ground state is populated with the $k + 1$ largest eigenvalues \tilde{p}_i in non-increasing order with respect

to non-decreasing energy eigenvalues of the subspace's basis vectors $|0, i\rangle$. That is, for all $i \in \{0, 1, 2, \dots, k\}$ and for all $j \in \mathbb{N}$ with $j > i$, we have $\langle 0, i | \varrho_{\mathcal{SM}}^{(k+1)} | 0, i \rangle \geq \langle 0, j | \varrho_{\mathcal{SM}}^{(k+1)} | 0, j \rangle$.

Since the Hilbert spaces of both \mathcal{S} and \mathcal{M} are infinite-dimensional, we can continue with such a sequence of two-level exchanges indefinitely, starting with $k = 1$ and continuing step-by-step as $k \rightarrow \infty$. Here we note that the choice of (m_k, n_k) is generally not unique at the k -th step, but as $k \rightarrow \infty$, the resulting final state is independent of the particular choices of (m_k, n_k) made along the way. In particular, in a fashion that is reminiscent of the famed Hilbert hotel paradox (see, e.g., Ref. [364, p. 17]), this sequence places all of the infinitely many eigenvalues \tilde{p}_k of the joint state of \mathcal{SM} (which must hence sum to one) into the subspace where \mathcal{S} is in the ground state. In other words, in the limit of infinitely many steps, the population of the ground-state subspace evaluates to

$$\sum_{k=0}^{\infty} (k+1) \tilde{p}_k = \sum_{k=0}^{\infty} (k+1) e^{-\beta\omega 2k} (1 - e^{-\beta\omega})^2 = 1, \quad (\text{A.117})$$

where we have taken into account the $(k+1)$ -fold degeneracy of the k^{th} eigenvalue \tilde{p}_k . We thus have $\lim_{k \rightarrow \infty} \text{tr}_{\mathcal{M}} [\varrho_{\mathcal{SM}}^{(k)}] = |0\rangle\langle 0|_{\mathcal{S}}$, the reduced state of the system is asymptotically the pure state $|0\rangle_{\mathcal{S}}$.

As per our requirement on the structural complexity (see Appendix A.4), the Hilbert space of the machine required to achieve this is infinite-dimensional, and since each step of the protocol is assumed to take a finite amount of time, the overall time for reaching the ground state diverges. At the same time, the control complexity for each individual step is finite, since each U_k acts nontrivially only on a two-dimensional subspace. To see that also the energy cost for this protocol is finite, we first note that the protocol results in a final state of the machine that is diagonal in the energy eigenbasis $|n\rangle_{\mathcal{M}}$, with probability weights \tilde{p}_k decreasing (but not strictly) with increasing energy. Due to the degeneracy of the eigenvalues \tilde{p}_k , each one appears $(k+1)$ times on the diagonal (w.r.t. the energy eigenbasis) of the resulting machine state, populating adjacent energy levels. The label $n(k)$ of the lowest energy level that is populated by a particular value \tilde{p}_k can be calculated as

$$\tilde{n}(k) := \sum_{n=0}^{k-1} (n+1) = \frac{1}{2}k(k+1), \quad (\text{A.118})$$

while the largest energy populated by \tilde{p}_k is given by $\tilde{n}(k+1) - 1$. With this, we calculate the energy of the machine after the protocol, which evaluates to

$$\begin{aligned} \frac{E_{\mathcal{M}}^{\text{final}}}{\omega} &= \sum_{k=1}^{\infty} e^{-\beta\omega 2k} (1 - e^{-\beta\omega})^2 \sum_{n=\tilde{n}(k)}^{\tilde{n}(k+1)-1} n = \sum_{k=1}^{\infty} e^{-\beta\omega 2k} (1 - e^{-\beta\omega})^2 \frac{1}{2}k(k+1)(k+2) \\ &= \frac{3}{4} \text{cosech}^2\left(\frac{\beta\omega}{2}\right). \end{aligned} \quad (\text{A.119})$$

Since the energy of the initial thermal state is given by

$$\frac{E[\tau(\beta)]}{\omega} = \sum_{n=0}^{\infty} n e^{-\beta\omega 2n} (1 - e^{-\beta\omega}) = \frac{e^{-\beta\omega}}{1 - e^{-\beta\omega}}, \quad (\text{A.120})$$

we thus arrive at the energy cost

$$\frac{\Delta E_{\mathcal{M}}}{\omega} = \frac{E_{\mathcal{M}}^{\text{final}} - E[\tau(\beta)]}{\omega} = \frac{e^{-\beta\omega}(2 + e^{-\beta\omega})}{(1 - e^{-\beta\omega})^2}. \quad (\text{A.121})$$

We thus see that this energy cost is finite for all finite initial temperatures (although note that the energy cost diverges when $\beta \rightarrow 0$).

However, as we shall show next, the energy cost for attaining the ground state is not minimal, i.e., the protocol achieves perfect cooling (with finite energy and control complexity, but infinite time) but not at the Landauer limit. To see this, we first observe that the entropy of the final pure state of the system \mathcal{S} vanishes, such that $\tilde{\Delta}S_{\mathcal{S}} = S[\tau(\beta)]$. Evaluating this entropy, one obtains

$$\begin{aligned} S[\tau(\beta)] &= -\text{tr}[\tau \log(\tau)] = -\sum_{n=0}^{\infty} e^{-\beta\omega 2n} (1 - e^{-\beta\omega}) \log[e^{-\beta\omega 2n} (1 - e^{-\beta\omega})] \\ &= -\sum_{n=0}^{\infty} e^{-\beta\omega 2n} (1 - e^{-\beta\omega}) [-\beta\omega 2n + \log(1 - e^{-\beta\omega})] \\ &= \frac{\beta\omega e^{-\beta\omega}}{1 - e^{-\beta\omega}} + \beta\omega + \log\left(\frac{e^{-\beta\omega}}{1 - e^{-\beta\omega}}\right) \\ &= \frac{\beta\omega}{1 - e^{-\beta\omega}} + \log\left(\frac{e^{-\beta\omega}}{1 - e^{-\beta\omega}}\right). \end{aligned} \quad (\text{A.122})$$

Using the results from Eqs. (A.121) and (A.122), we can thus compare the expressions for $\beta\Delta E_{\mathcal{M}}$ and $\tilde{\Delta}S_{\mathcal{S}}$, and we find that $\beta\Delta E_{\mathcal{M}} - \tilde{\Delta}S_{\mathcal{S}} > 0$ for all nonzero initial temperatures. The origin of this difference is easily identified: Although the protocol results in an uncorrelated final state because the system is left in a pure state, that is, $I(\mathcal{S} : \mathcal{M})_{\rho'_{\mathcal{S}\mathcal{M}}} = 0$, the last term $D(\rho'_{\mathcal{M}} \parallel \tau_{\mathcal{M}})$ in Eq. (1.8) is nonvanishing for nonzero temperatures because the protocol does not result in a thermal state of the machine.

With this, we have thus shown that perfect cooling is indeed possible using a finite energy cost and a finite control complexity in every one of infinitely many steps (thus using diverging time). As we have seen, the structural requirement of an infinite-dimensional effective machine Hilbert space can be met by realising \mathcal{M} as a single harmonic oscillator. Although the presented protocol does not minimise the energy cost to saturate the Landauer bound, we cannot at this point conclusively say that it is not possible to do so in this setting. However, we suspect that a more complicated energy-level structure of the machine is necessary.

Finally, let us comment again on the notion of control complexity in terms of effective machine dimension as opposed to the notion of complexity that is often (loosely) associated with the distinction between Gaussian and non-Gaussian operations. As we see from the protocols presented here, the concept of control complexity based on the nontrivially accessed Hilbert-space dimension of the machine indeed captures the resource that must diverge in order to reach the ground state, while the intuition of complexity associated with (non)-Gaussian operations, albeit valid as a characterisation of a certain practical

difficulty in realising such operations, seems to be irrelevant for determining if the ground state can be reached. In the protocol presented in this section, non-Gaussian operations with finite control complexity are used in each step to reach the ground state. Infinitely many steps (i.e., diverging time) could then be traded for a single (also non-Gaussian) operation of infinite control complexity, performed in unit time. In the previous protocol based on Gaussian operations (Appendix A.5.2), the control complexity diverges in every single step of the cooling protocol, but only when there are infinitely many such steps (diverging time) or one operation in unit time on infinitely many modes (see below), can we reach the ground state. However, in the latter case, the operation, although acting on a diverging number of harmonic oscillators, remains Gaussian, as we now show explicitly.

A.5.3 Diverging Control Complexity Cooling Protocol for Harmonic Oscillators

Here we give a protocol for perfectly cooling a harmonic oscillator in unit time and with the minimum energy cost, but with diverging control complexity. In accordance with Theorem 1.3, the machines used to cool the target system will likewise be harmonic oscillators. Let the operators a (a^\dagger) and b_k (b_k^\dagger) respectively denote the annihilation (creation) operators of the target system and a machine subsystem labelled k . We then consider the the unitary transformation in Eq. (A.105), namely

$$U_k := e^{i\frac{\pi}{2}(a^\dagger b_k + ab_k^\dagger)}. \quad (\text{A.123})$$

One can then apply the diverging-time cooling protocol from Appendix A.5.2 to cool the system to the ground state at the Landauer limit via the total unitary transformation

$$U_{\text{tot}} := \lim_{N \rightarrow \infty} U_{(N)}, \quad \text{with} \quad U_{(N)} := \prod_{k=1}^N U_k. \quad (\text{A.124})$$

We now seek the Hamiltonian that generates U_{tot} . First note that $U_{(N)} a U_{(N)}^\dagger = ib_1$ and

$$U_{(N)} b_k U_{(N)}^\dagger = \begin{cases} -b_{k+1} & \text{for } k < N \\ ia & \text{for } k = N \\ b_k & \text{for } k > N, \end{cases} \quad (\text{A.125})$$

which can be proven by induction. In contrast with Appendix A.5.2, here we use the complex representation of the symplectic group to describe the transformation, i.e., the set of matrices S satisfying $SKS^\dagger = K$, where $K := \mathbb{1}_N \oplus (-\mathbb{1}_N)$. Gathering the raising and lowering operators of the target system and the first N machines into the vector $\vec{\xi} := \left(a \quad b_1 \quad b_2 \quad \dots \quad b_N \quad a^\dagger \quad b_1^\dagger \quad b_2^\dagger \quad \dots \quad b_N^\dagger \right)^\text{T}$, we can write the transformation above

as $U_{(N)}\vec{\xi}U_{(N)}^\dagger = S^T\vec{\xi}$ [365], where

$$S = \begin{pmatrix} \alpha_{(N)} & 0 \\ 0 & \alpha_{(N)} \end{pmatrix}, \quad \text{with} \quad \alpha_{(N)} := \begin{pmatrix} 0 & 0 & 0 & \dots & 0 & i \\ i & 0 & 0 & \dots & 0 & 0 \\ 0 & -1 & 0 & \dots & 0 & 0 \\ 0 & 0 & -1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & -1 & 0 \end{pmatrix}. \quad (\text{A.126})$$

Now, defining the matrix of Hamiltonian coefficients $h_{(N)}$ implicitly by $U_{(N)} =: \exp(-i\vec{\xi}^\dagger \cdot h_{(N)} \cdot \vec{\xi})$, we have that $S = \exp(-iK h_{(N)})$ [365], i.e., $h_{(N)} = iK \log(S^T) = iK \log(S)^T$, where we take the principal logarithm. To calculate this, we must diagonalise the matrix $\alpha_{(N)}$ in Eq. (A.126). The eigenvalues of $\alpha_{(N)}$ are

$$\lambda_k := -e^{-i\pi \frac{2k-1}{N+1}}, \quad \text{with} \quad k \in \{1, 2, \dots, N+1\}, \quad (\text{A.127})$$

i.e., the negative of the $(N+1)^{\text{th}}$ roots of -1 , and it is diagonalised by the unitary matrix V constructed from the eigenvectors \vec{v}_k :

$$V := \begin{pmatrix} \vec{v}_1 & \vec{v}_2 & \vec{v}_3 & \dots & \vec{v}_{N+1} \end{pmatrix} \quad \text{with} \quad \vec{v}_k := \frac{-1}{\sqrt{N+1}} \begin{pmatrix} i(-\lambda_k)^{-1} \\ (-\lambda_k)^{-2} \\ (-\lambda_k)^{-3} \\ \vdots \\ (-\lambda_k)^{-(N+1)} \end{pmatrix}. \quad (\text{A.128})$$

Specifically, $\alpha_{(N)} = VDV^\dagger$, where $D := \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_{N+1})$, and thus

$$\begin{aligned} h_{(N)}^T &= iK \log \begin{pmatrix} VDV^\dagger & 0 \\ 0 & VDV^\dagger \end{pmatrix} = iK \begin{pmatrix} V & 0 \\ 0 & W \end{pmatrix} \begin{pmatrix} \log(D) & 0 \\ 0 & \log(D) \end{pmatrix} \begin{pmatrix} V^\dagger & 0 \\ 0 & V^\dagger \end{pmatrix} \\ &=: \begin{pmatrix} A & 0 \\ 0 & -A \end{pmatrix} \end{aligned} \quad (\text{A.129})$$

for some matrix A . By direct calculation, one finds that

$$A_{jk} = i^{\delta_{j1}} i^{\delta_{k1}} \frac{\pi}{(N+1)^2} \sum_{p=1}^{N+1} (2p-2-N) e^{-i\pi \frac{2p-1}{N+1}(j-k)}. \quad (\text{A.130})$$

Now, considering the identity

$$\sum_{p=1}^{N+1} e^{i\theta p} = \frac{e^{i\theta(N+1)} - 1}{1 - e^{i\theta}} \quad (\text{A.131})$$

for $\theta \in \mathbb{R}$, as well as its derivative with respect to θ , one can calculate the sum in Eq. (A.130). We then have

$$\lim_{N \rightarrow \infty} A_{jk} = \begin{cases} 0, & \text{for } j = k \\ ii^{\delta_{j1}} i^{\delta_{k1}} \frac{1}{j-k}, & \text{for } j \neq k \end{cases}. \quad (\text{A.132})$$

Then, finally, we have that $U_{\text{tot}} = e^{-iH_{\text{tot}}}$, where $H_{\text{tot}} = \lim_{N \rightarrow \infty} (\vec{v}^\dagger \cdot h_{(N)} \cdot \vec{v})$, i.e.,

$$H_{\text{tot}} = - \sum_{j=2}^{\infty} \left(\frac{1}{j-1} b_j^\dagger a + \text{H.c.} \right) + \sum_{j,k=1; j \neq k}^{\infty} \frac{i}{j-k} b_j^\dagger b_k. \quad (\text{A.133})$$

Thus, the system is cooled to the ground state at an energy cost saturating the Landauer bound, and in unit time, but via a procedure that implements a multi-mode Gaussian unitary on a diverging number of modes.

A.6 Cooling Protocols in the Incoherent-Control Paradigm

In this section, we investigate the required resources to cool the target system within the incoherent-control paradigm. For simplicity, we consider only the finite-dimensional setting. Here, we have a qudit target system \mathcal{S} interacting resonantly (i.e., in an energy-conserving manner) with a qudit machine \mathcal{M} , which is partitioned into one part, \mathcal{C} , in thermal contact with the ambient environment at inverse temperature β and another part, \mathcal{H} , in contact with a hot bath at inverse temperature $\beta_H < \beta$. The Hamiltonians for each subsystem are $H_{\mathcal{X}} = \sum_{n=0}^{d_{\mathcal{X}}-1} n \omega_{\mathcal{X}} |n\rangle\langle n|_{\mathcal{X}}$; the energy resonance condition enforces that $\omega_{\mathcal{H}} = \omega_{\mathcal{C}} - \omega_{\mathcal{S}}$. For the most part in this section, we focus on equally spaced Hamiltonians for simplicity; we comment specifically whenever we consider otherwise.

In order to cool the target system, we aim to compress as much population as possible into the its lowest energy eigenstates via interactions that are restricted to the energy-degenerate subspaces of the joint \mathcal{SCH} system. Thus we are restricted to global energy-conserving unitaries U_{EC} that satisfy

$$[H_{\mathcal{S}} + H_{\mathcal{C}} + H_{\mathcal{H}}, U_{\text{EC}}] = 0. \quad (\text{A.134})$$

In Ref. [83], it was shown that for the case where all three subsystems are qubits, the optimal global unitary in this setting (inasmuch as they render the target system in the coldest state possible given the restrictions) is

$$U_{\text{EC}} = |0, 1, 0\rangle\langle 1, 0, 1|_{\mathcal{SCH}} + |1, 0, 1\rangle\langle 0, 1, 0|_{\mathcal{SCH}} + \bar{\mathbb{1}}, \quad (\text{A.135})$$

where $\bar{\mathbb{1}}$ denotes the identity matrix on all subspaces that are not energy-degenerate. Considering the generalisation to qudit subsystems, it is straightforward to see that, for equally spaced Hamiltonians, the optimal global unitaries must be of the form

$$U_{\text{EC}} = \left[\sum_{m,n,l=0}^{d-2} |m, n+1, l\rangle\langle m+1, n, l+1|_{\mathcal{SCH}} + |m+1, n, l+1\rangle\langle m, n+1, l|_{\mathcal{SCH}} \right] + \bar{\mathbb{1}}. \quad (\text{A.136})$$

For the most general case where the Hamiltonians of each subsystem are arbitrary, it is not possible to write down a generic form of the optimal unitary, since the energy-resonant transitions that lead to cooling the target now depend on the microscopic details of the energetic structure. Nonetheless, in Appendix A.7, we provide a protocol (i.e., not the unitary per se, but a sequence of steps) in this setting that attains perfect cooling and saturates the Carnot-Landauer limit.

Intuitively, the above types of unitaries simply reshuffle populations that are accessible through resonant transitions. For the purpose of cooling, one wishes to do this in such a way that the largest population is placed in the lowest energy eigenstate of the target system, the second largest in the second lowest energy eigenstate, and so on (in line with the optimal unitaries in the coherent-control setting); indeed, on the energy-degenerate subspaces accessible, such unitaries act precisely in this way. It is straightforward to show that interactions of this form satisfy Eq. (A.134).

For the sake of simplicity, we now focus on the case where all systems are qubits, although the results generalise to the qudit setting. Consider the initial joint state $\varrho_{SC\mathcal{H}} = \sum_{m,n,l=0}^1 p_{mnl} |m, n, l\rangle\langle m, n, l|_{SC\mathcal{H}}$. By applying a unitary U_{EC} of the form given in Eq. (A.136), the post-transformation joint state is

$$\varrho'_{SC\mathcal{H}} = U_{EC} \varrho_{SC\mathcal{H}} U_{EC}^\dagger = \varrho_{SC\mathcal{H}} + \Delta p |0, 1, 0\rangle\langle 0, 1, 0|_{SC\mathcal{H}} - \Delta p |1, 0, 1\rangle\langle 1, 0, 1|_{SC\mathcal{H}}, \quad (\text{A.137})$$

where $\Delta p := p_{101} - p_{010}$ indicates the amount of population that has been transferred from the excited state of the target system to the ground state throughout the interaction. Naturally, in order to cool the target system, $\Delta p \geq 0$, i.e., the initial population p_{101} must be at least as large as p_{010} .

Due to the energy-conserving nature of the global interaction, the energy exchanged between the subsystems throughout a single such interaction, $\Delta E_x = \text{tr} [H_x(\varrho'_x - \varrho_x)]$, can be calculated via

$$\Delta E_S = -\omega_S \Delta p, \quad \Delta E_C = \omega_C \Delta p, \quad \Delta E_{\mathcal{H}} = -\omega_{\mathcal{H}} \Delta p. \quad (\text{A.138})$$

Thus, for a fixed energy-level structure of all subsystems (i.e., given the local Hamiltonians), one only requires knowledge of the pre- and post-transformation state of any one of the subsystems to calculate the energy change for all of them.

A.6.1 Diverging Energy: Proof of Theorem 1.5

The first thing to note is that in the incoherent-control paradigm, even when one allows for the energy cost, i.e., the heat drawn from the hot bath, to be diverging, it is not possible to perfectly cool the target system, as presented in Theorem 1.5. The intuition behind this result is that the target system can only interact with energy-degenerate subspaces of the hot and cold machine subsystems. The optimal transformation that one

can do here to achieve cooling is to transfer the highest populations of any such subspace to the lowest energy eigenstate of the target system; however, any such subspace has population strictly less than one for any $0 \leq \beta_H \leq \beta < \infty$ independently of the energy structure. Moreover, the difference from one can be bounded by a finite amount that does not vanish independent of the energy level structure of any machine of finite dimension. This makes it impossible to attain a subspace population of one even as the energy cost diverges for any fixed and finite control complexity. It follows that the ground-state population of the target system can never reach unity in a single operation of finite control complexity and hence perfect cooling cannot be achieved.

Precisely, we show the following. Let \mathcal{S} be a finite-dimensional system of dimension $d_{\mathcal{S}}$ with associated Hamiltonian with finite but otherwise arbitrary energy gaps $H_{\mathcal{S}} = \sum_{i=0}^{d_{\mathcal{S}}-1} \omega_{\mathcal{S}}^i |i\rangle\langle i|_{\mathcal{S}}$, and let $d_{\mathcal{C}}$ and $d_{\mathcal{H}}$ be integers denoting the dimensions of the cold and hot parts of the machine respectively. Then it is impossible to cool the system \mathcal{S} in the incoherent-control paradigm, i.e., using energy-conserving unitaries involving \mathcal{C} and \mathcal{H} at some initial inverse temperatures β, β_H respectively, arbitrarily close to the ground state. Note that, in particular, this result holds irrespective of the energy level structure of \mathcal{C} and \mathcal{H} and no matter how much energy is drawn from the hot bath as a resource.

In order to set notation for the following, we assume $\omega_{\mathcal{X}}^i \geq \omega_{\mathcal{X}}^j$ for $i \geq j$ and $\omega_{\mathcal{X}}^0 = 0$, where $\omega_{\mathcal{X}}^i$ denotes the i^{th} energy eigenvalue of system \mathcal{X} with $\mathcal{X} \in \{\mathcal{S}, \mathcal{C}, \mathcal{H}\}$. We also assume the initial states on \mathcal{S} and \mathcal{C} to be thermal at inverse temperature β , and \mathcal{H} is assumed to be initially in a thermal state at inverse temperature $\beta_H \leq \beta$. We denote by $p_{\mathcal{X}}^i$ the i^{th} population of system \mathcal{X} in a given state, i.e., $p_{\mathcal{X}}^i = \langle i | \rho_{\mathcal{X}} | i \rangle$, where $|i\rangle$ denotes the i^{th} energy eigenstate of $\rho_{\mathcal{X}}$. We will also write $p_{ijk} := p_{\mathcal{S}}^i p_{\mathcal{C}}^j p_{\mathcal{H}}^k$.

The intuition behind the proof is as follows. The global ground state level of the joint hot-and-cold machine has some non-zero initial population for any finite-dimensional machine; in particular it can always be lower-bounded by $\frac{1}{d_{\mathcal{C}} d_{\mathcal{H}}}$ for any Hamiltonians and initial temperatures, which is strictly greater than zero as long as the dimensions remain finite. Fixing the control complexity of any protocol considered here to be finite in value thus implies a lower bound on the initial ground state population of the total machine that is larger than zero by a finite amount. Depending on the energy level structure of the hot and cold parts of the machine, there may be other non-zero initial populations, but in order to cool the target system \mathcal{S} perfectly, at least all of the previously mentioned populations must be transferred into spaces spanned by energy eigenstates of the form $|0jk\rangle_{\mathcal{S}\mathcal{C}\mathcal{H}}$. This intuition is formalised via Lemma A.2, where we show that independent of the energy structure of \mathcal{C} and \mathcal{H} , one must be able to make such transfers of population in order to perfectly cool \mathcal{S} . However, in order to make such transfers in an energy conserving manner, all energy eigenstates of the form $|i00\rangle_{\mathcal{S}\mathcal{C}\mathcal{H}}$ must be degenerate with some of the form $|0jk\rangle_{\mathcal{S}\mathcal{C}\mathcal{H}}$. This degeneracy condition, in turn, also implies that every energy eigenstate of the form $|0jk\rangle_{\mathcal{S}\mathcal{C}\mathcal{H}}$ has an associated initial population p_{0jk} that is

non-vanishing for all machines of finite dimension (i.e., for all protocols with finite control complexity). Thus, upon transferring some population p_{i00} into the subspace spanned by $|0jk\rangle_{\mathcal{SCH}}$, i.e., one of a relevant form for the population to contribute to the final ground state population of the target, one inevitably transfers some finite amount of population away from the relevant space and into $|i00\rangle_{\mathcal{SCH}}$, which does not contribute to the final ground state population of the target. We formalise this intuition in the discussion following Lemma A.2. In this way, no matter what one does, there is always a finite amount of population, which is lower-bounded by some strictly positive number due to the constraint on control complexity, that does not contribute to the final ground state population of the target, implying that perfect cooling is not possible.

The formal proof occurs in two steps. We first show that some specific degeneracies in the joint \mathcal{SCH} system must be present in order to be able to even potentially cool \mathcal{S} arbitrarily close to the ground state. We then prove that, given the above degeneracies, one cannot cool the system \mathcal{S} beyond a fixed ground-state population that is independent of the energy structure of \mathcal{C} and \mathcal{H} ; in particular, one can draw as much energy from the hot bath as they like and still do no better. We begin with the following lemma.

Lemma A.2. *Given \mathcal{S} , $d_{\mathcal{C}}$ and $d_{\mathcal{H}}$ as above, one can reach a final ground-state population of the system \mathcal{S} arbitrarily close to one in the incoherent-control setting only if each $|i00\rangle_{\mathcal{SCH}}$, where $i \in \{1, \dots, d_{\mathcal{S}} - 1\}$, energy eigenstate is degenerate with at least one $|0jk\rangle_{\mathcal{SCH}}$ energy eigenstate, where $j \in \{0, \dots, d_{\mathcal{C}} - 1\}, k \in \{0, \dots, d_{\mathcal{H}} - 1\}$.*

Proof. Suppose that there exists an $i^* \in \{1, \dots, d_{\mathcal{S}} - 1\}$ such that $|i^*00\rangle_{\mathcal{SCH}}$ is not degenerate with any $|0jk\rangle_{\mathcal{SCH}}$, where $j \in \{0, \dots, d_{\mathcal{C}} - 1\}, k \in \{0, \dots, d_{\mathcal{H}} - 1\}$. We show that, then, one cannot cool \mathcal{S} arbitrarily close to zero.

Let B_i denote the degenerate subspace of the total Hamiltonian $H_{\mathcal{S}} + H_{\mathcal{C}} + H_{\mathcal{H}}$, where $H_{\mathcal{X}}$ denotes the Hamiltonian of system $\mathcal{X} \in \{\mathcal{S}, \mathcal{C}, \mathcal{H}\}$, that contains the eigenvector $|i00\rangle_{\mathcal{SCH}}$. Then, any energy conserving unitary U_{EC} used to cool the system in the incoherent-control paradigm must act within such B_i subspaces, i.e., $U_{\text{EC}} = \bigoplus_i U_{B_i}$ (this is a direct consequence of $[U_{\text{EC}}, H_{\mathcal{S}} + H_{\mathcal{C}} + H_{\mathcal{H}}] = 0$, see, e.g., Lemma 5 of Ref. [96]). This means in particular that the initial population of $|i^*00\rangle_{\mathcal{SCH}}$ can only be distributed within B_{i^*} , and as no eigenvector of the form $|0jk\rangle_{\mathcal{SCH}}$ is contained in B_{i^*} by assumption, it can never contribute to the final ground state population of \mathcal{S} , which we denote $\tilde{p}_{\mathcal{S}}^0$. So we have

$$\tilde{p}_{\mathcal{S}}^0 \leq 1 - p_{i^*00}. \quad (\text{A.139})$$

Now, as for $\mathcal{X} \in \{\mathcal{C}, \mathcal{H}\}$, with any $\{\omega_{\mathcal{X}}^i\}$ such that each $\omega_{\mathcal{X}}^i \geq 0$ with $\omega_{\mathcal{X}}^0 = 0$ and any inverse temperature $\beta \geq 0$, we have for the partition function $\mathcal{Z}_{\mathcal{S}}$ that

$$\mathcal{Z}_{\mathcal{X}} = 1 + e^{-\beta\omega_{\mathcal{X}}^1} + \dots + e^{-\beta\omega_{\mathcal{X}}^{d_{\mathcal{X}}-1}} \leq d_{\mathcal{X}}, \quad (\text{A.140})$$

and so we have the following bound on the initial populations associated to each eigenvector $|i00\rangle_{SC\mathcal{H}}$

$$p_{i00} = \frac{e^{-\beta\omega_S^i}}{\mathcal{Z}_S \mathcal{Z}_C \mathcal{Z}_\mathcal{H}} \geq \frac{e^{-\beta\omega_S^i}}{\mathcal{Z}_S d_C d_\mathcal{H}} > 0 \quad \forall i \in \{1, \dots, d_S - 1\}. \quad (\text{A.141})$$

Combining the above, we have that

$$\tilde{p}_S^0 \leq 1 - \frac{e^{-\beta\omega_S^{i^*}}}{\mathcal{Z}_S d_C d_\mathcal{H}} < 1. \quad (\text{A.142})$$

So as desired, we have shown that one cannot cool beyond $1 - \frac{e^{-\beta\omega_S^{i^*}}}{\mathcal{Z}_S d_C d_\mathcal{H}}$, a bound strictly smaller than 1 for any finite-dimensional machine (i.e., for any protocol using only finite control complexity) and independent of the energies of \mathcal{C} and \mathcal{H} . \square

We can now proceed to the second step of the proof of Theorem 1.5.

Proof. To this end, consider any $i^* \in \{1, \dots, d_S - 1\}$. If $|i^*00\rangle_{SC\mathcal{H}}$ is not degenerate with any $|0jk\rangle_{SC\mathcal{H}}$, our assertion is proven by Lemma A.2. On the other hand, if there is a $j^* \in \{0, \dots, d_C - 1\}$ and a $k^* \in \{0, \dots, d_\mathcal{H} - 1\}$ for which $|i^*00\rangle_{SC\mathcal{H}}$ and $|0j^*k^*\rangle_{SC\mathcal{H}}$ are degenerate, then B_{i^*} , the degenerate subspace containing $|i^*00\rangle_{SC\mathcal{H}}$, also contains $|0j^*k^*\rangle$. Now B_{i^*} may also contain other eigenvectors of the form $|0jk\rangle_{SC\mathcal{H}}$, i.e., some other $|0j'k'\rangle_{SC\mathcal{H}}$ with $j' \in \{0, \dots, d_C - 1\}$, $k' \in \{0, \dots, d_\mathcal{H} - 1\}$. Crucially, each such eigenvector in B_{i^*} must have an associated minimal amount of initial population as long as the machine is finite-dimensional. Indeed, for any such $|0j^*k^*\rangle_{SC\mathcal{H}}$ in B_{i^*} , we have the condition $\omega_C^{j^*} + \omega_\mathcal{H}^{k^*} = \omega_S^{i^*}$ and so $\omega_C^{j^*} \leq \omega_S^{i^*}$, $\omega_\mathcal{H}^{k^*} \leq \omega_S^{i^*}$, implying that $\beta\omega_C^{j^*} \leq \beta\omega_S^{i^*}$ and $\beta\omega_\mathcal{H}^{k^*} \leq \beta\omega_S^{i^*}$. Thus we have the bound

$$p_{0j^*k^*} = \frac{e^{-\beta\omega_C^{j^*}} e^{-\beta\omega_\mathcal{H}^{k^*}}}{\mathcal{Z}_S \mathcal{Z}_C \mathcal{Z}_\mathcal{H}} \geq \frac{e^{-2\beta\omega_S^{i^*}}}{\mathcal{Z}_S \mathcal{Z}_C \mathcal{Z}_\mathcal{H}} \geq \frac{e^{-2\beta\omega_S^{i^*}}}{\mathcal{Z}_S d_C d_\mathcal{H}} =: q_{i^*}. \quad (\text{A.143})$$

Now, take any particular $i^* \in \{1, \dots, d_S - 1\}$ and let π_{i^*} be the dimension of B_{i^*} , μ the number of energy eigenstates of the form $|0jk\rangle_{SC\mathcal{H}}$ that B_{i^*} contains and $\nu = \pi_{i^*} - \mu$ the number of energy eigenstates of the form $|ijk\rangle_{SC\mathcal{H}}$, where $i \neq 0$, that B_{i^*} contains. So

$$B_{i^*} = \text{span}\{|0jk\rangle, |0j_2k_2\rangle, \dots, |0j_\mu k_\mu\rangle, |i^*00\rangle, |i_2\ell_2 m_2\rangle, \dots, |i_\nu \ell_\nu m_\nu\rangle\}. \quad (\text{A.144})$$

Let $\mathbf{v} = \{p_{0jk}, p_{0j_2k_2}, \dots, p_{0j_\mu k_\mu}, p_{i^*00}, p_{i_2\ell_2 m_2}, \dots, p_{i_\nu \ell_\nu m_\nu}\}$ be the vector of initial populations associated to the eigenvectors of B_{i^*} , and \mathbf{v}^\uparrow be the vector whose components are those of \mathbf{v} arranged in non-decreasing order. Using Schur's theorem [362], we know that after applying any unitary transformation $U_{B_{i^*}}$ on the relevant energy-degenerate subspace, then the vector of transformed populations, $\tilde{\mathbf{v}}$, is majorised by \mathbf{v} . In particular, labelling the vector elements by \mathbf{v}_α , we have

$$\tilde{p}_{i^*00} + \sum_{\alpha=2}^{\nu} \tilde{p}_{i_\alpha \ell_\alpha m_\alpha} \geq \sum_{\alpha=1}^{\nu} \mathbf{v}_\alpha^\uparrow. \quad (\text{A.145})$$

We now claim that $\sum_{\alpha=1}^{\nu} \mathbf{v}_{\alpha}^{\dagger} \geq q_{i^*}$ from Eq. (A.143). Indeed, as \mathbf{v} has at most $\nu - 1$ elements that do not belong to the set $A := \{p_{0jk}, p_{0j_2k_2}, \dots, p_{0j_{\mu}k_{\mu}}, p_{i^*00}\}$, at least one element of A must contribute to the sum $\sum_{\alpha=1}^{\nu} \mathbf{v}_{\alpha}^{\dagger}$. Let x be that element. As $\mathbf{v}_{\alpha}^{\dagger} \geq 0$ for all $\alpha = 1, \dots, \pi = \mu + \nu$, we have

$$\sum_{\alpha=1}^{\nu} \mathbf{v}_{\alpha}^{\dagger} \geq x. \quad (\text{A.146})$$

Now as $p_{0j_{\gamma}k_{\gamma}} \geq q_{i^*}$ for all $\gamma = 2, \dots, \mu$, we have

$$x \geq \min(q_{i^*}, p_{i^*00}) = q_{i^*}, \quad (\text{A.147})$$

where $p_{i^*00} \geq q_{i^*}$ can be seen from Eq. (A.143), as claimed.

As the l.h.s. of Eq. (A.145) represents the amount of population in the subspace B_{i^*} that does not contribute to the final ground-state population of the target system, we have

$$\tilde{p}_S^0 \leq 1 - \left(\tilde{p}_{i^*00} + \sum_{\alpha=2}^{\nu} \tilde{p}_{i_{\alpha}l_{\alpha}m_{\alpha}} \right) \leq 1 - q_{i^*} = 1 - \frac{e^{-2\beta\omega_S^*}}{\mathcal{Z}_S d_C d_{\mathcal{H}}}. \quad (\text{A.148})$$

So, for any finite-dimensional machine, one cannot cool the system \mathcal{S} beyond $1 - \frac{e^{-\beta\omega_S^*}}{\mathcal{Z}_S d_C d_{\mathcal{H}}}$, a bound strictly smaller than 1 and independent of the energy structure of \mathcal{C} and \mathcal{H} , as desired. \square

As a concrete example, consider the case where all systems are qubits. The initial joint state is

$$\varrho_{SC\mathcal{H}}^{(0)} = \frac{(|0\rangle\langle 0| + e^{-\beta\omega_S}|1\rangle\langle 1|)_S \otimes (|0\rangle\langle 0| + e^{-\beta\omega_C}|1\rangle\langle 1|)_C \otimes (|0\rangle\langle 0| + e^{-\beta_H\omega_{\mathcal{H}}}|1\rangle\langle 1|)_{\mathcal{H}}}{\mathcal{Z}_S(\beta, \omega_S)\mathcal{Z}_C(\beta, \omega_C)\mathcal{Z}_{\mathcal{H}}(\beta_H, \omega_{\mathcal{H}})}. \quad (\text{A.149})$$

The only energy-conserving unitary interaction that is relevant for cooling is the one that exchanges the populations in the levels spanned by $|010\rangle$ and $|101\rangle$, which have initial populations $\frac{e^{-\beta\omega_C}}{\mathcal{Z}_S(\beta, \omega_S)\mathcal{Z}_C(\beta, \omega_C)\mathcal{Z}_{\mathcal{H}}(\beta_H, \omega_{\mathcal{H}})}$ and $\frac{e^{-\beta\omega_S}e^{-\beta_H\omega_{\mathcal{H}}}}{\mathcal{Z}_S(\beta, \omega_S)\mathcal{Z}_C(\beta, \omega_C)\mathcal{Z}_{\mathcal{H}}(\beta_H, \omega_{\mathcal{H}})}$ respectively, which are both strictly less than one. The necessary condition for any cooling to be possible implies that $e^{-\beta\omega_S}e^{-\beta_H\omega_{\mathcal{H}}} \geq e^{-\beta\omega_C}$; now, performing the optimal cooling unitary leads to the final ground-state population of the target system

$$p'_S(0) = \langle 0 | \text{tr}_{C\mathcal{H}} [U \varrho_{SC\mathcal{H}}^{(0)} U^{\dagger}] | 0 \rangle_S = \frac{1 + e^{-\beta_H\omega_{\mathcal{H}}}(1 + e^{-\beta\omega_S} + e^{-\beta\omega_C})}{\mathcal{Z}_S(\beta, \omega_S)\mathcal{Z}_C(\beta, \omega_C)\mathcal{Z}_{\mathcal{H}}(\beta_H, \omega_{\mathcal{H}})} < 1. \quad (\text{A.150})$$

Indeed, using $e^{-\beta\omega_S}e^{-\beta_H\omega_{\mathcal{H}}} \geq e^{-\beta\omega_C}$,

$$p'_S(0) \leq \frac{1 + e^{-\beta_H\omega_{\mathcal{H}}}e^{-\beta\omega_S}}{\mathcal{Z}_S(\beta, \omega_S)\mathcal{Z}_C(\beta, \omega_C)} \leq \frac{1}{\mathcal{Z}_C(\beta, \omega_C)} \leq 1. \quad (\text{A.151})$$

The second inequality is strict unless $\beta_{\mathcal{H}} = 0$ or $\omega_{\mathcal{H}} = 0$. In the both cases, for equality in the first inequality, we need $\beta\omega_S = \beta\omega_C$. If $\beta = 0$, then $\mathcal{Z}_C(\beta, \omega_C) = 2$ and the last inequality is strict. If $\omega_S = \omega_C$, no cooling is possible; hence $p'_S(0) = p_S(0) < 1$.

A.6.2 Diverging Time / Control Complexity

We now move to analyse the case where diverging time is allowed, where we wish to minimise the energy cost and control complexity throughout the protocol over a diverging number of energy-conserving interactions between the target system and the hot and cold subsystems of the machine. We again consider all three systems to be qubits, but the results generalise to arbitrary (finite) dimensions. Here, the machines and ancillas begin as thermal states with initial inverse temperatures β and $\beta_H \leq \beta$ respectively. Just as in the diverging time cooling protocol in the coherent-control setting presented in Appendix A.3, we will consider a diverging number of machines, with slightly increasing energy gaps, in a configuration such that the target system interacts with the n^{th} machine at time step n . Suppose that after n steps of the protocol, the target qubit has been cooled to some inverse temperature $\beta_n > \beta$; equivalently, this can be expressed as a thermal state with corresponding energy gap $\omega_n = \frac{\beta_n}{\beta} \omega_S$. We now wish to interact the target system $\tau_S(\beta_n, \omega_S)$ with a machine \mathcal{M}_{n+1} with slightly increased energy gaps with respect to the most recent one \mathcal{M}_n , i.e., we increase the energy gaps of the cold subsystem \mathcal{C} from ω_n to $\omega_{n+1} = \omega_n + \epsilon_n$; the resonance condition enforces the energy gap of the hot subsystem \mathcal{H} to be similarly increased to $\omega_n + \epsilon_n - \omega_S$. Thus, the next step of the protocol is a unitary acting on the global state

$$\varrho_{SC\mathcal{H}}^{(n)} = \tau_S(\beta_n, \omega_S) \otimes \tau_C(\beta, \omega_n + \epsilon_n) \otimes \tau_H(\beta_H, \omega_n + \epsilon_n - \omega_S). \quad (\text{A.152})$$

In order to cool the target system via said unitary, we must have that $p_{101} \geq p_{010}$ for the state in Eq. (A.152), which implies that ϵ_n must satisfy the following condition:

$$e^{-\beta\omega_n - \beta_H(\omega_n + \epsilon_n - \omega_S)} \geq e^{-\beta(\omega_n + \epsilon_n)} \Rightarrow \epsilon_n \geq \gamma(\omega_n - \omega_S) \quad \text{where} \quad \gamma := \frac{\beta_H}{\beta - \beta_H}. \quad (\text{A.153})$$

This condition is crucial. It means that if the hot subsystem \mathcal{H} is coupled to a heat bath at any finite temperature, i.e., $\beta_H > 0$, ϵ_n depends linearly on the inverse temperature of the target system at the previous step β_n , and can thus not be taken to be arbitrarily small. As we will now show, this condition prohibits the ability to perfectly cool the target system at the Landauer limit for the energy cost whenever the heat bath is at finite temperature.

On the other hand, for infinite-temperature heat baths, perfect cooling at the Landauer limit is seemingly achievable; here, $\beta_H \rightarrow 0$ and so $\gamma \rightarrow 0$, leading to the trivial constraint $\epsilon_n \geq 0$ which allows it to be arbitrarily small, as is required. Nonetheless, the explicit construction of any protocol doing so in the incoherent-control setting is a priori unclear, as the restriction of energy conservation makes for a fundamentally different setting from the coherent-control paradigm. We now explicitly derive the optimal diverging-time protocol to perfectly cool at the Landauer limit for an infinite-temperature heat bath, thereby proving Theorem 1.6.

A.6.3 Saturating the Landauer Limit with an Infinite-Temperature Heat Bath

Before calculating the energy cost, we briefly discuss the attainability of the optimally cool target state. We begin with all subsystems as qubits, for the sake of simplicity, but the logic generalises to higher dimensions. In the incoherent paradigm, the target system \mathcal{S} interacts with a virtual qubit of the total machine $\mathcal{M} = \mathcal{CH}$ that consists of the energy eigenstates $|0, 1\rangle_{\mathcal{CH}}$ and $|1, 0\rangle_{\mathcal{CH}}$, with populations $p_{0_{\mathcal{C}1_{\mathcal{H}}}}$ and $p_{1_{\mathcal{C}0_{\mathcal{H}}}}$ respectively. Suppose that at step $n + 1$ the cold subsystem involved in the interaction has energy gap $\omega_n + \epsilon_n$. In Ref. [83], it is shown that by repeating the incoherent cooling process [i.e., implementing the unitary in Eq. (A.136)] and taking the limit of infinite cycles, this scenario equivalently corresponds to the general (coherent) setting where arbitrary unitaries are permitted and the target system interacts with a virtual qubit machine with effective energy gap ω_n^{eff} given by

$$\begin{aligned} e^{-\beta\omega_n^{\text{eff}}} &:= \frac{p_{1_{\mathcal{C}0_{\mathcal{H}}}}}{p_{0_{\mathcal{C}1_{\mathcal{H}}}}} = e^{-\beta(\omega_n + \epsilon_n)} e^{\beta_H(\omega_n + \epsilon_n - \omega_S)} \\ \therefore \omega_n^{\text{eff}} &= \omega_n + \epsilon_n - \frac{\beta_H}{\beta}(\omega_n + \epsilon_n - \omega_S). \end{aligned} \quad (\text{A.154})$$

It is clear that for finite-temperature heat baths, i.e., $\beta_H > 0$, the effective energy gap ω_n^{eff} is always smaller than the energy gap of the machine at any given step, i.e., $\omega_n^{\text{eff}} \leq \omega_n + \epsilon_n$; on the other hand, equality holds iff the heat bath is at infinite temperature, i.e., $\beta_H \rightarrow 0$. Thus, in the infinite-temperature case, given a target system beginning at some step of the protocol in the state $\varrho_S^*(\beta, \omega_n)$, it is possible to get close to the asymptotic state $\varrho_S^*(\beta, \omega_n + \epsilon_n)$; if the temperature is finite, however, this state is not attainable (even asymptotically). Following the arguments in Appendix A.3, i.e., considering a diverging number of machines, each of which having the part connected to the cold bath with energy gap $\omega_{c_n} = \omega_n + \epsilon_n$ and taking the limit of $\epsilon_n \rightarrow 0$, which one can only do if the hot bath temperature is infinite, allows one to cool perfectly in diverging time in the incoherent paradigm at the Landauer limit.

