



universität
wien

DISSERTATION / DOCTORAL THESIS

Titel der Dissertation / Title of the Doctoral Thesis

From Resource to Constraint:
Thermodynamic Implications of Control in Timekeeping,
Cooling and Measurement

verfasst von / submitted by

Emanuel Schwarzhans, MSc

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

Wien, 2023 / Vienna 2023

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:

Mag. Dr. Marcus Huber, Privatdoz

Abstract

Throughout the last century, significant advancements have been made in the control of quantum systems. Although the concept of control is frequently used in literature, its definition as a resource remains vague and poorly understood. This thesis is a collection of scientific contributions made during the my PhD studies, focused on investigating the role of control in the achievability of various tasks in physics and thus taking a step towards a more concise definition of the concept of control. To obtain a comprehensive understanding of control, we analyse a diverse range of tasks and settings through a thermodynamic lens, leading to the distinction of two fundamental concepts, *temporal control* and *structural control*.

In Chapter III, we establish a connection between these two concepts by examining the limitations of clocks as imposed by structural control. By identifying the process of *temporal probability concentration* as a fundamental part of any “ticking clock” we are able to identify the conditions that are necessary for precise time keeping and thus for maintaining high levels of temporal control. This, furthermore, allows us to identify fundamental limitations to the task of time keeping independent of the amount of control that can be exerted.

Chapter IV focuses on the paradigmatic task of ground-state cooling, delving into the details and complexities involved in the pursuit of zero-temperature states. Detailed analysis allows us to identify control complexity as a resource whose divergence is crucial for the achievability of ground-state cooling in finite time and with finite energy cost. The findings from this chapter are not only a precise clarification and expansion of the third law of thermodynamics, but also highlight the significant role that control plays in shaping the feasibility of various tasks.

In Chapter II, we thoroughly investigate the influence that control is allowed to have on spontaneous processes and link these to the process of measurement. This results in the development of the *Measurement-Equilibration Hypothesis*, which offers a thermodynamically consistent model for measurements by making realistic assumptions on the extent of control that can be exerted throughout the measurement process. This allows us to draw fundamental conclusions about the interactions that can in principle lead to objective results such as we know them from our every day world. Furthermore, this approach gives us the tools to investigate long standing problems in quantum mechanics, such as the measurement problem from a new perspective.

Kurzfassung

Im Laufe des letzten Jahrhunderts wurden bedeutende Fortschritte in der Kontrolle quantenmechanischer Systeme gemacht. Obwohl das Konzept der Kontrolle in der Literatur häufig erwähnt wird, bleibt seine Definition als Ressource vage und unklar. Diese Dissertation ist eine Sammlung wissenschaftlicher Beiträge, die ich während meiner Doktorarbeit gemacht habe und welche sich auf die Untersuchung der Rolle der Kontrolle bei der Erreichbarkeit verschiedener Aufgaben in der Physik konzentrieren und somit einen Schritt in Richtung einer präziseren Definition des Konzepts der Kontrolle unternehmen. Um ein umfassendes Verständnis der Kontrolle zu erlangen, analysieren wir eine Vielzahl von Aufgaben und Settings durch eine thermodynamische Brille, was zur Unterscheidung zweier grundlegender Konzepte, der *temporalen Kontrolle* und der *strukturellen Kontrolle*, führt.

In Kapitel III etablieren wir eine Verbindung zwischen diesen beiden Konzepten, indem wir die Einschränkungen von Uhren unter Berücksichtigung struktureller Kontrolle untersuchen. Durch die Identifizierung des Prozesses der *temporalen Wahrscheinlichkeitskonzentration*, als ein grundlegendes Merkmal jeder “tickenden Uhr” sind wir in der Lage, die Bedingungen zu identifizieren, die für präzises Zeitmessen und somit für ein hohes Maß an temporaler Kontrolle notwendig sind. Darüber hinaus ermöglicht uns dies, fundamentale Einschränkungen für die Aufgabe der Zeitmessung unabhängig von der Menge an Kontrolle, die ausgeübt werden kann, zu identifizieren.

Kapitel IV konzentriert sich auf die paradigmatische Aufgabe der Grundzustandskühlung und geht auf die Details und Komplexitäten ein, die mit dem Streben nach Nulltemperaturzuständen verbunden sind. Eine detaillierte Analyse ermöglicht es uns, die *Kontrollkomplexität* als eine Ressource zu identifizieren, deren Divergenz entscheidend für die Erreichbarkeit einer Grundzustandskühlung in endlicher Zeit und mit endlichen Energiekosten ist. Die Erkenntnisse aus diesem Kapitel sind nicht nur eine Präzisierung und Erweiterung des dritten Hauptsatzes der Thermodynamik, sondern verdeutlichen auch die bedeutende Rolle, die die Kontrolle bei der Gestaltung der Erreichbarkeit verschiedener Aufgaben spielt.

In Kapitel II untersuchen wir gründlich den Einfluss, den Kontrolle auf spontane Prozesse haben darf, und verknüpfen diese mit dem Prozess der Messung. Dies führt zur Entwicklung der *Measurement-Equilibration Hypothesis*, die ein thermodynamisch konsistentes Modell für Messungen bietet, indem sie realistische Annahmen über das Ausmaß der Kontrolle trifft, die während des gesamten Messprozesses ausgeübt werden kann. Daraus lassen sich grundsätzliche Rückschlüsse auf die Wechselwirkungen ziehen, die prinzipiell zu objektiven Ergebnissen führen können, wie wir sie aus unserer Alltagswelt kennen. Darüber hinaus

gibt uns dieser Ansatz die Werkzeuge an die Hand, um langjährige Probleme in der Quantenmechanik, wie z. B. das Messproblem, aus einer neuen Perspektive zu untersuchen.

Acknowledgements

I would like to begin by acknowledging the fact that academic success and luck do not always treat everyone equally. I was born in Austria as a privileged white man who never had to worry about fulfilling basic needs and had access to free education throughout my life. This is a fortunate situation that is not shared by the majority of people worldwide, a reality which is reflected in the lack of diversity in science both in terms of gender balance as well as ethnicity. Despite being grateful for the opportunities I have had, I recognise that this is not the case for many people, and that it is a fact that must change. While single individuals, and even academia as a whole, may have limited power to address the structural inequalities that cause these disparities, we can and should strive to make a difference.

I would like to express my deepest gratitude to my supervisor, Marcus Huber, for welcoming me into his research group when as a master student and for offering me the opportunity to continue my scientific development as a PhD candidate. In particular I am thankful for his utterly relaxed way of doing science, which provided a secure environment for open discussions and creative exploration of all kinds of scientific ideas. Even more, I am thankful for his approach to academia, and the way he lives science, namely focusing on the fun of doing good science rather than playing the career-maximisation game. It was a lot of fun doing science with you and I hope it will continue!

Furthermore I'd like to give my sincerest thanks to Max Lock, who supported me as a mentor since the beginnings of my academic journey, even more so as he himself just had finished his PhD in the beginning. Both his scientific advice, but in particular also his mental support in moments of despair helped me to get through the ups and downs of my academic upbringing. Especially his relentless optimism, and his ability to see progress in every dead end helped me to find joy also in the more difficult aspects of science. Thanks for everything Max!

I also want to thank Paul Erker, who co-supervised me throughout my masters degree and for the first half of my PhD. Thanks for being that stochastic element that breaks old patterns allowing for new ideas to emerge!

I'd like to thank all the members of the QUITPhysics group, those who are still actively in the group as well as those who moved on, but never truly left. As they say: once a quitter, always a quitter! I was incredibly much fun to share this time of my life with you and to grow as a group together. I came to see you as my academic family rather than just a research group. In particular I want to thank my academic siblings Philip Taranto, Simon Morelli, Jessica Bavaresco, Fabien Clivaz, Parnam Bakhshinezhad, Mirdit Doda, Paul Appel and Xiaoqin Gao

and all my “younger” PhD siblings, having you around made it all so much more fun.

I would also like to thank all my collaborators for countless fruitful discussions and meetings. Your scientific expertise and insights helped me to tackle problems that otherwise might have eluded me. Science, after all, is a collaborative endeavour, and that is precisely how it felt with you. I am looking forward to continuing to work with you, be it projects that are ongoing or new ones. In particular I want to thank Nicolai Friis and Felix Binder, who’s in depth scientific insights have helped me to expand my scientific horizon allowing me to gain a better understanding of many concepts.

Special thanks goes to my siblings, in particular Angelika and Christian Schwarzhans. Geli, thanks for living with me through the struggles of life of the last three decades and for sharing the struggles of your PhD journey with me as well as allowing me to share mine with you. Chrisi, thanks for being the best little brother that one can have! Both of you shaped the way I am, and I am incredibly thankful for having you.

I’d like to thank all the people who helped me throughout the writing of this thesis. Thanks to Marcus Huber and Max Lock for giving me valuable input and support both in terms of the science but also by telling me to “relax” every once in a while when I was getting too stressed. Additionally, I am grateful to Eliza Agudelo for her constructive writing advice and for simply being there to listen to my challenges. Lastly, I would like to thank Tom Rivlin for not only proofreading my English in this thesis, but also for enhancing my language abilities throughout the process.

Finally, I want acknowledge the Austrian Academy of Sciences (ÖAW) and the Institute of Quantum Optics and Quantum Information as well as the Atominstitut of the Technical University for their infrastructural support, and the the Austrian Science Fund (FWF) through the START project Y879-N27, the ESQ Discovery grant “Emergence of physical laws: From mathematical foundations to applications in many body physics” and the ERC consolidator grant COCOQUEST (Grant No. 101043705) for the monetary support.

List of Publications

This is a list of all my publications. The publications that are part of this thesis are framed in red.

Peer reviewed

Emanuel Schwarzhans, Maximilian P. E. Lock, Paul Erker, Nicolai Friis, and Marcus Huber. “Autonomous Temporal Probability Concentration: Clockworks and the Second Law of Thermodynamics”. In: *arXiv:2007.01307 [cond-mat, physics:quant-ph]* (2020-07). DOI: [10.1103/PhysRevX.11.011046](https://doi.org/10.1103/PhysRevX.11.011046). arXiv: [2007.01307 \[cond-mat, physics:quant-ph\]](https://arxiv.org/abs/2007.01307)

Flavio Del Santo and Emanuel Schwarzhans. ““Philosophysics” at the University of Vienna: The (Pre-)History of Foundations of Quantum Physics in the Viennese Cultural Context”. In: *Physics in Perspective* 24.2-3 (2022-10), pp. 125–153. ISSN: 1422-6944, 1422-6960. DOI: [10.1007/s00016-022-00290-y](https://doi.org/10.1007/s00016-022-00290-y). arXiv: [2110.05217 \[physics, physics:quant-ph\]](https://arxiv.org/abs/2110.05217)

Under review

Philip Taranto et al. “Landauer vs. Nernst: What Is the True Cost of Cooling a Quantum System?” In: *arXiv:2106.05151 [quant-ph]* (2021-06). arXiv: [2106.05151 \[quant-ph\]](https://arxiv.org/abs/2106.05151)

Pre-print

Florian Meier, Emanuel Schwarzhans, Paul Erker, and Marcus Huber. “Fundamental Accuracy-Resolution Trade-off for Timekeeping Devices”. In: *arXiv:2301.05173* (2023-01). arXiv: [2301.05173 \[quant-ph\]](https://arxiv.org/abs/2301.05173)

Emanuel Schwarzhans, Felix Binder, Marcus Huber, and Maximilian P. E. Lock. “Quantum Measurements and Equilibration: The Emergence of Objective Reality via Entropy Maximisation”. In preparation, manuscript not yet submitted for publication

Contents

I	Introduction	1
I.1	Setting the stage	1
I.2	The task of timekeeping	7
I.3	The task of ground-state cooling	13
I.4	The emergence of measurements	18
II	Measurements and the emergence of Objectivity	23
II.1	Quantum measurements and equilibration	23
II.1.1	Contribution	23
III	Fundamental limitations of time keeping	39
III.1	Autonomous Temporal Probability Concentration	39
III.1.1	Contribution	39
III.2	Fundamental accuracy-resolution trade-off	63
III.2.1	Contribution	63
IV	The Role of Control for Ground-State Cooling	75
IV.1	What is the True Cost of Cooling a Quantum System?	75
IV.1.1	Contribution	75
V	Conclusion	137

Chapter I

Introduction

I.1 Setting the stage

At the heart of scientific progress is the notion of precise control, a concept that is notoriously challenging to define. Even long before modern science had emerged, we continuously probed and interacted with nature in more or less controlled ways, allowing us to understand certain aspects of it while others kept eluding us. We learned to understand how nature reacts when we act upon it and, as a result, encountered limitations of our own ability to elicit certain responses from nature. All the while this endeavor never was a lonely one, but rather a collective effort that was shaped and broadcast via socio-cultural means. With the advent of written language this process of knowledge transfer was accelerated significantly, eventually leading to the emergence of modern science as we know it today. While we continue to probe the world with the tools at our disposal, which naturally limit the amount of control we can exert upon it, during the last century, and in particular in the last few decades, the amount of control we have over physical systems has increased drastically. This allowed us to uncover novel fundamental properties of the world while also prompting us to develop increasingly precise theories, helping us to understand its foundations in ever greater detail. At the beginning of the 20th century these developments lead to the genesis of quantum theory, allowing us to describe the microscopic world reliably, and, at that time, far beyond our experimental capabilities. However, technology kept evolving and, in the last few decades, allowed us to reach unprecedented levels of control over the microscopic world.

As we approach a domain where fundamental limitations may become increasingly relevant we must consider the role of control in these contexts, which prompts us to ask the following: *What is the precise interplay between control and our ability to achieve certain tasks? What are the fundamental physical limitations to the amount of control we can exert and how can we*

quantify these? And are there fundamental limitations that cannot be surpassed regardless of the amount of control we have? While quantum mechanics already allows us to identify certain fundamental limitations, such as those imposed by the uncertainty relations, on its own, tackling questions of control and agency exceeds its capabilities, as it was not designed as an agent-based, but rather as a fundamental one¹. Fortunately there are two fields of research (which are highly related to each other) dealing with questions of control and resources: information theory and thermodynamics.

The first formal appearance of basic thermodynamic concepts was as early as 1650, when Otto von Guericke built the world’s first vacuum pump [7]. He was the first to use the concepts of pressure, temperature and volume in this context. The invention of the steam engine in the following centuries and its increasing economic and scientific importance in the 19th century ultimately lead to the formalisation of the theory of thermodynamics, most notably by Sadi Carnot [8]. The initial motivation for developing thermodynamics was to provide a notion of efficiency for steam engines, with the goal to understand how much work an engine could perform given certain resources. As such, thermodynamics can be seen as an agent-based theory, describing the capability of agents to manipulate physical systems in order to extract work from them. While thermodynamics was first properly formalised more than half a century before the dawn of quantum mechanics and general relativity, not only does it remain universally valid, but, furthermore it is still at the center of active research on a wide range of fundamental research questions [9–22]. Its major strength leading to this universality can be seen as a direct consequence of its ignorance. Thermodynamics does not care for describing the fundamental building blocks of the world, but rather takes a “coarse-grained” perspective, ignoring the microscopic details, allowing it to describe emergent properties that are in some sense no less fundamental than the properties derived by microscopic theories. This coarse-grained description was not intentionally put into the theory but rather originated from the technical limitations at that time, that being for example the ability to heat up a reservoir or to change the volume of a container by moving a piston, but not interacting with individual particles of the gas (more so as, at that time, the existence of gas particles was still debated [23]). This inbuilt regard for the practical limitations of agents renders thermodynamics, at heart, a theory about control.

Which operations agents can perform on physical systems, or more precisely, the probability for achieving a certain task by performing these operations, crucially depends on the

¹It could be argued that the measurement postulate suggests otherwise, depending on which interpretation of quantum mechanics one adopts [6]. The exact role of measurements in quantum mechanics has been debated since its beginnings with no conclusive answer yet. However, fundamental questions regarding our control over certain resources and the ability to define what even constitutes a resource cannot come from quantum mechanics alone.

information agents have about the systems. However, while thermodynamics was developed to become a theory about control, its relation to the concept of information remained elusive until the first half of the 20th century, when information theory was born and with it the idea that entropy corresponds to a notion of information. Historically, the concept of entropy in thermodynamics was long understood as a property of the physical systems under investigation, rather than relating to the agents interacting with them. Its original intent was to quantify the amount the “useless energy” in a system, *i.e.* energy that cannot be used to perform work. One of the goals of statistical physics was to connect this concept to microscopic properties, which led to the realisation that entropy can also be viewed as related to the number of micro-states that can realise a given macro-state, thus cementing this picture of objective entropy even more. The advent of information theory challenged this conception of entropy, giving it a subjective connotation as a measure of information agents have about physical systems. Rather than being an absolute property of physical systems, entropy became a relational notion. The subjective character of entropy was exemplified most notably by the well-known paradox of James C. Maxwell, later known as Maxwell’s demon [24], in which an all-knowing being could create a heat flow from a cold to a hot bath by performing simple operations, seemingly contradicting the second law of thermodynamics. As was shown by Leó Szilárd [25], precise control of the demon requires detailed information about the system that is being controlled, which in turn requires the demon to perform a measurement of all the relevant properties of the system (*e.g.* the momentum and position of all particles) and to store this information, thus requiring an entropy increase due to the storage of information equal to the reduction of entropy in the system. The realisation that this storage of information comes at an inevitable cost of heat dissipation was possible due to the work of Rolf Landauer [26], which originally was aimed at describing the thermodynamic cost of erasing a bit of information. Landauer showed that performing irreversible operations on bits of information, such as setting a bit from any initial value to zero, *i.e.* erasing it, requires a finite, non-zero amount of heat to be dissipated into the environment, also known as the Landauer Principle. Charles Bennett later realised that this principle establishes a direct relation between information and physical properties, allowing us to identify the cost of resetting the demon’s brain as the source of entropy required to satisfy the second law of thermodynamics² [28]. Information theory therefore can be seen as a completion of thermody-

²The aim of Maxwell’s demon was to show that the second law of thermodynamics does not necessarily hold universally if analysed from a microscopic perspective. In [27] the authors argue that Bennett’s argument in fact does not resolve the paradox due to a circular line of reasoning: The original Landauer principle is based on phase space arguments. Landauer showed that the change of entropy of setting a thermal ensemble of bits all to one is at least $k_B T \log(2)$ per bit (where k_B is Boltzmann’s constant). He then uses the second law of thermodynamics to

namics as a theory of control. The combination of these two theories allows us to describe the limitations of agents not only based on fixed thermodynamic properties of the physical world, but also, from a relational perspective, based on the information agents have about the world and, thus their ability to interact with it.

Apart from thermodynamics, information theory has become an important tool for a wide variety of research fields with applications ranging from finding limitations of data compression and information transmission rates, to the study of algorithmic complexity, to the discussion of foundational philosophical concepts such as Occam’s Razor and many more [29]. One of the major insights that Landauer and Bennett provided is that, although the abstraction of information is a useful approach that allows for a wide range of applications, in the end, information must be encoded in physical systems, or as Landauer put it: *Information is physical* [30]. While the most fundamental quantity of information in classical information theory is the bit, an abstract entity which can either take the value 0 or 1, assuming that the world is governed by the laws of quantum mechanics allows us to encode information in quantum systems. This opens up a wide range of new possibilities but also puts new constraints on the way this information can be processed and stored. In contrast to classical information carriers, quantum bits (qubits) allow for infinitely many different configurations in which one bit of information can be stored by allowing for superpositions of 0 and 1. The process of reading out information, which is so trivial in classical physics that it usually is not even mentioned, must be incorporated in the description of quantum information processing ³, as it has an impact on the state of the system itself. Furthermore, this vast range of possible ways to encode information allows for correlations between systems that could not arise classically, *i.e.* quantum correlations. And finally, while classical bits can simply be observed and accordingly copied and broadcast, unknown qubits can neither fully be copied, nor can they be broadcast [31]. Classical information theory is only equipped to deal with probability distributions, which cannot feature any of the above novelties that quantum mechanical systems bring with them. The development of quantum information theory allowed us to incorporate these novel properties of quantum systems by replacing probability distributions by density matrices and, in turn, Shannon entropy by von Neumann entropy as a measure of information.

Since classical systems have well-defined properties with given values at all times, speaking

argue that, that same exact amount of entropy must appear elsewhere as a heating effect. Therefore the universality of the second law cannot be established via such an argument alone, as it must be assumed in the first place for Landauer’s principle to hold.

³Promoting the relational character of information to the next level, as the information agents have access to needs to be viewed as a result of a measurement they perform on the system.

of entropy as a property that is only defined relationally, *i.e.* with respect to an agent, makes sense, as it represents the ignorance of an agent towards these well defined properties. However, the birth of quantum mechanics prompts us to reconsider this view and ask: Can this relational understanding of entropy be maintained for quantum information, or is entropy an objective property after all? Quantum correlations arising from the phenomenon of entanglement in fact can be seen as a conceptual game changer on this question. Classical systems always have well-defined properties, which means that agents can learn arbitrarily much about them. This, in turn means that their entropy only depends on the agents' knowledge. However, quantum systems that are entangled with parts of the universe that the agent can't access have properties that the agent can't know, or that are not well defined, seemingly independent of the agent. Even though this seems to lead to the conclusion that entropy, rather than being relational, is truly an objective property of the system, incorporating measurements into the picture reveals once more such a relational character. The crucial realisation here is that it matters which part of the universe the agent has access to and how exactly, in other words, it matters which observables the agent can interrogate. Different agents will ascribe different amounts of entropy to a system depending on the way in which they can observe the system. Thus, while not in the sense of classical entropy, the relational character of entropy remains. This connection between attainable information and accessible degrees of freedom (related to the observers ability to interact with them) suggests a relation between control and the observable structure of physical systems, prompting us to tentatively think of a kind of control that relates to this in terms of *structural control*. Clearly, this concept is yet to be defined properly, so rather than viewing the above as a clear definition of this notion, one should view it as a motivation that inspires further investigation, which we will engage with throughout this thesis.

The ability of agents to gain information about a system and to perform operations on it, is not only limited by the structural properties of the observables that the agents have access to, but also by their temporal resolution. Since physical systems dynamically evolve, gaining information about them requires precisely timed interactions. In practice there are no perfectly temporally sharp measurements (as we will discuss in more detail later). Similar to the way in which it is not possible to observe all parts of a universe that features quantum correlation, having limited temporal resolution results in fundamentally reduced obtainable information about this universe[17, 32–34]. In fact, the lower the temporal resolution, the higher is the entropy of the observed system, which in turn reduces the amount of control the agent can exert over the system. This serves as an inspiration to think about a second notion of control connected to information, which we will for now refer to as *temporal control*.

Structural and temporal control share an intricate relationship, which will be explored throughout the main part of this thesis. On the one side, a high degree of structural control allows for the extraction of precise temporal information which in turn allows for better temporal control, which will be explored in more depth in Chapter III. On the other side many tasks critically rely on the access to temporal information, such as for example the task of ground state cooling (explored in Chapter IV), which in turn can help to unlock higher degrees of structural control [14].

While control is one factor that influences an agent's ability to achieve a task, depending on the resources available to the agent, achieving the same goal may require different amounts of control. As a rather trivial example, take the task of cooling. If an agent already has as a resource a system that is at zero temperature, ground state-cooling becomes an almost trivial task that only requires the ability to perform a swap operation on the system and the resource. While control is always exerted over the available resources, the former can in principle be used to produce the latter, making them inter-convertible. However, this task of converting control into resources is one that is highly non-trivial in many cases.

This brings us to one of the primary questions of this thesis: *Given only resources that require little to no control to be prepared, i.e. minimal resources, how does the amount and type of control influence how well we are able to achieve a specific task?* We will elucidate this question by analysing two rather distinct, but highly prevalent tasks, the task of precise time keeping (Chapter III) and the task of ground state cooling (Chapter IV).

One of the prime directives of the works in this thesis is the demand of fair bookkeeping in accordance with the thermodynamic laws. To uphold this directive we chose an operational approach towards problems, trying to give explicit descriptions of all processes involved rather than reverting to abstraction. While the assumption of quantum mechanics is at the base of all the research in this thesis, this assumption, and more precisely the assumption of the measurement postulate of quantum mechanics opens up a big conceptual hole when it comes to thermodynamic consistency, *i.e.* fair bookkeeping. Naively taking the measurement postulate at face value seemingly allows us to break all thermodynamic laws. As an example, say one measures a system and finds it to be in its energetic ground state. Naively this would mean that we now have access to a perfectly cool system without expending any resources, thus contradicting the first⁴, second⁵ and third⁶ laws of thermodynamics. Of course this is not actually

⁴As it changes the energy of the system without accounting for it elsewhere

⁵As it reduces the system's entropy

⁶As it seemingly allows for ground state cooling

the case. First, because measurements are never perfectly precise and second, because they are not for free [35]. The measurement postulate consequently cannot be taken at face value, as it conceals resources. Uncovering these resources requires an explicit and thermodynamically consistent description of measurements. In particular, modelling the dynamics of measurements, which brings us to the second primary question of this thesis: *How can measurement be modelled explicitly in a thermodynamically consistent manner?* We will investigate this question in Chapter II

This thesis presents a collection of theoretical research papers developed during my time as a PhD candidate, in which I provide detailed analysis of different aspects of control in quantum mechanics by using thermodynamic concepts. To this end I investigate the impact of control on an agents ability to keep track of time (Chapter III), develop a deeper understanding of the notion of control itself and ways to quantify it on the example of ground-state cooling (Chapter IV), as well as the ramifications of the lack of control on the measurement process (Chapter II).

I.2 The task of timekeeping

Time has been an elusive concept throughout the history of science. In classical physics, as well as in quantum theory time appears as a parameter of the respective dynamical equations. While the same is true for the theory of relativity, the paradigm shift that followed the advent of the theory of relativity, *i.e.* operationalism [36] gave rise to an operationally-motivated definition of time: *Time is what clocks measure*. Such a definition, of course, by itself, does not yet answer any fundamental questions about the nature of time, but rather allows us to move away from an axiomatic definition that does not allow for any conceptually new insights about time from physical principles, to one that can be defined in terms of a task, *i.e.* the task of time keeping, thus giving us (in principle) the possibility to draw fundamental conclusions based on practical observations and limitations.

Rather than asking what time is, we can now frame the question as a relational one: *How well can an agent tell the time by probing a specific physical system and what must this system look like in order to make a good clock?* This in turn gives rise to questions regarding the quantification of the performance of clocks, as well as the fundamental limitations that clocks are faced with.

In order to analyse these questions we focus on the analysis of autonomous clocks, *i.e.* we

assume that clocks are physical systems that output temporal information spontaneously by themselves without the need for an agent to actively intervene, thus providing a temporal reference. As such they don't allow for any external control, *i.e.* their dynamics are generated by time-independent and energy-conserving Hamiltonians. It might seem like a weirdly specific choice to focus on such a specialised notion of clocks when trying to derive general statements about the limitations of clocks. However, the assumption of autonomy is hardly a specific assumption. In fact, any non-autonomous description of a clock hides resources that might require themselves temporal information to be prepared, thus presupposing another clock. A thermodynamically consistent description thus must render any clock, at heart, autonomous. By focusing our analysis on autonomous clocks we ensure fair bookkeeping of all resources, which is a prerequisite for the derivation of truly fundamental limitations. Note that current atomic clocks are usually far from autonomous, as they rely on feedback mechanisms, require power input and are subject to losses, non of which can be fully accounted for by their analysis⁷. However, the limitations we will arrive at qualitatively apply to any clock, including such clocks. Due to highly challenging experimental constraints, the regime in which current clocks operate is far from reaching the fundamental bounds derived here. Our work should thus be considered of foundational interest rather than being targeted at direct applicability. To avoid a common misconception it should furthermore be noted that autonomous clocks by definition exclude *stop-watches*, *i.e.* physical systems that require an agent to interact with them on two given points in time to “start” and “stop” them in order to extract information about the time that has passed in between these two instances. Such clocks are incapable of providing a temporal reference.

Keeping track of the passage of time presupposes that it is possible to tell the difference between future and past and thus to associate a preferred direction to the passage of time, also known as the *arrow of time* [38–41]. However, the parameter t that appears in the equations of motion does not allow for the distinction between forwards- and backwards-flowing time, it features *time-reversal symmetry*. This symmetry can only be broken by introducing irreversibility into the dynamics, which comes with an inevitable entropy production. We therefore conceptualise clocks not only as physical systems that undergo some periodic evolution, but we include the irreversible nature of temporal information as a necessary part of the description of any clock. To cultivate some intuition on why this must be the case, let us look at a commonly-made, albeit non-autonomous counter example: Consider a pendulum which evolves according

⁷Note that such an analysis usually is of little interest in these settings in the first place, with few exceptions (*e.g.* [37]), as it is practically impossible to keep track of all the sources of control, and because the focus lies on optimisation of frequency stabilisation rather than the analysis of fundamental limitations.

to classical mechanics, and thus is only subject to reversible dynamics. By observing the pendulum it is possible to gain temporal information by counting one unit of time, or one “tick”, whenever it reaches its maximum height. Thus it seems like reversible systems can just as well serve as clocks, and there is no need for irreversible processes. There are, however, two hidden sources of irreversibility in this example, one coming from the process of observing the pendulum, and the other coming thermodynamic cost associated with driving the pendulum. The process of observing the pendulum, *i.e.* the process of measurement, in itself is necessarily irreversible. In this example, measuring the position of the pendulum requires an agent to shine light onto the pendulum such that the reflected photons could be measured by some kind of photon detector which transmits the acquired information to an agent. The process by which this information becomes accessible to the agent is, as I discuss in Chapter II, intricately connected to the process of equilibration and is thus necessarily irreversible. Any system that undergoes some periodic evolution while producing an output of some sort, meaning that energy, and more importantly, information escapes the system in some way, will eventually come to a halt if no further energy is supplied. It equilibrates with its environment, coming to a state where no further non-trivial dynamics exist [39]. The same is the case for the pendulum. While theoretically one could think of a frictionless mounted pendulum in perfect vacuum that should theoretically swing forever, the principal possibility of observing the pendulum, that is, *e.g.* the existence of photons that are reflected from it, will introduce friction and thus eventually drain all kinetic energy from the pendulum. Whether these photons are then measured or not makes no difference in this⁸, which is why it can be seen as the second source of irreversibility. While keeping the first source of irreversibility in mind, in Chapter III we focus on the limitations that are imposed on clocks due to the second one.

Knowing that any irreversible process inevitably has to produce entropy, we can formulate the task of timekeeping in terms of a thermodynamic process that utilises an out-of-equilibrium resource by explicitly modelling clocks as quantum thermal machines. Akin to the spirit of Carnot’s bound, this allows us to derive fundamental limitations that, in principle, should apply to any timekeeping device.

As any entropy-producing process is fundamentally subject to fluctuations [42–44], the time signal produced by clocks is subject to these same fluctuations. To characterise them we use the notion of *Accuracy*, a quantity akin to the well-studied Alan variance [45–47], with the difference that it is not subject to drifts. In essence it quantifies the (in)stability of clocks due

⁸The information about the pendulum that the photons carry into the universe after they interacted with it might eventually become so diluted that it does not allow for any conclusions about the pendulum anymore, thus not constituting a measurement as defined in Chapter II.

to these fundamental fluctuations. At the same time, the evolution of any quantum system is subject to the so called *quantum speed limits* (QSL), which bound from below the time it takes for a system to change from one state to another, due to properties of the Hamiltonian that governs this change [48–50]. Assuming that clocks adhere to quantum physics, this suggests that they can provide temporal information only at a finite rate. In other words, they cannot “tick” with infinite frequency, prompting us to define a second quantity to characterise the performance of clocks called *Resolution*. This quantity relates to the number of “ticks” a clock can produce per second⁹.

While the rate of entropy production is intricately related to the performance of clocks [39, 40, 52], it is not the only factor that constrains their performance. In Section III.1 of we analyse the role of control and complexity for the performance of clocks. Starting from the premise that any clock must feature an irreversible process as outlined above, we explicitly model clocks as thermal machines that utilise a heat flow from a hot bath to a cold bath in order to produce ticks, which allows for precise bookkeeping of all resources. In a first conceptual step we formalise the task of timekeeping as a combination of two stages that any precise clock has to feature: An *irreversible process* as mentioned above, which is necessary to keep track of the passage of time, and an *internal clockwork* which acts as a filter for the irreversible process, mitigating the fluctuations associated to the irreversible process and thus temporally concentrating the probability of individual ticks. We call this *Temporal Probability Concentration* (TPC). By performing this conceptual split we are able to analyse the microscopic structure of clocks and relate clock performance back to the complexity of the microscopic structure of the clock and thus to the structural control an agent can exert. To that end we explicitly model the clockwork as a quantum thermal machine [53] that drives a target system, in our case a d-dimensional quantum system (qudit) we call a “ladder”, to its maximally excited state. We assume that this state is unstable and coupled to a photon field such that it has a chance of spontaneously emitting a photon. This photon can then be measured and registered as a *tick*. The task of the clockwork in this setting is to prepare the ladder in its unstable state such that the probability for it being in that state is temporally concentrated to a narrow time window. With the premise that an ideal clockwork should in principle be capable of concentrating this probability arbitrarily well¹⁰, meaning that it should be capable of approximating a infinitely sharp peak of top level

⁹Here we compare to Schrödinger time, *i.e.* the parameter t that appears in the Schrödinger equation. Note, however, that resolution can also be defined relative to another clock, thus avoiding the definition of a background time [51]

¹⁰A premise which has been refined since the conception of this work, see [4]

population around one point in time and zero probability everywhere else, *i.e.*

$$P_{\text{top}}(t) = \begin{cases} 1, & \text{if } t = t_0 \\ 0, & \text{otherwise} \end{cases}, \quad (\text{I.1})$$

we set out to explicitly construct such a clockwork. The expectation that a thermal machine can achieve such a task perfectly might seem implausible, as it is itself fundamentally stochastic in nature. However, we show that it is possible to explicitly construct an ideal clockwork by assuming a high degree of complexity of the micro structure of the clockwork. In fact, we derive an analytical expression that allows us to directly relate a notion of complexity¹¹ to the degree to which we are able to achieve TPC. In essence, this informs us that the “better” we want TPC to be, *i.e.* the closer we want it to resemble Eq. I.1, the higher must be the dimension of the Hilbert space that the interaction Hamiltonian acts upon in a fine tuned way. A clear quantifier of complexity has yet to be found, and this is a question we discuss in more detail in Chapter IV, which is why the connection we establish here must be viewed as a qualitative one. With such a clockwork at hand we numerically investigate the effect of TPC on *accuracy* and *resolution* and relate it to the complexity of the clockwork. Doing this we are able to draw three major conclusions¹²: First, based on the assumption that there is a minimal time scale $\frac{1}{c}$ at which irreversible processes can occur in actual physical systems, which highly depend on the physical systems under consideration, we show that no matter how well we can achieve TPC, ultimately the combination of *accuracy* and *resolution* is subject to an upper bound based on this timescale. Second, we find that TPC allows us to establish a connection between *accuracy* and *resolution*, that is, there exists a trade-off relation between those two quantities. In particular this means that increasing the TPC, which necessitates an increase in clockwork complexity, allows us to achieve higher *accuracy* at the cost of *resolution*. Third, we recover the known relation between entropy production and *accuracy*, telling us that the achievable accuracy is determined by the entropy dissipated. In particular we find that for a given timescale at which the irreversible event happens there exists a linear regime such that $N \propto \Delta S$ – a feature which seems universal in classical quantum clocks [37, 54, 55] – and establish a numerical relation to the complexity of the clockwork, which qualitatively suggests that increasing the accuracy comes at an increased entropy dissipation as well as requiring increased complexity of the clockwork.

¹¹Namely the number of minimal thermal machines that the clockwork comprises. This is of course only a proxy for complexity. For an in depth discussion of complexity see Chapter IV

¹²Albeit assuming no back action between the environment and the ladder, which unfortunately is implicit in the paper, an oversight which we corrected in [4].

Having this framework at hand we then turn our attention away from the microscopic details of the clockwork and focus on the question of ultimate fundamental limitations of timekeeping and ask: *Are there fundamental limitations to time keeping independent of the specific physical system that acts as a clock?* In Section III.2 we conduct a detailed analysis of the trade-off relation between *accuracy* and *resolution* in general, that is, not focusing on any particular physical regime, or on thermodynamic consistency. We thus leave the paradigm of autonomous clocks behind us and focus purely on the temporal signal, *i.e.* the *ticks* produced by clocks and the probability distributions associated to them. The premise of this work builds upon the main conceptual insight of [1], that is that any clock fundamentally can be split into an irreversible process and TPC. Within this framework we model the irreversible process as an exponential decay process with characteristic timescale $\frac{1}{\Gamma}$, as it is the only time-continuous process that is memoryless. At the same time we put no physical constraints on our ability to perform TPC, thus giving up on exact bookkeeping in order to determine whether there exists an optimal TPC that must bound all other processes¹³.

In order to determine bounds on the trade-off between *accuracy* N and *resolution* ν we first ask about the possibility of utilising the statistics of many *ticks* in order to increase one of these quantities. For this we assume no TPC, that is, we assume that the *ticks* are generated solely according to an exponentially decaying function with rate Γ . By averaging over many *ticks* we show that it is possible to increase the *accuracy* at the cost of *resolution*, such that the *accuracy* is inversely proportional to the *resolution* with a proportionality factor of Γ , *i.e.* $N = \frac{\Gamma}{\nu}$. In terms of control, this can be seen as one extreme, that is, having no control whatsoever. Comparing this scenario to the other extreme, *i.e.* that of, say, unbounded control, is what constitutes the main result of this work. We show that the optimal trade-off relation between *accuracy* and *resolution* is given by $N = \frac{\Gamma^2}{\nu^2}$. However, in order to achieve such a trade-off relation one is required to have a perfect clock to begin with, as it can only be achieved by the TPC in the form of a step function, *i.e.* activating the decay process exactly at one point in time. Such a TPC would therefore require an unbounded amount of control and is therefore practically not achievable, but only approachable.

From a thermodynamic perspective this result is particularly interesting as it shows another difference between post-processing and TPC. In order to achieve an accuracy that is M -fold higher, post-processing will lead to an M -fold decrease in resolution. At the same time, the irreversible events producing the *ticks* still happen at the original rate. The thermodynamic cost

¹³Again, assuming a memoryless irreversible process that produces the ticks. This constraint might be lifted in future work.

per *tick* thus has to increase M -fold as well. On the other side, performing TPC trades those two quantities without “consuming ticks”, it simply increases the accuracy of each irreversible event while decreasing the rate at which these events happen, making it thermodynamically less expensive.

In Chapter III we develop these ideas in more depth and provide detailed analysis and derivations of all the mentioned results, in particular we discuss control and complexity aspects of timekeeping in Section III.1 fundamental trade-off limitations for *accuracy* and *resolution* in Section III.2.

I.3 The task of ground-state cooling

As outlined in the introduction, thermodynamics can be seen as a theory about control. But what exactly do we mean by control, and how can we quantify the degree of control we have over a system? One way of approaching this question is to take an operational approach. Given a limited amount of resources we can ask about the exact operations that allow us to achieve a specific task to a certain degree. Defining a general notion of control is notoriously hard, but considering it in the context of specific tasks makes it easier to understand and more tractable. However, specialisation usually comes at the price of generality.

In order to gain as much general understanding as possible about the notions of control and complexity we therefore have to choose a task that is “maximally useful”. From a resource-theoretic perspective, achieving certain tasks can be regarded as more valuable than others. For instance, the resource theory of athermality [56] tells us which states can be reached from a given initial state by using only free resources, such as thermal states that require minimal control and information to prepare [57], and by only performing energy-conserving operations (albeit arbitrarily well-controlled ones). While the ability to convert between states is influenced by their specific properties, as demonstrated by the use of majorisation and thermo-majorisation, which do not follow a strict ordering, limiting the states to thermal states creates a one-parameter family that does follow a partial order. Furthermore, the state at the top of this order, which can be converted to all other states using thermal operations, is the zero-temperature state, making it a particularly valuable resource. Accessibility to the such a state or an approximation of it is highly valuable for a range of quantum technologies [58–64], as well as a prerequisite for the fundamental achievability of tasks such as precise timekeeping [1, 52] and ideal measurement [35]. Therefore, to gain a comprehensive understanding of the role of control in a specific task with as much generality as possible, it is beneficial to examine the

exemplary case of cooling to the ground state or zero temperature.

In his formulation of the third law of thermodynamics, Walter Nernst showed that it is, in principle, impossible to attain zero temperature. He argues that the entropy of the system asymptotically approaches a constant as the system approaches the ground state [65], thus making it practically impossible to reach zero temperature. However, this suggests that it is possible to approach the ground-state at the cost of some diverging quantity¹⁴. To see this we can look at the simple example of a qubit which is to be cooled. Say that the qubit is initially in a thermal state at the temperature of its environment. Performing a swap operation between the system and two energy levels of the environment with an energy gap that is much larger than that of the qubit (and which is in a thermal state since we have minimal knowledge about it), results in a final state of the system that is colder than before, as the thermal distribution prescribes a decrease of population exponential in the energy eigenvalue. As the energy gap of the environment gets larger the final temperature of the system approaches zero, thus asymptotically allowing for ground-state cooling. However, performing such an operation is not energy-conserving. In order to actually reach the ground state one would have to invest a diverging amount of energy in order to perform this swap.

Interestingly, there is a somewhat different perspective on the problem of ground-state cooling that originated in the work of Rolf Landauer. As mentioned previously, Landauer's principle establishes a thermodynamically motivated connection between physical systems and their information content by introducing a bound on the minimal amount of energy dissipation in order to erase one bit of information, given by $k_B T \ln(2)$. While its original derivation makes use of the second law of thermodynamics, this principle also bears an intricate relation to the third law of thermodynamics. If applied to microscopic systems, erasing information, *i.e.* reducing the rank of a system to 1, is unitarily equivalent to ground-state cooling. What this tells us is that it could in principle be possible to cool a system to the ground state by only expending a finite amount of energy, leading to an apparent contradiction with Nernst's principle. However, explicit examination of the protocols that implement Landauer erasure shows that such protocols, in fact, would take an infinite amount of time to complete. This uncovers the time as the quantity that becomes unbounded in this scenario.

However, this raises a follow-up question: Is it possible to develop a protocol that can accomplish this task within a limited amount of time and with limited energy dissipation? In Chapter IV, we demonstrate that this question can be answered affirmatively, and provide con-

¹⁴I explicitly refrain from calling it resource here, as it could be confused with the resource-theoretic notion. As the reader will see later, defining what exactly should count as a resource (beyond resource theory) might still prove useful.

crete protocols that are capable of doing so. Once again, we find ourselves in the seemingly paradoxical situation of being able to achieve ground state cooling without any divergent quantity. To understand how this is possible and how this paradox can be resolved, consider the set of operations $U_k = e^{-iH_k t_k}$ $k = 1, \dots, N$ that make up the Landauer erasure protocol. Consecutive application of an infinite number of these operations allows for ground state cooling (with the addition of rethermalisation steps of the helper systems between each step, as detailed in Section IV.1). Alternatively, this concatenation of operations can be written as a single operation $U_{\text{tot}} = \lim_{N \rightarrow \infty} \prod_{k=1}^N U_k = e^{-iH_{\text{tot}} t_{\text{tot}}}$, which can be implemented in finite time t_{tot} , allowing for ground-state cooling in finite time and with finite energy. The key to the resolution of this lies in H_{tot} . While each individual H_k in the Landauer protocol may be of relatively simple form, creating the same result with a single combined Hamiltonian constitutes a highly complex operation, in fact it requires divergence of a quantity we call *control complexity*. This brings us to one of the primary outcomes of this research, the characterisation of *control complexity* as a quantity that (a) aligns with our intuitive understanding of complexity as “hard to implement”, and (b) adheres to the third law of thermodynamics, meaning that its divergence is essential for reaching the ground state (with finite time and energy dissipation).

We analyse this question in two distinct paradigms: The *coherent paradigm* and the *incoherent paradigm*. The setup for both of them consists of a *system* and a *machine* which interacts with the system, both of which are initially in a thermal state at the temperature of the environment. After each interaction the machine is decoupled from the system and rethermalises, such that it is reset to the initial state before again interacting with the system. We call these two steps the *interaction step* and the *rethermalisation step*. The key difference between the *coherent paradigm* and the *incoherent paradigm* lies in the way energy is made available for operations to be performed. In the *coherent paradigm* energy is provided by a work source, thus allowing for arbitrary global unitary operations on the *system* and the *machine*, which implicitly assumes access to a coherent energy source in the form of a *quantum battery* [66–69]. The *incoherent paradigm* constitutes the minimal control scenario. Here the only operation that is allowed is to activate an interaction for a certain amount of time. Therefore, the only control that is necessary is temporal control¹⁵. As this excludes any external work source, the interactions implemented this way are required to be energy-conserving. The machine must consist of two parts in this setting, one at thermal equilibrium with the environment and one at a higher temperature in thermal equilibrium with a hot bath, as thermal operations need to utilise the heat flow from the hot bath to the cold bath in order to perform any non-trivial op-

¹⁵Which comes with its own costs in terms of complexity and control, see Chapter III

erations. It should be noted that this paradigm, as opposed to the coherent one, constitutes a thermodynamically consistent framework in which all resources are precisely monitored and accounted for of. Thus it is the paradigm which allows us to draw fundamental conclusions about the achievability of ground-state cooling for given diverging quantities.

This brings us to the main results of this work. Before we can look at any of the actual cooling results we must determine question of how to quantify complexity. While on the one side the goal here is to recover the third law of thermodynamics – that is the necessity of divergence this quantity in order to cool to the ground state in finite time and at finite energy cost – on the other side we also want this quantity to represent our intuitive notion of complexity as “being hard to implement”. With the achievability of ground-state cooling in mind we find that the notion of complexity splits into two separate parts, each of which are required to diverge to achieve zero temperature: *Structural Complexity* and *Control Complexity*. *Structural Complexity* refers to the structure of the machine that is being used to cool the system. In itself it refers to two factors that play into the achievability of ground state cooling. The first represents a necessary condition for ground-state cooling, namely that the maximal energy gap of the machine diverges, and the second represents a necessary condition for achieving it at finite energy cost, namely that the machine must be infinite-dimensional. Given a machine that fulfills these requirements we turn our attention to requirements on the control of the machine. We find that ground-state cooling (for finite time and energy) can only be achieved if the unitary that implements the cooling protocol acts non-trivially on an infinite dimensional subspace of the Hilbert space of the machine. As a consequence its effective dimension, which is defined as $d_{\text{eff}} = \min \dim(\mathcal{A}) \text{ s.t. } U_{\mathcal{SM}} = U_{\mathcal{A}} \oplus \mathbb{1}_{\neg\mathcal{A}}$, must diverge¹⁶. Therefore, the effective dimension, while serving as an proxy of *Control Complexity*, is not a comprehensive measure of it. This becomes evident as we present explicit cases where its divergence does not equate to successful ground-state cooling.

With this proxy quantity identified our main results separate into the mentioned paradigms. In a first step we derive explicit protocols in the *coherent paradigm* that show the achievability of ground state cooling both for discrete and continuous variables for all cases where one quantity (energy, time or complexity) diverges while the others stay finite and minimal. In particular we show that it is possible to cool with diverging complexity in a single time step at the cost of the Landauer bound. Conversely we also show that if one of the other quantities diverges, ground-state cooling can be achieved with finite complexity. However, in the case of

¹⁶Intuitively the effective dimension tells us, for a sufficiently large machine dimension, how much of the machine takes part in the interaction [3]

continuous variables, *e.g.* where the system is modeled as a harmonic oscillator, the effective dimension in fact might diverge for operations that are easy to implement, such as Gaussian operations. In these cases we explicitly show that the protocols can be implemented with only Gaussian operations if time or energy diverge. In the case of diverging time we are able to show that non-Gaussian operations of effective dimension 1 also allow for ground state cooling at finite energy costs that are bigger than the Landauer bound. For an overview of this see Table I in Chapter IV. This prompts us to think of an alternative, more general notion of complexity that can incorporate more potential scenarios. In fact, a notion that applies universally to all our settings and fulfills the required criteria is the *energy gap variety*, which in essence quantifies the number of distinct energy gaps of the machine in a given energy interval. We find that the unitaries that achieve ground-state cooling at finite time and energy in any setting have to couple the system to an infinite variety of energy gaps that densely cover the range from the smallest energy gap of the machine ω_0 to infinity $[\omega_0, \infty)$. This suggests energy gap variety as the best candidate for the notion of *control complexity*.

However, the *coherent paradigm* presupposes access to a coherent energy source, thus potentially hiding resources that are required to prepare such a energy source, which makes it more difficult to draw fundamental conclusions about the notion of *control complexity*. Applying the same analysis to the *Incoherent Paradigm* we find that this assumption in fact drastically influences our understand of which resources are sufficient for ground-state cooling. In particular we find that perfect cooling, while still achievable for diverging time and complexity, no longer can be achieved alone by letting the energy diverge. Furthermore, in the case of time or complexity diverging Landauer's bound is no longer achievable in this paradigm. Instead we are required to incorporate the efficiency of the machine into the picture, leading to a minimal energy that needs to be drawn from the hot bath given by a bound we "name" the *Carnot-Landauer bound*

$$\Delta F_S^{(\beta)} + \eta \Delta E_H \leq 0, \quad (\text{I.2})$$

where $\Delta F_S^{(\beta)}$ is the change in free energy of the system, ΔE_H the energy change in the hot part of the machine and $\eta = 1 - \frac{\beta_H}{\beta}$ the Carnot efficiency of the machine.

While this shows that the incoherent and coherent settings are clearly different paradigms which are hard to compare, it also emphasises the importance of investigating the role that control plays for the achievability of thermodynamic tasks by using thermodynamically consistent analysis.

I.4 The emergence of measurements

Of all the postulates of quantum mechanics, the measurement postulate certainly is the most contentious. As mentioned previously, it stands in contradiction with quantum evolution given by the Schrödinger equation as well as the laws of thermodynamics. While it is a useful tool to derive probabilities, it cannot be taken at face value. It must be seen as a simplification of underlying processes that give rise to measurements. A thermodynamically consistent account of measurements thus requires explicit modelling and detailed analysis of these processes.

One of the first attempts to develop an explicit model of measurements was made as early as 1932 by John von Neumann, resulting in the well known von Neumann measurement scheme [70]. In it, measurements are modelled explicitly by letting a system interact with a probe that acts as a classical pointer which can be read out to reveal the measurement outcome. This measurement model later gave rise to the so called “standard measurement interactions” given by an interaction Hamiltonian of the form $H = X_S \otimes Y$, where X_S acts on the system and Y on the pointer. However, how exactly information is read out from the pointer, *i.e.* how observers themselves gain information about the system, remained an open problem. With the advent of decoherence theory [71, 72] and the subsequent development of Quantum Darwinism [73–75] this problem was confronted explicitly for the first time.

While decoherence theory focuses on properties of the system subject to interactions with the environment, Quantum Darwinism takes a somewhat opposite approach by focusing on the information about the system that is stored in the environment through these interactions and the accessibility of this information. In particular according to QD, for information about a system to be classically accessible, it must be redundantly encoded in the environment. Information encoded this way can be called *objective*, meaning that it can be accessed by multiple observers that have access to different parts of the environment such that they come to the same conclusions about the system without perturbing each others information about the system, it can be called *objective* [76].

As known from Landauer and Bennett, information cannot exist on its own, it needs to be stored in physical systems, in this case, quantum states. The existence of objective information thus must be reflected in a corresponding property of these quantum systems. When combined with the requirement of *Strong Independence*, which states that different parts of the environment should be uncorrelated except via their correlations with the system, it can be shown to be equivalent to a specific state structure known as the Spectrum Broadcast Structure (SBS) [76–78].

Although Quantum Darwinism and von-Neumann-type measurement schemes both seem

to tackle the same question, namely to explicitly model measurements, they, in fact, focus on different parts of this process. To understand this we can split the process of measurement into two parts: *Pre-Measurement* and *Objectification*. On one side, Pre-Measurement refers to highly controlled interactions of different degrees of freedom, for example system and pointer, similar to the model of von Neumann measurements. As a practical example one could think of a Stern-Gerlach apparatus where the interaction between spin (*i.e.* system) and momentum (*i.e.* pointer) degrees of freedom induced by the magnetic gradient creates correlations between these degrees of freedom. As mentioned before, however, this part alone does not fully describe the measurement process as it does not yet allow for objective information. On the other side objectification refers to the process by which the information in the pointer becomes classical, *i.e.* by which this information becomes accessible to observers. In the Stern-Gerlach example, this part of the measurement takes place when the particle hits the screen. Assuming *e.g.* that the screen is an Avalanche Photon Detector (APD), when a particle hits it this triggers a chain reaction eventually resulting in a macroscopic current of electrons in a wire that eventually signals to a computer screen that something has been detected, which then can be displayed such that the information is objectively available.

In contrast to Pre-Measurement, objectification must be seen as a process over which we have very little control, as it includes too many degrees of freedom to realistically have control over. In fact, the only control we can exert over this part of the measurement process is in form of the preparation of the initial conditions. An APD for example must be driven out of equilibrium into a metastable state to work, thus setting the initial conditions for the chain reaction to occur, while detailed control over every individual electron is of course far from achievable.

Knowing from thermodynamics that freely-evolving systems tend towards states of higher entropy, thus equilibrating, this lack of control suggests that objectivity, in fact, is the result of a spontaneous, entropy increasing process. This perspective prompts us to formulate a hypothesis we call the *Measurement-Equilibration Hypothesis*, stating that *any measurement must feature an entropy-increasing process towards an equilibrium state*.

From a unitary quantum mechanics standpoint, it may seem strange to reject the measurement update rule due to its irreversible nature and then introduce a notion of measurement that requires another irreversible process. At first glance, it appears that quantum mechanics is incompatible with thermodynamics. While thermodynamics predicts that the entropy of the universe as a whole increases monotonically, the unitary dynamics of quantum mechanics keeps the entropy of the universe constant. This apparent contradiction led to the development

of the field of closed-system equilibration, a sub-field of quantum thermodynamics [17–19], within which this problem has been approached from multiple angles. These approaches can be broadly divided into two groups: kinematic approaches [79, 80], which make probabilistic arguments about the state space related to the typicality of states in sub-spaces of this space, and dynamical approaches, which have brought forth two equivalent solutions by constraining the set of observables [33, 79], either through locality constraints or precision constraints¹⁷ [81, 82]. As our goal is to explicitly model measurements, that is, the process by which information becomes objective, dynamical models of equilibration provide us with the tools that are necessary to develop a first framework for this.

In Chapter II, we develop a framework for modelling measurements that focuses on the objectification aspect of measurements which operates under the premise of limited control, thereby invoking the Measurement-Equilibration Hypothesis. As was mentioned above, it is possible to identify a state structure that implies objectivity, *i.e.* Spectrum Broadcast Structure. This allows us to formalise the Measurement-Equilibration Hypothesis in the following way: Given an arbitrary system and environment, initially uncorrelated with each other, measurement occurs when the total system-environment state has come to an equilibrium state that is approximately of Spectrum Broadcast Structure.

To thoroughly examine the circumstances in which equilibrium is achieved, we will utilize a method of dynamic equilibration known as *equilibration on average*. Essentially, this approach recognises that it is not possible for a system to unitarily evolve into a time-independent equilibrium state unless it was initially in one. Instead, in order for a system to be in equilibrium it must spend the vast majority of its time in a state that is almost indistinguishable from some time-independent state. The emergence of such an indistinguishability is coupled to certain requirements on the initial state, the Hamiltonian, and the operators employed to distinguish the states: Given a non-degenerate, time-independent Hamiltonian H such that the time-evolved state of a system is given by $\rho(t) = e^{-iHt}\rho(0)e^{iHt}$, and an operator O with operator norm $\|O\|$, the following holds:

$$\langle |tr(\rho(t)O) - tr(\omega O)|^2 \rangle_\infty \leq \frac{\|O\|}{d_{\text{eff}}}, \quad (\text{I.3})$$

where $d_{\text{eff}} = \frac{1}{\sum_j (p_j)^2}$ is the effective dimension of the initial state and p_j are the diagonal entries of the initial state with respect to the energy eigenbasis [33]. Large effective dimension of the initial state compared to the operator norm of O ¹⁸, implies that the observable O does not allow

¹⁷This distinction resembles the idea of different types of control mentioned in the introduction, such as temporal and structural control.

¹⁸which essentially also means that state explores a large part of the state space over time. This is typically the

for a distinction between the equilibrium state and the evolving state for most points in time, thus making it indistinguishable from the equilibrium state for all practical purposes. Similar results have also been derived for sub-systems, and in terms of the trace distance [32]. With this methodology in mind we investigate the circumstances that have to be fulfilled such that system and environment equilibrate towards states that resemble Spectrum Broadcast Structure and thus exhibit objective properties. In particular, we are able to constrain the form of interactions which are capable of producing equilibrium states and show that, while it is in principle impossible to achieve Spectrum Broadcast Structure exactly, it is possible to approximate it via coarse-graining, *i.e.* by considering increasingly macroscopic observers.

This allows us to arrive at three main results. Starting with a weaker requirement on post-measurement states, being that the initial system-environment state equilibrates to a “classical-quantum state” [83] given by $\rho = \sum_{i=1}^{d_S} p_i |i\rangle\langle i|_S \otimes \rho_E^{(i)}$ rather than an objective state, we first show that the form of interaction necessary to achieve this is given by $H = \sum_i |i\rangle\langle i| \otimes H_E^{(i)}$, where $H_E^{(i)}$ is some Hermitian matrix acting on the environment. Note that, for any other form of interaction, “classical-quantum states” cannot be achieved, which implies that SBS cannot be achieved either as a result of equilibration. Furthermore we show that there exists no Hamiltonian that equilibrates to SBS exactly. For this, note that exact SBS requires the environment to unambiguously encode information about the system, which in turn implies zero overlap between states of the environment that encode different outcomes of the system, *i.e.* zero overlap between $\rho_E^{(i)}$ and $\rho_E^{(j)}$ for $i \neq j$. By placing a lower bound on the fidelity between these states we are able to show that there exists no Hamiltonian that equilibrates to such a state exactly.

The standard measurement Hamiltonian, *i.e.* a Hamiltonian of the form $H = X_S \otimes Y$, was shown to yield SBS under the assumption of precise temporal control [77, 84]. As mentioned, this Hamiltonian has found broad application throughout the literature as an interaction used to model measurements [85]. However, we show that the equilibrium states achievable by such Hamiltonians must be of the form $\rho_\infty = \rho_S \otimes \rho_E$, meaning that they don’t feature any correlations between environment and system. Such states cannot approximate objective states.

Finally, to understand how well equilibration can approximate SBS we construct a toy model for measurements by assuming a star-shaped interaction Hamiltonian of the form $H = \sum_i |i\rangle\langle i| \otimes \bigotimes_{k=1}^N c_k H_k^{(i)}$, which can be seen as a coupling of the system and the k -th environment with coupling strength c_k . Such an interaction yields an equilibrium state that is given by $\rho_\infty = \sum_i |i\rangle\langle i| \otimes \bigotimes_{k=1}^N \rho_k^{(i)}$. By grouping environment systems together into “macro-environments”

case in realistic scenarios [79]

N_q , such that $\rho_q = \bigotimes_{k \in N_q} \rho_k^i$ we show that the equilibrium states of such interactions exponentially tend towards objective states in the number of micro-environments per macro-environment. This provides an example of the emergence of Spectrum Broadcast Structure as a result of equilibration, proving that such states can in principle be achieved in this way. The requirement for coarse-graining combined with the Measurement-Equilibration Hypothesis highlights the role of the aforementioned notions of *temporal control* and *structural control* in this context. That is, observers are naturally constrained with respect to both of these kinds of control, which is what allows them to come to objective results in the first place.

This work should be seen as a first step towards such a model rather than a complete model for measurements. It provides the framework within which we can continue to develop our understanding of what constitutes a measurement. In Chapter II detailed analysis of the above results can be found.

Chapter II

Measurements and the emergence of Objectivity

II.1 Quantum measurements and equilibration: the emergence of objective reality via entropy maximisation

II.1.1 Contribution

I contributed significantly to the conception of the general theoretic framework of this research project. I was involved in the conception of all theorems and the calculation of all mathematical proofs thereof. In particular, I derived the proofs given in Appendix B and D and contributed to the proof in Appendix C. I wrote the first draft of the main text of the manuscript and took part in its revisions. I designed Figure 1.

Quantum measurements and equilibration: the emergence of objective reality via entropy maximisation

Emanuel Schwarzhans,^{1,2,*} Felix C. Binder,^{3,†} Marcus Huber,^{2,1,‡} and Maximilian P. E. Lock^{2,1,§}

¹*Institute for Quantum Optics and Quantum Information - IQOQI Vienna,
Austrian Academy of Sciences, Boltzmanngasse 3, 1090 Vienna, Austria*

²*Atominstitut, Technische Universität Wien, 1020 Vienna, Austria*

³*School of Physics, Trinity College Dublin, Dublin 2, Ireland*

(Dated: February 1, 2023)

Textbook quantum physics features two types of dynamics, reversible unitary dynamics and irreversible measurements. The latter stands in conflict with the laws of thermodynamics and has evoked debate on what actually constitutes a measurement. With the help of modern quantum statistical mechanics, we take the first step in formalising the hypothesis that quantum measurements are instead driven by the natural tendency of closed systems to maximize entropy, a notion that we call the Measurement-Equilibration Hypothesis. In this paradigm, we investigate how objective measurement outcomes can emerge within a purely unitary framework, and find that: (i) the interactions used in standard measurement models fail to spontaneously feature emergent objectivity and (ii) while ideal projective measurements are impossible, we can (for a given form of Hamiltonian) approximate them exponentially well as we collect more physical systems together into an “observer” system. We thus lay the groundwork for self-contained models of quantum measurement, proposing improvements to our simple scheme, such as by including pre-measurements and the phenomenon of pre-thermalisation.

Introduction.— The mathematical formulation of quantum theory includes two types of dynamics for a closed quantum system: unitary evolution, and measurement. The latter is often referred to as a collapse of the wavefunction, instantaneously updating the state of the system, and is at the heart of many interpretational discussions of quantum mechanics (e.g. the “measurement problem”). For projective measurements, this is an apparent contradiction of the laws of thermodynamics, as it does not conserve energy, allows for an arbitrary decrease in entropy [1], and allows for cooling to zero-temperature with finite resources, in violation of Nernst’s formulation of the 3rd law [2, 3]. In this work, we investigate a potential resolution to this apparent contradiction by modelling the process of measurement as the system of interest and its environment undergoing an equilibration process, thus being driven by an *increase* in entropy.

In the widely-used von Neumann measurement scheme [4], subsequently developed into what is sometimes referred to as the standard model of measurement [5], a carefully-timed interaction correlates the system of interest with some probe or pointer system. This allows, for example, an investigation of the restrictions placed by quantum theory on incompatible observables. Quantum Darwinism [6–8] broadens this model to include the environment, and thus the observer itself, highlighting the role of the redundant encoding of information about the system state in the measurement basis (the “pointer basis”) into the environment degrees of freedom, arguably removing the need for an explicit “Heisenberg cut”. Demanding that this redundancy be such that multiple observers (i.e. parts of the environment) independently agree on the system state, and that the only corre-

lations between them be their shared information about the system implies that the total state of the system and observers necessarily exhibits Spectrum Broadcast Structure [9–12].

The process of measurement, specifically the transformation from quantum indeterminacy to an objective fact, is irreversible, and occurs without precise control of all microscopic degrees of freedom of the system’s environment. A dynamical model of measurement must therefore explain how the system and its environment evolve spontaneously and irreversibly towards a state structure exhibiting objectivity. Since spontaneous and irreversible processes in nature are those associated with the process of equilibration, we hypothesise that any measurement is an entropy-increasing transition towards an equilibrium state, and refer to this as the Measurement-Equilibration Hypothesis. The notion that measurements correspond to an increase in entropy has been discussed at various points since the early years of quantum theory, e.g. [4, 13–15], but is absent from dynamical models of measurement. Recent advances in quantum statistical mechanics, specifically in the equilibration of closed quantum systems [16] allow us to address this absence, and as a first step in this direction, we seek to determine the circumstances under which the system and environment equilibrate to states exhibiting objective properties, including the Spectrum Broadcast Structure mentioned above.

In this work, we constrain the form of the interactions capable of producing such equilibrium states, and show that while it is impossible to reach this state structure exactly without violating the laws of thermodynamics, one can approach it exponentially as we increase the macro-

scopcity of “observers” for whom the measurement outcome should be objective. We find that the usual form of interactions employed in the standard model of measurement in fact does not result in an equilibrium state exhibiting objectivity, or even any correlations whatsoever between the system and its environment.

We begin by revisiting the concepts of objectivity and Spectrum Broadcast Structure and how they connect, as well as the role of the time-averaged state in equilibration phenomena. We then derive our main results and discuss the consequences of our findings as well as future directions and open questions.

Equilibration of quantum systems.— To investigate the conditions under which objectivity emerges as a result of equilibration, we will employ the notion of equilibration on average [16], whereby the equilibrium state is given by the infinite-time average. Given a system evolving as $\rho(t)$, the infinite-time average is given by

$$\rho_\infty := \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \rho(t) dt. \quad (1)$$

This is the unique state which maximises the von Neumann entropy given all conserved quantities [16, 17]. To illustrate the sense in which states equilibrate to their time average, let us review two examples of equilibration on average.

First, consider a system evolving unitarily via some time-independent Hamiltonian H , so that $\rho(t) = U(t)\rho(0)U^\dagger(t)$ with $U(t) = e^{iHt}$, and denote by $\{p_j\}_j$ the diagonal elements of $\rho(t)$ in the eigenbasis of H . The effective dimension – sometimes called the inverse participation ratio – is then defined as $d_{\text{eff}}[\rho(t)] := (\sum_j p_j^2)^{-1}$ (note the lack of time-dependence). Assuming that no two energy gaps of H are equal in magnitude (which also implies that H is non-degenerate), then the average squared difference between $\text{Tr}[\rho(t)O]$ and $\text{Tr}[\rho_\infty O]$ is bounded above by $\|O\|/d_{\text{eff}}[\rho(t)]$ for any bounded operator O [18]. As a second example, consider $\rho(t)$ to be some subsystem of a larger, unitarily evolving system, *i.e.* $\rho(t) = \text{Tr}_X[U(t)\tilde{\rho}(0)U^\dagger(t)]$ for some $\tilde{\rho}(0)$ and tensor factor X , and again assume that no two energy gaps of H are equal. In this case, the average trace distance between $\rho(t)$ and $\text{Tr}_X[\tilde{\rho}_\infty]$ is bounded by $\frac{1}{2}\sqrt{d^2/d_{\text{eff}}[\tilde{\rho}(t)]}$, where d is the dimensionality of the Hilbert space to which $\rho(t)$ belongs [19].

We note three issues for both of these examples to be considered equilibration. First, both rely on an effective dimension $d_{\text{eff}}[\rho(t)]$ being in some sense large. This quantity measures the number of significantly occupied energy levels participating in the evolution of $\rho(t)$, and its typically large size for large systems can be argued for on both physical and mathematical grounds. Second, there is the necessity for some finite equilibration time, and third, the possibility of Poincaré recurrences. All three of these issues, as well as generalisations beyond

the condition of no equal energy gaps, are addressed in detail in [16].

Objectivity and quantum states.— Let $\mathcal{H}_S$ denote the Hilbert space corresponding to some d_S -dimensional system of interest, and let $\{\mathcal{H}_k\}$ with $k = 1, \dots, N$ denote the Hilbert spaces corresponding to N systems which form the environment of the system of interest, and which we will associate with observers, for whom we seek objectivity. Note that we do not *a priori* impose any criteria for what is to be considered as an observer. In order for a fact about the system to be considered objective (and thus classical), multiple observers must be able to discover it in such a way that they agree with each other and that none of them changes the fact through their observation. As noted above, formalising this notion in the context of quantum theory [20] and combining it with the requirement that observers be uncorrelated except via the system (“strong independence”), is equivalent to demanding that the state of the system and the observers exhibits Spectrum Broadcast Structure [10]. A state has Spectrum Broadcast Structure with respect to a pointer basis $\{|i\rangle\}_i$, where $i = 1, \dots, d_S$, if it is of the form

$$\rho_{\text{SBS}} = \sum_{i=1}^{d_S} p_i |i\rangle\langle i|_S \bigotimes_{k=1}^N \rho_k^{(i)}, \quad (2)$$

where p_i is the probability associated with outcome $|i\rangle$, and $\rho_k^{(i)}$ is the state on $\mathcal{H}_k$ conditioned on that outcome, and where $\rho_k^{(i)} \rho_k^{(j)} = 0$ for $i \neq j$. As in [21], we will find it more convenient to express this orthogonality in terms of the fidelity, namely $F(\rho_k^{(i)}, \rho_k^{(j)}) = 0$ for $i \neq j$, where $F(\sigma, \rho) := \text{Tr} \left(\sqrt{\sqrt{\sigma} \rho \sqrt{\sigma}} \right)^2$.

The structure in Eq. (2) thus represents a state where quantum information has become redundantly encoded, and through this can be considered classical [22]. Note, however, that this is only true in the pointer basis, in accordance with no-cloning theorem [23]. This state structure holds if and only if the conditions of Quantum Darwinism and Strong Independence are met [11]. A projective measurement can then be understood as a post-selection performed on the process by which an initially uncorrelated state between system and observers tends to a state with Spectrum Broadcast Structure in the measurement basis. Consequently, to say that a projective measurement has been performed, one must first arrive at a state structure exhibiting objectivity.

Measurement as an equilibration process.— With this framework, we can now state precisely what it means to demand that objectivity emerge as a consequence of equilibration, and thus the maximization of entropy. To wit, we assume that system and observers are initially uncorrelated and seek the conditions under which they equilibrate to a state of objectivity, specifically when $\rho_\infty = \rho_{\text{SBS}}$. Note that this demand does not restrict the observer states $\rho_k^{(i)}$ in Eq. (2), beyond the orthogonality

requirement described above. As above, we denote the measurement basis on $\mathcal{H}_S$ by $\{|i\rangle\}_i$, and the outcome probabilities by $\{p_i\}$, *i.e.* $p_i = \langle i | \rho_{S,0} | i \rangle$, where $\rho_{S,0}$ is the initial state of the system.

Note that the infinite-time average is a pinching map with respect to the Hamiltonian generating the evolution [17], *i.e.*

$$\rho_\infty = \sum_E \Pi_E \rho(0) \Pi_E, \quad (3)$$

where $H = \sum_E E \Pi_E$ is the spectral decomposition of H , and we therefore have $[H, \rho_\infty] = 0$. Consequently, to equilibrate to a state with a given structure, the Hamiltonian must commute with all states of that structure. In particular, for Spectrum Broadcast Structure, we must have $[H, \rho_{\text{SBS}}] = 0$, which can be interpreted as an extension of the commutation relation defining the pointer basis [24] to the observers' systems. As we discuss below, no such Hamiltonian exists, which precludes exact equilibration to states with this structure.

Exact projective measurements are impossible.— Let us first examine a weaker requirement than that of Spectrum Broadcast Structure. We write the initial state $\rho(0) = \rho_{S,0} \otimes \rho_{E,0}$, where $\rho_{E,0}$ is the state on $\bigotimes_{k=1}^N \mathcal{H}_k$, and we demand a “classical-quantum” [25] equilibrium state between the system and the observers, *i.e.* $\rho_\infty = \sum_{i=1}^{d_S} p_i |i\rangle\langle i|_S \otimes \rho_E^{(i)}$, where the $\{\rho_E^{(i)}\}_i$ are some density matrices on $\bigotimes_{k=1}^N \mathcal{H}_k$. Such a state may violate the Strong Independence condition. As shown in Appendix B, this requirement implies that H takes the form

$$H = \sum_i |i\rangle\langle i|_S \otimes H_E^{(i)}, \quad (4)$$

for some Hermitian matrices $\{H_E^{(i)}\}_i$ on $\bigotimes_{k=1}^N \mathcal{H}_k$. Note that Eq. (4) can only be considered as a sum of a free system Hamiltonian H_S and some interaction H_{int} in the case where the pointer basis consists of eigenstates of H_S . For measurement in any other basis, it corresponds to an interaction term alone, necessitating the so-called measurement limit, often used in dynamical models of measurement, where free terms in the Hamiltonian are neglected *i.e.* $H \approx H_{\text{int}}$ (see *e.g.* [5]).

When we now demand, in addition to the above, that the observers' states unambiguously encode the system state, *i.e.* that $F(\rho_E^{(i)}, \rho_E^{(j)}) = 0$ for $i \neq j$, we find that this only holds when $\rho_E^{(i)} = 0 \ \forall i$ (see Appendix B for proof). In particular, $F(\rho_E^{(i)}, \rho_E^{(j)})$ is bounded below by $1/d_E^2$, where d_E is the dimensionality of $\bigotimes_{k=1}^N \mathcal{H}_k$. Consequently, it is impossible for an initially uncorrelated state to equilibrate to a state with the properties listed above. This precludes equilibration to a state exhibiting objectivity, as defined in [10]. The stronger condition of Spectrum Broadcast Structure (which combines objectivity with Strong Independence) is then likewise

impossible. In [1], the authors found that it was possible to perform an ideal projective measurement given access to a zero-temperature pointer state, but here we see that, adopting the Measurement-Equilibration Hypothesis, it is impossible even given access to such a resource. That the lower bound of $F(\rho_k^{(i)}, \rho_k^{(j)})$ decreases with dimensionality however, does allow the approximation of approximate an objective state as the size of the environment increases, as we will discuss later.

Equilibration and the standard model of measurement.— As discussed above, dynamical models of measurement usually assume a degree of temporal control such that a particular unitary transformation can be performed on the system and some apparatus/environment, generated the product of observables acting on the system and the apparatus [5]. For example, one such model was used to investigate the transition to Spectrum Broadcast Structure in [21]. In our context, let us denote such a Hamiltonian by $H = X_S \otimes Y$, where X_S is any operator whose eigenbasis is the pointer basis, and where Y is an observable on $\bigotimes_{k=1}^N \mathcal{H}_k$. In this case, as shown in Appendix C, the equilibrium state is given by

$$\rho_\infty = \sum_{i=1}^{d_S} p_i |i\rangle\langle i|_S \otimes \rho_E, \quad (5)$$

for some state ρ_E on $\bigotimes_{k=1}^N \mathcal{H}_k$. This highlights the distinction between the standard model of measurement and the equilibration paradigm; given temporal control, one can of course generate the desired correlations between the system and the observers, but considering a spontaneous and uncontrolled process, one finds that on average there are no correlations whatsoever between the system and the observers.

Approximating projective measurements.— While exact equilibration to a state with Spectrum Broadcast Structure is impossible, there remains the question of how well an equilibrium state can approximate this structure. We saw earlier that, with increasing dimension of $\bigotimes_{k=1}^N \mathcal{H}_k$, one can approach the desired orthogonality condition required for objectivity (though recall that this was a weaker condition than Spectrum Broadcast Structure). In the same vein, we will now see that by collecting the observers' systems together into composite systems, which we call macro-observers, one can approach Spectrum Broadcast Structure exponentially as the size of these macro-observers increases.

As discussed above, equilibration to a state exhibiting the appropriate statistics in the pointer basis, with no coherences between different pointer states, and allowing correlations between the observers and the system in this basis requires a conditional structure of the Hamiltonian given in Eq. (4). With the Strong Independence condition in mind, which demands that the observers should be uncorrelated except via their correlations with the sys-

tem, we expand this to consider an interaction of the form

$$H = \sum_i |i\rangle\langle i|_S \otimes \sum_{k=1}^N c_k H_k^{(i)}. \quad (6)$$

We can see this structure as coupling k to S with coupling strength c_k , such that for the system state i , the action $H_k^{(i)}$ is performed. Assuming H in Eq. (6) to be non-degenerate and starting with the state $\rho(0) = \rho_{S,0} \otimes \bigotimes_{k=1}^N \rho_{k,0}$, we arrive at the equilibrium state

$$\rho_\infty = \sum_{i=1}^{d_S} p_i |i\rangle\langle i|_S \bigotimes_{k=1}^N \rho_k^{(i)}, \quad (7)$$

where $\rho_k^{(i)}$ are obtained from $\rho_{k,0}$ by applying the pinching map with respect to $H_k^{(i)}$. This state has the form required for Spectrum Broadcast Structure, but it remains to check the orthogonality condition. To that end, let $|E_{m,k}^{(i)}\rangle$ denote the eigenvector of $H_k^{(i)}$ corresponding to the eigenvalue $E_{m,k}^{(i)}$. We will find it useful to define the quantity

$$\eta_k^{(ij)} := \sum_m \sqrt{\langle E_{m,k}^{(i)} | \rho_k^{(i)} \rho_k^{(j)} | E_{m,k}^{(i)} \rangle}. \quad (8)$$




Figure 1. This is a schematic representation of the conditional Hamiltonian given in Eq. 6. It depicts the interaction of system $\mathcal{H}_S$ (in yellow) with individual parts of the environment $\mathcal{H}_k$ (in red). The coupling strength between the system and each environment is indicated by the dashed lines and the labels c_k . The light-red frames exemplify ways of coarse-graining the environments into macro-environments N_k .

Now, let $M < N$ denote the number of macro-observers, and let N_q denote the set of k labels to

be collected together to form a given macro-observer q (see Fig. 1). We can then write Eq. (7) in terms of macro-observers as $\rho_\infty = \sum_{i=1}^{d_S} p_i |i\rangle\langle i|_S \bigotimes_{q=1}^M \tilde{\rho}_q^{(i)}$, where $\{\tilde{\rho}_q^{(i)} := \bigotimes_{k \in N_q} \rho_k^{(i)}\}$ are states on $\bigotimes_{k \in N_q} \mathcal{H}_k$. As we show in Appendix D, the fidelity between the conditional states of the macro-observers satisfies

$$F(\tilde{\rho}_q^{(i)}, \tilde{\rho}_q^{(j)}) \leq e^{-\gamma_q^{(ij)} |N_q^{(ij)}|}, \quad (9)$$

where

$$\gamma_q^{(ij)} := -2 \ln \left(\max_{k \in N_q} \eta_k^{(ij)} \right) \geq 0, \quad (10)$$

and $|N_q^{(ij)}|$ is proportional to $|N_q|$ under certain non conspiracy assumptions (see Appendix D). Thus as we increase each $|N_q|$, *i.e.* the the number of degrees of freedom collected together to form each macro-observer q , the equilibrium state tends exponentially towards one with Spectrum Broadcast Structure for the macro-observers. In other words, as we coarse-grain our view of observers towards ever-larger collections of microscopic systems, we tend exponentially towards a state of objectivity commensurate with the occurrence of a measurement; macroscopic observers will see objective facts about the system.

Discussion and conclusions.— Quantum measurement is often described as the ‘collapse of the wavefunction’ when a quantum system makes contact with an observer. One is then led to the question, what exactly can constitute an observer? [26] The Measurement-Equilibration Hypothesis allows us to constrain this question, giving the answer that an observer is whatever system interacts with the system of interest such that the result equilibrates to a state that approximates objective state structure asymptotically. In this paradigm, we have shown the impossibility of exact projective measurements (in accordance with known results [1]). However, they can be approximated exponentially as more and more microscopic systems are incorporated into a given macroscopic observer. This accords with the intuition that small systems should not be considered observers [26], but that larger ones can be.

In the simple version of an equilibrating measurement presented above, the Strong Independence condition on the equilibrium state led to the demand that observers interact via a Hamiltonian with a conditional form, and surprisingly, that the particular conditional form given by the standard model of measurement does not satisfy the most minimal requirement for measurement, namely correlation of observer and system. More broadly, the conditional form imposed by Strong Independence does not resemble the usual forms of interaction found in nature, suggesting that a the model needs to be expanded if the Measurement-Equilibration Hypothesis is to be employed in the modelling of real-world experiments.

In this regard, two immediate generalisations of our work may suffice. First, one might allow for vanishingly small correlations violating the Strong Independence condition. It seems likely that the requirement that such violations be vanishingly small will nonetheless strongly constrain the form of the interaction. A second, and in our view more promising generalisation, would be the inclusion of an initial “pre-measurement” step, whereby one allows some time-controlled interaction to correlate an apparatus with the system of interest, whose subsequent equilibration with some observers leads to a state of objectivity between the system of interest and the observers; in this case the requirement that the system of interest be undisturbed in the measurement basis would be less restrictive in the equilibrium step. Indeed many measurements modelled via projection operators do not leave the system in an eigenstate corresponding to the measurement outcome, *e.g.* photon detection.

Beyond simply improving the model, our work provides a new thermodynamically consistent framework for tackling questions regarding the role of the driving/preparation of measurement apparatuses and the characterisation of the necessary out-of-equilibrium resources. This is reinforced by the work of [27], showing that thermal equilibrium and objectivity is often at odds. Given such investigations, one may go beyond the simple statement of the Measurement-Equilibration Hypothesis above, and investigate whether the quality of the measurement depends on the amount of entropy dissipated in the process. Furthermore, this framework allows us to analyse the problem of instantaneous measurements [28] by analysing the limitations of such dissipative processes.

While ours is only the first step towards a thermodynamically self-contained model for quantum measurements, we hope that this will resolve many of the foundational issues that manifest through the apparent inconsistency of text-book postulates of quantum measurements and thermodynamic laws.

-
- [1] Yelena Guryanova, Nicolai Friis, and Marcus Huber, *Ideal projective measurements have infinite resource costs*, [Quantum](#) **4**, 222 (2020).
 - [2] Walter Nernst, *über die Beziehungen zwischen Wärmeerzeugung und maximaler Arbeit bei kondensierten Systemen*. in *Sitzungsberichte Der Königlich Preussischen Akademie Der Wissenschaften* (Berlin, 1906) pp. 933–940.
 - [3] 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, *Landauer vs. nernst: What is the true cost of cooling a quantum system?* (2021).
 - [4] John Von Neumann, *Mathematische Grundlagen der Quantenmechanik* (Springer Berlin, 1930).
 - [5] Paul Busch and Pekka J. Lahti, *The standard model of quantum measurement theory: History and applications*, [Foundations of Physics](#) **26**, 875 (1996).
 - [6] Wojciech Hubert Zurek, *Decoherence, einselection, and the quantum origins of the classical*, [Reviews of modern physics](#) **75**, 715 (2003).
 - [7] Wojciech Hubert Zurek, *Decoherence and the transition from quantum to classical – revisited*, in [Quantum Decoherence](#) (Springer, 2006) pp. 1–31.
 - [8] Wojciech Hubert Zurek, *Quantum darwinism*, [Nature physics](#) **5**, 181 (2009).
 - [9] JK Korbicz, Paweł Horodecki, and Ryszard Horodecki, *Objectivity in a noisy photonic environment through quantum state information broadcasting*, [Physical review letters](#) **112**, 120402 (2014).
 - [10] R. Horodecki, J. K. Korbicz, and P. Horodecki, *Quantum origins of objectivity*, [Physical Review A](#) **91** (2015), [10.1103/physreva.91.032122](#).
 - [11] Thao P. Le and Alexandra Olaya-Castro, *Strong quantum darwinism and strong independence are equivalent to spectrum broadcast structure*, [Physical Review Letters](#) **122** (2019), [10.1103/physrevlett.122.010403](#).
 - [12] J. K. Korbicz, *Roads to objectivity: Quantum darwinism, spectrum broadcast structures, and strong quantum darwinism—a review*, [Quantum](#) **5**, 571 (2021).
 - [13] Leo Szilard, *Über die entropieverminderung in einem thermodynamischen system bei eingriffen intelligenter wesen*, [Zeitschrift für Physik](#) **53**, 840 (1929).
 - [14] Baidyawath Misra, Ilya Prigogine, and Maurice Courbage, *Lyapounov variable: Entropy and measurement in quantum mechanics*, [Proceedings of the National Academy of Sciences](#) **76**, 4768 (1979).
 - [15] Adriana Daneri, Angelo Loinger, and Giovanni Maria Prosperi, *Quantum theory of measurement and ergodicity conditions*, [Nuclear physics](#) **33**, 297 (1962).
 - [16] Christian Gogolin and Jens Eisert, *Equilibration, thermalisation, and the emergence of statistical mechanics in closed quantum systems*, [Reports on Progress in Physics](#) **79**, 056001 (2016).
 - [17] Christian Gogolin, Markus P Müller, and Jens Eisert, *Absence of thermalization in nonintegrable systems*, [Physical review letters](#) **106**, 040401 (2011).
 - [18] Peter Reimann, *Foundation of statistical mechanics under experimentally realistic conditions*, [Physical review letters](#) **101**, 190403 (2008).
 - [19] Noah Linden, Sandu Popescu, Anthony J Short, and Andreas Winter, *Quantum mechanical evolution towards thermal equilibrium*, [Physical Review E](#) **79**, 061103 (2009).
 - [20] Harold Ollivier, David Poulin, and Wojciech H Zurek, *Objective properties from subjective quantum states: Environment as a witness*, [Physical review letters](#) **93**, 220401 (2004).
 - [21] J Tuziemiński and J K Korbicz, *Analytical studies of spectrum broadcast structures in quantum brownian motion*, [Journal of Physics A: Mathematical and Theoretical](#) **49**, 445301 (2016).
 - [22] Robin Blume-Kohout and Wojciech H Zurek, *Quantum darwinism: Entanglement, branches, and the emergent classicality of redundantly stored quantum information*, [Physical review A](#) **73**, 062310 (2006).
 - [23] W K Wootters and W H Zurek, *Nature*, [Tech. Rep.](#) 9 (1982).

- [24] Wojciech H Zurek, *Pointer basis of quantum apparatus: Into what mixture does the wave packet collapse?* [Physical review D](#) **24**, 1516 (1981).
- [25] Igor Devetak and Andreas Winter, *Distilling common randomness from bipartite quantum states*, [IEEE Transactions on Information Theory](#) **50**, 3183 (2004).
- [26] Caslav Brukner, *On the quantum measurement problem*, (2015), number: arXiv:1507.05255 [arXiv:1507.05255 \[quant-ph\]](#).
- [27] Thao P. Le, Andreas Winter, and Gerardo Adesso, *Thermality versus objectivity: can they peacefully coexist?* (2021), [arXiv:2109.13265 \[quant-ph\]](#).
- [28] Philipp Strasberg, Kavan Modi, and Michalis Skotiniotis, *How long does it take to perform a projective measurement?* [arXiv preprint arXiv:2011.13254](#) (2020).
- [29] P. Mironowicz, J. K. Korbicz, and P. Horodecki, *Monitoring of the Process of System Information Broadcasting in Time*, [Physical Review Letters](#) **118**, 150501 (2017).
- [30] P. Mironowicz, P. Nalezyty, P. Horodecki, and J. K. Korbicz, *System information propagation for composite structures*, [Physical Review A](#) **98** (2018), 10.1103/PhysRevA.98.022124.
- [31] Alexandre Feller, Benjamin Roussel, Irénée Frérot, and Pascal Degiovanni, *Comment on “strong quantum darwinism and strong independence are equivalent to spectrum broadcast structure”*, [Physical Review Letters](#) **126**, 188901 (2021).
- [32] Cyril Elouard, David Herrera-Martí, Benjamin Huard, and Alexia Auffeves, *Extracting work from quantum measurement in maxwell’s demon engines*, [Physical Review Letters](#) **118**, 260603 (2017).
- [33] Cyril Elouard and Andrew N Jordan, *Efficient quantum measurement engines*, [Physical review letters](#) **120**, 260601 (2018).
- [34] Maximilian Schlosshauer, *Decoherence, the measurement problem, and interpretations of quantum mechanics*, [Reviews of Modern physics](#) **76**, 1267 (2005).
- [35] Barış Çakmak, Özgür E Müstecaplıoğlu, Mauro Paternostro, Bassano Vacchini, and Steve Campbell, *Quantum darwinism in a composite system: Objectivity versus classicality*, [Entropy](#) **23**, 995 (2021).
- [36] C Jess Riedel, Wojciech H Zurek, and Michael Zvolak, *The rise and fall of redundancy in decoherence and quantum darwinism*, [New Journal of physics](#) **14**, 083010 (2012).
- [37] Loïc Herviou, Jens H Bardarson, and Nicolas Regnault, *Many-body localization in a fragmented hilbert space*, [Physical Review B](#) **103**, 134207 (2021).
- [38] Maron Anka, Thiago Rodrigues de Oliveira, and Daniel Jonathan, *Measurement based quantum heat engine in a three-level system*, [arXiv preprint arXiv:2108.02229](#) (2021).
- [39] Armen E Allahverdyan, Roger Balian, and Theo M Nieuwenhuizen, *Understanding quantum measurement from the solution of dynamical models*, [Physics Reports](#) **525**, 1 (2013).
- [40] Wojciech H Zurek, *Environment-induced superselection rules*, [Physical review D](#) **26**, 1862 (1982).
- [41] Rafael D Sorkin, *Impossible measurements on quantum fields*, in *Directions in general relativity: Proceedings of the 1993 International Symposium, Maryland, Vol. 2* (1993) pp. 293–305.
- [42] Marek Żukowski and Marcin Markiewicz, *Physics and metaphysics of wigner’s friends: Even performed premeasurements have no results*, [Physical Review Letters](#) **126**, 130402 (2021).
- [43] Salvatore Lorenzo, Mauro Paternostro, and G Massimo Palma, *Anti-zeno-based dynamical control of the unfolding of quantum darwinism*, [Physical Review Research](#) **2**, 013164 (2020).
- [44] Katarzyna Roszak and Jarosław K Korbicz, *Entanglement and objectivity in pure dephasing models*, [Physical Review A](#) **100**, 062127 (2019).
- [45] Steve Campbell, Barış Çakmak, Özgür E Müstecaplıoğlu, Mauro Paternostro, and Bassano Vacchini, *Collisional unfolding of quantum darwinism*, [Physical Review A](#) **99**, 042103 (2019).
- [46] Carlo Maria Scandolo, Roberto Salazar, Jarosław K Korbicz, and Paweł Horodecki, *Universal structure of objective states in all fundamental causal theories*, [Physical Review Research](#) **3**, 033148 (2021).
- [47] Jarosław K Korbicz, Edgar A Aguilar, Piotr Ówikliński, and P Horodecki, *Generic appearance of objective results in quantum measurements*, [Physical Review A](#) **96**, 032124 (2017).
- [48] Thomas K Uden, Daniel Louzon, Michael Zvolak, Wojciech H Zurek, and Fedor Jelezko, *Revealing the emergence of classicality using nitrogen-vacancy centers*, [Physical review letters](#) **123**, 140402 (2019).
- [49] Ming-Cheng Chen, Han-Sen Zhong, Yuan Li, Dian Wu, Xi-Lin Wang, Li Li, Nai-Le Liu, Chao-Yang Lu, and Jian-Wei Pan, *Emergence of classical objectivity of quantum darwinism in a photonic quantum simulator*, [Science Bulletin](#) **64**, 580 (2019).
- [50] Mario A Ciampini, Giorgia Pinna, Paolo Mataloni, and Mauro Paternostro, *Experimental signature of quantum darwinism in photonic cluster states*, [Physical Review A](#) **98**, 020101 (2018).
- [51] Eoin O’Connor, Giacomo Guarnieri, and Steve Campbell, *Action quantum speed limits*, (2020), [arXiv:2011.05232 \[quant-ph\]](#).
- [52] Sebastian Deffner, *Quantum speed limits and the maximal rate of information production*, [Phys. Rev. Research](#) **2**, 013161 (2020).
- [53] Sebastian Deffner and Eric Lutz, *Quantum speed limit for non-markovian dynamics*, [Physical Review Letters](#) **111** (2013), 10.1103/physrevlett.111.010402.
- [54] Sebastian Deffner and Steve Campbell, *Quantum speed limits: from heisenberg’s uncertainty principle to optimal quantum control*, [Journal of Physics A: Mathematical and Theoretical](#) **50**, 453001 (2017).
- [55] Philipp Kammerlander and Janet Anders, *Coherence and measurement in quantum thermodynamics*, [Scientific reports](#) **6**, 1 (2016).
- [56] Jordan M Horowitz and Massimiliano Esposito, *Thermodynamics with continuous information flow*, [Physical Review X](#) **4**, 031015 (2014).
- [57] Wojciech H Zurek, *Reduction of the wavepacket: How long does it take?* in [Frontiers of Nonequilibrium Statistical Physics](#) (Springer, 1986) pp. 145–149.
- [58] Giuseppe Vitagliano, Claude Klöckl, Marcus Huber, and Nicolai Friis, *Trade-off between work and correlations in quantum thermodynamics*, in [Thermodynamics in the Quantum Regime](#) (Springer, 2018) pp. 731–750.
- [59] Felix C. Binder, Sai Vinjanampathy, Kavan Modi, and John Goold, *Quantacell : powerful charging of quantum*

- batteries, *New Journal of Physics* **17**, 75015 (2015).
- [60] Felix C. Binder, *Work, Heat, and Power of Quantum Processes*, Ph.D. thesis, University of Oxford (2016).
- [61] Francesco Campaioli, Felix A. Pollock, Felix C. Binder, Lucas Céleri, John Goold, Sai Vinjanampathy, and Kavan Modi, *Enhancing the Charging Power of Quantum Batteries*, *Physical Review Letters* **118**, 150601 (2017), [arXiv:1612.04991](#).
- [62] Francesco Campaioli, *Tightening Time-Energy Uncertainty Relations*, Ph.D. thesis (2020), [arXiv:2004.08384](#).
- [63] Francesco Campaioli, Felix A. Pollock, Felix C. Binder, and Kavan Modi, *Tightening Quantum Speed Limits for Almost All States*, *Physical Review Letters* **120**, 060409 (2018), [arXiv:1709.08941](#).
- [64] Gerhard C. Hegerfeldt, *High-speed driving of a two-level system*, *Physical Review A - Atomic, Molecular, and Optical Physics* **90**, 1 (2014), [arXiv:1407.6502](#).
- [65] Gerhard C. Hegerfeldt, *Driving at the quantum speed limit: Optimal control of a two-level system*, *Physical Review Letters* **111**, 1 (2013), [arXiv:1305.6403](#).
- [66] R. Blume-Kohout and W.H. Zurek, *A Simple Example of “Quantum Darwinism”: Redundant Information Storage in Many-Spin Environments*, *Found Phys* **35**, 1857–1876 (2005).
- [67] Katarzyna Macieszczak, Dominic C. Rose, Igor Lesanovsky, and Juan P. Garrahan, *Theory of classical metastability in open quantum systems*, (2020), [arXiv:2006.01227 \[cond-mat.stat-mech\]](#).
- [68] P F Romero, *Equilibration and typicality in quantum processes*, (2021).
- [69] Fernando G. S. L. Brandão, Marco Piani, and Paweł Horodecki, *Generic emergence of classical features in quantum darwinism*, *Nature Communications* **6** (2015), 10.1038/ncomms8908.
- [70] Paul A. Knott, Tommaso Tufarelli, Marco Piani, and Gerardo Adesso, *Generic emergence of objectivity of observables in infinite dimensions*, *Physical Review Letters* **121** (2018), 10.1103/physrevlett.121.160401.
- [71] Mateusz Kiciński and Jarosław K. Korbicz, *Decoherence and objectivity in higher spin environments*, (2021), [arXiv:2105.09093 \[quant-ph\]](#).
- [72] Thao P Le and Alexandra Olaya-Castro, *Witnessing non-objectivity in the framework of strong quantum darwinism*, *Quantum Science and Technology* **5**, 045012 (2020).
- [73] Anthony J Short and Terence C Farrelly, *Quantum equilibration in finite time*, *New Journal of Physics* **14**, 013063 (2012).
- [74] Fernando Galve, Roberta Zambrini, and Sabrina Maniscalco, *Non-Markovianity hinders Quantum Darwinism*, *Scientific Reports* **6**, 1 (2016).
- [75] Salvatore Lorenzo, Mauro Paternostro, and G. Massimo Palma, *Reading a Qubit Quantum State with a Quantum Meter: Time Unfolding of Quantum Darwinism and Quantum Information Flux*, *Open Systems and Information Dynamics* **26**, 1 (2019), [arXiv:2001.11558](#).
- [76] Nadia Milazzo, Salvatore Lorenzo, Mauro Paternostro, and G. Massimo Palma, *Role of information backflow in the emergence of quantum darwinism*, *Physical Review A* **100**, 1 (2019), [arXiv:1901.05826](#).
- [77] Fabio Anza, Christian Gogolin, and Marcus Huber, *Eigenstate thermalization for degenerate observables*, *Physical Review Letters* **120** (2018), 10.1103/physrevlett.120.150603.
- [78] Dmitry A. Abanin and Zlatko Papić, *Recent progress in many-body localization*, *Annalen der Physik* **529**, 1700169 (2017).
- [79] Richard Jozsa, *Fidelity for Mixed Quantum States*, *Journal of Modern Optics* **41**, 2315 (1994).
- [80] Marco Tomamichel, *Quantum Information Processing with Finite Resources – Mathematical Foundations*, Vol. 5 (2016) [arXiv:1504.00233 \[math-ph, physics:quant-ph\]](#).
- [81] John Watrous, *The Theory of Quantum Information*, 1st ed. (Cambridge University Press, 2018)

APPENDICES

Appendix A: Setting and notation

In the following appendices, we prove the results stated in the main text. First, for ease of reference, we recall the mathematical setting as well as some notation.

We assume a Hilbert space structure given by $\mathcal{H} = \mathcal{H}_S \otimes \mathcal{H}_E = \mathcal{H}_S \otimes_{k=1}^N \mathcal{H}_k$, where we denote the system that is being measured by the subscript S, the total environment by the subscript E and the observers’ systems which constitute the environment by k . The Hamiltonian of the system will be denoted by H , and the measurement (*i.e.* “pointer”) basis by $\{|i\rangle\}_i$. The probability of measurement outcome i is then given by

$$p_i := \text{Tr}[\rho_{S,0}|i\rangle\langle i|], \quad (\text{A1})$$

where $\rho_{S,0}$ is the initial state on $\mathcal{H}_S$. The time average of a quantity $A(t)$ is denoted by

$$\langle A(t) \rangle_T := \frac{1}{T} \int_0^T A(t) dt.$$

The equilibrium, *i.e.* infinite-time average state is therefore given by

$$\rho_\infty = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T e^{-iHt} \rho(0) e^{iHt} dt.$$

For simplicity, when an operator acts non-trivially only on a given tensor factor, we will suppress tensor products with the identity operator, *e.g.* writing $\bigotimes_{k=1}^{j-1} \mathbb{1}_k \otimes H_j \otimes \bigotimes_{k'=j+1}^N \mathbb{1}_{k'}$ as simply H_j .

We recall the definition of Spectrum Broadcast Structure: a state has this structure with respect to the measurement basis $\{|i\rangle\langle i|\}_i$ if it is of the form

$$\rho_{\text{SBS}} = \sum_{i=1}^{d_S} p_i |i\rangle\langle i|_S \bigotimes_{k=1}^N \rho_k^{(i)}, \quad (\text{A2})$$

where $\rho_k^{(i)}$ is the state on $\mathcal{H}_k$ conditioned on the state $|i\rangle$ on $\mathcal{H}_S$, and where $\rho_k^{(i)} \rho_k^{(j)} = 0$ for $i \neq j$.

Appendix B: The impossibility of ideal projective measurements

In this appendix, we demonstrate the impossibility of performing an exact projective measurement according to the Measurement Equilibration Hypothesis. In particular, we prove that an initially uncorrelated state cannot equilibrate to one with Spectrum Broadcast structure; this follows from a stronger result, described in the following. Firstly, we allow the possibility of initial correlations between the observers' systems, but not with the system to be measured *i.e.* $\rho(0) = \rho_{S,0} \otimes \rho_{E,0}$, where $\rho_{S,0}$ is the initial state of the system and $\rho_{E,0}$ the initial state of the environment. Secondly, we demand equilibration to a classical-quantum state [25] with respect to the measurement basis, with the appropriate probabilities, *i.e.* a state of the form

$$\rho_{\text{CQ}} = \sum_i p_i |i\rangle\langle i|_S \otimes \rho_E^{(i)} \quad (\text{B1})$$

for some $\rho_E^{(i)}$ on $\mathcal{H}_E$. Thirdly, we demand orthogonality of the conditional states $\{\rho_E^{(i)}\}_i$, *i.e.* $\rho_E^{(i)} \rho_E^{(j)} = 0$, or equivalently $F(\rho_E^{(i)}, \rho_E^{(j)}) = 0$, for $i \neq j$. These are weaker requirements than for Spectrum Broadcast Structure, as can be seen by comparing Eqs. (B1) and (A2). Thus when we prove that they cannot be simultaneously satisfied, the impossibility of equilibration to a state with Spectrum Broadcast Structure follows as a consequence.

The proof proceeds by first constraining the form of the Hamiltonian that would be necessary to equilibrate to a classical-quantum state, and then showing that the resulting fidelity between the conditional states $\{\rho_E^{(i)}\}_i$ is necessarily non-zero.

B.1. Constraining the Hamiltonian

Recalling that if a system equilibrates on average, it equilibrates to its infinite-time average [16], our first step is to find which Hamiltonians can give rise to an infinite-time average of the required classical-quantum form, *i.e.* Eq. (B1). Writing the spectral decomposition of the Hamiltonian that we seek as $H = \sum_n E_n \Pi_n$, we recall that the infinite-time average is a pinching map with respect to the Hamiltonian [17]

$$\rho_\infty = \sum_n \Pi_n \rho(0) \Pi_n, \quad (\text{B2})$$

and thus $[\rho_\infty, H] = 0$. Consequently, if $\rho_\infty = \rho_{\text{CQ}}$, then $[\rho_{\text{CQ}}, H] = 0$. We now seek the constraints that this latter equation imposes on the form of H . Letting $\{|\lambda\rangle\}$ denote an arbitrary basis of $\mathcal{H}_E$, where $\dim(\mathcal{H}_E) = d_E$, we can take the tensor product with the pointer basis on $\mathcal{H}_S$ to write any Hamiltonian in the form

$$H = \sum_{ij\lambda\lambda'} C_{ij\lambda\lambda'} |i\rangle\langle j| \otimes |\lambda\rangle\langle\lambda'|, \quad (\text{B3})$$

where $C_{ij\lambda\lambda'}$ are the matrix elements of H in this basis. Then the commutator gives us

$$\begin{aligned} [H, \rho_{\text{CQ}}] &= \sum_l p_l \sum_{ij\lambda\lambda'} C_{ij\lambda\lambda'} \left[|i\rangle\langle j| \otimes |\lambda\rangle\langle\lambda'|, |l\rangle\langle l| \otimes \rho_E^{(l)} \right] \\ &= \sum_j p_j \sum_{i\lambda\lambda'} \left(C_{ij\lambda\lambda'} |i\rangle\langle j| \otimes |\lambda\rangle\langle\lambda'| \rho_E^{(j)} - C_{ji\lambda\lambda'} |j\rangle\langle i| \otimes \rho_E^{(j)} |\lambda\rangle\langle\lambda'| \right) \\ &\stackrel{!}{=} 0. \end{aligned} \quad (\text{B4})$$

Now, since we seek a Hamiltonian which equilibrates to the SBS state for any initial system state, *i.e.* any $\{p_i\}$, this implies that

$$\sum_{i\lambda\lambda'} \left(C_{ij\lambda\lambda'} |i\rangle\langle j| \otimes |\lambda\rangle\langle\lambda'| \rho_E^{(j)} - C_{ji\lambda\lambda'} |j\rangle\langle i| \otimes \rho_E^{(j)} |\lambda\rangle\langle\lambda'| \right) = 0 \quad \forall j. \quad (\text{B5})$$

Noting that $\forall i, j$ with $i \neq j$, the tensor factors $|i\rangle\langle j|$, $|j\rangle\langle i|$ and $|j\rangle\langle j|$ are all linearly independent, we then see that Eq. (B5) implies the following

$$\sum_{\lambda\lambda'} \left(C_{ij\lambda\lambda'} |i\rangle\langle j| \otimes |\lambda\rangle\langle\lambda'| \rho_E^{(j)} - C_{ji\lambda\lambda'} |j\rangle\langle i| \otimes \rho_E^{(j)} |\lambda\rangle\langle\lambda'| \right) = 0 \quad \forall i, j, i \neq j, \quad (\text{B6})$$

$$\sum_{\lambda\lambda'} C_{ii\lambda\lambda'} |i\rangle\langle i| \otimes [|\lambda\rangle\langle\lambda'|, \rho_E^{(i)}] = 0 \quad \forall i. \quad (\text{B7})$$

Again, due to the linear independence of $|i\rangle\langle j|$ and $|j\rangle\langle i|$, we must have that $C_{ij\lambda\lambda'} = 0 \quad \forall i \neq j, \lambda, \lambda'$, in order for Eq. (B6) to be satisfied. For the infinite-time average state to have the structure of Eq. (B1) (for an arbitrary initial system state), the Hamiltonian must then be of the form

$$H = \sum_i |i\rangle\langle i| \otimes H_E^{(i)}, \text{ where } H_E^{(i)} := \sum_{\lambda\lambda'} C_{ii\lambda\lambda'} |\lambda\rangle\langle\lambda'| \quad (\text{B8})$$

This generates unitary evolution of the form

$$e^{-iHt} = \sum_i |i\rangle\langle i| \otimes U_E^{(i)}(t), \quad (\text{B9})$$

where $U_E^{(i)}(t) := \exp[-it \sum_{\lambda\lambda'} C_{ii\lambda\lambda'} |\lambda\rangle\langle\lambda'|]$. We can then write

$$\begin{aligned} \rho_\infty &= \left\langle \sum_i |i\rangle\langle i| \otimes U_E^{(i)}(t) \rho_{S,0} \otimes \rho_{E,0} \sum_j |j\rangle\langle j| \otimes U_E^{(j)\dagger}(t) \right\rangle_\infty \\ &= \sum_{i,j} \langle i | \rho_{S,0} | j \rangle |i\rangle\langle j| \otimes \left\langle U_E^{(i)}(t) \rho_{E,0} U_E^{(j)\dagger}(t) \right\rangle_\infty \stackrel{!}{=} \sum_i p_i |i\rangle\langle i| \otimes \rho_E^{(i)} \end{aligned} \quad (\text{B10})$$

which implies that

$$\left\langle U_E^{(i)}(t) \rho_{E,0} U_E^{(j)\dagger}(t) \right\rangle_\infty \stackrel{!}{=} \begin{cases} 0 & i \neq j \\ \rho_E^{(i)} & i = j. \end{cases} \quad (\text{B11})$$

Now, let us denote the eigenvalues of $H_E^{(i)}$ by $\{\varepsilon_n^{(i)}\}_n$, with associated eigenvectors $\{|\varepsilon_n^{(i)}\rangle\}_n$. Writing $U_E^{(i)}(t)$ and $U_E^{(j)}(t)$ in the left hand side of Eq. (B11) in the eigenbasis of $H_E^{(i)}$ and $H_E^{(j)}$ respectively, we have

$$\left\langle U_E^{(i)}(t) \rho_{E,0} U_E^{(j)\dagger}(t) \right\rangle_\infty = \sum_{n,m} \left\langle e^{-i(\varepsilon_n^{(i)} - \varepsilon_m^{(j)})t} \right\rangle_\infty \langle \varepsilon_n^{(i)} | \rho_{E,0} | \varepsilon_m^{(j)} \rangle |\varepsilon_n^{(i)}\rangle\langle \varepsilon_m^{(j)}|, \quad (\text{B12})$$

and so in order to satisfy Eq. (B11), we must have that

$$\left\langle e^{-i(\varepsilon_n^{(i)} - \varepsilon_m^{(j)})t} \right\rangle_\infty \langle \varepsilon_n^{(i)} | \rho_{E,0} | \varepsilon_m^{(j)} \rangle \stackrel{!}{=} 0 \quad \forall n, m \text{ where } n \neq m. \quad (\text{B13})$$

If the $\{H_E^{(i)}\}_i$ are each nondegenerate, and have no eigenvalues in common, *i.e.* $\varepsilon_n^{(i)} \neq \varepsilon_m^{(j)} \quad \forall n, m$ with $n \neq m$, then the above equation holds independently of the initial state $\rho_{E,0}$. Note that both of these conditions are true if and only if the total Hamiltonian is non-degenerate. In the case where there are values of $\varepsilon_n^{(i)}$ and $\varepsilon_m^{(j)}$ that coincide, *i.e.* there are some degeneracies, then we must have $\langle \varepsilon_n^{(i)} | \rho_{E,0} | \varepsilon_m^{(j)} \rangle = 0$ for each pair of coinciding eigenvalues, requiring a initially rank-deficient state. Assuming either non-degeneracy of the total Hamiltonian, or the appropriate rank-deficiency of the initial environmental states in the case of degeneracies, we arrive at a state of the form

$$\rho_\infty = \sum_i p_i |i\rangle\langle i| \otimes \rho_E^{(i)} \quad \text{with} \quad \rho_E^{(i)} = \sum_n \langle \varepsilon_n^{(i)} | \rho_{E,0} | \varepsilon_n^{(i)} \rangle |\varepsilon_n^{(i)}\rangle\langle \varepsilon_n^{(i)}|. \quad (\text{B14})$$

As an aside, we note that this state satisfies the faithfulness requirement in [1].