We now calculate the energy cost explicitly for the infinite-temperature heat bath case, precisely demonstrating attainability of the Landauer limit. We use a similar approach to that described in Appendix A.3: We have a diverging number of cold machines for each energy gap ω_n , with which the target system at the $n - 1^{\text{th}}$ time step interacts; for an infinite-temperature heat bath, i.e., \mathcal{H} is in the maximally mixed state independent of its energy structure, the state of the target system at each step $\varrho_S^*(\beta, \omega_{n-1})$ is achievable. From Eq. (A.138), the energy change between all subsystems for a given step of the

protocol, i.e., taking $\varrho_S^*(\beta, \omega_{n-1}) \rightarrow \varrho_S^*(\beta, \omega_n)$, can be calculated as

$$\begin{aligned}\Delta E_S^{(n)} &= \text{tr} \left[H_S(\omega_S) (\varrho_S^*(\beta, \omega_n) - \varrho_S^*(\beta, \omega_{n-1})) \right], \\ \Delta E_C^{(n)} &= -\text{tr} \left[H_C(\omega_n) (\varrho_S^*(\beta, \omega_n) - \varrho_S^*(\beta, \omega_{n-1})) \right], \\ \Delta E_{\mathcal{H}}^{(n)} &= \text{tr} \left[H_{\mathcal{H}}(\omega_n - \omega_S) (\varrho_S^*(\beta, \omega_n) - \varrho_S^*(\beta, \omega_{n-1})) \right].\end{aligned}\quad (\text{A.155})$$

In general, i.e., for finite-temperature heat baths, we would have $\omega_n = \omega_{n-1} + \epsilon_{n-1}$, with a lower bound on ϵ_{n-1} for cooling to be possible [in accordance with Eq. (A.153)]. However, for infinite-temperature heat baths, this lower bound trivialises since the energy structure of the hot machine subsystem plays no role in its state; thus we can choose the energy gap structure for the machines as $\{\omega_n = \omega_S + n\epsilon\}_{n=1}^N$ with ϵ arbitrarily small. Taking the limit $\epsilon \rightarrow 0$, the diverging time limit $N \rightarrow \infty$, and writing $\omega_N = \omega_{\max}$ for the maximum energy gap of the cold machine subsystems, the energy exchanged throughout the entire cooling protocol here is given by

$$\begin{aligned}\Delta E_S &= \lim_{N \rightarrow \infty} \sum_{n=1}^N \Delta E_S^{(n)} = \text{tr} \left[H_S(\omega_S) (\varrho_S^*(\beta, \omega_{\max}) - \varrho_S^*(\beta, \omega_S)) \right] \\ \Delta E_C &= \lim_{N \rightarrow \infty} \sum_{n=1}^N \Delta E_C^{(n)} = \frac{1}{\beta} \left\{ S[\varrho_S^*(\beta, \omega_S)] - S[\varrho_S^*(\beta, \omega_{\max})] \right\} = \frac{1}{\beta} \tilde{\Delta} S_S \\ \Delta E_{\mathcal{H}} &= \lim_{N \rightarrow \infty} \sum_{n=1}^N \Delta E_{\mathcal{H}}^{(n)} = -\Delta E_S - \Delta E_C.\end{aligned}\quad (\text{A.156})$$

Here, the expression for ΔE_C can be derived using the same arguments as presented in Appendix A.3.1. In particular, the heat dissipated by the cold part of the machine, which is naturally connected to the heat sink in the incoherent setting as an infinite-temperature heat-bath can be considered a work source since any energy drawn comes with no entropy change, is in accordance with the Landauer limit. It is straightforward to obtain the same result for qudit systems. Lastly, in a similar way to the other protocols we have presented, one could compress all of the diverging number of operations into a single one whose control complexity diverges, thereby trading off between time and control complexity.

A.6.4 An Analysis of Finite-Temperature Heat Baths

We now return to the more general consideration of finite-temperature heat baths, i.e., $0 < \beta_H \leq \beta$. In the case where $\beta_H = \beta$, from Eq. (A.154), it is straightforward to see that for any machine energy gap ω_n , the effective gap ω_n^{eff} is equal to the gap of the target system, which means that no cooling can be achieved in the incoherent paradigm. Nonetheless, for any \mathcal{H} subsystem coupled to a heat bath of inverse temperature $\beta_H < \beta$, cooling is possible. We first provide more detail regarding why cooling at the Landauer limit is not possible in this setting, before deriving the minimal energy cost in accordance with the Carnot-Landauer limit presented in Theorem 1.4; in Appendix A.7, we provide

explicit protocols that saturate this bound for any finite-temperature heat bath and arbitrary finite-dimensional systems and machines.

Suppose that at some step n one has the initial joint state of Eq. (A.152), where $\epsilon_n = \gamma(\omega_n - \omega_S) + \epsilon$ and $\omega_n = \omega_S + n\epsilon$. Here, γ is as in Eq. (A.153). We now wish to cool the target system to $\varrho_S^*(\beta, \omega_n + \epsilon)$. For cooling to be possible in the incoherent setting here, we need the cold machine subsystem to have an energy gap of at least $\omega_n + \epsilon_n$; moreover, with a finite-temperature heat bath, this energy gap is insufficient to achieve the desired transformation [see Eq. (A.153)]. Based on Eq. (A.138), we can see that nonetheless, if we calculate the hypothetical energy change in this scenario if it were possible, we can derive a lower bound for the actual energy cost incurred. Employing Eq. (A.155), we have

$$\begin{aligned}
 \Delta E_c^{(n+1)} &\geq -\text{tr}\{H_c(\omega_n + \epsilon_n)[\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\} \\
 &= -\text{tr}\{H_c[(\gamma + 1)\omega_n - \gamma\omega_S + \epsilon][\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\} \\
 &= -\text{tr}\{H_c[(\gamma + 1)\omega_n - \gamma\omega_S + \epsilon + \gamma\epsilon - \gamma\epsilon][\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\} \\
 &= -(\gamma + 1)\text{tr}\{H_c(\omega_n + \epsilon)[\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\} \\
 &\quad + \gamma\text{tr}\{H_c(\omega_S + \epsilon)[\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\} \\
 &= (\gamma + 1)\Delta E_c^{*(n+1)} + \gamma\Delta E_S^{*(n+1)} + \gamma\text{tr}\{H_c(\epsilon)[\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\},
 \end{aligned} \tag{A.157}$$

where we have made use of the fact that for equally spaced Hamiltonians, the structure of the Hamiltonians on each subsystem take the same form [i.e., we can write, with slight abuse of notation, $H_c(\omega + \omega_S) = H_c(\omega) + H_S(\omega_S)$]. We use the star in $\Delta E_{\mathcal{A}}^*$ to denote the idealised energy cost [i.e., that corresponding to what would be achievable in the infinite-temperature setting; see Eq. (A.155)] and the energy costs without the star to represent those for when the temperature of the heat bath is finite. The additional term $\text{tr}\{H(\gamma\epsilon)[\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\}$ vanishes for $\epsilon \rightarrow 0$.

Summing up these contributions for a diverging number of steps gives the lower bound for the heat dissipated throughout the entire protocol for cooling an initial state $\tau_S(\beta, \omega_S)$ to some final $\tau_S(\beta_{\max}, \omega_S)$ is given by

$$\begin{aligned}
 \Delta E_c &= \lim_{N \rightarrow \infty} \sum_{n=1}^N \Delta E_c^{(n+1)} \\
 &\geq (\gamma + 1) \frac{1}{\beta} \tilde{\Delta} S_S + \gamma \Delta E_S \\
 &= \frac{1}{\beta} \tilde{\Delta} S_S + \gamma \left(\Delta E_S + \frac{1}{\beta} \tilde{\Delta} S_S \right).
 \end{aligned} \tag{A.158}$$

Note that for infinite-temperature heat baths, $\gamma \rightarrow 0$ and the usual Landauer limit is recovered; nonetheless, for finite-temperature heat baths, $\gamma > 0$ and there is an additional energy contribution, implying that the Landauer limit cannot be achieved. Moreover, note that the expression inside the parenthesis in the second term above is always non-negative,

as it is the free energy difference of the system during the cooling process. Lastly, it is straightforward to show that this lower bound is equivalent to the Carnot-Landauer limit in Eq. (A.14), which was derived in a protocol-independent manner as the ultimate limit in the incoherent-control setting. We now present explicit protocols that saturate this bound.

A.7 Perfect Cooling at the Carnot-Landauer Limit in the Incoherent-Control Paradigm

The precise statement that we wish to prove regarding saturation of the Carnot-Landauer limit is the following:

Lemma A.3. *For any $\beta^* \geq \beta > \beta_H$ and $\epsilon_{1,2} > 0$, there exists a cooling protocol in the incoherent-control setting comprising a number of unitaries of finite control complexity which, when the number of operations diverges, cools to some final temperature β' that is arbitrarily close to the ideal temperature value β^* , i.e.,*

$$|\beta' - \beta^*| < \epsilon_1, \quad (\text{A.159})$$

with an energy cost, measured as heat drawn from the hot bath, that is arbitrarily close to the ideal Carnot-Landauer limit, i.e.,

$$|\Delta E_{\mathcal{H}} - \eta^{-1} \tilde{\Delta} F_S^{(\beta)}| < \epsilon_2, \quad (\text{A.160})$$

where $\eta = 1 - \frac{\beta_H}{\beta}$ and $\Delta F_S^{(\beta)} = F_\beta(\varrho'_S) - F_\beta(\varrho_S)$ is the free energy difference between the initial $\varrho_S = \tau_S(\beta, H_S)$ and final $\varrho'_S = \tau_S(\beta^*, H_S)$ system states (w.r.t. inverse temperature β).

We begin by presenting the diverging-time protocol that saturates the Carnot-Landauer limit when all three subsystems $\mathcal{S}, \mathcal{C}, \mathcal{H}$ are qubits. The simplicity of this special case allows us to calculate precisely bounds on the number of operations required to reach any chosen error threshold. Building on this intuition, we then present the generalisation to the case where all systems are qudits. The protocols with diverging control complexity follow directly via the same line of reasoning presented in the main text.

A.7.1 Qubit Case

We begin with setting some notation and intuition for the proof, before expanding on mathematical details.

Sketch of Protocol.—The protocol consists of the following. There are N stages, each labelled by $n \in \{1, 2, \dots, N\}$. Each stage proceeds as follows:

1. A qubit with energy gap $\omega_S + n\theta$ is taken from the cold part of the machine, and a qubit with energy gap $n\theta$ is taken from the hot part (see below). The initial state of the machine at the beginning of the n^{th} stage is thus $\tau_C(\beta, \omega_S + n\theta) \otimes \tau_H(\beta_H, n\theta)$.
2. The energy-preserving three qubit unitary cycle in the $\{010, 101\}_{\mathcal{SCH}}$ subspace is performed [see Eq. (A.136)], after which the cold and hot qubits are rethermalised to their respective initial temperatures.
3. The above steps are repeated m_n times.

The energy increment θ is defined as

$$\theta := \frac{\omega_S}{N} \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right), \quad (\text{A.161})$$

while the number of repetitions within each stage is given by

$$m_n = \left\lceil \frac{\log(\delta)}{\log(1 - N_V^{(n)})} \right\rceil. \quad (\text{A.162})$$

$\lceil \cdot \rceil$ is the ceiling function, and $N_V^{(n)}$ is the sum of the initial thermal populations in the $\{01, 10\}_{\mathcal{CH}}$ subspace of the machine, i.e.,

$$N_V^{(n)} := \langle 01 | \tau_C(\beta, \omega_S + n\theta) \otimes \tau_H(\beta_H, n\theta) | 01 \rangle + \langle 10 | \tau_C(\beta, \omega_S + n\theta) \otimes \tau_H(\beta_H, n\theta) | 10 \rangle. \quad (\text{A.163})$$

The parameter δ will be chosen appropriately to complete the proof ($\delta = \frac{1}{N^2}$ works).

The intuition for the proof is as follows. We first consider how the populations of the target system changes in the idealised protocol where $m_n \rightarrow \infty$, so that in each stage, the system reaches the virtual temperature determined by the \mathcal{CH} qubits. We can use this ideal setting to find expressions for the final temperature and energy cost, which serves as a baseline that we wish to attain to within arbitrary precision. We then consider the protocol as constructed above with a finite number of repetitions m_n in each stage, and show that its expressions for temperature and work cost are close (w.r.t. $1/N$) to the original expressions, and by taking N to be sufficiently large but still finite (i.e., in the diverging time limit), we prove that the protocol can be arbitrarily close in temperature and energy cost to the ideal values.

Proof. We label the population in the excited state level of the target system at the end of stage n as p_n . Thus p_0 is the initial population and p_N is the final population in the excited level of the target system qubit, i.e., that spanned by $|1\rangle\langle 1|_S$. We also label by q_n what the corresponding population p_n would hypothetically be in the limit $m_n \rightarrow \infty$. This value can be calculated by matching the temperature of the target system qubit to

the temperature of the $\{01, 10\}_{\mathcal{CH}}$ virtual qubit within the machine (see Appendix G in Ref. [83]). Thus q_n is defined via the Gibbs ratio

$$\frac{q_n}{1 - q_n} = e^{-\beta(\omega_S + n\theta)} e^{+\beta_H n\theta} = e^{-\beta\omega_S} e^{-(\beta - \beta_H)n\theta}. \quad (\text{A.164})$$

Note that

1. $\{p_n\}, \{q_n\}$ are both monotonically decreasing sequences, as each stage cools the target qubit further.
2. $p_n > q_n$ for all n , as more repetitions within each stage keep cooling the target qubit further.

To keep track of the energetic resource cost, which we will take here to be the total heat drawn from the hot bath, we must sum the energetic contribution from each time the hot qubit is rethermalised to β_H after the application of the three-party cycle unitary. Due to the fact that the only manner in which the population of the hot qubit changes is due to the $\{010, 101\}_{\mathcal{SCH}}$ exchange, it follows that any population change in the hot qubit is identical to the population change in the target system qubit.

Focusing on a single stage, where the machine qubits are fixed in energy gap, the total population change in the hot qubit that must be restored by the hot bath is therefore equal to the population change in the target system throughout that stage. The heat drawn from the hot bath throughout the entire stage is therefore

$$\tilde{\Delta}E_{\mathcal{H}}^{(n)} = \omega_{\mathcal{H}}^{(n)}(p_{n-1} - p_n) = n\theta(p_{n-1} - p_n). \quad (\text{A.165})$$

With these expressions derived, we can study the properties of the abstract protocol where the number of repetitions within each stage goes to infinity: $m_n \rightarrow \infty$. First, the final temperature asymptotically achieved here is given by finding the temperature $\tilde{\beta}$ associated with the qubit with excited state population q_N

$$\frac{q_N}{1 - q_N} = e^{-\tilde{\beta}\omega_S} \Rightarrow e^{-\beta\omega_S} e^{-(\beta - \beta_H)N\theta} = e^{-\tilde{\beta}\omega_S} \Rightarrow \tilde{\beta} = \beta^*, \quad (\text{A.166})$$

where we have made use of the definition of θ in Eq. (A.161). We can thus identify $q_N = q^*$, since it is the population associated with the ideal final temperature β^* .

We also have the following expression for the total energetic cost of the ideal protocol after N stages

$$\tilde{\Delta}E_{\mathcal{H}}^* = \sum_{n=1}^N n\theta(q_{n-1} - q_n), \quad (\text{A.167})$$

which can alternatively be expressed as

$$\tilde{\Delta}E_{\mathcal{H}}^* = \sum_{n=1}^N [(n-1)\theta(q_{n-1} - q_n)] + \theta(q_0 - q_N). \quad (\text{A.168})$$

The sums appearing in the two alternate expressions are the left and right Riemann sums of the integral of the variable $y = n\theta$ integrated with respect to the variable q , i.e.,

$$I := - \int_{q_0}^{q^*} y \, dq, \quad \text{where} \quad \frac{q(y)}{1 - q(y)} = e^{-\beta\omega_S} e^{-(\beta - \beta_H)y}, \quad (\text{A.169})$$

from Eq. (A.164). For $y > 0$, $q(y)$ is monotonically decreasing and so the converse is also true, i.e., y is monotonically decreasing w.r.t. $q(y)$. This implies that the integral is bounded by the left and right Riemann sums, so we have

$$\sum_{n=1}^N (n-1)\theta(q_{n-1} - q_n) \leq \int_{q_0}^{q^*} y \, dq \leq \sum_{n=1}^N n\theta(q_{n-1} - q_n), \quad (\text{A.170})$$

from which we can deduce that the value of $\Delta E_{\mathcal{H}}^*$ is itself is bounded both ways from Eqs. (A.167) and (A.168):

$$\int_{q_0}^{q^*} y \, dq \leq \tilde{\Delta} E_{\mathcal{H}}^* \leq \int_{q_0}^{q^*} y \, dq + \theta(q_0 - q^*). \quad (\text{A.171})$$

The integral itself can be expressed in terms of the free energy of the qubit target system with respect to the environment inverse temperature β . Expressing the free energy as a function of the excited state population q and differentiating w.r.t. q gives

$$F(q) = \langle E \rangle(q) - \frac{S(q)}{\beta} = q \omega_S + \frac{1}{\beta} [q \log(q) + (1-q) \log(1-q)]. \quad (\text{A.172})$$

$$\frac{\partial F}{\partial q} = \omega_S + \frac{1}{\beta} \log\left(\frac{q}{1-q}\right) = \left(\omega_S + \frac{1}{\beta} (-\beta\omega_S - (\beta - \beta_H)y)\right) = -\frac{\beta - \beta_H}{\beta} y. \quad (\text{A.173})$$

Using the above expression, the definite integral in Eq. (A.169) amounts to

$$I = \frac{1}{\eta} [F(q^*) - F(q_0)] =: \frac{1}{\eta} (F^* - F_0), \quad (\text{A.174})$$

where we have identified the Carnot efficiency $\eta = 1 - \beta_H/\beta$ and for ease of notation written $F^* := F(q^*)$ and $F_0 := F(q_0)$. Thus we can bound $\tilde{\Delta} E_{\mathcal{H}}^*$ on both sides

$$\frac{1}{\eta} (F^* - F_0) \leq \tilde{\Delta} E_{\mathcal{H}}^* \leq \frac{1}{\eta} (F^* - F_0) + \theta(q_0 - q^*) \leq \frac{1}{\eta} (F^* - F_0) + \frac{\omega_S}{N} \left(\frac{\beta^* - \beta}{\beta - \beta_H}\right), \quad (\text{A.175})$$

where the rightmost inequality follows from the fact that $\{q_n\}$ forms a decreasing sequence.

We now proceed to consider the cooling protocol with a finite number of repetitions m_n within each stage. We first bound the difference between p_n and q_n . Using the properties of the exchange unitary under repetitions [83, 366] (in particular, see Appendix G in Ref. [83]), we have that in each stage

$$\frac{p_n - q_n}{p_{n-1} - q_n} = \left(1 - N_V^{(n)}\right)^{m_n}. \quad (\text{A.176})$$

Thus, the population difference to the asymptotically achievable population given by the virtual temperature shrinks as a power law w.r.t. the number of repetitions. Since $0 < N_V^{(n)} < 1$ (all strict inequalities), three points follow: First, the population q_n can never be attained with a finite number of steps within the stage n ; Second, that every repetition cools the system further by some finite amount; Third, that one can get arbitrarily close to q_n by taking m_n sufficiently large. In fact, by our definition of m_n , we have that

$$\frac{p_n - q_n}{p_{n-1} - q_n} \leq \delta. \quad (\text{A.177})$$

From this, we can prove that

$$p_n - q_n \leq \delta^n q_0 - \delta q_n + (1 - \delta) \delta \sum_{j=1}^{n-1} \delta^{n-j-1} q_j. \quad (\text{A.178})$$

The proof is by induction. For $n = 0$, $p_0 = q_0$ (initial state), and for $n = 1$, using Eq. (A.177)

$$p_1 - q_1 \leq \delta(p_0 - q_1) = \delta(q_0 - q_1). \quad (\text{A.179})$$

Suppose that the above statement holds true for p_k . Then from Eq. (A.177)

$$\begin{aligned} p_{k+1} - q_{k+1} &\leq \delta(p_k - q_{k+1}) \\ &= \delta(p_k - q_k + q_k - q_{k+1}) \\ &\quad \vdots \\ &\leq \delta^{k+1} q_0 - \delta q_{k+1} (1 - \delta) \delta + \sum_{j=1}^{(k+1)-1} \delta^{(k+1)-j-1} q_j. \end{aligned} \quad (\text{A.180})$$

With this result, we can now bound the difference between the energy cost of this finite-repetition protocol and that of the idealised one. We now proceed to prove that

$$\begin{aligned} \tilde{\Delta} E_{\mathcal{H}} - \tilde{\Delta} E_{\mathcal{H}}^* &= \sum_{n=1}^N n\theta(p_{n-1} - p_n) - \sum_{n=1}^N n\theta(q_{n-1} - q_n) \\ &\leq \theta \left(q_0 \sum_{j=1}^{N-1} \delta^{N-j} - \sum_{j=1}^{N-1} \delta^{N-j} q_j \right). \end{aligned} \quad (\text{A.181})$$

We again use proof by induction. First note that we can rewrite

$$\sum_{n=1}^N n\theta(f_{n-1} - f_n) = \theta \left(\sum_{n=1}^N f_{n-1} \right) - N\theta f_N, \quad (\text{A.182})$$

for $f_n \in \{p_n, q_n\}$. Therefore, we can rewrite the difference

$$\tilde{\Delta} E_{\mathcal{H}} - \tilde{\Delta} E_{\mathcal{H}}^* = \theta \sum_{n=1}^N (p_{n-1} - q_{n-1}) - N\theta(p_N - q_N) \leq \theta \left(\sum_{n=1}^N (p_{n-1} - q_{n-1}) \right), \quad (\text{A.183})$$

since the last subtracted term is always strictly positive. Consider now the partial sum

$$\mathcal{E}_k = \sum_{n=1}^k (p_{n-1} - q_{n-1}). \quad (\text{A.184})$$

For $k = 1$, $\mathcal{E}_1 = 0$, since $p_0 = q_0$. For $k = 2$, we have

$$\mathcal{E}_1 = (p_1 - q_1) \leq \delta(q_0 - q_1) = \left(q_0 \sum_{j=1}^1 \delta^{2-j} - \sum_{j=1}^1 \delta^{2-j} q_j \right), \quad (\text{A.185})$$

which matches the hypothesis of Eq. (A.181). Assuming that the same holds true for \mathcal{E}_k , then for \mathcal{E}_{k+1} , we have

$$\begin{aligned} \mathcal{E}_{k+1} &= \mathcal{E}_k + (p_k - q_k) \\ &\leq \left(q_0 \sum_{j=1}^{k-1} \delta^{k-j} - \sum_{j=1}^{k-1} \delta^{k-j} q_j \right) + \left(\delta^k q_0 + (1 - \delta) \delta \sum_{j=1}^{k-1} \delta^{k-j-1} q_j - \delta q_k \right) \\ &\quad \vdots \\ &= q_0 \sum_{j=1}^k \delta^{k+1-j} - \sum_{j=1}^k \delta^{k+1-j} q_j. \end{aligned} \quad (\text{A.186})$$

Then, by dropping the second sum, which is a strictly positive quantity, the difference in Eq. (A.181) can be further simplified to

$$\tilde{\Delta} E_{\mathcal{H}} - \tilde{\Delta} E_{\mathcal{H}}^* \leq \theta q_0 \sum_{j=1}^{N-1} \delta^{N-j} = \theta q_0 \delta \sum_{k=0}^{N-2} \delta^k < \theta q_0 \delta (N-1) < \theta q_0 \delta N < \omega_S \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right) \delta, \quad (\text{A.187})$$

where we have used that $\delta < 1$. Finally, to upper bound the number of operations required in the protocol, we bound the number of repetitions within each stage by bounding the total population of the virtual qubit spanned by the levels $\{01, 10\}_{\mathcal{CH}}$ as follows:

$$\begin{aligned} N_V^{(n)} &= \langle 01 | \tau_C(\beta, \omega_S + n\theta) \otimes \tau_H(\beta_H, n\theta) | 01 \rangle + \langle 10 | \tau_C(\beta, \omega_S + n\theta) \otimes \tau_H(\beta_H, n\theta) | 10 \rangle \\ &= \frac{e^{-\beta_H n\theta} + e^{-\beta(\omega_S + n\theta)}}{(1 + e^{-\beta_H n\theta})(1 + e^{-\beta(\omega_S + n\theta)})} \\ &> \frac{e^{-\beta(\omega_S + n\theta)}}{4}. \end{aligned} \quad (\text{A.188})$$

From this, it follows that

$$\begin{aligned} \log \left[1 - N_V^{(n)} \right] &< \log \left[1 - \frac{e^{-\beta(\omega_S + n\theta)}}{4} \right] \\ &< -\frac{e^{-\beta(\omega_S + n\theta)}}{4} \quad \text{if } x \in (0, 1) \Rightarrow \log(1 - x) < -x. \\ \Rightarrow -\frac{1}{\log \left[1 - N_V^{(n)} \right]} &< 4e^{+\beta(\omega_S + n\theta)}. \end{aligned} \quad (\text{A.189})$$

Thus we can bound the number of repetitions in each stage from Eq. (A.162). Noting that $\log(\delta) < 0$, we have

$$m_n < 4 \log\left(\frac{1}{\delta}\right) e^{+\beta(\omega_S + n\theta)} + 1. \quad (\text{A.190})$$

For a crude bound, we can replace n by its maximum value N , and sum over all the stages to find an upper bound on the total number of three-qubit exchange unitaries implemented throughout the entire protocol, which gives

$$M = \sum_{n=1}^N m_n < N \left[4 \log(1/\delta) e^{+\beta(\omega_S + N\theta)} + 1 \right] = N \left[4 \log(1/\delta) e^{\omega_S(\beta^* - \beta_H)/\eta} + 1 \right]. \quad (\text{A.191})$$

Also, note that $\lim_{\delta \rightarrow 0} p_N = q_N = q^*$. More precisely, using Eq. (A.178), we have

$$\begin{aligned} p_N - q^* &< \delta \left(\delta^{N-1} q_0 + (1 - \delta) \sum_{j=1}^{N-1} \delta^{n-j-1} q_j - q_N \right) \\ &< \delta (1 + (1 - \delta)(N - 1)) < \delta N. \end{aligned} \quad (\text{A.192})$$

In summary, we have the following bounds on the protocol in which each stage consists of a finite number of steps

$$\begin{aligned} p_N - q^* &< \delta N, \\ \tilde{\Delta} E_{\mathcal{H}} &< \frac{1}{\eta} (F^* - F_0) + \omega_S \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right) \left(\frac{1}{N} + \delta \right), \end{aligned} \quad (\text{A.193})$$

where we have combined Eqs. (A.175) and (A.187) for the second expression. For simplicity, we choose $\delta = \frac{1}{N^2}$, so that

$$\begin{aligned} p_N - q^* &< \frac{1}{N}, \\ \tilde{\Delta} E_{\mathcal{H}} &< \frac{1}{\eta} (F^* - F_0) + \omega_S \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right) \left(\frac{2}{N} \right). \end{aligned} \quad (\text{A.194})$$

Thus, given any final temperature (encoded by the population q^*), and allowed errors ϵ_1 and ϵ_2 for the final population and energy cost respectively, one can always choose N large enough so that both quantities are within the error threshold. Specifically, choosing N as

$$N = \left\lceil \max \left\{ \epsilon_1^{-1}, 2\omega_S \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right) \epsilon_2^{-1} \right\} \right\rceil, \quad (\text{A.195})$$

we automatically have that $p_N - q^* < \epsilon_1$ and $\Delta E_{\mathcal{H}} < \frac{F^* - F_0}{\eta} + \epsilon_2$. The total number of unitary operations (each of which is followed by rethermalisation of the machine) is then bounded by Eq. (A.191)

$$M < N \left(8 \log[N] e^{\omega_S(\beta^* - \beta_H)/\eta} + 1 \right). \quad (\text{A.196})$$

We can see from Theorem A.2 that the protocol is asymptotically optimal with respect to the energy extracted from the hot bath. \square

A.7.2 Qudit Case

The extension of the proof above to the case of qudits is nontrivial. This is because, while for qubits there is only one energy-resonant subspace that leads to cooling and hence a unique protocol [see Eq. (A.136)] that asymptotically attains perfect cooling at the Carnot-Landauer bound, this is no longer the case for higher-dimensional systems; here, there can be a number of energy-resonant subspaces that cool the target and the question of optimality hinges crucially on the complex energy-level structure of all systems involved. Hence, it is not possible to provide a unique unitary that generates the optimal protocol independently of the subsystem Hamiltonians. Nonetheless, we slightly modify the protocol for the qubit case above to be implemented on a number of particular three-qubit subspaces of the three-qudit global state such that, at the end of each stage, the state of the target system is arbitrarily close to the (known) state which would be achieved in an abstract protocol in the diverging-time limit. This asymptotically-attainable state is precisely that which would be achieved in the coherent-control paradigm with a machine the same dimension as the joint hot-cold qudits. Thus, we will first begin by presenting the necessary steps for the proof in the coherent-control setting, which we will then adapt as appropriate for the incoherent setting control. Finally, summing the overall energy cost of said protocol over all stages saturates the Carnot-Landauer bound, as required.

Proof. An idealised sequence of temperatures and system states.—We construct the incoherent protocol in the following manner. We will seek to take the system through a sequence of thermal states starting at inverse temperature β and ending at inverse temperature β^* with N equally spaced intermediary steps, i.e.,

$$\beta_n = \beta + n\theta(\beta - \beta_{\mathcal{H}}), \quad (\text{A.197})$$

$$\theta = \frac{1}{N} \left(\frac{\beta^* - \beta}{\beta - \beta_{\mathcal{H}}} \right), \quad (\text{A.198})$$

so that $\beta_N = \beta^*$ by construction. This corresponds to taking the system through the following sequence of thermal states

$$\varrho_S^{(n)} = \frac{e^{-\beta_n H_S}}{\mathcal{Z}_S(H_S, \beta_n)}. \quad (\text{A.199})$$

Note that, in contrast to the coherent protocol where such a sequence can be traversed by simply swapping the target system with a sequence of appropriate machines, in the incoherent setting such a protocol is generally not possible as such swaps are not energy conserving. Nonetheless, we will develop a modified protocol that is energy conserving and mimics this idealised one.

Corresponding to each step in the sequence, we define the following quantity, which we will eventually show to be related to the heat drawn from the hot bath:

$$G^{(n)} = -n\theta \Delta E_S^{(n)} = -n\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n)} - \varrho_S^{(n-1)} \right) \right]. \quad (\text{A.200})$$

We proceed to show that the total $\sum_n G^{(n)}$ that we label *the idealised heat cost* $\tilde{\Delta} E_{\mathcal{H}}^*$ is close to the free energy difference over the entire sequence. We have

$$\begin{aligned} \tilde{\Delta} E_{\mathcal{H}}^* &= \sum_{n=1}^N G^{(n)} \\ &= \sum_{n=1}^N n\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n-1)} - \varrho_S^{(n)} \right) \right] \end{aligned} \quad (\text{A.201})$$

$$= \left\{ \sum_{n=1}^N (n-1)\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n-1)} - \varrho_S^{(n)} \right) \right] \right\} + \theta \operatorname{tr} \left[H_S \left(\varrho_S^{(0)} - \varrho_S^{(N)} \right) \right]. \quad (\text{A.202})$$

The sums on the second and third lines above, (A.201) and (A.202) respectively, are the right and left Riemann sums corresponding to the following integral:

$$\begin{aligned} I &= \int_{q_i}^{q_f} q (-dx) = \int_{q_f}^{q_i} q dx, \\ \text{where } n\theta &\rightarrow q, \\ x &= \operatorname{tr} [H_S \varrho_S(q)], \\ \varrho_S(q) &= \frac{e^{-[\beta+q(\beta-\beta_{\mathcal{H}})]H_S}}{\operatorname{tr} [e^{-[\beta+q(\beta-\beta_{\mathcal{H}})]H_S}]}. \end{aligned} \quad (\text{A.203})$$

We observe that x is the average energy of the thermal state of temperature $\beta + q(\beta - \beta_{\mathcal{H}})$, and thus x and q are strictly monotonically decreasing w.r.t. each other (which explains why the left and right sums are switched). It follows that the Riemann sums bound the integral

$$\sum_{n=1}^N (n-1)\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n-1)} - \varrho_S^{(n)} \right) \right] \leq \int_{q_f}^{q_i} q dx \leq \sum_{n=1}^N n\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n-1)} - \varrho_S^{(n)} \right) \right]. \quad (\text{A.204})$$

We can thus bound the idealised heat cost in both directions via

$$I \leq \tilde{\Delta} E_{\mathcal{H}}^* \leq I + \theta \operatorname{tr} \left[H_S \left(\varrho_S^{(0)} - \varrho_S^{(N)} \right) \right]. \quad (\text{A.205})$$

The integral in Eq. (A.203) can be shown to be equal to the change in free energy of the target system (w.r.t. inverse temperature β)

$$\begin{aligned} F_{\beta}[\varrho_S(q)] &= \operatorname{tr} [H_S \varrho_S(q)] + \frac{1}{\beta} \operatorname{tr} [\varrho_S(q) \log \varrho_S(q)], \\ \frac{d}{dq} F_{\beta}[\varrho_S(q)] &= \operatorname{tr} \left[\left(H_S + \frac{\mathbb{1}_S + \log \varrho_S(q)}{\beta} \right) \frac{d\varrho_S(q)}{dq} \right]. \end{aligned} \quad (\text{A.206})$$

Note that $\varrho_S(q)$ and $d\varrho_S(q)$ are both always diagonal in H_S and full rank for all $q \in \mathbb{R}$, so we have no problems with $\log \varrho_S(q)$, and all of the operators in the expression are

well-defined and commute. Proceeding, we repeatedly use $\text{tr}[\text{d}\varrho_S(q)] = \text{d}\text{tr}[\varrho_S(q)] = 0$ and label the partition function $\mathcal{Z}(q) := \text{tr}\left[e^{-[\beta+q(\beta-\beta_{\mathcal{H}})]H_S}\right]$ to obtain

$$\begin{aligned} \frac{\text{d}}{\text{d}q} F_\beta[\varrho_S(q)] &= \text{tr}\left[\left(H_S + \frac{\log \varrho_S(q)}{\beta}\right) \frac{\text{d}\varrho_S(q)}{\text{d}q}\right] \\ &= \text{tr}\left[\left(H_S - \frac{\beta + q(\beta - \beta_{\mathcal{H}})}{\beta} H_S - \mathbb{1}_S \frac{\log \mathcal{Z}(q)}{\beta}\right) \frac{\text{d}\varrho_S(q)}{\text{d}q}\right] \\ &= -q \left(1 - \frac{\beta_{\mathcal{H}}}{\beta}\right) \frac{\text{d}}{\text{d}q} \text{tr}[H_S \varrho_S(q)] = -q\eta \frac{\text{d}x}{\text{d}q}, \end{aligned} \quad (\text{A.207})$$

where we have identified the Carnot efficiency η for an engine operating between β and $\beta_{\mathcal{H}}$. The integral thus simplifies to

$$I = \eta^{-1} \left(F_\beta[\varrho_S(q_f)] - F_\beta[\varrho_S(q_i)]\right) =: \eta^{-1} \Delta F_S^{(\beta)}. \quad (\text{A.208})$$

The idealised heat cost is thus bounded by

$$\eta^{-1} \Delta F_S^{(\beta)} \leq \tilde{\Delta} E_{\mathcal{H}}^* \leq \eta^{-1} \Delta F_S^{(\beta)} + \theta \text{tr}\left[\mathcal{H}_S \left(\varrho_S^{(0)} - \varrho_S^{(N)}\right)\right]. \quad (\text{A.209})$$

The left inequality is Landauer's bound applied to cooling a target system with Hamiltonian H_S (see Theorem 1.4), and the error term on the right can be bounded quite easily; for instance, for $\beta > 0$, we have

$$\begin{aligned} \text{tr}\left[\mathcal{H}_S \left(\varrho_S^{(0)} - \varrho_S^{(N)}\right)\right] &= \text{tr}\left[\left(\mathcal{H}_S - E_S^{\min} \mathbb{1}_S\right) \left(\varrho_S^{(0)} - \varrho_S^{(N)}\right)\right] \\ &\leq \text{tr}\left[\left(\mathcal{H}_S - E_S^{\min} \mathbb{1}_S\right) \varrho_S^{(0)}\right] \\ &\leq \text{tr}\left[\left(\mathcal{H}_S - E_S^{\min} \mathbb{1}_S\right) \frac{\mathbb{1}_S}{d_S}\right] \leq \frac{\omega_S^{\max}}{d_S}, \end{aligned} \quad (\text{A.210})$$

where $\omega_S^{\max} := E_S^{\max} - E_S^{\min}$ is the largest energy gap in the target system Hamiltonian, d_S is the system dimension, and the second line follows from the fact that $\mathcal{H}_S - E_S^{\min} \mathbb{1}_S$ is a positive semidefinite operator. We have used the fact that since $\varrho_S^{(0)}$ is a thermal state of positive temperature, its average energy is less than that of the infinite temperature thermal state, $\frac{\mathbb{1}_S}{d_S}$. Since $\theta \propto \frac{1}{N}$, it follows that one can always find an N large enough such that the error is smaller than a given value, thereby saturating the Landauer bound.

A sequence of machine Hamiltonians to mimic the idealised sequence.—Next we construct a protocol that mimics the above sequence and obeys the global energy conservation condition imposed in the incoherent-control setting. The protocol is split into N stages (like above). In each stage, the Hamiltonian of the machine is fixed. The machine here comprises to two parts: The “cold” part and the “hot” part. The cold part is chosen to begin in a thermal state at temperature β of the Hamiltonian

$$H_c = (1 + n\theta) H_S. \quad (\text{A.211})$$

At this point we note that this sequence of cold machine states is exactly the same as in the coherent protocol, which would proceed by simply swapping the full state of target

system and machine in each stage. However, that is not possible here since this is not an energy preserving operation. To allow for energy preserving operations, the hot part of the machine consists of $\frac{d_S(d_S-1)}{2}$ qubits, each corresponding to a pair of levels (i, j) of the target system (henceforth we take $i < j$ to avoid double-counting), whose energy gap is equal to the difference in energies of the target and cold qubit subspaces (hence rendering the desired exchange energy-resonant)

$$H_{\mathcal{H}}^{(ij)} = [\omega_i + (1 + n\theta)\omega_j - (\omega_j + (1 + n\theta)\omega_i)] |1\rangle\langle 1|_{\mathcal{H}}^{(ij)} = n\theta(\omega_j - \omega_i) |1\rangle\langle 1|_{\mathcal{H}}^{(ij)}, \quad (\text{A.212})$$

where we have labelled the energy eigenvalues of H_S by $\{\omega_i\}$. Each of these hot qubits begins at inverse temperature β_H . After every unitary operation, the cold and hot parts of the machine are rethermalised to their respective initial temperatures.

To understand the choice of machine Hamiltonians, consider the following two energy eigenstates of the machine: $|i\rangle_c \otimes |1\rangle_{\mathcal{H}}^{(ij)}$ and $|j\rangle_c \otimes |0\rangle_{\mathcal{H}}^{(ij)}$. The energy difference is

$$\Delta^{(ij)} = \omega_j(1 + n\theta) - \omega_i(1 + n\theta) - n\theta(\omega_j - \omega_i) = \omega_j - \omega_i, \quad (\text{A.213})$$

matching the energy difference between the corresponding pair of energy eigenstates of the target system. Furthermore, calculating the ratio of populations of the two levels we find

$$g^{(ij)} = \frac{e^{-\beta\omega_j(1+n\theta)}}{e^{-\beta\omega_i(1+n\theta)}e^{-\beta_{\mathcal{H}}n\theta(\omega_j-\omega_i)}} = e^{-(\omega_j-\omega_i)[\beta+n\theta(\beta-\beta_{\mathcal{H}})]}. \quad (\text{A.214})$$

This corresponds to the Gibbs ratio of a qubit at the temperature $\beta + n\theta(\beta - \beta_{\mathcal{H}})$, which is the temperature that defines stage n [see Eq. (A.197)]. In summary, we have constructed a machine featuring $\frac{d_S(d_S-1)}{2}$ qubit subspaces (or virtual qubits), each of the same energy gap as one pair of energy eigenstates of the system, and all of which have a Gibbs ratio (or virtual temperature) corresponding the n^{th} temperature of our desired sequence.

A single step of the protocol: The max-exchange operation.—Within each stage of the protocol, a single step consists of a unitary operation on \mathcal{SCH} , followed by the rethermalisation of the machine parts to their respective initial temperatures. We construct the unitary operation as follows: For every pair (i, j) of system energy levels, one can calculate the absolute value of the difference in populations of the following two degenerate eigenstates $|i\rangle_S |j\rangle_C |0\rangle_{\mathcal{H}}^{(ij)}$ and $|j\rangle_S |i\rangle_C |1\rangle_{\mathcal{H}}^{(ij)}$. This value corresponds to the amount of population that would move under an exchange $|i\rangle_S |j\rangle_C |0\rangle_{\mathcal{H}}^{(ij)} \leftrightarrow |j\rangle_S |i\rangle_C |1\rangle_{\mathcal{H}}^{(ij)}$. We then choose the pair with the largest absolute value of this difference and perform that exchange, with an identity operation applied to all other subspaces. We call this unitary operation the *max-exchange*. We proceed to prove two statements about the max-exchange operation. First, that the heat extracted from the hot bath is proportional to the change in average energy of the system; and second, that system state under repetition of said operation converges to the thermal state of the temperature that defines the stage n .

Consider the change in average energy of the target system under the exchange unitary. The only two populations that change are those of the $|i\rangle_S$ and $|j\rangle_S$. We label the increase

in the population of $|i\rangle_S$ as δp . Then, we have

$$\Delta E_S = \text{tr} \left[H_S \left(\varrho'_S - \varrho_S \right) \right] = -\delta p (\omega_j - \omega_i). \quad (\text{A.215})$$

On the other hand, the populations of the corresponding hot qubit (i.e., tracing out the target system and cold machine) change by the same amount, i.e., there is a move of δp from $|1\rangle_{\mathcal{H}}^{(ij)}$ to $|0\rangle_{\mathcal{H}}^{(ij)}$. In order to rethermalise the hot qubit, the heat drawn from the hot bath is thus

$$\tilde{\Delta} E_{\mathcal{H}} = \delta p n \theta (\omega_j - \omega_i) = -n \theta \Delta E_S. \quad (\text{A.216})$$

This is an expression conveniently independent of the pair (i, j) that applies after an arbitrary number of repetitions of the max-exchange operation (which will use different pairs in general).

Convergence of the max-exchange protocol to the virtual temperature.—To show that the max-exchange protocol indeed converges to the desired system state in each stage of the protocol, we first prove a rather general statement: Given a state ϱ diagonal in the energy eigenbasis, if we exchange any qubit subspace within this system with a virtual qubit of a particular virtual temperature, then the relative entropy of the target system w.r.t. the thermal state of that (virtual) temperature decreases.

To this end, consider the relative entropy of a state ϱ that is diagonal in the energy eigenbasis to a thermal state τ . Labelling the populations of ϱ as p_i and those of τ as q_i , this can be expressed as

$$D(\varrho||\tau) = \sum_k p_k \log \left(\frac{p_k}{q_k} \right). \quad (\text{A.217})$$

We now focus on a single qubit subspace labelled by $\{i, j\}$, which leads to

$$\begin{aligned} D(\varrho||\tau) &= p_i \log \left(\frac{p_i}{q_i} \right) + p_j \log \left(\frac{p_j}{q_j} \right) + \sum_{k \notin \{i, j\}} p_k \log \left(\frac{p_k}{q_k} \right) \\ &= (p_i + p_j) \left[\frac{p_i}{p_i + p_j} \log \left(\frac{\frac{p_i}{p_i + p_j} p_i + p_j}{\frac{q_i}{q_i + q_j} q_i + q_j} \right) + \frac{p_j}{p_i + p_j} \log \left(\frac{\frac{p_j}{p_i + p_j} p_i + p_j}{\frac{q_j}{q_i + q_j} q_i + q_j} \right) \right] \\ &\quad + \sum_{k \notin \{i, j\}} p_k \log \left(\frac{p_k}{q_k} \right) \\ &= N \left(\bar{p}_i \log \frac{\bar{p}_i}{\bar{q}_i} + \bar{p}_j \log \frac{\bar{p}_j}{\bar{q}_j} + \log \frac{N}{N_V} \right) + \sum_{k \notin \{i, j\}} p_k \log \frac{p_k}{q_k}. \end{aligned} \quad (\text{A.218})$$

In the last line we have renormalised the populations within the qubit subspace and labelled the total populations of the system and thermal state qubit subspaces of interest by N and N_V respectively. Labelling the normalised states within these subspaces as ϱ_V and τ_V respectively, we have

$$D(\varrho||\tau) = N \left[D(\varrho_V||\tau_V) + \log \left(\frac{N}{N_V} \right) \right] + \sum_{k \notin \{i, j\}} p_k \log \left(\frac{p_k}{q_k} \right). \quad (\text{A.219})$$

Suppose now that this qubit subspace of the target system is exchanged with a qubit subspace of any machine that has the same temperature as the thermal state above. The only object that changes in the the above expression is ϱ_V , since the norm N remains the same. In addition, ϱ_V always gets closer to τ_V under such an exchange [83, 366], implying that the relative entropy always strictly decreases under such an operation.

Returning to the max-exchange protocol, note that by construction, every virtual qubit in the machine that is exchanged with the qubit subspace $\{i, j\}$ of the target system in a given stage n has the same virtual temperature, $\beta_n = \beta + n\theta(\beta - \beta_{\mathcal{H}})$. Thus the relative entropy of the system to the thermal state at this temperature always decreases under this operation, unless the operation does not shift any population, which only happens at the unique fixed point where every qubit subspace of the system is already at the virtual temperature β_n . By monotone convergence, the relative entropy must converge, and moreover converge to the value that it has at the fixed point of the operation, which is the thermal state at inverse temperature β_n . Note that rather than choosing the qubit subspace with maximum population difference to exchange we could also have picked at random from among the pairs $\{i, j\}$ and convergence would still hold; the max-exchange protocol simply ensured the fastest rate of convergence among these choices.