We have so far made no demands on the microscopic structure on $\mathcal{H}_E$, aside from the potential rank-deficient initial states required to deal with degeneracies in H . A Hamiltonian of the form of Eq. B8 is therefore necessary for the formation of classical-quantum states in general.

B.2. Bounding the fidelity between conditional states

We now show that the orthogonality requirement, *i.e.* $F(\rho_E^{(i)}, \rho_E^{(j)}) = 0$ for $i \neq j$, cannot be satisfied. In particular, we show that

$$F(\rho_E^{(i)}, \rho_E^{(j)}) \geq \frac{1}{d_E^2}. \quad (\text{B15})$$

First, note that the conditional time-averaged state $\rho_E^{(i)}$ in Eq. (B14) is a result of the pinching map with respect to the projectors $\{|\varepsilon_n^{(i)}\rangle\langle\varepsilon_n^{(i)}|\}_{n=1}^{d^{(i)}}$. Denoting the number of these projectors (*i.e.* the number of distinct eigenvalues of $H_E^{(i)}$) by $d^{(i)}$, we can rewrite the pinching map as a mixed-unitary channel [80, 81] as follows

$$\rho_E^{(i)} = \sum_{n=1}^{d^{(i)}} |\varepsilon_n^{(i)}\rangle\langle\varepsilon_n^{(i)}| \rho_{E,0} |\varepsilon_n^{(i)}\rangle\langle\varepsilon_n^{(i)}| = \frac{1}{d^{(i)}} \sum_{y=1}^{d^{(i)}} U_y^{(i)} \rho_{E,0} U_y^{(i)\dagger} \quad (\text{B16})$$

where $U_y := \sum_{x=1}^{m^{(i)}} e^{-\frac{i2\pi xy}{m^{(i)}}} \Pi_x^{(i)}$. Noting that $U_{d^{(i)}} = \mathbb{1}$, we write

$$\rho_E^{(i)} = \frac{\rho_{E,0}}{d^{(i)}} + \frac{1}{d^{(i)}} \sum_{y=1}^{d^{(i)}-1} U_y^{(i)} \rho_{E,0} U_y^{(i)\dagger} \quad (\text{B17})$$

and thus

$$\begin{aligned} \text{Tr}[(\rho_E^{(i)} \rho_E^{(j)})] &= \frac{1}{d^{(i)} d^{(j)}} \left\{ \text{Tr}[\rho_{E,0}^2] + \text{Tr}\left[\rho_{E,0} \sum_{y=1}^{d^{(i)}-1} U_y^{(j)} \rho_{E,0} U_y^{(j)\dagger}\right] + \text{Tr}\left[\sum_{y=1}^{d^{(i)}-1} U_y^{(i)} \rho_{E,0} U_y^{(i)\dagger} \rho_{E,0}\right] \right. \\ &\quad \left. + \text{Tr}\left[\sum_{y=1}^{d^{(i)}-1} U_y^{(i)} \rho_{E,0} U_y^{(i)\dagger} \sum_{y'=1}^{d^{(i)}-1} U_{y'}^{(j)} \rho_{E,0} U_{y'}^{(j)\dagger}\right] \right\}. \end{aligned} \quad (\text{B18})$$

Since that the trace of the product of two positive operators is itself positive, we can remove the second, third and fourth terms to obtain the inequality

$$\text{Tr}(\rho_E^{(i)} \rho_E^{(j)}) \geq \frac{\text{Tr}(\rho_{E,0}^2)}{d^{(i)} d^{(j)}}. \quad (\text{B19})$$

Now, note that the fidelity between two states satisfies $F(\rho, \sigma) \geq \text{Tr}(\rho\sigma)$ [79], allowing us to bound the Fidelity

$$F(\rho_E^{(i)}, \rho_E^{(j)}) \geq \frac{\text{Tr}(\rho_{E,0}^2)}{d^{(i)} d^{(j)}}. \quad (\text{B20})$$

Thus, the fidelity between any two time-averaged states with the same initial state cannot be zero. It is therefore impossible to equilibrate to a classical-quantum state, Eq. (B1), satisfying the orthogonality condition. As a consequence, it is impossible for an initially uncorrelated state to equilibrate to a state with Spectrum Broadcast Structure, as this is a stronger condition than the one proven to be impossible, as discussed above.

Finally, we note that since $d^{(i)}$ is bounded by the Hilbert space dimension of the environment, *i.e.* $d_E \geq d^{(i)}$, Eq. (B15) follows immediately from Eq. (B20).

Appendix C: The standard measurement model does not equilibrate to a correlated state

In this section we consider the Hamiltonian used in the von-Neumann measurement scheme (or the “standard model of measurement”) [5]:

$$H = X_S \otimes Y_E, \quad (\text{C1})$$

where the measurement basis is given by the eigenbasis of X_S . By carefully controlling the time for which this Hamiltonian is applied, one can correlate the system S with the environment E in this basis. However, without

temporal control – specifically in an equilibration paradigm – this Hamiltonian leads to an equilibrium state that does not feature any correlations between system and environment, as we now show.

The unitary operator generated by Eq. (C1) is given by

$$e^{-iHt} = e^{-iX_S \otimes Y_E t} = e^{-i(\sum_i x_i |i\rangle\langle i|_S \otimes Y_E)t} = \sum_i |i\rangle\langle i|_S \otimes e^{-ix_i Y_E t}, \quad (C2)$$

As in the previous section, we assume that system and environment are initially uncorrelated, such that the initials state is given by

$$\rho(0) = \rho_{S,0} \otimes \rho_{E,0}, \quad (C3)$$

and we calculate the time evolved state of the total system

$$\rho(t) = e^{-iHt} \rho(0) e^{iHt} = e^{-iX_S \otimes Y_E t} (\rho_{S,0} \otimes \rho_{E,0}) e^{iX_S \otimes Y_E t} = \sum_{i,j} p_{ij} |i\rangle\langle j| \otimes \rho_E^{(ij)}(t) \quad (C4)$$

where p_{ij} are the elements of the initial system state in the pointer basis, *i.e.* $\rho_{S,0} = \sum_{i,j} p_{ij} |i\rangle\langle j|$, and we have defined $\rho_E^{(ij)}(t) := e^{-ix_i Y_E t} \rho_{E,0} e^{ix_j Y_E t}$. Now, writing the spectral decomposition $Y_E = \sum_n \varepsilon_n |\varepsilon_n\rangle\langle \varepsilon_n|$, we have

$$\rho_E^{(ij)}(t) = \sum_{m,n} e^{-i(x_i \varepsilon_m - x_j \varepsilon_n)t} \rho_E^{mn} |\varepsilon_m\rangle\langle \varepsilon_n| \quad (C5)$$

where $\rho_E^{mn} := \langle \varepsilon_n | \rho_{E,0} | \varepsilon_m \rangle$. Examining this together with the total state in Eq. (C4), we see that the time average of the total state is determined by

$$g_{mn}^{(ij)}(T) := \langle e^{-i(x_i \varepsilon_m - x_j \varepsilon_n)t} \rangle_T = \begin{cases} 1 & \text{if } x_i \varepsilon_m = x_j \varepsilon_n \\ \frac{ie^{[i(x_i \varepsilon_m - x_j \varepsilon_n)T]} - 1}{(x_i \varepsilon_m - x_j \varepsilon_n)T} & \text{otherwise} \end{cases} \quad (C6)$$

Assuming $x_i \varepsilon_m \neq x_j \varepsilon_n \quad \forall i, j, m, n$ we find the equilibrium state

$$\langle \rho(t) \rangle_\infty = \sum_i p_i |i\rangle\langle i|_S \otimes \sum_m \rho_E^{mm} |\varepsilon_m\rangle\langle \varepsilon_m|. \quad (C7)$$

Thus the environment is independent of i , and this state features no correlations between system and environment. The environment therefore can not reveal any information about the system, and “standard model” Hamiltonians are incapable of (even approximately) equilibrating to states exhibiting objectivity.

Appendix D: Asymptotically approaching Spectrum Broadcast Structure

In this section we show that and how Spectrum Broadcast Structure can be achieved approximately by coarse-graining the environment. We will choose a particular form of interaction, *i.e.* a “star-shaped” interaction where every part of the environment interacts directly with the system. The purpose of this is to show one example of how Spectrum Broadcast Structure can emerge. Further research is needed to expand this framework into a realistic model of measurements. These calculations therefore should be considered as a first step and a proof of principle.

As was discussed above Eq. 6 achieving Spectrum Broadcast Structure as a time average state requires a conditional structure of the Hamiltonian. In conjunction with the condition of Strong Independence this prompts us to consider a “star-shaped” Hamiltonian (see Fig. 1) given by

$$H_{\text{int}} = \sum_i |i\rangle\langle i|_S \otimes \sum_{k=1}^N c_k H_k^{(i)} \quad (D1)$$

where $H_k^{(i)} \in \mathcal{L}(\mathcal{H}_k)$ and c_k are the different couplings between the system and the environments (eg. distance dependent couplings). This generates a unitary time evolution given by

$$e^{-iH_{\text{int}}t} = \sum_i |i\rangle\langle i|_S \bigotimes_{k=1}^N U_k^{(i)}(t), \quad (D2)$$

where $U_k^{(i)}(t) := e^{-ic_k H_k^{(i)} t}$.

Assuming a separable initial state, *i.e.* $\rho(0) = \rho_{S,0} \otimes_{k=1}^N \rho_{k,0}$, the state at time t can be written

$$\rho(t) = \sum_{i,j} \langle i | \rho_{S,0} | j \rangle | i \rangle \langle j |_S \bigotimes_{k=1}^N \rho_k^{(ij)}(t) \quad (\text{D3})$$

where

$$\rho_k^{(ij)}(t) := U_k^{(i)}(t) \rho_{k,0} U_k^{(j)}(t). \quad (\text{D4})$$

As was argued in the main text and the papers cited therein, given an initial systems with large effective dimension d_{eff} the trace distance between the infinite-time average state and the equilibrium state vanishes. What is therefore left to show is that the infinite-time average state $\langle \rho(0) \rangle_\infty$ approaches Spectrum Broadcast Structure under coarse-graining of the environment.

To this end we first calculate infinite-time average of $\rho(t)$. Noting that the system tensor factor in Equation D3 is time-independent, we calculate the time average of the remaining factor. First, letting $\varepsilon_{n_k}^{(i)}$ denote the n_k^{th} energy eigenvalue of $H_k^{(i)}$, we define

$$\rho_{k,0}^{mn(ij)} := \langle \varepsilon_{m_k}^{(i)} | \rho_{k,0} | \varepsilon_{n_k}^{(j)} \rangle. \quad (\text{D5})$$

Then the time average state of the total environment at time T is given by

$$\begin{aligned} \left\langle \bigotimes_{k=1}^N \rho_k^{(ij)}(t) \right\rangle_T &= \left\langle \bigotimes_{k=1}^N \sum_{m_k, n_k} \exp[-ic_k (\varepsilon_{m_k}^{(i)} - \varepsilon_{n_k}^{(j)}) t] \rho_{k,0}^{mn(ij)} | \varepsilon_{m_k}^{(i)} \rangle \langle \varepsilon_{n_k}^{(j)} | \right\rangle_T \\ &= \left\langle \left(\sum_{m_1, n_1} \exp[-ic_1 (\varepsilon_{m_1}^{(i)} - \varepsilon_{n_1}^{(j)}) t] \rho_{1,0}^{mn(ij)} | \varepsilon_{m_1}^{(i)} \rangle \langle \varepsilon_{n_1}^{(j)} | \right) \otimes \left(\sum_{m_2, n_2} \exp[-ic_2 (\varepsilon_{m_2}^{(i)} - \varepsilon_{n_2}^{(j)}) t] \rho_{2,0}^{mn(ij)} | \varepsilon_{m_2}^{(i)} \rangle \langle \varepsilon_{n_2}^{(j)} | \right) \otimes \dots \right\rangle_T \\ &= \sum_{m_1, n_1} \sum_{m_2, n_2} \dots \sum_{m_N, n_N} \left\langle \exp \left[-i \sum_{k=1}^N c_k (\varepsilon_{m_k}^{(i)} - \varepsilon_{n_k}^{(j)}) t \right] \right\rangle_T \bigotimes_{k=1}^N \rho_{k,0}^{mn(ij)} | \varepsilon_{m_k}^{(i)} \rangle \langle \varepsilon_{n_k}^{(j)} | \\ &= \sum_{\{m_k\}\{n_k\}} f_{\{m_k\}\{n_k\}}^{(ij)}(T) \bigotimes_{k=1}^N \rho_{k,0}^{mn(ij)} | \varepsilon_{m_k}^{(i)} \rangle \langle \varepsilon_{n_k}^{(j)} | \end{aligned} \quad (\text{D6})$$

where we have used the notation

$$\sum_{\{a_k\}} = \sum_{a_1} \sum_{a_2} \dots \sum_{a_N}, \quad (\text{D7})$$

and defined

$$f_{\{m_k\}\{n_k\}}^{(ij)}(T) := \begin{cases} 1 & \text{if } \sum_{k=1}^N c_k \varepsilon_{m_k}^{(i)} = \sum_{k=1}^N c_k \varepsilon_{n_k}^{(j)} \\ i^{\frac{\exp[-i \sum_{k=1}^N c_k (\varepsilon_{m_k}^{(i)} - \varepsilon_{n_k}^{(j)}) T] - 1}{\sum_{k=1}^N c_k (\varepsilon_{m_k}^{(i)} - \varepsilon_{n_k}^{(j)}) T}} & \text{otherwise.} \end{cases} \quad (\text{D8})$$

Note that the cases where $f_{\{m_k\}\{n_k\}}^{(ij)} = 1$ are particular "fine-tuned" cases, where the coupling constants and energy eigenvalues have to have to conspire in order to yield this result, requiring a highly degenerate total Hamiltonian. Assuming that the total Hamiltonian is non-degenerate, which is a rather natural assumption as it is fulfilled by most Hamiltonians that appear in nature, implies $\sum_{k=1}^N c_k \varepsilon_{m_k}^{(i)} \neq \sum_{k=1}^N c_k \varepsilon_{n_k}^{(j)} \forall i, j$. As a consequence, the first case in Eq. D8 never occurs and the time average results in

$$\left\langle \bigotimes_{k=1}^N \rho_k^{(ij)}(t) \right\rangle_\infty = 0 \quad \text{for } i \neq j, \quad (\text{D9})$$

combining this fact with Equations D3 and D6, we have that

$$\begin{aligned} \langle \rho(t) \rangle_\infty &= \sum_i \langle i | \rho_{S,0} | i \rangle | i \rangle \langle i |_S \left\langle \bigotimes_{k=1}^N \rho_k^{(ii)}(t) \right\rangle_\infty \\ &= \sum_i p_i | i \rangle \langle i |_S \sum_{\{m_k\}\{n_k\}} f_{\{m_k\}\{n_k\}}^{(ii)}(\infty) \bigotimes_{k=1}^N \rho_{k,0}^{mn(ii)} | \varepsilon_{m_k}^{(i)} \rangle \langle \varepsilon_{n_k}^{(i)} |, \end{aligned} \quad (\text{D10})$$

where we simplified the notation of $\lim_{T \rightarrow \infty} f_{\{m_k\}\{n_k\}}^{(ii)}(T) = f_{\{m_k\}\{n_k\}}^{(ii)}(\infty)$, and $p_i := \langle i | \rho_{S,0} | i \rangle$ is the probability of outcome i for the projective measurement $\{|i\rangle\langle i|_S\}_i$ on the system. Due to the assumption of non-degeneracy of the total Hamiltonian, the total conditional environment Hamiltonian $H_E^{(i)}$ has to be non-degenerate as well, consequently resulting in $\sum_{k=1}^N c_k \varepsilon_{m_k}^{(i)} \neq \sum_{k=1}^N c_k \varepsilon_{n_k}^{(i)} \forall m_k \neq n_k$, and thus only the $m_k = n_k$ terms remain. Then Equation (D8) implies

$$\begin{aligned} \langle \rho(t) \rangle_\infty &= \sum_i p_i |i\rangle\langle i|_S \sum_{m_1 \dots m_N} \bigotimes_{k=1}^N \rho_{k,0}^{mm(ii)} |\varepsilon_{m_k}^{(i)}\rangle\langle \varepsilon_{m_k}^{(i)}| \\ &= \sum_i p_i |i\rangle\langle i|_S \bigotimes_{k=1}^N \sum_{m_k} \rho_{k,0}^{mm(ii)} |\varepsilon_{m_k}^{(i)}\rangle\langle \varepsilon_{m_k}^{(i)}| \\ &= \sum_i p_i |i\rangle\langle i|_S \bigotimes_{k=1}^N \rho_{k,0}^{(i)}, \end{aligned} \quad (\text{D11})$$

where

$$\rho_{k,0}^{(i)} := \sum_{m_k} |\varepsilon_{m_k}^{(i)}\rangle\langle \varepsilon_{m_k}^{(i)}| \rho_{k,0} |\varepsilon_{m_k}^{(i)}\rangle\langle \varepsilon_{m_k}^{(i)}| \quad (\text{D12})$$

is the initial state of environmental fraction k (*i.e.* $\rho_{k,0}$) subject to full decoherence in the eigenbasis of $H_k^{(i)}$. Now let us gather together the N factors of the environment into M “fragments” or “macro-environments”, and let N_q denote the set of environment labels belonging to the q^{th} fragment, so that $\bigotimes_{k=1}^N \mathcal{H}_k = \bigotimes_{q=1}^M \mathcal{E}_q$ with $\mathcal{E}_q := \bigotimes_{k \in N_q} \mathcal{H}_k$. Then

$$\langle \rho(t) \rangle_\infty = \sum_i p_i |i\rangle\langle i|_S \bigotimes_{q=1}^M \tilde{\rho}_q^{(i)}, \quad \text{where} \quad \tilde{\rho}_q^{(i)} := \bigotimes_{k \in N_q} \rho_{k,0}^{(i)}, \quad (\text{D13})$$

If we can show that the Uhlmann Fidelity of two branches of each macro-environment, *i.e.* $F(\rho_q^{(i)}, \rho_q^{(j)})$ tends to 0 as the macro-environment size increases we know that in the limit they must have orthogonal support, which is due to the fact

$$F(\sigma, \rho) = 0 \iff \sigma \rho = 0 \quad (\text{D14})$$

where the Uhlmann fidelity is defined as $F(\sigma, \rho) = \text{Tr}(\sqrt{\sqrt{\sigma}\rho\sqrt{\sigma}})^2$.

To show this we use the inequality

$$F(\sigma, \rho) \leq \sum_m \sqrt{\text{Tr}[E_m \rho]} \sqrt{\text{Tr}[E_m \sigma]} \quad (\text{D15})$$

where the E_m are POVM elements from which follows that $\sum_m E_m = 1$. We choose them to be projective measurements on the Eigenbasis of the Hamiltonian of the i branch (w.l.o.g.), *i.e.* $E_m = |\varepsilon_{m_k}^{(i)}\rangle\langle \varepsilon_{m_k}^{(i)}|$. Since $\rho_{k,0}^{(i)} = \sum_{m_k} \rho_{k,0}^{mm(ii)} |\varepsilon_{m_k}^{(i)}\rangle\langle \varepsilon_{m_k}^{(i)}|$ we get

$$\begin{aligned} \sqrt{F(\rho_{k,0}^{(i)}, \rho_{k,0}^{(j)})} &\leq \sum_{m_k} \sqrt{\rho_{k,0}^{mm(ii)}} \sqrt{\text{Tr}[\sum_{n_k} \rho_{k,0}^{nn(jj)} |\varepsilon_{n_k}^{(j)}\rangle\langle \varepsilon_{n_k}^{(j)}| |\varepsilon_{m_k}^{(i)}\rangle\langle \varepsilon_{m_k}^{(i)}|]} \\ &= \sum_{m_k} \sqrt{\rho_{k,0}^{mm(ii)}} \sqrt{\sum_{n_k} \rho_{k,0}^{nn(jj)} |\langle \varepsilon_{m_k}^{(i)} | \varepsilon_{n_k}^{(j)} \rangle|^2} = \sum_{m_k} \sqrt{\sum_{n_k} \rho_{k,0}^{mm(ii)} \rho_{k,0}^{nn(jj)} |\langle \varepsilon_{m_k}^{(i)} | \varepsilon_{n_k}^{(j)} \rangle|^2} \end{aligned} \quad (\text{D16})$$

In the cases where $H^{(i)} = H^{(j)}$ or $\rho_{k,0} = \frac{1}{d_k} \mathbb{1}$ we see that the fidelity is 1, as one would expect.

To see that it approaches zero, we recall that $F(\bigotimes_k \sigma_k, \bigotimes_k \rho_k) = \prod_k F(\sigma_k, \rho_k)$. Then,

$$\sqrt{F(\bigotimes_{k \in N_q} \rho_{k,0}^{(i)}, \bigotimes_{k \in N_q} \rho_{k,0}^{(j)})} = \prod_{k \in N_q} \sqrt{F(\rho_{k,0}^{(i)}, \rho_{k,0}^{(j)})} \leq \prod_{k \in N_q} \left(\sum_{m_k} \sqrt{\sum_{n_k} \rho_{k,0}^{mm(ii)} \rho_{k,0}^{nn(jj)} |\langle \varepsilon_{m_k}^{(i)} | \varepsilon_{n_k}^{(j)} \rangle|^2} \right) \quad (\text{D17})$$

Now let's say that $\eta_k^{(ij)} = \sum_{m_k} \sqrt{\sum_{n_k} \rho_{k,0}^{mm(ii)} \rho_{k,0}^{nn(jj)} |\langle \varepsilon_{m_k}^{(i)} | \varepsilon_{n_k}^{(j)} \rangle|^2}$, $N_q^{(ij)} = \{k : \eta_k^{(ij)} < 1\}$ and $\eta_{\max}^{(ij)} = \max_{k \in N_q^{(ij)}} \eta_k^{(ij)}$ then

$$F(\tilde{\rho}_{q,0}^{(i)}, \tilde{\rho}_{q,0}^{(j)}) \leq (\eta_{\max}^{(ij)})^{2|N_q^{(ij)}|}, \quad (\text{D18})$$

which can be simplified by defining $\gamma_q^{(ij)} = -2\ln(\eta_{\max}^{(ij)})$ to finally get

$$F(\tilde{\rho}_{q,0}^{(i)}, \tilde{\rho}_{q,0}^{(j)}) \leq e^{-\gamma_q^{(ij)} |N_q^{(ij)}|} \quad (\text{D19})$$

which exponentially tends to 0 for increasing macro-environment size under the assumption that the conditional interaction Hamiltonians $H_k^{(i)}$ don't conspire to yield a non increasing $|N_q^{(ij)}|$, i.e. that the eigenvectors of branch i and j are not all the same, which is a reasonable assumption for most physical interactions. The infinite-time-average state therefore tends exponentially to the Spectrum Broadcast Structure state as the size of the macro-environments (and thus the environment as a whole) increases.

Chapter III

Fundamental limitations of time keeping

III.1 Autonomous Temporal Probability Concentration: Clockworks and the second law of thermodynamics

III.1.1 Contribution

I contributed significantly to the conception of the general theoretic framework of this research project. I was involved in the conception of all analytic calculations and I wrote the Mathematica code for the numerical analysis and produced all the plots. In particular I derived the main analytical result, *i.e.* the analytic form of the temporal probability concentration result of Equation 2 showing the approachability of an ideal clockwork (Appendices A, B and C), as well as appendix D. Furthermore I did the numerical analysis leading to appendix F. I wrote the main text of the manuscript and managed revisions it after peer review. I designed Figure 2 and Figure 3.

Autonomous Temporal Probability Concentration: Clockworks and the Second Law of Thermodynamics

Emanuel Schwarzhans,^{*} Maximilian P. E. Lock, Paul Erker, and Nicolai Friis
*Institute for Quantum Optics and Quantum Information - IQOQI Vienna,
Austrian Academy of Sciences, Boltzmannngasse 3, 1090 Vienna, Austria*

Marcus Huber[†]
*Institute for Quantum Optics and Quantum Information - IQOQI Vienna,
Austrian Academy of Sciences, Boltzmannngasse 3, 1090 Vienna, Austria and
Vienna Center for Quantum Science and Technology, Atominstitut, TU Wien, 1020 Vienna, Austria*

According to thermodynamics, the inevitable increase of entropy allows the past to be distinguished from the future. From this perspective, any clock must incorporate an irreversible process that allows this flow of entropy to be tracked. In addition, an integral part of a clock is a clockwork, that is, a system whose purpose is to temporally concentrate the irreversible events that drive this entropic flow, thereby increasing the accuracy of the resulting clock ticks compared to counting purely random equilibration events. In this article, we formalise the task of *autonomous temporal probability concentration* as the inherent goal of any clockwork based on thermal gradients. Within this framework, we show that a perfect clockwork can be approximated arbitrarily well by increasing its complexity. Furthermore, we combine such an idealised clockwork model, comprised of many qubits, with an irreversible decay mechanism to showcase the ultimate thermodynamic limits to the measurement of time.

POPULAR SUMMARY

The laws of physics governing the microscopic world from which our experience emerges know no directionality of time. A recording of an isolated physical system looks plausible played forwards and backwards. Nevertheless, an arrow of time can be identified in many complex processes, such as the breaking of an egg. The irreversibility of such processes can also be used to keep track of the passage of time. In our paper, we investigate the implications of this insight for building (quantum) clocks from first principles. Any system that can function as a clock needs an irreversible component whose dynamics indicate a clear directionality of time. However, typical irreversible processes found in nature, such as decaying atoms and equilibrating heat baths, have no discernable regular temporal structure and therefore make for rather bad clocks by themselves. It thus comes at no surprise that cooling coffee cups and similar items are not usually used for timekeeping. Instead, one relies on the combination of internal, periodic dynamics with an irreversible, temporally unstructured, process to construct well-functioning clocks. The same is fundamentally true when envisioning clocks at the quantum scale: periodic internal dynamics of a clockwork are combined with an irreversible process to yield temporally well-structured ‘ticks’ of a clock. In this work, we identify and systemat-

ically analyse the task performed by the clockwork, which we refer to as autonomous temporal probability concentration (ATPC). That is, the clockwork generates a periodic structure that determines at which times the irreversible events – the clock’s ticks – can occur. We show that ATPC is intimately related to the complexity of the clockwork and can be performed arbitrarily well, as long as complexity is not constrained. Nonetheless, the irreversible process itself sets a fundamental limit for the quality of clocks, beyond which they cannot be improved. In this sense, the second law of thermodynamics – the origin of irreversibility, generates but also limits the potential for any system to serve as a clock, while the complexity of the clockwork determines how well this can be achieved in practice.

I. INTRODUCTION

Time plays a special role in quantum physics. While other physical quantities of interest are represented as Hermitian operators, there is no observable corresponding to time itself. That is, it is not possible to find an operator conjugate to the Hamiltonian (representing energy) that may serve as ‘time observable’ in the same way as is done for position and momentum [1] (see e.g. [2] for some caveats to this statement). Time thus plays the role of a parameter in the equations of motion. Consequently, the passage of time is estimated via the evolution of a reference system — a *clock*. By tracking the dynamical evolution of (observable quantities related to) such a clock system it is possible to extract information about the flow of time, see, e.g. [3–10]. *But what makes a spe-*

^{*} emanuel.schwarzhans@oeaw.ac.at

[†] marcus.huber@univie.ac.at

cific system useful as a clock? To address this question, we consider time to be a continuously elapsing parameter t ('Schrödinger time') whose value is estimated by a clock in terms of discrete increments ('ticks'). According to quantum theory, the evolution of any closed system is time-reversal symmetric, and therefore any complete description of an instrument that measures time inevitably requires an irreversible part that breaks this symmetry. By definition, the equilibrium state of any system features no non-trivial evolution in time. Thus, the first necessary ingredient for building a clock is an out-of-equilibrium system, such that the clock can harness the irreversible transition to higher entropy to produce ticks.

Entropy-increasing processes are fundamentally stochastic. Consequently, individual events resulting from such a process provide little information about t and thus make for rather bad clocks. While one could, in principle, use any equilibrating system as a clock — such as a hot coffee mug cooling down on your desk — its ticks, e.g., the spontaneous emissions of thermal photons (which exhibit super-Poissonian statistics), come at highly irregular intervals with respect to Schrödinger time. Structuring this irregular entropy flow into a series of ticks to allow for a precise synchronisation of events is exactly the purpose of a clock. In this article we formalise the task of timekeeping by conceptually separating two stages:

- (i) an irreversible process that follows the second law of thermodynamics, *i.e.* an out-of-equilibrium system moving towards equilibrium by means of discrete and stochastic events,
- (ii) an internal clockwork that temporally concentrates the probability of an irreversible event occurring, thereby mitigating the fluctuations of the intervals between the equilibration events.

As we will see, the particular choice of (i) provides the context for evaluating clock performance because it represents a basic form of clock itself, while at the same time limiting the performance of a clock for any given clockwork. Stage (ii) gives rise to a clearly defined mathematical task that we will refer to as *autonomous temporal probability concentration* (ATPC).

Here, we consider clocks to be autonomous. That is, the Hamiltonian generating the evolution of the clockwork is energy-conserving and time-independent, and the irreversible process is memoryless and requires no external control, *i.e.*, *no active measurement*. Although current quantum clocks are usually far from autonomous, as they require power input and are subject to losses, both of which are usually not fully accounted for in their analysis, we focus on autonomous clocks in order to provide a full analysis of the resources that are fundamentally required to drive a clock.

To describe the performance of a clock, we use two quantities: *accuracy* and *resolution*. The accuracy $N = \left(\frac{\bar{t}}{\Delta t}\right)^2$, is the average number of ticks until the clock

is, on average, off by one tick with respect to Schrödinger time. The resolution $R = 1/\bar{t}$, is the average of the tick frequency with respect to Schrödinger time. Note that this choice for quantifying the resolution is not the only possibility. We choose the above definition since it represents a conservative figure of merit in the sense that it prevents statistical outliers of ticks occurring at small times t from unduly inflating the estimated resolution, as we discuss in more detail in Sec. III.

That there is a trade-off relation between accuracy and resolution, and that there is a proportional relation between the entropy dissipated in the process and the clock performance, was first noticed in a model of an autonomous quantum clock as an open quantum system in [11] and recently corroborated in a mesoscopic experiment in [12].

Here, we combine these aspects and provide a detailed investigation of the trade-offs between accuracy, resolution, and entropy production for given energy and complexity within the framework of autonomous quantum clocks [13, 14]. A central tenet for providing these trade-offs is the separation of timekeeping into two separate processes mentioned above: (i) the irreversible out-of-equilibrium transitions of the clockwork via interaction with an environment, resulting in distinguishable events registered as 'ticks', which we model with a decay mechanism, and (ii) the internal closed-system (unitary) dynamics that provide a clockwork and temporally concentrate the population of states from which an irreversible transition can emerge. That is, the clockwork ensures that the circumstances that allow for a tick to happen (*e.g.*, a specific energy level resonant with the out-of-equilibrium dynamics being highly populated) occur only within a very narrow time window.

We first find that a simple clockwork can only concentrate probability in a limited fashion, prompting the question of whether more complex designs could perform better. We answer this by finding an analytical relation between ATPC and complexity of a specific clockwork model. More generally, we identify the important features of a clockwork that lead to this improvement, and prove that for cold environments, ATPC can be performed arbitrarily well. Then we investigate whether perfect ATPC allows for perfect clocks and find that the answer is no. In fact, the irreversible process sets a limit to the clock quality, and while increasing the complexity (and thus the concentration of probability) first improves the quality of the clock, after a certain point a further increase will actually be detrimental. We thus illustrate the trade-offs between accuracy, resolution, entropy production, and clockwork complexity.

The specific clock model that we consider here consists of (1) external heat baths as out-of-equilibrium resources, (2) a quantum system representing the 'clockwork', and (3) an external field that the clockwork can emit energy ('ticks', *e.g.*, photons) into. In Sec. II, we first discuss the role and choice of the clockwork, and formalise the task of ATPC. In Sec. III, we then discuss mechanisms




Figure 1. Illustration of timekeeping at the level of individual irreversible events. The equilibration events that follow the second law of thermodynamics are inherently stochastic and irregular, in our example we use radiative decays from an excited energy level of a quantum system (which we label ‘top’). By inserting an autonomously operating clockwork between the two out-of-equilibrium systems (‘hot’ decaying quantum systems and a ‘cold’ environment), these decays are temporally structured by temporal variation of the population of the top level, a task that we refer to as Autonomous Temporal Probability Concentration (ATPC). This concentrates the probability of such a decay around the oscillatory peaks of excited population. The panel on the right showcases how greater clockwork complexity leads to a regularisation of individual thermalisation events, *i.e.*, clock ticks. Starting from a thermal population with randomly distributed ticks depicted above, continuing to a simple clockwork with limited population and still significant variance, resulting in ticks being more likely during peak populations and thus less frequent and more regular, and then finally a complex clockwork, increasing population while decreasing temporal variance, giving yet more regular ticks. The tick distribution on the right is exemplary and depicts random ticks, whose spacing approximates the cycle time of the clock work as the temporal probability becomes more concentrated around sharp peaks

for coupling the clockwork to an equilibrating process to produce ticks. In Sec. IV we combine the two, to showcase the limitations set by the irreversible process and how the complexity of a clockwork can be utilised to reach the maximal potential of a clock. We continue in Sec. V with a discussion of the implications and the relation to other literature on clocks and end with a short conclusion in Sec. VI.

II. THERMAL MACHINES AND THE CLOCKWORK

Let us now consider a clockwork in the sense discussed above, that is, a device that contains a target subsystem, which is to be prepared for an out-of-equilibrium transition, thus resulting in a ‘tick’. From a thermodynamic perspective, such a preparation requires work to be performed on the target, which can be achieved by a quantum thermal machine. Operating such a machine in turn requires an out-of-equilibrium resource, which we here consider to be provided by thermal baths at different temperatures, *i.e.* a thermal gradient. More specifically, we assume that two independent baths are available, a hot bath and a cold bath, at temperatures T_H and T_C , respectively, where the latter represents the environment. This setup is depicted in Fig. 1.

This choice is motivated, first, by the general availability of heat baths, *i.e.* it is the most common out-of-equilibrium resource found in nature, such as *e.g.* the sun. Second, because systems are usually expected to thermalise (eventually) without detailed external control or timing, *i.e.* preparing such heat baths does not require any timing device or detailed control of the system’s in-

ternal structure, just an increase in average energy. Consequently, heat baths allow for transparent bookkeeping of the relevant resources, *i.e.* of the average amount of entropy dissipated by the clockwork for each tick.

A specific focus of the analysis performed here lies on the identification of trade-offs between different figures of merit for the clock performance for fixed energy input and clock complexity. In principle, the performance of a given clock also depends on the (difference between the) temperatures T_C and T_H . However, since we are primarily interested in upper bounds on the relevant figures of merit, we will often concentrate on the case where the environment temperature is $T_C = 0$. For the sake of completeness, calculations for general T_C can be found in Appendices B and C.

Our clockwork model then consists of two parts, a d -dimensional ‘ladder’ target system (in the simplest case, a qubit, $d = 2$) and a machine, which itself has some substructure and couples to the ladder via unitary interaction. This interaction supplies work (which the machine draws from its coupling to the heat baths) to the ladder, driving it to its excited states. The ladder in turn couples irreversibly to an external field, and thus these excitations eventually result in ticks (*i.e.* energy emitted into the field). Here, we consider a model where only a non-zero population $P_{\text{top}}(t)$ of the ‘top level’ — the most highly excited state of the ladder — can lead to a tick. Barring some improbable combination of selection rules, such a single sharp energy transition can in practice of course only be approximated. However, as will become clear once we have introduced our model, allowing the possibility of clock ticks occurring due to other transitions would serve to spread the temporal profile of the ticks, decreasing probability concentration. In the spirit

of deriving idealised but fundamental bounds, we therefore focus on decays resulting from only one particular transition. As a consequence, the quality of the clockwork depends on the properties of the particular probability distribution $P_{\text{top}}(t)$ as a function of Schrödinger time t . In particular, an ideal clockwork should be capable of producing

$$P_{\text{top}}(t) = \begin{cases} 1, & \text{if } t = t_0 \\ 0, & \text{otherwise} \end{cases}. \quad (1)$$

While one would expect a perfect clockwork to be capable of producing this distribution, it is also clear that it is not always desirable in conjunction with an irreversible mechanism. If the probability is arbitrarily temporally concentrated, *i.e.*, it is only close to one for a short period of time, but the coupling of the ladder to the external field is of finite strength, then the emission of the ladder energy into the field has a chance not to occur during the peak, thus skipping this tick and worsening the clock performance. Nonetheless, an ideal clockwork should be capable of approximating this ideal distribution to the desired precision set by the irreversible mechanism. Arguably, it seems implausible that a heat engine itself, which intrinsically also harnesses the stochastic flow of energy from a hot to a cold bath, should be able to produce such a perfect signal. However, it may be reasonable to expect that a sufficiently complex clockwork, itself driven by a heat engine, could approximate the ideal ATPC of Eq. (1). In the following, we therefore investigate the role of the complexity of the internal structure of the machine in approximating the ideal ATPC. In order to do so, we decompose the machine into a set of elementary few-qubit machines, each realising an effective virtual qubit [15]. This allows the number of (elementary) machines to be used as a proxy for the complexity of the clockwork's microscopic structure. In terms of these quantifiers, *i.e.* the dimension d of the target system and the number $M(d-1)$ of virtual-qubit machines¹, a central result on autonomous probability concentration that we derive in this paper can be phrased as follows:

Result 1: Autonomous temporal probability concentration of qubit machines

Driving a d -dimensional target system at temperature $T_C = 0$, with M virtual-qubit machines per transition between neighbouring levels, autonomously allows a top-level probability of

$$P_{\text{top}}(t) = \left(1 - \left(1 - \left(\frac{\mathcal{Z}_H - 1}{\mathcal{Z}_H}\right)^{d-1}\right)^M\right) \sin^{2(d-1)}(gt) \quad (2)$$

to be reached. Here, $\mathcal{Z}_H$ is the partition function of a qubit coupled to the hot bath, and can thus take values between 1 and 2.

In other words, we show in the following that the behaviour of an ideal clockwork [*i.e.* Eq. (1)] can be approximated arbitrarily well by increasing the complexity of the clockwork, that is, by increasing M and d .

II.1. Two-Qubit Machine

We begin by considering the simplest possible heat-engine-driven clockwork: a 2-dimensional ladder coupled to a ‘cold’ bath (the environment) and to a two-qubit machine, *i.e.* $d = 2$ and $M = 1$. In terms of Hilbert space dimension, this is the smallest possible thermal machine [15], consisting of a ‘cold’ qubit and a ‘hot’ qubit, in contact with the ‘cold’ environment and a hot bath, respectively, as illustrated in Fig. 2.

Before the machine is activated, the qubits only interact with their respective baths. Under the assumption of weak coupling between the qubits and baths, each qubit thermalises to the corresponding bath temperature. Denoting the energy gaps of the hot, cold and ladder qubits as E_H , E_C and E_L respectively, the reduced states of the qubits can be represented by the thermal states

$$\rho_i = \frac{e^{-\beta_i H_i}}{\mathcal{Z}_i}, \quad (3)$$

with $i = H, C, L$, and where $\mathcal{Z}_i = 1 + e^{-\beta_i E_i}$ are the respective partition functions and H_i the corresponding free Hamiltonians with eigenstates $|0_i\rangle$ and $|1_i\rangle$. The total initial state of the clockwork — the machine and the ladder — thus takes the form $\rho_0 = \rho_H \otimes \rho_C \otimes \rho_L$.

We further assume that the timescale of the interaction between the machine and target qubits is much shorter than that of their thermalisation with the respective baths. Consequently, the relevant dynamics of the clockwork are well described by energy-conserving unitary processes on the clockwork Hilbert space $\mathcal{H} = \mathcal{H}_H \otimes \mathcal{H}_C \otimes \mathcal{H}_L$. This corresponds to assuming that the energy scales of the clockwork are much greater than that of its coupling to the environment. It is suited to our purpose of obtaining fundamental limits to the task of ATPC, as relaxing this assumption will result in the

¹ We consider each of the $d - 1$ transitions between neighbouring energy levels of the ladder to be coupled to M virtual-qubit machines.




Figure 2. Energy-level structure of the minimal thermal clockwork. The transitions induced by H_{int} are indicated by arrows. The green arrows indicate the transition where the ladder gets excited. The yellow arrows show the reverse transition. Coupling the qubit with the biggest energy gap to the hot bath introduces a bias towards the transition that is indicated by the green arrows.

regularity of the clockwork being impaired by the randomness of the bath. Now, since the purpose of the machine is to transfer energy to the target system, we are interested in designing the internal structure of the clockwork, namely the energy levels of the free Hamiltonian $H_0 = H_H + H_C + H_L$ and an interaction Hamiltonian H_{int} such that $[H_0, H_{\text{int}}] = 0$ and $[H_L, H_{\text{int}}] \neq 0$. This can be achieved by choosing the energy gaps to satisfy $E_H \geq E_C$ and $E_L = E_H - E_C$, which results in two degenerate energy levels of H_0 : $|0_C 1_H 0_L\rangle$ and $|1_C 0_H 1_L\rangle$. This, in turn, allows us to define an interaction Hamiltonian that acts non-trivially only within the degenerate subspace, given by

$$H_{\text{int}} = g(|1_C 0_H 1_L\rangle\langle 0_C 1_H 0_L| + |0_C 1_H 0_L\rangle\langle 1_C 0_H 1_L|), \quad (4)$$

where $g \in \mathbb{R}$ is a coupling constant. The unitary dynamics generated by the total Hamiltonian $H = H_0 + H_{\text{int}}$ hence conserves the total energy of the clockwork, since $[H_0, H_{\text{int}}] = 0$. However, since $[H_L, H_{\text{int}}] \neq 0$, the interaction, once activated, can perform work on the ladder.

The resulting dynamics leads to an increase of the population of the top energy level $|1_L\rangle$ of the ladder, which (in units where $\hbar = 1$) is given by

$$P_{\text{top}}(t) = \text{Tr}(|1_L\rangle\langle 1_L| e^{-iHt} \rho_0 e^{iHt}). \quad (5)$$

The maximally reachable population depends on the temperatures of the baths, as well as the energy gaps of the machine qubits [15]. The top-level probability in Eq. (5) evaluates to (see Appendix A)

$$P_{\text{top}}(t) = \left(\frac{Z_H - 1}{Z_C Z_H Z_L} \right) \sin^2(gt) + \left(\frac{(Z_C - 1)(Z_L - 1)}{Z_C Z_H Z_L} \right) \cos^2(gt) + \frac{Z_L - 1}{Z_L} - \frac{(Z_C - 1)(Z_L - 1)}{Z_C Z_H Z_L}. \quad (6)$$

For $T_C = 0$, this simplifies to

$$P_{\text{top}}(t) = \left(1 - \frac{1}{Z_H} \right) \sin^2(gt). \quad (7)$$

Thus, even when $T_C = 0$, this function is far away from the ideal shape in Eq. (1), both in terms of its maximal value and the width of the distribution around its peak. Even in the limit $T_H \rightarrow \infty$, the maximal value reached at $t = \frac{\pi}{2g}$ is only $\frac{1}{2}$. Moreover, this top-level population could also have been achieved by directly coupling the ladder to the hot bath. Thus, the two-qubit machine does not provide the desired ATPC by itself. However, in the following we present a generalisation of this framework which allow arbitrarily precise ATPC, and hence an ideal clockwork to be approximated to within any given error.

II.2. Generalised Machines

In the following we present a generalised clockwork model that allows both the ‘sharpness’ and the amplitude of $P_{\text{top}}(t)$ to be controlled, while keeping track of all the relevant resources. This can be achieved by two qualitatively different but compatible extensions that we refer to as ‘horizontal’ and ‘vertical’ extensions, as illustrated in Fig. 3. The horizontal extension allows the amplitude of $P_{\text{top}}(t)$ to be increased, while the vertical extension allows the width of the peak of $P_{\text{top}}(t)$ to be decreased, thus increasing its ‘sharpness’. Specifically, we add more levels to the target ladder and with it more two-qubit machines, interacting with each successive transition (vertical extension); to a given transition we add more machines (horizontal extension). We start by collecting all interactions along a vertical column (see Fig. 3) of machines interacting with the ladder into a term H_1 . This vertically extends the interaction of a single two-qubit machine, Eq. (4), along all ladder states, *i.e.*

$$H_1 = g \sum_{n=1}^{d-1} (|1_C 0_H\rangle\langle 0_C 1_H|_{M_1^n} \otimes |n+1_L\rangle\langle n_L| + \text{H.c.}), \quad (8)$$

for the first vertical column, where M_i^j denotes the Hilbert space of the j^{th} two-qubit machine acting on the i^{th} ladder transition. We then add another term H_2 , which does the same for the second vertical column, albeit with an additional projector onto the subspace orthogonal to the one on which H_1 acts non-trivially to ensure commutativity of H_1 and H_2 . This continues for M vertical columns, always projecting onto the orthogonal subspace of all previously added machines. Using $M_{(i)}$ to denote the Hilbert space of the vertical group of the i^{th} machine, *i.e.* $M_{(i)} := \bigotimes_{j=1}^{d-1} M_i^j$, we can then write our generalisation of the interaction Hamiltonian from the previous section in a compact notation as

$$H_{\text{int}} = \sum_{k=1}^M \bigotimes_{i=1}^{k-1} \mathbb{I}_{M_{(i)}} \otimes J_{M_{(k)}L} \otimes \bigotimes_{i=k+1}^M \Pi_{M_{(i)}} = \sum_{k=1}^M H_k. \quad (9)$$




Figure 3. On the right-hand side, the notion of ‘horizontal’ and ‘vertical’ extension is schematically illustrated. The ladder dimension d determines the number of vertical extensions (‘rows’). Each pair of neighbouring energy levels of the ladder couples to M copies of the two-qubit machine, where M determines the number of horizontal extensions (‘columns’). On the left-hand side, the energy-level structure of a horizontally extended machine for a two-machine clockwork coupling to a two-level ladder ($d = 2$) is illustrated. This is a special case of the generalised clockwork shown on the right-hand side. The different energy-conserving transitions induced by H_{int} are indicated by the differently coloured arrows and dots, where the dots represent a conditioning of the transitions of machine 2 either on the ground state (orange) or on the excited state (purple) of machine 1. The reverse transitions are not depicted here for sake of clarity.

Here, we have defined the projectors

$$\Pi_{M(i)} := \mathbb{1}_{M(i)} - \sum_{n=0}^{d-1} |n_{M(i)}\rangle \langle n_{M(i)}| \quad (10)$$

and the operator

$$J_{M(k)L} := ig \sum_{n=1}^{d-1} \sqrt{n(d-n)} \left(|n_{M(k)}, n_L\rangle \langle n-1_{M(k)}, n-1_L| - \text{H.c.} \right), \quad (11)$$

and the states $|n_{M(k)}\rangle$ are defined as

$$|n_{M(k)}\rangle := \bigotimes_{j=1}^n |1_C 0_H\rangle_{M_k^j} \bigotimes_{l=n+1}^{d-1} |0_C 1_H\rangle_{M_k^l}. \quad (12)$$

That is, the state $|n_{M(k)}\rangle$ can be considered to be the n^{th} excited state of the k^{th} vertical group $M(k)$ in the sense that the first n machines M_k^j for $j = 1, \dots, n$ are in the ‘used’ state $|1_C 0_H\rangle_{M_k^j}$, whereas the remaining $d - n + 1$ machines M_k^l , with $l = n + 1, \dots, d - 1$, are in the ‘unused’ state $|0_C 1_H\rangle_{M_k^l}$. The relative normalisation factor $\sqrt{n(d-n)}$ of the different interaction terms in the Hamiltonian is precisely chosen such that the different subspace rotations are in phase to single out a $\sin^{2(d-1)}(gt)$ scaling of $P_{\text{top}}(t)$ as opposed to a mixture of different powers of trigonometric functions. For further details, see Appendix C.

In the following we will briefly discuss the horizontal and vertical extensions separately to outline their physical impact.

II.2.1. Horizontal extension

As shown in Appendix B, for $T_C = 0$, the interaction Hamiltonian for $N = 2$ in Eq. (9) then modifies the top-level probability from Eq. (7) to

$$P_{\text{top}}(t) = \left(1 - \frac{1}{Z_H^M}\right) \sin^2(gt). \quad (13)$$

For finite T_C , the weight of this sinusoidal term changes, and there are additional constant and cosine terms, whose relative weight increase with increasing T_C (see Appendix B).

From Eq. (13) we see that the maximal value of $P_{\text{top}}(t)$ increases with increasing M , and total population inversion can be achieved in the limit $M \rightarrow \infty$. However, in order to achieve ATPC, only increasing the magnitude of $P_{\text{top}}(t)$ is not sufficient since this neglects the temporal concentration. In the next section we therefore introduce the ‘vertical’ extension, which allows us to temporally concentrate $P_{\text{top}}(t)$, leading to sharper peaks.

II.2.2. Vertical extension

For the vertical extension, we generalise the ladder to a non-degenerate system with d evenly spaced energy eigenstates, with the gap between neighbouring states equal to E_L . To each of the $d - 1$ pairs of neighbouring energy levels of the vertically extended ladder, a 2-qubit machine can be coupled in the way described in the previous section. In total, the vertically extended clockwork thus consists of a d -dimensional ladder and $d - 1$ two-qubit machines, as illustrated in Fig. 3. The resulting

top-level probability for $T_C = 0$ becomes

$$P_{\text{top}}(t) = \left(\frac{Z_H - 1}{Z_H} \right)^{(d-1)} \sin^{2(d-1)}(gt). \quad (14)$$

We thus see that just vertically extending the machine makes the temporal distribution sharper, but it also decreases the achievable top-level population.

II.2.3. General extended clockwork

Finally, by combining the horizontal and vertical extension, we can combine the advantages of both, *i.e.* simultaneously increase the top-level population and the sharpness of the temporal distribution. Straightforward calculation of the top-level probability for $T_C = 0$ (shown in detail in Appendix C) yields

$$P_{\text{top}}(t) = \left\{ 1 - \left[1 - \left(\frac{Z_H - 1}{Z_H} \right)^{d-1} \right]^M \right\} \sin^{2(d-1)}(gt). \quad (15)$$

For T_C only slightly greater than zero, Eq. (15) smoothly approximates the above top-level probability (see Appendix C).

A direct consequence of the particular form of the top-level probability in Eq. (15) is that the amplitude and temporal variance (‘sharpness’) of $P_{\text{top}}(t)$ can be optimised to within any desired error by controlling the number of machines (M) per neighbouring pair of energy levels in the horizontal extension and the dimension of the ladder (d) in the vertical extension, respectively.

Since we restricted the clockwork to consist of qubit machines that have up to $M(d-1)$ -body interactions, and the Hilbert space of the machine is $4^{M(d-1)}$ -dimensional, the most reasonable quantifier of complexity should be related simply to M and d . We therefore focus on elucidating the role of M and d separately. What we can now see is that, in order to decrease the temporal variance (increasing d) while increasing the amplitude (increasing M), the complexity necessarily has to increase. Thus, for a fixed complexity there exists a trade-off between temporal variance and probability amplitude.

In the following sections we will include the irreversible decay mechanism to numerically analyse how accuracy and resolution of clocks are influenced by changes in M and d .

III. IRREVERSIBILITY AND CLOCK TICKS

Any autonomous quantum clock (or any clock for that matter) inevitably produces entropy in order to tick [11], as it needs to be subject to an irreversible evolution. While the internal clockwork produces a temporally well-concentrated and repeating distribution, there needs to be an irreversible process that turns this into a measurable signal. For this to happen there needs to be a system that is driven out of equilibrium in order to relax back to

equilibrium while producing a tick. In our case the system that is driven out of equilibrium is the ladder. The system with respect to which the ladder is driven out of equilibrium is assumed to couple to the ladder such that the top level is unstable and decays, emitting energy into that system. As an example, one can take this system to be a photon field at the environment temperature T_C that couples to the ladder, such that when the top level population decays to the ground state it emits a photon of energy $E_\gamma = (d-1)E_L$. However, note that the assumption that only this particular ladder transition couples to the field is an idealization for the purpose of deriving fundamental bounds (as discussed in Sec. II). Since any irreversible process can be viewed as a reversible process on a larger Hilbert space, the presence of such a decay channel must in principle also allow for the reverse process of exciting the ladder while absorbing energy, *e.g.*, in the form of a photon in the example above. However the probability for this to happen can be made arbitrarily small by demanding that the background temperature of the field satisfies $\frac{E_\gamma}{k_B T_C} \gg 1$.

The number of possible decay processes is vast. However, since our aim is to capture all resources that are necessary to operate a clock, allowing for decay processes that require memory would miss the purpose, since the required resources are not clearly defined for them. We therefore require the photon field to be memoryless, *i.e.* that correlations with the ladder are diluted very quickly and are thus negligible. The resulting dynamics are governed by the law of exponential decay, and thus constitute an ideal case, giving an effective upper bound to the clock performance and allowing us to keep track of the resources that are invested. In particular, the probability density for a tick occurring at time t is given by (see Appendix D)

$$P_{\text{tick}}(t) = c P_{\text{top}}(t) e^{-c \int P_{\text{top}}(t') dt'}, \quad (16)$$

where c is the coupling strength of the photon field with the top level of the ladder.

Let us now consider the energetic resources required to run the clock. We first note that, taking $T_C = 0$, as the clockwork state evolves, each branch of its superposition where the ladder’s top level is excited corresponds to a transition of $d-1$ machines: $|0_{M(k)}\rangle \rightarrow |d-1_{M(k)}\rangle$, where the value of k differs between branches, and we recall that the $\{|n_{M(k)}\rangle\}$ were defined in Eq. (12). Thus, regardless of which branch is realised, if the ladder’s top level is excited then the heat flow from the hot bath into the total system is given by $Q_{\text{in}} = (d-1)E_H$. This heat flow does the work of driving the ladder from $|0\rangle_L$ to $|d-1\rangle_L$, *i.e.* $W = (d-1)E_L$. After the clock has ticked and the cold qubits of the machines re-thermalise, $Q_{\text{out}} = (d-1)E_C$ of heat will be dissipated into the cold bath. Since $E_H = E_L + E_C$, we thus have the usual relation for a thermal machine, *i.e.* $Q_{\text{in}} = W + Q_{\text{out}}$, and the thermal efficiency of the process is $\eta_{\text{th}} := \frac{W}{Q_{\text{in}}} = \frac{E_L}{E_L + E_C}$. From this, one can see that as E_C/E_L decreases, we approach the Carnot efficiency bound $\eta_{\text{th}} \leq 1$.