Choosing a large enough number of repetitions in each stage so that the overall heat cost is close to the idealised heat cost.—Given that the max-exchange protocol in stage n converges to the thermal state that we label $\varrho_S^{(n)}$, given any error δ_E , we choose a number of repetitions m_n that is large enough so that the difference between the average energy of the actual final state of this stage, which we label $\tilde{\varrho}_S^{(n)}$, and that of the ideal state $\varrho_S^{(n)}$ is less than δ_E . In this case, the total heat cost over all stages is close to the idealised heat cost

$$\begin{aligned} |\tilde{\Delta}E_{\mathcal{H}} - \tilde{\Delta}E_{\mathcal{H}}^*| &= \left| \sum_{n=1}^N \left\{ -n\theta \operatorname{tr} \left[H_S \left(\tilde{\varrho}_S^{(n)} - \tilde{\varrho}_S^{(n-1)} \right) \right] \right\} - \sum_{n=1}^N \left\{ -n\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n)} - \varrho_S^{(n-1)} \right) \right] \right\} \right| \\ &= \left| \sum_{n=0}^{N-1} \theta \operatorname{tr} \left[H_S \left(\tilde{\varrho}_S^{(n)} - \varrho_S^{(n)} \right) \right] - N\theta \left(\tilde{\varrho}_S^{(N)} - \varrho_S^{(N)} \right) \right| \\ &\leq 2N\theta\delta_E = 2 \left(\frac{\beta^* - \beta}{\beta - \beta_{\mathcal{H}}} \right) \delta_E. \end{aligned} \quad (\text{A.220})$$

The number of repetitions in each stage m_n required depends only upon the initial choice of β^* and N .

Completing the proof.—Finally, suppose that one is given any target temperature β^* and two arbitrarily small errors, ϵ_β for the cooling and ϵ_E for the heat cost, and asked to cool incoherently in such a way that achieves

$$|\beta' - \beta^*| \leq \epsilon_\beta, \quad (\text{A.221})$$

$$|\tilde{\Delta}E_{\mathcal{H}} - \eta^{-1} \Delta F_S^{(\beta)}| \leq \epsilon_E. \quad (\text{A.222})$$

We proceed by first choosing a number of stages N so that the idealised heat cost $\tilde{\Delta}E_{\mathcal{H}}^*$ is within $\frac{\epsilon_E}{2}$ to the Carnot-Landauer bound above. The idealised sequence of temperatures satisfies $\beta_N = \beta^*$ by construction. Once N is fixed, for each stage from $n = 1$ to $N - 1$ we choose a number of repetitions for each stage m_n such that the actual heat cost is within $\frac{\epsilon_E}{2}$ of the idealised heat cost, as discussed above. This ensures that the total heat cost is within ϵ_E of the bound. Finally, we check that the number of repetitions of the last stage m_N is large enough for us to be within ϵ_β of β^* . If not, we increase the number of repetitions (this can only decrease the error in the heat cost anyway) until we are close enough, as required.

□

A.8 Comparison of Cooling Paradigms and Resources for Imperfect Cooling

Although we have looked at a number of cooling protocols throughout to demonstrate the ability for perfect cooling in the asymptotic limit, here we focus on imperfect cooling behaviour, i.e., when all resources are restricted to be finite and thus a perfectly pure state cannot be attained. We have three main goals in doing so:

1. To illustrate the finite trade-offs between the trinity of resources (energy, time, control complexity),
2. To compare the behaviour of different constructions of the cooling unitary for machines of the same size (i.e., analysing the energy-time trade-off for fixed control complexity),
3. To demonstrate the increase in resources required for cooling in the thermodynamically self-contained paradigm of energy-preserving unitaries (i.e., incoherent control), as compared to coherently-driven unitaries.

A.8.1 Rates of Resource Divergence for Linear Qubit Machine Sequence

Consider cooling a qubit target system with energy gap ω_S by swapping it sequentially with a sequence of N machine qubits of linearly increasing energy gaps. In Appendix A.7.1, we derived the deviation from the idealised heat dissipation in the incoherent control setting for a sequence of N machines [see Eq. (A.175)], which we repeat below:

$$\frac{1}{\eta} (F^* - F_0) \leq \tilde{\Delta}E_{\mathcal{H}}^* \leq \frac{1}{\eta} (F^* - F_0) + \frac{\omega_S}{N} \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right). \quad (\text{A.223})$$

We can immediately adapt this result to the paradigm of coherent control by taking $\beta_H = 0$ and replacing the heat by work, which yields

$$\Delta F_S \leq W \leq \Delta F_S + \frac{\omega_S}{N} \left(\frac{\beta^*}{\beta} - 1 \right). \quad (\text{A.224})$$

Since the above inequalities are derived from the left and right Riemann sums of an integral, as N becomes large, one can expect that W lies roughly halfway between both extremes; we can thus cast the scaling in the approximate form

$$\left[\frac{W - \Delta F}{\omega_S} \right] N \sim \frac{1}{2} \left(\frac{\beta^*}{\beta} - 1 \right). \quad (\text{A.225})$$

Thus, we see that the relevant quantifier of the energy resource here is the extra work cost above the Landauer limit relative to the system energy. Additionally, the quantifier of how much said resource is required (per machine qubit) is $\frac{\beta^*}{\beta} - 1$, which, for cold enough final temperatures, is approximately the ratio $\frac{\beta^*}{\beta}$.

Returning to the incoherent control paradigm, analysing the scaling behaviour between energy and time is more complicated. On the one hand, the expression above is only slightly modified, with the work being replaced by the heat dissipated multiplied by the Carnot factor:

$$\left[\frac{\eta \Delta E_H - \Delta F}{\omega_S} \right] N \sim \frac{1}{2} \left(\frac{\beta^*}{\beta} - 1 \right), \quad (\text{A.226})$$

which is consistent with the work-to-heat efficiency of a Carnot engine. However, in the case of incoherent control, since the population swap only takes place within a subspace of the two-qubit machine, the total population is not completely exchanged in a single operation (in contrast to that in the coherent control setting). Thus the number of operations here required to transfer a desired amount of population to the ground state of the target is greater than the number of machine qubits N . To make a fair comparison, one could either compare the same number of machine qubits but swap repeatedly (with rethermalisation of the machine in between operations)—thereby fixing the control complexity at the expense of longer time—or one could increase the number of machine qubits and count time by the number of two-level swaps—thereby fixing time to be equal at the expense of increased control complexity overall. We investigate both methods in the coming section.

A.8.2 Comparison of Coherent and Incoherent Control

Intuitively, the incoherent control paradigm requires the utilisation of a greater amount of resources (albeit less overall control in general) than the coherent control counterpart because of two distinct disadvantages. First, the temperature of the baths plays a substantial role in cooling performance. Consider the example of a swap between a system

and machine qubit: In the coherent control case, this operation transforms the target system to the state of the thermal machine qubit, characterised by the Gibbs ratio of ground-state to excited-state population. In the incoherent control case, one requires the addition of a thermal qubit from the hot bath to render said operation energy-preserving; as a result, the Gibbs ratio of the virtual qubit that the target system swaps with is, in general, worse than that of the coherent control setting, and only becomes equal in the limit of an infinite temperature hot bath. This is the first disadvantage. The second disadvantage is that in the incoherent control setting, the target system swaps with only a subspace of the machine rather than the entire one, i.e., it is swapped with a virtual qubit. Thus, the exchange of population is only partial as compared to the coherent control case: In the limiting case of an infinite temperature hot bath, said factor goes to $\frac{1}{2}$ for all relevant two-level subspaces. This implies that a greater number of operations, and thus time, is required in the incoherent control paradigm in order to achieve a similar result as its coherent control counterpart.

We illustrate this behaviour via the following example. The system is a degenerate qubit (beginning in the maximally mixed state), and we fix the final target ground-state population ($p = 0.99$, corresponding to $\epsilon = 1 - p = 0.01$). Even in this simple case, the optimal finite-resource protocols with coherent and incoherent control are not known; we therefore compare protocols from each setting that make use of machines of a similar structure, namely swapping with machine qubits (virtual ones, in the incoherent control setting) of linearly increasing energy gaps.

More specifically, the coherent control cooling protocol employed is that of a sequence of swaps with machine qubits of linearly increasing energy gaps, and for the fixed target population, we can calculate the surplus work cost over the Landauer limit as a function of the number N of operations (which corresponds in this case to the number of machine qubits). In the incoherent control case, we take the hot bath to be at infinite temperature, allowing for the potential saturation of the Landauer limit as in the coherent case. In this way we isolate the disadvantage that arises due to working in degenerate subspaces in our analysis. Here too we take a linear sequence of energy gaps for the cold (and hot) baths, with a single operation step corresponding to a three-level energy-conserving exchange involving the qubit taken from each of the hot and cold parts of the machine, i.e., $|1\rangle_S|0\rangle_C|0\rangle_{\mathcal{H}} \leftrightarrow |0\rangle_S|1\rangle_C|1\rangle_{\mathcal{H}}$. As mentioned previously, for an incoherent control protocol of fixed overall machine size, there are essentially two extremal methods of implementation. The first is to identify N two-level subspaces of the total machine with distinct energy gaps and perform the sequence of virtual swaps between them and the target; in the language of Appendix A.7, we therefore have N different stages with a single step within each stage (no repetitions) before moving on to the next stage. The second is to take $\frac{N}{m}$ two-level subspaces and swap the target with each virtual qubit m times before moving on to the next; i.e., we here have $\frac{N}{m}$ different stages with m steps (repetitions) within each




Figure A.1: Imperfect Cooling with Coherent and Incoherent Control. We compare the performance of coherent and incoherent control protocols for cooling a degenerate qubit target by swapping it with machine qubits with linearly increasing energy. The final ground-state population is fixed to be 0.99. The inverse of the surplus work cost $W - \beta\tilde{\Delta}S_{\mathcal{S}}$ (with $\beta = 1$) is plotted against the total number of unitary operations, with the temperature of the hot bath in the incoherent control protocols set to $\beta_H = \infty$ in order to make meaningful comparison to the coherent control case. We see that the coherent control protocol (blue) outperforms the two incoherent ones (purple, red) at any given time. As discussed in the text, there are two choices for how to implement an incoherent control protocol of this type with fixed control complexity: The red line corresponds to a protocol in which a machine (subspace) with the same energy gap is reused 5 times before moving on to the next; on the other hand, the purple line depicts the case where there are no repetitions within each stage defined by a distinct energy gap in the machine. By inspection, the single-use incoherent protocol (purple) requires ~ 3 times more unitaries to achieve the same efficiency as the coherent one (blue), whereas the 5-repetition incoherent protocol (red) requires ~ 5.3 times as many unitaries as the coherent one.

stage. For the same fixed ground-state population, we plot the surplus work cost (energy drawn from the hot bath in the case of incoherent control) against the total machine size / number of two-level unitary swaps, as characterised by N , for both of these incoherent control adaptations, comparing them to the coherent control paradigm in Fig. A.1.

In both control paradigms, we see that the deviation of the energy cost above the Landauer limit scales inversely with the number of operations [as expected from Eqs. (A.225) and (A.226)], but the proportionality constant is worse in the case of incoherent control. Moreover, the incoherent control paradigm with no repetitions within stages outperforms that with multiple repetitions, as intuitively expected since the former protocol corresponds to one for which the spacing between distinct energy gaps that are utilised is smaller, allowing us to stay closer to the reversible limit in each step. In our example, the no repetition incoherent control protocol is around 3 times worse than the coherent control protocol and the incoherent control protocol with $m = 5$ repetitions is around 5.3 times worse, implying that one would require that many times the number of operations (i.e., that much more time) to achieve the same performance with incoherent control paradigm as with coherent control.

Supplemental Information for Chapter 2

B.1 Markovian Embedding of Collision Models with Memory

We are interested in exploring analytically the effects of memory regarding the task of cooling a quantum system. We do not wish to allow for arbitrary non-Markovianity, as this would lead to an infinite resource in a sense that it allows us to cool the system to the ground state perfectly. Rather aim to obtain a cooling bound in the limit of infinite cycles for a particular class of non-Markovian dynamics, namely a generalised collision model endowed with memory. Such collision models with memory are quite general, simply assuming that between each step of the dynamics, the system interacts with k constituent sub-machines (which altogether make up the full set of machines), of which some $\ell < k$ of these carry the memory forward. The Markovian setting is recovered for $\ell = 0$. A schematic is provided in Fig. B.1.

We consider a target system of dimension d_S and local Hamiltonian $H_S = \sum_{i=0}^{d_S-1} E_i |i\rangle\langle i|_S$, where $\{E_i\}$ are sorted in non-decreasing order and an environment comprising of a number of constituent identical machines of finite size d_M , each of which has the local Hamiltonian $H_M = \sum_{i=0}^{d_M-1} \mathcal{E}_i |i\rangle\langle i|_M$, where $\{\mathcal{E}_i\}$ are also sorted in non-decreasing order. The system and all machines (i.e., the entire environment) begin uncorrelated in a thermal state with the same inverse temperature β :

$$\varrho_S^{(0)}(\beta) = \tau_S(\beta) \quad \text{and} \quad \varrho_E^{(0)}(\beta) = \bigotimes_j \tau_{M_j}(\beta), \quad (\text{B.1})$$

where $\tau_X(\beta) := \mathcal{Z}_X^{-1} e^{-\beta H_X}$ with the partition function $\mathcal{Z}_X := \text{tr}[e^{-\beta H_X}]$.

Fixing k and ℓ provides a particular dynamical structure of the non-Markovian process: It stipulates that at each timestep there are k machines interacting with the system, of which ℓ are kept to perpetuate the memory. For example, after n steps, the system state




Figure B.1: Generalised Collision Model with Memory. In the collision model picture, the environment that the system (green) interacts with is assumed to be comprised of individual sub-units, which we call machines. Between each step of dynamics, the system interacts unitarily (blue outline) with a subset of these; the cardinality of this set is labelled by k throughout this work. A further subset of these previously-interacted-with machines of cardinality $\ell < k$ take part in the following interaction, becoming memory carriers (red). At each timestep, $k - \ell$ fresh machines are incorporated into the dynamics (yellow). Here we have shown $k = 2, \ell = 1$.

is

$$\varrho_S^{(n)}(\beta, k, \ell) = \text{tr}_M[U_{k,\ell}^{(n)} \dots U_{k,\ell}^{(1)} (\tau_S(\beta) \bigotimes_{j=1}^m \tau_{M_j}(\beta)) U_{k,\ell}^{(1)\dagger} \dots U_{k,\ell}^{(n)\dagger}], \quad (\text{B.2})$$

where $U_{k,\ell}^{(n)}$ is an arbitrary unitary transformation between the target system and the k machines labelled by $\{(n-1)(k-\ell) + 1, \dots, n(k-\ell) + \ell\}$ (an identity map is implied on the other machines) and $m := k + (n-1)(k-\ell)$ is the total number of machines used by the protocol up until timestep n , which will be important in making finite time comparisons, as we do in Appendix B.4 (see Fig. 2.1 for a graphical depiction in terms of a circuit diagram).

Importantly, the state of the system at any time is a function of the full microscopic energy structure $\{E_i\}$ and $\{\mathcal{E}_i\}$ (which we do not explicitly label for ease of notation), β, k and ℓ ; the latter two numbers specify a particular dynamical structure in terms of which systems the unitaries act upon between timesteps. If $\ell = 0$, the dynamics of the system is Markovian, since at each step, the system interacts with fresh machines that contain no memory of the past dynamics of the system. Otherwise, each of the machines interacts more than once with the target and only $k - \ell$ fresh machines are added into the interaction at each step.

Eq. (B.2) highlights the restriction imposed by the assumption of generalised collision model dynamics from the fully general case of non-Markovian dynamics where the full system-environment must be tracked; in particular, a subset of the environment (the $k - \ell$ rethermalising systems) is traced out between steps, rendering the dynamics tractable for small k, ℓ . However, it is important to note that on the level of the system, memory effects still play a role. We first show that for $\ell > 0$ the dynamics considered is indeed non-Markovian in general.

To analyse the proposed setting, we need to look at the evolution of the entire joint system and machines to consider the effect of the memory in the protocol. For instance, consider rewriting Eq. (B.2) as a dynamical map taking the initial system state $\varrho_S^{(0)}(\beta) = \tau_S(\beta)$ to the later one under a generic dynamical structure determined by the

choice of k and ℓ , i.e., define

$$\Lambda_{k,\ell}^{(n:0)}(\beta)[X_S] := \text{tr}_M[U_{k,\ell}^{(n)} \dots U_{k,\ell}^{(1)} (X_S \otimes_j \tau_{M_j}(\beta)) U_{k,\ell}^{(1)\dagger} \dots U_{k,\ell}^{(n)\dagger}], \quad (\text{B.3})$$

where we have now included all of the machines in the environment and an identity map on those not taking part in the interactions until timestep n is implied, such that

$$\varrho_S^{(n)}(\beta, k, \ell) = \Lambda_{k,\ell}^{(n:0)}(\beta)[\varrho_S^{(0)}(\beta)]. \quad (\text{B.4})$$

Linearity, complete positivity and trace preservation of the map $\Lambda_{k,\ell}^{(n:0)}(\beta)$ is guaranteed for any β, k, ℓ and most importantly n by the fact that S and E begin initially uncorrelated and the dynamics evolves unitarily on the global level, before a final partial trace is taken over the environment degrees of freedom. Complete positivity is particularly important to ensure that the map takes valid quantum states to valid quantum states. In general, the global state $\varrho_{SM}^{(n)}(\beta, k, \ell)$, where M labels the subset of the environment that has taken part non-trivially in the dynamics up until timestep n , involves correlations between S and M ; taking the partial trace over M destroys all such correlations. Thus, one cannot, in general, describe the evolution of the system between multiple times as a divisible concatenation of Completely Positive and Trace Preserving (CPTP) maps, i.e.,

$$\varrho_S^{(n)}(\beta, k, \ell) = \Lambda_{k,\ell}^{(n:0)}(\beta)[\varrho_S^{(0)}(\beta)] \neq \Lambda_{k,\ell}^{(n:t)}(\beta) \circ \Lambda_{k,\ell}^{(t:0)}(\beta)[\varrho_S^{(0)}(\beta)]. \quad (\text{B.5})$$

Here, we have defined $\Lambda_{k,\ell}^{(n:t)}(\beta)$ as the map that would be tomographically constructed if one were to discard the system at time t (which is generally correlated to M) and perform a quantum channel tomography by preparing a fresh basis of input states (see Fig. B.2); since the reprepared state is uncorrelated to M by construction, the map $\Lambda_{k,\ell}^{(n:t)}(\beta)$ is guaranteed to be CPTP for any choice of parameters [264]. Testing for equality in Eq. (B.5) then corresponds to the operational notion of Completely Positive (CP)-divisibility proposed in Ref. [151]; importantly, its breakdown acts as a valid witness for non-Markovianity that is stricter than other notions of CP-divisibility proposed throughout the literature (in particular, it is stronger than that based on invertible CP-divisibility in any case where the dynamical maps are invertible). Of course, the fact that Eq. (B.5) is generally an inequality for generic dynamics does not imply that it is so for the particular optimal cooling dynamics described throughout this article; however, it is simple to show that the optimal cooling protocol indeed generates correlations between the system and machines that lead to a breakdown of (operational) CP-divisibility, and hence the particular dynamics considered throughout is inherently non-Markovian.

Nonetheless, the collision model memory structure that we have introduced crucially allows for a Markovian embedding that permits a significant simplification in the analysis [95]. As mentioned previously, in general, one would need to track the total joint evolution throughout the entire protocol, which quickly becomes computationally exhaustive as k grows. However, for a choice of k and ℓ , we can group the system S and ℓ of




Figure B.2: Operational CP-Divisibility. A breakdown of operational CP-divisibility is a witness for non-Markovianity. The test consists of tomographically constructing a set of maps describing the dynamics and checking the validity of Eq. (B.5). We illustrate the scenario for a subset of times $t = 2$, $n = 4$, with $k = 2$, $\ell = 1$: The left panel depicts the map $\Lambda_{2,1}^{(4:0)}$, which comprises everything within the purple border, i.e., the initial states of all machines, all joint unitary interactions, and the final partial trace. Note that the final memory carrier M_5 should continue forward, but the map $\Lambda_{2,1}^{(4:0)}$ does not capture this and traces out that subsystem. The middle panel depicts the map $\Lambda_{2,1}^{(2:0)}$. Both maps can be tomographically reconstructed by preparing a basis of input states at the initial time and measuring the outputs at time $t = 4$ and $t = 2$, respectively. As the system begins initially uncorrelated with the machines, the unitary dilation guarantees that the maps constructed are CPTP. The final map needed, $\Lambda_{2,1}^{(4:2)}$, is shown in the right panel. In general, at time $t = 2$, the system is correlated to the machines, thus breaking the assumption of no initial correlations. An operational circumvention is to discard the system state at $t = 2$ and reprepare a fresh one, thereby erasing all system-machine correlations. This has the effect of rendering the memory carriers into a fixed quantum state, which can be included in the description of $\Lambda_{2,1}^{(4:2)}$ to ensure that it is CP. When memory is present, tracing out the system at the intermediary timestep generally conditions the state of the memory carriers into a state that generally differs from the initial thermal state $\tau(\beta)$, labelled here $\sigma(\beta)$, with the altered part of the evolution depicted in green; thus, the full dynamics generically differs from that described by the concatenation.

the machines into a larger joint system, which we label SL , which interacts with $k - \ell$ fresh machines at each timestep; we label these fresh machines with R as they model rethermalization of some of the machines with the environment. On the level of SL , the dynamics is Markovian, as the degrees of freedom carrying the memory have been included in the description of the target system. One can obtain the state of the overall SL target by tracing out the R machines at each step. We therefore have

$$\varrho_{SL}^{(n)}(\beta, k, \ell) = \text{tr}_R[\tilde{U}_{k,\ell}^{(n)} \dots \tilde{U}_{k,\ell}^{(1)} (\varrho_{SL}^{(0)}(\beta, \ell) \otimes \varrho_R^{(0)}(\beta, k, \ell)) \tilde{U}_{k,\ell}^{(1)\dagger} \dots \tilde{U}_{k,\ell}^{(n)\dagger}], \quad (\text{B.6})$$

where $\varrho_{SL}^{(0)}(\beta, \ell) := \tau_S(\beta) \otimes_{j=1}^{\ell} \tau_{M_j}(\beta)$, $\varrho_R^{(0)}(\beta, k, \ell) := \otimes_{j=\ell+1}^m \tau_{M_j}$ and $\tilde{U}_{k,\ell}^{(n)}$ is an arbitrary unitary interaction between SL and $k - \ell$ fresh machines occurring immediately prior to timestep n (see Fig. B.3). Due to the fact that no memory transportation occurs on the SL level throughout the protocol, the full dynamics of the system and memory carriers is captured by the following concatenation of dynamical maps:

$$\varrho_{SL}^{(n:0)}(\beta, k, \ell) = \tilde{\Lambda}_{k,\ell}^{(n:t)} \circ \tilde{\Lambda}_{k,\ell}^{(t:0)}[\varrho_{SL}^{(0)}(\beta, \ell)], \quad (\text{B.7})$$

where $\tilde{\Lambda}_{k,\ell}^{(n:t)}$ is a CPTP map that acts only upon SL and depends on the unitary operators $\tilde{U}_{k,\ell}^{(n)}, \dots, \tilde{U}_{k,\ell}^{(t)}$ and the initial state of the $k - \ell$ fresh machines taking part in each interaction. Thus the dynamics is (operationally) CP-divisible on the level of SL , and it is easy to see that it is even Markovian in the stronger sense provided in Refs. [128, 180].

This ‘‘Markovian embedding’’ of the non-Markovian dynamics provides an opportunity to investigate the problem at hand with a simplified Markovian dynamics on the larger SL




Figure B.3: Markovian Embedding of Generalised Collision Model. The generalised collision models considered in this work can be embedded as a Markovian process by grouping together the system and memory carriers into a larger target system. On the left, we show that by considering the original circuit shown in Fig. 2.1, plus allowing for a swap interaction between relevant machines, the dynamics can be transformed into the circuit in the right panel. Here, we identify the memory carrier systems as L ; the entire SL system now interacts with fresh machines between each timestep, which we label R . The full dynamics of SL can be described by a Markovian sequence of CPTP maps: Here we show only the first two, $\tilde{\Lambda}_{2,1}^{(1:0)}, \tilde{\Lambda}_{2,1}^{(2:1)}$ within the purple borders, although the dynamics between any steps can be described similarly. Note that, in contrast to the dynamics of the system itself, the dynamical maps on the level of SL contain the complete description of the process and Eq. (B.7) always holds; intuitively, this is because none of the systems carrying memory are artificially “cut” by the description of the dynamical map (see the partial traces on the red lines in Fig. B.2 for comparison).

system instead of complicated non-Markovian dynamics that occurs on the level of S . In the sense of Ref. [95], the number of memory carriers ℓ corresponds to the memory depth of the dynamics; intuitively, this is the number of additional subsystems that need to be included in the description of the system so that the dynamics is rendered Markovian.

B.2 Proof of Theorem 2.1

Here we prove Theorem 2.1. The proof makes use of the main result of Refs. [82, 83], which derive the ultimate cooling bounds for a Markovian protocol. We first embed the non-Markovian dynamics of S as a Markovian one by considering the target system SL , before finding the optimally cool SL state in the asymptotic limit, which we denote ϱ_{SL}^* . We then combine this result with the fact that there always exists a unitary that can finally be implemented on just SL such that the reduced state of S majorises all of the possible reduced states of the system SL , as long as ϱ_{SL} is majorised by ϱ_{SL}^* . This implies that the optimal asymptotic system state ϱ_S^* can be calculated from the reduced state of any ϱ_{SL} which has the same eigenvalue spectrum of the asymptotically optimal SL state.

Before we begin with the proof, we provide a definition of majorisation for completeness:

Definition B.1. Given a vector of real numbers $\mathbf{a} \in \mathbb{R}^d$, we denote by \mathbf{a}^\downarrow the vector with the same components but sorted in non-increasing order. Given $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$, we say that $\mathbf{a} \succ \mathbf{b}$ (\mathbf{a} majorises \mathbf{b}) iff

$$\sum_i^k a_i^\downarrow \geq \sum_i^k b_i^\downarrow \quad \forall k \in \{1, \dots, d\}. \quad (\text{B.8})$$

Proof of Theorem 2.1. We first perform a Markovian embedding of the non-Markovian collision model dynamics by considering the evolution of the larger target system SL (with dimension $d_S d_M^\ell$); for a given number ℓ of memory carriers, the Markovian embedding corresponds to a memory depth of ℓ in the sense of Ref. [95], which is to say that by including the description of the ℓ memory carrying systems with that of the original target system S , the dynamics is rendered Markovian. This is because, at each step of the protocol, SL interacts with $k - \ell$ fresh machine systems (with total dimension $d_M^{k-\ell}$), which are subsequently discarded and play no further role in the dynamics.

In the Markovian regime, we can use the theorem of universal cooling bound presented in Ref. [82], which holds for an arbitrary target system interacting with an arbitrary machine, which are initially in a thermal state with inverse temperature β , in the limit of infinite cycles.

Lemma B.1 (Markovian Asymptotic Cooling Limit [Theorem 1 in Ref. [82]]). *For any $d_{\tilde{S}}$ -dimensional system with Hamiltonian $H_{\tilde{S}} = \sum_{i=0}^{d_{\tilde{S}}-1} E_i |i\rangle\langle i|_{\tilde{S}}$ interacting with a $d_{\tilde{M}}$ dimensional machine with Hamiltonian $H_{\tilde{M}} = \sum_{i=0}^{d_{\tilde{M}}-1} \mathcal{E}_i |i\rangle\langle i|_{\tilde{M}}$ with $\{E_i\}, \{\mathcal{E}_i\}$ sorted in non-decreasing order, in the limit of infinite cycles,*

- *The ground state population of the target system \tilde{S} is upper bounded by*

$$\tilde{p}^*(\beta) = \left(\sum_{n=0}^{d_{\tilde{S}}-1} e^{-\beta n \tilde{\mathcal{E}}_{\max}} \right)^{-1} \quad (\text{B.9})$$

where $\tilde{\mathcal{E}}_{\max}$ is the largest energy gap of the machine.

- *In both coherent and incoherent control scenarios, the vectorized form of eigenvalues of the final state is majorised by that of the following state,*

$$\varrho_{\tilde{S}}^*(\beta) = \sum_{n=0}^{d_{\tilde{S}}-1} \frac{e^{-\beta n \tilde{\mathcal{E}}_{\max}}}{\mathbb{Z}_{\tilde{S}}(\beta, \tilde{\mathcal{E}}_{\max})} |n\rangle\langle n|_{\tilde{S}}, \quad (\text{B.10})$$

if the initial state $\varrho_{\tilde{S}}^{(0)}(\beta)$ is majorised by $\varrho_{\tilde{S}}^*(\beta)$.

- *In the coherent control paradigm, the asymptotically optimal state, which is also achievable, is given by $\varrho_{\tilde{S}}^*(\beta)$.*

In view of the fact that the final state $\varrho_{\tilde{S}}^*(\beta)$ has a unique eigenvalue distribution and is achievable in the infinite-cycle limit, it is possible to investigate this bound on the population of a particular subspace of dimension d , rather than just the ground state population. It is straightforward to show that its population is upper bounded by the d largest eigenvalues

$$\tilde{p}_{\mathcal{H}_d}^*(\beta) \leq \frac{\sum_{n=0}^{d-1} e^{-\beta n \tilde{\mathcal{E}}_{\max}}}{\mathbb{Z}_{\tilde{S}}(\beta, \tilde{\mathcal{E}}_{\max})}. \quad (\text{B.11})$$

With this knowledge, we are in the position to study the optimal cooling of the non-Markovian collision model protocol in the limit of infinite cycles by employing Lemma B.1. In our case, an arbitrary target system S interacts with k machines at each step and ℓ of these carries memory forward to be involved in the next interaction. Thus, the joint system SL corresponds to the target system \tilde{S} here, which undergoes Markovian dynamics with respect to the $k - \ell$ fresh machines added at each step, which comprise \tilde{M} ; hence, $\tilde{\mathcal{E}}_{\max}$ is equal to $(k - \ell)\mathcal{E}_{\max}$. It turns out that the maximum energy gap of the fresh machines and the total dimension and number of the memory carriers play an important role in the ultimate cooling bound.

Using our Markovian embedding of the dynamics and Lemma B.1, we see that in the limit of infinite cycles for any control paradigm, the vector of the eigenvalues of the asymptotic state is majorised by

$$\varrho_{SL}^*(\mathcal{E}_{\max}, \beta, k, \ell) = \sum_{n=0}^{d_S d_M^\ell - 1} \frac{e^{-\beta n(k-\ell)\mathcal{E}_{\max}}}{\mathbb{Z}_{SL}(\beta, (k-\ell)\mathcal{E}_{\max})} |n\rangle\langle n|_{SL} \quad (\text{B.12})$$

if $\varrho_{SL}^{(0)}(\mathcal{E}_{\max}, \beta, k, \ell) \prec \varrho_{SL}^*(\mathcal{E}_{\max}, \beta, k, \ell)$ and $\{|n\rangle_{SL}\}$ is the energy eigenbasis with respect to which the energy eigenvalues are sorted in non-decreasing order. So far, we have found the achievable passive state that majorises all other reachable states of SL via unitary operations on SLR . However, this state is not unique as the characterisation is based solely on its eigenstate distribution: One can indeed find a whole set of equally cool reachable states, i.e., those for which $\vec{\lambda}[\varrho_{SL}^*(\mathcal{E}_{\max}, \beta, k, \ell)] = \vec{\lambda}[U_{SL}\varrho_{SL}^*(\mathcal{E}_{\max}, \beta, k, \ell)U_{SL}^\dagger] \forall U_{SL}$, where $\vec{\lambda}[\varrho]$ indicates the vectorised form of the eigenvalues of ϱ . We now present another lemma which says that from any such state of SL , one can reach the optimally cool state of S , ϱ_S^* , helping us complete the proof.

Lemma B.2 (Reduced State Majorisation). *For any pair states ϱ_{AB} and σ_{AB} , if $\sigma_{AB} \prec \varrho_{AB}$, there exists a unitary U_{AB}^{opt} on \mathcal{H}_{AB} such that:*

$$\text{tr}_B[U_{AB}\sigma_{AB}U_{AB}^\dagger] \prec \text{tr}_B[U_{AB}^{\text{opt}}\varrho_{AB}U_{AB}^{\text{opt}\dagger}] \quad \forall U_{AB}. \quad (\text{B.13})$$

Proof. Without loss of generality, we assume that the eigenvalues of both states ϱ_{AB} and σ_{AB} are sorted in non-increasing order as follows

$$\mathbf{p}_{AB} = \{p_\alpha^\downarrow\}_{\alpha=0}^{d_A d_B - 1} \quad \text{and} \quad \mathbf{q}_{AB} = \{q_\alpha^\downarrow\}_{\alpha=0}^{d_A d_B - 1}. \quad (\text{B.14})$$

Based on the sorted eigenvalues, $\sigma_{AB} \prec \varrho_{AB}$ if and only if

$$\sum_{\alpha=0}^k q_\alpha^\downarrow \leq \sum_{\alpha=0}^k p_\alpha^\downarrow \quad \forall k \in \{0, 1, \dots, d_A d_B - 1\}. \quad (\text{B.15})$$

Now we aim to find the reduced state σ_A^{opt} majorising all of the achievable reduced states possible to generate by a unitary transformation of σ_{AB} , which we assume to be diagonal

in the orthonormal basis $\{|ij\rangle_{AB}\}$ without loss of generality:

$$\sigma_{AB} = \sum_{i=0}^{d_A-1} \sum_{j=0}^{d_B-1} q_{ij} |ij\rangle\langle ij|_{AB} \quad (\text{B.16})$$

One can show that it is possible to obtain σ_A^{opt} from a bipartite state that is diagonal in the same basis. Then we have,

$$\tilde{\sigma}_{AB} = U_{AB}^{\text{opt}} \sigma_{AB} U_{AB}^{\text{opt} \dagger} = \sum_{i=0}^{d_A-1} \sum_{j=0}^{d_B-1} \tilde{q}_{ij} |ij\rangle\langle ij|_{AB}, \quad (\text{B.17})$$

where U_{AB}^{opt} is simply a permutation matrix that reorders the eigenvalues appropriately. The final reduced state is then given by

$$\tilde{\sigma}_A = \sum_{i=0}^{d_A-1} \left(\sum_{j=0}^{d_B-1} \tilde{q}_{ij} \right) |i\rangle\langle i|_A. \quad (\text{B.18})$$

We now need to show that the appropriate unitary maximises the eigenvalues of the reduced state with respect to eigenvalues of σ_{AB} . If we rearrange the eigenvalues in such a way that $\tilde{q}_{ij} = q_{\alpha}^{\downarrow}$ where α is given by $\alpha := i d_B + j$, we obtain σ_A^{opt} as the following

$$\sigma_A^{\text{opt}} = \sum_{i=0}^{d_A-1} \eta_i^{\downarrow} |i\rangle\langle i|_A = \sum_{i=0}^{d_A-1} \left(\sum_{j=0}^{d_B-1} q_{\alpha=i d_B + j}^{\downarrow} \right) |i\rangle\langle i|_A, \quad (\text{B.19})$$

where, due to the sorting of $\{q_{\alpha}^{\downarrow}\}$, the eigenvalues of σ_A^{opt} are sorted in non-decreasing order. The final reduced state satisfies the following condition

$$\text{tr}_B \left[U_{AB} \sigma_{AB} U_{AB}^{\dagger} \right] \prec \text{tr}_B \left[U_{AB}^{\text{opt}} \sigma_{AB} U_{AB}^{\text{opt} \dagger} \right] = \sigma_A^{\text{opt}} \quad \forall U_{AB}. \quad (\text{B.20})$$

Similarly one can find ϱ_A^{opt} by applying a unitary V_{AB}^{opt} ,

$$\varrho_A^{\text{opt}} = \sum_{i=0}^{d_A-1} \xi_i^{\downarrow} |i\rangle\langle i|_A = \sum_{i=0}^{d_A-1} \left(\sum_{j=0}^{d_B-1} p_{\alpha=i d_B + j}^{\downarrow} \right) |i\rangle\langle i|_A, \quad (\text{B.21})$$

whose eigenvalues are also in non-increasing order by construction. The final step of the proof is to show that $\sigma_A^{\text{opt}} \prec \varrho_A^{\text{opt}}$ whenever $\sigma_{AB} \prec \varrho_{AB}$. This majorisation condition can be recast in the form of

$$\sum_{i=0}^k \eta_i^{\downarrow} \leq \sum_{i=0}^k \xi_i^{\downarrow} \Rightarrow \sum_{\alpha=0}^{(k+1)d_B-1} q_{\alpha}^{\downarrow} \leq \sum_{\alpha=0}^{(k+1)d_B-1} p_{\alpha}^{\downarrow} \quad \forall k \in \{0, 1, \dots, d_A - 1\}. \quad (\text{B.22})$$

Using inequality (B.15), one can easily show that inequality (B.22) always holds, i.e., $\sigma_A^{\text{opt}} \prec \varrho_A^{\text{opt}}$, completing the proof. \square

In the next step, we aim to maximise the population of the system ground state, i.e., the maximum population of the specific subspace of the SL target given in Eq. (B.11). One must therefore find the target state $\varrho_S^*(\beta, k, \ell)$ that can be achieved from the states

ϱ_{SL} with the same eigenvalues as $\varrho_{SL}^*(\mathcal{E}_{\max}, \beta, k, \ell)$, since, from Lemma B.2 we know that this state majorises the largest set of states in S . In order to do so, we maximise the eigenvalues of $\varrho_S^*(\mathcal{E}_{\max}, \beta, k, \ell)$ with respect to those of $\varrho_{SL}^*(\mathcal{E}_{\max}, \beta, k, \ell)$. One can appropriately sort the eigenvalues of the system S and the memory carrier machines L with the following unitary:

$$\begin{aligned}
 & U_{SL}^{\text{opt}} \varrho_{SL}^*(\mathcal{E}_{\max}, \beta, k, \ell) U_{SL}^{\text{opt}\dagger} \\
 &= \sum_{n=0}^{d_S-1} \sum_{j=0}^{d_M^\ell-1} \frac{e^{-\beta(nd_M^\ell+j)(k-\ell)\mathcal{E}_{\max}}}{\mathbb{Z}_{SL}(\beta, (k-\ell)\mathcal{E}_{\max})} |nj\rangle\langle nj|_{SL} \\
 &= \sum_{n=0}^{d_S-1} \sum_{j=0}^{d_M^\ell-1} \frac{e^{-\beta(nd_M^\ell+j)(k-\ell)\mathcal{E}_{\max}}}{\sum_{n'=0}^{d_S-1} \sum_{j'=0}^{d_M^\ell-1} e^{-\beta(n'd_M^\ell+j')(k-\ell)\mathcal{E}_{\max}}} |nj\rangle\langle nj|_{SL} \\
 &= \sum_{n=0}^{d_S-1} \frac{e^{-\beta nd_M^\ell(k-\ell)\mathcal{E}_{\max}}}{\sum_{n'=0}^{d_S-1} e^{-\beta n'd_M^\ell(k-\ell)\mathcal{E}_{\max}}} |n\rangle\langle n|_S \otimes \sum_{j=0}^{d_M^\ell-1} \frac{e^{-\beta j(k-\ell)\mathcal{E}_{\max}}}{\sum_{j'=0}^{d_M^\ell-1} e^{-\beta j'(k-\ell)\mathcal{E}_{\max}}} |j\rangle\langle j|_L, \quad (\text{B.23})
 \end{aligned}$$

Thus, beginning with the optimally cool SL state in Eq. (B.12), we can reorder the eigenvalues via U_{SL}^{opt} in Eq. (B.23) such that the subsystem S is optimally cool; finally applying Lemma B.2 then implies that

$$\text{tr}_L[U_{SL} \varrho_{SL}^*(\mathcal{E}_{\max}, \beta, k, \ell) U_{SL}^\dagger] \prec \varrho_S^*(\mathcal{E}_{\max}, \beta, k, \ell) \quad \forall U_{SL}, \quad (\text{B.24})$$

where $\varrho_S^*(\mathcal{E}_{\max}, \beta, k, \ell)$ is indeed given by taking the partial trace over L of the optimal joint state $\varrho_{SL}^*(\mathcal{E}_{\max}, \beta, k, \ell)$:

$$\begin{aligned}
 \varrho_S^*(\mathcal{E}_{\max}, \beta, k, \ell) &= \sum_{n=0}^{d_S-1} \frac{\sum_{j=0}^{d_M^\ell-1} e^{-\beta(nd_M^\ell+j)(k-\ell)\mathcal{E}_{\max}}}{\sum_{n'=0}^{d_S-1} \sum_{j'=0}^{d_M^\ell-1} e^{-\beta(n'd_M^\ell+j')(k-\ell)\mathcal{E}_{\max}}} |n\rangle\langle n|_S \\
 &= \sum_{n=0}^{d_S-1} \frac{e^{-\beta nd_M^\ell(k-\ell)\mathcal{E}_{\max}}}{\sum_{n'=0}^{d_S-1} e^{-\beta n'd_M^\ell(k-\ell)\mathcal{E}_{\max}}} |n\rangle\langle n|_S, \quad (\text{B.25})
 \end{aligned}$$

thus establishing $\varrho_S^*(\mathcal{E}_{\max}, \beta, k, \ell)$ as the optimal system state in the asymptotic limit.

In conclusion, from Lemma B.1, we know that the final state of the system SL , for any control paradigm in the infinite-cycle limit, is majorised by $\varrho_{SL}^*(\mathcal{E}_{\max}, \beta, k, \ell)$. Consequently, via Lemma B.2, the final state of S is also majorised by $\varrho_S^*(\mathcal{E}_{\max}, \beta, k, \ell)$. Then, the population of the ground state of the system is upper bounded by the sum of the d_M^ℓ largest eigenvalues of $\varrho_{SL}^*(\mathcal{E}_{\max}, \beta, k, \ell)$, i.e., $p^*(\mathcal{E}_{\max}, \beta, k, \ell) = (\sum_{n=0}^{d_S-1} e^{-\beta nd_M^\ell(k-\ell)\mathcal{E}_{\max}})^{-1}$.

We finally prove that in the coherent scenario, the state $\varrho_S^*(\mathcal{E}_{\max}, \beta, k, \ell)$ is achievable in the limit of infinite cycles. Using Lemma B.1, one can easily show that the final state of SL under optimal coherent operations converges to $\varrho_{SL}^*(\mathcal{E}_{\max}, \beta, k, \ell)$. To do so, we use the fact that in the coherent scenario, one can apply any unitary operation on the system SL at the final step. We then achieve the desired target state $\varrho_S^*(\mathcal{E}_{\max}, \beta, k, \ell)$ via employing the unitary U_{SL}^{opt} , completing the proof. \square

B.3 Role of System-Memory Carrier Correlations

We here remark on the correlations that can develop between the target system S and memory carriers L throughout the cooling protocols that have been discussed in the main text.

In particular, we have focused on two procedures. The first strategy optimally cools the joint SL system at each timestep, which does not necessarily ensure that S is locally optimally cool; it is only at the final step that the target system is further cooled by transferring entropy away from it and toward the memory carriers. More precisely, this protocol implements the unitary whose action is defined in Eq. (2.7) at each step (which globally cools SL), and only finally implements the unitary that ensures S to be locally cool, whose action is defined in Eq. (2.8). This protocol thus focuses in each step on cooling SL *globally*: In effect, it cools SL with respect to its own (global) energy eigenbasis ($|0\rangle_{SL}, |1\rangle_{SL}, \dots, |n\rangle_{SL}$), with $|i\rangle_{SL}$ denoting the i^{th} excited state of SL and $n = d_S d_M - 1$; as such, we refer to it here as the “global basis cooling protocol”.

While the above protocol eventually, i.e., at the last step, optimally cools S , it does not necessarily do so at each step. To this end, in the main text we presented a second cooling protocol which is step-wise optimal. Intuitively, this protocol globally cools SL optimally at each step (as does the strategy described above), and, given that optimally cool state, additionally performs a unitary on SL to furthermore optimally cool S *locally* at each step; as such, we refer to this scheme as the “local basis cooling protocol”. In practical terms, one can view this protocol as cooling SL optimally at each step in the locally-ordered energy eigenbasis ($|00\rangle_{SL}, |01\rangle_{SL}, \dots, |0, d_L - 1\rangle_{SL}, |10\rangle_{SL}, \dots, |d_S - 1, d_L - 1\rangle_{SL}$), where $|i\rangle_S$ is the i^{th} energy excited state of S and $|j\rangle_L$ the j^{th} energy excited state of L .

Note that, while related by a permutation of the basis elements, this local energy eigenbasis generally differs from the global energy eigenbasis of SL , which does not take local information regarding the energy structure into account. For instance, if the target and memory comprise a qubit each with (respectively) distinct energy gaps, $\{E_S, E_L\}$ such that $E_L > E_S$ without loss of generality, then cooling with respect to the global basis would order the eigenvalues $\lambda_0 \geq \dots \geq \lambda_3$ into the respective subspaces $\{|00\rangle, |10\rangle, |01\rangle, |11\rangle\}$. However, for S to be optimally cool, the eigenvalues need to be sorted in non-increasing order with respect to $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$; this type of ordering is only achieved at the last timestep via the final unitary in the global cooling protocol, but at every timestep in the local cooling protocol. Discrepancies between the locally-optimal and globally-optimal basis ordering typically become more pronounced as systems become more complex, i.e., multi-partite and high-dimensional, highlighting the necessity for careful accounting. In general, the logic above implies that the local cooling protocol will not reach the coldest SL possible state in particular at each step, but nonetheless always reaches one that is unitarily equivalent to it.