Curiously, for $T_C = 0$ and $M \rightarrow \infty$, the top-level population is just $P_{\text{top}}(t) = \sin^{2(d-1)}(gt)$. If interpreted as a heat engine whose purpose is to charge a battery (the ladder), then one can indeed reach an efficiency of $\eta_{\text{th}} \approx 1$ and still charge the battery in finite time $\tau = \frac{\pi}{2g}$. Even the task of ATPC can be achieved to arbitrary precision at perfect efficiency. One can interpret this as sufficient clockwork complexity permitting perfect efficiency at finite power.

In any case, the efficacy of ATPC and the resulting clock dynamics are essentially determined by the ladder dimension d and the number of driving machines M , which together correspond to a simple notion of *clockwork complexity*. In order to investigate how these affect the quality of the clock, we quantify this quality using two notions introduced in [11]. These are the accuracy, which is the average number of ticks until the next tick is off by the average time between two ticks, *i.e.*

$$N = \left(\frac{\bar{t}}{\Delta t} \right)^2, \quad (17)$$

and the resolution, which is the inverse average time between two ticks, *i.e.*

$$R = \frac{1}{\bar{t}}. \quad (18)$$

Here, $\bar{t} = \int_0^\infty t P_{\text{tick}}(t) dt$ and $(\Delta t)^2 = \overline{t^2} - (\bar{t})^2$ with $\overline{t^2} = \int_0^\infty t^2 P_{\text{tick}}(t) dt$.

Let us remark here that our choice of quantifier for the resolution R is not the only option. For instance, another candidate to quantify the resolution would be $1/\bar{t} = \int_0^\infty \frac{1}{t} P_{\text{tick}}(t) dt$. However, for this choice, small times would contribute much more strongly to the average than for the inverse of the average times. For example, already a single outlier at a very small value of t would result in a very large value of $1/\bar{t}$, and one would conclude that the resolution was very high, even if most of the events were observed at larger values of t . Conversely, choosing $1/\bar{t}$ as a quantifier means that even a few outliers at large values t would result in a low resolution. Therefore, our choice $R = 1/\bar{t}$ represents the more conservative of these choices, ensuring that our description results in upper bounds on the resolution.

In the following section, we present numerical calculations of the accuracy as a function of the resolution, the clockwork complexity and the energy dissipated per tick.

For comparison, let us take as a baseline an example where there are no qubit machines employed, and the ladder simply begins in equilibrium with the hot bath and emits this energy into the cold bath via the irreversible process, *i.e.* there is no ATPC. In that case, the top-level probability is constant, *i.e.* $P_{\text{top}}(t) = \exp[-\beta_H(d-1)E_L]/Z_L$, which results in $R = \frac{1}{\bar{t}} = \frac{1}{\Delta t} = c \exp[-\beta_H(d-1)E_L]/Z_L$, and thus the resolution is essentially determined by the decay rate c and the population of the decaying level. The accuracy is simply $N = 1$.

This highlights the main purpose of a clockwork: An individual event resulting from pure thermalisation would result in an accuracy of 1 and come at a work cost of $(d-1)E_L$, whereas the clockwork can increase the accuracy while keeping the work cost of one tick constant.

IV. NUMERICAL RESULTS

Since we are interested in upper bounds on the clock quality, for the following results we assume the temperature of the hot bath to be infinite, *i.e.* $T_H \rightarrow \infty$, as well as $T_C = 0$. The curves in the following figures are generated numerically by varying three free parameters, namely M , d and c (varying g has the inverse effect of varying c – see Appendix F). In particular, each curve corresponds to fixed values of M and c while d varies; we display the accuracy N on the vertical axes, and the ladder dimension d (Fig. 4), the resolution R (Fig. 5), and the energy dissipation rate ϵ (Fig. 6), respectively, on the horizontal axes.

First we analyse the relation between the ‘sharpness’ of the peak of $P_{\text{top}}(t)$ and the clock accuracy. Recalling the discussion in Sec. II.2.2, we note that the ‘sharpness’ of $P_{\text{top}}(t)$ increases with increasing d , and the latter may therefore be used as a measure of the ‘sharpness’ of $P_{\text{top}}(t)$. In Fig. 4, the behaviour of the accuracy as d increases is depicted for different ladder-bath coupling strengths c and different values of M . For small d we see that the accuracy increases linearly. However, increasing the ‘sharpness’ beyond a certain point leads to a decrease in accuracy (this behaviour will be discussed in detail in Sec. V). The value of c therefore puts a bound on the maximally achievable accuracy for all potential clocks. The same limiting behaviour is apparent if we fix c and vary g instead (see Appendix F), for reasons that we discuss below.

In order to analyse both accuracy and resolution with respect to the clockwork complexity in Fig. 5, we compare those two quantities for different fixed values of M while varying d . Increasing M allows us to reach a higher max-




Figure 4. The effect of the ladder dimension d on the clock accuracy N for different M . The top three curves show this in the limit $M \rightarrow \infty$ for different c , where $[c] = s^{-1}$.




Figure 5. The trade-off between clock accuracy N and resolution R for clockworks of various complexities, where R is increased by decreasing d . For finite horizontal extensions M , we show this behaviour for a fixed photon field coupling $c = 25$, where $[c] = \text{s}^{-1}$. Increasing M allows for higher maximal accuracy, and thus the orange line represents an upper bound of the accuracy for a given resolution for all clocks with $c = 25$. For $M \rightarrow \infty$, we see that increasing c increases the potentially achievable combinations accuracy and resolution. We have chosen $g = 1 E_C$ in all cases.




Figure 6. The accuracy N as a function of the energy dissipation rate $\varepsilon = Q_{\text{out}}/\bar{t}$ (with $[\varepsilon] = E_C \text{s}^{-1}$) for different numbers of horizontal extensions M , with $c = 10^5 \text{s}^{-1}$. As M increases, the maxima of the curves obtained by varying d are shifted towards higher values of N and ε , *i.e.*, up and to the right, reaching finite values (not shown) in the limit $M \rightarrow \infty$ (black curve). The individual peaks obtained for a given coupling strength c and number M of horizontal extensions correspond to the clocks that achieve maximal accuracy under these constraints. We exclude sub-optimal cases, *i.e.* cases where increasing d reduces resolution and accuracy.

imal accuracy, while increasing d (which increases from right to left in Fig. 5) allows us to trade resolution for accuracy up to the optimal point, after which the accuracy reduces again. We further observe that all clocks with the same c and g lie under a curve defined by the case of $M \rightarrow \infty$. Increasing c allows for clocks of higher quality, *i.e.* that have higher accuracy and resolution. Furthermore, the position of the maximum depends on the value of g . Here g was chosen to be equal to $1 E_C$.

Increasing this value shifts the peak to the right, *i.e.* to higher resolutions (see Appendix F).

Finally, in order to analyse the effect of the energy dissipation rate $\varepsilon = Q_{\text{out}}/\bar{t}$ on the clock accuracy with respect to different values of M and d , which are simply related to the complexity of the clockwork (see Sec. II.2.3), in Fig. 6 we plot the accuracy over the energy dissipation rate ε for clockworks of different complexity. In particular, we compare different values of M while varying d . What we can see in Fig. 6 is that for fixed M (at fixed c and g), increasing the energy dissipation rate (which is achieved by increasing d) increases the accuracy at a certain slope until a maximum is reached. Increasing d further decreases the accuracy. Furthermore, for a given c , increasing M leads to a lower slope (approaching the slope of $M \rightarrow \infty$) while allowing for higher maximal accuracy, which suggests that a greater maximal accuracy can be achieved at the cost of a greater energy dissipation rate. We also note that increasing c increases the maximally achievable accuracy N , as can also be seen in Figs. 4 and 5. At the same time, increasing c increases the resolution (which can be seen in Fig. 6). Since Q_{out} does not depend on c , this increase in resolution leads to a higher energy-dissipation rate for all values of d and M , thus increasing the slope of the N - ε curve (in the linear regime). For reasons of visual clarity this is not depicted in Fig. 6. One should note, however, that increasing c indefinitely would eventually break the assumptions inherent in our analysis and force us to explore deviations from a Markovian exponential decay towards memory effects in the irreversible dynamics.

V. DISCUSSION

Our results have two general implications. The first concerns the task of autonomous probability concentration. We show that, in principle, sufficiently increasing the clockwork complexity alone is enough to concentrate the temporal probability arbitrary well. In particular, this task can again be split into two conceptually different sub-tasks: maximizing the achievable population and improving the temporal ‘sharpness’ with which this (maximal) population can be reached. By splitting the clockwork into a target ladder and virtual-qubit machines coupled to the different ladder transitions, we were able to analyse how more complex clockworks can help to achieve the two respective sub-tasks. While we have worked with equally-spaced ladder systems, the same result (qualitatively) also holds for arbitrarily spaced target Hamiltonians, simply by redefining the respective coupling strengths (the g ’s) of the interaction Hamiltonians. We have equipped our clockwork with a particular tensor product structure, the division into two-qubit machines, for the sake of keeping track of its complexity. Our machine operates optimally within the framework set by this division, but whether more general machines could also achieve the same performance with smaller overall size

remains an open question.

In our analysis, we have optimised the internal structure of the clock, *i.e.* the clockwork, to concentrate the probability in a fashion that most closely resembles the temporal distribution of an ideal clockwork. For given c , this amplifies the clock quality only up to a limit, which we showcase in Fig. 4. This can intuitively be understood by considering the two key timescales of the clock, namely that of the clockwork's dynamics, and that of the irreversible decay. Increasing d while keeping c fixed, one eventually reaches a point where $P_{\text{top}}(t)$ is so well concentrated temporally that the comparative slowness of the decay mechanism reduces the probability that the clock will tick. In other words, it becomes more likely that the decay mechanism will skip that peak. We see the inverse of this behaviour if instead of c , we consider curves of fixed g (see Appendix F), as increasing g speeds up the clockwork, effectively making the limit imposed by c more restrictive.

This brings us to the second implication of our work. The irreversible mechanism, in our case characterised by the parameter c , puts an absolute upper bound on the achievable combinations of resolution and accuracy, *i.e.* the clock quality, and thus determines the potential for how well a particular physical process can be used as the basis for a clock. The question of how well this upper bound can be approximated brings us to the role of our two extensions. First of all, the horizontal extension, *i.e.* the coupling of multiple elementary machines to a single transition between neighbouring ladder levels primarily serves the purpose of increasing the possible population inversion and with it the achievable top-level population. As we see in Eq. (16), c always appears multiplied by the prefactor of the sine in Eq. (15), resulting in an *effective coupling*:

$$C_M = c \left\{ 1 - \left[1 - \left(\frac{Z_H - 1}{Z_H} \right)^{d-1} \right]^M \right\}. \quad (19)$$

From this we can see that increasing the horizontal extension M is physically equivalent to increasing the coupling c , and this is why they play the same role in Fig. 4 (though we note that C_M is bounded with respect to M but not with respect to c). One cannot make a similar statement to relate c with the vertical extension d , since the exponent of the sine in Eq. (15) will vary as d does. As noted above, d sharpens the temporal distribution, thus increasing the accuracy of the clock as long as the limit set by c and M [via Eq. (19)] is not surpassed, as demonstrated in Fig. 4.

In the regime where the accuracy grows linearly with the 'sharpness' (Fig. 4), which is determined by d (see Appendix F), there exists a trade-off relation between accuracy and resolution. To see this, note that the resolution decreases monotonically with d [see Fig. 7(b)]. For fixed c and g , the case of $M \rightarrow \infty$ represents an upper bound on the clockwork quality, *i.e.* on the achievable combinations of accuracy and resolution, which is illustrated in Fig. 5.

In our computations, we have focused on the limit $T_C \rightarrow 0$. State-of-the-art atomic clocks operate at optical frequencies [16], and at even higher frequencies in novel proposals [17]. The vacuum state of any optical-frequency mode of the electromagnetic field has a population of ≈ 1 at room temperature, and this situation is thus virtually indistinguishable from temperature 0. For clocks operating at much lower frequencies, such as those based on microwave transitions, cooling the environmental degrees of freedom into which the irreversible mechanism dissipates heat would be necessary in order to approach the fundamental limits we derive here.

It is nonetheless important to stress that we are interested in fundamental limits of timekeeping and the associated complexity and cost, which one of the reasons why we consider autonomous clocks in contrast with atomic clocks, for example, which require external control. For the practical purpose of building clocks for everyday use, atomic clocks can require as little as 30 mW [18], which is many orders of magnitude above the scale defined by the system energy and the timescale of the relevant processes, but still insignificant for global energy use. The majority of that cost, however, comes from the fact that at some point that single event needs to be amplified and registered by a measurement apparatus, whose inherent irreversible nature is also thermodynamic and comes with its own costs and limitations [19]. Conversely, the inevitable imperfections of clocks and the associated costs also limit the achievable quality of measurements and, consequently, of all estimation procedures, *e.g.*, of work itself [20].

As far as current and future prospects are concerned the limits derived in this paper have more fundamental relevance for the autonomous control of quantum systems by a quantum clock [21, 22]. Here, a small quantum system is envisioned to be controlled by an autonomous quantum clock. This is important for any type of unitary process requiring precise timing, from small machines operating in cycles [23] to general repeating unitary processes such as circuit-based cooling models [24–28].

Coming back to the actual energy cost, there are a number of interesting observations that follow from our clock model. First of all, the horizontal extension always comes at a finite energy cost and dissipation for each tick for any clock, while the vertical extensions linearly increase the energy cost and dissipation. The vertical extensions on the other hand linearly increases the energy cost and dissipation and, as long as the limit imposed by the c and M is not exceeded, also increases the accuracy. This limit can be qualitatively understood as a point after which further concentrating probability is actually detrimental to clock performance. We can observe that the increase in accuracy follows a linear behaviour in d , before switching to a sublinear behaviour close to the peak, after which it actually decreases accuracy again. Thus, in addition to recovering the observation that $N \propto \Delta S$, *i.e.* that a clock's accuracy is essentially determined by the entropy it dissipates (which seems to be a prevalent

feature in all classical and quantum clocks [12]), we have pinpointed which combination of the resources M , d and c allow us to maintain this linear regime. The proportionality factor itself can be identified numerically from the plots in Fig. 6.

Finally, let us put our clockwork model in context with recent literature on quantum clocks. In Refs. [29–32], the relationship between achievable clock accuracy as a function of ‘clock dimension’ is studied by means of repeated applications of maps from a clock system to a register. These works provide fundamental bounds for clock accuracy for fixed system dimension d , showing that the accuracy of classical (incoherent) clocks can at best scale linearly in d , whereas a quantum clock’s accuracy (with states featuring coherence) may scale as d^2 . The clock system considered in these works is exactly what we here refer to as the ladder system. Meanwhile, the map that Refs. [29–33] refer to as being responsible for creating a tick event in the register subsumes the interactions between the ladder, the qubit machines, and the heat baths, and also includes the irreversible mechanism and the subsequent read-out. In other words, our work provides a concrete physical realisation of the maps that effect the transfer of ticks to the register. In Fig. 4, we see that, in the regime of the clockwork not exceeding the clock potential dictated by the irreversible mechanism, the accuracy scales linearly with the ladder dimension d , which is already the optimal achievable scaling [32].

VI. CONCLUSION

In this article, we have put forward a framework for studying fundamental limits of timekeeping. The conceptual split of any such task into a *clockwork*, which creates a temporally concentrated probability distribution, and a mechanism for irreversibility, allowed us to derive an analytic formula for the achievable temporal probability concentration of the clockwork. The irreversible mechanism provides a context for the operation of the clock by allowing the passage of time to be tracked in the first place. Meanwhile, the chosen irreversible mechanism sets the reference timescale that ultimately constrains the potential of any clockwork that harnesses this mechanism to form a clock. But it is the clockwork that needs to

be appropriately tuned to achieve maximal performance given these constraints. By composing the clockwork of the smallest possible thermal machines, we were further able to conceptually split the task of autonomous probability concentration into two sub-tasks. First, by having more machines interact with a single transition, we can increase the maximum top-level probability and with it, the effective coupling to the irreversible mechanism. Second, by concatenating multiple transitions of this kind, we are able to sharpen the temporal distribution. This reveals the intricate ways in which the complexity of the clockwork determines its performance.

In the future, it will be interesting to study more exotic irreversible mechanisms beyond exponential decay and whether they could be harnessed to further improve the clock quality. Moreover, one might study increasing clockwork complexity in a manner other than the addition of qubits. Furthermore, beyond simply asking if optimal ATPC is achievable in some limit, it would be of interest to ask which temporal probability profile (*i.e.* $P_{\text{tick}}(t)$) optimizes clock performance under some constraints, and the extent to which such a profile is achievable with a quantum machine. Another interesting avenue would be to make a stronger connection to experimental implementations of clocks and how the notion of irreversibility versus temporal probability concentration can be more formally made on larger scales. One experiment in that direction is performed in Ref. [12], using a nano-mechanical membrane. Photoisomerisation could also provide a potential platform for implementing such a clock at molecular scales [34]. While there are many open paths to explore and questions to answer, our results consolidate the fact that perfect clocks are practically impossible when derived from first principles and that significant thermodynamic resources have to be invested to reach the potential of any physical system to act as a clock.

Acknowledgements. We acknowledge support from the Austrian Science Fund (FWF) through the START project Y879-N2, the Zukunftskolleg ZK03 and the project P 31339-N27. We further acknowledge funding from the ESQ Discovery grant “*Emergence of physical laws: From mathematical foundations to applications in many body physics*”.

-
- [1] Wolfgang Pauli, *Die allgemeinen Prinzipien der Wellenmechanik* (Springer, Berlin, Heidelberg, 1958).
 - [2] John C. Garrison and Jack Wong, *Canonically Conjugate Pairs, Uncertainty Relations, and Phase Operators*, *J. Math. Phys.* **11**, 2242 (1970).
 - [3] Samuel L. Braunstein, Carlton M. Caves, and Gerard J. Milburn, *Generalized Uncertainty Relations: Theory, Examples, and Lorentz Invariance*, *Ann. Phys. (N. Y.)* **247**, 135 (1996), [arXiv:quant-ph/9507004](#).
 - [4] G. J. Milburn and T. J. Milburn, *A quantum optomechanical Mach clock*, [arXiv:1708.02369](#) [quant-ph] (2017).
 - [5] Paul Erker, *The Quantum Hourglass - Approaching time measurement with quantum information theory*, [Master’s thesis](#), ETH Zurich (2014).
 - [6] Andre C. Barato and Udo Seifert, *Cost and Precision of Brownian Clocks*, *Phys. Rev. X* **6**, 041053 (2016), [arXiv:1610.07960](#).
 - [7] Shishir Khandelwal, Maximilian P. E. Lock, and Mischa P. Woods, *Universal quantum modifications to gen-*

- eral relativistic time dilation in delocalised clocks, *Quantum* **4**, 309 (2020), [arXiv:1904.02178](#).
- [8] Alexander R. H. Smith and Mehdi Ahmadi, *Quantum clocks observe classical and quantum time dilation*, *Nat. Commun.* **11**, 1 (2020), [1904.12390](#).
- [9] Alexander R. H. Smith and Mehdi Ahmadi, *Quantizing time: Interacting clocks and systems*, *Quantum* **3**, 160 (2019), [arXiv:1712.00081](#).
- [10] G. J. Milburn, *The thermodynamics of clocks*, *Contemp. Phys.* **61**, 69 (2020), [arXiv:2007.02217](#).
- [11] Paul Erker, Mark T. Mitchison, Ralph Silva, Mischa P. Woods, Nicolas Brunner, and Marcus Huber, *Autonomous Quantum Clocks: Does Thermodynamics Limit Our Ability to Measure Time?* *Phys. Rev. X* **7**, 031022 (2017), [arXiv:1609.06704](#).
- [12] Anna N. Pearson, Yelena Guryanova, Paul Erker, Edward A. Laird, G. Andrew Davidson Briggs, Marcus Huber, and Natalia Ares, *Measuring the thermodynamic cost of timekeeping*, [arXiv:2006.08670](#) [quant-ph] (2020).
- [13] John Goold, Marcus Huber, Arnau Riera, Lidia del Rio, and Paul Skrzypczyk, *The role of quantum information in thermodynamics — a topical review*, *J. Phys. A* **49**, 143001 (2016), [arXiv:1505.07835](#).
- [14] Mark T. Mitchison, *Quantum thermal absorption machines: refrigerators, engines and clocks*, *Contemp. Phys.* **60**, 164 (2019), [arXiv:1902.02672](#).
- [15] Nicolas Brunner, Noah Linden, Sandu Popescu, and Paul Skrzypczyk, *Virtual qubits, virtual temperatures, and the foundations of thermodynamics*, *Phys. Rev. E* **85**, 051117 (2012), [arXiv:1106.2138](#).
- [16] Andrew D. Ludlow, Martin M. Boyd, Jun Ye, E. Peik, and Piet O. Schmidt, *Optical atomic clocks*, *Rev. Mod. Phys.* **87**, 637 (2015), [arXiv:1407.3493](#).
- [17] Benedict Seiferle, Lars von der Wense, Pavlo V. Bilous, Ines Amersdorffer, Christoph Lemell, Florian Libisch, Simon Stellmer, Thorsten Schumm, Christoph E. Düllmann, Adriana Pálffy, and Peter G. Thirolf, *Energy of the ^{229}Th nuclear clock transition*, *Nature* **573**, 243 (2019), [arXiv:1905.06308](#).
- [18] R. Lutwak, J. Deng, W. Riley, M. Varghese, J. Leblanc, G. Tepolt, M. Mescher, D. K. Serkland, K. M. Geib, and G. M. Peake, *The chip-scale atomic clock-low-power physics package*, Tech. Rep. (Symmetricom – Technology Realization Center, Beverly, MA, 2004).
- [19] Yelena Guryanova, Nicolai Friis, and Marcus Huber, *Ideal Projective Measurements Have Infinite Resource Costs*, *Quantum* **4**, 222 (2020), [arXiv:1805.11899](#).
- [20] Tiago Debarba, Gonzalo Manzano, Yelena Guryanova, Marcus Huber, and Nicolai Friis, *Work estimation and work fluctuations in the presence of non-ideal measurements*, *New J. Phys.* **21**, 113002 (2019), [arXiv:1902.08568](#).
- [21] Artur S. L. Malabarba, Anthony J. Short, and Philipp Kammerlander, *Clock-driven quantum thermal engines*, *New J. Phys.* **17**, 045027 (2015), [arXiv:1412.1338](#).
- [22] Harrison Ball, William D. Oliver, and Michael J. Biercuk, *The role of master clock stability in scalable quantum information processing*, *npj Quantum Inf.* **2**, 16033 (2016), [1602.04551](#).
- [23] T. Feldmann and J.P. Palao, *Performance of Quantum Thermodynamic Cycles*, in *Thermodynamics in the Quantum Regime: Fundamental Aspects and New Directions*, edited by Felix Binder, Luis A Correa, Christian Gogolin, Janet Anders, and Gerardo Adesso (Springer International Publishing, Cham, Switzerland, 2019) Chap. 3, pp. 67–85.
- [24] Sadegh Raeisi, Mária Kieferová, and Michele Mosca, *Novel Technique for Robust Optimal Algorithmic Cooling*, *Phys. Rev. Lett.* **122**, 220501 (2019), [arXiv:1902.04439](#).
- [25] Álvaro M. Alhambra, Matteo Lostaglio, and Christopher Perry, *Heat-Bath Algorithmic Cooling with optimal thermalization strategies*, *Quantum* **3**, 188 (2019), [arXiv:1807.07974](#).
- [26] Fabien Clivaz, Ralph Silva, Géraldine Haack, Jonatan Bohr Brask, Nicolas Brunner, and Marcus Huber, *Unifying Paradigms of Quantum Refrigeration: A Universal and Attainable Bound on Cooling*, *Phys. Rev. Lett.* **123**, 170605 (2019), [arXiv:1903.04970](#).
- [27] Fabien Clivaz, Ralph Silva, Géraldine Haack, Jonatan Bohr Brask, Nicolas Brunner, and Marcus Huber, *Unifying paradigms of quantum refrigeration: fundamental limits of cooling and associated work costs*, *Phys. Rev. E* **100**, 042130 (2019), [arXiv:1710.11624](#).
- [28] Nayeli Azucena Rodríguez-Briones and Raymond Laflamme, *Achievable Polarization for Heat-Bath Algorithmic Cooling*, *Phys. Rev. Lett.* **116**, 170501 (2016), [arXiv:1412.6637](#).
- [29] Sandra Ranković, Yeong-Cherng Liang, and Renato Renner, *Quantum clocks and their synchronisation - the Alternate Ticks Game*, [arXiv:1506.01373](#) [quant-ph] (2015).
- [30] Sandra Stupar, Christian Klumpp, Renato Renner, and Nicolas Gisin, *Performance of stochastic clocks in the Alternate Ticks Game*, [arXiv:1806.08812](#) [quant-ph] (2018).
- [31] Mischa P. Woods, Ralph Silva, Gilles Pütz, Sandra Stupar, and Renato Renner, *Quantum clocks are more accurate than classical ones*, [arXiv:1806.00491](#) [quant-ph] (2018).
- [32] Yuxiang Yang and Renato Renner, *Ultimate limit on time signal generation*, [arXiv:2004.07857](#) [quant-ph] (2020).
- [33] Mischa P. Woods, *Autonomous Ticking Clocks from Axiomatic Principles*, *Quantum* **5**, 381 (2021), [arXiv:2005.04628](#).
- [34] Nicole Yunger Halpern and David T. Limmer, *Fundamental limitations on photoisomerization from thermodynamic resource theories*, *Phys. Rev. A* **101**, 042116 (2020), [arXiv:1811.06551](#).
- [35] Michael Morrison and Gregory Parker, *A Guide to Rotations in Quantum Mechanics*, *Aust. J. Phys.* **40**, 465 (1987).
- [36] E. P. Wigner and J. J. Griffin, *Group Theory and Its Application to the Quantum Mechanics of Atomic Spectra*, Pure and applied Physics (Academic Press, 1959).
- [37] Joseph Wiener, *Integrals of $\cos^{2n}(x)$ and $\sin^{2n}(x)$* , *The College Mathematics Journal* **31**, 60 (2000)

APPENDICES

In these appendices, we provide detailed derivations and background information for the results presented in the main text. In Appendix A we first present a derivation of the top-level probability for the two-qubit machine of Sec. II.1. In Appendix B, we then present the derivation of the top-level probability for the horizontal extension of the clockwork for arbitrary temperatures T_H and T_C . In Appendix C, we then again focus on the case $T_C = 0$, for which we derive the top-level probability in the full horizontal and vertical extension. Appendix D contains the derivation of the tick probability. Appendix E presents the details for the numerical computation of accuracy and resolution. Appendix F discusses the behaviour of clocks with changing ladder dimension d as well as changing coupling constant g .

Appendix A: Top-level probability of a two-qubit clockwork

Here, we present a derivation of the top-level probability $P_{\text{top}}(t)$ from Eq. (7). That is, we consider the minimal clockwork discussed in Sec. II.1, which consists of a hot qubit (coupling to the hot bath at temperature T_H , as well as a cold qubit and a ladder (both coupling to a cold bath at temperature T_C). The derivation presented here is a special case ($M = 1$ and $T_C = 0$) of the more general derivation of the top-level probability within the horizontal extension that we will present in Appendix B (where $M > 1$ and both T_H and T_C can take on arbitrary values). Nevertheless, we will first go through the much simpler derivation for $M = 1$ and $T_C = 0$ here, which will serve as a guiding example for the much more involved general calculation that is to follow.

Assuming that the systems have thermalised with their respective baths, the initial state of the clockwork is given by

$$\rho_0 = |0\rangle\langle 0|_C \otimes \tau_H \otimes |0\rangle\langle 0|_L = |0\rangle\langle 0|_C \otimes \frac{1}{Z_H} (|0\rangle\langle 0|_H + e^{-\beta_H E_H} |1\rangle\langle 1|_H) \otimes |0\rangle\langle 0|_L, \quad (\text{A1})$$

where $Z_H = 1 + e^{-\beta_H E_H}$ is the partition function of the hot qubit. Since the interaction term H_{int} in the total Hamiltonian $H = H_0 + H_{\text{int}}$ is chosen such that the free energy (that is, w.r.t. the free Hamiltonian H_0) is conserved, $[H_0, H_{\text{int}}] = 0$, and because the initial state ρ_0 is diagonal in the eigenbasis of H_0 , we can write the top-level probability from Eq. (5) as

$$P_{\text{top}}(t) = \text{Tr}(|1\rangle\langle 1|_L e^{-iHt} \rho_0 e^{iHt}) = \text{Tr}(|1\rangle\langle 1|_L e^{-iH_{\text{int}}t} \rho_0 e^{iH_{\text{int}}t}) = \text{Tr}(T_0 |0\rangle\langle 0|_C \otimes \tau_H T_0^\dagger), \quad (\text{A2})$$

where T_0 is a matrix encoding the transition amplitude between ground state and excited state of the ladder, given by

$$T_0 = {}_L\langle 1| e^{-iH_{\text{int}}t} |0\rangle_L = \sum_{k=0}^{\infty} \frac{(-it)^k}{k!} {}_L\langle 1| H_{\text{int}}^k |0\rangle_L. \quad (\text{A3})$$

Now, because H_{int}^2 is proportional to the identity on its support, that is,

$$H_{\text{int}}^2 = g^2 (|1\rangle\langle 0|_C \otimes |0\rangle\langle 1|_H \otimes |1\rangle\langle 0|_L + \text{H.c.})^2 = g^2 (|1\rangle\langle 1|_C \otimes |0\rangle\langle 0|_H \otimes |1\rangle\langle 1|_L + |0\rangle\langle 0|_C \otimes |1\rangle\langle 1|_H \otimes |0\rangle\langle 0|_L), \quad (\text{A4})$$

even powers of H_{int} do not contribute to T_0 . However, for odd powers we have $H_{\text{int}}^{2k+1} = g^{2k} H_{\text{int}}$, such that we find ${}_L\langle 1| H_{\text{int}}^{2k+1} |0\rangle_L = g^{2k+1} |1\rangle\langle 0|_C \otimes |0\rangle\langle 1|_H$. With this, we can evaluate the transition matrix T_0 , *i.e.*

$$T_0 = \sum_{k=0}^{\infty} \frac{(-it)^{2k+1}}{(2k+1)!} {}_L\langle 1| H_{\text{int}}^{2k+1} |0\rangle_L = \sum_{k=0}^{\infty} \frac{(-igt)^{2k+1}}{(2k+1)!} |1\rangle\langle 0|_C \otimes |0\rangle\langle 1|_H = \sin(gt) |1\rangle\langle 0|_C \otimes |0\rangle\langle 1|_H. \quad (\text{A5})$$

Inserting the result into Eq. (A2), we finally arrive at the top-level probability

$$P_{\text{top}}(t) = \text{Tr}(T_0 |0\rangle\langle 0|_C \otimes \tau_H T_0^\dagger) = \sin^2(gt) {}_L\langle 1| \tau_H |0\rangle_L = \sin^2(gt) \frac{e^{-\beta_H E_H}}{1 + e^{-\beta_H E_H}} = \sin^2(gt) \left(1 - \frac{1}{Z_H}\right). \quad (\text{A6})$$

Appendix B: The horizontal extension

In this appendix, we present more technical details of the horizontal extension of the autonomous clockwork from a single ($M = 1$) to multiple ($M > 1$) two-qubit machines interacting with the same two-level ($d = 2$) transition of

the ladder system. In particular we will derive the top-level probability for the horizontal extension for arbitrary temperatures T_C and T_H .

Following a similar approach as in Eq. (A2), we define transition operators

$$T_n := {}_{\text{L}}\langle 1 | e^{-iH_{\text{int}}t} | n \rangle_{\text{L}} = {}_{\text{L}}\langle 1 | \sum_{j=0}^{\infty} \frac{(-it)^j}{j!} H_{\text{int}}^j | n \rangle_{\text{L}} = {}_{\text{L}}\langle 1 | \sum_{j=0}^{\infty} \frac{(-it)^j}{j!} \sum_{k=1}^M H_k^j | n \rangle_{\text{L}}, \quad (\text{B1})$$

for $n = 0, 1$, where the last equality follows from the fact that the interaction terms H_k given by the terms in Eq. (9) have mutually disjoint support, *i.e.* $H_k H_{k'} = 0$ for $k \neq k'$. Before we calculate these transition operators, we note that H_k satisfies the cyclic property

$$H_k^{2q} = g^{2q} \bigotimes_{i=1}^{k-1} \mathbb{1}_{M_i} \otimes (|0_C 1_H\rangle\langle 0_C 1_H|_{M_k} \otimes |0\rangle\langle 0|_{\text{L}} + |1_C 0_H\rangle\langle 1_C 0_H|_{M_k} \otimes |1\rangle\langle 1|_{\text{L}}) \otimes \bigotimes_{j=k+1}^M \Pi_{M_j} \quad \text{for } q \in \mathbb{N}_{>0}, \quad (\text{B2})$$

$$H_k^{2q+1} = g^{2q+1} \bigotimes_{i=1}^{k-1} \mathbb{1}_{M_i} \otimes (\sigma_{M_k}^- \otimes \sigma_{\text{L}}^+ + \sigma_{M_k}^+ \otimes \sigma_{\text{L}}^-) \otimes \bigotimes_{j=k+1}^M \Pi_{M_j} = g^{2q+1} H_k \quad \text{for } q \in \mathbb{N}, \quad (\text{B3})$$

where $M_i = M_i^1 = C_i^1 \otimes H_i^1$ denotes the Hilbert space of the i^{th} horizontal extension, $\sigma_{\text{L}}^+ := |1\rangle\langle 0|_{\text{L}} = (\sigma_{\text{L}}^-)^{\dagger}$, $\sigma_{M_k}^+ := |0_C 1_H\rangle\langle 1_C 0_H|_{M_k} = (\sigma_{M_k}^-)^{\dagger}$, and $|l_C m_H\rangle\langle p_C q_H|_{M_k} := |l\rangle\langle p|_{C_k} \otimes |m\rangle\langle q|_{H_k}$, with $l, m, p, q = 0, 1$.

Now, for the transition operator T_0 , we note that only odd powers of H_k can map $|0\rangle_{\text{L}}$ to $|1\rangle_{\text{L}}$, and therefore

$$T_0 = \sum_{q=0}^{\infty} \sum_{k=1}^M \frac{(-it)^{2q+1}}{(2q+1)!} {}_{\text{L}}\langle 1 | H_k^{2q+1} | 0 \rangle_{\text{L}} = \sin(gt) \left(\sum_{k=1}^M \bigotimes_{j=1}^{k-1} \mathbb{1}_{M_j} \otimes \sigma_{M_k}^- \otimes \bigotimes_{l=k+1}^M \Pi_{M_l} \right). \quad (\text{B4})$$

We can calculate T_1 similarly, noting that only even powers of H_k contain the factor $|1\rangle\langle 1|_{\text{L}}$.

$$T_1 = \sum_{q=0}^{\infty} \sum_{k=1}^M \frac{(-it)^{2q}}{(2q)!} {}_{\text{L}}\langle 1 | H_k^{2q} | 1 \rangle_{\text{L}} = \cos(gt) \left(\sum_{k=1}^M \bigotimes_{j=1}^{k-1} \mathbb{1}_{M_j} \otimes |1_C 0_H\rangle\langle 1_C 0_H|_{M_k} \otimes \bigotimes_{l=k+1}^M \Pi_{M_l} \right) + \tilde{\Pi} \quad (\text{B5})$$

where $\tilde{\Pi}$ is a projection defined by

$$\tilde{\Pi} := \mathbb{1}_{\mathcal{H} \setminus \text{L}} - \sum_{k=1}^M \bigotimes_{j=1}^{k-1} \mathbb{1}_{M_j} \otimes |1_C 0_H\rangle\langle 1_C 0_H|_{M_k} \otimes \bigotimes_{l=k+1}^M \Pi_{M_l} = \bigotimes_{k=1}^M \Pi_{M_k} + \sum_{k=1}^M \bigotimes_{j=1}^{k-1} \mathbb{1}_{M_j} \otimes |0_C 1_H\rangle\langle 0_C 1_H|_{M_k} \otimes \bigotimes_{j'=k+1}^M \Pi_{M_{j'}}. \quad (\text{B6})$$

In order to evaluate the top-level probability, let us briefly inspect the initial state ρ_0 in this situation, which is given by

$$\begin{aligned} \rho_0 &= \bigotimes_{i=1}^M \tau_{C_i} \otimes \tau_{H_i} \otimes \tau_{\text{L}} \\ &= \underbrace{\bigotimes_{i=1}^M \left(\frac{1}{\mathcal{Z}_C} |0\rangle\langle 0|_{C_i} + \frac{\mathcal{Z}_C - 1}{\mathcal{Z}_C} |1\rangle\langle 1|_{C_i} \right) \otimes \left(\frac{1}{\mathcal{Z}_H} |0\rangle\langle 0|_{H_i} + \frac{\mathcal{Z}_H - 1}{\mathcal{Z}_H} |1\rangle\langle 1|_{H_i} \right)}_{\text{Tr}_{\text{L}}(\rho_0)} \otimes \left(\frac{1}{\mathcal{Z}_{\text{L}}} |0\rangle\langle 0|_{\text{L}} + \frac{\mathcal{Z}_{\text{L}} - 1}{\mathcal{Z}_{\text{L}}} |1\rangle\langle 1|_{\text{L}} \right). \end{aligned} \quad (\text{B7})$$

We split the initial state into two parts: one where the ladder is initially excited and one where it is not, *i.e.*

$$\rho_0 = \frac{1}{\mathcal{Z}_{\text{L}}} \text{Tr}_{\text{L}}(\rho_0) \otimes |0\rangle\langle 0|_{\text{L}} + \frac{\mathcal{Z}_{\text{L}} - 1}{\mathcal{Z}_{\text{L}}} \text{Tr}_{\text{L}}(\rho_0) \otimes |1\rangle\langle 1|_{\text{L}}. \quad (\text{B8})$$

The top-level probability can then be seen to split into two separate contributions, corresponding to the two terms in Eq. (B8), that is,

$$P_{\text{top}}(t) = \frac{1}{\mathcal{Z}_{\text{L}}} \text{Tr}(T_0 \text{Tr}_{\text{L}}(\rho_0) T_0^{\dagger}) + \frac{\mathcal{Z}_{\text{L}} - 1}{\mathcal{Z}_{\text{L}}} \text{Tr}(T_1 \text{Tr}_{\text{L}}(\rho_0) T_1^{\dagger}), \quad (\text{B9})$$

We will first consider the part of ρ_0 where the ladder is initially in the ground state. Considering $\bigotimes_{j=1}^{k-1} \mathbb{1}_{M_j} \otimes \sigma_{M_k}^- \otimes \bigotimes_{l=k+1}^M \Pi_{M_l}$ in Eq. (B4) we see that for each k there are $k-1$ machines that are acted upon only by identities, meaning

the partial trace over each of these machines simply contributes a factor 1. There are $M - k$ machines that are acted upon by an operator Π_{M_j} , each leading to a factor

$$\begin{aligned} & \text{Tr} \left[\frac{1}{\mathcal{Z}_H \mathcal{Z}_C} \Pi_{M_j} \left(|0\rangle\langle 0|_{C_j} + (\mathcal{Z}_C - 1) |1\rangle\langle 1|_{C_j} \right) \otimes \left(|0\rangle\langle 0|_{H_j} + (\mathcal{Z}_H - 1) |1\rangle\langle 1|_{H_j} \right) \Pi_{M_j}^\dagger \right] \\ &= \text{Tr} \left[\frac{1}{\mathcal{Z}_H \mathcal{Z}_C} \left(|0_C 0_H\rangle\langle 0_C 0_H|_{M_j} + (\mathcal{Z}_H - 1)(\mathcal{Z}_C - 1) |1_C 1_H\rangle\langle 1_C 1_H|_{M_j} \right) \right] = \frac{1 + (\mathcal{Z}_C - 1)(\mathcal{Z}_H - 1)}{\mathcal{Z}_H \mathcal{Z}_C}, \end{aligned} \quad (\text{B10})$$

in $\text{Tr}(T_0 \text{Tr}_L(\rho_0) T_0^\dagger)$. In addition, there is always exactly one machine which is acted upon by $\sigma_{M_k}^-$, contributing a factor of

$$\begin{aligned} & \text{Tr} \left[\frac{1}{\mathcal{Z}_H \mathcal{Z}_C} \sigma_{M_k}^- \left(|0\rangle\langle 0|_{C_k} + (\mathcal{Z}_C - 1) |1\rangle\langle 1|_{C_k} \right) \otimes \left(|0\rangle\langle 0|_{H_k} + (\mathcal{Z}_H - 1) |1\rangle\langle 1|_{H_k} \right) \sigma_{M_k}^+ \right] \\ &= \text{Tr} \left[\frac{1}{\mathcal{Z}_H \mathcal{Z}_C} (\mathcal{Z}_H - 1) |1_C 0_H\rangle\langle 1_C 0_H| \right] = \frac{\mathcal{Z}_H - 1}{\mathcal{Z}_H \mathcal{Z}_C} \end{aligned} \quad (\text{B11})$$

in $\text{Tr}(T_0 \text{Tr}_L(\rho_0) T_0^\dagger)$. The first term of $P_{\text{top}}(t)$ is thus given by the sum over all $k \in 1, 2, \dots, M$ [see Eq. (B4)], multiplied by the initial population of the ladder ground state, resulting in

$$\frac{1}{\mathcal{Z}_L} \text{Tr}(T_0 \text{Tr}_L(\rho_0) T_0^\dagger) = \frac{1}{\mathcal{Z}_L} \frac{\mathcal{Z}_H - 1}{\mathcal{Z}_H \mathcal{Z}_C} \sum_{k=1}^M \left(\frac{1 + (\mathcal{Z}_C - 1)(\mathcal{Z}_H - 1)}{\mathcal{Z}_H \mathcal{Z}_C} \right)^{k-1} \sin^2(gt). \quad (\text{B12})$$

The second part of $P_{\text{top}}(t)$ can be calculated in the same way. Comparing Eq. (B5) with Eq. (B4), one sees that (aside from the oscillating scalar factors) the first term of T_1 differs from T_0 by replacing $\sigma_{M_k}^-$ in the latter with $|1_C 0_H\rangle\langle 1_C 0_H|_{M_k}$, and the corresponding factor $\frac{\mathcal{Z}_H - 1}{\mathcal{Z}_H \mathcal{Z}_C}$ is thus replaced by

$$\begin{aligned} & \text{Tr} \left[\frac{1}{\mathcal{Z}_H \mathcal{Z}_C} |1_C 0_H\rangle\langle 1_C 0_H|_{M_k} \left(|0\rangle\langle 0|_{C_k} + (\mathcal{Z}_C - 1) |1\rangle\langle 1|_{C_k} \right) \otimes \left(|0\rangle\langle 0|_{H_k} + (\mathcal{Z}_H - 1) |1\rangle\langle 1|_{H_k} \right) |1_C 0_H\rangle\langle 1_C 0_H|_{M_k} \right] \\ &= \text{Tr} \left[\frac{1}{\mathcal{Z}_H \mathcal{Z}_C} (\mathcal{Z}_C - 1) |1_C 0_H\rangle\langle 1_C 0_H|_{M_k} \right] = \frac{\mathcal{Z}_C - 1}{\mathcal{Z}_H \mathcal{Z}_C}. \end{aligned} \quad (\text{B13})$$

The transition operator T_1 additionally contains the static term $\tilde{\Pi}$, which means that $P_{\text{top}}(t)$ contains the additional term

$$\begin{aligned} & \text{Tr} \left[\frac{1}{\mathcal{Z}_H^M \mathcal{Z}_C^M} \tilde{\Pi} \bigotimes_{k=1}^M \left(|0\rangle\langle 0|_{C_k} + (\mathcal{Z}_C - 1) |1\rangle\langle 1|_{C_k} \right) \otimes \left(|0\rangle\langle 0|_{H_k} + (\mathcal{Z}_H - 1) |1\rangle\langle 1|_{H_k} \right) \tilde{\Pi} \right] \\ &= \frac{(1 + (\mathcal{Z}_C - 1)(\mathcal{Z}_H - 1))^M}{\mathcal{Z}_H^M \mathcal{Z}_C^M} + \sum_{k=1}^M \frac{(1 + (\mathcal{Z}_C - 1)(\mathcal{Z}_H - 1))^{k-1} (\mathcal{Z}_H - 1)}{\mathcal{Z}_H^{k-1} \mathcal{Z}_C^{k-1}}. \end{aligned} \quad (\text{B14})$$

Thus, the second term of $P_{\text{top}}(t)$ becomes

$$\begin{aligned} \frac{\mathcal{Z}_L - 1}{\mathcal{Z}_L} \text{Tr}(T_1 \text{Tr}_L(\rho_0) T_1^\dagger) &= \frac{\mathcal{Z}_L - 1}{\mathcal{Z}_L} \frac{\mathcal{Z}_C - 1}{\mathcal{Z}_H \mathcal{Z}_C} \sum_{k=1}^M \left(\frac{1 + (\mathcal{Z}_C - 1)(\mathcal{Z}_H - 1)}{\mathcal{Z}_H \mathcal{Z}_C} \right)^{k-1} \cos^2(gt) \\ &\quad + \frac{\mathcal{Z}_L - 1}{\mathcal{Z}_L} \left(\frac{(1 + (\mathcal{Z}_C - 1)(\mathcal{Z}_H - 1))^M}{\mathcal{Z}_H^M \mathcal{Z}_C^M} + \sum_{k=1}^M \frac{(1 + (\mathcal{Z}_C - 1)(\mathcal{Z}_H - 1))^{k-1} (\mathcal{Z}_H - 1)}{\mathcal{Z}_H^{k-1} \mathcal{Z}_C^{k-1}} \right), \end{aligned} \quad (\text{B15})$$

and the total top-level probability of the horizontal extension is given by

$$\begin{aligned} P_{\text{top}}(t) &= \sum_{k=1}^M \left(\frac{1 + (\mathcal{Z}_C - 1)(\mathcal{Z}_H - 1)}{\mathcal{Z}_H \mathcal{Z}_C} \right)^{k-1} \left(\frac{1}{\mathcal{Z}_L} \frac{\mathcal{Z}_H - 1}{\mathcal{Z}_H \mathcal{Z}_C} \sin^2(gt) + \frac{\mathcal{Z}_L - 1}{\mathcal{Z}_L} \frac{\mathcal{Z}_C - 1}{\mathcal{Z}_H \mathcal{Z}_C} \cos^2(gt) \right) \\ &\quad + \frac{\mathcal{Z}_L - 1}{\mathcal{Z}_L} \left(\frac{(1 + (\mathcal{Z}_C - 1)(\mathcal{Z}_H - 1))^M}{\mathcal{Z}_H^M \mathcal{Z}_C^M} + \sum_{k=1}^M \frac{(1 + (\mathcal{Z}_C - 1)(\mathcal{Z}_H - 1))^{k-1} (\mathcal{Z}_H - 1)}{\mathcal{Z}_H^{k-1} \mathcal{Z}_C^{k-1}} \right) \end{aligned} \quad (\text{B16})$$

Taking the limit $T_C \rightarrow 0$, the top-level probability becomes

$$P_{\text{top}}(t) = \left(1 - \frac{1}{\mathcal{Z}_H^M} \right) \sin^2(gt). \quad (\text{B17})$$

Appendix C: Details on the vertical extension

Here, we present a detailed derivation of the top-level probability P_{top} for the vertical extension of the clockwork. That is, we consider $M(d-1)$ two-qubit machines coupled to the d -level ladder, such that M machines non-trivially couple to each of the $d-1$ ladder transitions. For the purpose of this derivation, we will consider the general case that both the cold and hot bath have finite temperatures, in particular, $T_C \geq 0$ and $T_H < \infty$, but we will assume that $T_C < T_H$. To label the different machines, we denote the Hilbert space of the j^{th} machine coupling to the i^{th} ladder transition (*i.e.*, the transition between the ladder levels $|i-1\rangle_L$ and $|i\rangle_L$) by M_j^i , where $i \in \{0, 1, \dots, d-1\}$ and $j \in \{1, 2, \dots, M\}$. Moreover, we denote the Hilbert space of the collection of all machines within the same ‘column’ (see, e.g., the illustration in Fig. 3), *i.e.* the collection of j^{th} machines for all ladder transitions, as $M_{(j)} := \bigotimes_{i=1}^{d-1} M_j^i$. Following these conventions, we define the fully (horizontally and vertically) extended interaction Hamiltonian as

$$H_{\text{int}} = \sum_{k=1}^M \bigotimes_{i=1}^{k-1} \mathbb{1}_{M_{(i)}} \otimes J_{M_{(k)}L} \otimes \bigotimes_{j=k+1}^M \Pi_{M_{(j)}} = \sum_{k=1}^M H_k. \quad (\text{C1})$$

Here, the operator H_k acts non-trivially on the joint Hilbert space $M_{(k)}$ of the k^{th} ‘column’ and the ladder via the operator

$$J_{M_{(k)}L} := ig \sum_{n=1}^{d-1} \sqrt{n(d-n)} \left(|n_{M_{(k)}}, n_L \rangle \langle n-1_{M_{(k)}}, n-1_L| - H.c. \right). \quad (\text{C2})$$

However, the action of $J_{M_{(k)}L}$ is conditioned on the states of the machines corresponding to the ‘columns’ $M_{(k+1)}$, $M_{(k+2)}$, ..., $M_{(M)}$ through the projectors

$$\Pi_{M_{(k)}} := \mathbb{1}_{M_{(k)}} - \sum_{n=0}^{N-1} |n_{M_{(k)}} \rangle \langle n_{M_{(k)}}| = \mathbb{1}_{M_{(k)}} - \bar{\Pi}_{M_{(k)}}, \quad (\text{C3})$$

where $|0_{M_{(k)}} \rangle := \bigotimes_{j=1}^{d-1} |0_C 1_H \rangle_{M_k^j}$, and the states $|n_{M_{(k)}} \rangle$ for $n = 1, \dots, d-1$ are defined as

$$|n_{M_{(k)}} \rangle := \bigotimes_{j=1}^n |1_C 0_H \rangle_{M_k^j} \bigotimes_{l=n+1}^{d-1} |0_C 1_H \rangle_{M_k^l}. \quad (\text{C4})$$

The state $|n_{M_{(k)}} \rangle$ can be considered to be the n^{th} excited state of the k^{th} vertical group $M_{(k)}$ in the sense that the first n machines M_k^j for $j = 1, \dots, n$ are in the ‘used’ state $|1_C 0_H \rangle_{M_k^j}$, whereas the remaining $d-n+1$ machines M_k^l for $l = n+1, \dots, d-1$ are in the ‘unused’ state $|0_C 1_H \rangle_{M_k^l}$. Similarly, the state $|0_{M_{(k)}} \rangle$ represents the corresponding ‘ground state’.

Further note that $J_{M_{(k)}L}$ acts as an effective generator of rotations on the states $|n_L, n_{M_{(k)}} \rangle$ for $n = 0, \dots, d-1$. To be more precise, $J_{M_{(k)}L}$ can be considered to be a spin- j representation (for $j = \frac{d-1}{2}$) of the generator of rotations around the y -axis on the subspace $\mathcal{W}_{(k)} := \text{span}(\{|n_{M_{(k)}} \rangle\}_{n=0, \dots, d-1}) \subset M_{(k)}$ of the Hilbert space $M_{(k)}$ of the k^{th} vertical group of machines. Let us further denote the orthogonal complement of $\mathcal{W}_{(k)}$ with respect to $M_{(k)}$ by $\mathcal{W}_{(k)}^\perp$, such that $M_{(k)} = \mathcal{W}_{(k)} \oplus \mathcal{W}_{(k)}^\perp$. We then observe that $\ker(J_{M_{(k)}L}) = \mathcal{W}_{(k)}^\perp \otimes \mathcal{H}_L$. Since $\Pi_{M_{(k)}}$ projects onto $\mathcal{W}_{(k)}^\perp$, we see that H_k only has support on the subspace $\text{supp}(H_k) = \bigotimes_{i=1}^{k-1} M_{(i)} \otimes \mathcal{W}_{(k)} \otimes \bigotimes_{j=k+1}^M \mathcal{W}_{(j)}^\perp \otimes \mathcal{H}_L$ of the total Hilbert space of the ladder and all machines. Moreover, these subspaces are orthogonal for different values of k , *i.e.*

$$H_k H_{k'} = 0 \quad \forall k \neq k'. \quad (\text{C5})$$

As a consequence, we have $H_{\text{int}}^q = \left(\sum_{k=1}^M H_k \right)^q = \sum_{k=1}^M H_k^q$, which we can use in the power expansion of $e^{-iH_{\text{int}}t}$, that is,

$$e^{-iH_{\text{int}}t} = \sum_{q=0}^{\infty} \frac{(-it)^q}{q!} H_{\text{int}}^q = \mathbb{1} + \sum_{q=1}^{\infty} \frac{(-it)^q}{q!} H_{\text{int}}^q = \mathbb{1} + \sum_{k=1}^M \sum_{q=1}^{\infty} \frac{(-it)^q}{q!} H_k^q = \mathbb{1} + \sum_{k=1}^M \bigotimes_{i=1}^{k-1} \mathbb{1}_{M_{(i)}} \otimes \left(\sum_{q=1}^{\infty} \frac{(-it)^q}{q!} J_{M_{(k)}L}^q \right) \otimes \bigotimes_{j=k+1}^M \Pi_{M_{(j)}}, \quad (\text{C6})$$

where we have isolated the leading order term ($q = 0$) in the expansion, and used the fact that the $\mathbb{1}_{M_{(i)}}$ and $\Pi_{M_{(j)}}$ are idempotent. Next, we observe that by definition $J_{M_{(k)}L}^q \bar{\Pi}_{M_{(k)}} = J_{M_{(k)}L}^q$ for all $q \geq 1$. We then define the operator

$U_{M(k)L}(t) := e^{-iJ_{M(k)L}t}$ and write

$$U_{M(k)L} \bar{\Pi}_{M(k)} = e^{-iJ_{M(k)L}t} \bar{\Pi}_{M(k)} = \left(\mathbb{1}_{M(k)L} + \sum_{q=1}^{\infty} \frac{(-it)^q}{q!} J_{M(k)L}^q \right) \bar{\Pi}_{M(k)} = \mathbb{1}_L \otimes \bar{\Pi}_{M(k)} + \sum_{q=1}^{\infty} \frac{(-it)^q}{q!} J_{M(k)L}^q. \quad (C7)$$

Inserting Eq. (C7) into Eq. (C6), we obtain

$$\begin{aligned} e^{-iH_{\text{int}}t} &= \mathbb{1} - \mathbb{1}_L \otimes \sum_{k=1}^M \bigotimes_{i=1}^{k-1} \mathbb{1}_{M(i)} \otimes \bar{\Pi}_{M(k)} \otimes \bigotimes_{j=k+1}^M \Pi_{M(j)} + \sum_{k=1}^M \bigotimes_{i=1}^{k-1} \mathbb{1}_{M(i)} \otimes U_{M(k)L} \bar{\Pi}_{M(k)} \otimes \bigotimes_{j=k+1}^M \Pi_{M(j)} \\ &= \mathbb{1} - \mathbb{1}_L \otimes \sum_{k=1}^M \tilde{\Pi}_{[k]} + \sum_{k=1}^M \tilde{U}_{[k]} = \bar{\mathbb{1}} + \sum_{k=1}^M \tilde{U}_{[k]}, \end{aligned} \quad (C8)$$

where we have defined $\bar{\mathbb{1}} := \mathbb{1} - \mathbb{1}_L \otimes \sum_{k=1}^M \tilde{\Pi}_{[k]}$. The projectors $\tilde{\Pi}_{[k]}$ and operators $\tilde{U}_{[k]}$ are defined as

$$\tilde{\Pi}_{[k]} := \bigotimes_{i=1}^{k-1} \mathbb{1}_{M(i)} \otimes \bar{\Pi}_{M(k)} \otimes \bigotimes_{j=k+1}^M \Pi_{M(j)}, \quad (C9)$$

$$\tilde{U}_{[k]} := \bigotimes_{i=1}^{k-1} \mathbb{1}_{M(i)} \otimes U_{M(k)L} \bar{\Pi}_{M(k)} \otimes \bigotimes_{j=k+1}^M \Pi_{M(j)}, \quad (C10)$$

such that $\tilde{U}_{[k]} \tilde{U}_{[k']} = 0$ and $\tilde{\Pi}_{[k]} \tilde{\Pi}_{[k']} = 0 \ \forall k \neq k'$, while $(\mathbb{1}_L \otimes \tilde{\Pi}_{[k']}) \tilde{U}_{[k]} = \tilde{U}_{[k]} (\mathbb{1}_L \otimes \tilde{\Pi}_{[k']}) = \delta_{kk'} \tilde{U}_{[k]}$ and $\bar{\mathbb{1}} \tilde{U}_{[k]} = \tilde{U}_{[k]} \bar{\mathbb{1}} = 0$.