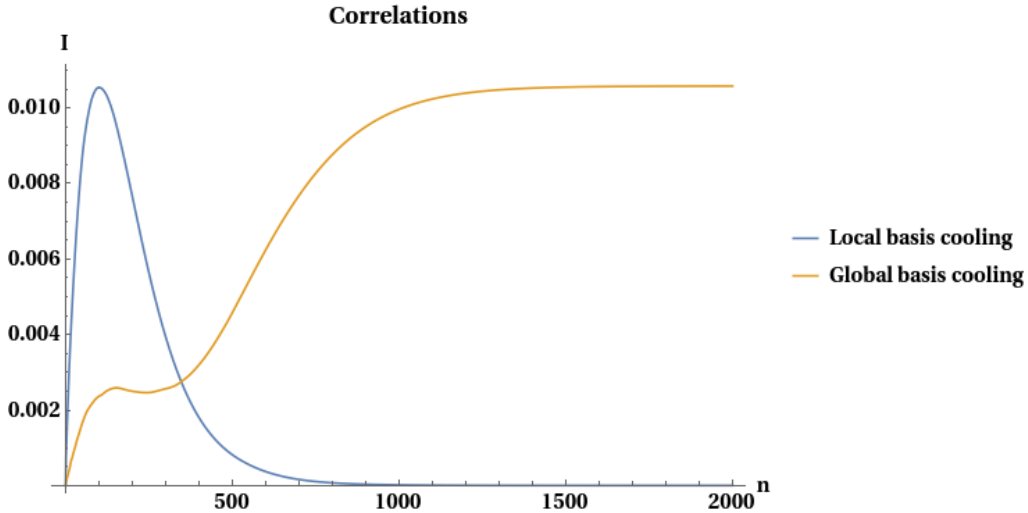


Figure B.4: Finite-Time Correlations for Global and Local Cooling Protocols. Here we compare the correlations generated in the global and local cooling protocols, as described in the text. The orange line depicts the protocol that cools SL globally, disregarding the local energy structure, whereas the blue line corresponds to the procedure that ensures S itself is locally optimally cool. Note that, strictly speaking, the global basis cooling protocol includes a final unitary to be implemented at the last timestep, which yields the optimally cool S state; this is not shown above (in order to highlight the distinction between the behaviour of correlations for both protocols) but it implies that at any finite time, the joint states achieved by either protocol are unitarily equivalent. In particular, for enough timesteps, the state achieved by the global cooling protocol is arbitrarily close to being unitarily equivalent to a correlation-free product state, which has the coolest possible state of S as its marginal (i.e., the asymptotic state achieved by the local basis cooling protocol). The values used for this simulation are: $d_S = 2$, $d_M = 3$, $k = 5$, $\ell = 3$, $E_1 = 1$, $\mathcal{E}_1 = 0.5$, $\mathcal{E}_2 = 1.2$ and $\beta = 0.2$, in natural units where k_B is set to 1.

We now analyse the evolution of correlations in SL , as measured by the mutual information $I(S : L) := H(\varrho_S) + H(\varrho_L) - H(\varrho_{SL})$, where $H(X) := -\text{tr}[X \log(X)]$ is the von Neumann entropy. In both protocols, the joint state begins as a tensor product and therefore has no correlations. Moreover, the asymptotic state of both protocols is also correlation-free, as was shown in Appendix B.2. It is of particular interest to note that the asymptotic state of the global protocol always has a product state in its unitary orbit that has the coldest possible state that S can be brought to as its marginal.

Nonetheless, although both protocols start and end with states that are completely decorrelated, correlations do build up for both protocols at finite steps, before decreasing asymptotically as shown in Figure B.4. The finite-time behaviour of the correlations generally depends on the full spectrum of SL and the number of steps performed (as does the cooling behaviour). In particular—in contrast to the coolness of S —the behaviour of correlations is non-monotonic, and one cannot even establish a hierarchy between the amount of correlations at any finite time of either protocol. Having presented these initial insights, we leave the full analysis of the role of correlations in quantum cooling as an interesting open avenue for pursuit.

B.4 Step-Wise Optimal Protocol and Finite Time Comparisons

Here we provide some analysis on the finite time behaviour for the cooling strategies discussed throughout the main text. It is important to note that the finite time properties in general depend upon the details of the full complex energy spectrum of the machines; nonetheless, we have the following observations.

We first detail the step-wise optimal protocol, briefly described in the main text and prove its optimality.

Definition B.2 (Step-Wise Optimal Cooling Unitary). Given a joint state ϱ_{SLR} , let V_{SLR}^{opt} be the unitary that reorders the eigenvalues of ϱ_{SLM} within each block partitioned by R such that the largest is in the subspace $|000\rangle\langle 000|_{SLR}$, second largest in $|001\rangle\langle 001|_{SLR}$, third largest in $|002\rangle\langle 002|_{SLR}$, \dots , $(d_M^{k-\ell})^{\text{th}}$ largest in $|010\rangle\langle 010|_{SLR}$, and so on until the smallest eigenvalue is in $|d_S - 1, d_M^\ell - 1, d_M^{k-\ell} - 1\rangle\langle d_S - 1, d_M^\ell - 1, d_M^{k-\ell} - 1|_{SLR}$, i.e., perform

$$U_{SLR}^{\text{opt}} \varrho_{SLR} U_{SLR}^{\text{opt} \dagger} = \sum_{\mu=0}^{d_S-1} \sum_{\nu=0}^{d_M^\ell-1} \sum_{\omega=0}^{d_M^{k-\ell}-1} \lambda_{\mu \cdot d_M^k + \nu \cdot d_M^{k-\ell} + \omega}^\downarrow |\mu\nu\omega\rangle\langle \mu\nu\omega|_{SLR}, \quad (\text{B.26})$$

where λ^\downarrow denotes the vector of eigenvalues of ϱ_{SLR} labelled in non-increasing order.

Theorem B.1 (Step-Wise Optimal Cooling Protocol). *By applying the unitary defined in Eq. (B.26) at each step, the cooling protocol is step-wise optimal.*

In the Markovian case, the step-wise optimal protocol simply considers all of the eigenvalues of the joint system-machine at each timestep and optimally reorders them such that the system is as cool as possible. However, such a protocol does not ensure step-wise optimality when memory is present: Here, not only must we optimally cool the system by rearranging the eigenvalues of the total accessible state at each step, but we must also ensure that this accessible state at each step is as cool as possible given its history. As the only information pertaining to the history is transmitted by the system SL , this means that the optimal protocol must at each step optimally cool S , and then subject to this constraint, optimally cool the memory carriers L which go on to further cool the system at later times.

Proof. We first need to show that $\varrho_S^{\text{opt}} := \text{tr}_{LR} [U_{SLR}^{\text{opt}} \varrho_{SLR} U_{SLR}^{\text{opt} \dagger}]$ obtained from Eq. (B.26) majorises all of the reachable marginal states of S ; this problem reduces to a constrained rearrangement of the eigenvalues of the entire system, i.e., the eigenvalues are to be arranged optimally with respect to certain eigenspaces. Since majorisation theory is independent of the eigenbasis, we choose the energy eigenbasis for simplicity.

To obtain the eigenspectrum of the system S that majorises all of the reachable states under unitary transformations on SLR , note that the output state of the entire system

can be written in the form of

$$\tilde{U}_{SLR}^{\text{opt}} \varrho_{SLR} \tilde{U}_{SLR}^{\text{opt} \dagger} = \sum_{\mu=0}^{d_{SL}-1} \sum_{\eta=0}^{d_M^k-1} \lambda_{\mu \cdot d_M^k + \eta}^{\downarrow} |\mu\eta\rangle \langle \mu\eta|_{SLR}, \quad (\text{B.27})$$

where $|\mu\eta\rangle_{SLR} = |\mu\rangle_S \otimes |\eta\rangle_{LR}$ and $d_{SL} = d_S d_M^\ell$. By the ordering of the eigenvalues that the unitary performs, it is straightforward to see that the S marginal following the optimal transformation majorises all others in the unitary orbit.

Second, we show that the state of the memory carriers after applying the optimal unitary, i.e., $\varrho_L^{\text{opt}} := \text{tr}_{SR} [U_{SLR}^{\text{opt}} \varrho_{SLR} U_{SLR}^{\text{opt} \dagger}]$, also majorises all of the reachable states of L given the mentioned majorisation condition. We must therefore rearrange the eigenvalues of ϱ_{SLR} within each block corresponding to a fixed μ , i.e., sort $\{\lambda_{\mu \cdot d_M^k + \eta}^{\downarrow}\}_{\eta=0}^{d_M^k-1}$ in such a way that the ν^{th} largest $d_M^{k-\ell}$ eigenvalues are placed in the ν^{th} eigenspace of the system L , which gives the state ϱ_L^{opt} that majorises all of those reachable via unitary transformations on SLR . To do so, we rearrange the eigenvalues of the joint SLR system as $\lambda_{\mu \cdot d_M^k + \nu \cdot d_M^{k-\ell} + \omega}^{\downarrow}$ via the unitary transformation defined in Eq. (B.26), where $|\mu\nu\omega\rangle_{SLR} = |\mu\rangle_S \otimes |\nu\rangle_L \otimes |\omega\rangle_R$. It is clear that the reduced state satisfies the required majorisation condition for ϱ_L , i.e., for all μ , we have

$$\nu > \nu' \Rightarrow \lambda_{\mu \cdot d_M^k + \nu \cdot d_M^{k-\ell} + \omega}^{\downarrow} \geq \lambda_{\mu \cdot d_M^k + \nu' \cdot d_M^{k-\ell} + \omega'}^{\downarrow} \quad \forall \omega, \omega' \in \{0, 1, \dots, d_M^{k-\ell} - 1\}, \quad (\text{B.28})$$

where this inequality holds due to the eigenvalue ordering of joint state of SLR .

Finally, we show that the output state of the system SL from Eq. (B.26) majorises all of those reachable states of SL . To do so, must show that ν^{th} largest $d_M^{k-\ell}$ eigenvalues of SLR only contribute to the ν^{th} eigenvalue of SL . This statement follows from

$$\mu \cdot d_M^k + \nu \cdot d_M^{k-\ell} > \mu' \cdot d_M^k + \nu' \cdot d_M^{k-\ell} \Rightarrow \lambda_{\mu \cdot d_M^k + \nu \cdot d_M^{k-\ell} + \omega}^{\downarrow} \geq \lambda_{\mu' \cdot d_M^k + \nu' \cdot d_M^{k-\ell} + \omega'}^{\downarrow} \quad \forall \omega, \omega'. \quad (\text{B.29})$$

Eq. (B.29) states that under the protocol considered, one achieves the SL state that majorises all other reachable states via unitaries on SLR . We now need to show that achieving this at every finite timestep j is necessary for subsequent optimal cooling, i.e., that any other protocol is suboptimal. By the stability of majorisation under tensor products [367], we know that $\varrho_{SL}^{(j)\text{opt}} \otimes \tau_R$, where $\varrho_{SL}^{(j)\text{opt}} := \text{tr}_R [\tilde{U}_{SLR}^{\text{opt}} \sigma_{SLR}^{(j)} \tilde{U}_{SLR}^{\text{opt} \dagger}]$ for any global state $\sigma_{SLR}^{(j)}$, majorises all of the states $\varrho_{SL}^{(j)} \otimes \tau_R$, where $\varrho_{SL}^{(j)}$ is generated by any other protocol and τ_R are the thermal bath machines to be added at said timestep. This majorisation relation cannot be changed by performing the optimal SLR unitary on $\varrho_{SL}^{(j)\text{opt}} \otimes \tau_R$ and any other unitary on $\varrho_{SL}^{(j)} \otimes \tau_R$ as the next step of the transformation, with the former therefore yielding $\varrho_{SLR}^{(j+1)\text{opt}}$ and the latter some suboptimal $\varrho_{SLR}^{(j+1)}$. Lastly, invoking the subspace majorisation result of Lemma B.2, it follows that $\varrho_{SL}^{(j+1)\text{opt}} \succ \varrho_{SL}^{(j+1)}$.

Thus, we have shown that at each step of the protocol, we have reached the optimal SL state possible given the history; it is important to note that at this level, the process is Markovian, allowing for an inductive extension of the above argumentation to hold. By

further invoking Lemma B.2 on the level of SL at each timestep, we yield the optimally cool state of the system S , thereby completing the proof. \square

B.5 Relation to Heat-Bath Algorithmic Cooling (HBAC) and State-Independent Asymptotically Optimal Protocol

Here, we propose a general and robust Heat-Bath Algorithmic Cooling (HBAC) technique, which we show to be a special case of our generalised collision model, to optimally cool down a target system in the limit of infinite cycles. To obtain the cooling limit most rapidly, in general one must adapt the operations based on the state of SL output by the dynamics at the most recent step. However, via the correspondence between the generalised collision model and HBAC, we can show that not only it is possible to cool down the system by a state-independent, fixed sequence of operations, but also that the protocol converges to the optimally cool state in the asymptotic limit. The result hence draws attention to the fact that in the limit of infinitely many repeated cycles, the dimension of the memory carriers of the protocol (not necessarily knowledge about the state at intermediate times) plays an important role and can already lead to exponential improvement over the Markovian case; in fact, perhaps surprisingly, the role of memory depth is more significant than that of the ability of the agent to implement multi-partite interactions between the system and machines at each step (although, of course, the number of memory carriers is upper bound by how multi-partite the interactions are allowed to be).

Here we will consider the effect of adding compression systems (in the terminology of the HBAC community) or a number of machines that carry memory forward (in the language of our generalised collision model) for a non-adaptive cooling protocol in which a fixed interaction between the target system and a subset of machines at each timestep is repeated infinitely many times. As we have previously, we assume that k machines interact with the system at each step and ℓ of them carry memory forward to the next step. This means that $(k - \ell)$ fresh machines and ℓ memory carriers participate in the interaction at each timestep. We fix at the outset, for any given choice of these parameters, the dimension of the machines d_M , which (along with k and ℓ) fixes the control complexity in each of the many cases we will look at, and we also fix the temperature at which everything begins, β .

In Appendix B.1, we showed how the dynamics of the system S in the non-Markovian collision model can be described by Markovian dynamics on the larger system SL (with total dimension $d_S d_M^\ell$); in the HBAC community, the larger system of such an embedding is known as the computation system, which comprises the original target and what are often referred to as compression or refrigerant systems. In this case, the system SL interact with $k - \ell$ fresh machines (with total dimension $d_M^{(k-\ell)}$) with maximum energy




Figure B.5: Equivalence Between HBAC and Generalised Collision Model. The circuit for a HBAC protocol applied to a quantum system S with one compression system (labelled L to be consistent with our notation) and one reset system R . The compression systems store memory of previous interactions (red), whereas the reset ones are assumed to rethermalise with a bath at inverse temperature β between each step of the protocol (orange). Noting that the “reset” step has the effect of tracing out the system and preparing a fresh one in the thermal state with the same temperature as the bath, it is clear that HBAC is equivalent to generalised collision model we consider (here, $k = 2, \ell = 1$). Further comparison with Fig. B.3 highlights that HBAC need not require the agent to control the compression systems with high fidelity for the entire duration of the protocol: By making appropriate swaps, one only needs to track the ℓ compression systems / memory carriers for ℓ timesteps.

gap $(k - \ell)\mathcal{E}_{\max}$; this is known as the reset system, since these are the machines that are discarded after each interaction step, modelling a rethermalisation with the external environment. One can decompose the total Hilbert space into the computation part and the reset part R , i.e., $\mathcal{H}_{SLR} = \mathcal{H}_{SL} \otimes \mathcal{H}_R$, where here R refers to all of the reset machines comprising the environment. At any timestep, the dynamics of the system SL , which arises from unitary evolution on the system SLR , is given by (with identity maps implied on the parts of R that do not yet take part in the interaction)

$$\varrho_{SL}^{(n)}(\beta, k, \ell) = \Lambda_{k,\ell}^{(n)}(\beta) [\varrho_{SL}^{(n-1)}(\beta, k, \ell)] := \text{tr}_R \left[U_{k,\ell}^{(n)} (\varrho_{SL}^{(n-1)}(\beta, k, \ell) \otimes \varrho_R^{(0)}(\beta, k, \ell)) U_{k,\ell}^{(n)\dagger} \right]. \quad (\text{B.30})$$

Note that $\varrho_R^{(0)}(\beta, k, \ell) = \bigotimes_{j=\ell+1}^k \tau_{M_j}$ is fixed and the same at each step of the protocol as it refers to the $k - \ell$ fresh machines taken from a thermal bath. In Fig. B.5, we depict the equivalence between the standard HBAC protocol and the generalised collision model formalism.

We now wish to consider a non-adaptive protocol, in which the agent is only allowed to repeatedly apply a fixed unitary operation, i.e., $U_{k,\ell}^{(n)} = U_{k,\ell} \forall n$. The dynamics can then be simplified to

$$\varrho_{SL}^{(n)}(\beta, k, \ell) = \circ^n \Lambda_{k,\ell}(\beta) [\varrho_{SL}^{(0)}(\beta, \ell)], \quad (\text{B.31})$$

where $\varrho_{SL}^{(0)}(\beta, \ell) = \tau_S(\beta) \bigotimes_{j=1}^{\ell} \tau_{M_j}(\beta)$ and $\circ^n \Lambda_{k,\ell}(\beta)$ is an n -fold concatenation of the dynamical map induced between any pair of timesteps, with $\Lambda_{k,\ell}(\beta)$ defined such that $\Lambda_{k,\ell}(\beta) [X_{SL}] := \text{tr}_R \left[U_{k,\ell} (X_{SL} \otimes \varrho_R^{(0)}(\beta, k, \ell)) U_{k,\ell}^\dagger \right]$. This dynamical map is thus independent of the timestep and fully determined by the unitary $U_{k,\ell}$ and the initial state of the fresh machines. In the following, we will show that it is possible to asymptotically reach the ultimate cooling limit via such a non-adaptive protocol.

The unitary U acts to swap every neighbouring element on the diagonal part of the global density matrix in the subspace $\mathcal{H}_{SL} \otimes \mathcal{H}_G$ and leave the other elements untouched. We now focus on the transformation of the diagonal elements on the global space SLR under such dynamics. We write the initial state as

$$\begin{aligned}\varrho_{SLR}^{(0)} &= \sum_{r=0}^{2d_S d_M^\ell - 1} \xi_r^{(0)} |r\rangle\langle r|_{SLR} + \sum_{r=2d_S d_M^\ell}^{d_S d_M^k - 1} \xi_r^{(0)} |r\rangle\langle r|_{SLR} \\ &= \alpha_{k\ell} \varrho_{SLG}^{(0)} \oplus (1 - \alpha_{k\ell}) \varrho_{SL\tilde{G}}^{(0)},\end{aligned}\quad (\text{B.35})$$

where $\varrho_{SLG}^{(0)}$ and $\varrho_{SL\tilde{G}}^{(0)}$ are normalised density matrices and $\alpha_{k\ell} = (1 + e^{-\beta(k-\ell)\mathcal{E}_{\max}})/\mathcal{Z}_R$. After applying the unitary U , we have

$$\begin{aligned}\varrho_{SLR}^{(1)} &= U \varrho_{SLR}^{(0)} U^\dagger = \alpha_{k\ell} V \varrho_{SLG}^{(0)} V^\dagger \oplus (1 - \alpha_{k\ell}) \varrho_{SL\tilde{G}}^{(0)} \\ &= \alpha_{k\ell} \left[\xi_0^{(0)} |0\rangle\langle 0|_{SLR} + \xi_{d_S d_M^k - 1}^{(0)} |d_S d_M^k - 1\rangle\langle d_S d_M^k - 1|_{SLR} \right. \\ &\quad \left. + \sum_{r=0}^{d_S d_M^\ell - 2} \left(\xi_{2r+2}^{(0)} |2r+1\rangle\langle 2r+1|_{SLR} + \xi_{2r+1}^{(0)} |2r+2\rangle\langle 2r+2|_{SLR} \right) \right] \oplus (1 - \alpha_{k\ell}) \varrho_{SL\tilde{G}}^{(0)}.\end{aligned}\quad (\text{B.36})$$

It is clear that the output state is also diagonal in the energy eigenbasis. One can easily obtain the reduced state of the system SL after one timestep from Eq. (B.36) by taking a partial trace over R :

$$\begin{aligned}\varrho_{SL}^{(1)} &= \text{tr}_R \left[\varrho_{SLR}^{(1)} \right] \\ &= \alpha_{k\ell} \left[\frac{(p_0^{(0)} + p_1^{(0)})}{1 + e^{-\beta(k-\ell)\mathcal{E}_{\max}}} |0\rangle\langle 0|_{SL} + \frac{(p_{d_S d_M^\ell - 2}^{(0)} + p_{d_S d_M^\ell - 1}^{(0)}) e^{-\beta(k-\ell)\mathcal{E}_{\max}}}{1 + e^{-\beta(k-\ell)\mathcal{E}_{\max}}} |d_S d_M^\ell - 1\rangle\langle d_S d_M^\ell - 1|_{SL} \right. \\ &\quad \left. + \sum_{r=1}^{d_S d_M^\ell - 2} \frac{(p_{r-1}^{(0)} e^{-\beta(k-\ell)\mathcal{E}_{\max}} + p_{r+1}^{(0)})}{1 + e^{-\beta(k-\ell)\mathcal{E}_{\max}}} |r\rangle\langle r|_{SL} \right] + (1 - \alpha_{k\ell}) \sum_{r=0}^{d_S d_M^\ell - 1} p_r^{(0)} |r\rangle\langle r|_{SL}.\end{aligned}\quad (\text{B.37})$$

Since the output state on \mathcal{H}_{SLR} has a block-diagonal structure with respect to this subspace decomposition, it is locally classical, i.e., has diagonal marginals with respect to the local energy eigenbasis. Therefore, the dynamics of the relevant part of the reduced state can be described in terms of a classical stochastic matrix acting on SL (instead of a CPTP map as would be required if coherences were relevant). In addition, this stochastic matrix is independent of the timestep (since the protocol is non-adaptive) and the SL state at each time. This allows us to describe the evolution of the target system under this protocol via a time-homogeneous Markov process.

Since the unitary applied does not create coherence in the marginals, it is convenient to introduce a notation for the vectorised form of the diagonal entries of the SL state, i.e., $\mathbf{p}_{SL} := \text{diag}\{p_r\}_{r=0}^{d_S d_M^\ell - 1}$, where p_r are the eigenvalues of the state ϱ_{SL} ; since the density matrix is a unit trace positive operator, it follows that the vector \mathbf{p}_{SL} has non-negative

entries that sum to 1, i.e., it is a stochastic vector. Then, the state transformation of the system SL between each step of the protocol can be written as

$$\mathbf{p}_{SL}^{(1)} = \left(\alpha_{k\ell} \mathbb{V} \left((k-\ell) \mathcal{E}_{\max} \right) + (1 - \alpha_{k\ell}) \mathbb{1} \right) \mathbf{p}_{SL}^{(0)} =: \mathbb{T} \mathbf{p}_{SL}^{(0)}. \quad (\text{B.38})$$

where \mathbb{T} describes the transition matrix for the Markovian process and the matrix $\mathbb{V}(\epsilon)$ is given by

$$\mathbb{V}(\epsilon) = \frac{1}{1 + e^{-\beta\epsilon}} \begin{bmatrix} 1 & 1 & \dots & 0 \\ e^{-\beta\epsilon} & 0 & 1 & \dots & 0 \\ 0 & e^{-\beta\epsilon} & 0 & \dots & 0 \\ 0 & 0 & \dots & \ddots & \vdots \\ 0 & 0 & \dots & e^{-\beta\epsilon} & e^{-\beta\epsilon} \end{bmatrix}_{d_S d_M^\ell \times d_S d_M^\ell}. \quad (\text{B.39})$$

Since we apply the fixed unitary at each step and the transition matrix is independent of the state of SL , the state transformation of SL after n steps can be written as

$$\mathbf{p}_{SL}^{(n)} = \mathbb{T}^n \mathbf{p}_{SL}^{(0)}. \quad (\text{B.40})$$

In order to obtain the asymptotic state of the system, we investigate the eigenvalues of the transition matrix given in terms of the two matrices \mathbb{V} and $\mathbb{1}$, which allows us to compute the eigenvalues of \mathbb{T} . The eigenvalues of the matrix \mathbb{V} are presented in Ref. [87]: \mathbb{V} has a unique eigenvalue $\nu_0 = 1$, with the remaining eigenvalues given by

$$\nu_q = \frac{2e^{-\frac{\beta}{2}(k-\ell)\mathcal{E}_{\max}} \cos\left(\frac{q\pi}{d_S d_M^\ell}\right)}{1 + e^{-\beta(k-\ell)\mathcal{E}_{\max}}} \quad \forall q \in \{1, \dots, d_S d_M^\ell - 1\}. \quad (\text{B.41})$$

Since $\mathbb{1}$ is diagonal with respect to any orthonormal basis and has uniform eigenvalues, it is straightforward to show that the eigenvalues of \mathbb{T} are obtained by:

$$\lambda_q = \alpha_{k\ell} \nu_q + (1 - \alpha_{k\ell}). \quad (\text{B.42})$$

Thus, \mathbb{T} also has a unique eigenvalue 1; the eigenvector associated to this value is the steady state solution of dynamics. Moreover, \mathbb{T} also has the same eigenvectors as \mathbb{V} , since those associated to $\mathbb{1}$ are trivial. We can then obtain the asymptotic state of the system SL under a constraint on its initial state, which turns out to only depends on the macroscopic properties of the system and the environment [87]:

$$\mathbf{p}_{SL}^* = \lim_{n \rightarrow \infty} \mathbb{T}^n \mathbf{p}_{SL}^{(0)} = \left\{ \frac{e^{-\beta(k-\ell)q\mathcal{E}_{\max}}}{\mathbb{Z}_{SL}(\beta, (k-\ell)\mathcal{E}_{\max})} \right\}_{q=0}^{d_S d_M^\ell - 1}. \quad (\text{B.43})$$

This steady state gives the eigenvalues of the optimally cool achievable state if $\mathbf{p}_{SL}^{(0)} \prec \mathbf{p}_{SL}^*$. So far, we have shown how one can reach the optimally asymptotic state of SL by employing the fixed unitary in Eq. (B.32) at each iteration. From this asymptotic state, one can easily obtain the coolest achievable reduced state for the system S , i.e., ϱ_S^* . \square

In the non-adaptive scenario, one can further investigate how many repetitions of the cycle are required to achieve the asymptotic state (within a given tolerance). One useful measure for the number of iterations is the mixing time of a Markov process to reach a distance less than η to the desired state, i.e., $t_{\text{mix}} := \min(n) : |\mathbf{p}_{SL}^{(n)} - \mathbf{p}_{SL}^*| \leq \eta$. This mixing time can be upper bounded by a function of difference between the largest and second largest eigenvalues, $\Delta := \lambda_0 - \lambda_1$, as follows

$$t_{\text{mix}}(\eta) \leq \frac{1}{\Delta} \log \left(\frac{1}{\eta p_{d_S d_M}^*} \right). \quad (\text{B.44})$$

For the protocol considered above, the spectral gap can be explicitly calculated

$$\begin{aligned} \Delta &= \lambda_0 - \lambda_1 = 1 - \alpha_{k\ell} \nu_1 - (1 - \alpha_{k\ell}) \\ &= \alpha_{k\ell} \left(1 - \frac{2e^{-\frac{\beta}{2}(k-\ell)\mathcal{E}_{\max}} \cos\left(\frac{\pi}{d_S d_M}\right)}{1 + e^{-\beta(k-\ell)\mathcal{E}_{\max}}} \right) \\ &\geq \frac{1 + e^{-\beta(k-\ell)\mathcal{E}_{\max}}}{(\mathcal{Z}_M(\beta))^{k-\ell}} \left(\frac{(1 - e^{-\frac{\beta}{2}(k-\ell)\mathcal{E}_{\max}}})^2}{1 + e^{-\beta(k-\ell)\mathcal{E}_{\max}}} \right) \\ &= \frac{(1 - e^{-\frac{\beta}{2}(k-\ell)\mathcal{E}_{\max}}})^2}{(\mathcal{Z}_M(\beta))^{k-\ell}}. \end{aligned} \quad (\text{B.45})$$

Then we have

$$t_{\text{mix}}(\eta) \leq \frac{(\mathcal{Z}_M(\beta))^{k-\ell}}{(1 - e^{-\frac{\beta}{2}(k-\ell)\mathcal{E}_{\max}}})^2} \log \left(\frac{1}{\eta p_0^* e^{-\beta(k-\ell)(d_S d_M - 1)\mathcal{E}_{\max}}} \right). \quad (\text{B.46})$$

This result provides an estimate for the number of iterations of the protocol to reach the optimally cool system.

We now compare the cooling performance between adaptive and non-adaptive strategies for a given choice of memory structure. In the non-adaptive strategy, the rate of cooling is determined completely by the spectral gap Δ in Eq. (B.45), as the same dynamics is repeated at each step. In the adaptive scenario, this is no longer the case and a single parameter does not dictate the rate of convergence to the asymptotic state. Instead, in general, the cooling rate depends upon the entire energy structure of the system and all machines, making a closed form expression difficult to derive. Nonetheless, we can describe the solution to the problem of reaching a step-wise provably optimal system state at finite times as a protocol, as done in the main text. This protocol converges to the same asymptotic value as the non-adaptive case, but offers a finite time advantage, as shown in Fig. B.6.

Lastly, in Fig. B.7, we compare various memory structures (i.e., values of k, ℓ) with respect to the optimal adaptive protocol. In order to do so in a meaningful way, we compute the ground state population of the system for after a fixed number $m = k + (n-1)(k-\ell)$ machines have been exhausted. If one were to compare the ground state populations

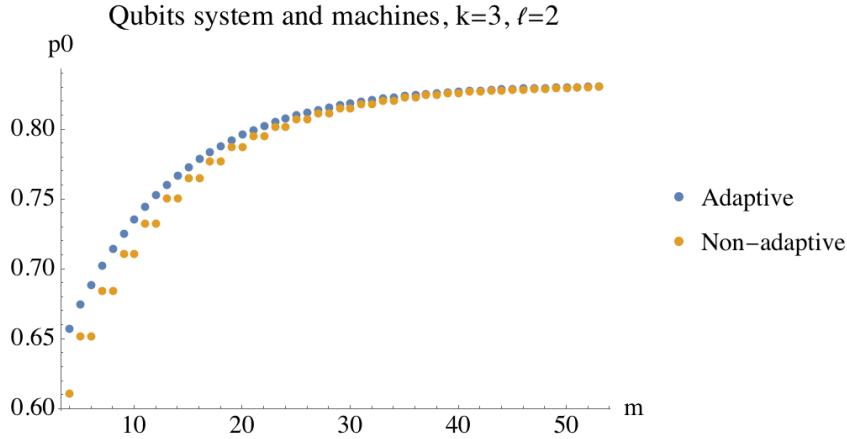


Figure B.6: Finite-Time Advantage with Adaptive Protocols. Here we compare the adaptive and non-adaptive protocols for a qubit system (with $E_{\max} = 1$) interacting with qubit machines (with $\mathcal{E}_{\max} = 2$) with initial temperature $\beta = 0.2$ and memory structure given by $k = 3, \ell = 2$. In the adaptive scenario, we make use of the step-wise optimal protocol described in the main text; in the non-adaptive, the unitary in Eq. (B.33) is repeatedly implemented. We see that, although both scenarios asymptotically converge to the same ground state population, the adaptive protocol outperforms the non-adaptive one at finite times. This behaviour is more pronounced for larger dimensions. On the other hand, the non-adaptive, state-independent protocol is more robust and as it rapidly converges to the asymptotically-optimal state, therefore better suited to practical implementations. Note that, in this case, every second step of the non-adaptive protocol does not act to cool the system.

after n unitaries had been implemented, for various k, ℓ , one would be making an unfair comparison with respect to the total resources at hand; e.g., after three unitaries with $k = 4, \ell = 3$ the experimenter has used six machines, whereas for $k = 3, \ell = 0$, they have used nine. Comparing various scenarios at fixed values of m provides insight into how cool the system can be prepared after all constituents of a finite sized environment are used up for the given memory structure. This change of perspective comes at the cost of the fact that the number of physical unitaries n needed to be implemented in order to exhaust the resources (quantified by m) now varies; e.g., to use six machines with $k = 4, \ell = 3$ takes three unitaries, whereas with $k = 1, \ell = 0$ it takes six. Lastly note that not all values of k, ℓ are valid for a given m , due to the restriction that n must be an integer.




Figure B.7: Finite-Time Behaviour of Step-Wise Optimal Protocol. Here we plot the ground state population after m machines have been completely exhausted, for various values of k, ℓ as described in the text (we use a qubit system and qubit machines, with $\beta = 0.2, E_{\max} = 1$ and $\mathcal{E}_{\max} = 2$). The finite time behaviour generally depends upon the full structure of the energy spectra, but already here we see some interesting effect. For instance, note that at $m = 7$ (first point shown), the $k = 7, \ell = 0$ case provides the best possible cooling, as it allows for a full 8-partite unitary between the system and seven machines. However, for more applications of such a unitary with in the memoryless scenario, the performance can be worse than other cases with more local interactions (smaller k) and longer memory (larger ℓ).



Supplemental Information for Chapter 3

C.1 Connection to Previous Results

In this section we show that the result derived in the main text for the Markovian case (that is, Theorem 3.1) implies the preceding one in Ref. [170]. For the ease of the reader, we restate both results here (changing slightly the terminology of the latter to the one used here).

Theorem C.1. *Let $\{\mathbb{P}_n(x_n, \dots, x_1)\}_{n \leq K}$ be a K -Markovian process (Definition 3.2). Then, the process is also K -classical (Definition 3.1) if and only if there exist a system state ρ_{t_0} (at a time $t_0 \leq t_1$) which is diagonal in the computational basis $\{|x\rangle\}_{x \in \mathcal{X}}$ and a set of propagators $\{\Lambda_{t_j, t_{j-1}}\}_{j=1, \dots, K}$ which are Non-Coherence-Generating-and-Detecting (NCGD) with respect to $\{|x\rangle\}_{x \in \mathcal{X}}$, such that ρ_{t_0} and $\{\Lambda_{t_j, t_{j-1}}\}_{j=1, \dots, K}$ yield $\{\mathbb{P}_n(x_n, \dots, x_1)\}_{n \leq K}$ via Eq. (3.16).*

Theorem C.2 (Theorem 2 of Ref. [170]). *Let $\{\mathbb{P}_n(x_n, \dots, x_1)\}_{n \leq K}$ be the process fixed by the Quantum Regression Formula (QRF) Eq. (3.16), with respect to a set of propagators forming a Completely Positive and Trace Preserving (CPTP) semigroup, i.e., $\Lambda_{t_l, t_j} = e^{\mathcal{L}(t_l - t_j)}$ for any $t_l \geq t_j$ with \mathcal{L} a Lindblad generator [281, 282], and an initial state ρ_{t_0} . Then the process $\{\mathbb{P}_n(x_n, \dots, x_1)\}_{n \leq K}$ is K -classical (Definition 3.1) for any ρ_{t_0} diagonal in the computational basis if and only if the family of propagators is NCGD in the sense that*

$$\Delta \circ \Lambda_{s_3, s_2} \circ \Delta \circ \Lambda_{s_2, s_1} \circ \Delta = \Delta \circ \Lambda_{s_3, s_1} \circ \Delta \quad \forall s_3 \geq s_2 \geq s_1 \geq t_0. \quad (\text{C.1})$$

While the two theorems are clearly related, there are two relevant differences. The new result is more operational in the sense that the statements only depend on the

statistics one obtains by making the measurements in the classical basis at the specified times, whereas the statement in Ref. [170] relies on two underlying assumptions on the Markovianity of the quantum dynamics. The first of these assumptions is that the system multi-time statistics satisfy the QRF [Eq. (3.16)], and the second is that the dynamics forms a semigroup. As we will see below, the second of these assumptions can be relaxed, but the first is crucial if one wants to have the benefit of the statement in Ref. [170], which not only relates possible models for the statistics,¹ but makes also a statement about how the possibility of modelling a process classically implies that the propagators referred to the actual underlying evolution have to be NCGD. To be able to make this connection between the statistics and the underlying quantum evolution, we need to restrict by assumption the types of evolutions we are considering. For the Markov case, considered here, the natural choice is the QRF [Eq. (3.16)], as we discussed in the main text that they are closely related.

To prove the connection between the two theorems, it is useful to consider the following corollary to Theorem 3.1 of the main text:

Corollary C.1. *Let $\{\mathbb{P}_n(x_n, \dots, x_1)\}_{n \leq K}$ be the process fixed by the QRF Eq. (3.16), with respect to a set of divisible propagators and an initial state ρ_{t_0} .*

Let the classical dynamics of this process be invertible, that is, $\mathbb{P}_1(x_j) \neq 0$ for an initial diagonal state that is full-rank, for any $t_j < \infty$. Then, the process $\{\mathbb{P}_n(x_n, \dots, x_1)\}_{n \leq K}$ is K -classical (Definition 3.1) for any ρ_{t_0} diagonal in the computational basis if and only if the family of propagators is NCGD, see Eq. (C.1).

Proof. Let $\{\mathbb{P}_n(x_n, \dots, x_1)\}_{n \leq K}$ be a process satisfying the QRF Eq. (3.16), with respect to a set of divisible propagators satisfying Eq. (C.1). Since the latter implies Eq. (3.21) and the QRF implies that the process is K -Markovian, for any initial diagonal state in the computational basis K -classicality follows from Theorem 3.1.

Conversely, let the assumptions hold and the process be K -classical, in particular for an initial diagonal full-rank state. The NCGD condition follows from the equation

$$\begin{aligned}
 & \text{tr} [\mathcal{P}_{x_3} \circ \Lambda_{s_3, s_2} \circ \Delta \circ \Lambda_{s_2, s_1} \circ \mathcal{P}_{x_1} \circ \Lambda_{s_1} [\rho_0]] \\
 &= \sum_{x_2} \text{tr} [\mathcal{P}_{x_3} \circ \Lambda_{s_3, s_2} \circ \mathcal{P}_{x_2} \circ \Lambda_{s_2, s_1} \circ \mathcal{P}_{x_1} \circ \Lambda_{s_1} [\rho_0]] \\
 &= \sum_{x_2} \mathbb{P}_3(x_3, x_2, x_1) \\
 &= \mathbb{P}_2(x_3, x_1) \\
 &= \text{tr} [\mathcal{P}_{x_3} \circ \Lambda_{s_3, s_1} \circ \mathcal{P}_{x_1} \circ \Lambda_{s_1} [\rho_0]]
 \end{aligned} \tag{C.2}$$

¹Namely that it is equivalent to be able to model the statistics as coming from an evolution satisfying NCGD, as to be able to model it as coming from a classical evolution.

(for $s_3 \geq s_2 \geq s_1$ in \mathcal{T}) by linearity, since from the assumptions (invertibility of the classical dynamics and taking a diagonal, full-rank initial state) we have that $\mathcal{P}_{x_1} \circ \Lambda_{s_1}[\rho_0] \neq 0 \forall x_1, s_1 < \infty$ (for $s_1, s_2, s_3 \rightarrow \infty, \Lambda_{s_i, s_j} \rightarrow \mathbb{1}$ and the NCGD property holds trivially). \square

The only difference between this corollary and Theorem 2 of Ref. [170] is that here we have the divisibility of the “full” propagators and invertibility of the classical propagators in the assumptions, while there the dynamics was assumed to be of Lindblad type. This latter assumption is however strictly stronger, as it implies divisibility and that $\mathcal{P}_{x_j} \circ e^{\mathcal{L}t_j}[\rho] \neq 0 \forall x_j, t_j < \infty$ and for any full-rank ρ , since (finite-dimensional) semigroup evolutions cannot decrease the rank of a state on a finite time [368].

In total, we have shown in this section that Theorem 2 of Ref. [170] can be interpreted as a corollary of Theorem 3.1 by using the connection between the QRF and Markovianity and further restricting to the case of Lindblad evolution. Moreover, Corollary C.1 shows how, by relaxing such restriction and assuming a proper invertibility condition on the classical dynamics, it is possible to establish a one-to-one correspondence between the classicality of a process satisfying the QRF and the NCGD property, where the latter is referred to the propagators of the actual dynamics.

C.2 Absence of Coherence for a Model System: Qubit Coupled to a Continuous Degree of Freedom

In this Appendix, we provide the mathematical details missing in the main text for Example 3.1. We begin with the expression of the global state at time t_1 , immediately before the first measurement:

$$\begin{aligned} \rho_{se}(t_1) = & \int_{-\infty}^{\infty} dp dp' f(p) f^*(p') \left(\rho_{00} e^{i(p-p')t_1} |0p\rangle \langle 0p'| + \rho_{01} e^{i(p+p')t_1} |0p\rangle \langle 1p'| \right. \\ & \left. + \rho_{10} e^{-i(p+p')t_1} |1p\rangle \langle 0p'| + \rho_{11} e^{-i(p-p')t_1} |1p\rangle \langle 1p'| \right). \end{aligned} \quad (\text{C.3})$$

After a measurement at time t_1 with outcome \pm , the state is subsequently given by

$$\rho_{se}^{(\pm)}(t_1) = |\pm\rangle \langle \pm| \otimes \int_{-\infty}^{\infty} dp dp' f_{1;t_1}^{(\pm)}(p, p') |p\rangle \langle p'|, \quad (\text{C.4})$$

where we emphasise that we have a tensor product state and have introduced the amplitude

$$f_{1;t_1}^{(\pm)}(p, p') := \frac{1}{C_{t_1}^{(\pm)}} f(p) f^*(p') \left(\rho_{00} e^{i(p-p')t_1} \rho_{01} e^{i(p+p')t_1} \pm \rho_{10} e^{-i(p+p')t_1} + \rho_{11} e^{-i(p-p')t_1} \right), \quad (\text{C.5})$$

as well as the normalisation factor $C_{t_1}^{(\pm)} = \int_{-\infty}^{\infty} dp |f(p)|^2 (1 \pm 2\text{Re}(\rho_{01} e^{2ipt_1}))$. Note that no $\hat{\sigma}_x$ -coherence is present at this stage.

If we now let the system-environment evolve up to a certain time $\tau > t_1$, the global state will be

$$\begin{aligned} \rho_{se}^{(\pm)}(\tau) &= \frac{1}{2} \int_{-\infty}^{\infty} dp dp' f_{1;t_1}^{(\pm)}(p, p') \left(e^{i(p-p')(\tau-t_1)} |0p\rangle\langle 0p'| \right. \\ &\quad \left. \pm e^{i(p+p')(\tau-t_1)} |0p\rangle\langle 1p'| \pm e^{-i(p+p')(\tau-t_1)} |1p\rangle\langle 0p'| \right. \\ &\quad \left. + e^{-i(p-p')(\tau-t_1)} |1p\rangle\langle 1p'| \right), \end{aligned} \quad (\text{C.6})$$

where the superscript \pm refers to the outcome of the first measurement at time t_1 . The corresponding system state at time τ is then given by tracing out the environmental degrees of freedom, resulting in

$$\rho_s^{(\pm)}(\tau) = \frac{1}{2} \begin{pmatrix} 1 & \pm k^{(\pm)}(\tau, t_1) \\ \pm k^{(\pm)*}(\tau, t_1) & 1 \end{pmatrix}, \quad (\text{C.7})$$

with

$$\begin{aligned} k^{(\pm)}(\tau, t_1) &= \int_{-\infty}^{\infty} dp f_{1;t_1}^{(\pm)}(p, p) e^{2ip(\tau-t_1)} \\ &= \frac{1}{C_{t_1}^{(\pm)}} \int_{-\infty}^{\infty} dp |f(p)|^2 \left(1 \pm \rho_{01} e^{2ipt_1} \pm \rho_{10} e^{-2ipt_1} \right) e^{2ip(\tau-t_1)} \\ &= \frac{1}{C_{t_1}^{(\pm)}} \left(k(\tau - t_1) \pm \rho_{01} k(\tau) \pm \rho_{10} k(\tau - 2t_1) \right). \end{aligned} \quad (\text{C.8})$$

Once again, we see that if the initial system state is a convex mixture of $|+\rangle$ and $|-\rangle$ and $k(t)$ is real (e.g., a Lorentzian distribution centered at 0) then no $\hat{\sigma}_x$ -coherence is present at any time τ . This can be seen because the reduced state can be written as in Eq. (3.34) for the real $\alpha = (\pm k^{(\pm)}(\tau, t_1) + 1)/2$. As a side remark, we note that even if the initial state had some coherences w.r.t. $\hat{\sigma}_x$, these would have been destroyed after the first measurement at time t_1 and, as long as $\rho_{01} \in \mathbb{R}$, would not have been “re-generated” by the subsequent evolution.

Indeed, the argument above can be reiterated for the subsequent measurements; for instance, if we consider the global state after the second measurement at time t_2 , we find

$$\rho_{se}^{(s)}(t_2) = |\pm\rangle\langle \pm| \otimes \int_{-\infty}^{\infty} dp dp' f_{2;t_2,t_1}^{(s)}(p, p') |p\rangle\langle p'| \quad (\text{C.9})$$

with

$$\begin{aligned} f_{2;t_2,t_1}^{(s)}(p, p') &= \frac{1}{C_{t_2,t_1}^{(s)}} f_{1;t_1}^{(\pm)}(p, p') \left(e^{i(p-p')(t_2-t_1)} + \text{sg}(s) e^{i(p+p')(t_2-t_1)} + \text{sg}(s) e^{-i(p+p')(t_2-t_1)} \right. \\ &\quad \left. + e^{-i(p-p')(t_2-t_1)} \right), \end{aligned} \quad (\text{C.10})$$

where s denotes the sequence of $+$ and $-$ obtained in the measurements and $\text{sg}(s)$ the sign of the corresponding product. The entire procedure can be iterated, by replacing $f_{1;t_1}^{(\pm)}(p, p')$ with $f_{2;t_2,t_1}^{(s)}(p, p')$, so that the state at any subsequent time would remain in the

form of Eq. (3.34), with the off-diagonal elements given by a linear combination with real coefficients of the real function $k(t)$ evaluated at different times. In Appendix C.3, we will show how Example 3.1 can be described using a comb representation as introduced in Section 3.5.

C.3 Comb Representation of a Model System: Qubit Coupled to a Continuous Degree of Freedom

In Appendix C.2, we showed the absence of coherence in the state of the system at all times for the dynamics of Example 3.1. To do so, we computed the full system-environment dynamics; however, the full knowledge of the system-environment dynamics is not necessary to understand the multi-time probabilities of observables of the system alone. Moreover, the state of the environment is often not experimentally accessible in practice, as it is typically highly complex. Therefore, it is convenient to only describe the influence that the environment has on the multi-time probabilities. Importantly, this influence, and the resulting correct descriptor of the underlying process, can be deduced by probing the system alone.

Such a descriptor can be derived using the concept of quantum combs [125, 126], which we briefly reviewed in Section 3.5. A quantum comb contains all statistical information that can be inferred about the process it describes (on the set of times upon which it is defined). While here we will construct the comb for Example 3.1 by explicitly solving the system-environment dynamics, it is important to note that it could be reconstructed experimentally by means of measurements on the system alone, without any access to or knowledge of the environmental degrees of freedom, through a generalised tomographic scheme [129].

In slight deviation from the notation of the main text, in this appendix, for better orientation, here we explicitly write the labels of the Hilbert spaces a comb acts on, and the times it is defined upon, as sub- and superscripts, respectively.

As described in Example 3.1, we start with a system-environment state $\eta_{se}(t_0 = 0) = \rho_s(t_0 = 0) \otimes |\varphi^e\rangle\langle\varphi^e|$ where $|\varphi^e\rangle$ is fixed. As shown in Fig. C.1, the initial system state $\rho_s(t_0)$ is associated with the Hilbert space with label 1. The channel

$$\mathcal{C}^{t_1:t_0}(\rho_s) = \mathcal{U}_{t_1,t_0} \rho_s \otimes |\varphi^e\rangle\langle\varphi^e| \quad (\text{C.11})$$

maps the initial system state to the full system-environment state at time t_1 directly before the intervention. The corresponding channel in comb description is given by

$$\begin{aligned} \mathcal{C}_{2\alpha 1}^{t_1:t_0} &= \sum_{i,j} U_{t_1,t_0} \left(|i\rangle\langle j|_2 \otimes \int_{-\infty}^{\infty} dp \int_{-\infty}^{\infty} dq f(p) f^*(q) |p\rangle\langle q|_{\alpha} \right) U_{t_1,t_0}^{\dagger} \otimes |i\rangle\langle j|_1 \\ &= \sum_{i,j} \int_{-\infty}^{\infty} dp \int_{-\infty}^{\infty} dq f(p) f^*(q) e^{i(\phi_i p - \phi_j q) t_1} |i\rangle\langle j|_2 \otimes |p\rangle\langle q|_{\alpha} \otimes |i\rangle\langle j|_1, \end{aligned} \quad (\text{C.12})$$




Figure C.1: Labelling of Hilbert Spaces used for the Comb Description of Example 3.1. The grey box contains the comb $C_{2\alpha 1}^{t_1:t_0}$ and the red box the comb $C_{\beta 43\alpha}^{t_2:t_1}$. The comb $C_{4\beta 321}^{t_2:t_1:t_0}$ corresponds to everything inside the dashed box and consists of the contraction of the two combs $C_{2\alpha 1}^{t_1:t_0}$ and $C_{\beta 43\alpha}^{t_2:t_1}$.

where the superscripts denote the intervention times and the subscripts the Hilbert spaces on which the comb is acting. The object $C_{2\alpha 1}^{t_1:t_0}$ above is nothing other than the Choi state associated with the channel. The dynamics from time t_1 to time t_2 is similarly given by the channel

$$C^{t_2:t_1}(\rho_{se}) = \mathcal{U}_{t_2, t_1} \rho_{se} \quad (\text{C.13})$$

applied to the combined system-environment state directly after the first intervention. Again, this channel admits a Choi state description

$$C_{4\beta 3\alpha}^{t_2:t_1} = \sum_{i,j} \int_{-\infty}^{\infty} dp \int_{-\infty}^{\infty} dq e^{i(\phi_i p - \phi_j q)(t_2 - t_1)} |ipip\rangle\langle jqjq|_{4\beta 3\alpha}. \quad (\text{C.14})$$

The next step is to eliminate the explicit description of the environment state on Hilbert space α . To do this, we contract the Choi states of the two channels described above using the link product \star described in Refs. [125, 126]. This leaves us with the comb describing the dynamics on both times

$$\begin{aligned} C_{4\beta 321}^{t_2:t_1:t_0} &= C_{4\beta 3\alpha}^{t_2:t_1} \star C_{2\alpha 1}^{t_1:t_0} \\ &= \text{tr}_{\alpha} \left[\left(\mathbb{1}_{4\beta 3} \otimes C_{2\alpha 1}^{t_1:t_0 T_{\alpha}} \right) \left(C_{4\beta 3\alpha}^{t_2:t_1} \otimes \mathbb{1}_{21} \right) \right] \\ &= \int_{-\infty}^{\infty} ds \langle s |_{\alpha} \sum_{i,j} \int_{-\infty}^{\infty} dp \int_{-\infty}^{\infty} dq f(p) f^*(q) e^{i(\phi_i p - \phi_j q) t_1} |iqi\rangle\langle jpj|_{2\alpha 1} \\ &\quad \sum_{k,l} \int_{-\infty}^{\infty} dr \int_{-\infty}^{\infty} dt e^{i(\phi_k r - \phi_l t)(t_2 - t_1)} |krkr\rangle\langle ltl|_{4\beta 3\alpha} |s\rangle_{\alpha} \\ &= \sum_{i,j,k,l} \int \dots \int_{-\infty}^{\infty} ds dp dq dr dt \delta(s-q) \delta(s-t) \delta(p-r) f(p) f^*(q) e^{i(\phi_i p - \phi_j q) t_1} |ii\rangle\langle jj|_{21} \\ &\quad e^{i(\phi_k r - \phi_l t)(t_2 - t_1)} |krk\rangle\langle ltl|_{4\beta 3} \\ &= \sum_{i,j,k,l} \iint_{-\infty}^{\infty} ds dp f(p) f^*(s) e^{i(\phi_i p - \phi_j s) t_1} |ii\rangle\langle jj|_{21} e^{i(\phi_k p - \phi_l s)(t_2 - t_1)} |kpk\rangle\langle lsl|_{4\beta 3} \\ &= \sum_{i,j,k,l} \iint_{-\infty}^{\infty} ds dp f(p) f^*(s) e^{i(\phi_i p - \phi_j s) t_1} e^{i(\phi_k p - \phi_l s)(t_2 - t_1)} |kpkii\rangle\langle lsljj|_{4\beta 321}. \quad (\text{C.15}) \end{aligned}$$

We can also describe the projectors corresponding to the observed measurement outcomes using Choi states, e.g., if we measured in the eigenbasis of $\hat{\sigma}_x$ and obtained

outcome +, the corresponding Choi state is given by

$$M^+ = |+\rangle\langle+| \otimes \mathbb{1} \sum_{i,j} |ii\rangle\langle jj| |+\rangle\langle+| \otimes \mathbb{1} = \frac{1}{4} \sum_{i,j,k,l} |ij\rangle\langle lk|. \quad (\text{C.16})$$

Again, using the link product, we can obtain the unnormalised joint system-environment state directly after the second intervention at time t_2 , conditioned on the initial state of the system $\rho_s(0)$ and the interventions M^{x_1}, M^{x_2} (where the superscripts x_i refer to the outcomes) as follows

$$\begin{aligned} \rho_{se}^{(x_2, x_1)}(t_2)_{5\beta} &= C_{4\beta 321}^{t_2:t_1:t_0} \star \rho_s(t_0)_1 \star M_{32}^{x_1} \star M_{54}^{x_2} \\ &= \text{tr}_{4321} \left[\rho_s(t_0)_1^T \otimes M_{32}^{x_1 T_2} \otimes M_{54}^{x_2 T_4} C_{4\beta 321}^{t_2:t_1:t_0} \right]. \end{aligned} \quad (\text{C.17})$$

For instance, if we observed the outcome + twice, the joint state after the second intervention is given by

$$\begin{aligned} \rho_{se}^{(+,+)}(t_2)_{5\beta} &= \sum_{\substack{i,j,k,l, \\ m,n,x,y, \\ a,b,c,d, \\ f,g,h,o}} \langle fgho|_{4321} \rho_{mn}|n\rangle\langle m|_1 \otimes \frac{1}{4} |cx\rangle\langle dy|_{32} \otimes \frac{1}{2} |+\rangle\langle+|_5 \otimes |a\rangle\langle b|_\beta \times \\ &\int_{-\infty}^{\infty} ds dp f(p) f^*(s) e^{i(\phi_i p - \phi_j s)t_1} |ii\rangle\langle jj|_{21} e^{i(\phi_k p - \phi_l s)(t_2 - t_1)} |kpk\rangle\langle lsl|_{4\beta 3} |fgho\rangle_{4321} \\ &= \frac{1}{8} |+\rangle\langle+|_5 \otimes \sum_{i,j,k,l} \rho_{ij} \int_{-\infty}^{\infty} ds dp f(p) f^*(s) e^{i(\phi_i p - \phi_j s)t_1} e^{i(\phi_k p - \phi_l s)(t_2 - t_1)} |p\rangle\langle s|_\beta \\ &= \frac{1}{8} |+\rangle\langle+|_5 \otimes \int_{-\infty}^{\infty} dp ds \tilde{f}_{2;t_2, t_1}^{(+,+)}(p, s) |p\rangle\langle s|_\beta, \end{aligned} \quad (\text{C.18})$$

where we have introduced

$$\begin{aligned} \tilde{f}_{2;t_2, t_1}^{(+,+)}(p, s) &= \sum_{i,j,k,l} \rho_{ij} f(p) f^*(s) e^{i(\phi_i p - \phi_j s)t_1} e^{i(\phi_k p - \phi_l s)(t_2 - t_1)} \\ &= f(p) f^*(s) \left(\rho_{00} e^{i(p-s)t_1} + \rho_{01} e^{i(p+s)t_1} + \rho_{10} e^{-i(p+s)t_1} + \rho_{11} e^{-i(p-s)t_1} \right) \\ &\quad \left(e^{i(p-s)(t_2 - t_1)} + e^{i(p+s)(t_2 - t_1)} + e^{-i(p+s)(t_2 - t_1)} + e^{-i(p-s)(t_2 - t_1)} \right) \\ &= \tilde{f}_{2;t_2, t_1}^{(+,+)}(p, s) \end{aligned} \quad (\text{C.19})$$

and checked consistency with the direct description in Appendix C.2.

Since we are mainly interested in the question of whether the obtained measurement statistics can be explained classically, we restrict our attention to the unnormalised state of the system alone, because the probability of obtaining a specific sequence of measurement outcomes is encoded in the trace of the corresponding system state. Therefore we eliminate the description of the environment by tracing over the Hilbert space β , which we can do directly at the level of the comb itself:

$$\begin{aligned} \tilde{C}_{4321}^{t_2:t_1:t_0} &= \text{tr}_\beta [C_{4\beta 321}^{t_2:t_1:t_0}] \\ &= \sum_{i,j,k,l} \int_{-\infty}^{\infty} dq |f(q)|^2 e^{i(\phi_i - \phi_j)qt_1} e^{i(\phi_k - \phi_l)q(t_2 - t_1)} |kkii\rangle\langle lljj|_{4321}. \end{aligned} \quad (\text{C.20})$$




Figure C.2: Dilation for Example 3.1. Pictorial representation of the quantum combs describing Example 3.1 with two interventions.

Following the same procedure as above, we then obtain the system state after the second intervention

$$\rho_s^{(x_2, x_1)}(t_2)_5 = \tilde{C}_{4\beta 321}^{t_2::t_1:t_0} \star \rho_s(t_0)_1 \star M_{32}^{x_1} \star M_{54}^{x_2}. \quad (\text{C.21})$$

Similarly, the probability to obtain, e.g., twice the measurement result + is given by

$$\mathbb{P}_2(+, t_2; +, t_1) = \text{tr} \left[\rho_s^{(+, +)}(t_2)_5 \right]. \quad (\text{C.22})$$

If we introduce $\tau_n := t_n - t_{n-1}$, by way of induction, we find that

$$\begin{aligned} C^{t_n::t_0} &= \sum_{i_{2n} \dots i_1, j_{2n} \dots j_1} \iint_{-\infty}^{\infty} dp dq f(p) f^*(q) |p\rangle\langle q| \bigotimes_{a=1}^{2n} e^{i(\phi_{i_a} p - \phi_{j_a} q) \tau_a} |i_a i_a\rangle\langle j_a j_a|_{2a, 2a-1}, \\ \tilde{C}^{t_n::t_0} &= \sum_{i_{2n} \dots i_1, j_{2n} \dots j_1} \int_{-\infty}^{\infty} dp |f(p)|^2 \bigotimes_{a=1}^{2n} e^{i(\phi_{i_a} - \phi_{j_a}) p \tau_a} |i_a i_a\rangle\langle j_a j_a|_{2a, 2a-1}, \end{aligned} \quad (\text{C.23})$$

where we suppressed the subscripts of the combs. As above, $C^{t_n::t_0}$ denotes the comb including the outgoing environment and $\tilde{C}^{t_n::t_0}$ the comb describing the system alone, see Fig. C.2 for a pictorial representation. Therefore, the joint probability distribution for sequences of measurement outcomes is given by

$$\mathbb{P}_n(x_n, t_n; \dots; x_1, t_1) = \text{tr} \left[\rho_s(0)_1^T \bigotimes_{a=1}^n (M_{2a+1, 2a}^{x_a})^T \tilde{C}^{t_n::t_0} \right]. \quad (\text{C.24})$$

C.4 Alternative Example for Non-Classical Dynamics that do not Create Coherences

Here, we provide an alternative example of a process where the state of the system is diagonal in the computational basis at all times but does not yield classical statistics. To this end, consider the following circuit (see Fig. C.3): Let the initial system-environment state at time t_0 be a maximally entangled two qubit state φ^+ that undergoes trivial evolution between t_0 and t_1 . At t_1 the system alone is thus in a maximally mixed state ρ_{t_1} . Between t_1 and t_2 , the system and the environment undergo a CPTP map \mathcal{E}_{t_2, t_1} (which




Figure C.3: Non-Classical Process that does not Display Coherences. The state of the system is classical, i.e., it does not contain coherences with respect to the classical basis, at any step of the process. The corresponding statistics do not satisfy the Kolmogorov conditions, though. Potential measurements are depicted as green circles. The blue dotted line signifies the comb of the process (see Sec. 3.5).

could—in principle—be dilated to a unitary map [369], but for conciseness, we restrict ourselves to the relevant part of it), that yields output $|0\rangle$ on the system, if system and environment are in the state φ^+ , and $|1\rangle$ otherwise, i.e., when the system-environment state is orthogonal to φ^+ . Consequently, its action can be written as

$$\mathcal{E}_{t_2, t_1}[\eta] = \text{tr}(\varphi^+ \eta) |0\rangle\langle 0| + \text{tr}[(\mathbb{1} - \varphi^+) \eta] |1\rangle\langle 1|. \quad (\text{C.25})$$

It is easy to check that \mathcal{E}_{t_3, t_2} is indeed CPTP, and the state of the system at t_2 is a convex mixture of $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$ for all possible experimental interventions at t_1 ; there are thus no coherences in the state of the system at any of the times $\{t_1, t_2\}$. However, this process does not satisfy the Kolmogorov condition.

To see this, consider the probabilities for a measurement in the computational basis at t_2 , with no operation performed at t_1 . In this case, the system-environment state before the action of \mathcal{E}_{t_2, t_1} is equal to φ^+ , which means that we have $\rho_{t_2} = |0\rangle\langle 0|$. Consequently, a measurement in the computational basis at t_2 yields the probabilities

$$\mathbb{P}_1(0, t_2) = 1 \quad \text{and} \quad \mathbb{P}_1(1, t_2) = 0. \quad (\text{C.26})$$

On the other hand, performing a measurement at t_1 and discarding the outcomes amounts to performing the completely dephasing map Δ_1 . Immediately after this map, i.e., right before \mathcal{E}_{t_2, t_1} , the system-environment state is of the form

$$\eta_{t_1}^{se} = \frac{1}{2} \sum_{x_1} |x_1\rangle\langle x_1| \otimes |x_1\rangle\langle x_1| = \frac{1}{2} (\varphi^+ + \varphi^-), \quad (\text{C.27})$$

where $\varphi^- = (\sigma_z \otimes \mathbb{1}) \varphi^+ (\sigma_z \otimes \mathbb{1})$ is a Bell state. Consequently, in this case the final system state ρ_{t_2} is of the form $\rho_{t_2} = \frac{1}{2} (|0\rangle\langle 0| + |1\rangle\langle 1|)$. Finally, the obtained probabilities for a measurement in the computational basis at t_2 are

$$\begin{aligned} \mathbb{P}_1^{\Delta_1}(0, t_2) &= \sum_{x_1} \mathbb{P}_2(x_1, t_1; 0, t_2) = \frac{1}{2} \\ \text{and} \quad \mathbb{P}_1^{\Delta_1}(1, t_2) &= \sum_{x_1} \mathbb{P}_2(x_1, t_1; 1, t_2) = \frac{1}{2}, \end{aligned} \quad (\text{C.28})$$

which does not coincide with (C.26). Even though the state of the system is incoherent at all considered times, i.e., appears to be classical, the multi-time statistics do not satisfy the Kolmogorov condition.

C.5 Measure of Non-Classicality

In this appendix we derive the optimal solution of the game which defines our measure of non-classicality $M(C)$ and show that it can be formulated as a linear program. We also derive the dual of this problem for completeness.

In our game, Bob can choose the points in time at which he wants Rudolph to perform projective measurements and those for which Rudolph should not interfere with the natural evolution of the system. This defines a sequence of measurements $T_i(\vec{x}) = \otimes_{t_j \in \tau_i} \Phi_j^+ \otimes_{t_k \in \tau_i^c} P_{x_k}$. Given the choice of any sequence of this form and labelling the obtained outcome sequence of the experiment by \vec{x} , the best strategy for Bob is to announce that the comb that was tested is C if the probability for measuring outcome \vec{x} with said sequence $T_i(\vec{x})$ is higher for C than for $C^{\text{Cl.}}$ (i.e., if $\text{tr}[(C^{\text{Cl.}} - C)T_i(\vec{x})] < 0$), and announcing $C^{\text{Cl.}}$ otherwise. The probability that he is correct when announcing C , given that the outcome obtained was \vec{x} , is given by:

$$\mathbb{P}(C|\vec{x}) = \frac{\mathbb{P}(C, \vec{x})}{\mathbb{P}(\vec{x})} = \frac{\mathbb{P}(\vec{x}|C)}{\mathbb{P}(\vec{x})} \mathbb{P}(C), \quad (\text{C.29})$$

where the prior probability is $\mathbb{P}(C) = 1/2$. Denoting by $S^{\text{Cl.}}$ the set of all \vec{x} such that $\text{tr}[(C^{\text{Cl.}} - C)T_i(\vec{x})] > 0$ and $S_c^{\text{Cl.}}$ its complement, the probability that Bob wins the game is given by

$$\begin{aligned} & \sum_{\vec{x} \in S_c^{\text{Cl.}}} \mathbb{P}(C|\vec{x})\mathbb{P}(\vec{x}) + \sum_{\vec{x} \in S^{\text{Cl.}}} \mathbb{P}(C^{\text{Cl.}}|\vec{x})\mathbb{P}(\vec{x}) \\ &= \frac{1}{2} \left(\sum_{\vec{x} \in S_c^{\text{Cl.}}} \mathbb{P}(\vec{x}|C) + \sum_{\vec{x} \in S^{\text{Cl.}}} \mathbb{P}(\vec{x}|C^{\text{Cl.}}) \right) \\ &= \frac{1}{2} \left(\sum_{\vec{x} \in S_c^{\text{Cl.}}} \text{tr}[CT_i(\vec{x})] + \sum_{\vec{x} \in S^{\text{Cl.}}} \text{tr}[C^{\text{Cl.}}T_i(\vec{x})] \right) \\ &= \frac{1}{2} \left(1 + \sum_{\vec{x} \in S_c^{\text{Cl.}}} (-\text{tr}[CT_i(\vec{x})] + \text{tr}[C^{\text{Cl.}}T_i(\vec{x})]) \right). \end{aligned} \quad (\text{C.30})$$

Assuming that both Alice and Bob play ideally, using Lemma 3.1, the probability $\mathbb{P}_B(C)$ that Bob wins is given by

$$\mathbb{P}_B(C) = \frac{1}{2} [1 + M(C)], \quad (\text{C.31})$$

where $M(C)$ is the solution of

$$\begin{aligned}
 & \text{minimise: } \max_i \sum_{\vec{x} \in S^{\text{Cl.}}} \text{tr}[(C^{\text{Cl.}} - C)T_i(\vec{x})] \\
 & \text{subject to: } S^{\text{Cl.}} = \{\vec{x} | \text{tr}[(C^{\text{Cl.}} - C)T_i(\vec{x})] \geq 0\}, \\
 & C^{\text{Cl.}} = \sum_{y_K, \dots, y_1} \mathbb{P}_K(\vec{y}) P_{y_K} \otimes \dots \otimes P_{y_1} + \chi, \\
 & \text{tr}[\chi \cdot (\otimes_{t_j \in \tau_i} A_j \otimes_{t_k \in \tau_i^c} P_{z_k})] = 0, \\
 & C^{\text{Cl.}} \geq 0, \\
 & \text{tr}_{K^i}[C^{\text{Cl.}}] = \mathbb{1}_{(K-1)^{\circ}} \otimes \Theta_{K-1}, \\
 & \vdots \\
 & \text{tr}_{2^i}[\Theta_2] = \mathbb{1}_{1^{\circ}} \otimes \rho_{1^i}, \\
 & \mathbb{P}_K(\vec{y}) \text{ joint probability distribution,} \tag{C.32}
 \end{aligned}$$

where we defined $A_j := \Phi_j^+ - D_j$ and ρ_{1^i} is a valid quantum state. The hierarchy of partial trace conditions on the comb written above ensure that the overall action of any instrument at a later time cannot influence previous statistics [125, 126].

Starting from the above program, we see that χ does not contribute to the trace, as $\text{tr}[\chi T_i(\vec{x})]$ is, by definition, a marginal of a zero-distribution (due to the third constraint above), see also the proof of Lemma 1). This leaves us with contributions only from the diagonal parts of the operator $C^{\text{Cl.}}$, where the non-zero entries are those that correspond to $\mathbb{P}_K(\vec{y}) P_{y_K} \otimes \dots \otimes P_{y_1}$, which must satisfy $\text{tr}[C^{\text{Cl.}}] = 1$ and $C^{\text{Cl.}} \geq 0$ due to the requirement that $\mathbb{P}_K(\vec{y})$ is a valid probability distribution. Note that for any such an operator, there exists a χ such that the total operator satisfies the additional requirements in the above program, since one simply must add terms of the form $\sum_{y_K, \dots, y_1} \mathbb{P}_K(\vec{y}) P_{y_K, z_K} \otimes \dots \otimes P_{y_1, z_1}$, where the P_{y_j, z_j} are projectors up to a permutation on the input basis (i.e., $P_{y_j, z_j} = |y_j\rangle\langle y_j|_{\circ} \otimes |z_j\rangle\langle z_j|_{\text{i}}$). We are then left with:

$$\begin{aligned}
 & \text{minimise: } \max_i \sum_{\vec{x} \in S^{\text{Cl.}}} \text{tr}[(C^{\text{Cl.}} - C)T_i(\vec{x})] \\
 & \text{subject to: } S^{\text{Cl.}} = \{\vec{x} | \text{tr}[(C^{\text{Cl.}} - C)T_i(\vec{x})] \geq 0\}, \\
 & C^{\text{Cl.}} = \sum_{y_K, \dots, y_1} \mathbb{P}_K(\vec{y}) P_{y_K} \otimes \dots \otimes P_{y_1}, \\
 & \mathbb{P}_K(\vec{y}) \text{ joint probability distribution.} \tag{C.33}
 \end{aligned}$$

Since both C and $C^{\text{Cl.}}$ represent (up to a non-contributing χ term) deterministic quantum combs, we have

$$\sum_{\vec{x}} \text{tr}[(C^{\text{Cl.}} - C)T_i(\vec{x})] = 0 \tag{C.34}$$

and thus

$$\sum_{\vec{x}} \left| \text{tr}[(C^{\text{Cl.}} - C)T_i(\vec{x})] \right| = 2 \sum_{\vec{x} \in S^{\text{Cl.}}} \text{tr}[(C^{\text{Cl.}} - C)T_i(\vec{x})]. \tag{C.35}$$

This allows us to express $M(C)$ as half of the solution of

$$\begin{aligned}
 & \text{minimise: } \max_i \sum_{\vec{x}} \left| \text{tr}[(C^{\text{Cl.}} - C)T_i(\vec{x})] \right| \\
 & \text{subject to: } C^{\text{Cl.}} = \sum_{y_K, \dots, y_1} \mathbb{P}_K(\vec{y}) P_{y_K} \otimes \dots \otimes P_{y_1}, \\
 & \mathbb{P}_K(\vec{y}) \text{ joint probability distribution.}
 \end{aligned} \tag{C.36}$$

In order to transform this program into a Linear Program (**LP**), for every testing sequence $\{T_i(\vec{x})\}_{\vec{x}}$, we define an arbitrary order of the outcomes \vec{x} , i.e, we label them as \vec{x}_j . Then

$$\max_i \sum_{\vec{x}} \left| \text{tr}[(C^{\text{Cl.}} - C)T_i(\vec{x})] \right| \tag{C.37}$$

is the solution of

$$\begin{aligned}
 & \text{minimise: } a \\
 & \text{subject to: } a \geq \sum_j \left| \text{tr}[(C^{\text{Cl.}} - C)T_i(\vec{x}_j)] \right|,
 \end{aligned} \tag{C.38}$$

which is equivalent to

$$\begin{aligned}
 & \text{minimise: } a \\
 & \text{subject to: } a \geq s_i, \\
 & s_i = \sum_j b_{ij}, \\
 & b_{ij} \geq c_{ij} \geq -b_{ij}, \\
 & c_{ij} = \text{tr} \left[(C^{\text{Cl.}} - C)T_i(\vec{x}_j) \right].
 \end{aligned} \tag{C.39}$$

Combining this with the outer minimisation, we finally have that $M(C)$ is half of the solution of

$$\begin{aligned}
 & \text{minimise: } a \\
 & \text{subject to: } a \geq s_i, \\
 & s_i = \sum_j b_{ij}, \\
 & b_{ij} \geq c_{ij} \geq -b_{ij}, \\
 & c_{ij} = \text{tr} \left[(C^{\text{Cl.}} - C)T_i(\vec{x}_j) \right], \\
 & C^{\text{Cl.}} = \sum_{y_K, \dots, y_1} \mathbb{P}_K(\vec{y}) P_{y_K} \otimes \dots \otimes P_{y_1}, \\
 & \mathbb{P}_K(\vec{y}) \text{ joint probability distribution,}
 \end{aligned} \tag{C.40}$$

which is a linear program.

In order to simplify the numerical implementation and the derivation of the dual program, we will also order the vectors \vec{y} (arbitrarily), identify p_k with $\mathbb{P}_K(\vec{y}(k))$, and

define α_{ijk}

$$\text{tr} \left[C^{\text{Cl.}} T_i(\vec{x}_j) \right] = \sum_k p_k \alpha_{ijk} \quad (\text{C.41})$$

for all p_k , i.e.,

$$\alpha_{ijk} = \text{tr} \left[P_{y_K(k)} \otimes \cdots \otimes P_{y_1(k)} T_i(\vec{x}_j) \right] \quad (\text{C.42})$$

for the sequence $y_K(k), \dots, y_1(k)$ corresponding to $\vec{y}(k)$. In addition, we define

$$\beta_{ij} = \text{tr} [C T_i(\vec{x}_j)], \quad (\text{C.43})$$

which allows us to write

$$\begin{aligned} c_{ij} &= \text{tr} \left[(C^{\text{Cl.}} - C) T_i(\vec{x}_j) \right] \\ &= \sum_k p_k \alpha_{ijk} - \beta_{ij}. \end{aligned} \quad (\text{C.44})$$

Then, the above optimisation problem is equivalent to

$$\begin{aligned} &\text{minimise: } a \\ &\text{subject to: } \sum_j b_{ij} - a \leq 0, \\ &\quad \sum_k p_k \alpha_{ijk} - \beta_{ij} - b_{ij} \leq 0, \\ &\quad - \sum_k p_k \alpha_{ijk} + \beta_{ij} - b_{ij} \leq 0, \\ &\quad \sum_k p_k - 1 = 0, \\ &\quad p_k, a, b_{ij} \geq 0. \end{aligned} \quad (\text{C.45})$$

The Lagrangian corresponding to this problem is

$$\begin{aligned} &L(a, p_k, b_{ij}, X_i, Y_{ij}, Z_{ij}, W) \\ &= a \left[1 - \sum_i X_i \right] + \sum_{ij} b_{ij} [X_i - Y_{ij} - Z_{ij}] \\ &\quad + \sum_k p_k \left[\sum_{ij} \alpha_{ijk} (Y_{ij} - Z_{ij}) - W \right] \\ &\quad + W + \sum_{ij} \beta_{ij} (Z_{ij} - Y_{ij}) \end{aligned} \quad (\text{C.46})$$

and the dual function explicitly written

$$\begin{aligned} &q(X_i, Y_{ij}, Z_{ij}, W) \\ &= \inf_{p_k \geq 0, a, b_{ij}} L(a, p_k, b_{ij}, X_i, Y_{ij}, Z_{ij}, W), \end{aligned} \quad (\text{C.47})$$

where we used that $a, b_{ij} \geq 0$ is implicit in the remaining conditions. The dual problem is then given by

$$\begin{aligned}
 & \text{maximise: } W + \sum_{ij} \beta_{ij} (Z_{ij} - Y_{ij}) \\
 & \text{subject to: } \sum_i X_i = 1, \\
 & \quad X_i - Y_{ij} - Z_{ij} = 0 \quad \forall ij, \\
 & \quad \sum_{ij} \alpha_{ijk} (Y_{ij} - Z_{ij}) - W \geq 0 \quad \forall k, \\
 & \quad X_i, Y_{ij}, Z_{ij} \geq 0, \\
 & \quad W \in \mathbb{R},
 \end{aligned} \tag{C.48}$$

which can straightforwardly be reformulated as

$$\begin{aligned}
 & \text{maximise: } \Omega \\
 & \text{subject to: } \Omega \leq \sum_{ij} (\alpha_{ijk} - \beta_{ij}) (2Y_{ij} - X_i) \quad \forall k, \\
 & \quad \sum_i X_i = 1, \\
 & \quad X_i, Y_{ij}, X_i - Y_{ij} \geq 0, \\
 & \quad \Omega \in \mathbb{R}.
 \end{aligned} \tag{C.49}$$

Evidently, the above considerations are amenable to many extensions but that is the matter of future work.

C.6 Non-Discord-Creating Maps

Here, for comprehensiveness, we characterise the set of maps $\Gamma : \mathcal{B}(\mathcal{H}_s^i \otimes \mathcal{H}_e^i) \rightarrow \mathcal{B}(\mathcal{H}_s^o \otimes \mathcal{H}_e^o)$ that map discord-zero states to discord-zero states, where we mean discord-zero with respect to the classical basis. Such system-environment maps form a subset of the Non-Discord-Generating-and-Detecting (**NDGD**) maps of Definition 3.3 [in the sense that a set of them would satisfy Eq. (3.69)] and would thus lead to classical statistics on the level of the system. However, for classical statistics, it is not necessary that the underlying maps do not create discord.

To facilitate notation, throughout this Appendix, we will denote discord-zero states as $D\emptyset$ states, and maps that do not create discord as $D\emptyset$ maps. We have the following lemma:

Lemma C.1 (Structure of DØ maps). *The Choi state G of a DØ map $\Gamma : \mathcal{B}(\mathcal{H}_s^i \otimes \mathcal{H}_e^i) \rightarrow \mathcal{B}(\mathcal{H}_s^o \otimes \mathcal{H}_e^o)$ is of the form*

$$G = \sum_{k,j=1}^{d_s} p_{k|j} \Pi_k^o \otimes \Pi_j^i \otimes O_{jk}^{oi} + G^\perp, \quad (\text{C.50})$$

where $\{\Pi_l^{i/o}\}$ are orthogonal rank-1 projectors on $\mathcal{H}_s^{i/o}$ that are diagonal in the computational basis, $O_{jk}^{oi} \in \mathcal{B}(\mathcal{H}_e^o \otimes \mathcal{H}_e^i)$ is the Choi state of a CPTP map $\Omega_{jk} : \mathcal{B}(\mathcal{H}_e^i) \rightarrow \mathcal{B}(\mathcal{H}_e^o)$, $p_{k|j}$ is a conditional probability distribution, i.e., $\sum_k p_{k|j} = 1$ and $p_{k|j} \geq 0$, and $G^\perp \in \mathcal{B}(\mathcal{H}_s^o \otimes \mathcal{H}_e^o \otimes \mathcal{H}_s^i \otimes \mathcal{H}_e^i)$ is orthogonal to the set of DØ states, i.e., $\text{tr}[(\mathbb{1} \otimes \rho)G^\perp] = 0$ for all DØ states $\rho \in \mathcal{B}(\mathcal{H}_s^i \otimes \mathcal{H}_e^i)$.

Before we prove this lemma, we emphasise its structural relation to the representation of any Maximally Incoherent Operation (MIO), i.e., the structure of maps $\mathcal{F} : \mathcal{B}(\mathcal{H}_s^i) \rightarrow \mathcal{B}(\mathcal{H}_s^o)$ that map incoherent states $\rho \in \Xi \subset \mathcal{B}(\mathcal{H}_s^i)$ onto incoherent states $\rho' = \mathcal{F}[\rho] \in \Xi \subset \mathcal{B}(\mathcal{H}_s^o)$, where Ξ denotes the set of incoherent states with respect to the computational basis. The Choi state F of the map \mathcal{F} is a positive element of $\mathcal{B}(\mathcal{H}_s^o \otimes \mathcal{H}_s^i)$. Choosing a basis $\{\tau_k^o\}_{k=1}^{d_s^2}$ and $\{\omega_j^i\}_{j=1}^{d_s^2}$ for $\mathcal{B}(\mathcal{H}_s^o)$ and $\mathcal{B}(\mathcal{H}_s^i)$, respectively, any F can be written as

$$F = \sum_{j,k} f_{jk} \tau_k^o \otimes \omega_j^i, \quad (\text{C.51})$$

where $f_{jk} \in \mathbb{R}$. We can choose the basis $\{\omega_j^i\}$ to consist of the d_s rank-1 projectors Π_j^i in the computational basis and $d_s(d_s - 1)$ elements $\tilde{\Pi}_s^i$ that are orthogonal to these projectors, i.e., such that $\text{tr}(\Pi_j^i \tilde{\Pi}_s^i) = 0$ (e.g., one could choose the off-diagonal elements $|m\rangle\langle n| + |n\rangle\langle m|$ and $i(|m\rangle\langle n| - |n\rangle\langle m|)$). With this choice of basis elements Eq. (C.51) reads

$$F = \sum_{j,k} f_{kj} \tau_k^o \otimes \Pi_j^i + \sum_{r,s} \tilde{f}_{rs} \tau_r^o \otimes \tilde{\Pi}_s^i. \quad (\text{C.52})$$

Imposing the requirement that \mathcal{F} does not create coherences with respect to the classical basis then yields

$$F = \sum_{j,k} p_{k|j} \Pi_k^o \otimes \Pi_j^i + \sum_{r,s} \tilde{f}_{rs} \tau_r^o \otimes \tilde{\Pi}_s^i, \quad (\text{C.53})$$

where $p_{k|j} \geq 0$, $\sum_k p_{k|j} = 1$, and $\tau_r^o \in \mathcal{B}(\mathcal{H}_s^o)$. Indeed, an F of the form of Eq. (C.53) yields an incoherent output state for any incoherent input state $\rho_{\text{cl}} = \sum_{r=1}^{d_s} q_r \Pi_r^i \in \Xi$:

$$\mathcal{F}[\rho_{\text{cl}}] = \text{tr}_i \left[\left(\mathbb{1}^o \otimes \rho_{\text{cl}}^T \right) F \right] = \sum_{kr} p_{k|r} q_r \Pi_k^o. \quad (\text{C.54})$$

Importantly, Eq. (C.53) constitutes a decomposition of the form $F = F^\parallel + F^\perp$, where $F^\parallel = \sum_{j,k} p_{k|j} \Pi_k^o \otimes \Pi_j^i$ encapsulates the action of \mathcal{F} on incoherent states, and F^\perp is such that all incoherent states lie in its kernel, i.e., $\text{tr}(\rho F^\perp) = 0$ for all $\rho \in \Xi$. The fact that

F^\perp does not have to vanish in order for \mathcal{F} to be an MIO demonstrates in a transparent way the (well-known) fact that there are MIOs that necessitate coherent resources for their implementation [184–186].

As emphasised throughout the main body of this paper, DØ states reduce to incoherent ones when the environment is trivial. Consequently, DØ maps are the natural extension of MIOs, and the proof of Lemma C.1 follows similar logic to the above proof for the structural properties of MIOs:

Proof. Employing the reasoning that led to Eq. (C.53), any DØ map Γ has a Choi state G of the form

$$G = \sum_{kj\mu\nu} g_{k\mu j\nu} \tau_k^\circ \otimes \Pi_j^i \otimes N_{\mu\nu}^{\circ i} + \sum_{rs\mu\nu} \tilde{g}_{r\mu s\nu} \tau_r^\circ \otimes \tilde{\Pi}_s^i \otimes N_{\mu\nu}^{\circ i}, \quad (\text{C.55})$$

where $g_{k\mu j\nu}, \tilde{g}_{r\mu s\nu} \in \mathbb{R}$ and $\{N_{\mu\nu}^{\circ i}\}_{\mu,\nu=1}^{d_e^2}$ is a basis of $\mathcal{B}(\mathcal{H}_e^\circ \otimes \mathcal{H}_e^i)$. As for the case of MIOs, Eq. (C.55) constitutes a decomposition $G = G^\parallel + G^\perp$, where G^\perp is orthogonal to the set of DØ states. Consequently, the action of Γ on any DØ state is entirely encapsulated in G^\parallel and it remains to show that this term is of the form given in the lemma. To this end, we note that a map Γ is DØ iff it maps any state of the form $\Pi_\ell^i \otimes \eta_\ell^i$ to a DØ state. Letting Γ act on such a product state, we obtain

$$\begin{aligned} \Gamma[\Pi_\ell^i \otimes \eta_\ell^i] &= \text{tr}_i \left\{ \left[\mathbb{1}^\circ \otimes (\Pi_\ell^i \otimes \eta_\ell^i)^\text{T} \right] G^\parallel \right\} \\ &= \sum_{k\mu\nu} g_{k\mu\ell\nu} \tau_k^\circ \otimes \text{tr}_i \left[\left(\mathbb{1}^\circ \otimes \eta_\ell^{i\text{T}} \right) N_{\mu\nu}^{\circ i} \right] \end{aligned} \quad (\text{C.56})$$

$$\stackrel{!}{=} \sum_r p_{r|\ell} \Pi_r^\circ \otimes \xi_{r|\ell}^\circ, \quad (\text{C.57})$$

where $\sum_r p_{r|\ell} = 1$ and $p_{r|\ell} \geq 0$, and $\xi_r^\circ \in \mathcal{B}(\mathcal{H}_e^\circ)$ are states of the environment. The last line of Eq. (C.57) stems from the requirement that Γ is a DØ map, and the remaining open index ℓ signifies that the resulting output state depends on the input state $\Pi_\ell^i \otimes \eta_\ell^i$. In the same way as above, we can choose the basis $\{\tau_k^\circ\}$ to consist of projectors $\{\Pi_k^\circ\}$ onto the computational basis and elements that are orthogonal to these projectors. Then, comparing Eqs. (C.56) and (C.57), we see that all of the terms of G^\parallel where τ_k° is not a projector onto the computational basis must vanish. Finally, the terms $N_{\mu\nu}^{\circ i}$ have to be such that $\text{tr}_i \left[\left(\mathbb{1} \otimes \eta_\ell^{i\text{T}} \right) \sum_{\mu\nu} g_{k\mu\ell\nu} N_{\mu\nu}^{\circ i} \right]$ yields the correct output state $p_{k|\ell} \xi_{k|\ell}^\circ$. Consequently, $\sum_{\mu\nu} g_{k\mu\ell\nu} N_{\mu\nu}^{\circ i}$ can be chosen to be (up to normalisation $p_{k|\ell}$) the Choi state $O_{k\ell}^{\circ i}$ of a CPTP map. Putting these observations together yields Eq. (C.50). \square

C.7 Proof that Non-Discord-Generating-and-Detecting (NDGD) Dynamics \Rightarrow Classical Process

For the proof of Theorem 3.4, we employ the fact that the completely dephasing map has no influence on the outcomes of a measurement in the computational basis, i.e.,

$$\mathcal{P}_{x_j} = \Delta_j \circ \mathcal{P}_{x_j} = \mathcal{P}_{x_j} \circ \Delta_j \quad \forall x_j. \quad (\text{C.58})$$

The probability $\mathbb{P}_k(x_k, \dots, x_1)$ to measure outcomes $\{x_k, \dots, x_1\}$ at times $\{t_k, \dots, t_1\}$ is given by [see Eq. (3.67)]

$$\text{tr}\{(\mathcal{P}_{x_k} \otimes \mathcal{I}^e) \circ \dots \circ \Gamma_{t_2, t_1} \circ (\mathcal{P}_{x_1} \otimes \mathcal{I}^e)[\eta_{t_1}^{se}]\}, \quad (\text{C.59})$$

where $\{\Gamma_{t_j, t_{j-1}}\}$ are system-environment CPTP maps and $\eta_{t_1}^{se}$ is the system-environment state at time t_1 . Summing this probability distribution over the outcomes at time t_j amounts to replacing \mathcal{P}_{x_j} in (C.59) by Δ_j . Zooming in on the relevant time (and leaving the \mathcal{I}^e implicit), we see that

$$\begin{aligned} & \mathcal{P}_{x_{j+1}} \circ \Gamma_{t_{j+1}, t_j} \circ \Delta_j \circ \Gamma_{t_j, t_{j-1}} \circ \mathcal{P}_{x_{j-1}} \\ &= \mathcal{P}_{x_{j+1}} \circ \Delta_{j+1} \circ \Gamma_{t_{j+1}, t_j} \circ \Delta_j \circ \Gamma_{t_j, t_{j-1}} \circ \Delta_{j-1} \circ \mathcal{P}_{x_{j-1}} \\ &= \mathcal{P}_{x_{j+1}} \circ \Gamma_{t_{j+1}, t_j} \circ \mathcal{I}_j \circ \Gamma_{t_j, t_{j-1}} \circ \mathcal{P}_{x_{j-1}}, \end{aligned} \quad (\text{C.60})$$

where we have used Eq. (C.58) in the first line, and both the fact that the dynamics is NDGD and Eq. (C.58) in the second line. As Eq. (C.60) holds for arbitrary times t_j , it implies that for NDGD dynamics, the completely dephasing map cannot be distinguished from the identity map when the process is probed by measurements in the computational basis, which implies that the Kolmogorov condition holds for any joint probabilities with at least 3 different times. For the 2-time joint probabilities, we can exploit, along with the NDGD property, the fact that the initial state is zero-discord. We have

$$\begin{aligned} \sum_{x_1} \mathbb{P}_2(x_2, x_1) &= \text{tr}\{\mathcal{P}_{x_2} \circ \Gamma_{t_2, t_1} \circ \Delta_1 \circ \Gamma_{t_1, t_0}[\eta_{t_0}^{se}]\} \\ &= \text{tr}\{\mathcal{P}_{x_2} \circ \Delta_2 \circ \Gamma_{t_2, t_1} \circ \Delta_1 \circ \Gamma_{t_1, t_0}[\eta_{t_0}^{se}]\} \\ &= \text{tr}\{\mathcal{P}_{x_2} \circ \Delta_2 \circ \Gamma_{t_2, t_1} \circ \Delta_1 \circ \Gamma_{t_1, t_0} \circ \Delta_0[\eta_{t_0}^{se}]\} \\ &= \text{tr}\{\mathcal{P}_{x_2} \circ \Gamma_{t_2, t_0}[\eta_{t_0}^{se}]\} = \mathbb{P}_1(x_2), \end{aligned} \quad (\text{C.61})$$

where we used Eq. (C.59) and $\sum_{x_1} \mathcal{P}_{x_1} = \Delta_1$ in the first line, Eq. (C.58) in the second line, the invariance of the initial zero-discord state with respect to Δ_0 in the third line, and finally the definition of NDGD dynamics, Eq. (C.58), and the invariance of $\eta_{t_0}^{se}$ in the fourth line. Consequently, the resulting statistics satisfy all of the Kolmogorov conditions and are thus classical.

C.8 Classicality \Rightarrow Non-Discord-Generating-and-Detecting (NDGD) Dynamics

Here, we provide an example of dynamics that are not NDGD, yet lead to classical dynamics, thus demonstrating that it is not necessary for a dynamics to be NDGD in order for it to appear classical. We consider the following situation (see Fig. C.4 for a graphical representation): Let the system of interest be a qubit that is initially in state $|0\rangle$ and let the initial environment be in a plus state, i.e., $\tau_{t_0}^e = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. The first evolution Γ_{t_1, t_0} from t_0 to t_1 is a Controlled-NOT (CNOT) gate, such that the system-environment state at t_1 is a maximally entangled state. The second evolution Γ_{t_2, t_1} from t_1 to t_2 is such that it yields a system-environment state $\frac{1}{2} \otimes |0\rangle\langle 0|$ if the se' input state is $\varphi_{se'}^+$, and $\frac{1}{2} \otimes |1\rangle\langle 1|$ otherwise. Consequently, when the completely dephasing map is applied at t_1 , the system-environment state at t_2 is $\frac{1}{2} \otimes \frac{1}{2}$, while it is equal to $\frac{1}{2} \otimes |0\rangle\langle 0|$ if the identity map was implemented, and as such, the dynamics is not NDGD. However, the system state is always maximally mixed, independent of whether Δ_1 or \mathcal{I}_1 was implemented at time t_1 . To make the example non-trivial, we add a third free dynamics Γ_{t_3, t_2} from t_2 to t_3 . We choose Γ_{t_3, t_2} such that it induces a unital dynamics on the level of the system, independent of the environment state at t_2 . This happens, e.g., when the corresponding system-environment Hamiltonian is of product form, i.e., $H_{se} = H_s \otimes H_e$, independent of the explicit form of the respective terms [286]. With this final dynamics, the system state at each of the times t_1, t_2 , and t_3 is maximally mixed, and the resulting statistics satisfy Kolmogorov conditions, i.e., they are classical.