With this, we are now in a position to provide a compact expression of the top-level probability $P_{\text{top}}(t)$, which takes the form

$$\begin{aligned} P_{\text{top}}(t) &= \text{Tr}[\langle d-1 | \langle d-1 | \rho \rangle] = \text{Tr}[\langle d-1 | e^{-iH_{\text{int}}t} \rho_0 e^{iH_{\text{int}}t} | d-1 \rangle_L] \\ &= \text{Tr}[\langle d-1 | (\bar{\mathbb{1}} + \sum_{k=1}^M \tilde{U}_{[k]}) \rho_0 (\bar{\mathbb{1}} + \sum_{k=1}^M \tilde{U}_{[k]}^\dagger) | d-1 \rangle_L] \\ &= \text{Tr}[\langle d-1 | \bar{\mathbb{1}} \rho_0 \bar{\mathbb{1}} | d-1 \rangle_L] + \sum_{k=1}^M \text{Tr}[\langle d-1 | \tilde{U}_{[k]} \rho_0 \tilde{U}_{[k]}^\dagger | d-1 \rangle_L], \end{aligned} \quad (C11)$$

where we have used the assumption that the initial state ρ_0 is diagonal with respect to the joint eigenbasis of the orthogonal projectors $\tilde{\Pi}_{[k]}$, which is the case here because the ladder and all machines qubits are initially thermal with respect to either the cold or hot bath.

For the first term in P_{top} we then have

$$\begin{aligned} \text{Tr}[\langle d-1 | \bar{\mathbb{1}} \rho_0 \bar{\mathbb{1}} | d-1 \rangle_L] &= \text{Tr}[\langle d-1 | \bar{\mathbb{1}} \rho_0 | d-1 \rangle_L] = \text{Tr}[\langle d-1 | (\mathbb{1} - \mathbb{1}_L \otimes \sum_{k=1}^M \tilde{\Pi}_{[k]}) \rho_0 | d-1 \rangle_L] \\ &= \langle d-1 | \tau_L(\beta_C) | d-1 \rangle_L \left(1 - \sum_{k=1}^M \text{Tr}[\tilde{\Pi}_{[k]} \text{Tr}_L(\rho_0)] \right). \end{aligned} \quad (C12)$$

Here, we further have

$$\begin{aligned} \text{Tr}[\tilde{\Pi}_{[k]} \text{Tr}_L(\rho_0)] &= \text{Tr}[\left(\bigotimes_{i=1}^{k-1} \mathbb{1}_{M(i)} \otimes \bar{\Pi}_{M(k)} \otimes \bigotimes_{j=k+1}^M \Pi_{M(j)} \right) \bigotimes_{l=1}^M \tau_{M(l)}] = \text{Tr}[\bar{\Pi}_{M(k)} \tau_{M(k)}] \prod_{j=k+1}^M \text{Tr}[\Pi_{M(j)} \tau_{M(j)}] \\ &= \text{Tr}[\bar{\Pi}_{M(k)} \tau_{M(k)}] \prod_{j=k+1}^M (1 - \text{Tr}[\bar{\Pi}_{M(j)} \tau_{M(j)}]), \end{aligned} \quad (C13)$$

where we can use Eq. (C3) to calculate

$$\begin{aligned} \text{Tr}[\bar{\Pi}_{M(j)} \tau_{M(j)}] &= \sum_{n=0}^{N-1} \langle n_{M(j)} | \tau_{M(j)} | n_{M(j)} \rangle = \sum_{n=0}^{N-1} \prod_{i=1}^n \langle 1_{C0H} | \tau_{M_j^i} | 1_{C0H} \rangle \prod_{l=n+1}^{N-1} \langle 0_{C1H} | \tau_{M_j^l} | 0_{C1H} \rangle \\ &= \sum_{n=0}^{N-1} \langle 1 | \tau_C | 1 \rangle^n \langle 0 | \tau_H | 0 \rangle^n \langle 0 | \tau_C | 0 \rangle^{d-n-1} \langle 1 | \tau_H | 1 \rangle^{d-n-1} = \frac{1}{\mathcal{Z}_C^{d-1} \mathcal{Z}_H^{d-1}} \sum_{n=0}^{d-1} (\mathcal{Z}_C - 1)^n (\mathcal{Z}_H - 1)^{d-n-1} \\ &= \frac{(\mathcal{Z}_H - 1)^d - (\mathcal{Z}_C - 1)^d}{\mathcal{Z}_H^{d-1} \mathcal{Z}_C^{d-1} (\mathcal{Z}_H - \mathcal{Z}_C)}. \end{aligned} \quad (C14)$$

Inserting Eqs. (C14) and (C13) into Eq. (C12) and evaluating the sum over k , we obtain

$$\text{Tr}[\text{L}\langle d-1 | \bar{1} \rho_0 \bar{1} | d-1 \rangle_{\text{L}}] = \text{L}\langle d-1 | \tau_{\text{L}}(\beta_{\text{C}}) | d-1 \rangle_{\text{L}} \left(1 - \frac{(\mathcal{Z}_{\text{H}}-1)^d - (\mathcal{Z}_{\text{C}}-1)^d}{\mathcal{Z}_{\text{H}}^{d-1} \mathcal{Z}_{\text{C}}^{d-1} (\mathcal{Z}_{\text{H}} - \mathcal{Z}_{\text{C}})} \right)^M. \quad (\text{C15})$$

Turning to the second term of P_{top} in Eq. (C11), we express the individual terms in the sum over k as

$$\text{Tr}[\text{L}\langle d-1 | \tilde{U}_{[k]} \rho_0 \tilde{U}_{[k]}^\dagger | d-1 \rangle_{\text{L}}] = \text{Tr}[\text{L}\langle d-1 | \bar{\Pi}_{\text{M}(k)} U_{\text{M}(k)\text{L}} \tau_{\text{L}} \otimes \tau_{\text{M}(k)} U_{\text{M}(k)\text{L}}^\dagger \bar{\Pi}_{\text{M}(k)} | d-1 \rangle_{\text{L}}] \prod_{j=k+1}^M \text{Tr}[\Pi_{\text{M}(j)} \tau_{\text{M}(j)}]. \quad (\text{C16})$$

Then, we note that we can write

$$\begin{aligned} \bar{\Pi}_{\text{M}(k)} U_{\text{M}(k)\text{L}} &= \bar{\Pi}_{\text{M}(k)} \otimes \mathbb{1}_{\text{L}} + \sum_{q=1}^{\infty} \frac{(-it)^q}{q!} J_{\text{M}(k)\text{L}}^q = \sum_{m,n=0}^{N-1} |m_{\text{M}(k)}, n_{\text{L}}\rangle \langle m_{\text{M}(k)}, n_{\text{L}}| + \sum_{q=1}^{\infty} \frac{(-it)^q}{q!} J_{\text{M}(k)\text{L}}^q \\ &= \sum_{\substack{m,n=0 \\ m \neq n}}^{d-1} |m_{\text{M}(k)}, n_{\text{L}}\rangle \langle m_{\text{M}(k)}, n_{\text{L}}| + \sum_{n=0}^{d-1} |n_{\text{M}(k)}, n_{\text{L}}\rangle \langle n_{\text{M}(k)}, n_{\text{L}}| + \sum_{q=1}^{\infty} \frac{(-it)^q}{q!} J_{\text{M}(k)\text{L}}^q, \\ &= \sum_{\substack{m,n=0 \\ m \neq n}}^{N-1} |m_{\text{M}(k)}, n_{\text{L}}\rangle \langle m_{\text{M}(k)}, n_{\text{L}}| + \sum_{n=0}^{d-1} |n_{\text{M}(k)}, n_{\text{L}}\rangle \langle n_{\text{M}(k)}, n_{\text{L}}| e^{-iJ_{\text{M}(k)\text{L}}t}, \end{aligned} \quad (\text{C17})$$

where we have separated the terms corresponding to projectors onto the kernel and support of the operator $J_{\text{M}(k)\text{L}}$ in the second step. With this, we can simplify the first factor appearing on the right-hand side of Eq. (C16) to

$$\begin{aligned} \text{Tr}[\text{L}\langle d-1 | \bar{\Pi}_{\text{M}(k)} U_{\text{M}(k)\text{L}} \tau_{\text{L}} \otimes \tau_{\text{M}(k)} U_{\text{M}(k)\text{L}}^\dagger \bar{\Pi}_{\text{M}(k)} | d-1 \rangle_{\text{L}}] &= \sum_{n=0}^{d-2} \langle n_{\text{M}(k)}, d-1_{\text{L}} | \tau_{\text{L}} \otimes \tau_{\text{M}(k)} | n_{\text{M}(k)}, d-1_{\text{L}} \rangle \\ &\quad + \text{Tr}[\text{L}\langle d-1 | \left(\sum_{n=0}^{d-1} |n_{\text{M}(k)}, n_{\text{L}}\rangle \langle n_{\text{M}(k)}, n_{\text{L}}| e^{-iJ_{\text{M}(k)\text{L}}t} \right) (\tau_{\text{L}} \otimes \tau_{\text{M}(k)}) \left(e^{iJ_{\text{M}(k)\text{L}}t} \sum_{n'=0}^{d-1} |n'_{\text{M}(k)}, n'_{\text{L}}\rangle \langle n'_{\text{M}(k)}, n'_{\text{L}}| \right) | d-1 \rangle_{\text{L}}] \\ &= \text{L}\langle d-1 | \tau_{\text{L}}(\beta_{\text{C}}) | d-1 \rangle_{\text{L}} \sum_{n=0}^{d-2} \langle n_{\text{M}(k)} | \tau_{\text{M}(k)} | n_{\text{M}(k)} \rangle \\ &\quad + \sum_{n=0}^{d-1} \langle n_{\text{M}(k)}, n_{\text{L}} | \tau_{\text{M}(k)} \otimes \tau_{\text{L}} | n_{\text{M}(k)}, n_{\text{L}} \rangle \left| \langle N-1_{\text{M}(k)}, d-1_{\text{L}} | e^{-iJ_{\text{M}(k)\text{L}}t} | n_{\text{M}(k)}, n_{\text{L}} \rangle \right|^2 \\ &= \text{L}\langle d-1 | \tau_{\text{L}}(\beta_{\text{C}}) | d-1 \rangle_{\text{L}} \sum_{n=0}^{d-2} \frac{(\mathcal{Z}_{\text{C}}-1)^n (\mathcal{Z}_{\text{H}}-1)^{d-n-1}}{\mathcal{Z}_{\text{C}}^{d-1} \mathcal{Z}_{\text{H}}^{d-1}} \\ &\quad + \sum_{n=0}^{d-1} \text{L}\langle n | \tau_{\text{L}}(\beta_{\text{C}}) | n \rangle_{\text{L}} \frac{(\mathcal{Z}_{\text{C}}-1)^n (\mathcal{Z}_{\text{H}}-1)^{d-n-1}}{\mathcal{Z}_{\text{C}}^{d-1} \mathcal{Z}_{\text{H}}^{d-1}} \left| \langle d-1_{\text{M}(k)}, d-1_{\text{L}} | e^{-iJ_{\text{M}(k)\text{L}}t} | n_{\text{M}(k)}, n_{\text{L}} \rangle \right|^2. \end{aligned} \quad (\text{C18})$$

Reinserting the first term appearing in the last step of Eq. (C18) back into Eq. (C16), and evaluating the sum over k in Eq. (C11), we obtain another [*i.e.* in addition to that in Eq. (C15)] time-independent contribution to the top-level probability, given by

$$\begin{aligned} \text{L}\langle d-1 | \tau_{\text{L}}(\beta_{\text{C}}) | d-1 \rangle_{\text{L}} \sum_{k=1}^M \left(1 - \frac{(\mathcal{Z}_{\text{H}}-1)^d - (\mathcal{Z}_{\text{C}}-1)^d}{\mathcal{Z}_{\text{H}}^{d-1} \mathcal{Z}_{\text{C}}^{d-1} (\mathcal{Z}_{\text{H}} - \mathcal{Z}_{\text{C}})} \right)^{M-k} \sum_{n=0}^{d-2} \frac{(\mathcal{Z}_{\text{C}}-1)^n (\mathcal{Z}_{\text{H}}-1)^{d-n-1}}{\mathcal{Z}_{\text{C}}^{d-1} \mathcal{Z}_{\text{H}}^{d-1}} \\ = \text{L}\langle d-1 | \tau_{\text{L}}(\beta_{\text{C}}) | d-1 \rangle_{\text{L}} \left(1 - \frac{(\mathcal{Z}_{\text{C}}-1)^{d-1} (\mathcal{Z}_{\text{H}} - \mathcal{Z}_{\text{C}})}{(\mathcal{Z}_{\text{H}}-1)^{d-1} (\mathcal{Z}_{\text{C}}-1)^d} \right) \left[1 - \left(1 - \frac{(\mathcal{Z}_{\text{H}}-1)^d - (\mathcal{Z}_{\text{C}}-1)^d}{\mathcal{Z}_{\text{H}}^{d-1} \mathcal{Z}_{\text{C}}^{d-1} (\mathcal{Z}_{\text{H}} - \mathcal{Z}_{\text{C}})} \right)^M \right]. \end{aligned} \quad (\text{C19})$$

For the second term appearing in the last step of Eq. (C18), we note that, since $J_{\text{M}(k)\text{L}}$ corresponds to the spin- j representation (with $j = \frac{d-1}{2}$) of the generator of rotations around the y -axis on the subspace spanned by the vectors $|n_{\text{M}(k)}, n_{\text{L}}\rangle$ for $n = 0, 1, \dots, d-1$, the matrix elements $\langle d-1_{\text{M}(k)}, d-1_{\text{L}} | e^{-iJ_{\text{M}(k)\text{L}}t} | n_{\text{M}(k)}, n_{\text{L}} \rangle$ coincide with the elements of the Wigner (small) d -matrix $d_{\mu,m}^j(\beta) := \langle j, \mu | e^{-i\beta J_y} | j, m \rangle$ for $\mu = j$, $m = n - j$, and $\beta = 2gt$, see, e.g., [35] or [36]. In particular, Eq. (B7) in [35, p. 485] lets us write

$$\left| \langle d-1_{\text{M}(k)}, d-1_{\text{L}} | e^{-iJ_{\text{M}(k)\text{L}}t} | n_{\text{M}(k)}, n_{\text{L}} \rangle \right|^2 = \binom{d-1}{n} \cos^{2n}(gt) \sin^{2(d-n-1)}(gt). \quad (\text{C20})$$

The prefactors of these sinusoidal contributions are then obtained by combining the second term in the last step of Eq. (C18) with Eq. (C16), and evaluating the sum over k in Eq. (C11), which yields

$$\begin{aligned} & {}_{\text{L}}\langle n | \tau_{\text{L}}(\beta_{\text{C}}) | n \rangle_{\text{L}} \frac{(\mathcal{Z}_{\text{C}}-1)^n (\mathcal{Z}_{\text{H}}-1)^{d-n-1}}{\mathcal{Z}_{\text{C}}^{d-1} \mathcal{Z}_{\text{H}}^{d-1}} \sum_{k=1}^M \left(1 - \frac{(\mathcal{Z}_{\text{H}}-1)^d - (\mathcal{Z}_{\text{C}}-1)^d}{\mathcal{Z}_{\text{H}}^{d-1} \mathcal{Z}_{\text{C}}^{d-1} (\mathcal{Z}_{\text{H}} - \mathcal{Z}_{\text{C}})} \right)^{M-k} \\ &= {}_{\text{L}}\langle n | \tau_{\text{L}}(\beta_{\text{C}}) | n \rangle_{\text{L}} \frac{(\mathcal{Z}_{\text{C}}-1)^n (\mathcal{Z}_{\text{H}}-1)^{d-n-1} (\mathcal{Z}_{\text{H}} - \mathcal{Z}_{\text{C}})}{(\mathcal{Z}_{\text{H}}-1)^d - (\mathcal{Z}_{\text{C}}-1)^d} \left[1 - \left(1 - \frac{(\mathcal{Z}_{\text{H}}-1)^d - (\mathcal{Z}_{\text{C}}-1)^d}{\mathcal{Z}_{\text{H}}^{d-1} \mathcal{Z}_{\text{C}}^{d-1} (\mathcal{Z}_{\text{H}} - \mathcal{Z}_{\text{C}})} \right)^M \right]. \end{aligned} \quad (\text{C21})$$

Finally, we can collect Eqs. (C20) and (C21), and combine them with the time-independent terms in P_{top} to arrive at

$$\begin{aligned} P_{\text{top}}(t) &= \sum_{n=0}^{d-2} {}_{\text{L}}\langle n | \tau_{\text{L}}(\beta_{\text{C}}) | n \rangle_{\text{L}} (\mathcal{Z}_{\text{C}}-1)^n (\mathcal{Z}_{\text{H}}-1)^{d-n-1} f(M, d, \beta_{\text{C}}, \beta_{\text{H}}) \binom{d-1}{n} \cos^{2n}(gt) \sin^{2(d-n-1)}(gt) \\ &+ {}_{\text{L}}\langle d-1 | \tau_{\text{L}}(\beta_{\text{C}}) | d-1 \rangle_{\text{L}} \left[1 - \left(1 - \cos^{2(d-1)}(gt) \right) (\mathcal{Z}_{\text{C}}-1)^{d-1} f(M, d, \beta_{\text{C}}, \beta_{\text{H}}) \right], \end{aligned} \quad (\text{C22})$$

where the coefficient $f(M, d, \beta_{\text{C}}, \beta_{\text{H}})$ is given by

$$f(M, d, \beta_{\text{C}}, \beta_{\text{H}}) = \frac{\mathcal{Z}_{\text{H}} - \mathcal{Z}_{\text{C}}}{(\mathcal{Z}_{\text{H}}-1)^d - (\mathcal{Z}_{\text{C}}-1)^d} \left[1 - \left(1 - \frac{(\mathcal{Z}_{\text{H}}-1)^d - (\mathcal{Z}_{\text{C}}-1)^d}{\mathcal{Z}_{\text{H}}^{d-1} \mathcal{Z}_{\text{C}}^{d-1} (\mathcal{Z}_{\text{H}} - \mathcal{Z}_{\text{C}})} \right)^M \right]. \quad (\text{C23})$$

The expression in Eq. (C22) holds for arbitrary temperatures T_{C} and $T_{\text{H}} > T_{\text{C}}$, and includes the desired term proportional to $\sin^{2(d-1)}(gt)$ in the sum for $n=0$. In particular, this term is the only term in $P_{\text{top}}(t)$ that remains when taking the limit $T_{\text{C}} \rightarrow 0$, in which case $\mathcal{Z}_{\text{C}} \rightarrow 1$, ${}_{\text{L}}\langle d-1 | \tau_{\text{L}}(\beta_{\text{C}}) | d-1 \rangle_{\text{L}} \rightarrow 0$, and ${}_{\text{L}}\langle 0 | \tau_{\text{L}}(\beta_{\text{C}}) | 0 \rangle_{\text{L}} = 1$, and we have

$$\lim_{T_{\text{C}} \rightarrow 0} P_{\text{top}}(t) = \left[1 - \left(1 - \left(\frac{\mathcal{Z}_{\text{H}}-1}{\mathcal{Z}_{\text{H}}} \right)^{d-1} \right)^M \right] \sin^{2(d-1)}(gt), \quad (\text{C24})$$

as stated in Eq. (15) of the main text.

To see that small deviations from the ideal case where $T_{\text{C}} = 0$ still allow for $P_{\text{top}}(t)$ to be close to the corresponding value of the ideal case, *i.e.* to show the stability of our approach to ATPC, we analyse the behaviour of $P_{\text{top}}(t)$ in the limits $M \rightarrow \infty$ and $d \rightarrow \infty$ at finite temperatures. To this end we first inspect Eq. (C23), and note that the term that is potentiated by M is smaller than 1. To see this, we first write

$$\frac{(\mathcal{Z}_{\text{H}}-1)^d - (\mathcal{Z}_{\text{C}}-1)^d}{\mathcal{Z}_{\text{H}}^{d-1} \mathcal{Z}_{\text{C}}^{d-1} (\mathcal{Z}_{\text{H}} - \mathcal{Z}_{\text{C}})} = \frac{(\mathcal{Z}_{\text{H}}-1)^d - (\mathcal{Z}_{\text{C}}-1)^d}{\mathcal{Z}_{\text{H}}^d \mathcal{Z}_{\text{C}}^d} \frac{\mathcal{Z}_{\text{H}} \mathcal{Z}_{\text{C}}}{\mathcal{Z}_{\text{H}} - \mathcal{Z}_{\text{C}}} = \frac{x^d - y^d}{x - y}, \quad (\text{C25})$$

where we have defined $x := \frac{\mathcal{Z}_{\text{H}}-1}{\mathcal{Z}_{\text{H}} \mathcal{Z}_{\text{C}}}$ and $y := \frac{\mathcal{Z}_{\text{C}}-1}{\mathcal{Z}_{\text{H}} \mathcal{Z}_{\text{C}}}$ with the property $0 \leq y < x \leq \frac{1}{2}$. The expression on the right-hand side of Eq. (C25) is smaller or equal than 1 if $x - x^d \geq y - y^d$, which is the case if $x - x^d$ is monotonically increasing on the interval $[0, \frac{1}{2}]$. Inspecting the derivative, we have $\frac{\partial}{\partial x}(x - x^d) = 1 - dx^{d-1} \geq 0$ since $dx^{d-1} \leq d/2^{d-1} \leq 1$ for $d \geq 2$. Consequently, we have $\frac{x^d - y^d}{x - y} \leq 1$ and $\lim_{M \rightarrow \infty} \left(1 - \frac{x^d - y^d}{x - y} \right)^M = 0$. Therefore, we see that

$$\lim_{M \rightarrow \infty} f(M, d, \beta_{\text{C}}, \beta_{\text{H}}) = \frac{\mathcal{Z}_{\text{H}} - \mathcal{Z}_{\text{C}}}{(\mathcal{Z}_{\text{H}}-1)^d - (\mathcal{Z}_{\text{C}}-1)^d}. \quad (\text{C26})$$

Since we know that $P_{\text{top}}(t)$ must lie in $[0, 1]$, showing that the first term of Eq. (C22) ($n=0$) remains close to 1 when M and N go to infinity is sufficient to show that our approach is stable with respect to deviations from $T_{\text{H}} \rightarrow \infty$ and $T_{\text{C}} \rightarrow 0$, *i.e.*

$$\begin{aligned} & \lim_{M \rightarrow \infty} \lim_{d \rightarrow \infty} {}_{\text{L}}\langle 0 | \tau_{\text{L}}(\beta_{\text{C}}) | 0 \rangle_{\text{L}} (\mathcal{Z}_{\text{H}}-1)^{d-1} f(M, d, \beta_{\text{C}}, \beta_{\text{H}}) \binom{d-1}{0} \sin^{2(d-1)}(gt) \\ &= \lim_{d \rightarrow \infty} \frac{1}{\mathcal{Z}_{\text{L}}} (\mathcal{Z}_{\text{H}}-1)^{d-1} \frac{\mathcal{Z}_{\text{H}} - \mathcal{Z}_{\text{C}}}{(\mathcal{Z}_{\text{H}}-1)^d - (\mathcal{Z}_{\text{C}}-1)^d} \sin^{2(d-1)}(gt) = \lim_{d \rightarrow \infty} \frac{1}{\mathcal{Z}_{\text{L}}} \frac{\mathcal{Z}_{\text{H}} - \mathcal{Z}_{\text{C}}}{\mathcal{Z}_{\text{H}} - 1} \frac{1}{1 - \left(\frac{\mathcal{Z}_{\text{C}}-1}{\mathcal{Z}_{\text{H}}-1} \right)^d} \sin^{2(d-1)}(gt) \\ &= \begin{cases} \frac{\mathcal{Z}_{\text{H}} - \mathcal{Z}_{\text{C}}}{\mathcal{Z}_{\text{L}}(\mathcal{Z}_{\text{H}}-1)}, & \text{if } t = \frac{\pi}{2g} \\ 0, & \text{otherwise} \end{cases}. \end{aligned} \quad (\text{C27})$$

The value of the expression in Eq. (C27) for $t = \frac{\pi}{2g}$ can further be written as

$$\begin{aligned} \frac{Z_H - Z_C}{Z_L(Z_H - 1)} &= \frac{1 + e^{-\beta_H E_H} - (1 + e^{-\beta_C E_C})}{e^{-\beta_H E_H} \sum_{n=0}^{d-1} e^{-n\beta_C(E_H - E_C)}} = \frac{1 - e^{\beta_H E_H - \beta_C E_C}}{1 + \sum_{n=1}^{\infty} e^{-n\beta_C(E_H - E_C)}} \\ &= 1 - e^{-\beta_C(E_H - E_C)} - e^{-(\beta_C E_C - \beta_H E_H)} + e^{-E_H(\beta_C - \beta_H)}. \end{aligned} \quad (C28)$$

Recalling that $E_H > E_C$, we can see that the quantity in Eq. (C28) remains close to 1 for finite temperatures when $\beta_C \gg \beta_H$ such that $\beta_H E_H < \beta_C E_C$, and when $k_B T_C \ll (E_H - E_C)$, which are both in keeping with the assumptions made in the main text.

Appendix D: Tick probability density

In this appendix we show how our derivation of the tick probability results in an exponential decay. The derivation should not be understood as a new result, but rather as a reminder and clarification. Treating the decay of the top level as a random event we can approximate its probability of occurring in the time interval Δt by

$$\Delta P = \Gamma \Delta t \quad (D1)$$

where Γ is given in terms of probability per unit time. In our case here, we have that Γ corresponds to the top-level population times the constant c , *i.e.* $\Gamma(t) = P_{\text{top}}(t)c$. Let us denote the cumulative probability that no decay occurred until time t as $P(0, t)$. We can then approximate the probability that no event occurred until $t + \Delta t$ as the probability of no event happening until t times the probability that no event happens in the interval Δt , *i.e.*

$$P(0, t + \Delta t) = P(0, t)(1 - \Gamma(t)\Delta t) \quad (D2)$$

which leads to

$$\frac{P(0, t + \Delta t) - P(0, t)}{\Delta t} = -\Gamma(t)P(0, t). \quad (D3)$$

If we further let $\Delta t \rightarrow dt$, we get that

$$\frac{dP(0, t)}{dt} = -\Gamma(t)P(0, t) = -cP_{\text{top}}(t)P(0, t) \quad (D4)$$

and consequently

$$P(0, t) = e^{-c \int_0^t P_{\text{top}}(t') dt'}. \quad (D5)$$

Given this expression for the cumulative probability that no event occurred until time t we can proceed to calculate the probability density of a decay event occurring between time t and $t + dt$. To do so, we differentiate the cumulative probability of having had a decay at time t with respect to t , *i.e.* $\frac{d[1 - P(0, t)]}{dt}$, which results in

$$P_{\text{tick}}(t) = cP_{\text{top}}(t)e^{-c \int P_{\text{top}}(t') dt'}. \quad (D6)$$

Appendix E: Numerical calculation of accuracy and resolution

In order to execute the numerical calculations of the resolution and the accuracy efficiently we need to simplify the necessary integrals [defined in Eqs. (18) and (17)]. In this appendix we present details on our approach to this problem. For simplicity we showcase the calculations for $T_C \rightarrow 0$ and $T_H \rightarrow \infty$.

Assessing resolution and accuracy for a given set of parameters d, M, c and g , breaks down to calculating the first and second moment of the tick distribution, *i.e.*

$$\bar{t} = \int_0^\infty t P_{\text{tick}}(t) dt \quad \bar{t}^2 = \int_0^\infty t^2 P_{\text{tick}}(t) dt \quad (E1)$$

where

$$P_{\text{tick}}(t) = c \left\{ 1 - \left[1 - \left(\frac{Z_H - 1}{Z_H} \right)^{d-1} \right]^M \right\} \sin^{2(d-1)}(gt) \exp \left[-c \left\{ 1 - \left[1 - \left(\frac{Z_H - 1}{Z_H} \right)^{d-1} \right]^M \right\} \int_0^t dt' \sin^{2(d-1)}(gt') \right]. \quad (E2)$$

In order to simplify the cumbersome expressions we will use the following notation. We will denote the k^{th} moment, as

$$I_k = c \left\{ 1 - \left[1 - \left(\frac{\mathcal{Z}_H - 1}{\mathcal{Z}_H} \right)^{d-1} \right]^M \right\} \int_0^\infty dt t^k \sin^{2(d-1)}(gt) \exp \left[-c \left\{ 1 - \left[1 - \left(\frac{\mathcal{Z}_H - 1}{\mathcal{Z}_H} \right)^{d-1} \right]^M \right\} \int_0^t dt' \sin^{2(d-1)}(gt') \right]. \quad (\text{E3})$$

The *effective coupling* is defined as

$$C_M := c \left\{ 1 - \left[1 - \left(\frac{\mathcal{Z}_H - 1}{\mathcal{Z}_H} \right)^{d-1} \right]^M \right\}. \quad (\text{E4})$$

Furthermore, let

$$f(t) = C_M \int_0^\infty dt' \sin^{2(d-1)}(gt'). \quad (\text{E5})$$

This leads to a much simpler form for the different moments,

$$I_k = C_M \int_0^\infty dt t^k \sin^{2d}(gt) e^{-f(t)}. \quad (\text{E6})$$

We can solve the integral in the term $f(t)$ with a solution introduced by Wiener [37],

$$\int_0^x dx' \sin^{2(d-1)}(x') = \frac{1}{4^{d-1}} \left[\binom{2(d-1)}{d-1} x + \sum_{p=1}^{d-1} \frac{(-1)^p}{p} \binom{2(d-1)}{d-1-p} \sin(2px) \right], \quad (\text{E7})$$

where $\binom{n}{m} = \frac{n!}{m!(n-m)!}$ is the binomial coefficient. Employing the solution we get

$$f(t) = \frac{C_M}{4^{d-1}} \left\{ \binom{2(d-1)}{d-1} t + \frac{1}{g} \sum_{p=1}^{d-1} \binom{2(d-1)}{d-1-p} \sin(2pgt) \right\} \quad (\text{E8})$$

and thus,

$$I_k = C_M \int_0^\infty dt t^k \sin^{2(d-1)}(gt) \exp \left[-\frac{C_M}{4^{d-1}} \binom{2(d-1)}{d-1} t \right] \times \exp \left[-\frac{C_M}{4^{d-1}g} \sum_{p=1}^d \frac{(-1)^p}{p} \binom{2(d-1)}{d-1-p} \sin(2pgt) \right]. \quad (\text{E9})$$

By introducing a ‘cycle’ counting variable $q = \lfloor \frac{gt}{\pi} \rfloor \in \mathbb{Z}$ and its residue $\Theta_q = qt - q\pi$, $\Theta_q \in [0, \pi)$, *i.e.* substituting with $t = \frac{q\pi + \Theta_q}{g}$, we arrive at

$$I_k = \frac{C_M}{g} \sum_{q=0}^\infty \int_0^\pi d\Theta_q \left(\frac{q\pi + \Theta_q}{g} \right)^k \sin^{2(d-1)}(\Theta_q) \times \exp \left[-\frac{C_M}{4^{d-1}g} \binom{2(d-1)}{d-1} (q\pi + \Theta_q) \right] \quad (\text{E10})$$

$$\times \exp \left[-\frac{C_M}{4^{d-1}g} \sum_{p=1}^{d-1} \frac{(-1)^p}{p} \binom{2(d-1)}{d-1-p} \sin(2p\Theta_q) \right], \quad (\text{E11})$$

where we have used that $\sin^{2(d-1)}(x \pm n\pi) = \sin^{2(d-1)}(x)$ for $n \in \mathbb{Z}$ as well as $\sin[2(x \pm n\pi)] = \sin(2x)$.

We are only interested in explicitly calculating the cases $k = 1$ and $k = 2$. Note that, k only appears in the term $\left(\frac{q\pi + \Theta_q}{g} \right)^k$ and that $\left(\frac{q\pi + \Theta_q}{g} \right)^2 = \frac{1}{g^2} [(q\pi)^2 + 2q\pi\Theta_q + \Theta_q^2]$, which allows us to define the function

$$E(\Theta_q) = \exp \left[-\frac{C_M}{4^{d-1}g} \binom{2(d-1)}{d-1} \Theta_q \right] \times \exp \left[-\frac{C_M}{4^{d-1}g} \sum_{p=1}^{d-1} \frac{(-1)^p}{p} \binom{2(d-1)}{d-1-p} \sin(2p\Theta_q) \right], \quad (\text{E12})$$

such that

$$I_k = \frac{C_M}{g} \sum_{q=0}^\infty \exp \left[-\frac{C_M}{4^{d-1}g} \binom{2(d-1)}{d-1} q\pi \right] \times \int_0^\pi d\Theta_q \left(\frac{q\pi + \Theta_q}{g} \right)^k \sin^{2(d-1)}(\Theta_q) E(\Theta_q). \quad (\text{E13})$$

As a last step we observe that $E(\Theta_q)$ does not depend on q directly, but only through Θ_q , so the only direct dependence on q in the integral comes from the term $\left(\frac{q\pi + \Theta_q}{g}\right)^k$. This leads us to define

$$\tilde{I}_j := \int_0^\pi d\Theta_q \Theta_q^j \sin^{2(d-1)}(\Theta_q) E(\Theta_q), \quad (\text{E14})$$

such that we can write the desired first and second moments of the tick distribution as

$$I_1 = \frac{C_M}{g^2} \sum_{q=0}^{\infty} \exp\left[-\frac{C_M}{4^{d-1}g} \binom{2(d-1)}{d-1} q\pi\right] \{q\pi \tilde{I}_0 + \tilde{I}_1\} \quad (\text{E15})$$

and

$$I_2 = \frac{C_M}{g^3} \sum_{q=0}^{\infty} \exp\left[-\frac{C_M}{4^{d-1}g} \binom{2(d-1)}{d-1} q\pi\right] \{(q\pi)^2 \tilde{I}_0 + 2q\pi \tilde{I}_1 + \tilde{I}_2\}, \quad (\text{E16})$$

respectively. In this way only $\tilde{I}_j$ needs to be calculated numerically for $j = \{0, 1, 2\}$, which decreases the effective computational costs enormously.

Appendix F: How the ‘sharpness’ of $P_{\text{top}}(t)$ influences accuracy and resolution

The aim of this appendix is to give further insight about the behaviour of clocks with changing ladder dimension d as well as changing coupling constant g , in particular with respect to Figs. 4, 5 and 6 in Sec. IV.

First, let us discuss the relationship between accuracy and ‘sharpness’ of $P_{\text{top}}(t)$. The intuition is that clockworks that are capable of producing a very ‘sharp’ temporal probability distribution should have the potential to give rise to highly accurate clocks, given a suitable irreversible process for the ‘tick’ production. Since the maximal amplitude of $P_{\text{top}}(t)$ is given by $1 - \left[1 - \left(\frac{Z_H - 1}{Z_H}\right)^{d-1}\right]^M$ (for $T_C = 0$), increasing M leads to an amplitude of $P_{\text{top}}(t)$ that approaches 1 very quickly. Assuming that M is chosen large enough so that the maximal amplitude is within a desired distance to the value 1, the only parameter left that influences the ‘sharpness’ of the probability distribution is the ladder dimension d . We can therefore use d as a proxy for ‘sharpness’. We then proceed by numerically calculating the accuracy in this situation for given values of c and g . The results are shown in Fig. 7 (a) and indicate that the accuracy grows linearly with d . In this regime, the ‘sharpness’ therefore determines the accuracy up to a constant factor. However, one should note that this linear relationship only holds in a regime where the decay process happens fast enough, *i.e.* assuming a sufficiently large value of c (or small enough value of g). If $P_{\text{top}}(t)$ is too ‘sharp’ compared to the time scale of the decay process increasing the ladder dimension leads to a reduction of the accuracy [as seen in Figs. 7(a) and 4]. This implies that for a given combination of c and g there are certain choices of d that lead to sub-optimal clocks. Considering the resolution as a function of d [Fig. 7(b)] in the limit $M \rightarrow \infty$ we do not observe an optimal configuration. The resolution simply decreases with increasing d indicating a trade-off relation between accuracy and resolution in the regime where the accuracy increases linearly with d . Thus plotting accuracy over resolution reveals the trade-off relation depicted in Fig. 5. However, considering finite M the resolution reaches a point at which it starts dropping to zero quickly. The reason for this can again be found in the amplitude of $P_{\text{top}}(t)$, which goes to zero for large enough d and fixed M . Thus, not only the accuracy (see Fig. 4), but also the resolution is bounded from above by the corresponding resolution obtained for $M \rightarrow \infty$. There c and g determine this upper bound.

Furthermore, considering only the cases where $M \rightarrow \infty$, Fig. 8 shows that increasing g at fixed c shifts the point of maximal accuracy to the right, *i.e.* towards higher resolutions. Thus the value of g determines the lowest resolution at which (optimal) clocks can operate *i.e.* the *maximal cycle-time*. However this increase in resolution comes at the cost of accuracy as increasing g reduces the maximally reachable accuracy.




Figure 7. (a) The accuracy is shown as a function of d for different values of g (and fixed $c = 10 \text{ s}^{-1}$), where $[g] = E_C$. The maximally achievable accuracy increases with decreasing g . (b) The resolution is shown as a function d (at fixed $g = 1 E_C$ and $c = 1000 \text{ s}^{-1}$). The lines with solid dots show cases of finite M . The line with orange circles illustrates the behaviour for the case $M \rightarrow \infty$, which provides an upper bound to the cases with finite M . For finite M , increasing d reduces the top-level population, eventually becoming so small that the decay event is considerably more likely to skip the first peak of $P_{\text{top}}(t)$. This leads to an additional reduction in resolution initiating a drop of the resolution eventually approaching 0 (with d).




Figure 8. The trade-off between clock accuracy N and resolution R for clockworks of various coupling constants g , where $[g] = E_C$ and R is increased by decreasing d . Here we only consider cases of $M \rightarrow \infty$ and $c = 25 \text{ s}^{-1}$. Increasing g shifts the peak towards the right, *i.e.* to higher resolutions while decreasing the maximum accuracy.

III.2 Fundamental accuracy-resolution trade-off for timekeeping devices

III.2.1 Contribution

In this this work, I played a significant role in the conceptualisation of the theoretical framework and the formalisation of the methods. I was involved in preliminary numerical investigations, the formulation of the main proofs of the main results, writing and revising the manuscript, and overall organisation of the project as well as supervision.

Fundamental accuracy-resolution trade-off for timekeeping devices

Florian Meier,^{1,2,*} Emanuel Schwarzhans,¹ Paul Erker,^{1,3} and Marcus Huber^{1,3}

¹*Vienna Center for Quantum Science and Technology,
Atominstitut, Technische Universität Wien, 1020 Vienna, Austria*

²*Institute for Theoretical Physics, ETH Zurich, 8093 Zürich, Switzerland*

³*Institute for Quantum Optics and Quantum Information,
Austrian Academy of Sciences, 1090 Vienna, Austria*

(Dated: January 31, 2023)

From a thermodynamic point of view, all clocks are driven by irreversible processes. Additionally, one can use oscillatory systems to temporally modulate the thermodynamic flux towards equilibrium. Focusing on the most elementary thermalization events, this modulation can be thought of as a temporal probability concentration for these events. There are two fundamental factors limiting the performance of clocks: On the one level, the inevitable drifts of the oscillatory system, which are addressed by finding stable atomic or nuclear transitions that lead to astounding precision of today's clocks. On the other level, there is the intrinsically stochastic nature of the irreversible events upon which the clock's operation is based. This becomes relevant when seeking to maximize a clock's resolution at high accuracy, which is ultimately limited by the number of such stochastic events per reference time unit. We address this essential trade-off between clock accuracy and resolution, proving a universal bound for all clocks whose elementary thermalization events are memoryless.

Clocks dominate our daily lives unlike any other technology – from ordinary things like catching the train in the morning to locating ourselves using the GPS – they are involved everywhere and have been omnipresent since the dawn of human civilization. From the sundials of Mayan times to the UTC standard of today, over the centuries, physics provided the theoretical and experimental foundation to develop accurate and stable clocks. In the 50s, this development has taken a sharp turn with the invention of atomic clocks [1]. Year by year, these clocks have been improving their relative accuracy. By now, if you run a state-of-the-art optical clock, it is so accurate that it accumulates less than a hundred milliseconds of error over the lifespan of the sun [2–4]. These advancements beg for the question whether there are any physical principles constraining a clock's performance. Driven by this open problem, quantum clocks are emerging into an independent field of research, unifying approaches from quantum information theory [5–8] and quantum thermodynamics [9–13]. Recent experiments showcase today's technology is capable of exploring these ultimate limits of timekeeping [14–16].

This letter explores fundamental limitations of clocks coming from thermodynamics, because clocks, like all other physical systems, are subject to thermodynamical laws [11]. Even worse, they very much rely on the increase of entropy originating in the second law of thermodynamics. This implies that (1) clocks are witnesses of the macroscopic breaking of time-reversal symmetry because they tick forwards in time. (2) Thus, they must be driven by irreversible processes that drain out-of-equilibrium resources to output temporal information. These processes, though, are thermodynamic,

therefore inherently stochastic and never perfectly predictable. (3) We conclude, even an idealized clock could never be perfect simply due to the fact that the clock itself is fundamentally driven by stochastic processes.

We look at clocks that use this flow by counting elementary thermalization events to define *ticks*. Instructive examples of such a process could range from grains of sand that pass through an hourglass to the photons that are reflected from a pendulum to ascertain its position. Even modern atomic clocks ultimately rely on a macroscopic number of photons to read out the laser frequency (which is stabilized by feedback from atomic transitions). In all these cases, the underlying thermalization processes are stochastic and have some intrinsic rate Γ . The time passing between two events, relative to assumed smooth parameter time, is probabilistic and described by a probability density function. We refer to its average as μ and its uncertainty as σ . We define the *resolution* ν of the clock to be the inverse average time between two ticks $\nu = 1/\mu$, and the *accuracy* N as the average number of times the clock ticks until it is off by one tick. For a sequence of independent and identically distributed ticks, the accuracy equals the signal-to-noise ratio $N = \mu^2/\sigma^2$. A clock usually modulates the temporal distribution of when the thermalization events occur as to improve its accuracy. For most practical clocks, the stability of this temporal modulation (for instance the laser frequency in atomic clocks) is the limiting factor to the accuracy. We show that even if this temporal modulation is perfect, the stochastic process underlying the tick generation limits the accuracy at given resolution.

Theorem (Accuracy-resolution trade-off). *Clocks with elementary ticking events generated by a memoryless stochastic process at rate Γ obey the trade-off relation*

$$N \leq \frac{\Gamma^2}{\nu^2}. \quad (1)$$

* florian.meier@tuwien.ac.at

This trade-off complements other established results in the field considering restrictions imposed on the clock through entropy production [9, 17] or Hilbertspace dimension [18, 19].

Temporal probability concentration. The most primitive clock consists of two out-of-equilibrium thermal reservoirs in contact with each other. Counting the individual stochastic thermalization events as ticks can serve as a way to measure time, we call this the *thermal reference clock*. Sufficiently large reservoirs are memoryless, therefore, these events become exponentially distributed [20–22]. The coupling of the two baths defines a characteristic thermalization rate Γ . For an exponential distribution the standard deviation equals the average, which leads to $\mu = \sigma = \Gamma^{-1}$. Consequently, such a clock has unit accuracy $N = 1$ and only by averaging over many of these ticks are we able to achieve higher accuracy, but at the expense of resolution. Averaging over M independent and identically distributed (i.i.d.) such events increases the variance and mean of the tick time M -fold, leading to a resolution $\nu \propto 1/M$ and accuracy $N \propto M$. This gives an inverse proportional accuracy-resolution scaling

$$N = \frac{\Gamma}{\nu}, \quad (2)$$

quadratically smaller than the upper bound in eq. (1). Upon closer inspection, we find a thermodynamic cost associated to this increase in accuracy: instead of having a single irreversible event producing a tick, now M irreversible events are required.

A natural question to ask is whether it is possible to increase the accuracy beyond what is achievable through averaging in eq. (2), while still using the same underlying stochastic tick generating process. The answer lies in the observation that all ticking clocks known to us use a combination of two processes to tell time:

- (a) irreversible processes that generate ticks, and
- (b) a filter process, *temporal probability concentration*, which modulates the probability of the irreversible ticking events to occur.

By means of temporal probability concentration (TPC), a (sensible) clock centers the probability distribution of the stochastic events with a periodic process such that the ticks occur closely around well-defined instants in time. A thermodynamically optimal strategy would be to sample the periodic process with a single equilibration event per cycle [9, 10, 23], as we show in Figure 1. One microscopic example of such a clock could be a leaky cavity where an autonomous process induces Rabi oscillations. These oscillations temporally concentrate the population of the cavity that can leak out and generate a tick. Many macroscopically sized clocks, be it pendulums or atomic clocks, operate in an oversampling regime, where multiple irreversible events occur per TPC cycle (see Figure 2). In this regime, the thermodynamic cost of a clockwork often becomes obscure, as both the dissipation due to the macroscopic number of irreversible




FIG. 1. The two plots qualitatively show the ticking behavior of two clocks with respect to parameter time t (horizontal axis). The ticks of such a clock are generated by individual thermalization events at rate Γ (vertical stripes with cross on top). Figure a) sketches a generic example, where these thermalization events are Poisson distributed. In b), temporal probability concentration is shown, another process through which the probability of an thermalization event can be suppressed at times (shown by the hat-shaped curve). As a result, the probability density for a tick can be concentrated around a desired average value, here $5 \Gamma^{-1}$, with the tick time uncertainty of order Γ^{-1} bounded by the width of the TPC window.




FIG. 2. We sketch the oversampling regime of an exemplary clock: a pendulum in a weakly lit environment. The two sources of entropy production for this clock are first of all the friction within the clockwork itself and secondly, the matter light interaction necessary to track the position of the pendulum. The plot shows the elementary ticking events of this clock as a function of time, i.e., the photons reflected off the pendulum when it is close to its maximum deflection. In the oversampling regime, the average time between two such ticks is much shorter than that of the period of the TPC (continuous line), which in the case of this pendulum is 2 s. Due to technical limitations, one does not count photons, but rather the TPC cycles through the averaged light intensity.

events and the TPC have to be accounted for. Atomic clocks, for example, do not count photons as a way to tell time, rather they use the oversampled coherent oscillations of a maser tuned some stable reference atomic transition to estimate the TPC frequency. The time-scale Γ of the fundamental ticking events appearing in eq. (1)

is therefore not the limiting factor to the accuracy of atomic clocks, the stability of the coherent oscillation of the electromagnetic field is, i.e., the TPC stability. In atomic clocks, accuracy is examined using Allan Variance which captures the stability of the TPC oscillation over many different time-scales [24–26]. Quantum projection noise, thermal noise but also natural drifts in the experimental setup are what affect the stability of atomic clocks [27, 28].

So far, we have established that every clock is subject to irreversible processes and that through TPC, they can increase their accuracy. In the following, we introduce a mathematical model to describe clocks on a quantum scale and where the two contributions (a) irreversible ticks and (b) temporal probability concentrations have an explicit representation in the equations of motion. Eventually, this framework allows us to formulate the fundamental trade-off between the accuracy and resolution of clocks, dictated by thermodynamics.

Model. Quantum clocks [8–10, 18, 23] only weakly coupled to a memoryless environment (sufficiently large thermal baths are one such environment) can be described by a Lindblad master-equation [29]. If we are again talking about ticking clocks, their state can be Fourier-decomposed by introducing a free counting field χ [30],

$$\rho_C(t, \chi) = \sum_n \rho_C^{(n)}(t) e^{in\chi}. \quad (3)$$

Each non-normalized density matrix $\rho^{(n)}(t)$ can be thought of as the system's state conditioned on n ticks having already occurred. Clocks that reset to the same state after they tick regardless of how many times they have ticked before give rise to an i.i.d. sequence of ticks. This class of clocks is particularly interesting, as their ticking statistics are entirely determined by the first tick, and it is sufficient to consider the evolution of the non-normalized state $\rho^{(0)}$. In a setting where the interactions with the environment are memoryless, the (first) ticks are generated by linear jump operators J_j . Aside from this, the clock is subject to a general open quantum system's evolution with Lindblad operator $\mathcal{L}_0$, which leads us to write the equations of motion of the clock's state (where no tick has yet occurred) as

$$\dot{\rho}_C^{(0)}(t) = \mathcal{L}_0[\rho_C^{(0)}(t)] - \frac{1}{2} \sum_j \left\{ J_j^\dagger J_j, \rho_C^{(0)}(t) \right\}. \quad (4)$$

Ticking statistics. The right-most term of eq. (4) is responsible for the drain of population from the no-tick state $\rho_C^{(0)}$ to the state of the clock that has ticked once $\rho_C^{(1)}$. If we take the random variable T to describe the time of the first tick, the cumulative probability $P[t \leq T]$ that no tick has occurred up to time t equals the trace $P[t \leq T] = \text{tr} \rho_C^{(0)}$. There is an equivalent way of writing the cumulative tick probability by looking at the clock's state conditioned on not having decayed $\rho_C^{\text{no tick}} =$

$\rho_C^{(0)} / \text{tr} \rho_C^{(0)}$. Using this state, the tick probability can alternatively be written as the exponential

$$P[t \leq T] = \exp \left(- \int_0^t dt' \text{tr} (V \rho_C^{\text{no tick}}(t')) \right), \quad (5)$$

where $V = \sum_j J_j^\dagger J_j$ is the positive operator generating the clock's ticks [31]. This formulation faithfully represents the two processes that run a clock. The irreversible tick generating process (a) is represented by the underlying exponential and the operator V , whereas the TPC (b) lies in the time-evolution of the clock's no-tick state, i.e., $\rho_C^{\text{no tick}}(t')$. In this sense, one may attempt to separate (a) and (b) into two independent processes, as in [10]; however the back-action of the tick channel always affects the clock evolution which is an inherent feature of clocks that operate on a quantum scale.

The master-equation description of quantum ticking clocks allows us to formalize the statement made in the accuracy-resolution trade-off theorem. The tick channel is governed by the positive operator V , whose spectral decomposition reveals the time-scales involved in the decay. We define the rate Γ as the fastest one of them,

$$\Gamma := \|V\|_{\max} = \max_{\rho \in \mathcal{S}(\mathcal{H}_C)} \text{tr} V \rho, \quad (6)$$

where the maximum is taken over all possible clock states $\rho \in \mathcal{S}(\mathcal{H}_C)$. This is consistent with the special case of exponential decay (as used in [9, 17]) with rate Γ , since there, the tick generator is given by a single jump operator $J = \sqrt{\Gamma} |0\rangle\langle 1|$. If a clock is described by eq. (4) and produces its ticks by means of the generators J_j , then the accuracy is limited by the resolution, regardless of how well a possible clockwork in the background works. The resulting bound is eq. (1) from the accuracy-resolution trade-off theorem which we recall here for completeness,

$$N \leq \frac{\Gamma^2}{\nu^2},$$

and prove in the following.

Proof. The key observation is that no clock can on average tick faster than the decay process that mediates the ticks allows for. If the elementary ticks are generated by an ensemble of jump operators J_j , then the fastest such rate is given by Γ (see eq. (6)). A single such channel produces ticks that are exponentially distributed and as a consequence, we can reduce the generic form from eq. (5) to the special case of exponential decay, where the exponent reduces to $-\Gamma \int dt' p(t')$, and $p(t)$ is a measure of the clock's state population that can decay. Regardless of how the TPC modulates $p(t)$, the variance σ^2 of the tick can never be smaller than Γ^2 , the variance of exponential decay. This is a manifestation of the fact that no clock can tick faster than its underlying decay process. Now that the variance σ^2 of the tick time distribution is bounded from above by Γ^2 , the main theorem follows. We refer the reader to Appendix A for a detailed account of the proof. $\square$




FIG. 3. In this figure, we plot the accuracy N on the vertical axis versus the resolution ν on the horizontal axis. The red curve highlights the upper bound for the accuracy $N = \Gamma^2/\nu^2$ given by the accuracy-resolution trade-off theorem. Below, we have the blue curve, indicating the accuracy and resolution that can be achieved by averaging the thermal reference clock with fundamental tick rate Γ . The green shaded area between the two curves contains all the accuracy-resolution tuples that cannot be achieved by classically averaging a rate Γ stochastic process, but which are still allowed by the upper bound. How close one can get to the red curve with quantum clocks and whether classical clocks are constrained to be below the blue line are open questions.