Figure C.4: Non-NDGD Dynamics that leads to Classical Statistics. The first map Γ_{t_1, t_0} (blue transparent box) performs a CNOT gate on the system and the environment. The subsequent CPTP map Γ_{t_2, t_1} maps φ^+ and $\frac{1}{4}$ onto two different system-environment states with the same reduced system state $\rho_{t_3} = \frac{1}{2}$. The final CPTP map Γ_{t_3, t_2} is such that it induces a unital dynamics on the system. Consequently, the system state at t_1, t_2 , and t_3 is maximally mixed independent of whether the completely dephasing, or the identity map was implemented at t_1 and t_2 .

C.9 Example of a Genuinely Quantum Process

Consider the following process, depicted in Fig. C.5, which is a variation on that presented in Section 3.7. The process begins with a two-qubit system-environment state in the Bell pair φ_{se}^+ , the system part of which the experimenter has access to measure at t_1 . Following this, the process “performs” the CPTP system-environment map $\Gamma_z^{sez} : \mathcal{B}(\mathcal{H}^{s^i} \otimes \mathcal{H}^{e^i}) \rightarrow \mathcal{B}(\mathcal{H}^{s^o} \otimes \mathcal{H}^{e^o} \otimes \mathcal{H}^{z^o})$, whose action is as follows: It measures its joint inputs in the Bell basis and if the measurement outcome corresponds to φ_{se}^+ , it outputs a φ_{se}^+ system-environment state as well as a classical *flag* state $|0\rangle_z$; on the other hand, if the measurement outcome does not correspond to φ_{se}^+ , it outputs a system-environment state whose system part is a pure state in the z -basis, and sets the flag state to $|1\rangle_z$ to indicate that the system state has been biased in the z -basis. The action of the map is thus as follows:

$$\Gamma_z^{sez}[\eta_{se}] = \text{tr}[\varphi_{se}^+ \eta_{se}] \varphi_{se}^+ \otimes |0\rangle\langle 0|_z + \text{tr}[(\mathbb{1}_{se} - \varphi_{se}^+) \eta_{se}] |0\rangle\langle 0|_s \otimes \tau_e \otimes |1\rangle\langle 1|_z. \quad (\text{C.62})$$

For this map (and all that follow in this example), the output state of the environment when the φ_{se}^+ outcome is not recorded is irrelevant for our argument; as such, we simply write a generic quantum state τ_e .

Following this part of the dynamics, the experimenter has access to measure the system at time t_2 . The subsequent dynamics of the process is controlled on the state of the classical z flag: If it is in the state $|0\rangle_z$, the system-environment is subject to a similar dynamics as before, $\Gamma_y^{sey} : \mathcal{B}(\mathcal{H}^{s^i} \otimes \mathcal{H}^{e^i}) \rightarrow \mathcal{B}(\mathcal{H}^{s^o} \otimes \mathcal{H}^{e^o} \otimes \mathcal{H}^{y^o})$; however, this time if the Bell basis measurement outcome does not correspond to φ_{se}^+ , the system is biased in the y -basis, e.g., set to the -1 eigenstate of $\sigma^{(y)}$, $|-(y)\rangle := \frac{1}{\sqrt{2}}(|0\rangle - i|1\rangle)$, with a classical y flag set to the state $|1\rangle_y$ and sent forward. If, on the other hand, the z flag is in the state $|1\rangle_z$, the system-environment undergoes trivial dynamics (i.e., is subject to the identity map) and the y flag is set to $|0\rangle_y$. In either case, the previous z flag state is also sent forward unperturbed. Thus, between t_2 and t_3 , the system-environment evolves conditionally according to

$$\begin{aligned} z = 0 : \\ \Gamma_y^{sey}[\eta_{se}] &= \text{tr}[\varphi_{se}^+ \eta_{se}] \varphi_{se}^+ \otimes |0\rangle\langle 0|_y + \text{tr}[(\mathbb{1}_{se} - \varphi_{se}^+) \eta_{se}] |-(y)\rangle\langle -(y)|_s \otimes \tau_e \otimes |1\rangle\langle 1|_y \\ z = 1 : \\ \mathcal{I}^{se}[\eta_{se}] &\otimes |0\rangle\langle 0|_y. \end{aligned} \quad (\text{C.63})$$

Following this, the experimenter has access to the system at t_3 .

The final portion of the dynamics between t_3 and t_4 follows a similar construction to above, but the implementation of the map $\Gamma_x^{sex} : \mathcal{B}(\mathcal{H}^{s^i} \otimes \mathcal{H}^{e^i}) \rightarrow \mathcal{B}(\mathcal{H}^{s^o} \otimes \mathcal{H}^{e^o} \otimes \mathcal{H}^{x^o})$ is controlled on the joint state of the z and y classical flags. If $zy = 00$, the system-environment is measured in the Bell basis: If the measurement outcome does not correspond




Figure C.5: Genuinely Quantum Process. The system-environment begin in a maximally entangled Bell state φ_{se}^+ . As described in the text, the process dynamics consists of a sequence of maps, $\Gamma_z^{sez}, \Gamma_y^{sey}, \Gamma_x^{sex}$, that either output φ_{se}^+ or else bias the system in either the z -, y - or x -basis respectively [see Eqs. (C.62) – (C.64)]. The overall implementation of each of these maps is controlled on the joint state of all previous classical flag states z, y, x , which encode whether or not the system has already been biased. We show that this process is genuinely quantum by tracking the system-environment state throughout the dynamics, conditioned on whether the identity map \mathcal{I} or an arbitrary CPTP map \mathcal{M}_i was implemented at time t_i ; the labels $\eta_2^{sez}, \eta_3^{sezy}$ and η_4^{sezyz} refer to the overall joint state immediately prior to the interrogation at the relevant time [see Eqs. (C.65), (C.67), (C.69), (C.70), (C.72) and (C.73)]. In particular, we show that there does not exist a non-pathological Positive Operator-Valued Measure (POVM) Π_4 that an experimenter can implement at t_4 such that the four sequences $\{\mathcal{I}_1, \mathcal{I}_2, \mathcal{I}_3\}$, $\{\mathcal{M}_1, \mathcal{I}_2, \mathcal{I}_3\}$, $\{\mathcal{I}_1, \mathcal{M}_2, \mathcal{I}_3\}$ and $\{\mathcal{I}_1, \mathcal{I}_2, \mathcal{M}_3\}$ cannot be distinguished, thereby proving that the process is genuinely quantum.

to φ^+ , the system is biased in the x -basis, e.g., set to the -1 eigenstate of $\sigma^{(x)}$, $|-(x)\rangle := \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$, with a classical x flag set to the state $|1\rangle_x$ and sent forward. If, on the other hand, $zy \neq 00$, the system-environment undergoes trivial dynamics (i.e., is subject to the identity map) and the x flag is set to $|0\rangle_x$. Mathematically, the controlled dynamics is described as

$$\begin{aligned}
 & \underline{zy = 00} : \\
 & \Gamma_x^{sex}[\eta_{se}] = \text{tr}[\varphi_{se}^+ \eta_{se}] \varphi_{se}^+ \otimes |0\rangle\langle 0|_x + \text{tr}[(\mathbb{1}_{se} - \varphi_{se}^+) \eta_{se}] |-(x)\rangle\langle -(x)|_s \otimes \tau_e \otimes |1\rangle\langle 1|_x \\
 & \underline{zy = 10, 01} : \\
 & \mathcal{I}^{se}[\eta_{se}] \otimes |0\rangle\langle 0|_x.
 \end{aligned} \tag{C.64}$$

Note that the flag state $zy = 11$ cannot occur. Finally, the environment and all flag states are discarded and the experimenter has access to the system at t_4 , concluding the process.

We now show that there exist no unrestricted measurement scheme for this process such that the statistics observed are classical, i.e., we prove that the process is genuinely quantum. As in the main text, we do this by considering the state of the system to be measured at the final time t_4 conditioned on a history of identity maps and arbitrary CPTP maps $\{\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3\}$ implemented at various sets of earlier times. In each case, by demanding classicality we end up with a different constraint on the structure of the POVM at the final time, and the only valid POVMs that simultaneously satisfy all conditions are the pathological ones that do not reveal anything about the process. The conclusion is that any non-pathological POVM at t_4 will be able to distinguish between previous implementations of the identity map or an arbitrary non-pathological instrument at a

given time, therefore picking up on the invasiveness of (at least some of) the previous interrogations and leading to non-classical statistics.

Consider first the scenario where the experimenter implements identity maps at the first three times, $\mathcal{I}_1, \mathcal{I}_2, \mathcal{I}_3$. In this case, the overall state immediately prior to the measurement at t_4 is

$$\eta_4^{sezyx}(\mathcal{I}_1, \mathcal{I}_2, \mathcal{I}_3) = \varphi_{se}^+ \otimes |000\rangle\langle 000|_{zyx}. \quad (\text{C.65})$$

The reduced system state is then maximally mixed:

$$\eta_4^s(\mathcal{I}_1, \mathcal{I}_2, \mathcal{I}_3) = \frac{\mathbb{1}}{2}. \quad (\text{C.66})$$

Next, consider the case where the experimenter implements the identity map at the first two times, $\mathcal{I}_1, \mathcal{I}_2$, followed by an arbitrary CPTP map $\mathcal{M}_3 \neq \mathcal{I}_3$ at t_3 . The system-environment joint state immediately prior to t_3 is φ_{se}^+ , since the previous identity maps on the system and the dynamics $\Gamma_z^{sez}, \Gamma_y^{sey}$ leading up to t_3 preserve the initial state; moreover, the zy flag is in the joint state 00 , since both previous Bell basis measurements are necessarily successful. Now, the system-local CPTP map $\mathcal{M}_3 \neq \mathcal{I}_3$ will perturb the joint system-environment state, and so the map Γ_x^{sex} (which is implemented due to the joint state of the input flags) only successfully records the outcome corresponding to φ_{se}^+ with some probability $r = \text{tr}[\varphi_{se}^+(\mathcal{M}_3^s \otimes \mathcal{I}^e)[\varphi_{se}^+]] < 1$; otherwise, the system is biased in the x -basis. The total joint state immediately prior to t_4 in this scenario is then

$$\eta_4^{sezyx}(\mathcal{I}_1, \mathcal{I}_2, \mathcal{M}_3) = r\varphi_{se}^+ \otimes |000\rangle\langle 000|_{zyx} + (1-r)|-(x)\rangle\langle -(x)|_s \otimes \tau_e \otimes |001\rangle\langle 001|_{zyx}. \quad (\text{C.67})$$

The reduced system state is thus biased in the x -basis:

$$\eta_4^s(\mathcal{I}_1, \mathcal{I}_2, \mathcal{M}_3) = \frac{r}{2}\mathbb{1} + (1-r)|-(x)\rangle\langle -(x)|. \quad (\text{C.68})$$

Next, consider the case where the experimenter implements the identity map at the first and third time, $\mathcal{I}_1, \mathcal{I}_3$, with an arbitrary CPTP map $\mathcal{M}_2 \neq \mathcal{I}_2$ implemented in between at time t_2 . The system-environment joint state immediately prior to t_2 is φ_{se}^+ , since the previous identity map on the system and the dynamics Γ_z^{sez} prior to t_2 again preserve the initial state; moreover, the z flag is in the state 0 , since the earlier Bell basis measurement is necessarily successful. Again, the system-local CPTP map $\mathcal{M}_2 \neq \mathcal{I}_2$ will perturb the joint system-environment state, and so the map Γ_y^{sey} (which is implemented due to the state of the input flag) only successfully records the outcome corresponding to φ_{se}^+ with some probability $q = \text{tr}[\varphi_{se}^+(\mathcal{M}_2^s \otimes \mathcal{I}^e)[\varphi_{se}^+]] < 1$; otherwise, the system is biased in the y -basis. The total joint state immediately prior to t_3 in this scenario is then

$$\eta_3^{sezy}(\mathcal{I}_1, \mathcal{M}_2) = q\varphi_{se}^+ \otimes |00\rangle\langle 00|_{zy} + (1-q)|-(y)\rangle\langle -(y)|_s \otimes \tau_e \otimes |01\rangle\langle 01|_{zy}. \quad (\text{C.69})$$

In this case, the experimenter then implements the identity map to the system at t_3 , which leaves the overall state invariant. The subsequent system-environment dynamics

Γ_x^{sex} will be enacted when $zy = 00$, i.e., with probability q : In each such run, the system-environment state is guaranteed to be in the state φ_{se}^+ , thus the system-environment state output by Γ_x^{sex} will be also. In the other cases, when $zy \neq 00$, the subsequent dynamics will be trivial. Thus, the total joint state immediately prior to t_4 in this scenario is

$$\eta_4^{seyx}(\mathcal{I}_1, \mathcal{M}_2, \mathcal{I}_3) = q\varphi_{se}^+ \otimes |00\rangle\langle 00|_{zyx} + (1-q)|-\langle -^{(y)}|_s \otimes \tau_e \otimes |010\rangle\langle 010|_{zyx}. \quad (\text{C.70})$$

The final reduced system state is thus biased in the y -basis:

$$\eta_4^s(\mathcal{I}_1, \mathcal{M}_2, \mathcal{I}_3) = \frac{q}{2}\mathbb{1} + (1-q)|-\langle -^{(y)}|. \quad (\text{C.71})$$

Lastly, consider the scenario where the experimenter first implements an arbitrary CPTP map $\mathcal{M}_1 \neq \mathcal{I}_1$ at t_1 , followed by identity maps at the second and third time, $\mathcal{I}_2, \mathcal{I}_3$. Just as in the main text, $\mathcal{M}_1 \neq \mathcal{I}_1$ will perturb the initial system-environment state and so the map Γ_z^{sez} will only successfully record the outcome corresponding to φ_{se}^+ with some probability $p = \text{tr}[\varphi_{se}^+(\mathcal{M}_1^s \otimes \mathcal{I}^e)[\varphi_{se}^+]] < 1$; otherwise, the system will be biased in the z -basis. The total joint state immediately prior to t_2 in this scenario is then

$$\eta_2^{sez}(\mathcal{M}_1) = p\varphi_{se}^+ \otimes |0\rangle\langle 0|_z + (1-p)|0\rangle\langle 0|_s \otimes \tau_e \otimes |1\rangle\langle 1|_z. \quad (\text{C.72})$$

The identity map implemented by the experimenter on the system at t_2 does not change this state. Thus, Γ_y^{sey} will subsequently be enacted with probability p , i.e., when $z = 0$: In such cases, the system-environment is in the state φ_{se}^+ and the output of the map Γ_y^{sey} will be so also, accompanied by the classical y flag with the value 0. In the other cases, the system-environment undergoes trivial dynamics. Again, at t_3 implementation of the identity map on the system leaves the joint state unperturbed. Only when the joint state of zy is 00 will the map Γ_x^{sex} be implemented: In each such run, the system-environment is guaranteed to be in the state φ_{se}^+ , and thus so too will be the output of the map. In the other cases, trivial dynamics ensues. The overall joint state in this scenario immediately prior to t_4 is then

$$\eta_4^{seyx}(\mathcal{M}_1, \mathcal{I}_2, \mathcal{I}_3) = p\varphi_{se}^+ \otimes |000\rangle\langle 000|_{zyx} + (1-p)|0\rangle\langle 0|_s \otimes \tau_e \otimes |100\rangle\langle 100|_{zyx}, \quad (\text{C.73})$$

and so the reduced system state is biased in the z -basis:

$$\eta_4^s(\mathcal{M}_1, \mathcal{I}_2, \mathcal{I}_3) = \frac{p}{2}\mathbb{1} + (1-p)|0\rangle\langle 0|. \quad (\text{C.74})$$

We are now in a position to prove the claim that we set out to, namely that the process considered is genuinely quantum. Demanding classicality means that the experimenter cannot distinguish whether an identity map or a dephasing map was implemented at any subset of previous times: To allow for arbitrary and possibly unrestricted interrogation schemes, here we have considered the more general case where the experimenter is allowed to implement arbitrary CPTP maps, of which any POVM measurement followed by an

arbitrary preparation is a special case. This more general notion of classicality (with respect to a general, possibly unrestricted, interrogation scheme) means that the experimenter cannot distinguish between the implementation of the identity map or the CPTP map at any subset of previous times and thereby provides a valid notion of a genuinely quantum process. Above, in Eqs. (C.68), (C.71) and (C.74), we have calculated the system state that would be measured at t_4 conditioned on the fact that a CPTP map was implemented at each one of the previous three times [as well as the case where only a sequence of identity maps was implemented in Eq. (C.66)]. Intuitively, in each of the three cases, the system is biased in one of the x -, y - or z -basis directions, and in the case where the experimenter interacts only trivially with the system, it is completely unbiased. The only way that a measurement at t_4 cannot distinguish between these four scenarios is if it is blind to biases in every basis. The only types of POVM that can achieve this are trivial, with all elements proportional to the identity matrix, $\{\Pi_4^{(x_4)}\} \propto \mathbb{1} \forall x_4$. Thus, there is no (non-trivial) measurement scheme for this process such that the full statistics appears classical, and thus it is a genuinely quantum process.



Supplemental Information for Chapter 5

D.1 Classical Processes

D.1.1 Memoryless Classical Processes and Markovian Statistics

Here, we prove Observation 1 of the main text:

Observation 1. *In the classical setting, memoryless processes are equivalent to Markovian statistics.*

Naturally, this equivalence is well-known, but its explicit discussion exposes many of the subtleties with respect to marginalisation that play a crucial role in the quantum case. For the proof, in the forwards direction, beginning with Eq. (5.2), we have

$$\mathbb{P}(x_n, \dots, x_1) = \langle x_n | S_{n:n-1} | x_{n-1} \rangle \langle x_{n-1} | \dots | x_2 \rangle \langle x_2 | S_{2:1} | x_1 \rangle \langle x_1 | \mathbf{p}_1, \quad (\text{D.1})$$

where $\{S_{j:j-1}\}$ and \mathbf{p}_1 are, respectively, the stochastic matrices and the initial probability vector that define the memoryless classical process. Computing the conditional probability for an arbitrary time t_j given the entire sequence of historic outcomes up until that time explicitly gives

$$\begin{aligned} \mathbb{P}(x_j | x_{j-1}, \dots, x_1) &= \frac{\langle x_j | S_{j:j-1} | x_{j-1} \rangle \langle x_{j-1} | S_{j-1:j-2} | x_{j-2} \rangle \langle x_{j-2} | \dots | x_2 \rangle \langle x_2 | S_{2:1} | x_1 \rangle \langle x_1 | \mathbf{p}_1}{\langle x_{j-1} | S_{j-1:j-2} | x_{j-2} \rangle \langle x_{j-2} | \dots | x_2 \rangle \langle x_2 | S_{2:1} | x_1 \rangle \langle x_1 | \mathbf{p}_1} \\ &= \langle x_j | S_{j:j-1} | x_{j-1} \rangle \quad \forall x_{j-2}, \dots, x_1. \end{aligned} \quad (\text{D.2})$$

This expression is independent of all x_1, \dots, x_{j-2} and it is indeed equivalent to the conditional probability $\mathbb{P}(x_j | x_{j-1})$ an experimenter would observe when only making measurements at t_{j-1} and t_j , i.e., they do not measure (which we denote below by

$\mathcal{I}_{j-2:1}$). Unlike in quantum mechanics, this conditional probability can be expressed by marginalising the full joint probability distribution [see Eq. (5.4)] as follows

$$\begin{aligned}
 \mathbb{P}(x_j|x_{j-1}) &= \frac{\mathbb{P}(x_j, x_{j-1}, \mathcal{I}_{j-2:1})}{\mathbb{P}(x_{j-1}, \mathcal{I}_{j-2:1})} \\
 &= \frac{\sum_{x_{j-2}, \dots, x_1} \mathbb{P}(x_j, x_{j-1}, x_{j-2}, \dots, x_1)}{\sum_{x_{j-2}, \dots, x_1} \mathbb{P}(x_{j-1}, x_{j-2}, \dots, x_1)} \\
 &= \frac{\sum_{x_{j-2}, \dots, x_1} \langle x_j | S_{j:j-1} | x_{j-1} \rangle \langle x_{j-1} | S_{j-1:j-2} | x_{j-2} \rangle \langle x_{j-2} | \dots | x_2 \rangle \langle x_2 | S_{2:1} | x_1 \rangle \langle x_1 | p_1 | x_0 \rangle}{\sum_{x_{j-2}, \dots, x_1} \langle x_{j-1} | S_{j-1:j-2} | x_{j-2} \rangle \langle x_{j-2} | \dots | x_2 \rangle \langle x_2 | S_{2:1} | x_1 \rangle \langle x_1 | p_1 | x_0 \rangle} \\
 &= \langle x_j | S_{j:j-1} | x_{j-1} \rangle. \tag{D.3}
 \end{aligned}$$

Thus we have that for a memoryless process, the conditional probabilities $\mathbb{P}(x_j|x_{j-1}, \dots, x_1)$ and $\mathbb{P}(x_j|x_{j-1})$ coincide (and amount to $\langle x_j | S_{j:j-1} | x_{j-1} \rangle$), leading to Markovianity of the statistics and consequently the decomposition of the joint probability distribution expressed in Eq. (5.3).

Conversely, any Markovian statistics can be faithfully reproduced via a memoryless classical model: Given any joint probability distribution over measurement outcomes for a classical stochastic process, one can always write

$$\mathbb{P}(x_n, \dots, x_1) = \mathbb{P}(x_n|x_{n-1}, \dots, x_1) \mathbb{P}(x_{n-1}|x_{n-2}, \dots, x_1) \dots \mathbb{P}(x_2|x_1) \mathbb{P}(x_1). \tag{D.4}$$

Equation (D.4) holds true for any probability distribution (by definition of conditional probabilities), with the decomposition on the r.h.s. encoding potential memory effects. For Markovian statistics, the above expression simplifies to Eq. (5.3). Then, one can simply define a set of matrices $\{S_{j:j-1}\}$ via

$$\langle x_j | S_{j:j-1} | x_{j-1} \rangle := \mathbb{P}(x_j|x_{j-1}). \tag{D.5}$$

These matrices are stochastic (as they contain only non-negative entries and each of the columns to unity since $\sum_{x_j} \mathbb{P}(x_j|x_{j-1}) = 1 \forall x_{j-1}$). One can also define the initial state via $\langle x_1 | p_1 := \mathbb{P}(x_1)$. From these objects, one can reproduce the joint statistics faithfully via the memoryless dynamical model expressed in Eq. (5.2).

D.1.2 Sub-Statistics of Memoryless Classical Processes

Here, we prove Corollary 5.1 of the main text:

Corollary 5.1. *All sub-statistics of a memoryless classical process are Markovian and the corresponding conditional probabilities are mutually compatible.*

For the proof, consider a memoryless process on $\mathcal{T}_n = \{t_1, \dots, t_n\}$ and an arbitrary sub-statistics where the experimenter measures at time t_j and any subset of earlier times $\Gamma^{(i)}$ (with corresponding sequence of outcomes $\mathbf{x}_{\Gamma^{(i)}}$), where $t_j > t_i = \max(\Gamma^{(i)})$. Below,

we both assume that $\Gamma^{(i)}$ does not “skip times” (e.g., it can be of the form $\{t_3, t_2, t_1\}$, but not $\{t_3, t_1\}$) and that $\min(\Gamma^{(i)}) = t_1$. These assumptions are not crucial and do not affect the generality of the results, but significantly simplifies notation. We denote by M the set of all times between t_i and t_j and by F that of all times after t_j , with corresponding outcome sequences \mathbf{x}_M and \mathbf{x}_F and do-nothing operations \mathcal{I}_M and \mathcal{I}_F , respectively. For convenience, we also introduce the do-nothing operation \mathcal{I}_{FjM} for all times after t_i . With this, we explicitly calculate the probability of observing x_j conditioned on the sequence of previous outcomes $\mathbf{x}_{\Gamma^{(i)}}$ as

$$\begin{aligned}
 & \mathbb{P}(x_j | \mathbf{x}_{\Gamma^{(i)}}) \\
 &= \frac{\mathbb{P}(\mathcal{I}_F, x_j, \mathcal{I}_M, x_i, \dots, x_1)}{\mathbb{P}(\mathcal{I}_{FjM}, x_i, \dots, x_1)} \\
 &= \frac{\sum_{\mathbf{x}_F \mathbf{x}_M} \mathbb{P}(x_n, \dots, x_1)}{\sum_{\mathbf{x}_F x_j \mathbf{x}_M} \mathbb{P}(x_n, \dots, x_1)} \\
 &= \frac{\sum_{\mathbf{x}_F \mathbf{x}_M} \mathbb{P}(x_n | x_{n-1}) \dots \mathbb{P}(x_2 | x_1)}{\sum_{\mathbf{x}_F x_j \mathbf{x}_M} \mathbb{P}(x_n | x_{n-1}) \dots \mathbb{P}(x_2 | x_1)} \\
 &= \frac{\left[\sum_{\mathbf{x}_F} \mathbb{P}(x_n | x_{n-1}) \dots \mathbb{P}(x_{j+1} | x_j) \right] \left[\sum_{\mathbf{x}_M} \mathbb{P}(x_j | x_{j-1}) \dots \mathbb{P}(x_{i+1} | x_i) \right] \{ \mathbb{P}(x_i | x_{i-1}) \dots \mathbb{P}(x_2 | x_1) \}}{\left[\sum_{\mathbf{x}_F x_j \mathbf{x}_M} \mathbb{P}(x_n | x_{n-1}) \dots \mathbb{P}(x_{i+1} | x_i) \right] \{ \mathbb{P}(x_i | x_{i-1}) \dots \mathbb{P}(x_2 | x_1) \}} \\
 &= \sum_{\mathbf{x}_M} \mathbb{P}(x_j | x_{j-1}) \dots \mathbb{P}(x_{i+1} | x_i) =: \mathbb{P}(x_j | x_i), \tag{D.6}
 \end{aligned}$$

where in the second line we employed the marginalisation rule to compute the sub-statistics from the full process on \mathcal{T}_n , in the third line we invoked the Markovianity condition (on the full statistics), in the fourth line we split the sums into independent parts, in the fifth line we used the fact that the first sum in the numerator and the sum in the denominator both evaluate to unity, and the final line only depends on x_j and x_i and satisfies the properties of a conditional probability distribution. Thus we see that any sub-statistics of a memoryless classical stochastic process are also Markovian, i.e., $\mathbb{P}(x_j | \mathbf{x}_{\Gamma^{(i)}}) = \mathbb{P}(x_j | x_i)$ for all $t_j > t_i$. As mentioned, this reasoning also holds for more “complicated” sets $\Gamma^{(i)}$, albeit with a slightly more cumbersome notation than used in the proof above.

Moreover, the Markovian sub-statistics are compatible in the sense that it does not matter what occurred at any time prior to that of the most recent conditioning argument, i.e., t_i . For instance, if one computes $\mathbb{P}(x_j | x_i, \mathcal{I}_{i-1:\ell+1}, x_\ell, \mathcal{I}_{\ell-1:1})$, this should also be independent of x_ℓ (i.e., Markovian sub-statistics) *and* equal to $\mathbb{P}(x_j | \mathbf{x}_{\Gamma^{(i)}}) = \mathbb{P}(x_j | x_i)$ computed above (i.e., compatible). This can be seen by noting that for any historic sequence (of either measuring or not at any times t_1, \dots, t_{i-1} , which we denote with $(x \cup \mathcal{I})_{i-1:1}$), the logic of Eq. (D.6) holds, since the only changes would appear in the terms in curly parentheses in the fourth line, which always cancel. Hence, we have the compatibility $\mathbb{P}(x_j | x_i, (x \cup \mathcal{I})_{i-1:1}) = \mathbb{P}(x_j | x_i)$ for all possible combinations of measuring or not in the history leading up to time t_i . Again, this argument can be run in exactly the

same vein for any two subsets of times $\Gamma^{(i)}$ and $\Gamma^{(i)'}$ satisfying $\max(\Gamma^{(i)}) = \max(\Gamma^{(i)'}) = t_i$, with the result that $\mathbb{P}(x_j | \mathbf{x}_{\Gamma^{(i)}}) = \mathbb{P}(x_j | \mathbf{x}_{\Gamma^{(i)'}}) = \mathbb{P}(x_j | x_i)$ for all $t_j > t_i$.

D.2 Quantum Processes

D.2.1 Memoryless Quantum Processes and Markovian Statistics

Here, we prove Lemma 5.1 of the main text:

Lemma 5.1. *Any memoryless quantum process leads to Markovian statistics (for sharp, projective measurements).*

Beginning with Eq. (5.6), we have that for any memoryless quantum process

$$\mathbb{P}(x_n, \dots, x_1) = \text{tr} \left[\mathcal{P}_n^{(x_n)} \Lambda_{n:n-1} \dots \Lambda_{2:1} \mathcal{P}_1^{(x_1)} \rho_1 \right], \quad (\text{D.7})$$

where $\{\Lambda_{j:j-1}\}$ are mutually independent Completely Positive and Trace Preserving (CPTP) maps, ρ_1 is an initial quantum state, and $\mathcal{P}_j^{(x_j)}[\bullet] = |x_j\rangle\langle x_j| \bullet |x_j\rangle\langle x_j|$. The statistics up to t_j is given by $\mathbb{P}(\mathcal{I}_{n:j+1}, x_j, \dots, x_1) =: \mathbb{P}(x_j, \dots, x_1) = \text{tr} \left[\mathcal{P}_j^{(x_j)} \Lambda_{j:j-1} \dots \Lambda_{2:1} \mathcal{P}_1^{(x_1)} \rho_1 \right]$ (this can be seen either by direct computation or by invoking causality), where $\mathcal{I}_{n:j+1}$ denotes “do-nothing” operations from t_j to t_n . With this, computing the conditional probability for an arbitrary time t_j given the entire sequence of historic outcomes up until that time explicitly gives

$$\begin{aligned} & \mathbb{P}(x_j | x_{j-1}, \dots, x_1) \\ &= \frac{\text{tr} \left[\mathcal{P}_j^{(x_j)} \Lambda_{j:j-1} \dots \Lambda_{2:1} \mathcal{P}_1^{(x_1)} \rho_1 \right]}{\text{tr} \left[\mathcal{P}_{j-1}^{(x_{j-1})} \Lambda_{j-1:j-2} \dots \Lambda_{2:1} \mathcal{P}_1^{(x_1)} \rho_1 \right]} \\ &= \frac{\text{tr} \left[\sum_{\ell} |x_j\rangle\langle x_j| L_{j:j-1}^{\ell} |x_{j-1}\rangle\langle x_{j-1}| \left(\Lambda_{j-1:j-2} \dots \Lambda_{2:1} \mathcal{P}_1^{(x_1)} \rho_1 \right) |x_{j-1}\rangle\langle x_{j-1}| L_{j:j-1}^{\ell\dagger} |x_j\rangle\langle x_j| \right]}{\text{tr} \left[|x_{j-1}\rangle\langle x_{j-1}| \left(\Lambda_{j-1:j-2} \dots \Lambda_{2:1} \mathcal{P}_1^{(x_1)} \rho_1 \right) |x_{j-1}\rangle\langle x_{j-1}| \right]} \\ &= \frac{\sum_{\ell} \langle x_j | L_{j:j-1}^{\ell} |x_{j-1}\rangle\langle x_{j-1}| L_{j:j-1}^{\ell\dagger} |x_j\rangle\langle x_{j-1}| \left(\Lambda_{j-1:j-2} \dots \Lambda_{2:1} \mathcal{P}_1^{(x_1)} \rho_1 \right) |x_{j-1}\rangle}{\langle x_{j-1} | \left(\Lambda_{j-1:j-2} \dots \Lambda_{2:1} \mathcal{P}_1^{(x_1)} \rho_1 \right) |x_{j-1}\rangle} \\ &= \sum_{\ell} \langle x_j | L_{j:j-1}^{\ell} |x_{j-1}\rangle\langle x_{j-1}| L_{j:j-1}^{\ell\dagger} |x_j\rangle \\ &= \langle x_j | \Lambda_{j:j-1} [|x_{j-1}\rangle\langle x_{j-1}|] |x_j\rangle, \end{aligned} \quad (\text{D.8})$$

where we wrote $\Lambda_{j:j-1}[\bullet] := \sum_{\ell} L_{j:j-1}^{\ell} \bullet L_{j:j-1}^{\ell\dagger}$ in Kraus operator form in the second line, and then made use of the cyclicity of the trace and the fact that the measurements are sharp (rank-1) projectors in the third line (importantly, if the projectors are not rank-1, corresponding, e.g., to the measurement of an observable with degeneracies, then memoryless processes do *not* necessarily lead to Markovian statistics [106, 131], since in

this case the state after the measurement is not fully determined by the outcome; this fact is also true in the classical setting). This expression is independent of all x_1, \dots, x_{j-2} and therefore the conditional probabilities are Markovian. We now show that it is indeed equivalent to the conditional probability $\mathbb{P}(x_j|x_{j-1}, \mathcal{I}_{j-2:1}) =: \mathbb{P}(x_j|x_{j-1})$, where the experimenter does *not* measure at all on times t_1, \dots, t_{j-2} . Explicitly, we have

$$\begin{aligned}
 & \mathbb{P}(x_j|x_{j-1}) \\
 &= \frac{\text{tr} \left[\mathcal{P}_j^{(x_j)} \Lambda_{j:j-1} \mathcal{P}_{j-1}^{(x_{j-1})} \Lambda_{j-1:j-2} \mathcal{I}_{j-2} \dots \Lambda_{2:1} \mathcal{I}_1 \rho_1 \right]}{\text{tr} \left[\mathcal{P}_{j-1}^{(x_{j-1})} \Lambda_{j-1:j-2} \mathcal{I}_{j-2} \dots \Lambda_{2:1} \mathcal{I}_1 \rho_1 \right]} \\
 &= \frac{\text{tr} \left[\sum_{\ell} |x_j\rangle\langle x_j| L_{j:j-1}^{\ell} |x_{j-1}\rangle\langle x_{j-1}| (\Lambda_{j-1:j-2} \mathcal{I}_{j-2} \dots \Lambda_{2:1} \mathcal{I}_1 \rho_1) |x_{j-1}\rangle\langle x_{j-1}| L_{j:j-1}^{\ell\dagger} |x_j\rangle\langle x_j| \right]}{\text{tr} \left[|x_{j-1}\rangle\langle x_{j-1}| (\Lambda_{j-1:j-2} \mathcal{I}_{j-2} \dots \Lambda_{2:1} \mathcal{I}_1 \rho_1) |x_{j-1}\rangle\langle x_{j-1}| \right]} \\
 &= \frac{\sum_{\ell} \langle x_j | L_{j:j-1}^{\ell} |x_{j-1}\rangle\langle x_{j-1} | L_{j:j-1}^{\ell\dagger} |x_j\rangle\langle x_{j-1} | (\Lambda_{j-1:j-2} \mathcal{I}_{j-2} \dots \Lambda_{2:1} \mathcal{I}_1 \rho_1) |x_{j-1}\rangle}{\langle x_{j-1} | (\Lambda_{j-1:j-2} \mathcal{I}_{j-2} \dots \Lambda_{2:1} \mathcal{I}_1 \rho_1) |x_{j-1}\rangle} \\
 &= \sum_{\ell} \langle x_j | L_{j:j-1}^{\ell} |x_{j-1}\rangle\langle x_{j-1} | L_{j:j-1}^{\ell\dagger} |x_j\rangle \\
 &= \langle x_j | \Lambda_{j:j-1} [|x_{j-1}\rangle\langle x_{j-1}|] |x_j\rangle. \tag{D.9}
 \end{aligned}$$

Thus, we see that the conditional statistics in both situations above coincide and are indeed Markovian $\mathbb{P}(x_j|x_{j-1}, \dots, x_1) = \mathbb{P}(x_j|x_{j-1}, \mathcal{I}_{j-2:1}) = \mathbb{P}(x_j|x_{j-1}) = \langle x_j | \Lambda_{j:j-1} [|x_{j-1}\rangle\langle x_{j-1}|] |x_j\rangle$.

D.2.2 Sub-Statistics of Memoryless Quantum Processes

Here, we prove Lemma 5.2 from the main text:

Lemma 5.2. *Any memoryless quantum process leads to Markovian sub-statistics (for sharp, projective measurements) that are mutually compatible.*

Similar to Appendix D.1.2, we will restrict the discussion again to subsets of \mathcal{T}_n of the form $\Gamma^{(i)} = \{t_1, \dots, t_i\}$ and show that $\mathbb{P}(x_j|\mathbf{x}_{\Gamma^{(i)}}) = \mathbb{P}(x_j|x_i)$ holds for all $t_j > t_i$ and $t_i = \max(\Gamma^{(i)})$. Let $\mathcal{I}_{j-1:i+1}$ denote the “do-nothing” operation at all times between t_i and t_j . With this, we obtain

$$\begin{aligned}
 & \mathbb{P}(x_j|\mathcal{I}_{j-1:i+1}, \mathbf{x}_{\Gamma^{(i)}}) \\
 &= \frac{\text{tr} \left[\mathcal{P}_j^{(x_j)} \Lambda_{j:j-1} \mathcal{I}_{j-1} \Lambda_{j-1:j-2} \dots \mathcal{I}_{i+1} \Lambda_{i+1:i} \mathcal{P}_i^{(x_i)} \Lambda_{i:i-1} \mathcal{P}_{i-1}^{(x_{i-1})} \dots \mathcal{P}_1^{(x_1)} \rho_1 \right]}{\text{tr} \left[\mathcal{P}_i^{(x_i)} \Lambda_{i:i-1} \mathcal{P}_{i-1}^{(x_{i-1})} \dots \mathcal{P}_1^{(x_1)} \rho_1 \right]} \\
 &= \frac{\langle x_j | \Lambda_{j:j-1} \mathcal{I}_{j-1} \Lambda_{j-1:j-2} \dots \mathcal{I}_{i+1} \Lambda_{i+1:i} [|x_i\rangle\langle x_i|] |x_j\rangle \langle x_i | \Lambda_{i:i-1} \mathcal{P}_{i-1}^{(x_{i-1})} \dots \mathcal{P}_1^{(x_1)} \rho_1 |x_i\rangle}{\langle x_i | \Lambda_{i:i-1} \mathcal{P}_{i-1}^{(x_{i-1})} \dots \mathcal{P}_1^{(x_1)} \rho_1 |x_i\rangle} \\
 &= \langle x_j | \Lambda_{j:j-1} \mathcal{I}_{j-1} \Lambda_{j-1:j-2} \dots \mathcal{I}_{i+1} \Lambda_{i+1:i} [|x_i\rangle\langle x_i|] |x_j\rangle \quad \forall x_{i-1}, \dots, x_1. \tag{D.10}
 \end{aligned}$$

In the case where *no* measurements are made until time t_i , we similarly have

$$\begin{aligned}
 & \mathbb{P}(x_j | \mathcal{I}_{j-1:i+1}, x_i, \mathcal{I}_{i-1:1}) \\
 &= \frac{\text{tr} \left[\mathcal{P}_j^{(x_j)} \Lambda_{j:j-1} \mathcal{I}_{j-1} \Lambda_{j-1:j-2} \dots \mathcal{I}_{i+1} \Lambda_{i+1:i} \mathcal{P}_i^{(x_i)} \Lambda_{i:i-1} \mathcal{I}_{i-1} \dots \mathcal{I}_1 \rho_1 \right]}{\text{tr} \left[\mathcal{P}_k^{(x_i)} \Lambda_{i:i-1} \mathcal{I}_{i-1} \dots \mathcal{I}_1 \rho_1 \right]} \\
 &= \frac{\langle x_j | \Lambda_{j:j-1} \mathcal{I}_{j-1} \Lambda_{j-1:j-2} \dots \mathcal{I}_{i+1} \Lambda_{i+1:i} [|x_i\rangle\langle x_i|] |x_j\rangle \langle x_i | \Lambda_{i:i-1} \mathcal{I}_{i-1} \dots \mathcal{I}_1 \rho_1 |x_i\rangle}{\langle x_i | \Lambda_{i:i-1} \mathcal{I}_{i-1} \dots \mathcal{I}_1 \rho_1 |x_i\rangle} \\
 &= \langle x_j | \Lambda_{j:j-1} \mathcal{I}_{j-1} \Lambda_{j-1:j-2} \dots \mathcal{I}_{i+1} \Lambda_{i+1:i} [|x_i\rangle\langle x_i|] |x_j\rangle. \tag{D.11}
 \end{aligned}$$

Thus, we see that both conditional probabilities are equal and independent of all measurement outcomes prior to t_i i.e., we have

$$\mathbb{P}(x_j | \mathcal{I}_{j-1:i+1}, x_i, x_{i-1}, \dots, x_1) = \mathbb{P}(x_j | \mathcal{I}_{j-1:i+1}, x_i, \mathcal{I}_{i-1:1}) =: \mathbb{P}(x_j | x_i) \tag{D.12}$$

and the sub-statistics are indeed Markovian. Regarding consistency, note that for any combination of measuring or not in the times prior to t_i , the only changes to the above expressions occur in the numerator term that always cancels with the corresponding part in the denominator, and so compatibility holds true, i.e., $\mathbb{P}(x_j | \mathcal{I}_{j-1:i+1}, x_i, (x \cup \mathcal{I})_{i-1:1}) = \mathbb{P}(x_j | x_i)$ for all possible choices of $(x \cup \mathcal{I})_{i-1:1}$, i.e., all possible choices of measuring or not at times $\{t_1, \dots, t_{i-1}\}$ in the history. As for the classical case we demonstrated in Appendix D.1.2, the argument above can be run in exactly the same way for more “complicated” subsets $\Gamma^{(i)} \subset \mathcal{T}_n$, with the only difference that the notation becomes slightly more cumbersome.

D.3 Hidden Quantum Memory and Incompatibility

D.3.1 Hidden Quantum Memory

Here we explicitly calculate all sub-statistics of the example used in Observation 2 and show that, while the full statistics is Markovian, there are non-Markovian sub-statistics, i.e., we uncover hidden quantum memory. This phenomenon acts as a witness to the impossibility of a memoryless quantum process description by way of contradiction with (the first part of) Lemma 5.2, which states that any memoryless process leads to Markovian sub-statistics. The circuit corresponding to the process we discuss is shown in Fig. D.1, where, for convenience, the states we explicitly calculate in the discussion below are annotated.

We begin with the full statistics. The probability over measurement outcomes at time t_1 are set by the initial state of the process, i.e.,

$$\mathbb{P}(x_1) = \text{tr} [|x_1\rangle\langle x_1| \rho_1], \tag{D.13}$$

with the post-measurement (subnormalised) state $\rho_2(x_1) = \mathbb{P}(x_1) |x_1\rangle\langle x_1|$; without loss of generality, we choose $\rho_1 = \frac{1}{2}$ and thereby set $\mathbb{P}(x_1 = 0) = \mathbb{P}(x_1 = 1) = \frac{1}{2}$. The process

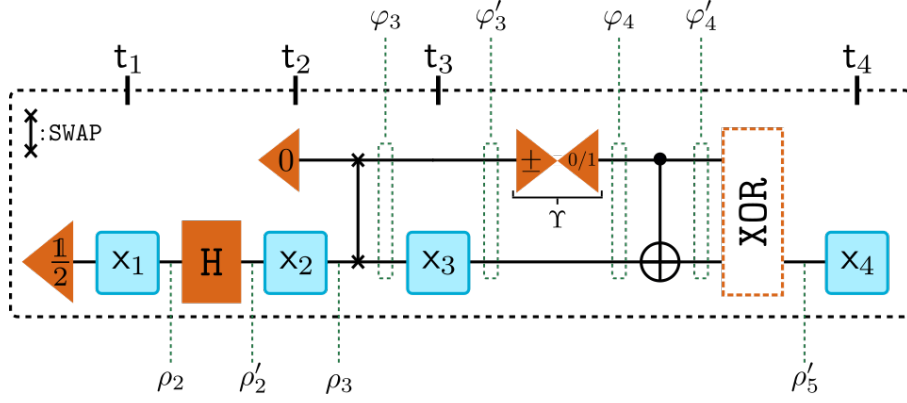


Figure D.1: Circuit with Hidden Quantum Memory / Incompatible Markovian Statistics.