In this letter, we have analyzed quantum timekeeping devices through a thermodynamic lens, where their tick generation can be decomposed into two processes, (a) a stochastic process which irreversibly produces the ticks and (b) temporal probability concentration (TPC) through a clockwork which controls when these elementary ticks occur. Then, we asserted that there is a fundamental trade-off between the clock's accuracy and resolution (see trade-off theorem but also Figure 3), stating that the number of times a clock can tick until it goes wrong by one tick is universally bounded by the inverse of the clock's resolution squared.

Atomic and optical clocks. These results have their validity in the regime where the number of fundamental thermalization events is directly used to define ticks. Macroscopic clocks such as atomic and optical ones work in a different regime, where the TPC is oversampled by irreversible events and ticks are defined not by event numbers directly, but rather those are used to estimate the oscillatory TPC process frequency. Atomic clocks working with a 10 dBm maser at the 21 cm line of hydrogen, for example, emit of the order of 10^{12} photons per period of the electromagnetic field oscillation. Given that the individual photons are exponentially distributed, a

(hypothetical) clock which counts the maser's photons and averages over one second would achieve an accuracy of $N \sim 10^{21}$ compared to $N \sim 10^{13}$ that current-day commercial masers reach in terms of relative precision per second [32]. Accuracy restrictions coming from the elementary thermalization events are not the limiting factor when it comes to atomic clocks because of the macroscopic number of such events over which is averaged. The main sources of error are drifts due to changes in temperature, environmental electromagnetic fields and other experimental imperfections. Moreover, it is also important to keep in mind that atomic clocks do not generate independent and identically distributed ticks, as their real advantage comes from long-time stability achieved through feedback-control using an ultra-stable atomic transition which is not captured by the i.i.d. formalism herein presented.

Achievability. Good clocks excel by resolving time well while at the same time being highly accurate, i.e., maximizing both ν and N ; ideally, they do so at optimal thermodynamic costs far away from the oversampled regime and only with a single thermalization event per tick. As a foundational question, we may ask, does there exist an appropriate clock Lindblad operator $\mathcal{L}$, through which it is possible to saturate the bound $N = \Gamma^2/\nu^2$? And the answer is: No, at least not for finite systems. As it turns out, saturating this bound amounts to a time-dependent Lindblad-operator, which instantaneously rotates the clock state from a subspace of states where it can't tick onto a state from which it decays at rate Γ (see details of the proof for the accuracy-resolution trade-off theorem in Appendix A). The TPC in this case would be ideal as it concentrates the decay event to the most narrow time-window possible, the one given by the underlying stochastic process, but quantum speed limits prohibit such an instantaneous state rotation for systems finite in energy and dimension [33–35]. A weaker question we can ask is whether at least the scaling $N = O(\Gamma^2/\nu^2)$ can be reached, and which resources are required for this. Without imposing any restriction on the clock's Lindbladian aside from time-independence, there are quantum clocks which asymptotically reach the squared scaling of N for large dimensions of the clock Hilbert space [23]. To do so, they require highly coherent states, whose generation using only thermal resources is technically infinitely expensive from an entropic perspective [36] and the approximation with finite resources is an open problem. Of interest in the field of thermodynamics are the autonomous quantum clocks which only require thermal resources to run [9, 10, 12, 31], in particular no external control but also no coherence in the initial state. It is ongoing research, whether it is possible for such clocks to approach the optimal accuracy-resolution scaling.

Questions about fundamental precision limits are generally of great interest. The accuracy of clocks falls into this category and closely related problems have been examined in the field of (quantum) stochastic thermodynamics under the name of thermodynamic uncertainty

relations (TUR) [37–39] and kinematic uncertainty relations (KUR) [40–42]. Future work exploring connections between timekeeping and the TUR / KUR has to reveal whether a quantum thermodynamic advantage close to the optimal accuracy-resolution bound $N = \Gamma^2/\nu^2$ is in principle possible. Both superconducting circuits and optomechanical systems are promising platforms to test the achievability of the optimal accuracy-resolution relation while accounting for the resources and ensuring that they do not introduce a hidden clock through a backdoor.

We wish to acknowledge fruitful discussions and feedback from Ralph Silva, Nuriya Nurgalieva, Renato Renner, Maximilian Edward Lock and Jake Xuereb. F.M. and M.H. acknowledge the financial support from the Eu-

ropean Commission under the Horizon Europe project ASPECTS (Grant No. 101080167) and the ERC consolidator grant COCOQUEST (Grant No. 101043705). F.M. acknowledges the SEMP scholarship from Movetia for his research stay at Atominstitut, Technische Universität Wien. E.S. acknowledges the support from the Austrian Science Fund (FWF) through the START project Y879-N27, the ESQ Discovery grant “Emergence of physical laws: From mathematical foundations to applications in many body physics” and the ERC consolidator grant COCOQUEST (Grant No. 101043705). P.E. and M.H. further acknowledge funds from the FQXi (FQXi-IAF19-03-S2) within the project “Fueling quantum field machines with information”.

-
- [1] L. Essen and J. V. L. Parry, An Atomic Standard of Frequency and Time Interval: A Cesium Resonator, *Nature* **176**, 280 (1955).
 - [2] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. Schmidt, Optical atomic clocks, *Reviews of Modern Physics* **87**, 637 (2015).
 - [3] F. Riehle, Towards a redefinition of the second based on optical atomic clocks, *Comptes Rendus Physique* **16**, 506 (2015), the measurement of time / La mesure du temps.
 - [4] T. Bothwell, C. J. Kennedy, A. Aeppli, D. Kedar, J. M. Robinson, E. Oelker, A. Staron, and J. Ye, Resolving the gravitational redshift across a millimetre-scale atomic sample, *Nature* **602**, 420 (2022).
 - [5] S. Ranković, Y.-C. Liang, and R. Renner, *Quantum clocks and their synchronisation - the Alternate Ticks Game* (2015).
 - [6] R. Renner and S. Stupar, eds., *Time in Physics* (Springer International Publishing, 2017).
 - [7] S. Stupar, C. Klumpp, R. Renner, and N. Gisin, *Performance of stochastic clocks in the Alternate Ticks Game* (2018).
 - [8] M. P. Woods, Autonomous Ticking Clocks from Axiomatic Principles, *Quantum* **5**, 381 (2021).
 - [9] 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?, *PRX* **7**, 031022 (2017).
 - [10] E. Schwarzhans, M. P. E. Lock, P. Erker, N. Friis, and M. Huber, Autonomous Temporal Probability Concentration: Clockworks and the Second Law of Thermodynamics, *PRX* **11**, 011046 (2021).
 - [11] G. J. Milburn, The thermodynamics of clocks, *Contemporary Physics* **61**, 69 (2020), <https://doi.org/10.1080/00107514.2020.1837471>.
 - [12] S. K. Manikandan, *Autonomous quantum clocks using athermal resources* (2022).
 - [13] D. Cilluffo, Statistical time-domain characterization of non-periodic optical clocks, *Quantum* **6**, 764 (2022).
 - [14] A. N. Pearson, Y. Guryanova, P. Erker, E. A. Laird, G. A. D. Briggs, M. Huber, and N. Ares, Measuring the Thermodynamic Cost of Timekeeping, *Phys. Rev. X* **11**, 021029 (2021).
 - [15] M. Berholts, R. Knut, R. Stefanuik, H. Wikmark, S. Saha, and J. Söderström, Quantum watch and its intrinsic proof of accuracy, *Phys. Rev. Res.* **4**, 043041 (2022).
 - [16] X. He, P. Pakkiam, A. Gangat, G. Milburn, and A. Fedorov, *Measurement driven quantum clock implemented with a superconducting qubit* (2022).
 - [17] A. C. Barato and U. Seifert, Cost and precision of brownian clocks, *Phys. Rev. X* **6**, 041053 (2016).
 - [18] M. P. Woods, R. Silva, G. Pütz, S. Stupar, and R. Renner, Quantum Clocks are More Accurate Than Classical Ones, *PRXQUANTUM* **3**, 010319 (2022).
 - [19] Y. Yang and R. Renner, *Ultimate limit on time signal generation* (2020).
 - [20] R. S. Ingarden, A. Kossakowski, and M. Ohya, *Information Dynamics and Open Systems* (Springer Netherlands, 1997).
 - [21] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, 2007).
 - [22] A. Klenke, *Probability Theory* (Springer International Publishing, 2020).
 - [23] M. P. Woods, R. Silva, and J. Oppenheim, Autonomous Quantum Machines and Finite-Sized Clocks, *Annales Henri Poincaré* **20**, 125 (2019).
 - [24] D. W. Allan, Statistics of atomic frequency standards, *Proceedings of the IEEE* **54**, 221 (1966).
 - [25] D. W. Allan, N. Ashby, and C. C. Hodge, *The Science of Timekeeping Application Note 1289*, Hewlett Packard (1997), packard Application Note 1289.
 - [26] W. Riley and D. Howe, *Handbook of Frequency Stability Analysis* (2008).
 - [27] N. Shiga, M. Mizuno, K. Kido, P. Phoonthong, and K. Okada, Accelerating the averaging rate of atomic ensemble clock stability using atomic phase lock, *New Journal of Physics* **16**, 073029 (2014).
 - [28] M. Schulte, C. Lisdar, P. O. Schmidt, U. Sterr, and K. Hammerer, Prospects and challenges for squeezing-enhanced optical atomic clocks, *Nature Communications* **11**, 5955 (2020).
 - [29] G. Lindblad, On the generators of quantum dynamical semigroups, *Communications in Mathematical Physics* **48**, 119 (1976).
 - [30] G. Schaller, *Open Quantum Systems Far from Equilibrium* (Springer International Publishing, 2014).
 - [31] F. Meier, *Performance limits of decay clocks: Fundamental accuracy and resolution limits of quantum clocks with*

- [exponential ticking mechanism](#) (2022).
- [32] *MHM-2020 Active Hydrogen Maser*, Microchip Technology Inc. (2019), rev. 1.
 - [33] L. Mandelstam, I. Tamm, and I. E. Tamm, The Uncertainty Relation Between Energy and Time in Non-relativistic Quantum Mechanics, in *Selected Papers*, edited by B. M. Bolotovskii, V. Y. Frenkel, and R. Peierls (Springer Berlin Heidelberg, Berlin, Heidelberg, 1991) pp. 115–123.
 - [34] N. Margolus and L. B. Levitin, The maximum speed of dynamical evolution, *Physica D: Nonlinear Phenomena* **120**, 188 (1998), proceedings of the Fourth Workshop on Physics and Consumption.
 - [35] S. Deffner and S. Campbell, Quantum speed limits: from heisenberg’s uncertainty principle to optimal quantum control, *Journal of Physics A: Mathematical and Theoretical* **50**, 453001 (2017).
 - [36] P. Taranto, F. Bakhshinezhad, A. Bluhm, R. Silva, N. Friis, M. P. E. Lock, G. Vitagliano, F. C. Binder, T. Debarba, E. Schwarzhans, F. Clivaz, and M. Huber, *Landauer vs. Nernst: What is the True Cost of Cooling a Quantum System?* (2021), [arXiv:2106.05151 \[quant-ph\]](#).
 - [37] A. C. Barato and U. Seifert, Thermodynamic Uncertainty Relation for Biomolecular Processes, *Phys. Rev. Lett.* **114**, 158101 (2015).
 - [38] J. M. Horowitz and T. R. Gingrich, Thermodynamic uncertainty relations constrain non-equilibrium fluctuations, *Nature Physics* **16**, 15 (2020).
 - [39] T. V. Vu and K. Saito, Thermodynamics of Precision in Markovian Open Quantum Dynamics, *Physical Review Letters* **128**, 140602 (2022).
 - [40] I. D. Terlizzi and M. Baiesi, Kinetic uncertainty relation, *Journal of Physics A: Mathematical and Theoretical* **52**, 02LT03 (2018).
 - [41] V. T. Vo, T. V. Vu, and Y. Hasegawa, Unified thermodynamic–kinetic uncertainty relation, *Journal of Physics A: Mathematical and Theoretical* **55**, 405004 (2022).
 - [42] K. Prech, P. Johansson, E. Nyholm, G. T. Landi, C. Verdozzi, P. Samuelsson, and P. P. Potts, *Entanglement and thermo-kinetic uncertainty relations in coherent mesoscopic transport* (2022).
 - [43] J. Goold, M. Huber, A. Riera, L. d. Rio, and P. Skrzypczyk, The role of quantum information in thermodynamics—a topical review, *Journal of Physics A: Mathematical and Theoretical* **49**, 143001 (2016).

APPENDICES

Appendix A: Proof of the accuracy-resolution trade-off

Before we get started with the proof of the accuracy-resolution trade-off theorem (which is a generalization of the one shown in [31]), some preliminary notation has to be established. When it comes to the cumulative non-tick probability $P[t < T]$ as in eq. (5), the trace $\text{tr } V\rho_C(t)$ can be rewritten as

$$\text{tr}(V\rho_C(t)) = \Gamma p(t), \quad (\text{A1})$$

where $p : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ is a smooth function that takes values between 0 and 1. This comes from the fact that Γ is the maximum of the expression over all states ρ and the function can only take positive values because V is a positive operator. This expression puts us into the position, where we can write

$$P[t \leq T] = \exp\left(-\Gamma \int_0^t dt' p(t')\right), \quad (\text{A2})$$

which is formally the exact same expression that one would obtain in the case where the clock's elementary ticks are exponential decay. We will from now on refer to $p(t)$ as the *top-level population*. The remaining appendices are an adapted version of the proof of the accuracy-resolution theorem originally derived in [31].

1. The Heaviside- Θ population

Definition 1 (Heaviside populations). *The family of Heaviside populations comprises all the functions*

$$p^\Theta : \mathbb{R}_{\geq 0} \rightarrow [0, 1] \quad (\text{A3})$$

$$t \mapsto \Theta(t - t_0), \quad (\text{A4})$$

where $t_0 \geq 0$.

Properties of the Heaviside population. Such a top-level population is unphysical because of its discontinuity at time $t = t_0$. This is not an obstacle, though, as we only use the Heaviside population to prove an upper bound for the accuracy of a decay clock. No claim is made whether we can physically obtain the Heaviside population. The non-tick probability $P_\Theta[t \leq T]$ and tick probability density $p_{\text{tick}}^\Theta(t)$ associated to this population are given by (see eq. (A2))

$$P_\Theta[t \leq T] = \begin{cases} 1 & t \leq t_0 \\ e^{-\Gamma(t-t_0)} & t \geq t_0, \end{cases} \quad (\text{A5})$$

and

$$p_{\text{tick}}^\Theta(t) = \begin{cases} 0 & t \leq t_0 \\ \Gamma e^{-\Gamma(t-t_0)} & t \geq t_0. \end{cases} \quad (\text{A6})$$

See Figure 4 for a visualization of these functions. Using the analytic solution of the integral over an exponential, $\int_0^\infty dx e^{-ax} = 1/a$, we can calculate the accuracy and resolution for clocks working the Heaviside population.

Lemma 2. *The Heaviside population $p^\Theta(t) = \Theta(t - t_0)$ has an accuracy and resolution given by*

$$N = (1 + \Gamma t_0)^2, \text{ and } \nu = \frac{1}{t_0 + \frac{1}{\Gamma}}, \quad (\text{A7})$$

leading to an accuracy-resolution relation

$$N = \frac{\Gamma^2}{\nu^2}. \quad (\text{A8})$$

Before we prove Lemma 2, let us put the Heaviside population into perspective. Given the accuracy-resolution trade-off theorem the equality $N = \Gamma^2/\nu^2$ for clocks with Heaviside population tells us that this top-level population profile is ideal in the accuracy-resolution sense. This is no coincidence: a Heaviside population with offset at t_0 requires a clock in the background which perfectly knows t_0 . The population then describes a ladder conditioned on not having decayed, whose top-level is populated precisely at $t = t_0$, and stays there. This would require instantaneous external driving at $t = t_0$, i.e., a perfect background clock. The exponential decay channel coupled to the top-level then smears out the Heaviside profile of $p^\Theta(t)$ into an exponentially decaying tick probability density $p_{\text{tick}}^\Theta(t)$ of width $1/\Gamma$, giving rise to the results of Lemma 2.

Proof. The average tick time is given by an integral over $tp_{\text{tick}}^\Theta(t)$ and splits into two parts

$$\mu = \int_0^\infty dt t p_{\text{tick}}^\Theta(t) = 0 + \int_{t_0}^\infty dt t e^{-\Gamma(t-t_0)} \quad (\text{A9})$$

$$= t_0 + \frac{1}{\Gamma}. \quad (\text{A10})$$




FIG. 4. This graph visualizes the ticking statistics for the Heaviside population. The solid line in the above figure visualizes the tick probability density, whereas the dashed line stands for the cumulative non-tick probability. The time axis scales in inverse units of the decay rate Γ and the population in the example is chosen such that on average the decay occurs at $\mu = 4.5 \Gamma$.

Thus, the resolution is given by $\nu = (t_0 + \frac{1}{\Gamma})^{-1}$. For the accuracy, which is defined as $N = (\mu/\sigma)^2$, we only need to calculate the variance σ^2 . The variance, however, is invariant under translations of the tick probability density and, thus, one can calculate the variance without loss of generality for $t_0 = 0$,

$$\sigma^2 = \int_0^\infty dt t^2 e^{-\Gamma t} - \mu^2 \quad (\text{A11})$$

$$= \frac{2}{\Gamma^2} - \frac{1}{\Gamma^2} = \frac{1}{\Gamma^2}. \quad (\text{A12})$$

Expressing the accuracy in terms of t_0 , we find $N = (1 + \Gamma t_0)^2$, that is, the accuracy increases with higher offset t_0 . There is a tradeoff, though: The greater t_0 , the lower the resolution. Eliminating the t_0 -dependency in the equations, we can establish an accuracy-resolution relation for the family of Heaviside populations given by $N = \Gamma^2/\nu^2$. $\square$

Analytic preliminaries. The tick probability density $p_{\text{tick}}(t)$ reveals how likely it is for a tick to happen during a given time interval. By the fundamental theorem of calculus, the cumulative non-tick probability $P[t \leq T]$ (whose negative derivative is $p_{\text{tick}}(t)$) contains the same information, and for proving the accuracy-resolution trade-off, the latter function turns out to be useful. Let us collect some general identities for the non-tick probability $P[t \leq T]$ which for some function $p(t)$ (be reminded, that the superscript *no tick* is suppressed) is given by

$$P[t \leq T] = e^{-\Gamma \int_0^t d\tau p(\tau)}. \quad (\text{A13})$$

There is a general constraint (Lemma 3) on how fast the non-tick probability decays. The fact that $p(t) \in [0, 1]$ ensures that $P[t \leq T]$ is monotonically decreasing (because the integral $\int_0^t d\tau p(\tau)$ is monotonically increasing) but does this not faster than exponentially.

Lemma 3. *For all $t > 0$ and $s > 0$, the cumulative non-tick probability satisfies the following inequalities:*

$$P[t \leq T]e^{-\Gamma s} \leq P[t + s \leq T] \leq P[t \leq T]. \quad (\text{A14})$$

Proof. This is a consequence of eq. (A13) and the fact that $0 \leq p(t) \leq 1$. $\square$

The cumulative non-tick probability can be used to calculate the moments of the probability density $p_{\text{tick}}(t)$ (see Definition 4 and Lemma 5) via the relation $p_{\text{tick}}(t) = -\dot{P}[t \leq T]$ and partial integration. It is important to note, that this discussion only makes sense for clocks that tick with certainty, i.e., $\lim_{t \rightarrow \infty} P[t \leq T] = 0$. This need not be true for all clocks, however, whenever we talk about a *tick probability density*, we implicitly assume that we have a properly normalized probability density in the probability theoretic sense [22].

Definition 4 (Moments). *For a given tick probability density $p_{\text{tick}}(t)$, define its k -th moment as*

$$t_k := \int_0^\infty dt t^k p_{\text{tick}}(t). \quad (\text{A15})$$

Lemma 5 (Tick probability moments). *The k -th moment of $p_{\text{tick}}(t)$ is related to $P[t \leq T]$ by the integral*

$$t_k = k \int_0^\infty dt t^{k-1} P[t \leq T]. \quad (\text{A16})$$

Proof. This is partial integration and for the boundary conditions, we use the assumption that $\lim_{t \rightarrow \infty} P[t \leq T] = 0$. $\square$

In particular, we can apply this result to $\mu = t_1$ and $\sigma = t_2 - \mu^2$, two expressions that are essential in Section A 2, where we prove the upper bound. The average tick time $\mu = t_1$ can be written as the area below the graph of $P[t \leq T]$,

$$\mu = t_1 = \int_0^\infty dt t p_{\text{tick}}(t) = \int_0^\infty dt P[t \leq T]. \quad (\text{A17})$$

The second moment t_2 , on the other hand, is related to the center of mass of the graph for $P[t \leq T]$, up to normalization,

$$t_2 = \int_0^\infty dt t^2 p_{\text{tick}}(t) = 2 \int_0^\infty dt t P[t \leq T]. \quad (\text{A18})$$

2. Proof construction

The introduction of the Heaviside population in Definition 1 together with the result in Lemma 2, that these populations achieved the (claimed) optimal accuracy resolution relation, leads us to the following approach in proving the inequality $N \leq \Gamma^2/\nu^2$: we try to show that all clocks are worse than the one with a Heaviside population, which has essentially a perfect background clock. In that sense, what we are showing is that no clock is better than the one that is already perfect. The only premise we have for the proof is that the clocks we consider eventually tick (i.e., $p_{\text{tick}}(t)$ is a valid probability density) and that their evolution is continuous,¹ we call this *well-behaved*. That being said, our strategy to prove the upper bound consists of the following three steps:

- (i) We show that for any well-behaved top-level population $p(t)$, there exists a Heaviside- Θ type top-level population with the same average tick time.

¹ This ensures we are doing a fair comparison, otherwise, we'd have to discuss how to compare clocks that possibly never tick to ones that always tick. We reserve that discussion for future work.




FIG. 5. The sketch shows the time-dependency of a generic top-level population $p(t)$ together with the associated cumulative non-tick probability $P[t \leq T]$ compared to the non-tick probability $e^{-\Gamma t}$ of the top-level population which is constantly 1. A top-level population smaller than unity leads to a slower decay of the non-tick probability, which is responsible for a lower bound on the average tick-time given by $1/\Gamma$.

- (ii) Then, we argue that the variance of the tick probability density coming from the Heaviside- Θ top-level population lower-bounds the variance coming from the generic population $p(t)$.
- (iii) We conclude that any well-behaved top-level population $p(t)$ must have an accuracy upper bounded by that of the Heaviside population and by using Lemma 2 on the properties of the Heaviside population, we have $N \leq \Gamma^2/\nu^2$. Once we are there, we have proven the accuracy-resolution trade-off theorem.

Step (i). Begin with a generic, but well behaved top-level population $p(t)$. Let μ be its average tick time and define $t_0 := \mu - 1/\Gamma$. The top-level population

$$p^\Theta(t) = \Theta(t - t_0) \quad (\text{A19})$$

has the same first moment $t_1 = \mu$ as the generic top-level population.² For the results from Lemma 2 to carry over, we need $t_0 > 0$. This is generally true as Lemma 6 guarantees. In fact, it tells us (see Figure 5) that the average tick time of any well-behaved top-level probability can not be smaller than $1/\Gamma$ and that the best resolution is achieved by the top-level population which is constantly one.

Lemma 6 (Resolution upper bound). *For any well-behaved top-level population $p: \mathbb{R}_{\geq 0} \rightarrow [0, 1]$, the induced resolution ν cannot be greater than Γ .*

Proof. Equivalently to the statement in the Lemma, we can prove that $\mu \geq 1/\Gamma$. For this matter, use eq. (A17), to estimate the average tick time

$$\mu = \int_0^\infty dt P[t \leq T] \geq \int_0^\infty dt e^{-\Gamma t} = 1/\Gamma. \quad (\text{A20})$$

² The average tick time is given by $\mu^\Theta = t_0 + 1/\Gamma$, with t_0 coming from the time translation of the Heaviside- Θ function and $1/\Gamma$ coming from the exponential decay.

This concludes the proof. $\square$

Step (ii). If $p(t) = p^\Theta(t)$, there is nothing to be shown, because $p^\Theta(t)$ achieves the minimal tick time variance $\sigma^2 = 1/\Gamma^2$ (see eq. (A12)). Hence, we assume from now on $p(t) \neq p^\Theta(t)$ to avoid this pathological case. Proposition 8 in Appendix A3 further discusses, that there exists a unique $t_* > t_0 = \mu - 1/\Gamma$, such that $P[t_* \leq T] = e^{-\Gamma(t_* - t_0)}$ (see Figure 6). For all times $t < t_*$, the non-tick probability for the generic top-level population is smaller than that of the Heaviside population, i.e., $P[t \leq T] \leq P_\Theta[t \leq T]$. In the generic case, a tick before t_* is therefore more likely than for the Heaviside case. On the other hand, for all $t > t_*$, we have $P[t \leq T] \geq P_\Theta[t \leq T]$. That is, in the generic case, it is also more likely that after t_* the tick did not happen. Formally, the variance of the tick signal is bigger for a generic population than for the Heaviside one (Proposition 7).

Proposition 7. *For any well-behaved top-level population $p: \mathbb{R}_{\geq 0} \rightarrow [0, 1]$, the variance of the tick probability density $p_{\text{tick}}(t)$ is lower-bounded by that of $p_{\text{tick}}^\Theta(t)$, the tick probability density coming from the Heaviside top-level population $p^\Theta(t)$ defined in eq. (A19).*

Proof. By construction both $p_{\text{tick}}(t)$ and $p_{\text{tick}}^\Theta(t)$ have the same average tick time μ . Because the variance of the tick time is given by the difference $\sigma^2 = t_2 - \mu^2$, the following two statements are equivalent,

$$\sigma^2 \geq \sigma_\Theta^2 \Leftrightarrow t_2 \geq t_2^\Theta. \quad (\text{A21})$$

The second moments on the right-hand side can be calculated as the center of mass of the respective non-tick probabilities $P[t \leq T]$ and $P_\Theta[t \leq T]$ (up to constant prefactors, see also eq. (A18)). Therefore, the question we have to answer is, whether the graph of $P[t \leq T]$ has a center of mass at larger values of t than that of $P_\Theta[t \leq T]$? The two areas A_1 and A_2 in Figure 6 are equal and they correspond to the difference in the area below the graphs of $P[t \leq T]$ and $P_\Theta[t \leq T]$. The graph of $P[t \leq T]$ has less ‘mass’ (i.e., area) on the interval $0 \leq t \leq t_*$ than $P_\Theta[t \leq T]$. For $t_* \leq t \leq \infty$, however, $P[t \leq T]$ has more ‘mass’, resulting in a center of mass further to the right for $P[t \leq T]$ than for $P_\Theta[t \leq T]$. Formally,

$$\frac{t_2}{2} = \int_0^\infty dt t P[t \leq T] \quad (\text{A22})$$

$$= \int_0^\infty dt t P_\Theta[t \leq T] + \int_0^\infty dt t \underbrace{(P[t \leq T] - P_\Theta[t \leq T])}_{:= \Delta P(t)} \quad (\text{A23})$$

$$\geq \int_0^\infty dt t P_\Theta[t \leq T] = \frac{t_2^\Theta}{2}, \quad (\text{A24})$$

where the inequality comes about because the integral




FIG. 6. The dashed lines show the top-level populations for the generic (blue) and the Heaviside case (green). The induced cumulative non-tick probabilities $P[t \leq T]$ and $P_\Theta[t \leq T]$ are drawn in the same color but as solid lines. The time t_* marks the unique crossing point of the two non-tick probabilities. On its right and left, the two areas A_1 and A_2 are defined by the difference between the two non-tick probabilities.

over $\Delta P(t)$ is greater equal than zero,

$$\begin{aligned} \int_0^\infty dt t \Delta P(t) &= \int_0^{t_*} dt t \Delta P(t) + \int_{t_*}^\infty dt t \Delta P(t) \quad (\text{A25}) \\ &\geq \int_0^{t_*} dt t_* \Delta P(t) + \int_{t_*}^\infty dt t_* \Delta P(t) = 0. \end{aligned} \quad (\text{A26})$$

All in all, this shows that $t_2 \geq t_2^\Theta$, which, by the initial remark, proves the proposition. $\square$

Step (iii). Combining all the results from the previous two steps, we proof the Theorem.

Proof. The accuracy N is defined as $N = (\mu/\sigma)^2$, and the resolution ν as $\nu = 1/\mu$. We can therefore express

$$N = \left(\frac{\mu}{\sigma}\right)^2 = \frac{1}{\nu^2 \sigma^2}. \quad (\text{A27})$$

From Proposition 7 we know that $\sigma^2 \geq \sigma_\Theta^2$, but $\sigma_\Theta^2 = 1/\Gamma^2$, as we derived in Lemma 2 (see eq. (A12) in the proof). In conclusion,

$$N = \frac{1}{\nu^2 \sigma^2} \leq \frac{1}{\nu^2 \sigma_\Theta^2} = \frac{\Gamma^2}{\nu^2}, \quad (\text{A28})$$

which proves the claim, and thereby the optimality of the family of Heaviside- Θ top-level populations. $\square$

3. Details on the optimality of the Heaviside population

In Step (ii) of Section A 2, the claim is made that there exists a unique $t_* \geq t_0 = \mu - 1/\Gamma$, such that the two non-tick probabilities coincide $P[t_* \leq T] = P_\Theta[t_* \leq T]$. To be more precise, we formulate




FIG. 7. This sketch visualizes the first step of the the proof of Proposition 8. The solid green line indicates the non-tick probability for the Heaviside population. We see that for the red dashed line which does not fall below 1 before t_0 , the total integral is greater than that of $P_\Theta[t \leq T]$, contradicting the construction, which ensures that both functions have the same integral. The blue, dashed line shows a plausible function for $P[t \leq T]$.




FIG. 8. Here, we visualize step two of the proof of Proposition 8. Again, the solid, green line indicates the non-tick probability for the Heaviside population. To satisfy the construction that $P[t \leq T]$ and $P_\Theta[t \leq T]$ have the same integral on $\mathbb{R}_{\geq 0}$, the area below the graph of $P[t \leq T]$ for $t > t_0$ must be bigger than that of $P_\Theta[t \leq T]$. Thus, at some point after t_0 , $P[t \leq T]$ must be bigger than $P_\Theta[t \leq T]$ (see the blue area denoted by A_{good}), and not always smaller (see red area denoted by A_{bad}).

Proposition 8. The set $\mathcal{S} = \{t \geq 0 : P[t \leq T] = P_\Theta[t \leq T]\}$ is of the form

$$\mathcal{S} = [0, a] \sqcup \{t_*\}, \quad (\text{A29})$$

where $t_* > t_0$, and $[0, a]$ is an interval with upper bound $a < t_0$.

Proof. We divide the proof of this proposition into three steps (a) - (c), with visual aids in Figures 7 and 8.

Step (a). We claim that there exists a $0 < t < t_0$, such that $P[t \leq T] < 1$. Suppose that this was not the case, then $P[t \leq T] = 1$ for all $t \leq t_0$ (see Figure 7). By requirement, there must exist a t such that $p(t) \neq p^\Theta(t)$, and because by our (contradictory) assumption t cannot be smaller than t_0 , it must be bigger than t_0 . Since, $P[t \leq T] \geq e^{-\Gamma(t-t_0)}$, for all $t > t_0$, there must exist a $t > t_0$, such that $P[t \leq T] > e^{-\Gamma(t-t_0)}$. But this contradicts the

assumption that $\int_0^\infty dt P[t \leq T] = t_0 + 1/\Gamma$. Therefore, it must be that there is a $t < t_0$ with $P[t \leq T] < 1$.

Step (b). Our next claim is, that *there exists an* $0 < a < t_0$ *such that on* $[0, a]$, *we have* $P[t \leq T] = 1$ *and for all other* $a < t \leq t_0$, $P[t \leq T] < 1$. By our previous claim (1), there exists $0 < t < t_0$ with $P[t \leq T] < 1$. Let a be the infimum of all such t .³ Then, $[0, a] \subset \mathcal{S}$ (continuity of P). Moreover, by the properties of an infimum, for all $\varepsilon > 0$, there exists a $\delta > 0$, such that $P[a + \delta \leq T] < 1$. But $P[t \leq T]$ is monotonically decreasing, hence, this is true for all $\delta > 0$, such that $a + \delta \leq t_0$, proving the claim.

Step (c). In this last step, we show *the existence of a unique* $t_* > t_0$ *such that* $t_* \in \mathcal{S}$. For a visualization of this step, see Figure 8. The integrals of $P[t \leq T]$

and $P_\Theta[t \leq T]$ on $\mathbb{R}_{\geq 0}$ coincide. However, we have just argued that on a non-empty interval between 0 and t_0 , $P[t \leq T]$ is strictly smaller than $P_\Theta[t \leq T]$. This leads to the inequality $\int_0^{t_0} dt P[t \leq T] < \int_0^{t_0} dt P_\Theta[t \leq T]$. In order to ensure that the integrals over all of $\mathbb{R}_{\geq 0}$ are equal, there must exist a $t' > t_0$ such that $P[t' \leq T] > P_\Theta[t' \leq T]$. For all $t > t'$, this inequality must be true too, because $P[t \leq T]$ can not drop faster than $e^{-\Gamma(t-t_0)}$, due to boundedness of the top-level population $p(t) \leq 1$. Set t_* to be the infimum of all such t' . Note that t_* must be strictly greater than t_0 due to continuity of $P[t \leq T]$. This concludes the proof of the last step and therefore of the proposition. $\square$

³ The set $\{t : P[t \leq T] < 1\}$ is lower bounded by 0 and non-empty by the claim made in step (a), thus, the infimum exists.

Chapter IV

The Role of Control for Ground-State Cooling

IV.1 Landauer vs. Nernst: What is the True Cost of Cooling a Quantum System?

IV.1.1 Contribution

I contributed to the conception of the project idea and the general theoretical framework. I contributed to early versions of some of the analytic calculations concerning the incoherent regime. I contributed to writing and revising the main text. I produced Figure 1 and Figure 2 of the paper.

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

Philip Taranto,^{1,2,*} Faraj Bakhshinezhad,^{3,2,†} Andreas Bluhm,^{4,‡} Ralph Silva,^{5,‡} Nicolai Friis,^{1,2} Maximilian P. E. Lock,^{1,2} Giuseppe Vitagliano,² Felix C. Binder,^{6,2,1} Tiago Debarba,⁷ Emanuel Schwarzhans,^{2,1} Fabien Clivaz,^{8,2} and Marcus Huber^{1,2,§}

¹Atominstitut, Technische Universität Wien, 1020 Vienna, Austria

²Institute for Quantum Optics and Quantum Information - IQOQI Vienna, Austrian Academy of Sciences, Boltzmannngasse 3, 1090 Vienna, Austria

³Department of Physics and Nanolund, Lund University, Box 118, 221 00 Lund, Sweden

⁴QMATH, Department of Mathematical Sciences, University of Copenhagen, Universitetsparken 5, 2100 Copenhagen, Denmark

⁵Institute for Theoretical Physics, ETH Zürich, Wolfgang-Pauli-Str. 27, Zürich, Switzerland

⁶School of Physics, Trinity College Dublin, Dublin 2, Ireland

⁷Departamento Acadêmico de Ciências da Natureza, Universidade Tecnológica Federal do Paraná (UTFPR), Campus Cornélio Procopio, Avenida Alberto Carazzai 1640, Cornélio Procopio, Paraná 86300-000, Brazil

⁸Institut für Theoretische Physik und IQST, Universität Ulm, Albert-Einstein-Allee 11, D-89069 Ulm, Germany
(Dated: February 1, 2023)

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.

I. 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 [1, 2], high-precision quantum parameter estimation [3–5], and fault-tolerant quantum computation [6, 7]. Fundamentally, pure states are prerequisites for ideal measurements [8] and precise timekeeping [9, 10]. 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 $k_B T \log(2)$ [11]. Alternatively, one could consult Nernst’s unattainability principle (the third law of thermodynamics) [12], 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 [13–17]. *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, the control complexity of a protocol refers to the structure of machine energy gaps that the cooling unitary must couple the system to; we demonstrate that this energy-level spectrum must approximate a continuum in order to cool with minimal time and energy costs. In short, 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 [18] 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 explicit protocols that saturate the ultimate limits. We demonstrate that mitigating overall heat dissipation comes at the practical cost of controlling fine-tuned interactions that require a *coherent* external work source, i.e., a quantum battery [19–23]. From a thermodynamic perspective, this may seem somewhat 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 overall control, where interactions need only be switched on and off to generate transformations, i.e., a heat engine alone drives the dynamics [24–28]. The incoherent-control setting is therefore fully thermodynamically consistent inasmuch as both the machine state is assumed to be thermal (and to rethermalize between control

* philip.taranto@tuwien.ac.at; P. T. and F. B. contributed equally.

† P. T. and F. B. contributed equally.

‡ A. B. and R. S. contributed equally.

§ marcus.huber@tuwien.ac.at




FIG. 1. *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 S and machine 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).

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, subsequently derive a novel limit—which we dub the *Carnot-Landauer* bound—and construct protocols that saturate it, thereby establishing its significance. 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. [29] 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 additionally 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, which, by mitigating the cost of control at the very outset, could lead to tangible advantages.

Overview & Summary of Results

Loosely speaking, 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. From a physics perspective, said task can be rephrased as *perfectly cooling* a system to the ground

state, or more generally, preparing a pure state.¹

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, whose asymptotic trade-off relation is relatively well-established: 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.²

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 do such protocols require and how difficult are they to implement in practice? We address these questions by considering the task of cooling a quantum system in two extremal control paradigms (see Fig. 1): One driven by a *coherent* work source and the other by an *incoherent* heat engine.

After laying out the framework, we proceed to analyse the relationship between the aforementioned three resources for cooling. A core idea of this article originates from the observation that it is possible to perfectly cool a physical system with both finite energy and time. Although said observation is simple in nature inasmuch as it can be obtained by a shift in perspective of Landauer's original protocol, its consequences run deep: Indeed, the apparent tension between Landauer cooling and Nernst's unattainability principle that arises when only energy and time are considered as resources is resolved via the inclusion of control complexity as a consideration. Subsequently, we define a meaningful notion of control complexity in terms of the energy-level structure of the machine that the system must be coupled to throughout the cooling protocol and demonstrate its thermodynamic consistency by showing that it indeed must diverge to cool the system to the ground state at minimal energy cost, thereby reconciling the viewpoints of Landauer and Nernst.

Having established the relevant trinity of resources, we present 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 in two extremal control paradigms, as well as develop insight into the finite-resource regime for some special cases. We begin by outlining the framework.

II. FRAMEWORK: COOLING A PHYSICAL SYSTEM

Consider a target system $\mathcal{S}$ in an initial state $\rho_{\mathcal{S}}$ 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,

$$\rho_{\mathcal{M}} = \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}}) := \frac{e^{-\beta H_{\mathcal{M}}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})}, \quad (1)$$

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, a single 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 primitives³. 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

¹ Low temperature thermal states correspond to those with low information content, as they have low entropy or small effective support; viewing cooling more broadly (i.e., not restricting to thermal states and allowing for arbitrary Hamiltonians), we see that cooling indeed encompasses information erasure: States with smaller effective support are "colder" than those with greater support according to any meaningful notion of "cool" (see Ref. [30]).

² Note, however, that although the asymptotic trade-off relationship is known, the connection between energy and time in the finite-resource setting remains unresolved: For instance, if one uses twice the amount of energy, it is not clear how much faster a given protocol can be implemented; we provide some preliminary insight to such questions in Sec. VI.

³ One 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.

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.

III. COHERENT CONTROL

We begin by considering cooling with coherently-controlled resources (see Fig. 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 more specifically on the nature and role of control complexity.

A. Energy, Time, and Control Complexity as Resources

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 $\varrho_M = \tau_M(\beta, H_M)$, i.e.,

$$\varrho'_S := \text{tr}_M [U(\varrho_S \otimes \varrho_M)U^\dagger]. \quad (2)$$

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 $\Delta E_M := \text{tr}[H_M(\varrho'_M - \varrho_M)]$, where $\varrho'_M := \text{tr}_S [U(\varrho_S \otimes \varrho_M)U^\dagger]$. The latter is associated with the heat dissipated into the environment and is given by [29]

$$\beta \Delta E_M = \tilde{\Delta} S_S + I(\mathcal{S} : \mathcal{M})_{\varrho'_{SM}} + D(\varrho'_M \| \varrho_M), \quad (3)$$

where $S(\varrho) := -\text{tr}[\varrho \log(\varrho)]$ is the von Neumann entropy, $\tilde{\Delta} S_A := S(\varrho_A) - S(\varrho'_A)$ ⁴, $I(\mathcal{A} : \mathcal{B})_{\varrho_{AB}} := S(\varrho_A) + S(\varrho_B) - S(\varrho_{AB})$ (with marginals $\varrho_{A/B} := \text{tr}_{B/A}[\varrho_{AB}]$) is the mutual 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. (3) and its generalisation to the incoherent-control setting in Appendix A. The mutual information is non-negative and vanishes iff $\varrho_{AB} = \varrho_A \otimes \varrho_B$; similarly, the relative entropy

	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 I. *Coherent-control cooling protocols for finite-dimensional (qudit) and harmonic oscillator systems.* Landauer energy cost refers to saturation of Eq. (4) and complexity refers to the proxy measure effective dimension (see Def. 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: $d_S = d_M =: d$.

is non-negative and vanishes iff $\varrho = \sigma$. Dropping these terms leads to the Landauer bound [11]

$$\beta \Delta E_M \geq \tilde{\Delta} S_S. \quad (4)$$

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 $\varrho_S = \tau_S(\beta, H_S)$, 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. [31] 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. (4) with ϱ'_S pure as *perfect cooling at the Landauer limit*.

Although 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 (see below). As summarised in Table I, we now present coherently-controlled protocols that perfectly cool an arbitrary finite-dimensional target system using thermal machines when any one of the three considered resources—energy, time or control complexity—diverges; moreover, the resources that are kept finite saturate protocol-independent ultimate bounds. The following thus provides a comprehensive analysis of cooling with respect to the trinity of resources that can be traded off amongst each other.

⁴ Note the differing sign conventions (denoted by the tilde) that we use for changes in energies, $\Delta E_X := E'_X - E_X$, and in entropies, $\tilde{\Delta} S_X := S_X - S'_X$, such that energy *increases* and entropy *decreases* are positive.

B. Perfect Cooling at the Ultimate Limits with Infinite Resources

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

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 subspace of the target system. 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. Moreover, in addition to being sufficient for perfect cooling with both finite time and control complexity, any protocol that cools perfectly with both finite time and control complexity requires diverging energy. See Appendix B for details.

We now move to consider the situations in which the energy cost is minimised at the expense of either diverging time or control complexity. Equation (3) provides insight for understanding the conditions required for saturating the Landauer bound. Although for finite-dimensional machines only trivial processes of the form $U_{S\mathcal{M}} = U_S \otimes \mathbb{1}_{\mathcal{M}}$ saturate the Landauer limit [29], 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 $I(S : \mathcal{M})_{\rho'_{S\mathcal{M}}} \rightarrow 0$ and effectively not disturb the machine, i.e., yield $\rho'_{\mathcal{M}} \rightarrow \rho_{\mathcal{M}}$ such that $D(\rho'_{\mathcal{M}} \| \rho_{\mathcal{M}}) \rightarrow 0$. Indeed, any correlations created between initially thermal systems would come at the expense of an additional energetic cost [32–34] whose minimisation is a problem that has so far only been partially resolved [35]. However, it has been shown that for any (strictly) rank non-decreasing process, there exists a thermal machine and joint unitary such that for any $\epsilon > 0$, the heat dissipated satisfies $\beta \Delta E_{\mathcal{M}} \leq \tilde{\Delta} S_S + \epsilon$ [29], 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.

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 [20, 29, 36]:

Theorem 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 $\rho_S = \tau_S(\beta, H_S)$ 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 C, along with a more detailed dimension-dependent energy cost function for the special case of equally spaced Hamiltonians.

Through the protocol described above, we see that given a diverging amount of time, the target system can be sequentially coupled with a machine of finite complexity that rethermalizes between control steps in such a way that the final target system state is arbitrarily close to the ground state for any initial temperature. This trade-off between energy and time is well-known, and we discuss it only briefly in order to help build intuition and highlight the versatility of our framework. 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 control complexity.

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. 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 control 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 energy requirements follows from our constructive proof of Theorem 2. We therefore have:

Corollary 1. *With diverging control complexity, any finite-dimensional quantum system can be perfectly cooled at the Landauer limit in finite time.*

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. Indeed, the inclusion of an explicit quantifier of control complexity regarding thermodynamic tasks—which, although

crucial for practical purposes, is oftentimes overlooked—is one of the main novelties of our present work.

IV. CONTROL COMPLEXITY IN QUANTUM THERMODYNAMICS

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. 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*. We provide a detailed study of structural and control complexity in Appendix D, and here summarise the key methods.

A. Structural & Dynamical Notions of 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 (see Fig. 2).

1. Structural Complexity

Regarding the former, first note that one can lower-bound the smallest eigenvalue $\lambda_{\min}$ of the final state ϱ'_S (and hence how cold the system can become) after *any* unitary interaction with a thermal machine by [29]

$$\lambda_{\min}(\varrho'_S) \geq e^{-\beta \omega_{\mathcal{M}}^{\max}} \lambda_{\min}(\varrho_S), \quad (5)$$

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




FIG. 2. *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. (5)]. 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.

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. [29], 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 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 energy gap 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. 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 dynamical properties of the unitary that represents a particular protocol.

2. 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.

In the following sections, we begin by demonstrating that any cooling protocol that achieves perfect cooling with minimal time and energy resources requires coupling the target system to an infinite-dimensional machine, thereby capturing a notion of control complexity that satisfies the second point above. However, by subsequently analysing the sufficient conditions for such optimal cooling, we see that such a condition is in general insufficient to achieve said goal; furthermore, coupling to an infinite-dimensional machine is not necessarily difficult to implement in practice in certain experimental platforms. The insights gained here finally motivate our more refined notion of control complexity, namely that the system must be coupled to a spectrum of machine energy gaps that approximate a continuum. This condition is indeed difficult to achieve in all experimental settings and therefore provides a reasonable definition of control complexity inasmuch as it satisfies both desiderata outlined above.

B. Effective Dimension as a Notion of Control Complexity

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. 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 $U_{S\mathcal{M}} = U_{\mathcal{A}} \oplus \mathbb{1}_{\mathcal{A}^\perp}$:

$$d^{\text{eff}} := \min \dim(\mathcal{A}) : U_{S\mathcal{M}} = U_{\mathcal{A}} \oplus \mathbb{1}_{\mathcal{A}^\perp}. \quad (6)$$

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.

We begin by demonstrating that the effective dimension (nontrivially) accessed by a unitary (see Def. 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 and provides a necessary condition. 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 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. [29] then implies that there are finite-dimensional correction terms such that the Landauer limit cannot be saturated.

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 without any restrictions on the types of individual operations allowed. This highlights a certain level of generality regarding the definition put forth, inasmuch as it is not tied to any presupposed structure of the systems at hand or the ability of the agent to control them. Additionally, as we will discuss below, in many situations of interest, such as a machine comprising a collection of qubits

and/or natural gate-set limitations, said definition also corresponds to protocols that are difficult to implement in practice, therefore also satisfying the above point 1. However, such additional restrictions are by no means generic. Moreover, it is a priori unclear if having a diverging effective dimension is enough to permit perfect cooling with minimal time and energy cost. We now move on to discuss the connection to practical difficulty in general before analysing sufficient conditions regarding control complexity.

1. Correspondence to Practical Difficulty

Importantly, if one supposes that the system and machines are finite-dimensional, then diverging effective dimension implies diverging circuit complexity, where the latter is defined in terms of the minimum number of gates (from a predetermined set of possibilities) required to implement the overall circuit representing a particular protocol. For instance, considering a qubit system and machines, and the ability to perform arbitrary two-qubit gates, the effective dimension is simply the logarithm of the number of distinct machine qubits that the system interacts with throughout the protocol. For any cooling protocol that achieves Landauer energy cost, it is clear that every one of a diverging number of qubit machines must take part in the overall transformation. Moreover, the particular interactions applied can be taken to be swap gates, which require the ability for the agent to be able to perform a CNOT gate, which in turn permits universal quantum computation with two-qubit interactions. Thus, given a universal two-qubit gate set, the circuit required to perform perfect cooling at minimal energy cost has a complexity that scales with the number of machine qubits. For higher-dimensional architectures or further restrictions on the gate set, any meaningful notion of control complexity will increase accordingly. This means that the task of cooling a finite-dimensional system with finite-dimensional machines at the Landauer limit is—even with a perfect quantum computer—an impossibly difficult task.

However, although our proposed definition of effective dimension as a notion of control complexity is flexible inasmuch as it applies to arbitrary system-machine structures, the price of such generality comes with the drawback that it tends to overestimate the difficulty of implementing a particular protocol in practice. In other words, without imposing any additional assumptions regarding the situation at hand, the effective dimension does not necessarily satisfy the above point 1. For example, whilst the effective dimension and the circuit complexity coincide for qubits, in higher-dimensional settings, the former overestimates the latter because not all system-machine subspaces are necessarily required to implement a particular protocol (i.e., although using all such subspaces provides one way to achieve it, this is not unique). Thus, the extent to which the circuit complexity is overestimated depends on the allowed gate set that is considered “simple” in general. At the extreme end, i.e., for harmonic oscillator systems and machines, this can be seen from the fact that a single beam-splitter operation (which is a two-mode Gaussian operation, corresponding to a simple circuit com-

plexity in the usual sense considered for infinite-dimensional quantum circuit architectures) already has infinite effective dimension, but is far from sufficient to achieve perfect cooling at Landauer cost.

As a representative for infinite-dimensional systems, we treat harmonic oscillator target systems separately in Appendix E. 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 Eq. (6): As mentioned above, 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. [22, 37]), but nonetheless with infinite effective dimensionality according to Def. 1]. This result highlights that the polynomial degree of the generators of a particular protocol would—somewhat counterintuitively, since operations corresponding to high polynomial degree are difficult to achieve in practice—not provide a suitable measure of control complexity inasmuch as its divergence is not necessary for Landauer-cost cooling. 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 similar protocol that saturates the Landauer limit exists in this setting remains an open question.

2. Sufficiency for Optimal Cooling

Thus, in general, accessing an infinite-dimensional machine subspace is not sufficient for reaching the Landauer limit. 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, which intuitively corresponds to complicated multipartite gates and demonstrates that an operationally meaningful notion of control complexity must take into account factors beyond the effective dimensionality accessed by an operation. 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 eigenval-

ues depends upon properties of both the target and machine. To highlight this, in Appendix D, 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. Drawing from these insights, in the coming section we propose a stronger notion of control complexity (in the sense that it bounds the effective dimension from below and that it corresponds to practical difficulty in virtual every setting imaginable) in terms of the energy-gap structure of the machine and demonstrate that this measure too must diverge to cool perfectly with minimal time and energy costs. 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. [38–40]).

C. Energy-Gap Variety as a Notion of Control Complexity

This analysis motivates searching for a more detailed notion of control complexity that takes the energy-level structure of the machine into account which holds across all platforms and dimension scales. 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 introduced above 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 control complexity can be extended in order to more faithfully represent what permits saturation of the ultimate limitations and is difficult to achieve in practice.

Looking at all of our cooling protocols, a common property that seems to be important in minimising the energy cost of cooling is that the system is coupled to a set of machine energy gaps that 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 sets the final achievable temperature of the system (for perfect cooling to the ground state, note that one requires $\omega^* \rightarrow \infty$). Let us denote the number of distinct energy gaps in a (fixed) interval as the *energy-gap variety*. More formally, we have:

Definition 2. Consider an interval $[\omega_a, \omega_b) \subseteq \mathbb{R}$. We denote the energy-gap variety $\mathcal{E}_{[\omega_a, \omega_b)}$ in terms of the set of machine

energy gaps that lie in said interval, i.e.,

$$\mathcal{E}_{[\omega_a, \omega_b)} := \{\omega_\gamma := \omega_i - \omega_j \mid \omega_i - \omega_j \in [\omega_a, \omega_b)\}_{\gamma}. \quad (7)$$

The number of distinct elements in such a set is the *energy-gap variety*.

On the one hand, it is clear that coupling a system to a large number of distinct and/or closely-spaced energy gaps requires fine-tuned control that is difficult in any experimental setting. On the other, the energy-gap variety lower bounds the effective dimension, and thus it is not clear that it needs to diverge in order to cool at Landauer energy cost. In Appendix D, we demonstrate that the energy-gap variety must indeed diverge and, additionally, that the set of energy gaps must densely cover a relevant interval (whose endpoints set the amount of cooling possible) in order to perfectly cool at the Landauer limit. Taken in combination with its sufficiency to achieve said task, this result posits the energy-gap variety as a better quantifier of control complexity than the effective dimension, constituting the best thermodynamically meaningful notion of control complexity that we have put forth so far.

Aside from introducing and highlighting the important role of control complexity, we now take a step back to consider the notion of overall control at a higher level. It is clear that the protocols that saturate the Landauer limit for the energy cost of cooling require highly controlled microstate interactions between the system and machine; in turn, such transformations necessitate that the agent has access to a versatile *work source*, i.e., either a quantum battery [19–23] or a classical work source with a precise clock [9, 10]. Such control is reminiscent of Maxwell’s demon, who can indeed address all microscopic configurations at hand. This level of control is, however, in some sense at odds with the true spirit of thermodynamics. Indeed, the very reason that the machine is taken to begin as a thermal (Gibbs) state in thermodynamics is precisely because it provides the microscopic description that is *both* consistent with macroscopic observations (in particular, average energy) and makes minimal assumptions regarding the information that the agent has about the initial state; thermodynamics as a whole is largely concerned with what can be done with minimal information requirements. Beginning with this, and then going on to permit dynamical interactions that address the full complex microstructure is somewhat contradictory, at least in essence; indeed, it has been argued that “Maxwell’s demon cannot operate” [41] as an autonomous thermal being. Thus, a more thermodynamically sound setting would also restrict the transformations themselves to be ones that can be driven with minimal overall control. We now move to analyse the task of cooling within such a context.

V. INCOHERENT CONTROL (HEAT ENGINE)

The 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 in general not possible to cool a target system with machines that are initially thermal at a single temperature, as was considered in the coherent-control paradigm [42]. 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 $\beta_H < \beta$ [38, 42] (see Fig. 1, bottom panel). In other words, one uses a hot and a 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 and serves to highlight the fact that the incoherent-control setting is a fundamentally distinct paradigm that must be considered independently. Here, we focus on finite-dimensional systems and leave the analysis of infinite-dimensional ones to future work.

A. Ultimate Limits in the Incoherent Control Paradigm

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 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 $\eta := 1 - \beta_H/\beta \in (0, 1)$ be the Carnot efficiency with respect to the hot and cold baths. In the incoherent-control setting, the quantity

$$\begin{aligned} \Delta F_s^{(\beta)} + \eta \Delta E_{\mathcal{H}} \\ = -\frac{1}{\beta} [\Delta S_s + \Delta S_c + \Delta S_{\mathcal{H}} + D(\varrho'_c || \varrho_c) + D(\varrho'_{\mathcal{H}} || \varrho_{\mathcal{H}})] \end{aligned} \quad (8)$$

satisfies the inequality

$$\Delta F_s^{(\beta)} + \eta \Delta E_{\mathcal{H}} \leq 0. \quad (9)$$

Equation (9) 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, 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). On the other hand, in the incoherent-control setting, we have the following no-go theorem (see Appendix F for a proof):

Theorem 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.

B. Saturating the Carnot-Landauer Limit

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 F, 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 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 G, we detail cooling protocols that saturate the Carnot-Landauer limit for any finite-temperature heat bath. More precisely, we prove:

Theorem 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.

VI. 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 [42]. 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.

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 utilized. 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, $\omega_{\max}$ [see Eq. (5)]. 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 [42]. 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 D4 and inspired by one presented in Ref. [29] (see Appendix D therein); we hence refer to it as the 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,




FIG. 3. *Imperfect Cooling.* We compare the cooling performance of a degenerate qubit target system using either N machine qubits of linearly increasing energy accessed sequentially or a single unitary on a 2^N -dimensional machine, the latter being a finite adaptation of a protocol presented Ref. [29]. We set $\beta = 1$, choose units such that $\hbar = k_B = 1$, 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, $W - \beta \tilde{\Delta} S_S$ (in units of the smallest machine energy gap, $\omega_{\mathcal{M}}^{\min}$), 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.

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 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. 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 H, 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 H explicitly. Indeed, determining the optimal cooling protocols for a range of realistic assumptions remains a major open avenue.

VII. DISCUSSION

Relation to Previous Works

A vast amount of literature concerning quantum thermodynamics considers resource theories (see Refs. [43, 44] 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. [20, 36] and [29], studied the potential saturation of the second law of thermodynamics and Landauer’s limit, respectively. Refs. [20] and [36] 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. [29] 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, as we present in terms of the “energy-gap variety”. 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 [29] and inequality [11], respectively, where the protocol can only be run via a heat engine.

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 [18], 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 energy-gap variety accessed by a unitary protocol as a meaningful quantifier of control complexity that is both fully consistent with the third law of thermodynamics and chimes well with what is difficult to achieve in practice. 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 [45]. 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 [38–40]. 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 [8, 46] and non-negligible back-action on the operation of the machine itself [47]. 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.

ACKNOWLEDGMENTS

The authors thank Elizabeth Agudelo and Paul Erker for very insightful discussions at the early stages of the project. P. T. is supported by the Austrian Science Fund (FWF) project: Y879-N27 (START). F. B. is supported by grant number FQXi Grant Number: FQXi-IAF19-07 from the Foundational Questions Institute Fund, a donor advised fund of Silicon Valley Community Foundation. A. B. acknowledges support from the VILLUM FONDEN via the QMATH Centre of Excellence (Grant no. 10059) and from the QuantERA ERA-NET Cofund in Quantum Technologies implemented within the European Union’s Horizon 2020 Programme (QuantAlgo project) via the Innovation Fund Denmark. R. S. acknowledges funding from the Swiss National Science Foundation via an Ambizione grant PZ00P2_185986. N. F. is supported by the Austrian Science Fund (FWF) project: P 31339-N27. M. P. E. L. acknowledges financial support by the ESQ (Erwin Schrödinger Center for Quantum Science & Technology) Discovery programme, hosted by the Austrian Academy of Sciences (ÖAW) and TU Wien. G. V. is supported by the Austrian Science Fund (FWF) projects ZK 3 (Zukunftskolleg) and M 2462-N27 (Lise-Meitner). F. C. B. acknowledges support from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 801110 and the Austrian Federal Ministry of Education, Science and Research (BMBWF). This project reflects only the authors’ view, the EU Agency is not responsible for any use that may be made of the information it contains. T. D. acknowledges support from the Brazilian agency CNPq INCT-IQ through the project (465469/2014-0). E. S. is supported by the Austrian Science Fund (FWF) project: Y879-N27 (START). F. C. is supported by the ERC Synergy grant HyperQ (Grant No. 856432). M. H. is supported by the Austrian Science Fund (FWF) project: Y879-N27 (START), and acknowledges financial support by the ESQ (Erwin Schrödinger Center for Quantum Science & Technology) Discovery programme, hosted by the Austrian Academy of Sciences (ÖAW).

- [1] Nicolas Gisin, Grégoire Ribordy, Wolfgang Tittel, and Hugo Zbinden, *Quantum cryptography*, *Rev. Mod. Phys.* **74**, 145 (2002), [arXiv:quant-ph/0101098](#).
- [2] Stefano Pirandola, Ulrik L. Andersen, Leonardo Banchi, Mario Berta, Darius Bunandar, Roger Colbeck, Dirk Englund, Tobias Gehring, Cosmo Lupo, Carlo Ottaviani, Jason L. Pereira, Mohsen Razavi, Jesni Shamsul Shaari, Marco Tomamichel, Vladyslav C. Usenko, Giuseppe Vallone, Paolo Villoresi, and Petros Wallden, *Advances in quantum cryptography*, *Adv. Opt. Photon.* **12**, 1012 (2020), [arXiv:1906.01645](#).
- [3] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone, *Advances in quantum metrology*, *Nat. Photonics* **5**, 222 (2011), [arXiv:1102.2318](#).
- [4] Géza Tóth and Iagoba Apellaniz, *Quantum metrology from a quantum information science perspective*, *J. Phys. A: Math. Theor.* **47**, 424006 (2014), [arXiv:1405.4878](#).
- [5] Rafał Demkowicz-Dobrzański, Marcin Jarzyna, and Janek Kołodyński, *Quantum limits in optical interferometry*, *Prog. Optics* **60**, 345 (2015), [arXiv:1405.7703](#).
- [6] John Preskill, *Fault-tolerant quantum computation*, in *Introduction to Quantum Computation*, edited by H.-K. Lo, Sandu Popescu, and T. P. Spiller (World-Scientific, 1997) Chap. 8, pp. 213–269, [arXiv:quant-ph/9712048](#).
- [7] John Preskill, *Quantum Computing in the NISQ era and beyond*, *Quantum* **2**, 79 (2018), [arXiv:1801.00862](#).
- [8] Yelena Guryanova, Nicolai Friis, and Marcus Huber, *Ideal Projective Measurements Have Infinite Resource Costs*, *Quantum* **4**, 222 (2020), [arXiv:1805.11899](#).
- [9] Paul Erker, Mark T. Mitchison, Ralph Silva, Mischa P. Woods, Nicolas Brunner, and Marcus Huber, *Autonomous Quantum Clocks: Does Thermodynamics Limit Our Ability to Measure Time?* *Phys. Rev. X* **7**, 031022 (2017).
- [10] Emanuel Schwarzhans, Maximilian P. E. Lock, Paul Erker, Nicolai Friis, and Marcus Huber, *Autonomous Temporal Probability Concentration: Clockworks and the Second Law of Thermodynamics*, *Phys. Rev. X* **11**, 011046 (2021), [arXiv:2007.01307](#).
- [11] Rolf Landauer, *Irreversibility and Heat Generation in the Computing Process*, *IBM J. Res. Dev.* **5**, 183 (1961).
- [12] Walther Nernst, *Über die Beziehung zwischen Wärmeentwicklung und maximaler Arbeit bei kondensierten Systemen*, in *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften* (Berlin, 1906) pp. 933–940.
- [13] Francesco Ticozzi and Lorenza Viola, *Quantum resources for purification and cooling: fundamental limits and opportunities*, *Sci. Rep.* **4**, 5192 (2014), [arXiv:1403.8143](#).
- [14] Lluís Masanes and Jonathan Oppenheim, *A general derivation and quantification of the third law of thermodynamics*, *Nat. Commun.* **8**, 14538 (2017), [arXiv:1412.3828](#).
- [15] Henrik Wilming and Rodrigo Gallego, *Third Law of Thermodynamics as a Single Inequality*, *Phys. Rev. X* **7**, 041033 (2017), [arXiv:1701.07478](#).
- [16] Nahuel Freitas, Rodrigo Gallego, Lluís Masanes, and Juan Pablo Paz, *Cooling to Absolute Zero: The Unattainability Principle*, in *Thermodynamics in the Quantum Regime*, edited by Felix Binder, Luis A. Correa, Christian Gogolin, Janet Anders, and Gerardo Adesso (Springer International Publishing, Cham, Switzerland, 2018) Chap. 25, pp. 597–622, [arXiv:1911.06377](#).
- [17] Jakob Scharlau and Markus P. Müller, *Quantum Horn’s lemma, finite heat baths, and the third law of thermodynamics*, *Quantum* **2**, 54 (2018), [arXiv:1605.06092](#).
- [18] Antonio Acín, Immanuel Bloch, Harry Buhrman, Tommaso Calarco, Christopher Eichler, Jens Eisert, Daniel Esteve, Nicolas Gisin, Steffen J. Glaser, Fedor Jelezko, Stefan Kuhr, Maciej Lewenstein, Max F. Riedel, Piet O. Schmidt, Rob Thew, Andreas Wallraff, Ian Walmsley, and Frank K. Wilhelm, *The quantum technologies roadmap: a European community view*, *New J. Phys.* **20**, 080201 (2018), [arXiv:1712.03773](#).
- [19] J. Åberg, *Truly work-like work extraction via a single-shot analysis*, *Nat. Commun.* **4**, 1925 (2013), [arXiv:1110.6121](#).
- [20] Paul Skrzypczyk, Anthony J. Short, and Sandu Popescu, *Work extraction and thermodynamics for individual quantum systems*, *Nat. Commun.* **5**, 4185 (2014), [arXiv:1307.1558](#).
- [21] Matteo Lostaglio, David Jennings, and Terry Rudolph, *Description of quantum coherence in thermodynamic processes requires constraints beyond free energy*, *Nat. Commun.* **6**, 6383 (2015), [arXiv:1405.2188](#).
- [22] Nicolai Friis and Marcus Huber, *Precision and Work Fluctuations in Gaussian Battery Charging*, *Quantum* **2**, 61 (2018), [arXiv:1708.00749](#).
- [23] Francesco Campaioli, Felix A. Pollock, and Sai Vinjanampathy, *Quantum Batteries*, in *Thermodynamics in the Quantum Regime*, edited by Felix Binder, Luis A. Correa, Christian Gogolin, Janet Anders, and Gerardo Adesso (Springer, 2019) Chap. 8, pp. 207–225, [arXiv:1805.05507](#).
- [24] Henry E. D. Scovil and Erich O. Schulz-DuBois, *Three-Level Masers as Heat Engines*, *Phys. Rev. Lett.* **2**, 262 (1959).
- [25] Ronnie Kosloff and Amikam Levy, *Quantum Heat Engines and Refrigerators: Continuous Devices*, *Annu. Rev. Phys. Chem.* **65**, 365 (2014), [arXiv:1310.0683](#).
- [26] Raam Uzdin, Amikam Levy, and Ronnie Kosloff, *Equivalence of Quantum Heat Machines, and Quantum-Thermodynamic Signatures*, *Phys. Rev. X* **5**, 031044 (2015), [arXiv:1502.06592](#).
- [27] Mark T. Mitchison, *Quantum thermal absorption machines: refrigerators, engines and clocks*, *Contemp. Phys.* **60**, 164 (2019), [arXiv:1902.02672](#).
- [28] Mischa P. Woods, Nelly Huei Ying Ng, and Stephanie Wehner, *The maximum efficiency of nano heat engines depends on more than temperature*, *Quantum* **3**, 177 (2019), [arXiv:1506.02322](#).
- [29] David Reeb and Michael M. Wolf, *An improved Landauer principle with finite-size corrections*, *New J. Phys.* **16**, 103011 (2014), [arXiv:1306.4352](#).
- [30] Fabien Clivaz, *Optimal Manipulation Of Correlations And Temperature In Quantum Thermodynamics*, *Ph.D. thesis*, University of Geneva (2020), [arXiv:2012.04321](#).
- [31] Paul M. Riechers and Mile Gu, *Impossibility of achieving Landauer’s bound for almost every quantum state*, *Phys. Rev. A* **104**, 012214 (2021), [arXiv:2103.02337](#).
- [32] Marcus Huber, Martí Perarnau-Llobet, Karen V. Hovhannisyan, Paul Skrzypczyk, Claude Klöckl, Nicolas Brunner, and Antonio Acín, *Thermodynamic cost of creating correlations*, *New J. Phys.* **17**, 065008 (2015), [arXiv:1404.2169](#).
- [33] David E. Bruschi, Martí Perarnau-Llobet, Nicolai Friis, Karen V. Hovhannisyan, and Marcus Huber, *The thermodynamics of creating correlations: Limitations and optimal protocols*, *Phys. Rev. E* **91**, 032118 (2015), [arXiv:1409.4647](#).
- [34] Giuseppe Vitagliano, Claude Klöckl, Marcus Huber, and Nicolai Friis, *Trade-off between work and correlations in quantum thermodynamics*, in *Thermodynamics in the Quantum Regime*, edited by Felix Binder, Luis A. Correa, Christian Gogolin, Janet Anders, and Gerardo Adesso (Springer, 2019) Chap. 30, pp.

- 731–750, [arXiv:1803.06884](#).
- [35] Faraj Bakhshinezhad, Fabien Clivaz, Giuseppe Vitagliano, Paul Erker, Ali T. Rezakhani, Marcus Huber, and Nicolai Friis, *Thermodynamically optimal creation of correlations*, *J. Phys. A: Math. Theor.* **52**, 465303 (2019), [arXiv:1904.07942](#).
 - [36] Janet Anders and Vittorio Giovannetti, *Thermodynamics of discrete quantum processes*, *New J. Phys.* **15**, 033022 (2013), [arXiv:1211.0183](#).
 - [37] E. G. Brown, N. Friis, and Marcus Huber, *Passivity and practical work extraction using Gaussian operations*, *New J. Phys.* **18**, 113028 (2016), [arXiv:1608.04977](#).
 - [38] Fabien Clivaz, Ralph Silva, Géraldine Haack, Jonatan Bohr Brask, Nicolas Brunner, and Marcus Huber, *Unifying paradigms of quantum refrigeration: fundamental limits of cooling and associated work costs*, *Phys. Rev. E* **100**, 042130 (2019), [arXiv:1710.11624](#).
 - [39] Philip Taranto, Faraj Bakhshinezhad, Philipp Schüttelkopf, Fabien Clivaz, and Marcus Huber, *Exponential Improvement for Quantum Cooling through Finite-Memory Effects*, *Phys. Rev. Appl.* **14**, 054005 (2020), [arXiv:2004.00323](#).
 - [40] Yi-Zheng Zhen, Dario Egloff, Kavan Modi, and Oscar Dahlsten, *Universal Bound on Energy Cost of Bit Reset in Finite Time*, *Phys. Rev. Lett.* **127**, 190602 (2021), [arXiv:2106.00580](#).
 - [41] L. Brillouin, *Maxwell's Demon Cannot Operate: Information and Entropy. I*, *J. Appl. Phys.* **22**, 334 (1951).
 - [42] Fabien Clivaz, Ralph Silva, Géraldine Haack, Jonatan Bohr Brask, Nicolas Brunner, and Marcus Huber, *Unifying Paradigms of Quantum Refrigeration: A Universal and Attainable Bound on Cooling*, *Phys. Rev. Lett.* **123**, 170605 (2019), [arXiv:1903.04970](#).
 - [43] Nelly H. Y. Ng and Mischa P. Woods, *Resource Theory of Quantum Thermodynamics: Thermal Operations and Second Laws*, in *Thermodynamics in the Quantum Regime*, edited by Felix Binder, Luis A Correa, Christian Gogolin, Janet Anders, and Gerardo Adesso (Springer International Publishing, Cham, Switzerland, 2018) Chap. 26, pp. 625–650, [arXiv:1805.09564](#).
 - [44] Matteo Lostaglio, *An introductory review of the resource theory approach to thermodynamics*, *Rep. Prog. Phys.* **82**, 114001 (2019), [arXiv:1807.11549](#).
 - [45] Charles H. Bennett, *The thermodynamics of computation – a review*, *Int. J. Theor. Phys.* **21**, 905 (1982).
 - [46] Tiago Debarba, Gonzalo Manzano, Yelena Guryanova, Marcus Huber, and Nicolai Friis, *Work estimation and work fluctuations in the presence of non-ideal measurements*, *New J. Phys.* **21**, 113002 (2019), [arXiv:1902.08568](#).
 - [47] Gonzalo Manzano, Gian-Luca Giorgi, Rosario Fazio, and Roberta Zambrini, *Boosting the performance of small autonomous refrigerators via common environmental effects*, *New J. Phys.* **21**, 123026 (2019), [arXiv:1908.10259](#).
 - [48] Wim van Dam and Patrick Hayden, *Rényi-entropic bounds on quantum communication*, [arXiv:quant-ph/0204093](#) (2002).
 - [49] Masanori Ohya and Dénes Petz, *Quantum Entropy and Its Use*, 2nd ed. (Springer, 1993) Chap. 2, pp. 37–46.
 - [50] Masanori Ohya and Noboru Watanabe, *Quantum Entropy and Its Applications to Quantum Communication and Statistical Physics*, *Entropy* **12**, 1194 (2010).
 - [51] Armen E. Allahverdyan, Karen V. Hovhannisyan, Dominik Janzing, and Guenter Mahler, *Thermodynamic limits of dynamic cooling*, *Phys. Rev. E* **84**, 041109 (2011).
 - [52] James Ladyman, James Lambert, and Karoline Wiesner, *What is a complex system?* *Eur. J. Philos.* **3**, 33 (2013).
 - [53] Yuriy Holovatch, Ralph Kenna, and Stefan Thurner, *Complex systems: physics beyond physics*, *Eur. J. Phys.* **38**, 023002 (2017), [arXiv:1610.01002](#).
 - [54] Albert W. Marshall, Ingram Olkin, and Barry C. Arnold, *Inequalities: Theory of Majorization and its Applications*, 2nd ed., Vol. 143 (Springer, 2011) Chap. 9, pp. 297–365.
 - [55] Walther Thirring, *Quantum Mathematical Physics: Atoms, Molecules and Large Systems* (Springer, 2002) Chap. 7, pp. 423–499.
 - [56] George Gamow, *One Two Three... Infinity: Facts and Speculations of Science* (Viking Press, New York, 1947).
 - [57] Gerardo Adesso, Sammy Ragy, and Antony R. Lee, *Continuous Variable Quantum Information: Gaussian States and Beyond*, *Open Syst. Inf. Dyn.* **21**, 1440001 (2014), [arXiv:1401.4679](#).
 - [58] Ralph Silva, Gonzalo Manzano, Paul Skrzypczyk, and Nicolas Brunner, *Performance of autonomous quantum thermal machines: Hilbert space dimension as a thermodynamical resource*, *Phys. Rev. E* **94**, 032120 (2016), [arXiv:1604.04098](#).

Supplemental Material

CONTENTS

A. Equality Forms of the (Carnot-)Landauer Limit	16
A1. Coherent-Control Paradigm: The Landauer Limit	17
A2. Incoherent-Control Paradigm: The Carnot-Landauer Limit	17
B. Diverging Energy	19
B1. Sufficiency: Diverging Energy Cooling Protocol	19
B2. Necessity of Diverging Energy for Protocols with Finite Time and Control Complexity	19
C. Diverging Time Cooling Protocol for Finite-Dimensional Systems	20
C1. Proof of Theorem 2	20
C2. Special Case: Equally Spaced Hamiltonians	22
D. Conditions for Structural and Control Complexity	25
D1. Necessary Complexity Conditions	25
D1a. Necessary Structural Conditions	25
D1b. Necessary Control Complexity Conditions	25
D2. Proof of Theorem 3, Corollary 2, and Theorem 10	26
D3. Sufficient Complexity Conditions	28
D4. Fine-Tuned Control Conditions	29
D5. Energy-Gap Variety as a Notion of Control Complexity	36
E. Diverging-Time / Control Complexity Cooling Protocols for Harmonic Oscillators	37
E1. Preliminaries	38
E2. Diverging-Time Cooling Protocol for Harmonic Oscillators	39
E2a. Diverging-Time Protocol using Gaussian Operations (with Diverging Control Complexity)	39
E2b. Diverging-Time Protocol using Non-Gaussian Operations (with Finite Control Complexity)	41
E3. Diverging Control Complexity Cooling Protocol for Harmonic Oscillators	43
F. Cooling Protocols in the Incoherent-Control Paradigm	45
F1. Diverging Energy: Proof of Theorem 5	45
F2. Diverging Time / Control Complexity	48
F3. Saturating the Landauer Limit with an Infinite-Temperature Heat Bath	49
F4. An Analysis of Finite-Temperature Heat Baths	50
G. Perfect Cooling at the Carnot-Landauer Limit in the Incoherent-Control Paradigm	50
G1. Qubit Case	51
G2. Qudit Case	55
H. Comparison of Cooling Paradigms and Resources for Imperfect Cooling	59
H1. Rates of Resource Divergence for Linear Qubit Machine Sequence	59
H2. Comparison of Coherent and Incoherent Control	60

Appendix A: 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 [11], 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., $\varrho_x \rightarrow p_x := (p_0, \dots, p_d)$ with $p_n = e^{-\beta E_n}$, $\text{tr}[H\varrho_x] \rightarrow \langle E \rangle_{p_x} := \sum_n p_n E_n$, $S(\varrho_x) \rightarrow S(p_x) := -\sum_n p_n \log(p_n)$, $\mathcal{Z}(\beta, H_x) = \sum_n e^{-\beta E_n}$, and so on. Nonetheless, all of the results presented hold for the more general “quantum” properties.

A1. Coherent-Control Paradigm: The Landauer Limit

The coherent setting was already studied in detail in Ref. [29], 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 1 ([29, Lemma 2]). *Let the setting be as above. Then*

$$[S(\varrho'_s) - S(\varrho_s)] + [S(\varrho'_m) - S(\varrho_m)] = I(\mathcal{S} : \mathcal{M})_{\varrho'_{\mathcal{SM}}} \geq 0. \quad (\text{A1})$$

Proof. We note that

$$[S(\varrho'_s) - S(\varrho_s)] + [S(\varrho'_m) - S(\varrho_m)] = S(\varrho'_s) + S(\varrho'_m) - S(\varrho'_{\mathcal{SM}}), \quad (\text{A2})$$

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 8 (Equality form of Landauer’s principle, [29, Theorem 3]). *Let the setting be as above. Then*

$$\beta \text{tr}[H_{\mathcal{M}}(\varrho'_m - \varrho_m)] - [S(\varrho_s) - S(\varrho'_s)] = I(\mathcal{S} : \mathcal{M})_{\varrho'_{\mathcal{SM}}} + D(\varrho'_m || \varrho_m) \geq 0. \quad (\text{A3})$$

Proof. From Lemma 1, it follows that

$$[S(\varrho_s) - S(\varrho'_s)] + I(\mathcal{S} : \mathcal{M})_{\varrho'_{\mathcal{SM}}} = S(\varrho'_m) - S(\varrho_m). \quad (\text{A4})$$

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 into Eq. (A4) yields the claimed equality. The inequality results from non-negativity of relative entropy and mutual information. This completes the proof. $\square$

A2. 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 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 9. *In the above setting, it holds that*

$$\Delta F_{\mathcal{S}}^{(\beta)} + \eta \Delta E_{\mathcal{H}} = -\frac{1}{\beta} [\Delta S_{\mathcal{S}} + \Delta S_{\mathcal{C}} + \Delta S_{\mathcal{H}} + D(\varrho'_c || \varrho_c) + D(\varrho'_h || \varrho_h)] \leq 0, \quad (\text{A5})$$

where $(0, 1) \ni \eta := 1 - \beta_H/\beta$ is the Carnot efficiency and $\Delta F_{\mathcal{S}}^{(\beta)} = F_{\beta}(\varrho'_{\mathcal{S}}) - F_{\beta}(\varrho_{\mathcal{S}})$.

Proof. Let us consider

$$I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{\mathcal{SCH}}} := S(\varrho'_{\mathcal{S}}) + S(\varrho'_c) + S(\varrho'_h) - S(\varrho'_{\mathcal{SCH}}) \geq 0. \quad (\text{A6})$$

Note that the quantity $I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{\mathcal{SCH}}}$, which quantifies the tripartite mutual information of the state $\varrho'_{\mathcal{SCH}}$, is non-negative via subadditivity $S(\varrho_{\mathcal{A}}) + S(\varrho_{\mathcal{B}}) \geq S(\varrho_{\mathcal{AB}})$ for any state $\varrho_{\mathcal{AB}}$. Furthermore, since the von Neumann entropy is invariant under unitary transformations and additive for tensor product states, we have

$$I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{\mathcal{SCH}}} = \Delta S_{\mathcal{S}} + \Delta S_{\mathcal{C}} + \Delta S_{\mathcal{H}}. \quad (\text{A7})$$

We also have that

$$\Delta S_{\mathcal{C}} = \beta \Delta E_{\mathcal{C}} - D(\varrho'_c || \varrho_c) \quad (\text{A8})$$

and

$$\Delta S_{\mathcal{H}} = \beta_H \Delta E_{\mathcal{H}} - D(\varrho'_h || \varrho_h). \quad (\text{A9})$$

Thus,

$$I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{\mathcal{SCH}}} = \Delta S_{\mathcal{S}} + \beta \Delta E_{\mathcal{C}} - D(\varrho'_c || \varrho_c) + \beta_H \Delta E_{\mathcal{H}} - D(\varrho'_h || \varrho_h). \quad (\text{A10})$$

Since the unitary is energy conserving, we infer that $\Delta E_{\mathcal{S}} + \Delta E_{\mathcal{C}} + \Delta E_{\mathcal{H}} = 0$. Hence, we have

$$\Delta S_{\mathcal{S}} - \beta \Delta E_{\mathcal{S}} + (\beta_H - \beta) \Delta E_{\mathcal{H}} = I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{\mathcal{SCH}}} + D(\varrho'_c || \varrho_c) + D(\varrho'_h || \varrho_h). \quad (\text{A11})$$

Using the free energy, we can rewrite this as

$$-\beta[F_{\beta}(\varrho'_{\mathcal{S}}) - F_{\beta}(\varrho_{\mathcal{S}})] - (\beta - \beta_H) \Delta E_{\mathcal{H}} = I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{\mathcal{SCH}}} + D(\varrho'_c || \varrho_c) + D(\varrho'_h || \varrho_h). \quad (\text{A12})$$

Dividing by $-\beta$, we obtain the assertion, since, in particular, $I(\mathcal{S} : \mathcal{C} : \mathcal{H})_{\varrho'_{\mathcal{SCH}}} + D(\varrho'_c || \varrho_c) + D(\varrho'_h || \varrho_h) \geq 0$ 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:

$$\text{tr}[H_{\mathcal{H}}(\varrho_{\mathcal{H}} - \varrho'_{\mathcal{H}})] \geq \frac{1}{\eta} [F_{\beta}(\varrho'_S) - F_{\beta}(\varrho_S)]. \quad (\text{A13})$$

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. (A10) leads to

$$-\tilde{\Delta}S_S - \beta_H \Delta E_S + (\beta - \beta_H) \Delta E_C \geq 0, \quad (\text{A14})$$

which recovers the standard Landauer bound for the dissipated heat in the limit of an infinitely-hot heat bath, i.e., $\beta_H \rightarrow 0$.

Appendix B: Diverging Energy

B1. 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.

B2. 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

$$\text{rank}[\tau_S(\beta, H_S)] \text{rank}[\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})] = \text{rank}[\varrho'_{S\mathcal{M}}] \leq \text{rank}[\varrho'_S] \text{rank}[\varrho'_{\mathcal{M}}], \quad (\text{B1})$$

where the inequality follows from the subadditivity of the Rényi-zero entropy [48], which is the logarithm of the rank. To achieve perfect cooling of the target, one must (at least asymptotically) attain $\text{rank}[\varrho'_S] < \text{rank}[\tau_S(\beta, H_S)]$, which implies that $\text{rank}[\varrho'_{\mathcal{M}}] > \text{rank}[\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})]$. However, if this condition is achieved, then $D[\varrho'_{\mathcal{M}} \| \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})]$ diverges, implying a diverging energy cost by Eq. (3). The above argument already appears in Ref. [29].

The other situation that one must consider is the case where one attains a ϱ'_S such that $\text{rank}[\varrho'_S] = \text{rank}[\tau_S(\beta, H_S)]$ but nonetheless ϱ'_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 ϱ_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 D). Since the relative entropy $(\varrho, \sigma) \mapsto D(\varrho \| \sigma)$ is lower semicontinuous [49, 50] and since $D(\varrho'_{\mathcal{M}} \| \varrho_{\mathcal{M}}) \rightarrow \infty$ by the arguments above, we infer that $D[(\varrho'_{\mathcal{M}})^{(i)} \| \varrho_{\mathcal{M}}^{(i)}] \rightarrow \infty$ as $i \rightarrow \infty$.

This argument holds independently of $\text{rank}[\rho'_S]$; in particular, for the special case $\text{rank}[\rho'_S] = \text{rank}[\tau_S(\beta, H_S)]$ 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 $\rho_{\mathcal{M}}$ and ρ'_S 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.

Appendix C: Diverging Time Cooling Protocol for Finite-Dimensional Systems

C1. Proof of Theorem 2

Proof. Consider a target system S of dimension d with associated Hamiltonian

$$H_S = \sum_{k=0}^{d-1} \omega_k |k\rangle\langle k|_S, \quad (\text{C1})$$

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

$$H_{\mathcal{M}}^{(n)} = (1 + n\epsilon)H_S, \quad (\text{C2})$$

where $\epsilon = (\beta_{\max} - \beta)/(N\beta)$. We first cool the system initially at non-zero β to some fixed, finite $\beta_{\max}$, 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 $\tau_S(\beta_{\max}, H_S)$. 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 $\tilde{\Delta}S := S[\tau_S(\beta, H_S)] - S[\tau_S(\beta_{\max}, H_S)]$. 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

$$\mathbb{S}_{S\mathcal{M}_n}^d := \sum_{i,j=0}^{d-1} |i, j\rangle\langle j, i|_{S\mathcal{M}_n}. \quad (\text{C3})$$

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

$$\Delta E_S^{(n)} = \text{tr} \left[H_S \tau(\beta, H_{\mathcal{M}}^{(n)}) \right] - \text{tr} \left[H_S \tau(\beta, H_{\mathcal{M}}^{(n-1)}) \right], \quad (\text{C4})$$

and so the total energy change of the system over the entire N -step protocol is given by

$$\Delta E_S = \sum_{n=1}^N \Delta E_S^{(n)} = \text{tr} \left[H_S \tau(\beta, H_{\mathcal{M}}^{(N)}) \right] - \text{tr} \left[H_S \tau(\beta, H_{\mathcal{M}}^{(0)}) \right]. \quad (\text{C5})$$

The energy change of the machine subsystem that is swapped with the target system at each step is given by

$$\Delta E_{\mathcal{M}}^{(n)} = \text{tr} \left[H_{\mathcal{M}}^{(n)} \tau(\beta, H_{\mathcal{M}}^{(n-1)}) \right] - \text{tr} \left[H_{\mathcal{M}}^{(n)} \tau(\beta, H_{\mathcal{M}}^{(n)}) \right] = \sum_{k=0}^{d-1} (1 + n\epsilon)\omega_k \left[p_k(\beta, H_{\mathcal{M}}^{(n-1)}) - p_k(\beta, H_{\mathcal{M}}^{(n)}) \right], \quad (\text{C6})$$

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 $\mathcal{Z}_{\mathcal{M}_n}(\beta, H_{\mathcal{M}}^{(n)}) = \text{tr} [e^{-\beta H_{\mathcal{M}}^{(n)}}]$ 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:

$$\Delta E_{\mathcal{M}}^{(N)} = \sum_{n=1}^N \Delta E_{\mathcal{M}}^{(n)} = \sum_{n=1}^N \sum_{k=0}^{d-1} (1+n\epsilon)\omega_k [p_k(\beta, H_{\mathcal{M}}^{(n-1)}) - p_k(\beta, H_{\mathcal{M}}^{(n)})], \quad (\text{C7})$$

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

$$\Delta p_k^{(n)} := p_k(\beta, H_{\mathcal{M}}^{(n)}) - p_k(\beta, H_{\mathcal{M}}^{(n-1)}), \quad (\text{C8})$$

and

$$\Delta x_n := x_n - x_{n-1} \quad \text{with} \quad x_n := 1 + n\epsilon. \quad (\text{C9})$$

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

$$-\frac{\Delta p_k^{(n)}}{\Delta x_n} \Delta x_n \rightarrow -\frac{\partial p_k(\beta, xH_s)}{\partial x} dx \quad (\text{C10})$$

where $x := 1 + n\epsilon$ has become a continuous parameter. This way we can express the limit $N \rightarrow \infty$ of Eq. (C7) as a Riemann integral in the following form

$$\lim_{N \rightarrow \infty} \Delta E_{\mathcal{M}}^{(N)} = - \int_1^{x_{\text{max}}} \sum_{k=0}^{d-1} x\omega_k \frac{\partial p_k(\beta, xH_s)}{\partial x} dx, \quad (\text{C11})$$

where $x_{\text{max}} := \beta_{\text{max}}/\beta$. Both the summation and the integral converge, so one can swap the order of their evaluation. Integrating by parts then gives

$$\begin{aligned} \lim_{N \rightarrow \infty} \Delta E_{\mathcal{M}}^{(N)} &= \sum_{k=0}^{d-1} \left[-x\omega_k p_k(\beta, xH_s) \Big|_1^{x_{\text{max}}} + \int_1^{x_{\text{max}}} \omega_k p_k(\beta, xH_s) dx \right] \\ &= \sum_{k=0}^{d-1} \left[-x\omega_k p_k(\beta, xH_s) \Big|_1^{x_{\text{max}}} \right] - \int_1^{x_{\text{max}}} \frac{1}{\beta} \frac{\partial}{\partial x} [\log \mathcal{Z}(\beta, xH_s)] dx \\ &= E[\tau(\beta, H_s)] - E[\tau(\beta, x_{\text{max}}H_s)] - \frac{1}{\beta} \log \mathcal{Z}(\beta, x_{\text{max}}H_s) + \frac{1}{\beta} \log \mathcal{Z}(\beta, H_s), \end{aligned} \quad (\text{C12})$$

where in the second line we again swapped the order of the integral and the sum to write $\sum_{k=0}^{d-1} \omega_k p_k(\beta, xH_s) = -\frac{1}{\beta} \frac{\partial}{\partial x} [\log \mathcal{Z}(\beta, xH_s)]$ and in the last line we invoke $E[\tau(\beta, xH)] = \text{tr} [xH \tau(\beta, xH)]$. Finally, writing the partition function in terms of the average energy and entropy, i.e., $\log [\mathcal{Z}(\beta, xH)] = -\beta E[\tau(\beta, xH)] + S[\tau(\beta, xH)]$, the total energy change of the machine is given by

$$\begin{aligned} \lim_{N \rightarrow \infty} \Delta E_{\mathcal{M}}^{(N)} &= E[\tau(\beta, H_s)] - E[\tau(\beta, x_{\text{max}}H_s)] + E[\tau(\beta, x_{\text{max}}H_s)] - \frac{1}{\beta} S[\tau(\beta, x_{\text{max}}H)] - E[\tau(\beta, H_s)] + \frac{1}{\beta} S[\tau(\beta, H_s)] \\ &= \frac{1}{\beta} \{ S[\tau(\beta, H_s)] - S[\tau(\beta_{\text{max}}, H_s)] \} = \frac{1}{\beta} \tilde{\Delta} S_s, \end{aligned} \quad (\text{C13})$$

where we have made use of the property $\tau_s(\beta, x_{\text{max}}H_s) = \tau_s(\beta_{\text{max}}, H_s)$ and the entropy decrease of the target system corresponds to that associated with the transformation $\tau(\beta, H_s) \rightarrow \tau(\beta_{\text{max}}, H_s)$. 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 2, we can now take the limit $\beta_{\text{max}} \rightarrow \infty$. $\square$

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.

C2. 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

$$\tau_S(\beta, \omega_\alpha) := \frac{e^{-\beta H_S(\omega_\alpha)}}{\mathcal{Z}_S(\beta, \omega_\alpha)}, \quad (\text{C14})$$

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

$$\begin{aligned} \Delta E_{\mathcal{M}_\alpha}(\omega_\alpha) &= -\left\{ \text{tr}[H_{\mathcal{M}_\alpha}(\omega_\alpha) \tau_{\mathcal{M}_\alpha}(\beta, \omega_\alpha)] + \text{tr}[H_{\mathcal{M}_\alpha}(\omega_\alpha) \tau_{\mathcal{M}_\alpha}(\beta, \omega_{\alpha-1})] \right\} \\ &= -\text{tr}[H_S(\omega_\alpha) [\tau_S(\beta, \omega_\alpha) - \tau_S(\beta, \omega_{\alpha-1})]]. \end{aligned} \quad (\text{C15})$$

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 S , which is characterised via its local Hamiltonian H_S and calculated via

$$\Delta E_S(\omega_\alpha) = \text{tr}[H_S(\omega_S) \tau_S(\beta, \omega_\alpha)] - \text{tr}[H_S(\omega_S) \tau_S(\beta, \omega_{\alpha-1})], \quad (\text{C16})$$

in which we set $\omega_0 = \omega_S$. Using Eqs. (C15, C16), the total energy cost for each stage of cooling is given by

$$\Delta E_{S\mathcal{M}}(\omega_\alpha) = \Delta E_S(\omega_\alpha) + \Delta E_{\mathcal{M}}(\omega_\alpha) = \text{tr}\left\{ [H_S(\omega_S) - H_S(\omega_\alpha)] [\tau_S(\beta, \omega_\alpha) - \tau_S(\beta, \omega_{\alpha-1})] \right\}, \quad (\text{C17})$$

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

$$\Delta E_{S\mathcal{M}}^{(N)} = \sum_{\alpha=1}^N \Delta E_{S\mathcal{M}}(\omega_\alpha) = \sum_{\alpha=1}^N \text{tr}\left\{ [H_S(\omega_S) - H_S(\omega_\alpha)] [\tau_S(\beta, \omega_\alpha) - \tau_S(\beta, \omega_{\alpha-1})] \right\}. \quad (\text{C18})$$

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. (C18), we have

$$\begin{aligned}
\Delta E_{S\mathcal{M}}^{(N)} &= \sum_{\alpha=1}^N \text{tr} \left\{ [H_S(\omega_S) - H_S(\omega_\alpha)] [\tau_S(\beta, \omega_\alpha) - \tau_S(\beta, \omega_{\alpha-1})] \right\} \\
&= \sum_{\alpha=1}^N (\omega_S - \omega_\alpha) \left[\left(\frac{e^{-\beta\omega_\alpha}}{1 - e^{-\beta\omega_\alpha}} - \frac{e^{-\beta\omega_{\alpha-1}}}{1 - e^{-\beta\omega_{\alpha-1}}} \right) - \left(\frac{d e^{-\beta d\omega_\alpha}}{1 - e^{-\beta d\omega_\alpha}} - \frac{d e^{-\beta d\omega_{\alpha-1}}}{1 - e^{-\beta d\omega_{\alpha-1}}} \right) \right] \\
&= \lim_{K \rightarrow \infty} \sum_{\alpha=1}^N (\omega_S - \omega_\alpha) \sum_{k=0}^K [(e^{-\beta(k+1)\omega_\alpha} - e^{-\beta(k+1)\omega_{\alpha-1}}) - d(e^{-\beta(k+1)d\omega_\alpha} - e^{-\beta(k+1)d\omega_{\alpha-1}})] \\
&= \lim_{K \rightarrow \infty} \sum_{\alpha=1}^N (\omega_S - \omega_\alpha) \sum_{k=0}^K [e^{-\beta(k+1)\omega_\alpha} (1 - e^{-\beta(k+1)(\omega_{\alpha-1} - \omega_\alpha)}) - d e^{-\beta d(k+1)\omega_\alpha} (1 - e^{-\beta d(k+1)(\omega_{\alpha-1} - \omega_\alpha)})].
\end{aligned} \tag{C19}$$

Here, since both $H_{\mathcal{M}_\alpha}$ and H_S are equally spaced Hamiltonians, the average energy can be written as

$$E(\omega_x, \omega_y) = \text{tr} [H_S(\omega_x) \tau_S(\beta, \omega_y)] = \frac{\sum_{n=0}^{d-1} n \omega_x e^{-n\beta\omega_y}}{\sum_{n=0}^{d-1} e^{-n\beta\omega_y}} = \omega_x \left(\frac{e^{-\beta\omega_y}}{1 - e^{-\beta\omega_y}} - \frac{d e^{-\beta d\omega_y}}{1 - e^{-\beta d\omega_y}} \right) \tag{C20}$$

by evaluating the geometric series

$$\mathcal{Z}(\beta, \omega_y) = \sum_{n=0}^{d-1} e^{-\beta n \omega_y} = \frac{1 - e^{-\beta d \omega_y}}{1 - e^{-\beta \omega_y}} \tag{C21}$$

and writing

$$E(\omega_x, \omega_y) = \sum_{n=0}^{d-1} n \omega_x \frac{e^{-\beta n \omega_y}}{\mathcal{Z}(\beta, \omega_y)} = \frac{\omega_x}{\omega_y} \left\{ -\frac{\partial}{\partial \beta} \log [\mathcal{Z}(\beta, \omega_y)] \right\} = -\frac{\omega_x}{\omega_y} \frac{\partial}{\partial \beta} [\log (1 - e^{-\beta d \omega_y}) - \log (1 - e^{-\beta \omega_y})] \tag{C22}$$

as we do in the second line of Eq. (C19) 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 E2a, 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. (C19). 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:

$$\begin{aligned}
\Delta E_{S\mathcal{M}}^{(N)} &= - \lim_{K \rightarrow \infty} \sum_{\alpha=1}^N \alpha \epsilon \sum_{k=0}^K e^{-\beta k(\omega_S + \alpha \epsilon)} (1 - e^{\beta k \epsilon}) + \lim_{K \rightarrow \infty} \sum_{\alpha=1}^N \alpha d \epsilon \sum_{k=0}^K e^{-\beta k d(\omega_S + \alpha \epsilon)} (1 - e^{\beta k d \epsilon}) \\
&= \lim_{K \rightarrow \infty} \sum_{k=0}^K [e^{-\beta k \omega_S} (e^{\beta k \epsilon} - 1) (\sum_{\alpha=1}^N \alpha \epsilon e^{-\beta k \alpha \epsilon})] - \lim_{K \rightarrow \infty} \sum_{k=0}^K e^{-\beta k d \omega_S} [(e^{\beta k d \epsilon} - 1) (\sum_{\alpha=1}^N d \alpha \epsilon e^{-\beta k d \alpha \epsilon})], \tag{C23}
\end{aligned}$$

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 .

$$\begin{aligned}
\sum_{\alpha=1}^N \alpha \epsilon e^{-\beta \alpha \epsilon} &= -\frac{\partial}{\partial \beta} \sum_{\alpha=0}^N e^{-\beta \alpha \epsilon} = -\frac{\partial}{\partial \beta} \left(\frac{1 - e^{-\beta(N+1)\epsilon}}{1 - e^{-\beta \epsilon}} \right) \\
&= - \left(\frac{(N+1)\epsilon e^{-\beta(N+1)\epsilon} - (N+1)\epsilon e^{-\beta(N+2)\epsilon} - \epsilon e^{-\beta \epsilon} + \epsilon e^{-\beta(N+2)\epsilon}}{(1 - e^{-\beta \epsilon})^2} \right) \\
&= \frac{\epsilon e^{-\beta \epsilon}}{(1 - e^{-\beta \epsilon})^2} (1 - (N+1)e^{-\beta N \epsilon} + N e^{-\beta(N+1)\epsilon}) \\
&= \frac{\epsilon e^{-\beta \epsilon}}{(1 - e^{-\beta \epsilon})^2} (1 - e^{-\beta N \epsilon} - N e^{-\beta N \epsilon} (1 - e^{-\beta \epsilon})). \tag{C24}
\end{aligned}$$

Combining Eqs. (C23) and (C24), we arrive at

$$\begin{aligned}\Delta E_{s\mathcal{M}}^{(N)} &= \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k \omega_S}}{k} \frac{k \epsilon (1 - e^{-\beta N k \epsilon})}{(1 - e^{-\beta k \epsilon})} - N \epsilon e^{-\beta k (\omega_S + N \epsilon)} \right] \\ &\quad - \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k d \omega_S}}{k} \frac{k d \epsilon (1 - e^{-\beta N k d \epsilon})}{(1 - e^{-\beta k d \epsilon})} - N d \epsilon e^{-\beta k d (\omega_S + N \epsilon)} \right].\end{aligned}\quad (\text{C25})$$

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 $\epsilon = \Delta\omega/N := (\omega_{\max} - \omega_S)/N$. Substituting this expression for ϵ into the above equation yields

$$\begin{aligned}\Delta E_{s\mathcal{M}}^{(N)} &= \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k \omega_S}}{k} \frac{k \Delta\omega (1 - e^{-\beta k \Delta\omega})}{N (1 - e^{-\beta k \frac{\Delta\omega}{N}})} - \Delta\omega e^{-\beta k (\omega_S + \Delta\omega)} \right] \\ &\quad - \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k d \omega_S}}{k} \frac{k d \Delta\omega (1 - e^{-\beta k d \Delta\omega})}{N (1 - e^{-\beta k d \frac{\Delta\omega}{N}})} - d \Delta\omega e^{-\beta k d (\omega_S + \Delta\omega)} \right].\end{aligned}\quad (\text{C26})$$

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

$$\begin{aligned}\Delta E_{s\mathcal{M}}^{\text{tot}} &= \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k \omega_S}}{\beta k} - \frac{e^{-\beta k \omega_{\max}}}{\beta k} - (\omega_{\max} - \omega_S) e^{-\beta k \omega_{\max}} \right] \\ &\quad - \lim_{K \rightarrow \infty} \sum_{k=0}^K \left[\frac{e^{-\beta k d \omega_S}}{\beta k} - \frac{e^{-\beta k d \omega_{\max}}}{\beta k} - d (\omega_{\max} - \omega_S) e^{-\beta k d \omega_{\max}} \right].\end{aligned}\quad (\text{C27})$$

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 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:

$$\begin{aligned}\Delta E_{s\mathcal{M}}^{\text{tot}} &= -\frac{1}{\beta} \log(1 - e^{-\beta \omega_S}) + \frac{1}{\beta} \log(1 - e^{-\beta \omega_{\max}}) - \frac{(\omega_{\max} - \omega_S) e^{-\beta \omega_{\max}}}{1 - e^{-\beta \omega_{\max}}} \\ &\quad + \frac{1}{\beta} \log(1 - e^{-\beta d \omega_S}) - \frac{1}{\beta} \log(1 - e^{-\beta d \omega_{\max}}) + \frac{d(\omega_{\max} - \omega_S) e^{-\beta d \omega_{\max}}}{1 - e^{-\beta d \omega_{\max}}} \\ &= \frac{1}{\beta} \log \left(\frac{1 - e^{-\beta d \omega_S}}{1 - e^{-\beta \omega_S}} \right) - \frac{1}{\beta} \log \left(\frac{1 - e^{-\beta d \omega_{\max}}}{1 - e^{-\beta \omega_{\max}}} \right) - (\omega_{\max} - \omega_S) \left(\frac{e^{-\beta \omega_{\max}}}{1 - e^{-\beta \omega_{\max}}} - \frac{d e^{-\beta d \omega_{\max}}}{1 - e^{-\beta d \omega_{\max}}} \right) \\ &= \frac{1}{\beta} \log[\mathcal{Z}_S(\beta, \omega_S)] - \frac{1}{\beta} \log[\mathcal{Z}_S(\beta, \omega_{\max})] - \text{tr}[H_S(\omega_{\max}) \tau_S(\beta, \omega_{\max})] + \text{tr}[H_S(\omega_S) \tau_S(\beta, \omega_{\max})] \\ &= \frac{1}{\beta} \log[\mathcal{Z}_S(\beta, \omega_S)] - \frac{1}{\beta} \log[\mathcal{Z}_S(\beta, \omega_{\max})] \\ &\quad - \text{tr}[H_S(\omega_{\max}) \tau_S(\beta, \omega_{\max})] + \text{tr}[H_S(\omega_S) \tau_S(\beta, \omega_S)] - \text{tr}[H_S(\omega_S) \tau_S(\beta, \omega_S)] + \text{tr}[H_S(\omega_S) \tau_S(\beta, \omega_{\max})] \\ &= \frac{1}{\beta} \Delta S_S + \Delta E_S,\end{aligned}\quad (\text{C28})$$

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.

Appendix D: 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.

D1. Necessary Complexity Conditions

D1a. 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 $\lambda_{\min}$ of the final state ϱ'_S (and hence how cold the system can become) after *any* unitary interaction with a thermal machine by [29]

$$\lambda_{\min}(\varrho'_S) \geq e^{-\beta\omega_{\mathcal{M}}^{\max}} \lambda_{\min}(\varrho_S), \quad (\text{D1})$$

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. [42, 51]). 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. [29], 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 10. *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 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?*

D1b. 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. 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 $\mathcal{H}_{S\mathcal{M}}$ in terms of which the unitary can be decomposed as $U_{S\mathcal{M}} = U_{\mathcal{A}} \oplus \mathbb{1}_{\mathcal{A}^\perp}$, i.e.,

$$d^{\text{eff}} := \min \dim(\mathcal{A}) : U_{S\mathcal{M}} = U_{\mathcal{A}} \oplus \mathbb{1}_{\mathcal{A}^\perp}. \quad (\text{D2})$$

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 11), 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 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. [29], there are finite-dimensional correction terms that then imply that the Landauer limit cannot be saturated.

Note finally that in Theorem 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. [52, 53]). 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 $\omega_{\mathcal{M}_n} \in \{\omega_S + n\epsilon\}_{n=1,\dots,N}$, 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 .

D2. Proof of Theorem 3, Corollary 2, and Theorem 10

Here we prove Theorem 3, which implies Theorem 10 and leads to Corollary 2.