For convenience, we reproduce the circuit provided in Fig. 5.2 in the main text. Additionally, to better facilitate orientation, the states that are explicitly mentioned throughout the proof are marked in green, i.e., the points in the circuit where the states $\rho_2, \rho'_2, \rho_3, \varphi_3, \dots$ occur. For the proof of Observation 2, the final Exclusive-OR (XOR) gate is not part of the circuit, but it is included in the proof of Observation 3.

then consists of a Hadamard gate, which rotates said post-measurement state (which is diagonal in the σ_z -basis) to the σ_x -basis, and we have

$$\rho'_2(x_1) := H\rho_2(x_1)H = \begin{cases} \frac{1}{2}|+\rangle\langle+| & \text{for } x_1 = 0 \\ \frac{1}{2}|-\rangle\langle-| & \text{for } x_1 = 1. \end{cases} \quad (\text{D.14})$$

The experimenter then measures this state again in the computational (σ_z) basis, yielding the joint two-time statistics

$$\mathbb{P}(x_2, x_1) = \frac{1}{4} \quad \forall x_2, x_1. \quad (\text{D.15})$$

The state after the second measurement is independent of x_1 and given by

$$\rho_3(x_2, x_1) = \frac{1}{4}|x_2\rangle\langle x_2| \quad \forall x_2, x_1. \quad (\text{D.16})$$

This state is then swapped with the environment, which is prepared in an arbitrary fiducial state τ , which we set as the blank state $|0\rangle$. The joint system-environment state φ_3 immediately prior to the measurement at t_3 is given by

$$\varphi_3(x_2, x_1) = \text{SWAP}[\rho_3(x_2, x_1) \otimes \tau] = \tau \otimes \rho_3(x_2, x_1) = \frac{1}{4}|0\rangle\langle 0| \otimes |x_2\rangle\langle x_2|. \quad (\text{D.17})$$

The experimenter then measures the system at time t_3 , recording the probabilities

$$\mathbb{P}(x_3, x_2, x_1) = \begin{cases} \frac{1}{4} & \text{for } x_3 = 0 \\ 0 & \text{for } x_3 = 1 \end{cases} \quad \forall x_2, x_1. \quad (\text{D.18})$$

This distribution is Markovian, as we have the conditional probabilities

$$\mathbb{P}(x_3|x_2, x_1) = \begin{cases} 1 & \text{for } x_3 = 0 \\ 0 & \text{for } x_3 = 1 \end{cases} \quad \forall x_2, x_1, \quad (\text{D.19})$$

which is independent of x_1 [indeed, the statistics are “super”-Markovian as the conditional probabilities are also independent of x_2 , so we have $\mathbb{P}(x_3|x_2, x_1) = \mathbb{P}(x_3|x_2) = \mathbb{P}(x_3)$]. The system-environment state $\varphi'_3(x_3, x_2, x_1)$ following the measurement at t_3 is

$$\varphi'_3(x_3, x_2, x_1) = (\mathcal{P}_3^{(x_3)} \otimes \mathcal{I})[\varphi_3(x_2, x_1)] = \frac{1}{4}|0\rangle\langle 0| \otimes |x_2\rangle\langle x_2|, \quad (\text{D.20})$$

i.e., the measurement at t_3 is non-invasive (note that the outcome $x_3 = 1$ cannot occur). Subsequently, a channel occurs that measures the environment in the σ_x -basis and feeds forward $|0\rangle$ ($|1\rangle$) whenever the measurement outcome is $+$ ($-$). The corresponding CPTP map is given by $\Upsilon[\bullet] = \sum_k Y^k \bullet Y^{k\dagger}$ with Kraus operators $Y^0 = |0\rangle\langle +|$ and $Y^1 = |1\rangle\langle -|$. Since $\langle \pm|x_2\rangle\langle x_2|\pm\rangle = \frac{1}{2} \forall x_2$, this yields the system-environment state

$$\varphi_4(x_3, x_2, x_1) = (\mathcal{I} \otimes \Upsilon)[\varphi'_3(x_3, x_2, x_1)] = \frac{1}{8}|0\rangle\langle 0| \otimes \mathbb{1}. \quad (\text{D.21})$$

After this, a Controlled-NOT (**CNOT**) gate on system and environment occurs (with the environment qubit acting as control), leading to

$$\varphi'_4(x_3, x_2, x_1) = \text{CNOT}[\varphi_4(x_3, x_2, x_1)] = \frac{1}{16}\mathbb{1} \otimes \mathbb{1}. \quad (\text{D.22})$$

The experimenter performs the final measurement at t_4 , recording the probabilities

$$\mathbb{P}(x_4, x_3, x_2, x_1) = \begin{cases} \frac{1}{8} & \text{for } x_3 = 0 \\ 0 & \text{for } x_3 = 1 \end{cases} \quad \forall x_4, x_2, x_1. \quad (\text{D.23})$$

This distribution is indeed Markovian, as the conditional probabilities are

$$\mathbb{P}(x_4|x_3, x_2, x_1) = \begin{cases} \frac{1}{2} & \text{for } x_3 = 0 \\ 0 & \text{for } x_3 = 1 \end{cases} \quad \forall x_4, x_2, x_1, \quad (\text{D.24})$$

[where we take the convention that conditioning on an event that cannot occur (i.e., $x_3 = 1$) gives conditional probability 0]. Thus, the full joint statistics $\mathbb{P}(x_4, x_3, x_2, x_1)$ is Markovian.

On the other hand, consider the situation in which the experimenter does not measure at time t_2 , i.e., the sub-statistics $\mathbb{P}(x_4, x_3, \mathcal{I}_2, x_1)$. Everything until Eq. (D.14) remains the same, but without measurement at t_2 we have the state

$$\rho_3(\mathcal{I}_2, x_1) = \mathcal{I}_2[\rho'_2(x_1)] = \frac{1}{2} \begin{cases} |+\rangle\langle +| & \text{for } x_1 = 0 \\ |-\rangle\langle -| & \text{for } x_1 = 1. \end{cases} \quad (\text{D.25})$$

The system is then swapped with the environment, yielding the joint state

$$\varphi_3(\mathcal{I}_2, x_1) = \text{SWAP}[\rho_3(\mathcal{I}_2, x_1) \otimes \tau] = \tau \otimes \rho_3(\mathcal{I}_2, x_1) = \frac{1}{2} \begin{cases} |0\rangle\langle 0| \otimes |+\rangle\langle +| & \text{for } x_1 = 0 \\ |0\rangle\langle 0| \otimes |-\rangle\langle -| & \text{for } x_1 = 1. \end{cases} \quad (\text{D.26})$$

Measurement of the system at t_3 leads to the joint statistics

$$\mathbb{P}(x_3, \mathcal{I}_2, x_1) = \begin{cases} \frac{1}{2} & \text{for } x_3 = 0 \\ 0 & \text{for } x_3 = 1 \end{cases} \quad \forall x_1. \quad (\text{D.27})$$

Thus, we have the conditional probabilities

$$\mathbb{P}(x_3 | \mathcal{I}_2, x_1) = \begin{cases} 1 & \text{for } x_3 = 0 \\ 0 & \text{for } x_3 = 1 \end{cases} \quad \forall x_1. \quad (\text{D.28})$$

The system-environment state $\varphi'_3(x_3, \mathcal{I}_2, x_1)$ following the measurement at t_3 is

$$\varphi'_3(x_3, \mathcal{I}_2, x_1) = (\mathcal{P}_3^{(x_3)} \otimes \mathcal{I})[\varphi_3(\mathcal{I}_2, x_1)] = \frac{1}{2} \begin{cases} |0\rangle\langle 0| \otimes |+\rangle\langle +| & \text{for } x_1 = 0 \\ |0\rangle\langle 0| \otimes |-\rangle\langle -| & \text{for } x_1 = 1. \end{cases} \quad (\text{D.29})$$

After the map Υ on the environment, we have the system-environment state

$$\begin{aligned} \varphi'_4(x_3, \mathcal{I}_2, x_1) &= (\mathcal{I} \otimes \Upsilon)[\varphi'_3(x_3, \mathcal{I}_2, x_1)] = \frac{1}{2} \begin{cases} |0\rangle\langle 0| \otimes |0\rangle\langle 0| & \text{for } x_1 = 0 \\ |0\rangle\langle 0| \otimes |1\rangle\langle 1| & \text{for } x_1 = 1 \end{cases} \\ &= \frac{1}{2} |0\rangle\langle 0| \otimes |x_1\rangle\langle x_1|. \end{aligned} \quad (\text{D.30})$$

Upon application of the CNOT gate, the system-environment state is

$$\varphi'_4(x_3, \mathcal{I}_2, x_1) = \text{CNOT}[\varphi_4(x_3, \mathcal{I}_2, x_1)] = \frac{1}{2} |x_1\rangle\langle x_1| \otimes |x_1\rangle\langle x_1|. \quad (\text{D.31})$$

The experimenter finally performs the measurement at t_4 , recording the statistics

$$\mathbb{P}(x_4, x_3, \mathcal{I}_2, x_1) = \begin{cases} \frac{1}{2} \delta_{x_4 x_1} & \text{for } x_3 = 0 \\ 0 & \text{for } x_3 = 1. \end{cases} \quad (\text{D.32})$$

This sub-statistics is, however, non-Markovian, since the conditional probability at time t_4 depends on x_1 . Explicitly, we have

$$\mathbb{P}(x_4 | x_3, \mathcal{I}_2, x_1) = \frac{\mathbb{P}(x_4, x_3, \mathcal{I}_2, x_1)}{\mathbb{P}(x_3, \mathcal{I}_2, x_1)} = \begin{cases} \delta_{x_4 x_1} & \text{for } x_3 = 0 \\ 0 & \text{for } x_3 = 1 \end{cases} \neq \mathbb{P}(x_4 | x_3). \quad (\text{D.33})$$

As we have discussed in the main text, such non-Markovian sub-statistics cannot arise for a memoryless quantum process probed by sharp projective measurements (as is the case in this example), and therefore we conclude that the statistics observed—although Markovian on the whole—cannot be faithfully reproduced by a memoryless quantum process.

D.3.2 Incompatibility

Here we explicitly calculate all sub-statistics of the example used in Observation 3 and show that although they are all Markovian (i.e., unlike in Observation 2, there is no explicit hidden quantum memory), they are nonetheless incompatible, therefore serving to witness the impossibility of a memoryless quantum process description by way of contradiction with (the second part of) Lemma 5.2.

The process is the same as above, with an additional Exclusive-OR (**XOR**): $|00\rangle \mapsto |0\rangle, |11\rangle \mapsto |0\rangle, |01\rangle \mapsto |1\rangle, |10\rangle \mapsto |1\rangle$. being performed on the final system-environment state, so all the statistics computed until time t_3 remain unchanged. In the case where measurements are made at all times, we have [beginning at Eq. (D.22)]

$$\rho'_5(x_3, x_2, x_1) = \text{XOR}[\varphi'_4(x_3, x_2, x_1)] = \frac{1}{8}\mathbb{1}. \quad (\text{D.34})$$

The system is then measured, leading to the same joint statistics as in the earlier example [i.e., when measurements are made at each time, the XOR gate has no influence on the statistics; see Eqs. (D.23) and (D.24)]

$$\mathbb{P}(x_4, x_3, x_2, x_1) = \begin{cases} \frac{1}{8} & \text{for } x_3 = 0 \\ 0 & \text{for } x_3 = 1 \end{cases} \quad \forall x_4, x_2, x_1. \quad (\text{D.35})$$

These statistics are indeed Markovian, as we have

$$\mathbb{P}(x_4|x_3, x_2, x_1) = \begin{cases} \frac{1}{2} & \text{for } x_3 = 0 \\ 0 & \text{for } x_3 = 1 \end{cases} \quad \forall x_4, x_2, x_1. \quad (\text{D.36})$$

In the case of the sub-statistics where no measurement is made at time t_2 , we have [starting from Eq. (D.31)]

$$\rho'_5(x_3, \mathcal{I}_2, x_1) = \text{XOR}[\varphi'_4(x_3, \mathcal{I}_2, x_1)] = \frac{1}{2}|0\rangle\langle 0|, \quad (\text{D.37})$$

since the system-environment state before the XOR gate is perfectly classically correlated. The measurement at t_4 then yields the statistics

$$\mathbb{P}(x_4, x_3, \mathcal{I}_2, x_1) = \begin{cases} \frac{1}{2} & \text{for } x_4 = 0, x_3 = 0 \\ 0 & \text{for } x_4, x_3 \neq 0 \end{cases} \quad \forall x_1. \quad (\text{D.38})$$

In contrast to the previous example, these sub-statistics are also Markovian, since we have

$$\mathbb{P}(x_4|x_3, \mathcal{I}_2, x_1) = \begin{cases} 1 & \text{for } x_4 = 0, x_3 = 0 \\ 0 & \text{for } x_4, x_3 \neq 0 \end{cases} \quad \forall x_1. \quad (\text{D.39})$$

It is also straightforward to check that the only other relevant sub-statistics, namely $\mathbb{P}(x_4, \mathcal{I}_3, x_2, x_1)$ when the experimenter does not measure at time t_3 , is also Markovian

[which can be deduced from the Markovianity of the full statistics $\mathbb{P}(x_4, x_3, x_2, x_1)$ and noting that the measurement at t_3 is non-invasive, so $\mathbb{P}(x_4, \mathcal{I}_3, x_2, x_1) = \sum_{x_3} \mathbb{P}(x_4, x_3, x_2, x_1)$].

Thus, in this example, all sub-statistics are Markovian. However, noting the discrepancy between Eqs. (D.36) and (D.39), we see that they are nonetheless incompatible: As we showed in Lemma 5.2, any memoryless quantum process must lead to Markovian sub-statistics that are also compatible, i.e., such that $\mathbb{P}(x_4|x_3, x_2, x_1) = \mathbb{P}(x_4|x_3, \mathcal{I}_2, x_1)$. As this is not the case here, we see a contradiction with the possibility for a memoryless description that faithfully reproduces the observed statistics.

■

Bibliography

- [1] J. Cage, *For the Birds* (Marion Boyars Publishers, London, UK, 1981).
- [2] D. Smyth, *The Bridges of the Yarra* (Schurmann Trevlyn Pty. Ltd., Melbourne, Australia, 1979).
- [3] R. Landauer, *Irreversibility and Heat Generation in the Computing Process*, *IBM J. Res. Dev* **5**, 183 (1961).
- [4] R. Landauer, *Information is Physical*, *Phys. Today* **44**, 23 (1991).
- [5] C. E. Shannon, *A mathematical theory of communication*, *Bell Syst. Tech. J.* **27**, 379 (1948).
- [6] R. V. L. Hartley, *Transmission of information*, *Bell Syst. Tech. J.* **7**, 535 (1928).
- [7] H. Nyquist, *Certain Topics in Telegraph Transmission Theory*, *Trans. Am. Inst. Electr. Eng.* **47**, 617 (1928).
- [8] C. E. Shannon, *A Mathematical Theory of Cryptography* (1945).
- [9] C. E. Shannon and W. Weaver, *The Mathematical Theory of Communication* (University of Illinois Press, Urbana, IL, USA, 1949).
- [10] A. N. Kolmogorov, *Three approaches to the quantitative definition of information*, *Int. J. Comput. Math.* **2**, 157 (1968).
- [11] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. (John Wiley & Sons, Hoboken, NJ, USA, 2005).
- [12] C. H. Bennett, *The thermodynamics of computation—a review*, *Int. J. Theor. Phys.* **21**, 905 (1982).
- [13] L. Brillouin, *Science and Information Theory*, 2nd ed. (Academic Press Inc., London, UK, 1956).
- [14] L. Hardy, *Quantum Theory From Five Reasonable Axioms*, [arXiv:quant-ph/0101012](https://arxiv.org/abs/quant-ph/0101012) (2001).
- [15] C. A. Fuchs, *Quantum Foundations in the Light of Quantum Information*, [arXiv:quant-ph/0106166](https://arxiv.org/abs/quant-ph/0106166) (2001).
- [16] C. A. Fuchs, *Quantum Mechanics as Quantum Information (and only a little more)*, [arXiv:quant-ph/0205039](https://arxiv.org/abs/quant-ph/0205039) (2002).
- [17] L. Hardy, *Reformulating and Reconstructing Quantum Theory*, [arXiv:1104.2066](https://arxiv.org/abs/1104.2066) (2011).
- [18] L. Masanes and M. P. Müller, *A derivation of quantum theory from physical requirements*, *New J. Phys.* **13**, 063001 (2011), [arXiv:1004.1483](https://arxiv.org/abs/1004.1483).

- [19] G. Chiribella and R. W. Spekkens, eds., *Quantum Theory: Informational Foundations and Foils* (Springer, Dordrecht, Netherlands, 2016).
- [20] G. M. D’Ariano, G. Chiribella, and P. Perinotti, *Quantum Theory from First Principles* (Cambridge University Press, Cambridge, UK, 2017).
- [21] S. Popescu and D. Rohrlich, *Quantum nonlocality as an axiom*, *Found. Phys.* **24**, 379 (1994), [arXiv:quant-ph/9508009](#).
- [22] Y. Guryanova, R. Silva, A. J. Short, P. Skrzypczyk, N. Brunner, and S. Popescu, *Exploring the limits of no backwards in time signalling*, *Quantum* **3**, 211 (2019), [arXiv:1708.00669](#).
- [23] S. Aaronson, Y. Atia, and L. Susskind, *On the Hardness of Detecting Macroscopic Superpositions*, [arXiv:2009.07450](#) (2020).
- [24] J. Goold, M. Huber, A. Riera, L. del Rio, and P. Skrzypczyk, *The role of quantum information in thermodynamics—a topical review*, *J. Phys. A* **49**, 143001 (2016), [arXiv:1505.07835](#).
- [25] J. W. Gibbs, *Elementary Principles in Statistical Mechanics: Developed with Especial Reference to the Rational Foundation of Thermodynamics* (Charles Scribner’s Sons, New York, NY, USA, 1902).
- [26] E. T. Jaynes, *Information Theory and Statistical Mechanics*, *Phys. Rev.* **106**, 620 (1957).
- [27] E. T. Jaynes, *Information Theory and Statistical Mechanics. II*, *Phys. Rev.* **108**, 171 (1957).
- [28] J. M. R. Parrondo, J. M. Horowitz, and T. Sagawa, *Thermodynamics of information*, *Nat. Phys.* **11**, 131 (2015).
- [29] M. P. Frank, *The physical limits of computing*, *Comput. Sci. Eng.* **4**, 16 (2002).
- [30] E. Pop, *Energy dissipation and transport in nanoscale devices*, *Nano Res.* **3**, 147 (2010), [arXiv:1003.4058](#).
- [31] A. Bérut, A. Arakelyan, A. Petrosyan, S. Ciliberto, R. Dillenschneider, and E. Lutz, *Experimental verification of Landauer’s principle linking information and thermodynamics*, *Nature* **483**, 187 (2012).
- [32] Y. Jun, M. Gavrilov, and J. Bechhoefer, *High-Precision Test of Landauer’s Principle in a Feedback Trap*, *Phys. Rev. Lett.* **113**, 190601 (2014), [arXiv:1408.5089](#).
- [33] L. Martini, M. Pancaldi, M. Madami, P. Vavassori, G. Gubbiotti, S. Tacchi, F. Hartmann, M. Emmerling, S. Höfling, L. Worschech, and G. Carlotti, *Experimental and theoretical analysis of Landauer erasure in nano-magnetic switches of different sizes*, *Nano Energy* **19**, 108 (2016).
- [34] L. L. Yan, T. P. Xiong, K. Rehan, F. Zhou, D. F. Liang, L. Chen, J. Q. Zhang, W. L. Yang, Z. H. Ma, and M. Feng, *Single-Atom Demonstration of the Quantum Landauer Principle*, *Phys. Rev. Lett.* **120**, 210601 (2018), [arXiv:1803.10424](#).
- [35] L. Szilard, *Über die Entropieverminderung in einem thermodynamischen System bei Eingriffen intelligenter Wesen*, *Z. Phys.* **53**, 840 (1929), English version: *On the decrease of entropy in a thermodynamic system by the intervention of intelligent beings*; *Syst. Res.*, **9**, 301 (1964).
- [36] M. B. Plenio and V. Vitelli, *The physics of forgetting: Landauer’s erasure principle and information theory*, *Contemp. Phys.* **42**, 25 (2001), [arXiv:quant-ph/0103108](#).

- [37] H. S. Leff and A. R. Rex, *Maxwell's Demon 2: Entropy, Classical and Quantum Information, Computing*, 2nd ed. (CRC Press, Boca Raton, FL, USA, 2003).
- [38] K. Maruyama, F. Nori, and V. Vedral, *Colloquium: The physics of Maxwell's demon and information*, *Rev. Mod. Phys.* **81**, 1 (2009), [arXiv:0707.3400](#).
- [39] S. Ciliberto, Landauer's Bound and Maxwell's Demon, in *Information Theory: Poincaré Seminar 2018*, edited by B. Duplantier and V. Rivasseau (Springer International Publishing, Cham, Switzerland, 2021) p. 87.
- [40] W. Nernst, *Über die Beziehung zwischen Wärmeentwicklung und maximaler Arbeit bei kondensierten Systemen*, *Ber. Kgl. Pr. Akad.* **52**, 933 (1906).
- [41] J. Wilks, *The Third Law of Thermodynamics* (Oxford University Press, Oxford, UK, 1961).
- [42] L. Masanes and J. Oppenheim, *A general derivation and quantification of the third law of thermodynamics*, *Nat. Comm.* **8**, 14538 (2017), [arXiv:1412.3828](#).
- [43] N. Freitas, R. Gallego, L. Masanes, and J. P. Paz, Cooling to Absolute Zero: The Unattainability Principle, in *Thermodynamics in the Quantum Regime: Fundamental Aspects and New Directions*, edited by F. Binder, L. A. Correa, C. Gogolin, J. Anders, and G. Adesso (Springer International Publishing, Cham, Switzerland, 2018) Chap. 25, p. 597, [arXiv:1911.06377](#).
- [44] J. Baugh, O. Moussa, C. A. Ryan, A. Nayak, and R. Laflamme, *Experimental implementation of heat-bath algorithmic cooling using solid-state nuclear magnetic resonance*, *Nature* **438**, 470 (2005), [arXiv:quant-ph/0512024](#).
- [45] G. N. Price, S. T. Bannerman, K. Viering, E. Narevicius, and M. G. Raizen, *Single-Photon Atomic Cooling*, *Phys. Rev. Lett.* **100**, 093004 (2008), [arXiv:0802.0683](#).
- [46] F. Binder, L. A. Correa, C. Gogolin, J. Anders, and G. Adesso, eds., *Thermodynamics in the Quantum Regime: Fundamental Aspects and New Directions* (Springer International Publishing, Cham, Switzerland, 2018).
- [47] C. H. Bennett, *Logical Reversibility of Computation*, *IBM J. Res. Dev.* **17**, 525 (1973).
- [48] T. Toffoli, *Reversible Computing*, *MIT Lab. Comp. Sci. MIT/LCS/TM-151* (1980).
- [49] E. Fredkin and T. Toffoli, *Conservative logic*, *Int. J. Theor. Phys.* **21**, 219 (1982).
- [50] A. Acín, I. Bloch, H. Buhrman, T. Calarco, C. Eichler, J. Eisert, D. Esteve, N. Gisin, S. J. Glaser, F. Jelezko, S. Kuhr, M. Lewenstein, M. F. Riedel, P. O. Schmidt, R. Thew, A. Wallraff, I. Walmsley, and F. K. Wilhelm, *The quantum technologies roadmap: a European community view*, *New J. Phys.* **20**, 080201 (2018), [arXiv:1712.03773](#).
- [51] J. Preskill, *Quantum Computing in the NISQ era and beyond*, *Quantum* **2**, 79 (2018), [arXiv:1801.00862](#).
- [52] F. C. Binder, S. Vinjanampathy, K. Modi, and J. Goold, *Quantacell: Powerful charging of quantum batteries*, *New J. Phys.* **17**, 075015 (2015), [arXiv:1503.07005](#).
- [53] F. Campaioli, F. A. Pollock, F. C. Binder, L. C. Céleri, J. Goold, S. Vinjanampathy, and K. Modi, *Enhancing the Charging Power of Quantum Batteries*, *Phys. Rev. Lett.* **118**, 150601 (2017), [arXiv:1612.04991](#).

- [54] N. Friis and M. Huber, *Precision and Work Fluctuations in Gaussian Battery Charging*, [Quantum](#) **2**, 61 (2018), [arXiv:1708.00749](#).
- [55] F. Campaioli, F. A. Pollock, and S. Vinjanampathy, Quantum Batteries, in *Thermodynamics in the Quantum Regime*, edited by F. Binder, L. A. Correa, C. Gogolin, J. Anders, and G. Adesso (Springer International Publishing, Cham, Switzerland, 2019) Chap. 8, p. 207, [arXiv:1805.05507](#).
- [56] A. Peres, *Measurement of time by quantum clocks*, [Am. J. Phys.](#) **48**, 552 (1980).
- [57] P. Erker, M. T. Mitchison, R. Silva, M. P. Woods, N. Brunner, and M. Huber, *Autonomous Quantum Clocks: Does Thermodynamics Limit Our Ability to Measure Time?*, [Phys. Rev. X](#) **7**, 031022 (2017), [arXiv:1609.06704](#).
- [58] E. Schwarzthans, M. P. E. Lock, P. Erker, N. Friis, and M. Huber, *Autonomous Temporal Probability Concentration: Clockworks and the Second Law of Thermodynamics*, [Phys. Rev. X](#) **11**, 011046 (2021), [arXiv:2007.01307](#).
- [59] L. Brillouin, *Maxwell's Demon Cannot Operate: Information and Entropy. I*, [J. Appl. Phys.](#) **22**, 334 (1951).
- [60] H. E. D. Scovil and E. O. Schulz-DuBois, *Three-Level Masers as Heat Engines*, [Phys. Rev. Lett.](#) **2**, 262 (1959).
- [61] R. Kosloff and A. Levy, *Quantum Heat Engines and Refrigerators: Continuous Devices*, [Annu. Rev. Phys. Chem.](#) **65**, 365 (2014), [arXiv:1310.0683](#).
- [62] R. Uzdin, A. Levy, and R. Kosloff, *Equivalence of Quantum Heat Machines, and Quantum-Thermodynamic Signatures*, [Phys. Rev. X](#) **5**, 031044 (2015), [arXiv:1502.06592](#).
- [63] M. T. Mitchison, *Quantum thermal absorption machines: refrigerators, engines and clocks*, [Contemp. Phys.](#) **60**, 164 (2019), [arXiv:1902.02672](#).
- [64] J. Preskill, Fault-tolerant quantum computation, in *Introduction to Quantum Computation*, edited by H.-K. Lo, S. Popescu, and T. P. Spiller (World-Scientific, Singapore, 1997) Chap. 8, p. 213, [arXiv:quant-ph/9712048](#).
- [65] N. van Kampen, *Stochastic Processes in Physics and Chemistry*, 3rd ed. (Elsevier, New York, NY, USA, 2011).
- [66] H. Wio, R. Deza, and J. López, *An Introduction to Stochastic Processes and Nonequilibrium Statistical Physics*, Revised ed. (World Scientific, Singapore, 2012).
- [67] A. Ishizaki and G. R. Fleming, *Unified treatment of quantum coherent and incoherent hopping dynamics in electronic energy transfer: Reduced hierarchy equation approach*, [J. Chem. Phys.](#) **130**, 234111 (2009).
- [68] M. J. Biercuk, H. Uys, A. P. VanDevender, N. Shiga, W. M. Itano, and J. J. Bollinger, *Optimized dynamical decoupling in a model quantum memory*, [Nature](#) **458**, 996 (2009), [arXiv:0812.5095](#).
- [69] J. T. Barreiro, P. Schindler, O. Gühne, T. Monz, M. Chwalla, C. F. Roos, M. Hennrich, and R. Blatt, *Experimental multiparticle entanglement dynamics induced by decoherence*, [Nat. Phys.](#) **6**, 943 (2010), [arXiv:1005.1965](#).

- [70] S. F. Huelga and M. B. Plenio, *Vibrations, quanta and biology*, *Contemp. Phys.* **54**, 181 (2013), [arXiv:1307.3530](#).
- [71] K. Geerlings, Z. Leghtas, I. M. Pop, S. Shankar, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi, and M. H. Devoret, *Demonstrating a Driven Reset Protocol for a Superconducting Qubit*, *Phys. Rev. Lett.* **110**, 120501 (2013), [arXiv:1211.0491](#).
- [72] R. Schmidt, M. F. Carusela, J. P. Pekola, S. Suomela, and J. Ankerhold, *Work and heat for two-level systems in dissipative environments: Strong driving and non-Markovian dynamics*, *Phys. Rev. B* **91**, 224303 (2015), [arXiv:1412.6442](#).
- [73] B. Bylicka, M. Tukiainen, D. Chruściński, J. Piilo, and S. Maniscalco, *Thermodynamic power of non-Markovianity*, *Sci. Rep.* **6**, 27989 (2016), [arXiv:1504.06533](#).
- [74] J. Iles-Smith, A. G. Dijkstra, N. Lambert, and A. Nazir, *Energy transfer in structured and unstructured environments: Master equations beyond the Born-Markov approximations*, *J. Chem. Phys.* **144**, 44110 (2016), [arXiv:1511.05181](#).
- [75] A. Kato and Y. Tanimura, *Quantum heat current under non-perturbative and non-Markovian conditions: Applications to heat machines*, *J. Chem. Phys.* **145**, 224105 (2016), [arXiv:1609.08783](#).
- [76] M. Pezzutto, M. Paternostro, and Y. Omar, *Implications of non-Markovian quantum dynamics for the Landauer bound*, *New J. Phys.* **18**, 123018 (2016), [arXiv:1608.03497](#).
- [77] D. Basilewitsch, R. Schmidt, D. Sugny, S. Maniscalco, and C. P. Koch, *Beating the limits with initial correlations*, *New J. Phys.* **19**, 113042 (2017), [arXiv:1703.04483](#).
- [78] M. Naghiloo, J. J. Alonso, A. Romito, E. Lutz, and K. W. Murch, *Information Gain and Loss for a Quantum Maxwell's Demon*, *Phys. Rev. Lett.* **121**, 030604 (2018), [arXiv:1802.07205](#).
- [79] J. Fischer, D. Basilewitsch, C. P. Koch, and D. Sugny, *Time-optimal control of the purification of a qubit in contact with a structured environment*, *Phys. Rev. A* **99**, 033410 (2019), [arXiv:1901.05756](#).
- [80] Z.-X. Man, Y.-J. Xia, and R. Lo Franco, *Validity of the Landauer principle and quantum memory effects via collisional models*, *Phys. Rev. A* **99**, 042106 (2019), [arXiv:1811.11355](#).
- [81] N. Linden, S. Popescu, and P. Skrzypczyk, *How Small Can Thermal Machines Be? The Smallest Possible Refrigerator*, *Phys. Rev. Lett.* **105**, 130401 (2010), [arXiv:0908.2076](#).
- [82] F. Clivaz, R. Silva, G. Haack, J. B. Brask, N. Brunner, and M. Huber, *Unifying Paradigms of Quantum Refrigeration: A Universal and Attainable Bound on Cooling*, *Phys. Rev. Lett.* **123**, 170605 (2019), [arXiv:1903.04970](#).
- [83] F. Clivaz, R. Silva, G. Haack, J. B. Brask, N. Brunner, and M. Huber, *Unifying paradigms of quantum refrigeration: Fundamental limits of cooling and associated work costs*, *Phys. Rev. E* **100**, 042130 (2019), [arXiv:1710.11624](#).
- [84] L. J. Schulman, T. Mor, and Y. Weinstein, *Physical Limits of Heat-Bath Algorithmic Cooling*, *Phys. Rev. Lett.* **94**, 120501 (2005).
- [85] S. Raeisi and M. Mosca, *Asymptotic Bound for Heat-Bath Algorithmic Cooling*, *Phys. Rev. Lett.* **114**, 100404 (2015), [arXiv:1407.3232](#).

- [86] N. A. Rodríguez-Briones and R. Laflamme, *Achievable Polarization for Heat-Bath Algorithmic Cooling*, *Phys. Rev. Lett.* **116**, 170501 (2016), [arXiv:1412.6637](#).
- [87] S. Raeisi, M. Kieferová, and M. Mosca, *Novel Technique for Robust Optimal Algorithmic Cooling*, *Phys. Rev. Lett.* **122**, 220501 (2019), [arXiv:1902.04439](#).
- [88] E. Köse, S. Çakmak, A. Gençten, I. K. Kominis, and O. E. Müstecaplıoğlu, *Algorithmic quantum heat engines*, *Phys. Rev. E* **100**, 012109 (2019), [arXiv:1901.01166](#).
- [89] N. H. Y. Ng and M. P. Woods, Resource Theory of Quantum Thermodynamics: Thermal Operations and Second Laws, in *Thermodynamics in the Quantum Regime: Fundamental Aspects and New Directions*, edited by F. Binder, L. A. Correa, C. Gogolin, J. Anders, and G. Adesso (Springer International Publishing, Cham, Switzerland, 2018) Chap. 26, p. 625, [arXiv:1805.09564](#).
- [90] M. Lostaglio, *An introductory review of the resource theory approach to thermodynamics*, *Rep. Prog. Phys.* **82**, 114001 (2019), [arXiv:1807.11549](#).
- [91] F. Ciccarello, G. M. Palma, and V. Giovannetti, *Collision-model-based approach to non-Markovian quantum dynamics*, *Phys. Rev. A* **87**, 040103(R) (2013), [arXiv:1207.6554](#).
- [92] S. Kretschmer, K. Luoma, and W. T. Strunz, *Collision model for non-Markovian quantum dynamics*, *Phys. Rev. A* **94**, 012106 (2016), [arXiv:1603.00408](#).
- [93] S. Lorenzo, F. Ciccarello, G. M. Palma, and B. Vacchini, *Quantum Non-Markovian Piecewise Dynamics from Collision Models*, *Open Sys. Info. Dyn.* **24**, 1740011 (2017), [arXiv:1706.09025](#).
- [94] S. Lorenzo, F. Ciccarello, and G. M. Palma, *Composite quantum collision models*, *Phys. Rev. A* **96**, 032107 (2017), [arXiv:1705.03215](#).
- [95] S. Campbell, F. Ciccarello, G. M. Palma, and B. Vacchini, *System-environment correlations and Markovian embedding of quantum non-Markovian dynamics*, *Phys. Rev. A* **98**, 012142 (2018), [arXiv:1805.09626](#).
- [96] F. Clivaz, *Optimal Manipulation Of Correlations And Temperature In Quantum Thermodynamics*, *Ph.D. thesis*, University of Geneva (2020), [arXiv:2012.04321](#).
- [97] R. Cohen and S. Havlin, *Complex Networks: Structure, Robustness and Function* (Cambridge University Press, Cambridge, UK, 2010).
- [98] G. J. Chaitin, *On the Length of Programs for Computing Finite Binary Sequences: Statistical Considerations*, *J. ACM* **16**, 145 (1969).
- [99] M. Sipser, *Introduction to the Theory of Computation*, 3rd ed. (Thomson, Boston, MA, USA, 2006).
- [100] J. Watson, *Strategy: An Introduction to Game Theory* (Norton, New York, NY, USA, 2007).
- [101] Á. M. Alhambra, *Quantum many-body systems in thermal equilibrium*, [arXiv:2204.08349](#) (2022).
- [102] D. A. Lidar, Review of Decoherence-Free Subspaces, Noiseless Subsystems, and Dynamical Decoupling, in *Quantum Information and Computation for Chemistry* (John Wiley & Sons, Hoboken, NJ, USA, 2014) Chap. 11, p. 295.
- [103] G. Lindblad, *Non-Markovian quantum stochastic processes and their entropy*, *Commun. Math. Phys.* **65**, 281 (1979).

- [104] L. Accardi, A. Frigerio, and J. T. Lewis, *Quantum Stochastic Processes*, *Publ. Rest. Inst. Math. Sci.* **18**, 97 (1982).
- [105] N. G. van Kampen, *Remarks on Non-Markov Processes*, *Braz. J. Phys.* **28**, 90 (1998).
- [106] P. Taranto, *Memory Effects in Quantum Processes*, *Int. J. Quantum Inf.* **18**, 1941002 (2020), [arXiv:1909.05245](#).
- [107] Á. Rivas, S. F. Huelga, and M. B. Plenio, *Quantum non-Markovianity: Characterization, quantification and detection*, *Rep. Prog. Phys.* **77**, 094001 (2014), [arXiv:1405.0303](#).
- [108] H.-P. Breuer, E.-M. Laine, J. Piilo, and B. Vacchini, *Colloquium: Non-Markovian dynamics in open quantum systems*, *Rev. Mod. Phys.* **88**, 021002 (2016), [arXiv:1505.01385](#).
- [109] I. de Vega and D. Alonso, *Dynamics of non-Markovian open quantum systems*, *Rev. Mod. Phys.* **89**, 015001 (2017), [arXiv:1511.06994](#).
- [110] K. Banaszek, A. Dragan, W. Wasilewski, and C. Radzewicz, *Experimental Demonstration of Entanglement-Enhanced Classical Communication over a Quantum Channel with Correlated Noise*, *Phys. Rev. Lett.* **92**, 257901 (2004), [arXiv:quant-ph/0403024](#).
- [111] C. Sayrin, I. Dotsenko, X. Zhou, B. Peaudecerf, T. Rybarczyk, S. Gleyzes, P. Rouchon, M. Mirrahimi, H. Amini, M. Brune, J.-M. Raimond, and S. Haroche, *Real-time quantum feedback prepares and stabilizes photon number states*, *Nature* **477**, 73 (2011), [arXiv:1107.4027](#).
- [112] A. L. Grimsmo, *Time-Delayed Quantum Feedback Control*, *Phys. Rev. Lett.* **115**, 060402 (2015), [arXiv:1502.06959](#).
- [113] D. M. Reich, N. Katz, and C. P. Koch, *Exploiting Non-Markovianity for Quantum Control*, *Sci. Rep.* **5**, 12430 (2015), [arXiv:1409.7497](#).
- [114] J. Cerrillo, M. Buser, and T. Brandes, *Nonequilibrium quantum transport coefficients and transient dynamics of full counting statistics in the strong-coupling and non-Markovian regimes*, *Phys. Rev. B* **94**, 214308 (2016), [arXiv:1606.05074](#).
- [115] S. J. Whalen, A. L. Grimsmo, and H. J. Carmichael, *Open quantum systems with delayed coherent feedback*, *Quantum Sci. Technol.* **2**, 044008 (2017), [arXiv:1702.05776](#).
- [116] I. A. Luchnikov, S. V. Vintskevich, H. Ouerdane, and S. N. Filippov, *Simulation Complexity of Open Quantum Dynamics: Connection with Tensor Networks*, *Phys. Rev. Lett.* **122**, 160401 (2019), [arXiv:1812.00043](#).
- [117] M. R. Jørgensen and F. A. Pollock, *Exploiting the Causal Tensor Network Structure of Quantum Processes to Efficiently Simulate Non-Markovian Path Integrals*, *Phys. Rev. Lett.* **123**, 240602 (2019), [arXiv:1902.00315](#).
- [118] M. R. Jørgensen and F. A. Pollock, *Discrete memory kernel for multitime correlations in non-Markovian quantum processes*, *Phys. Rev. A* **102**, 052206 (2020), [arXiv:2007.03234](#).
- [119] L. Magrini, P. Rosenzweig, C. Bach, A. Deutschmann-Olek, S. G. Hofer, S. Hong, N. Kiesel, A. Kugi, and M. Aspelmeyer, *Real-time optimal quantum control of mechanical motion at room temperature*, *Nature* **595**, 373 (2021), [arXiv:2012.15188](#).

- [120] H. Ball, T. M. Stace, S. T. Flammia, and M. J. Biercuk, *Effect of noise correlations on randomized benchmarking*, *Phys. Rev. A* **93**, 022303 (2016), [arXiv:1504.05307](#).
- [121] P. Figueroa-Romero, K. Modi, R. J. Harris, T. M. Stace, and M.-H. Hsieh, *Randomized Benchmarking for Non-Markovian Noise*, *PRX Quantum* **2**, 040351 (2021), [arXiv:2107.05403](#).
- [122] P. Figueroa-Romero, K. Modi, and M.-H. Hsieh, *Towards a general framework of Randomized Benchmarking for non-Markovian Noise*, [arXiv:2202.11338](#) (2022).
- [123] L. Viola, E. Knill, and S. Lloyd, *Dynamical Decoupling of Open Quantum Systems*, *Phys. Rev. Lett.* **82**, 2417 (1999), [arXiv:quant-ph/9809071](#).
- [124] C. Arenz, D. Burgarth, P. Facchi, and R. Hillier, *Dynamical decoupling of unbounded Hamiltonians*, *J. Math. Phys.* **59**, 032203 (2018), [arXiv:1704.06143](#).
- [125] G. Chiribella, G. M. D’Ariano, and P. Perinotti, *Quantum Circuit Architecture*, *Phys. Rev. Lett.* **101**, 060401 (2008), [arXiv:0712.1325](#).
- [126] G. Chiribella, G. M. D’Ariano, and P. Perinotti, *Theoretical framework for quantum networks*, *Phys. Rev. A* **80**, 022339 (2009), [arXiv:0904.4483](#).
- [127] S. Mavadia, C. L. Edmunds, C. Hempel, H. Ball, F. Roy, T. M. Stace, and M. J. Biercuk, *Experimental quantum verification in the presence of temporally correlated noise*, *npj Quantum Inf.* **4**, 7 (2018), [arXiv:1706.03787](#).
- [128] F. A. Pollock, C. Rodríguez-Rosario, T. Frauenheim, M. Paternostro, and K. Modi, *Operational Markov Condition for Quantum Processes*, *Phys. Rev. Lett.* **120**, 040405 (2018), [arXiv:1801.09811](#).
- [129] F. A. Pollock, C. Rodríguez-Rosario, T. Frauenheim, M. Paternostro, and K. Modi, *Non-Markovian quantum processes: Complete framework and efficient characterization*, *Phys. Rev. A* **97**, 012127 (2018), [arXiv:1512.00589](#).
- [130] P. Taranto, F. A. Pollock, S. Milz, M. Tomamichel, and K. Modi, *Quantum Markov Order*, *Phys. Rev. Lett.* **122**, 140401 (2019), [arXiv:1805.11341](#).
- [131] P. Taranto, S. Milz, F. A. Pollock, and K. Modi, *Structure of quantum stochastic processes with finite Markov order*, *Phys. Rev. A* **99**, 042108 (2019), [arXiv:1810.10809](#).
- [132] P. Taranto, F. A. Pollock, and K. Modi, *Non-Markovian memory strength bounds quantum process recoverability*, *npj Quantum Inf.* **7**, 149 (2021), [arXiv:1907.12583](#).
- [133] G. A. L. White, C. D. Hill, F. A. Pollock, L. C. L. Hollenberg, and K. Modi, *Demonstration of non-Markovian process characterisation and control on a quantum processor*, *Nat. Commun.* **11**, 6301 (2020), [arXiv:2004.14018](#).
- [134] G. A. L. White, F. A. Pollock, L. C. L. Hollenberg, K. Modi, and C. D. Hill, *Non-Markovian Quantum Process Tomography*, [arXiv:2106.11722](#) (2021).
- [135] G. A. L. White, F. A. Pollock, L. C. L. Hollenberg, C. D. Hill, and K. Modi, *From many-body to many-time physics*, [arXiv:2107.13934](#) (2021).
- [136] Y. Guo, P. Taranto, B.-H. Liu, X.-M. Hu, Y.-F. Huang, C.-F. Li, and G.-C. Guo, *Experimental Demonstration of Instrument-Specific Quantum Memory Effects and Non-Markovian Process Recovery for Common-Cause Processes*, *Phys. Rev. Lett.* **126**, 230401 (2021), [arXiv:2003.14045](#).

- [137] C. A. Rodríguez-Rosario, K. Modi, and A. Aspuru-Guzik, *Linear assignment maps for correlated system-environment states*, *Phys. Rev. A* **81**, 012313 (2010), arXiv:0910.5568.
- [138] K. Modi, *Operational approach to open dynamics and quantifying initial correlations*, *Sci. Rep.* **2**, 581 (2012), arXiv:1011.6138.
- [139] K. Modi, C. A. Rodríguez-Rosario, and A. Aspuru-Guzik, *Positivity in the presence of initial system-environment correlation*, *Phys. Rev. A* **86**, 064102 (2012), arXiv:1203.5209.
- [140] L. Li, M. J. Hall, and H. M. Wiseman, *Concepts of quantum non-Markovianity: A hierarchy*, *Phys. Rep.* **759**, 1 (2018), arXiv:1712.08879.
- [141] M. Siefert, A. Kittel, R. Friedrich, and J. Peinke, *On a quantitative method to analyze dynamical and measurement noise*, *EPL* **61**, 466 (2003), arXiv:physics/0108034.
- [142] F. Böttcher, J. Peinke, D. Kleinhans, R. Friedrich, P. G. Lind, and M. Haase, *Reconstruction of Complex Dynamical Systems Affected by Strong Measurement Noise*, *Phys. Rev. Lett.* **97**, 090603 (2006), arXiv:nlin/0607002.
- [143] D. Kleinhans, R. Friedrich, M. Wächter, and J. Peinke, *Markov properties in presence of measurement noise*, *Phys. Rev. E* **76**, 041109 (2007), arXiv:0705.1222.
- [144] B. Lehle, *Analysis of stochastic time series in the presence of strong measurement noise*, *Phys. Rev. E* **83**, 021113 (2011), arXiv:1010.5641.
- [145] A. N. Kolmogorov, *Grundbegriffe der Wahrscheinlichkeitsrechnung* (Springer, Berlin, Germany, 1933) [*Foundations of the Theory of Probability* (Chelsea, New York, NY, USA1956)].
- [146] S. Milz, F. Sakuldee, F. A. Pollock, and K. Modi, *Kolmogorov extension theorem for (quantum) causal modelling and general probabilistic theories*, *Quantum* **4**, 255 (2020), arXiv:1712.02589.
- [147] A. J. Leggett and A. Garg, *Quantum mechanics versus macroscopic realism: Is the flux there when nobody looks?*, *Phys. Rev. Lett.* **54**, 857 (1985).
- [148] A. J. Leggett, *Realism and the physical world*, *Rep. Prog. Phys.* **71**, 022001 (2008).
- [149] C. Emary, N. Lambert, and F. Nori, *Leggett–Garg inequalities*, *Rep. Prog. Phys.* **77**, 016001 (2014), arXiv:1304.5133.
- [150] S. Milz, F. A. Pollock, and K. Modi, *Reconstructing non-Markovian quantum dynamics with limited control*, *Phys. Rev. A* **98**, 012108 (2018), arXiv:1610.02152.
- [151] S. Milz, M. S. Kim, F. A. Pollock, and K. Modi, *Completely Positive Divisibility Does Not Mean Markovianity*, *Phys. Rev. Lett.* **123**, 040401 (2019), arXiv:1901.05223.
- [152] H.-P. Breuer, E.-M. Laine, and J. Piilo, *Measure for the Degree of Non-Markovian Behavior of Quantum Processes in Open Systems*, *Phys. Rev. Lett.* **103**, 210401 (2009), arXiv:0908.0238.
- [153] E.-M. Laine, J. Piilo, and H.-P. Breuer, *Witness for initial system-environment correlations in open-system dynamics*, *EPL* **92**, 60010 (2010), arXiv:1004.2184.
- [154] Á. Rivas, S. F. Huelga, and M. B. Plenio, *Entanglement and Non-Markovianity of Quantum Evolutions*, *Phys. Rev. Lett.* **105**, 050403 (2010), arXiv:0911.4270.

- [155] B. Vacchini, A. Smirne, E.-M. Laine, J. Piilo, and H.-P. Breuer, *Markovianity and non-Markovianity in quantum and classical systems*, *New J. Phys.* **13**, 093004 (2011), [arXiv:1106.0138](#).
- [156] S. C. Hou, X. X. Yi, S. X. Yu, and C. H. Oh, *Alternative non-Markovianity measure by divisibility of dynamical maps*, *Phys. Rev. A* **83**, 062115 (2011), [arXiv:1102.4659](#).
- [157] L. Mazzola, C. A. Rodríguez-Rosario, K. Modi, and M. Paternostro, *Dynamical role of system-environment correlations in non-Markovian dynamics*, *Phys. Rev. A* **86**, 010102(R) (2012), [arXiv:1203.3723](#).
- [158] C. A. Rodríguez-Rosario, K. Modi, L. Mazzola, and A. Aspuru-Guzik, *Unification of witnessing initial system-environment correlations and witnessing non-Markovianity*, *EPL* **99**, 20010 (2012), [arXiv:1204.2197](#).
- [159] S. Kochen and E. Specker, *The Problem of Hidden Variables in Quantum Mechanics*, *Indiana Univ. Math. J.* **17**, 59 (1968).
- [160] C. Budroni, A. Cabello, O. Gühne, M. Kleinmann, and J.-Å. Larsson, *Quantum Contextuality*, [arXiv:2102.13036](#) (2021).
- [161] J. S. Bell, *On the Einstein Podolsky Rosen paradox*, *Phys. Phys. Fiz.* **1**, 195 (1964).
- [162] J. S. Bell, *On the Problem of Hidden Variables in Quantum Mechanics*, *Rev. Mod. Phys.* **38**, 447 (1966).
- [163] N. Brunner, D. Cavalcanti, S. Pironio, V. Scarani, and S. Wehner, *Bell nonlocality*, *Rev. Mod. Phys.* **86**, 419 (2014), [arXiv:1303.2849](#).
- [164] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, *Proposed Experiment to Test Local Hidden-Variable Theories*, *Phys. Rev. Lett.* **23**, 880 (1969).
- [165] B. S. Cirel'son, *Quantum generalizations of Bell's inequality*, *Lett. Math. Phys.* **4**, 93 (1980).
- [166] M. Horodecki, P. Horodecki, and R. Horodecki, *Separability of mixed states: necessary and sufficient conditions*, *Phys. Lett. A* **223**, 1 (1996), [arXiv:quant-ph/9605038](#).
- [167] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, *Quantum entanglement*, *Rev. Mod. Phys.* **81**, 865 (2009), [arXiv:quant-ph/0702225](#).
- [168] G. Lüders, *Concerning the state-change due to the measurement process*, *Ann. Phys.* **15**, 663 (2006), [*Ann. Phys.* **443**, 322 (1950)], [arXiv:quant-ph/0403007](#).
- [169] A. Peres, *Quantum Theory: Concepts and Methods* (Kluwer Academic Publishers, New York, NY, USA, 2002).
- [170] A. Smirne, D. Egloff, M. G. Díaz, M. B. Plenio, and S. F. Huelga, *Coherence and non-classicality of quantum Markov processes*, *Quantum Sci. Technol.* **4**, 01LT01 (2019), [arXiv:1709.05267](#).
- [171] S. Nakajima, *On Quantum Theory of Transport Phenomena: Steady Diffusion*, *Progr. Theo. Phys.* **20**, 948 (1958).
- [172] R. Zwanzig, *Ensemble Method in the Theory of Irreversibility*, *J. Chem. Phys.* **33**, 1338 (1960).

- [173] P. Gehring, J. M. Thijssen, and H. S. J. van der Zant, *Single-molecule quantum-transport phenomena in break junctions*, *Nat. Rev. Phys.* **1**, 381 (2019).
- [174] W. Heisenberg, *Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik*, *Z. Phys.* **43**, 172 (1927).
- [175] M. Ozawa, *Universally valid reformulation of the Heisenberg uncertainty principle on noise and disturbance in measurement*, *Phys. Rev. A* **67**, 042105 (2003), [arXiv:quant-ph/0207121](#).
- [176] P. Busch, P. Lahti, and R. F. Werner, *Proof of Heisenberg's Error-Disturbance Relation*, *Phys. Rev. Lett.* **111**, 160405 (2013), [arXiv:1306.1565](#).
- [177] P. Busch and J. Singh, *Lüders theorem for unsharp quantum measurements*, *Phys. Lett. A* **249**, 10 (1998), [arXiv:1304.0054](#).
- [178] A. Arias, A. Gheondea, and S. Gudder, *Fixed points of quantum operations*, *J. Math. Phys.* **43**, 5872 (2002).
- [179] T. Heinosaari and M. M. Wolf, *Nondisturbing quantum measurements*, *J. Math. Phys.* **51**, 092201 (2010), [arXiv:1005.5659](#).
- [180] F. Costa and S. Shrapnel, *Quantum causal modelling*, *New J. Phys.* **18**, 063032 (2016), [arXiv:1512.07106](#).
- [181] M. B. Plenio and S. Virmani, *An Introduction to Entanglement Measures*, *Quantum Info. Comput.* **7**, 1 (2007), [arXiv:quant-ph/0504163](#).
- [182] E. Chitambar, J. I. de Vicente, M. W. Girard, and G. Gour, *Entanglement manipulation beyond local operations and classical communication*, *J. Math. Phys.* **61**, 042201 (2020), [arXiv:1711.03835](#).
- [183] D. Beckman, D. Gottesman, M. A. Nielsen, and J. Preskill, *Causal and localizable quantum operations*, *Phys. Rev. A* **64**, 052309 (2001), [arXiv:quant-ph/0102043](#).
- [184] E. Chitambar and G. Gour, *Critical Examination of Incoherent Operations and a Physically Consistent Resource Theory of Quantum Coherence*, *Phys. Rev. Lett.* **117**, 030401 (2016), [arXiv:1602.06969](#).
- [185] E. Chitambar and G. Gour, *Comparison of incoherent operations and measures of coherence*, *Phys. Rev. A* **94**, 052336 (2016), [arXiv:1602.06969](#).
- [186] I. Marvian and R. W. Spekkens, *How to quantify coherence: Distinguishing speakable and unspeakable notions*, *Phys. Rev. A* **94**, 052324 (2016), [arXiv:1602.08049](#).
- [187] N. Gisin, G. Ribordy, W. Tittel, and H. Zbinden, *Quantum cryptography*, *Rev. Mod. Phys.* **74**, 145 (2002), [arXiv:quant-ph/0101098](#).
- [188] S. Pirandola, U. L. Andersen, L. Banchi, M. Berta, D. Bunandar, R. Colbeck, D. Englund, T. Gehring, C. Lupo, C. Ottaviani, J. L. Pereira, M. Razavi, J. Shamsul Shaari, M. Tomamichel, V. C. Usenko, G. Vallone, P. Villoresi, and P. Wallden, *Advances in quantum cryptography*, *Adv. Opt. Photon.* **12**, 1012 (2020), [arXiv:1906.01645](#).
- [189] V. Giovannetti, S. Lloyd, and L. Maccone, *Advances in quantum metrology*, *Nat. Photonics* **5**, 222 (2011), [arXiv:1102.2318](#).

- [190] G. Tóth and I. Apellaniz, *Quantum metrology from a quantum information science perspective*, *J. Phys. A: Math. Theor.* **47**, 424006 (2014), [arXiv:1405.4878](#).
- [191] R. Demkowicz-Dobrzański, M. Jarzyna, and J. Kołodyński, *Quantum Limits in Optical Interferometry*, *Progr. Opt.* **60**, 345 (2015), [arXiv:1405.7703](#).
- [192] Y. Guryanova, N. Friis, and M. Huber, *Ideal Projective Measurements Have Infinite Resource Costs*, *Quantum* **4**, 222 (2020), [arXiv:1805.11899](#).
- [193] F. Ticozzi and L. Viola, *Quantum resources for purification and cooling: fundamental limits and opportunities*, *Sci. Rep.* **4**, 5192 (2014), [arXiv:1403.8143](#).
- [194] H. Wilming and R. Gallego, *Third Law of Thermodynamics as a Single Inequality*, *Phys. Rev. X* **7**, 041033 (2017), [arXiv:1701.07478](#).
- [195] J. Scharlau and M. P. Müller, *Quantum Horn's lemma, finite heat baths, and the third law of thermodynamics*, *Quantum* **2**, 54 (2018), [arXiv:1605.06092](#).
- [196] J. Åberg, *Truly work-like work extraction via a single-shot analysis*, *Nat. Commun.* **4**, 1925 (2013), [arXiv:1110.6121](#).
- [197] P. Skrzypczyk, A. J. Short, and S. Popescu, *Work extraction and thermodynamics for individual quantum systems*, *Nat. Commun.* **5**, 4185 (2014), [arXiv:1307.1558](#).
- [198] M. Lostaglio, D. Jennings, and T. Rudolph, *Description of quantum coherence in thermodynamic processes requires constraints beyond free energy*, *Nat. Commun.* **6**, 6383 (2015), [arXiv:1405.2188](#).
- [199] M. P. Woods, N. H. Y. Ng, and S. Wehner, *The maximum efficiency of nano heat engines depends on more than temperature*, *Quantum* **3**, 177 (2019), [arXiv:1506.02322](#).
- [200] D. Reeb and M. M. Wolf, *An improved Landauer principle with finite-size corrections*, *New J. Phys.* **16**, 103011 (2014), [arXiv:1306.4352](#).
- [201] P. M. Riechers and M. Gu, *Impossibility of achieving Landauer's bound for almost every quantum state*, *Phys. Rev. A* **104**, 012214 (2021), [arXiv:2103.02337](#).
- [202] M. Huber, M. Perarnau-Llobet, K. V. Hovhannisyan, P. Skrzypczyk, C. Klöckl, N. Brunner, and A. Acín, *Thermodynamic cost of creating correlations*, *New J. Phys.* **17**, 065008 (2015), [arXiv:1404.2169](#).
- [203] D. E. Bruschi, M. Perarnau-Llobet, N. Friis, K. V. Hovhannisyan, and M. Huber, *The thermodynamics of creating correlations: Limitations and optimal protocols*, *Phys. Rev. E* **91**, 032118 (2015), [arXiv:1409.4647](#).
- [204] G. Vitagliano, C. Klöckl, M. Huber, and N. Friis, *Trade-off Between Work and Correlations in Quantum Thermodynamics*, in *Thermodynamics in the Quantum Regime*, edited by F. Binder, L. A. Correa, C. Gogolin, J. Anders, and G. Adesso (Springer International Publishing, Cham, Switzerland, 2019) Chap. 30, p. 731, [arXiv:1803.06884](#).
- [205] F. Bakhshinezhad, F. Clivaz, G. Vitagliano, P. Erker, A. T. Rezakhani, M. Huber, and N. Friis, *Thermodynamically optimal creation of correlations*, *J. Phys. A: Math. Theor.* **52**, 465303 (2019), [arXiv:1904.07942](#).

- [206] J. Anders and V. Giovannetti, *Thermodynamics of discrete quantum processes*, *New J. Phys.* **15**, 033022 (2013), [arXiv:1211.0183](#).
- [207] P. Taranto, F. Bakhshinezhad, P. Schüttelkopf, F. Clivaz, and M. Huber, *Exponential Improvement for Quantum Cooling through Finite-Memory Effects*, *Phys. Rev. Appl.* **14**, 054005 (2020), [arXiv:2004.00323](#).
- [208] Y.-Z. Zhen, D. Egloff, K. Modi, and O. Dahlsten, *Universal Bound on Energy Cost of Bit Reset in Finite Time*, *Phys. Rev. Lett.* **127**, 190602 (2021), [arXiv:2106.00580](#).
- [209] E. G. Brown, N. Friis, and M. Huber, *Passivity and practical work extraction using Gaussian operations*, *New J. Phys.* **18**, 113028 (2016), [arXiv:1608.04977](#).
- [210] T. Debarba, G. Manzano, Y. Guryanova, M. Huber, and N. Friis, *Work estimation and work fluctuations in the presence of non-ideal measurements*, *New J. Phys.* **21**, 113002 (2019), [arXiv:1902.08568](#).
- [211] G. Manzano, G.-L. Giorgi, R. Fazio, and R. Zambrini, *Boosting the performance of small autonomous refrigerators via common environmental effects*, *New J. Phys.* **21**, 123026 (2019), [arXiv:1908.10259](#).
- [212] W. T. Strunz, L. Diósi, and N. Gisin, *Open System Dynamics with Non-Markovian Quantum Trajectories*, *Phys. Rev. Lett.* **82**, 1801 (1999), [arXiv:quant-ph/9803079](#).
- [213] M. W. Jack and M. J. Collett, *Continuous measurement and non-Markovian quantum trajectories*, *Phys. Rev. A* **61**, 062106 (2000).
- [214] P. Figueroa-Romero, K. Modi, and F. A. Pollock, *Almost Markovian processes from closed dynamics*, *Quantum* **3**, 136 (2019), [arXiv:1802.10344](#).
- [215] P. Figueroa-Romero, K. Modi, and F. A. Pollock, *Equilibration on average in quantum processes with finite temporal resolution*, *Phys. Rev. E* **102**, 032144 (2020), [arXiv:1905.08469](#).
- [216] P. Strasberg, *Operational approach to quantum stochastic thermodynamics*, *Phys. Rev. E* **100**, 022127 (2019), [arXiv:1810.00698](#).
- [217] P. Strasberg and A. Winter, *Stochastic thermodynamics with arbitrary interventions*, *Phys. Rev. E* **100**, 022135 (2019), [arXiv:1905.07990](#).
- [218] P. Strasberg, *Repeated Interactions and Quantum Stochastic Thermodynamics at Strong Coupling*, *Phys. Rev. Lett.* **123**, 180604 (2019), [arXiv:1907.01804](#).
- [219] P. Strasberg, G. Schaller, N. Lambert, and T. Brandes, *Nonequilibrium thermodynamics in the strong coupling and non-Markovian regime based on a reaction coordinate mapping*, *New J. Phys.* **18**, 73007 (2016), [arXiv:1602.01340](#).
- [220] R. S. Whitney, *Non-Markovian quantum thermodynamics: Laws and fluctuation theorems*, *Phys. Rev. B* **98**, 085415 (2018), [arXiv:1611.00670](#).
- [221] P. O. Boykin, T. Mor, V. Roychowdhury, F. Vatan, and R. Vrijen, *Algorithmic cooling and scalable NMR quantum computers*, *Proc. Natl. Acad. Sci.* **99**, 3388 (2002), [arXiv:quant-ph/0106093](#).
- [222] Á. M. Alhambra, M. Lostaglio, and C. Perry, *Heat-Bath Algorithmic Cooling with optimal thermalization strategies*, *Quantum* **3**, 188 (2019), [arXiv:1807.07974](#).

- [223] T. Rybár, S. N. Filippov, M. Ziman, and V. Bužek, *Simulation of indivisible qubit channels in collision models*, *J. Phys. B* **45**, 154006 (2012), [arXiv:1202.6315](#).
- [224] N. K. Bernardes, A. R. R. Carvalho, C. H. Monken, and M. F. Santos, *Environmental correlations and Markovian to non-Markovian transitions in collisional models*, *Phys. Rev. A* **90**, 032111 (2014), [arXiv:1404.0019](#).
- [225] R. McCloskey and M. Paternostro, *Non-Markovianity and system-environment correlations in a microscopic collision model*, *Phys. Rev. A* **89**, 052120 (2014), [arXiv:1402.4639](#).
- [226] S. Lorenzo, F. Ciccarello, and G. M. Palma, *Class of exact memory-kernel master equations*, *Phys. Rev. A* **93**, 052111 (2016), [arXiv:1603.00248](#).
- [227] B. Çakmak, M. Pezzutto, M. Paternostro, and O. E. Müstecaplıoğlu, *Non-Markovianity, coherence, and system-environment correlations in a long-range collision model*, *Phys. Rev. A* **96**, 022109 (2017), [arXiv:1702.05323](#).
- [228] F. Ciccarello, *Collision models in quantum optics*, *Quantum Meas. Quantum Metrol.* **4**, 53 (2017), [arXiv:1712.04994](#).
- [229] J. Rau, *Relaxation Phenomena in Spin and Harmonic Oscillator Systems*, *Phys. Rev.* **129**, 1880 (1963).
- [230] M. Ziman, P. Štelmachovič, V. Bužek, M. Hillery, V. Scarani, and N. Gisin, *Diluting quantum information: An analysis of information transfer in system-reservoir interactions*, *Phys. Rev. A* **65**, 042105 (2002).
- [231] V. Scarani, M. Ziman, P. Štelmachovič, N. Gisin, and V. Bužek, *Thermalizing Quantum Machines: Dissipation and Entanglement*, *Phys. Rev. Lett.* **88**, 097905 (2002), [arXiv:quant-ph/0110088](#).
- [232] M. Ziman, P. Štelmachovič, and V. Bužek, *Description of Quantum Dynamics of Open Systems Based on Collision-Like Models*, *Open Sys. Info. Dyn.* **12**, 81 (2005), [arXiv:quant-ph/0410161](#).
- [233] J. Åberg, *Catalytic Coherence*, *Phys. Rev. Lett.* **113**, 150402 (2014), [arXiv:1304.1060](#).
- [234] B. P. Lanyon, M. Barbieri, M. P. Almeida, T. Jennewein, T. C. Ralph, K. J. Resch, G. J. Pryde, J. L. O'Brien, A. Gilchrist, and A. G. White, *Simplifying quantum logic using higher-dimensional Hilbert spaces*, *Nat. Phys.* **5**, 134 (2009), [arXiv:0804.0272](#).
- [235] A. Babazadeh, M. Erhard, F. Wang, M. Malik, R. Nouroozi, M. Krenn, and A. Zeilinger, *High-Dimensional Single-Photon Quantum Gates: Concepts and Experiments*, *Phys. Rev. Lett.* **119**, 180510 (2017), [arXiv:1702.07299](#).
- [236] P. Imany, J. A. Jaramillo-Villegas, M. S. Alshaykh, J. M. Lukens, O. D. Odele, A. J. Moore, D. E. Leaird, M. Qi, and A. M. Weiner, *High-dimensional optical quantum logic in large operational spaces*, *npj Quant. Inf.* **5**, 59 (2019), [arXiv:1805.04410](#).
- [237] A. S. L. Malabarba, A. J. Short, and P. Kammerlander, *Clock-driven quantum thermal engines*, *New J. Phys.* **17**, 045027 (2015), [arXiv:1412.1338](#).
- [238] M. P. Woods and M. Horodecki, *The Resource Theoretic Paradigm of Quantum Thermodynamics with Control*, [arXiv:1912.05562](#) (2019).

- [239] M. T. Mitchison, M. Huber, J. Prior, M. P. Woods, and M. B. Plenio, *Realising a quantum absorption refrigerator with an atom-cavity system*, *Quantum Sci. Technol.* **1**, 015001 (2016), arXiv:1603.02082.
- [240] G. Chiribella, Y. Yang, and R. Renner, *Fundamental Energy Requirement of Reversible Quantum Operations*, *Phys. Rev. X* **11**, 021014 (2021), arXiv:1908.10884.
- [241] A. Streltsov, U. Singh, H. S. Dhar, M. N. Bera, and G. Adesso, *Measuring Quantum Coherence with Entanglement*, *Phys. Rev. Lett.* **115**, 020403 (2015), arXiv:1502.05876.
- [242] A. Streltsov, G. Adesso, and M. B. Plenio, *Colloquium: Quantum coherence as a resource*, *Rev. Mod. Phys.* **89**, 041003 (2017), arXiv:1609.02439.
- [243] M.-L. Hu, X. Hu, J. Wang, Y. Peng, Y.-R. Zhang, and H. Fan, *Quantum coherence and geometric quantum discord*, *Phys. Rep.* **762-764**, 1 (2018), arXiv:1703.01852.
- [244] C. Napoli, T. R. Bromley, M. Cianciaruso, M. Piani, N. Johnston, and G. Adesso, *Robustness of Coherence: An Operational and Observable Measure of Quantum Coherence*, *Phys. Rev. Lett.* **116**, 150502 (2016), arXiv:1601.03781.
- [245] M. Hillery, *Coherence as a resource in decision problems: The Deutsch-Jozsa algorithm and a variation*, *Phys. Rev. A* **93**, 012111 (2016), arXiv:1512.01874.
- [246] J. M. Matera, D. Egloff, N. Killoran, and M. B. Plenio, *Coherent control of quantum systems as a resource theory*, *Quantum Sci. Technol.* **1**, 01LT01 (2016), arXiv:1512.07486.
- [247] G. D. Scholes, G. R. Fleming, L. X. Chen, A. Aspuru-Guzik, A. Buchleitner, D. F. Coker, G. S. Engel, R. van Grondelle, A. Ishizaki, D. M. Jonas, J. S. Lundeen, J. K. McCusker, S. Mukamel, J. P. Ogilvie, A. Olaya-Castro, M. A. Ratner, F. C. Spano, K. B. Whaley, and X. Zhu, *Using coherence to enhance function in chemical and biophysical systems*, *Nature* **543**, 647 (2017).
- [248] L. Wang, M. A. Allodi, and G. S. Engel, *Quantum coherences reveal excited-state dynamics in biophysical systems*, *Nat. Rev. Chem.* **3**, 477 (2019).
- [249] J. Åberg, *Quantifying Superposition*, arXiv:0612146 (2006).
- [250] T. Baumgratz, M. Cramer, and M. B. Plenio, *Quantifying Coherence*, *Phys. Rev. Lett.* **113**, 140401 (2014), arXiv:1311.0275.
- [251] F. Levi and F. Mintert, *A quantitative theory of coherent delocalization*, *New J. Phys.* **16**, 033007 (2014), arXiv:1310.6962.
- [252] M. Piani, M. Cianciaruso, T. R. Bromley, C. Napoli, N. Johnston, and G. Adesso, *Robustness of asymmetry and coherence of quantum states*, *Phys. Rev. A* **93**, 042107 (2016), arXiv:1601.03782.
- [253] A. Winter and D. Yang, *Operational Resource Theory of Coherence*, *Phys. Rev. Lett.* **116**, 120404 (2016), arXiv:1506.07975.
- [254] E. Chitambar and G. Gour, *Quantum resource theories*, *Rev. Mod. Phys.* **91**, 025001 (2019), arXiv:1806.06107.
- [255] T. Theurer, D. Egloff, L. Zhang, and M. B. Plenio, *Quantifying Operations with an Application to Coherence*, *Phys. Rev. Lett.* **122**, 190405 (2019), arXiv:1806.07332.

- [256] K.-D. Wu, T. Theurer, G.-Y. Xiang, C.-F. Li, G.-C. Guo, M. B. Plenio, and A. Streltsov, *Quantum coherence and state conversion: theory and experiment*, [npj Quantum Inf. **6**, 22 \(2020\)](#), [arXiv:1903.01479](#).
- [257] M. M. Wilde, J. M. McCracken, and A. Mizel, *Could light harvesting complexes exhibit non-classical effects at room temperature?*, [Proc. Roy. Soc. A **466**, 1347 \(2010\)](#), [arXiv:0911.1097](#).
- [258] J. S. Briggs and A. Eisfeld, *Equivalence of quantum and classical coherence in electronic energy transfer*, [Phys. Rev E **83**, 051911 \(2011\)](#), [arXiv:1104.2191](#).
- [259] W. H. Miller, *Perspective: Quantum or classical coherence?*, [J. Chem. Phys. **136**, 210901 \(2012\)](#).
- [260] R. de J. León-Montiel and J. P. Torres, *Highly Efficient Noise-Assisted Energy Transport in Classical Oscillator Systems*, [Phys. Rev. Lett. **110**, 218101 \(2013\)](#), [arXiv:1301.1605](#).
- [261] E. J. O'Reilly and A. Olaya-Castro, *Non-classicality of the molecular vibrations assisting exciton energy transfer at room temperature*, [Nat. Commun. **5**, 3012 \(2014\)](#), [arXiv:1301.6970](#).
- [262] C. Budroni, G. Fagundes, and M. Kleinmann, *Memory cost of temporal correlations*, [New J. Phys. **21**, 093018 \(2019\)](#), [arXiv:1902.06517](#).
- [263] M. Asano, T. Hashimoto, A. Khrennikov, M. Ohya, and Y. Tanaka, *Violation of contextual generalization of the Leggett–Garg inequality for recognition of ambiguous figures*, [Phys. Scr. **2014**, 014006 \(2014\)](#), [arXiv:1401.2897](#).
- [264] S. Milz, F. A. Pollock, and K. Modi, *An Introduction to Operational Quantum Dynamics*, [Open Sys. Info. Dyn. **24**, 1740016 \(2017\)](#), [arXiv:1708.00769](#).
- [265] W. Zurek, *Einselction and decoherence from an information theory perspective*, [Ann. Phys. **9**, 855 \(2000\)](#), [arXiv:quant-ph/0011039](#).
- [266] L. Henderson and V. Vedral, *Classical, quantum and total correlations*, [J. Phys. A **34**, 6899 \(2001\)](#), [arXiv:quant-ph/0105028](#).
- [267] H. Ollivier and W. H. Zurek, *Quantum Discord: A Measure of the Quantumness of Correlations*, [Phys. Rev. Lett. **88**, 017901 \(2001\)](#), [arXiv:quant-ph/0105072](#).
- [268] K. Modi, A. Brodutch, H. Cable, T. Paterek, and V. Vedral, *The classical-quantum boundary for correlations: Discord and related measures*, [Rev. Mod. Phys. **84**, 1655 \(2012\)](#), [arXiv:1112.6238](#).
- [269] A. Smirne, T. Nitsche, D. Egloff, S. Barkhofen, S. De, I. Dhand, C. Silberhorn, S. F. Huelga, and M. B. Plenio, *Experimental control of the degree of non-classicality via quantum coherence*, [Quantum Sci. Technol. **5**, 04LT01 \(2020\)](#), [arXiv:1910.11830](#).
- [270] G. Chiribella, G. M. D'Ariano, and P. Perinotti, *Probabilistic theories with purification*, [Phys. Rev. A **81**, 062348 \(2010\)](#), [arXiv:0908.1583](#).
- [271] W. Feller, *An Introduction to Probability Theory and Its Applications*, 2nd ed. (Wiley, New York, NY, USA, 1971).
- [272] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, Oxford, UK, 2007).
- [273] T. Tao, *An Introduction to Measure Theory* (American Mathematical Society, USA, 2011).

- [274] C. Piron, *Ideal measurement and probability in quantum mechanics*, [Erkenntnis](#) **16**, 397 (1981).
- [275] A. Fine, *Hidden Variables, Joint Probability, and the Bell Inequalities*, [Phys. Rev. Lett.](#) **48**, 291 (1982).
- [276] T. Heinosaari and M. Ziman, *The Mathematical Language of Quantum Theory* (Cambridge University Press, Cambridge, UK, 2012).
- [277] D. Egloff, J. M. Matera, T. Theurer, and M. B. Plenio, *Of local Operations and Physical Wires*, [Phys. Rev. X](#) **8**, 031005 (2018), [arXiv:1802.02021](#).
- [278] M. Lax, *Quantum Noise. XI. Multitime Correspondence between Quantum and Classical Stochastic Processes*, [J. Math. Phys.](#) **172**, 350 (1968).
- [279] H. Carmichael, *An Open Systems Approach to Quantum Optics* (Springer-Verlag, Berlin, Germany, 1993).
- [280] C. Gardiner and P. Zoller, *Quantum Noise: A Handbook of Markovian and Non-Markovian Quantum Stochastic Methods with Applications to Quantum Optics* (Springer, Berlin, Germany, 2004).
- [281] V. Gorini, A. Kossakowski, and G. Sudarshan, *Completely positive semigroups of N -level systems*, [J. Math. Phys.](#) **17**, 821 (1976).
- [282] G. Lindblad, *On the generators of quantum dynamical semigroups*, [Commun. Math. Phys.](#) **48**, 119 (1976).
- [283] Á. Rivas and S. F. Huelga, *Open Quantum Systems. An Introduction* (Springer, New York, USA, 2012).
- [284] M. M. Wolf and J. I. Cirac, *Dividing Quantum Channels*, [Commun. Math. Phys.](#) **279**, 147 (2008), [arXiv:math-ph/0611057](#).
- [285] G. Dümcke, *Convergence of multitime correlation functions in the weak and singular coupling limits*, [J. Math. Phys.](#) **24**, 311 (1983).
- [286] F. Sakuldee, S. Milz, F. A. Pollock, and K. Modi, *Non-Markovian quantum control as coherent stochastic trajectories*, [J. Phys. A](#) **51**, 414014 (2018), [arXiv:1802.03190](#).
- [287] G. A. Paz-Silva, M. J. W. Hall, and H. M. Wiseman, *Dynamics of initially correlated open quantum systems: Theory and applications*, [Phys. Rev. A](#) **100**, 042120 (2019), [arXiv:1810.12540](#).
- [288] G. Lindblad, [Response of Markovian and non-Markovian quantum stochastic systems to time-dependent forces](#) (1980), preprint, Stockholm.
- [289] C. Arenz, R. Hillier, M. Fraas, and D. Burgarth, *Distinguishing decoherence from alternative quantum theories by dynamical decoupling*, [Phys. Rev. A](#) **92**, 022102 (2015), [arXiv:1405.7644](#).
- [290] A. A. Budini, *Conditional past-future correlation induced by non-Markovian dephasing reservoirs*, [Phys. Rev. A](#) **99**, 052125 (2019), [arXiv:1903.05259](#).
- [291] D. Kretschmann and R. F. Werner, *Quantum channels with memory*, [Phys. Rev. A](#) **72**, 062323 (2005), [arXiv:quant-ph/0502106](#).

- [292] F. Caruso, V. Giovannetti, C. Lupo, and S. Mancini, *Quantum channels and memory effects*, *Rev. Mod. Phys.* **86**, 1203 (2014), [arXiv:1207.5435](#).
- [293] A. Jamiołkowski, *Linear transformations which preserve trace and positive semidefiniteness of operators*, *Rep. Math. Phys.* **3**, 275 (1972).
- [294] M.-D. Choi, *Completely positive linear maps on complex matrices*, *Linear Algebra Its Appl.* **10**, 285 (1975).
- [295] S. Shrapnel, F. Costa, and G. Milburn, *Updating the Born rule*, *New J. Phys.* **20**, 053010 (2018), [arXiv:1702.01845](#).
- [296] A. Ferraro, L. Aolita, D. Cavalcanti, F. M. Cucchietti, and A. Acín, *Almost all quantum states have nonclassical correlations*, *Phys. Rev. A* **81**, 052318 (2010), [arXiv:0908.3157](#).
- [297] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, Cambridge, UK, 2004).
- [298] K. M. R. Audenaert and M. B. Plenio, *When are correlations quantum?—Verification and quantification of entanglement by simple measurements*, *New J. Phys.* **8**, 266 (2006), [arXiv:quant-ph/0608067](#).
- [299] J. Eisert, F. G. S. L. Brandão, and K. M. R. Audenaert, *Quantitative entanglement witnesses*, *New J. Phys.* **9**, 46 (2007), [arXiv:quant-ph/0607167](#).
- [300] O. Gühne, M. Reimpell, and R. F. Werner, *Estimating Entanglement Measures in Experiments*, *Phys. Rev. Lett.* **98**, 110502 (2007), [arXiv:quant-ph/0607163](#).
- [301] C. M. Li, N. Lambert, Y. N. Chen, G. Y. Chen, and F. Nori, *Witnessing Quantum Coherence: from solid-state to biological systems*, *Sci. Rep.* **2**, 885 (2012), [arXiv:1212.0194](#).
- [302] A. Datta, *A Condition for the Nullity of Quantum Discord*, [arXiv:1003.5256](#) (2010).
- [303] B. Yadin, J. Ma, D. Girolami, M. Gu, and V. Vedral, *Quantum Processes Which Do Not Use Coherence*, *Phys. Rev. X* **6**, 041028 (2016), [arXiv:1512.02085](#).
- [304] A. Acín, N. Brunner, N. Gisin, S. Massar, S. Pironio, and V. Scarani, *Device-Independent Security of Quantum Cryptography against Collective Attacks*, *Phys. Rev. Lett.* **98**, 230501 (2007), [arXiv:quant-ph/0702152](#).
- [305] V. Scarani, *The device-independent outlook on quantum physics*, *Acta Phys. Slovaca* **62**, 347 (2012), [arXiv:1303.3081](#).
- [306] H. Liu, W. Wang, K. Wei, X.-T. Fang, L. Li, N.-L. Liu, H. Liang, S.-J. Zhang, W. Zhang, H. Li, L. You, Z. Wang, H.-K. Lo, T.-Y. Chen, F. Xu, and J.-W. Pan, *Experimental Demonstration of High-Rate Measurement-Device-Independent Quantum Key Distribution over Asymmetric Channels*, *Phys. Rev. Lett.* **122**, 160501 (2019), [arXiv:1808.08584](#).
- [307] J. Kołodyński, A. Máttar, P. Skrzypczyk, E. Woodhead, D. Cavalcanti, K. Banaszek, and A. Acín, *Device-independent quantum key distribution with single-photon sources*, *Quantum* **4**, 260 (2020), [arXiv:1803.07089](#).
- [308] O. Nairz, M. Arndt, and A. Zeilinger, *Quantum interference experiments with large molecules*, *Am. J. Phys.* **71**, 319 (2003).

- [309] A. D. O’Connell, M. Hofheinz, M. Ansmann, R. C. Bialczak, M. Lenander, E. Lucero, M. Neeley, D. Sank, H. Wang, M. Weides, J. Wenner, J. M. Martinis, and A. N. Cleland, *Quantum ground state and single-phonon control of a mechanical resonator*, *Nature* **464**, 697 (2010).
- [310] S. Eibenberger, S. Gerlich, M. Arndt, M. Mayor, and J. Tüxen, *Matter–wave interference of particles selected from a molecular library with masses exceeding 10 000 amu*, *Phys. Chem. Chem. Phys.* **15**, 14696 (2013), arXiv:1310.8343.
- [311] W. H. Zurek, *Decoherence, einselection, and the quantum origins of the classical*, *Rev. Mod. Phys.* **75**, 715 (2003), arXiv:quant-ph/0105127.
- [312] M. Schlosshauer, *Decoherence, the measurement problem, and interpretations of quantum mechanics*, *Rev. Mod. Phys.* **76**, 1267 (2005), arXiv:quant-ph/0312059.
- [313] W.-M. Zhang, P.-Y. Lo, H.-N. Xiong, M. W.-Y. Tu, and F. Nori, *General Non-Markovian Dynamics of Open Quantum Systems*, *Phys. Rev. Lett.* **109**, 170402 (2012), arXiv:1206.4490.
- [314] A. Chakraborty and R. Sensarma, *Power-law tails and non-Markovian dynamics in open quantum systems: An exact solution from Keldysh field theory*, *Phys. Rev. B* **97**, 104306 (2018), arXiv:1709.04472.
- [315] A. Smirne, J. Kołodyński, S. F. Huelga, and R. Demkowicz-Dobrzański, *Ultimate Precision Limits for Noisy Frequency Estimation*, *Phys. Rev. Lett.* **116**, 120801 (2016), arXiv:1511.02708.
- [316] J. F. Haase, A. Smirne, S. F. Huelga, J. Kołodyński, and R. Demkowicz-Dobrzański, *Precision Limits in Quantum Metrology with Open Quantum Systems*, *Quantum Meas. Quantum Metr.* **5**, 13 (2016), arXiv:1807.11882.
- [317] Y. Yang, *Memory Effects in Quantum Metrology*, *Phys. Rev. Lett.* **123**, 110501 (2019), arXiv:1904.07267.
- [318] J. Gemmer, A. Otte, and G. Mahler, *Quantum Approach to a Derivation of the Second Law of Thermodynamics*, *Phys. Rev. Lett.* **86**, 1927 (2001), arXiv:quant-ph/0101140.
- [319] J. Prior, A. W. Chin, S. F. Huelga, and M. B. Plenio, *Efficient Simulation of Strong System-Environment Interactions*, *Phys. Rev. Lett.* **105**, 050404 (2010), arXiv:1003.5503.
- [320] D. Tamascelli, R. Rosenbach, and M. B. Plenio, *Improved scaling of time-evolving block-decimation algorithm through reduced-rank randomized singular value decomposition*, *Phys. Rev. E* **91**, 063306 (2015), arXiv:1504.00992.
- [321] T. Farrelly, F. G. S. L. Brandão, and M. Cramer, *Thermalization and Return to Equilibrium on Finite Quantum Lattice Systems*, *Phys. Rev. Lett.* **118**, 140601 (2017), arXiv:1610.01337.
- [322] D. Tamascelli, A. Smirne, J. Lim, S. F. Huelga, and M. B. Plenio, *Efficient Simulation of Finite-Temperature Open Quantum Systems*, *Phys. Rev. Lett.* **123**, 090402 (2019), arXiv:1811.12418.
- [323] R. Alicki and N. V. Ryn, *A simple test of quantumness for a single system*, *J. Phys. A* **41**, 062001 (2008), arXiv:0704.1962.
- [324] R. Alicki, M. Piani, and N. V. Ryn, *Quantumness witnesses*, *J. Phys. A* **41**, 495303 (2008), arXiv:0807.2615.

- [325] P. Facchi, S. Pascazio, V. Vedral, and K. Yuasa, *Quantumness and entanglement witnesses*, *J. Phys. A* **45**, 105302 (2012), [arXiv:1111.5257](#).
- [326] P. Facchi, L. Ferro, G. Marmo, and S. Pascazio, *Defining quantumness via the Jordan product*, *J. Phys. A* **47**, 035301 (2013), [arXiv:1309.4635](#).
- [327] R. Fazio, K. Modi, S. Pascazio, V. Vedral, and K. Yuasa, *Witnessing the quantumness of a single system: From anticommutators to interference and discord*, *Phys. Rev. A* **87**, 052132 (2013), [arXiv:1201.1212](#).
- [328] V. V. Dodonov, ‘Nonclassical’ states in quantum optics: a ‘squeezed’ review of the first 75 years, *J. Opt. B* **4**, R1 (2002).
- [329] A. Kenfack and K. Życzkowski, *Negativity of the Wigner function as an indicator of non-classicality*, *J. Opt. B* **6**, 396 (2004), [arXiv:quant-ph/0406015](#).
- [330] C. Cormick, E. F. Galvão, D. Gottesman, J. P. Paz, and A. O. Pittenger, *Classicality in discrete Wigner functions*, *Phys. Rev. A* **73**, 012301 (2006), [arXiv:quant-ph/0506222](#).
- [331] M. Bohmann and E. Agudelo, *Phase-Space Inequalities Beyond Negativities*, *Phys. Rev. Lett.* **124**, 133601 (2020), [arXiv:1909.10534](#).
- [332] M. Bohmann, E. Agudelo, and J. Sperling, *Probing nonclassicality with matrices of phase-space distributions*, *Quantum* **4**, 343 (2020), [arXiv:2003.11031](#).
- [333] H. Barnum, C. M. Caves, C. A. Fuchs, R. Jozsa, and B. Schumacher, *Noncommuting Mixed States Cannot Be Broadcast*, *Phys. Rev. Lett.* **76**, 2818 (1996), [arXiv:quant-ph/9511010](#).
- [334] M. Piani, P. Horodecki, and R. Horodecki, *No-Local-Broadcasting Theorem for Multipartite Quantum Correlations*, *Phys. Rev. Lett.* **100**, 090502 (2008), [arXiv:0707.0848](#).
- [335] W. H. Zurek, *Quantum Darwinism*, *Nat. Phys.* **5**, 181 (2009), [arXiv:0903.5082](#).
- [336] R. Horodecki, J. K. Korbicz, and P. Horodecki, *Quantum origins of objectivity*, *Phys. Rev. A* **91**, 032122 (2015), [arXiv:1312.6588](#).
- [337] T. P. Le and A. Olaya-Castro, *Objectivity (or lack thereof): Comparison between predictions of quantum Darwinism and spectrum broadcast structure*, *Phys. Rev. A* **98**, 032103 (2018), [arXiv:1803.00765](#).
- [338] T. P. Le and A. Olaya-Castro, *Strong Quantum Darwinism and Strong Independence are Equivalent to Spectrum Broadcast Structure*, *Phys. Rev. Lett.* **122**, 010403 (2019), [arXiv:1803.08936](#).
- [339] J. K. Korbicz, *Roads to objectivity: Quantum Darwinism, Spectrum Broadcast Structures, and Strong quantum Darwinism – a review*, *Quantum* **5**, 571 (2021), [arXiv:2007.04276](#).
- [340] P. Strasberg and M. G. Díaz, *Classical quantum stochastic processes*, *Phys. Rev. A* **100**, 022120 (2019), [arXiv:1905.03018](#).
- [341] S. Milz, D. Egloff, P. Taranto, T. Theurer, M. B. Plenio, A. Smirne, and S. F. Huelga, *When Is a Non-Markovian Quantum Process Classical?*, *Phys. Rev. X* **10**, 041049 (2020), [arXiv:1907.05807](#).
- [342] J. Pearl, *Causality*, 2nd ed. (Cambridge University Press, Cambridge, UK, 2009).

- [343] J. Jeknić-Dugić, M. Arsenijević, and M. Dugić, *Invertibility as a witness of Markovianity of the quantum dynamical maps*, [arXiv:2012.08360](#) (2021).
- [344] P. Taranto and S. Milz, *Hidden Quantum Memory: Is Memory There When Somebody Looks?*, [arXiv:2204.08298](#) (2022).
- [345] P. Szańkowski and Ł. Cywiński, *Noise representations of open system dynamics*, *Sci. Rep.* **10**, 22189 (2020), [arXiv:2003.09688](#).
- [346] A. Metelmann and A. A. Clerk, *Nonreciprocal quantum interactions and devices via autonomous feedforward*, *Phys. Rev. A* **95**, 013837 (2017), [arXiv:1610.06621](#).
- [347] C. Addis, F. Ciccarello, M. Cascio, G. M. Palma, and S. Maniscalco, *Dynamical decoupling efficiency versus quantum non-Markovianity*, *New J. Phys.* **17**, 123004 (2015), [arXiv:1502.02528](#).
- [348] D. Manzano, *A short introduction to the Lindblad master equation*, *AIP Adv.* **10**, 025106 (2020), [arXiv:1906.04478](#).
- [349] M. Ringbauer, C. J. Wood, K. Modi, A. Gilchrist, A. G. White, and A. Fedrizzi, *Characterizing Quantum Dynamics with Initial System-Environment Correlations*, *Phys. Rev. Lett.* **114**, 090402 (2015), [arXiv:1410.5826](#).
- [350] M. Lax, *Formal Theory of Quantum Fluctuations from a Driven State*, *Phys. Rev.* **129**, 2342 (1963).
- [351] P. Figueroa-Romero, F. A. Pollock, and K. Modi, *Markovianization with approximate unitary designs*, *Commun. Phys.* **4**, 127 (2021), [arXiv:2004.07620](#).
- [352] N. Dowling, P. Figueroa-Romero, F. A. Pollock, P. Strasberg, and K. Modi, *Relaxation of Multitime Statistics in Quantum Systems*, [arXiv:2108.07420](#) (2021).
- [353] N. Dowling, P. Figueroa-Romero, F. A. Pollock, P. Strasberg, and K. Modi, *Equilibration of Non-Markovian Quantum Processes in Finite Time Intervals*, [arXiv:2112.01099](#) (2021).
- [354] C. Giarmatzi and F. Costa, *Witnessing quantum memory in non-Markovian processes*, *Quantum* **5**, 440 (2021), [arXiv:1811.03722](#).
- [355] M. Nery, M. T. Quintino, P. A. Guérin, T. O. Maciel, and R. O. Vianna, *Simple and maximally robust processes with no classical common-cause or direct-cause explanation*, *Quantum* **5**, 538 (2021), [arXiv:2101.11630](#).
- [356] W. van Dam and P. Hayden, *Renyi-entropic bounds on quantum communication*, [arXiv:quant-ph/0204093](#) (2002).
- [357] M. Ohya and D. Petz, *Quantum Entropy and Its Use*, 2nd ed. (Springer, Berlin, Germany, 1993).
- [358] M. Ohya and N. Watanabe, *Quantum Entropy and Its Applications to Quantum Communication and Statistical Physics*, *Entropy* **12**, 1194 (2010).
- [359] A. E. Allahverdyan, K. V. Hovhannisyan, D. Janzing, and G. Mahler, *Thermodynamic limits of dynamic cooling*, *Phys. Rev. E* **84**, 041109 (2011), [arXiv:1107.1044](#).
- [360] J. Ladyman, J. Lambert, and K. Wiesner, *What is a complex system?*, *Eur. J. Philos.* **3**, 33 (2013).

- [361] Y. Holovatch, R. Kenna, and S. Thurner, *Complex systems: Physics beyond physics*, *Eur. J. Phys.* **38**, 023002 (2017), [arXiv:1610.01002](#).
- [362] A. W. Marshall, I. Olkin, and B. C. Arnold, *Inequalities: Theory of Majorization and its Applications*, 2nd ed. (Springer, Berlin, Germany, 2011).
- [363] W. Thirring, *Quantum Mathematical Physics: Atoms, Molecules and Large Systems* (Springer, Berlin, Germany, 2002).
- [364] G. Gamow, *One Two Three... Infinity: Facts and Speculations of Science* (Viking Press, New York, NY, USA, 1947).
- [365] G. Adesso, S. Ragy, and A. R. Lee, *Continuous Variable Quantum Information: Gaussian States and Beyond*, *Open Syst. Inf. Dyn.* **21**, 1440001 (2014), [arXiv:1401.4679](#).
- [366] R. Silva, G. Manzano, P. Skrzypczyk, and N. Brunner, *Performance of autonomous quantum thermal machines: Hilbert space dimension as a thermodynamical resource*, *Phys. Rev. E* **94**, 032120 (2016), [arXiv:1604.04098](#).
- [367] J. V. Bondar, *Schur majorization inequalities for symmetrized sums with applications to tensor products*, *Linear Algebra Its Appl.* **360**, 1 (2003).
- [368] B. Baumgartner and H. Narnhofer, *Analysis of quantum semigroups with GKS–Lindblad generators: II. General*, *J. Phys. A* **41**, 395303 (2008), [arXiv:0806.3164](#).
- [369] W. F. Stinespring, *Positive functions on C^* -algebras*, *Proc. Amer. Math. Soc.* **6**, 211 (1955).