Proof. Let $\mathcal{H}_{\mathcal{X}}$ be a separable Hilbert space associated with the system $\mathcal{X}$. Consider

$$H_{\mathcal{M}} = \sum_{n=0}^{\infty} \omega_n |n\rangle\langle n| \quad \text{and} \quad \mathcal{H}_{\mathcal{M}'} = \text{span}_{n \leq m} \{|n\rangle\}, \quad (\text{D3})$$

for some finite m . In other words, $\mathcal{H}_{\mathcal{M}'}$ is a finite-dimensional restriction of $\mathcal{H}_{\mathcal{M}}$. We will show that any unitary that (non-trivially) 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 $\mathcal{H}_{\mathcal{S}}$ with $\mathcal{H}_{\mathcal{M}'}$; whenever we talk of an operation with finite effective dimension in this article, we mean specifically such a U , and by diverging effective dimension we mean a unitary that couples the target to any subspace of $\mathcal{H}_{\mathcal{M}}$ whose dimension diverges. Since

$$\mathcal{H}_{\mathcal{S}} \otimes \mathcal{H}_{\mathcal{M}} = \mathcal{H}_{\mathcal{S}} \otimes (\mathcal{H}_{\mathcal{M}'} \oplus \mathcal{H}_{\mathcal{M}'}^{\perp}) \simeq (\mathcal{H}_{\mathcal{S}} \otimes \mathcal{H}_{\mathcal{M}'}) \oplus (\mathcal{H}_{\mathcal{S}} \otimes \mathcal{H}_{\mathcal{M}'}^{\perp}), \quad (\text{D4})$$

we can associate the subspace $\mathcal{H}_{\mathcal{S}} \otimes \mathcal{H}_{\mathcal{M}'}$ with the label $\mathcal{A}$ and $\mathcal{H}_{\mathcal{S}} \otimes \mathcal{H}_{\mathcal{M}'}^{\perp}$ with $\mathcal{B}$ and write $U = U_{\mathcal{A}} \oplus \mathbb{1}_{\mathcal{B}}$. Then the initial configuration can be expressed as:

$$\varrho_{\mathcal{S}} \otimes \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}}) = \begin{bmatrix} \varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'} & 0 \\ 0 & \varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'}^{\perp} \end{bmatrix}, \quad (\text{D5})$$

where

$$\varrho_{\mathcal{M}'} := \frac{1}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \sum_{n \leq m} e^{-\beta \omega_n} |n\rangle\langle n| \quad \text{and} \quad \varrho_{\mathcal{M}'}^{\perp} := \frac{1}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \sum_{n > m} e^{-\beta \omega_n} |n\rangle\langle n| \quad (\text{D6})$$

add up to a (normalised) thermal state. Now consider the state

$$\tilde{\varrho}_{\mathcal{M}} = \begin{bmatrix} \varrho_{\mathcal{M}'} & 0 \\ 0 & \text{tr}[\varrho_{\mathcal{M}'}^{\perp}] \end{bmatrix}. \quad (\text{D7})$$

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

$$\tilde{H}_{\mathcal{M}} = \sum_{n \leq m} \omega_n |n\rangle\langle n| - \frac{1}{\beta} \log \left(\sum_{n > m} e^{-\beta \omega_n} \right) |m+1\rangle\langle m+1|. \quad (\text{D8})$$

To see this, note that $\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}}) = \mathcal{Z}_{\mathcal{M}}(\beta, \tilde{H}_{\mathcal{M}})$ and that

$$\exp \left\{ -\beta \left[-\frac{1}{\beta} \log \left(\sum_{n > m} e^{-\beta \omega_n} \right) \right] \right\} = \sum_{n > m} e^{-\beta \omega_n}. \quad (\text{D9})$$

Thus $\tilde{\varrho}_{\mathcal{M}} = \tau_{\mathcal{M}}(\beta, \tilde{H}_{\mathcal{M}})$. 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

$$\tilde{U} (\varrho_{\mathcal{S}} \otimes \tilde{\varrho}_{\mathcal{M}}) \tilde{U}^{\dagger} = \begin{bmatrix} U_{\mathcal{A}} (\varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'}) U_{\mathcal{A}}^{\dagger} & 0 \\ 0 & \frac{e^{-\beta \tilde{\omega}_{m+1}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \varrho_{\mathcal{S}} \end{bmatrix} \quad (\text{D10})$$

and

$$\text{tr}_{\mathcal{M}} [\tilde{U} (\varrho_{\mathcal{S}} \otimes \tilde{\varrho}_{\mathcal{M}}) \tilde{U}^{\dagger}] = \text{tr}_{\mathcal{M}'} [U_{\mathcal{A}} (\varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'}) U_{\mathcal{A}}^{\dagger}] + \frac{e^{-\beta \tilde{\omega}_{m+1}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \varrho_{\mathcal{S}} \quad (\text{D11})$$

Compare this to the expression

$$\text{tr}_{\mathcal{M}} [U (\varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}}) U^{\dagger}] = \text{tr}_{\mathcal{M}'} \begin{bmatrix} U_{\mathcal{A}} (\varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'}) U_{\mathcal{A}}^{\dagger} & 0 \\ 0 & \varrho_{\mathcal{S}} \otimes \varrho_{\mathcal{M}'}^{\perp} \end{bmatrix}$$

$$\begin{aligned}
&= \text{tr}_{\mathcal{M}'} [U_{\mathcal{A}}(\varrho_S \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^\dagger] + \text{tr} [\varrho_{\mathcal{M}'}^\perp] \varrho_S \\
&= \text{tr}_{\mathcal{M}'} [U_{\mathcal{A}}(\varrho_S \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^\dagger] + \frac{e^{-\beta\tilde{\omega}_{m+1}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \varrho_S,
\end{aligned} \tag{D12}$$

since $\text{tr} [\varrho_{\mathcal{M}'}^\perp] = \frac{1}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} \sum_{n>m} e^{-\beta\omega_n}$. 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

$$\begin{aligned}
\tilde{\varrho}'_{\mathcal{M}} &= \text{tr}_S [\tilde{U}(\varrho_S \otimes \tilde{\varrho}_{\mathcal{M}})\tilde{U}^\dagger] = \text{tr}_S [U_{\mathcal{A}}(\varrho_S \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^\dagger] + \frac{e^{-\beta\tilde{\omega}_{m+1}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} |m+1\rangle\langle m+1| \\
\tilde{\varrho}_{\mathcal{M}} &= \varrho_{\mathcal{M}'} + \frac{e^{-\beta\tilde{\omega}_{m+1}}}{\mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})} |m+1\rangle\langle m+1|.
\end{aligned} \tag{D13}$$

Thus, we have

$$\text{tr} [\tilde{H}_{\mathcal{M}}(\tilde{\varrho}'_{\mathcal{M}} - \tilde{\varrho}_{\mathcal{M}})] = \text{tr} \{H_{\mathcal{M}} [\text{tr}_S [U_{\mathcal{A}}(\varrho_S \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^\dagger] - \varrho_{\mathcal{M}'}]\}, \tag{D14}$$

since $U_{\mathcal{A}}$ only acts on $\mathcal{H}_S \otimes \mathcal{H}_{\mathcal{M}'}$ and $\tilde{H}_{\mathcal{M}|\mathcal{M}'} = H_{\mathcal{M}|\mathcal{M}'}$. In the same way, we have

$$\begin{aligned}
\text{tr}_S [U(\varrho_S \otimes \varrho_{\mathcal{M}})U^\dagger] &= \text{tr}_S [U_{\mathcal{A}}(\varrho_S \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^\dagger] + \varrho_{\mathcal{M}'}^\perp \\
\varrho_{\mathcal{M}} &= \varrho_{\mathcal{M}'} + \varrho_{\mathcal{M}'}^\perp.
\end{aligned} \tag{D15}$$

Thus, the energy difference is also

$$\text{tr} \{H_{\mathcal{M}} [\text{tr}_S [U_{\mathcal{A}}(\varrho_S \otimes \varrho_{\mathcal{M}'})U_{\mathcal{A}}^\dagger] - \varrho_{\mathcal{M}'}]\}. \tag{D16}$$

Hence, we have shown that we can replace (a potentially infinite-dimensional) $\mathcal{M}$ by some (finite) $m+1$ dimensional machine $\tilde{\mathcal{M}}$ if the joint unitary U only acts on m levels of $H_{\mathcal{M}}$. By Theorem 6 of Ref. [29], 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 3, which implies Theorem 10. $\square$

D3. 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 C 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. (C2), i.e.,

$$H_{\mathcal{M}}^{\text{tot}} = \sum_{n=1}^N H_{\mathcal{M}_n}^{(n)} = \sum_n (1 + n\epsilon) H_S^{(n)}, \tag{D17}$$

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 10, 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

$$\begin{aligned} U_{\mathcal{SM}} &= \Pi_{n=1}^N \mathbb{S}_{\mathcal{SM}_n}^2 = \Pi_n \left(\sum_{i,j_n=0}^1 |i, j_1, \dots, j_n, \dots, j_N\rangle \langle j_n, j_1, \dots, i, \dots, j_N|_{\mathcal{SM}} \right) \\ &= \sum_{i,j_1 \dots j_N=0}^1 |i, j_1, \dots, j_N\rangle \langle j_N, i, j_1, \dots, j_{N-1}|_{\mathcal{SM}}. \end{aligned} \quad (\text{D18})$$

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 $|i = 1, j_1 = 1, \dots, j_N = 1\rangle$.

D4. Fine-Tuned Control Conditions

Theorem 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 10 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 diverging effective dimensionality. In particular, we note that 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. In this section, we begin by outlining the general conditions that perfect cooling at the Landauer limit entails, before presenting a more nuanced notion of control complexity in terms of the variety of distinct energy gaps in the machine in Appendix D5.

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 both 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. In Appendix D5, we demonstrate that indeed this condition is necessary for minimal-energy cost cooling.

Before doing so, we here first derive 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 similarly derived.

Consider an arbitrary thermal target system and machine of finite dimensions, with respective spectra $\lambda_S := \{\lambda_S^0, \dots, \lambda_S^{d_S-1}\}$ and $\lambda_M := \{\lambda_M^0, \dots, \lambda_M^{d_M-1}\}$. The states begin uncorrelated, so the global spectrum of the initial joint state is $\lambda_{SM} := \{\lambda_{SM}^0, \dots, \lambda_{SM}^{d_S d_M-1}\} = \{\lambda_S^0 \lambda_M^0, \lambda_S^0 \lambda_M^1, \dots, \lambda_S^{d_S-1} \lambda_M^{d_M-1}\}$. 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 λ'_{SM} and λ_{SM} .

The transformation that cools the target system as much as possible⁵ 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 $\lambda^\downarrow$ the non-increasing ordering of the set λ . Since the target and machine begin thermal, the local spectra λ_S and λ_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

$$\varrho'_S = \sum_{i=0}^{d_S-1} \left(\sum_{j=0}^{d_M-1} \lambda_{SM}^{\downarrow i d_M + j} \right) |i\rangle\langle i|_S. \quad (\text{D19})$$

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 of the protocol. 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 render the machine energetically passive, leading to the form:

$$\varrho'_{SM} = \sum_{i=0}^{d_S-1} \sum_{j=0}^{d_M-1} \lambda_{SM}^{\downarrow i d_M + j} |ij\rangle\langle ij|_{SM}. \quad (\text{D20})$$

We can derive the above form of the final joint state as follows. Consider the following ordering for the energy eigenbasis of SM chosen to match the above form

$$\{|00\rangle_{SM}, |01\rangle_{SM}, \dots, |0, d_M - 1\rangle_{SM}, |10\rangle_{SM}, \dots, |1, d_M - 1\rangle_{SM}, \dots, |d_S - 1, 0\rangle_{SM}, \dots, |d_S - 1, d_M - 1\rangle_{SM}\}. \quad (\text{D21})$$

This ordering is monotonically non-decreasing primarily with respect to the energy of S , and secondarily w.r.t. M . We take the final state ρ'_{SM} to be expressed in this basis. To maximise the cooling in a single unitary operation, we maximise the sum of the first $k \cdot d_M$ diagonal elements, for each $k \in \{1, 2, \dots, d_S\}$, as each sum corresponds to the total population in the k^{th} lowest energy eigenstate of S . The initial state ϱ_{SM} is diagonal in this basis, so the vector of initial diagonal elements, which we label $\theta := \text{diag}(\varrho_{SM})$, is also the vector of eigenvalues, λ_{SM} , i.e., $\theta = \lambda_{SM}$. Furthermore, since the unitary operation leaves the set of eigenvalues invariant, we have via the Schur-Horn lemma [54] that the vector of final diagonal elements, which we label $\theta' := \text{diag}(\varrho'_{SM})$, is majorised by the vector of initial ones, i.e., $\theta' \prec \theta$. It follows that the partial sums we wish to maximise are upper bounded by the corresponding partial sums of the $k \cdot d_M$ 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. (D21).

More precisely, via the Schur-Horn lemma, one can always write $\theta' = D\theta$, with D a doubly stochastic matrix. The partial sums of the $k \cdot d_M$ first elements are linear functions of the elements of θ . Thus the maximum values are obtained at the extremal

⁵ We 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.

points of the convex set of doubly stochastic matrices, which are the permutation matrices, via the Birkhoff-von Neumann theorem [54]. 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 unique⁶ stochastic matrix D that 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 energetically 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

$$\varrho'_{\mathcal{M}} = \sum_{j=0}^{d_{\mathcal{M}}-1} \left(\sum_{i=0}^{d_{\mathcal{S}}-1} \lambda_{\mathcal{S}\mathcal{M}}^{\downarrow i d_{\mathcal{M}} + j} \right) |j\rangle\langle j|_{\mathcal{M}}. \quad (\text{D22})$$

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 11. *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*

$$\sum_{i=0}^{\frac{d_{\mathcal{M}}}{2}-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 1, \quad \sum_{i=\frac{d_{\mathcal{M}}}{2}}^{d_{\mathcal{M}}-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 0, \quad (\text{D23})$$

and

$$\frac{\frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2} + \lfloor \frac{i}{2} \rfloor} \right)}{\lambda_{\mathcal{M}}^{\downarrow i}} \rightarrow 1 \quad (\text{D24})$$

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 $d_{\mathcal{M}} \rightarrow \infty$ ⁷.

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

$$\lambda_{\mathcal{S}\mathcal{M}} = \frac{1}{2} \{ \lambda_{\mathcal{M}}^{\downarrow}, \lambda_{\mathcal{M}}^{\downarrow} \} = \frac{1}{2} \{ \lambda_{\mathcal{M}}^{\downarrow 0}, \lambda_{\mathcal{M}}^{\downarrow 1}, \dots, \lambda_{\mathcal{M}}^{\downarrow d_{\mathcal{M}}-1}, \lambda_{\mathcal{M}}^{\downarrow 0}, \lambda_{\mathcal{M}}^{\downarrow 1}, \dots, \lambda_{\mathcal{M}}^{\downarrow d_{\mathcal{M}}-1} \}. \quad (\text{D25})$$

⁶ Note that degeneracies in energy eigenvalues would lead to sets of equal diagonal elements, and prevent one 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 av-

erage energy change.

⁷ Strictly speaking, in the limit $d_{\mathcal{M}} \rightarrow \infty$ the conditions in Eq. (D24) 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.

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

$$\lambda_{s\mathcal{M}}^{\downarrow} = \frac{1}{2} \{ \lambda_{\mathcal{M}}^{\downarrow 0}, \lambda_{\mathcal{M}}^{\downarrow 0}, \lambda_{\mathcal{M}}^{\downarrow 1}, \lambda_{\mathcal{M}}^{\downarrow 1}, \dots, \lambda_{\mathcal{M}}^{\downarrow d_{\mathcal{M}}-1}, \lambda_{\mathcal{M}}^{\downarrow d_{\mathcal{M}}-1} \}. \quad (\text{D26})$$

Expressing the global states with respect to the product of local energy eigenbases, we have that the initial joint state is $\frac{1}{2} \otimes \tau_{\mathcal{M}}(\beta, H_{\mathcal{M}}) = \text{diag}(\lambda_{s\mathcal{M}}^{\downarrow})$ [see Eq. (D25)] 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 $\varrho'_{s\mathcal{M}} = \text{diag}(\lambda_{s\mathcal{M}}^{\downarrow})$. This leads to the following reduced states

$$\varrho'_s = \left(\sum_{i=0}^{\frac{d_{\mathcal{M}}}{2}-1} \lambda_{\mathcal{M}}^{\downarrow i} \right) |0\rangle\langle 0|_s + \left(\sum_{i=\frac{d_{\mathcal{M}}}{2}}^{d_{\mathcal{M}}-1} \lambda_{\mathcal{M}}^{\downarrow i} \right) |1\rangle\langle 1|_s, \quad (\text{D27})$$

$$\begin{aligned} \varrho'_{\mathcal{M}} &= \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow 0} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2}} \right) |0\rangle\langle 0|_{\mathcal{M}} + \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow 0} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2}} \right) |1\rangle\langle 1|_{\mathcal{M}} + \\ &+ \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow 1} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2}+1} \right) |2\rangle\langle 2|_{\mathcal{M}} + \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow 1} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2}+1} \right) |3\rangle\langle 3|_{\mathcal{M}} + \dots \end{aligned} \quad (\text{D28})$$

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 $\lambda_{\mathcal{M}}^{\downarrow} := \{ \lambda_{\mathcal{M}}^{\downarrow 0}, \lambda_{\mathcal{M}}^{\downarrow 1}, \dots, \lambda_{\mathcal{M}}^{\downarrow d_{\mathcal{M}}-1} \} = \frac{1}{2} \{ \lambda_{\mathcal{M}}^{\downarrow 0} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2}}, \lambda_{\mathcal{M}}^{\downarrow 0} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2}+1}, \dots, \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2}-1} + \lambda_{\mathcal{M}}^{\downarrow d_{\mathcal{M}}-1} \}$. 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 the terms on the r.h.s. of Eq. (3) 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 $I(S : M)_{\varrho'_{s\mathcal{M}}} \rightarrow 0$. In terms of the reduced states above, this means

that $\sum_{i=0}^{\frac{d_{\mathcal{M}}}{2}-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 1$ and $\sum_{i=\frac{d_{\mathcal{M}}}{2}}^{d_{\mathcal{M}}-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 0$, which can only occur if the largest half of energy eigenvalues of the machine, i.e., ω_i for all $i \geq \frac{d_{\mathcal{M}}}{2}$, 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, $D(\varrho'_{\mathcal{M}} \| \varrho_{\mathcal{M}})$. Here one can already see that an infinite-dimensional machine is required to saturate the Landauer bound: from Ref. [29], $D(\varrho'_{\mathcal{M}} \| \varrho_{\mathcal{M}}) \geq f(\Delta S_{\mathcal{M}}, d_{\mathcal{M}})$, 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 $\frac{1}{2} \|\varrho_{\mathcal{M}} - \varrho'_{\mathcal{M}}\|_1^2 \leq D(\varrho_{\mathcal{M}} \| \varrho'_{\mathcal{M}})$, 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. (D28) leads to a generic term that must be asymptotically satisfied of the form:

$$\frac{\frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor} + \lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2} + \lfloor \frac{i}{2} \rfloor} \right)}{\lambda_{\mathcal{M}}^{\downarrow i}} \rightarrow 1 \quad \forall i \in \{0, \dots, d_{\mathcal{M}} - 1\}. \quad (\text{D29})$$

□

In order to achieve perfect cooling at the Landauer limit, one thus must simultaneously satisfy the conditions outlined in Theorem 11. In other words, to minimise the relative-entropy term with the additional constraints $\sum_{i=0}^{\frac{d_{\mathcal{M}}}{2}-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 1$ and $\sum_{i=\frac{d_{\mathcal{M}}}{2}}^{d_{\mathcal{M}}-1} \lambda_{\mathcal{M}}^{\downarrow i} \rightarrow 0$. 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 $\{0, \dots, \frac{d_{\mathcal{M}}}{2} - 1\}$ or smaller half $\{\frac{d_{\mathcal{M}}}{2}, \dots, d_{\mathcal{M}}\}$, one cannot have $\lambda_{\mathcal{M}}^{\downarrow \frac{d_{\mathcal{M}}}{2} + \lfloor \frac{i}{2} \rfloor} = \lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor} \forall i$ (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. (D24) 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. (D23)]. 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 necessarily 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. [29] (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. [29].

The Hamiltonian of the machine is $d_{\mathcal{M}} := 2^{N+1}$ dimensional,

$$H_{\mathcal{M}} = \sum_{n=0}^N \sum_{j=1}^{2^n} \left(n\Delta |n; j\rangle\langle n; j|_{\mathcal{M}} \right) + N\Delta |N; 2^N+1\rangle\langle N; 2^N+1|_{\mathcal{M}} \quad (\text{D30})$$

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

$$\lambda_{\mathcal{M}}^{\downarrow i} = e^{-\beta\Delta} \lambda_{\mathcal{M}}^{\downarrow \lfloor \frac{i}{2} \rfloor} \quad \forall i \in \{2, \dots, d_{\mathcal{M}} - 1\}, \quad (\text{D31})$$

$$\lambda_{\mathcal{M}}^{\downarrow n} = e^{-\beta\Delta} \lambda_{\mathcal{M}}^{\downarrow n-1} \quad \forall n \in \{1, \dots, N\}. \quad (\text{D32})$$

We introduce a parameter ϵ to express the Gibbs ratio as

$$e^{-\beta\Delta} = \frac{1 - \epsilon}{2}, \quad (\text{D33})$$

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

$$p'_0 = \frac{1}{\mathcal{Z}_{\mathcal{M}}} \left(\sum_{n=0}^{N-1} 2^n \left(\frac{1 - \epsilon}{2} \right)^n + \left(\frac{1 - \epsilon}{2} \right)^N \right), \quad (\text{D34})$$

$$\text{where } \mathcal{Z}_{\mathcal{M}} = \sum_{n=0}^N 2^n \left(\frac{1 - \epsilon}{2} \right)^n + \left(\frac{1 - \epsilon}{2} \right)^N \quad (\text{D35})$$

is the partition function of the machine. The geometric series above evaluates to

$$p'_0 = \left(1 + \frac{\epsilon(1-\epsilon)^N}{1 - (1-\epsilon)^N + \epsilon(1-\epsilon)^N 2^{-N}} \right)^{-1}. \quad (\text{D36})$$

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

$$p'_0 = 1 - \frac{1}{N} \left(\frac{\theta}{e^\theta - 1} \right) + O\left(\frac{1}{N^2}\right), \quad (\text{D37})$$

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. [29] 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 the maximum excited energy level of the machine spectrum. The spectrum of the final state of the machine is then given by

$$\lambda_{\mathcal{M}}'^{i} = \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\lfloor \frac{i}{2} \rfloor} + \lambda_{\mathcal{M}}^{\lfloor \frac{i}{2} \rfloor + \frac{d_{\mathcal{M}}}{2}} \right) \quad \forall i \in \{0, \dots, d_{\mathcal{M}} - 1\}, \quad (\text{D38})$$

which leads to

$$\begin{aligned} \lambda_{\mathcal{M}}'^{0} &= \frac{1}{2} \left(\lambda_{\mathcal{M}}^0 + \lambda_{\mathcal{M}}^{2^N} \right) = \lambda_{\mathcal{M}}^0, & \lambda_{\mathcal{M}}'^{1} &= \frac{1}{2} \left(\lambda_{\mathcal{M}}^0 + \lambda_{\mathcal{M}}^{2^N} \right) = \lambda_{\mathcal{M}}^0, \\ \lambda_{\mathcal{M}}'^{i} &= \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\lfloor \frac{i}{2} \rfloor} + \lambda_{\mathcal{M}}^{\lfloor \frac{i}{2} \rfloor + \frac{d_{\mathcal{M}}}{2}} \right) & \forall i \in \{2, \dots, d_{\mathcal{M}} - 1\} \\ &= \frac{1}{2} \left(\frac{2}{1-\epsilon} \lambda_{\mathcal{M}}^i + \lambda_{\mathcal{M}}^{n=N} \right) = \frac{1}{2} \frac{1}{\mathcal{Z}_{\mathcal{M}}} \left[\left(\frac{1-\epsilon}{2} \right)^{n-1} + \left(\frac{1-\epsilon}{2} \right)^N \right], \end{aligned} \quad (\text{D39})$$

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. (D31) 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

$$\lambda_{\mathcal{M}}'^{\downarrow n} = \frac{1}{2} \left(\lambda_{\mathcal{M}}^{\downarrow n-1} + \lambda_{\mathcal{M}}^{\downarrow N} \right), \quad \forall n \in \{1, 2, \dots, N\}. \quad (\text{D40})$$

The energy cost can now be simply calculated from the difference in the average energy of the machine state,

$$\Delta E_{\mathcal{M}} = \sum_{i=0}^{d_{\mathcal{M}}-1} (\lambda_{\mathcal{M}}'^i - \lambda_{\mathcal{M}}^i) \omega_i, \quad (\text{D41})$$

where we denote the i^{th} energy eigenvalue by ω_i . $\lambda_{\mathcal{M}}'^0$ is unchanged, and although $\lambda_{\mathcal{M}}'^1$ 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

$$\Delta E_{\mathcal{M}} = \sum_{n=1}^N (\lambda_{\mathcal{M}}'^{\downarrow n} - \lambda_{\mathcal{M}}^{\downarrow n}) \omega_n = \frac{1}{\beta} \left[1 - \frac{2(1-\epsilon)^N}{1 - (1-\epsilon)^N + (1-2^{-N})(1-\epsilon)^N \epsilon} \right] \log \left(\frac{2}{1-\epsilon} \right). \quad (\text{D42})$$

As we did above, we parameterise $\epsilon = \frac{\theta}{N}$. The asymptotic behaviour of the energy cost is then

$$\beta \Delta E_{\mathcal{M}} = \log(2) + \frac{1}{N} \left(1 - \frac{2 \log(2)}{e^\theta - 1} \right) \theta + O\left(\frac{1}{N^2}\right), \quad (\text{D43})$$

or in terms of the decrease in entropy of the system,

$$\beta \Delta E_{\mathcal{M}} = \tilde{\Delta} S_{\mathcal{M}} + \frac{\log N}{N} \left(\frac{\theta}{e^\theta - 1} \right) + O\left(\frac{1}{N}\right). \quad (\text{D44})$$

Combining (D37) and (D43), 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. (D24), note that in the limit $N \rightarrow \infty$ (recalling that $\epsilon = \frac{\theta}{N}$),

$$\frac{\lambda'^{\downarrow n}}{\lambda^{\downarrow n}} = \lim_{N \rightarrow \infty} \frac{\frac{1}{2} \left[\left(\frac{1-\epsilon}{2} \right)^{n-1} + \left(\frac{1-\epsilon}{2} \right)^N \right]}{\left(\frac{1-\epsilon}{2} \right)^n} = \lim_{N \rightarrow \infty} \left[\frac{1}{1-\epsilon} + \frac{1}{2^{N-n+1}} \left(\frac{e^{-\theta}}{(1-\epsilon)^n} \right) \right] = 1, \quad (\text{D45})$$

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. (D24) 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. (D24) is satisfied and $S_{\pm}$ containing the all n for which Eq. (D24) 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 $\lambda'_n = \lambda(1 + \Delta_n)$ 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

$$D(\varrho' \parallel \varrho) = - \sum_{n \in S_{\pm}} \lambda'_n \log(1 + \Delta_n) = -N_{\pm} \sum_{n \in S_{\pm}} p_n \log(1 + \Delta_n), \quad (\text{D46})$$

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

$$\frac{1}{2} \log(1 + \theta) + \frac{1}{2} \log(1 - \theta) \geq \sum_{n \in S_{\pm}} p_n \log(1 + \Delta_n). \quad (\text{D47})$$

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

$$D(\varrho' \parallel \varrho) \geq -N_{\pm} \left[\frac{1}{2} \log(1 + \theta) + \frac{1}{2} \log(1 - \theta) \right] = -\frac{N_{\pm}}{2} \log(1 - \theta^2) \geq \frac{N_{\pm}}{2} \theta^2, \quad (\text{D48})$$

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. (D24) 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.

D5. Energy-Gap Variety as a Notion of Control Complexity

The insights drawn above regarding sufficient conditions for cooling a system at the Landauer limit lead us to propose a more nuanced notion of control complexity than the preliminary effective dimension that satisfies the natural desiderata outlined in the main text. In particular, here we will demonstrate that the *energy-gap variety* (see 2) provides a good measure of control complexity, both from a theoretical, thermodynamic standpoint as well as a practical one.

Firstly, it is quite clear that coupling a system to a diverging number of distinct machine energy gaps is a difficult task to achieve in almost any conceivable physical platform, especially when the energy gaps are closely spaced; thus, this definition indeed corresponds to our intuitive understanding of “complex” as an operation that is inherently difficult to perform in practice. Secondly, from all of the optimal cooling protocols that we have outlined in this article, we see that, in contrast to the effective dimension, having a diverging energy-gap variety that densely covers an appropriate interval is sufficient for saturating the Landauer limit, thereby making it a better quantifier of control complexity. The remaining point is to show that its divergence is necessary to cool a system to the ground state using a single control operation with energy cost saturating the Landauer bound, so that it is fully consistent also with Nernst’s unattainability principle. We argue that this is indeed the case below.

Claim: In order to cool $\varrho_S \mapsto |0\rangle\langle 0|$ with a thermal machine $\tau_M(\beta, H_M)$ at Landauer energy cost with a single control operation, the global unitary U must couple the system to a diverging number of distinct energy gaps that densely cover the interval $[\omega_0, \infty)$, where ω_0 is the smallest energy gap of the target system.

Proof: First of all, note that how cold the final system state can be made is bounded by the inequality:

$$\lambda_{\min}(\varrho'_S) \geq e^{-\beta \omega_M^{\max}} \lambda_{\min}(\varrho_S), \quad (\text{D49})$$

where $\lambda_{\min}$ denotes the minimal eigenvalue. For a pure final system state, the l.h.s. of the above equation goes to 0; thus, for any non-trivial initial system state (i.e., such that $\lambda_{\min}(\varrho_S) > 0$) and finite temperature $\beta < \infty$, we must have $\omega_M^{\max} \rightarrow \infty$. This determines the upper limit of the required interval of energy gaps. The lower limit of the required interval comes from the fact that the only subspaces of the machine that are relevant for cooling the target system are those associated to energy gaps that are at least as large as the smallest energy gap of the target, ω_0 [38].

Next, recall the equality form of the Landauer limit, which holds true for *any* global unitary transformation with a thermal machine:

$$\beta \Delta E_M = \tilde{\Delta} S_S + I(S : \mathcal{M})_{\varrho'_{SM}} + D(\varrho'_M \| \varrho_M), \quad (\text{D50})$$

Cooling the target system to a pure state necessitates that the final system and machine are uncorrelated and we therefore have $I(S : \mathcal{M})_{\varrho'_{SM}} = 0$ for the optimal process. We thus need to focus on minimising the relative-entropy term, which we do in the following steps.

Consider for simplicity the target system to be a qubit initially in the maximally mixed state. A generic cooling machine should be able to cool *any* system state, include the maximally mixed one; therefore the following insights pertaining to this special case apply generically. In this case, the initial joint spectrum of the system and machine is given by Eq. (D25). Cooling the target system to the ground state necessitates taking a subset $\mathcal{A}$ of these global eigenvalues such that $\sum_{i \in \mathcal{A}} \lambda_M^{\downarrow i} = 1 - \epsilon$ for arbitrarily small ϵ and placing these into the ground-state subspace of the target, with the remaining (small) ϵ amount of population only contributing to any higher-energy eigenstates [this is essentially a generalisation of the conditions put forth in Eqs. (D23), accounting for an arbitrarily small cooling error].

As discussed previously, there are many possible ways to achieve such a configuration, but there is a *unique* one that minimises the total energy cost of doing so: namely, that in which the reduced final state of the machine is rendered passive. This is because if one compares two protocols achieving the same cooling for the target system, one in which the final machine is passive and any other in which it is not, then the former protocol has the smaller energy cost since a positive amount of energy can be (unitarily) extracted from the latter machine in order to render it passive.

Thus, for any protocol saturating the Landauer limit, the final machine state must be arbitrarily close to a passive state, which implies that it must be approximately diagonal in the local machine energy eigenbasis with the globally ordered spectrum as per Eq. (D26). Moreover, in order to minimise the relative-entropy term and therefore saturate the Landauer limit, the final machine state must be arbitrarily close to the initial (i.e., thermal) machine state; following the argumentation put forth in the previous appendix, this leads to the set of conditions outlined in Eq. (D29), which must be satisfied up to arbitrary precision.

Since we have the exact relationship between the initial and final machine eigenvalues, the contribution to the energy cost from the relative-entropy term can be calculated explicitly, i.e., $D(\varrho' \parallel \varrho) = \sum_n \lambda'_n \log \left(\frac{\lambda'_n}{\lambda_n} \right)$. Following the argumentation from Eq. (D46) until Eq. (D48) in the previous appendix, we see that by assuming a finite deviation from *any* of the conditions of Eq. (D29), i.e., writing $\lambda'_n = \lambda_n(1 + \Delta_n)$ with $|\Delta_n| \geq \theta > 0$ for some θ , one can derive a lower bound on the relative entropy:

$$D(\varrho' \parallel \varrho) \geq \frac{N_{\pm}}{2} \theta^2, \quad (\text{D51})$$

where $N_{\pm}$ is the total population of the subspaces corresponding to the terms in the sum such that $\frac{\lambda'_n}{\lambda_n}$ differs from unity by at least θ . In other words, these are the relevant additional contributions to the energy cost; whenever $N_{\pm}$ is non-zero, the Landauer bound *cannot* be approached arbitrarily closely.

The final piece is to relate the machine eigenvalues to its energy-gap spectrum, which can be done straightforwardly due to the initial thermality of the machine, i.e., $\lambda_{\mathcal{M}}^{\downarrow i} = e^{-\beta \omega_i} / \mathcal{Z}_{\mathcal{M}}(\beta, H_{\mathcal{M}})$. We now argue that if there is ever a finite “jump” in the energy-gap structure of the machine, then one cannot achieve both a ground-state population of the target that is arbitrarily close to unity *and* have $N_{\pm}$ be arbitrarily close to zero concurrently. Suppose now that one has a machine with a dense energy-gap structure from ω_0 up until some (finite) ω_a , followed by a finite jump until the energy level $\omega_a + \Omega$ (for some strictly finite $\Omega > 0$), and then again a dense set of energy gaps throughout the interval $[\omega_a + \Omega, \infty)$. Then, one can utilise the energy-gap structure in the “lower band” $[\omega_0, \omega_a)$ in an optimal fashion in order to cool the target system to a minimum temperature (set by ω_a) at Landauer energy cost [38, 42]. However, assuming that the jump in the energy-gap structure begins at some finite ω_a , then there is always a finite amount of population in the machine that is supported on the energy levels corresponding to the “upper band” $[\omega_a + \Omega, \infty)$. To cool the target to arbitrarily close to the ground state, one must therefore access this population and transfer it to the ground-state subspace of the target; this precisely corresponds to the $N_{\pm}$ that contributes to the excess energy cost in a non-negligible manner for finite population exchanges. In particular, we have the bound $N_{\pm} \geq \min\left(\frac{e^{-\beta \omega_a}}{1 + e^{-\beta \omega_a}}, \frac{1}{1 + e^{-\beta \omega_a}}\right)$. Thus, whenever ω_a takes a finite value, $N_{\pm}$ is a strictly positive number. The only way that the relative-entropy term can vanish then is if θ vanishes; however, this can only occur if $\Omega \rightarrow 0$, because for any finite Ω , the ratio $\frac{\lambda'_n}{\lambda_n}$ for at least one value of n differs from 1 by a finite amount as argued above, which finally leads to a non-zero lower bound in Eq. (D51) and implies that the Landauer limit cannot be saturated. In other words, the endpoints of the lower and upper energy gap intervals considered above must coincide (upto arbitrary precision) in order to saturate the Landauer bound. This implies that the energy-gap variety must diverge and moreover, since the above logic holds for arbitrary ω_a which can be smoothly varied as a parameter, it follows that the diverging number of energy gaps must additionally approximately densely cover the interval in question.

Appendix E: 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 E2a, 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. (6)], such Gaussian operations have infinite control complexity. Subsequently, in Appendix E2b, 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 E3, we reconsider the protocol from Appendix E2b 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.

E1. 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 (resp: lowering, raising) mode operators $\{a_k, a_k^\dagger\}$ satisfying the bosonic commutation relations:

$$[a_k, a_{k'}^\dagger] = \delta_{kk'}, \quad [a_k, a_{k'}] = 0, \quad \forall k, k' = 1, 2, \dots, N. \quad (\text{E1})$$

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)

$$q_k := \frac{1}{\sqrt{2}}(a_k + a_k^\dagger), \quad p_k := \frac{1}{i\sqrt{2}}(a_k - a_k^\dagger). \quad (\text{E2})$$

As a consequence of the commutation relations in Eq. (E1), the generalised position and momentum operators satisfy the canonical commutation relations

$$[q_k, p_l] = i\delta_{kl}. \quad (\text{E3})$$

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

$$[\mathbb{X}_k, \mathbb{X}_l] = i\Omega_{kl}, \quad (\text{E4})$$

where the Ω_{kl} are the components of the symplectic form

$$\Omega = \bigoplus_{j=1}^N \Omega_j, \quad \Omega_j = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}. \quad (\text{E5})$$

The density operator associated to N harmonic oscillators can be written in the so-called *phase-space representation* as

$$\varrho = \frac{1}{(2\pi)^N} \int \chi(\Omega\xi) \mathcal{W}(-\Omega\xi) d^{2N}\xi, \quad (\text{E6})$$

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

$$\chi(\xi) = e^{-\frac{1}{4}\xi^T \Gamma \xi + i\bar{\mathbb{X}}^T \xi}. \quad (\text{E7})$$

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

$$\Gamma_{mn} := \langle \mathbb{X}_m \mathbb{X}_n + \mathbb{X}_n \mathbb{X}_m \rangle_\varrho - 2\langle \mathbb{X}_n \rangle_\varrho \langle \mathbb{X}_m \rangle_\varrho. \quad (\text{E8})$$

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. [55] 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) ,

$$\mathbb{X} \mapsto M\mathbb{X} + \kappa, \quad (\text{E9})$$

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.,

$$M \Omega M^T = \Omega. \quad (\text{E10})$$

Under such a mapping, the first and second moments transform according to

$$\overline{\mathbf{X}} \mapsto M \overline{\mathbf{X}} + \kappa, \quad \Gamma \mapsto M \Gamma M^T. \quad (\text{E11})$$

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 [22]

$$E(\varrho_G) = \sum_k \omega_k \left(\frac{1}{4} \text{tr} [\Gamma^{(k)} - 2] + \frac{1}{2} \|\overline{\mathbf{X}}^{(k)}\|^2 \right), \quad (\text{E12})$$

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 $\overline{\mathbf{X}}^{(k)}$ denotes the two-component subvector of first moments for the k^{th} mode of the displacement vector $\overline{\mathbf{X}}$.

E2. Diverging-Time Cooling Protocol for Harmonic Oscillators

E2a. 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 ω_S interacting with a harmonic oscillator machine $\mathcal{M}$ at frequency $\omega_M \geq \omega_S$ via a (non-energy-conserving) unitary acting on the joint system $\mathcal{SM}$ initialised as a tensor product of thermal states $\tau_S(\beta, H_S) \otimes \tau_M(\beta, H_M)$ at inverse temperature β with respect to their local Hamiltonians H_S and H_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.,

$$\Gamma[\tau_S(\beta, H_S) \otimes \tau_M(\beta, H_M)] = \Gamma[\tau_S(\beta, H_S)] \oplus \Gamma[\tau_M(\beta, H_M)], \quad (\text{E13})$$

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 $T_{\min} = \frac{\omega_S}{\omega_M} T$ (for the case $\omega_M \geq \omega_S$) [38]. The non-energy-conserving unitary transformation that achieves this is of the form

$$U = e^{-i\frac{\pi}{2}(a^\dagger b + ab^\dagger)}, \quad (\text{E14})$$

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. (E11)]

$$\begin{bmatrix} \coth\left(\frac{\beta\omega_S}{2}\right) \mathbb{1}_2 & 0 \\ 0 & \coth\left(\frac{\beta\omega_M}{2}\right) \mathbb{1}_2 \end{bmatrix} \xrightarrow{\text{SWAP}} \begin{bmatrix} \coth\left(\frac{\beta\omega_M}{2}\right) \mathbb{1}_2 & 0 \\ 0 & \coth\left(\frac{\beta\omega_S}{2}\right) \mathbb{1}_2 \end{bmatrix}. \quad (\text{E15})$$

This means that both the output target system and machine are thermal states at different temperatures $T'_S = \frac{\omega_S}{\omega_M} T$ and $T'_M = \frac{\omega_M}{\omega_S} T$. Making use of Eq. (E12), we can calculate the energy change for the system and machine as

$$\Delta E_S = E\left[\tau_S\left(\frac{\omega_M}{\omega_S}\beta, H_S\right)\right] - E[\tau_S(\beta, H_S)] = \frac{\omega_S}{2} \left[\coth\left(\frac{\beta\omega_M}{2}\right) - \coth\left(\frac{\beta\omega_S}{2}\right) \right],$$

$$\Delta E_{\mathcal{M}} = E \left[\tau_{\mathcal{M}} \left(\frac{\omega_s}{\omega_{\mathcal{M}}} \beta, H_{\mathcal{M}} \right) \right] - E [\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}})] = \frac{\omega_{\mathcal{M}}}{2} \left[\coth \left(\frac{\beta \omega_s}{2} \right) - \coth \left(\frac{\beta \omega_{\mathcal{M}}}{2} \right) \right]. \quad (\text{E16})$$

The total energy cost associated to such a swap operation is thus

$$\Delta E_{s,\mathcal{M}} = \Delta E_s + \Delta E_{\mathcal{M}} = \frac{(\omega_{\mathcal{M}} - \omega_s)}{2} \left[\coth \left(\frac{\beta \omega_s}{2} \right) - \coth \left(\frac{\beta \omega_{\mathcal{M}}}{2} \right) \right] = (\omega_{\mathcal{M}} - \omega_s) \frac{e^{-\beta \omega_s} (1 - e^{-\beta(\omega_{\mathcal{M}} - \omega_s)})}{(1 - e^{-\beta \omega_s})(1 - e^{-\beta \omega_{\mathcal{M}}})}. \quad (\text{E17})$$

Note that this form is similar to that for finite-dimensional systems with equally spaced Hamiltonian [cf., Eq. (C19)]; 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. (6) 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 $\omega_{\mathcal{M}_n} = \omega_s + n\epsilon$. 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 $\tau_{\mathcal{M}}(\beta, H_{\mathcal{M}}) = \bigotimes_{n=1}^N \tau_{\mathcal{M}_n}(\beta, H_{\mathcal{M}_n})$.

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. (E15), the covariance matrix transformation of the two-mode process at the first time step takes the form

$$\Gamma^{(1)}(\tau_s(\beta) \otimes \tau_{\mathcal{M}_1}(\beta)) = \begin{bmatrix} \coth \left(\frac{\beta \omega_s}{2} \right) \mathbb{1}_2 & 0 \\ 0 & \coth \left(\frac{\beta(\omega_s + \epsilon)}{2} \right) \mathbb{1}_2 \end{bmatrix} \xrightarrow{\text{SWAP}} \Gamma_{\text{opt}}^{(1)} = \begin{bmatrix} \coth \left(\frac{\beta(\omega_s + \epsilon)}{2} \right) \mathbb{1}_2 & 0 \\ 0 & \coth \left(\frac{\beta \omega_s}{2} \right) \mathbb{1}_2 \end{bmatrix}. \quad (\text{E18})$$

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 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

$$\begin{aligned} \Gamma^{(n)}(\tau_s(\beta) \otimes \tau_{\mathcal{M}_n}(\beta)) &= \begin{bmatrix} \coth \left(\frac{\beta(\omega_s + (n-1)\epsilon)}{2} \right) \mathbb{1}_2 & 0 \\ 0 & \coth \left(\frac{\beta(\omega_s + n\epsilon)}{2} \right) \mathbb{1}_2 \end{bmatrix} \\ \xrightarrow{\text{SWAP}} \Gamma_{\text{opt}}^{(n)} &= \begin{bmatrix} \coth \left(\frac{\beta(\omega_s + n\epsilon)}{2} \right) \mathbb{1}_2 & 0 \\ 0 & \coth \left(\frac{\beta(\omega_s + (n-1)\epsilon)}{2} \right) \mathbb{1}_2 \end{bmatrix}. \end{aligned} \quad (\text{E19})$$

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. (E16), one can calculate the energy changes of the target system and the machine at each time step as

$$\begin{aligned} \Delta E_s^{(n)} &= \frac{\omega_s}{2} \left[\coth \left(\frac{\beta(\omega_s + n\epsilon)}{2} \right) - \coth \left(\frac{\beta(\omega_s + (n-1)\epsilon)}{2} \right) \right], \\ \Delta E_{\mathcal{M}_n}^{(n)} &= \frac{(\omega_s + n\epsilon)}{2} \left[\coth \left(\frac{\beta(\omega_s + (n-1)\epsilon)}{2} \right) - \coth \left(\frac{\beta(\omega_s + n\epsilon)}{2} \right) \right]. \end{aligned} \quad (\text{E20})$$

The total energy change for the target system during the overall process (i.e., throughout the N steps) is thus given by

$$\Delta E_s = \sum_{n=1}^N \Delta E_s^{(n)} = \sum_{n=1}^N \frac{\omega_s}{2} \left[\coth \left(\frac{\beta(\omega_s + n\epsilon)}{2} \right) - \coth \left(\frac{\beta(\omega_s + (n-1)\epsilon)}{2} \right) \right]$$

$$= \frac{\omega_s}{2} \left[\coth \left(\frac{\beta(\omega_s + N\epsilon)}{2} \right) - \coth \left(\frac{\beta\omega_s}{2} \right) \right] = \omega_s \left[\frac{e^{-\beta(\omega_s + N\epsilon)}}{1 - e^{-\beta(\omega_s + N\epsilon)}} - \frac{e^{-\beta\omega_s}}{1 - e^{-\beta\omega_s}} \right]. \quad (\text{E21})$$

Here, we have written $\coth(x) = 1 + (2e^{-2x})/(1 - e^{-2x})$. Similarly, one can obtain the total energy change of the overall machine

$$\begin{aligned} \Delta E_{\mathcal{M}} &= \sum_{n=1}^N \Delta E_{\mathcal{M}_n}^{(n)} = \sum_{n=1}^N \frac{\omega_s + n\epsilon}{2} \left[\coth \left(\frac{\beta(\omega_s + (n-1)\epsilon)}{2} \right) - \coth \left(\frac{\beta(\omega_s + n\epsilon)}{2} \right) \right] \\ &= \sum_{n=1}^N (\omega_s + n\epsilon) \left[\frac{e^{-\beta(\omega_s + (n-1)\epsilon)}}{1 - e^{-\beta(\omega_s + (n-1)\epsilon)}} - \frac{e^{-\beta(\omega_s + n\epsilon)}}{1 - e^{-\beta(\omega_s + n\epsilon)}} \right]. \end{aligned} \quad (\text{E22})$$

It is straightforward to check that the total energy change, i.e., the sum of Eqs. (E21) and (E22), is equal to the energy cost obtained in Eq. (C19) with $d \rightarrow \infty$. In particular, this can be seen by considering the second line of Eq. (C19), 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 $\epsilon = (\omega_{\max} - \omega_s)/N \rightarrow 0$, 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. (6), 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. [22, 37]). 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 E3). 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.

E2b. 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

$$\tau(\beta) = \frac{e^{-\beta H}}{\text{tr}[e^{-\beta H}]} = \sum_{n=0}^{\infty} e^{-\beta\omega n} (1 - e^{-\beta\omega}) |n\rangle\langle n| = \sum_{n=0}^{\infty} p_n |n\rangle\langle n|_{\mathcal{S}\mathcal{M}}, \quad (\text{E23})$$

where the Hamiltonian H is given by $H = \sum_{n=0}^{\infty} n\omega |n\rangle\langle n|$ and the $p_n = e^{-\beta\omega n} (1 - e^{-\beta\omega})$ are the eigenvalues of τ . The joint initial state is a product state that we can then write as

$$\tau_{\mathcal{S}}(\beta) \otimes \tau_{\mathcal{M}}(\beta) = \sum_{m,n=0}^{\infty} p_m p_n |m\rangle\langle m|_{\mathcal{S}} \otimes |n\rangle\langle n|_{\mathcal{M}} = \sum_{m,n=0}^{\infty} \tilde{p}_{m+n} |m, n\rangle\langle m, n|, \quad (\text{E24})$$

where we have defined $\tilde{p}_k := e^{-\beta\omega k} (1 - e^{-\beta\omega})^2$. 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

$$U_{\mathcal{SM}}^{(k)} = \mathbb{1}_{\mathcal{SM}} - |0, k\rangle\langle 0, k| - |m_k, n_k\rangle\langle m_k, n_k| + (|0, k\rangle\langle m_k, n_k| + |m_k, n_k\rangle\langle 0, k|). \quad (\text{E25})$$

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. [56, 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

$$\sum_{k=0}^{\infty} (k+1) \tilde{p}_k = \sum_{k=0}^{\infty} (k+1) e^{-\beta\omega k} (1 - e^{-\beta\omega})^2 = 1, \quad (\text{E26})$$

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 D), 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

$$\tilde{n}(k) := \sum_{n=0}^{k-1} (n+1) = \frac{1}{2} k(k+1), \quad (\text{E27})$$

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

$$\frac{E_{\mathcal{M}}^{\text{final}}}{\omega} = \sum_{k=1}^{\infty} e^{-\beta\omega k} (1 - e^{-\beta\omega})^2 \sum_{n=\tilde{n}(k)}^{\tilde{n}(k+1)-1} n = \sum_{k=1}^{\infty} e^{-\beta\omega k} (1 - e^{-\beta\omega})^2 \frac{1}{2} k(k+1)(k+2) = \frac{3}{4} \text{cosech}^2\left(\frac{\beta\omega}{2}\right). \quad (\text{E28})$$

Since the energy of the initial thermal state is given by

$$\frac{E[\tau(\beta)]}{\omega} = \sum_{n=0}^{\infty} n e^{-\beta\omega n} (1 - e^{-\beta\omega}) = \frac{e^{-\beta\omega}}{1 - e^{-\beta\omega}}, \quad (\text{E29})$$

we thus arrive at the energy cost

$$\frac{\Delta E_{\mathcal{M}}}{\omega} = \frac{E_{\mathcal{M}}^{\text{final}} - E[\tau(\beta)]}{\omega} = \frac{e^{-\beta\omega} (2 + e^{-\beta\omega})}{(1 - e^{-\beta\omega})^2}. \quad (\text{E30})$$

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

$$\begin{aligned} S[\tau(\beta)] &= -\text{tr}[\tau \log(\tau)] = -\sum_{n=0}^{\infty} e^{-\beta\omega n} (1 - e^{-\beta\omega}) \log[e^{-\beta\omega n} (1 - e^{-\beta\omega})] = -\sum_{n=0}^{\infty} e^{-\beta\omega n} (1 - e^{-\beta\omega}) [-\beta\omega n + \log(1 - e^{-\beta\omega})] \\ &= \frac{\beta\omega e^{-\beta\omega}}{1 - e^{-\beta\omega}} + \beta\omega + \log\left(\frac{e^{-\beta\omega}}{1 - e^{-\beta\omega}}\right) = \frac{\beta\omega}{1 - e^{-\beta\omega}} + \log\left(\frac{e^{-\beta\omega}}{1 - e^{-\beta\omega}}\right). \end{aligned} \quad (\text{E31})$$

Using the results from Eqs. (E30) and (E31), 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 $D(\rho'_{\mathcal{S}\mathcal{M}} \| \tau_{\mathcal{M}})$ in Eq. (3) 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 E2a), 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.

E3. 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 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 unitary transformation in Eq. (E14), namely

$$U_k := e^{i\frac{\pi}{2}(a^\dagger b_k + a b_k^\dagger)}. \quad (\text{E32})$$

One can then apply the diverging-time cooling protocol from Appendix E2a to cool the system to the ground state at the Landauer limit via the total unitary transformation

$$U_{\text{tot}} := \lim_{N \rightarrow \infty} U_{(N)}, \quad \text{with} \quad U_{(N)} := \prod_{k=1}^N U_k. \quad (\text{E33})$$

We now seek the Hamiltonian that generates U_{tot} . First note that $U_{(N)} a U_{(N)}^\dagger = i b_1$ and

$$U_{(N)} b_k U_{(N)}^\dagger = \begin{cases} -b_{k+1}, & \text{for } k < N \\ i a, & \text{for } k = N \\ b_k, & \text{for } k > N \end{cases}, \quad (\text{E34})$$

which can be proven by induction. In contrast with Appendix E2a, 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 $\vec{\xi} := (a \ b_1 \ b_2 \ \dots \ b_N \ a^\dagger \ b_1^\dagger \ b_2^\dagger \ \dots \ b_N^\dagger)^\top$, we can write the transformation above as $U_{(N)}\vec{\xi}U_{(N)}^\dagger = S^\top\vec{\xi}$ [57], where

$$S = \begin{pmatrix} \alpha_{(N)} & 0 \\ 0 & \alpha_{(N)} \end{pmatrix}, \quad \text{with} \quad \alpha_{(N)} := \begin{pmatrix} 0 & 0 & 0 & \dots & 0 & i \\ i & 0 & 0 & \dots & 0 & 0 \\ 0 & -1 & 0 & \dots & 0 & 0 \\ 0 & 0 & -1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & -1 & 0 \end{pmatrix}. \quad (\text{E35})$$

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 $S = \exp(-iKh_{(N)})$ [57], i.e., $h_{(N)} = iK \log(S^\top) = iK \log(S)^\top$, where we take the principal logarithm. To calculate this, we must diagonalise the matrix $\alpha_{(N)}$ in Eq. (E35). The eigenvalues of $\alpha_{(N)}$ are

$$\lambda_k := -e^{-i\pi \frac{2k-1}{N+1}}, \quad \text{with} \quad k \in \{1, 2, \dots, N+1\}, \quad (\text{E36})$$

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 $\vec{v}_k$:

$$V := \begin{pmatrix} \vec{v}_1 & \vec{v}_2 & \vec{v}_3 & \dots & \vec{v}_{N+1} \end{pmatrix} \quad \text{with} \quad \vec{v}_k := \frac{-1}{\sqrt{N+1}} \begin{pmatrix} i(-\lambda_k)^{-1} \\ (-\lambda_k)^{-2} \\ (-\lambda_k)^{-3} \\ \vdots \\ (-\lambda_k)^{-(N+1)} \end{pmatrix}. \quad (\text{E37})$$

Specifically, $\alpha_{(N)} = VDV^\dagger$, where $D := \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_{N+1})$, and thus

$$h_{(N)}^\top = iK \log \begin{pmatrix} VDV^\dagger & 0 \\ 0 & VDV^\dagger \end{pmatrix} = iK \begin{pmatrix} V & 0 \\ 0 & W \end{pmatrix} \begin{pmatrix} \log(D) & 0 \\ 0 & \log(D) \end{pmatrix} \begin{pmatrix} V^\dagger & 0 \\ 0 & V^\dagger \end{pmatrix} =: \begin{pmatrix} A & 0 \\ 0 & -A \end{pmatrix} \quad (\text{E38})$$

for some matrix A . By direct calculation, one finds that

$$A_{jk} = i^{\delta_{j1}} i^{\delta_{k1}} \frac{\pi}{(N+1)^2} \sum_{p=1}^{N+1} (2p-2-N) e^{-i\pi \frac{2p-1}{N+1}(j-k)}. \quad (\text{E39})$$

Now, considering the identity

$$\sum_{p=1}^{N+1} e^{i\theta p} = \frac{e^{i\theta(N+1)} - 1}{1 - e^{i\theta}} \quad (\text{E40})$$

for $\theta \in \mathbb{R}$, as well as its derivative with respect to θ , one can calculate the sum in Eq. (E39). We then have

$$\lim_{N \rightarrow \infty} A_{jk} = \begin{cases} 0, & \text{for } j = k \\ ii^{\delta_{j1}} i^{\delta_{k1}} \frac{1}{j-k}, & \text{for } j \neq k \end{cases}. \quad (\text{E41})$$

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.,

$$H_{\text{tot}} = - \sum_{j=2}^{\infty} \left(\frac{1}{j-1} b_j^\dagger a + \text{H.c.} \right) + \sum_{j,k=1; j \neq k}^{\infty} \frac{i}{j-k} b_j^\dagger b_k. \quad (\text{E42})$$

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.

Appendix F: 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

$$[H_{\mathcal{S}} + H_{\mathcal{C}} + H_{\mathcal{H}}, U_{\text{EC}}] = 0. \quad (\text{F1})$$

In Ref. [38], 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

$$U_{\text{EC}} = |0, 1, 0\rangle\langle 1, 0, 1|_{\mathcal{SCH}} + |1, 0, 1\rangle\langle 0, 1, 0|_{\mathcal{SCH}} + \bar{\mathbb{I}}, \quad (\text{F2})$$

where $\bar{\mathbb{I}}$ 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

$$U_{\text{EC}} = \left[\sum_{m,n,l=0}^{d-2} |m, n+1, l\rangle\langle m+1, n, l+1|_{\mathcal{SCH}} + |m+1, n, l+1\rangle\langle m, n+1, l|_{\mathcal{SCH}} \right] + \bar{\mathbb{I}}. \quad (\text{F3})$$

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 G, 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. (F1).

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 $\varrho_{\mathcal{SCH}} = \sum_{m,n,l=0}^1 p_{mnl} |m, n, l\rangle\langle m, n, l|_{\mathcal{SCH}}$. By applying a unitary U_{EC} of the form given in Eq. (F3), the post-transformation joint state is

$$\varrho'_{\mathcal{SCH}} = U_{\text{EC}} \varrho_{\mathcal{SCH}} U_{\text{EC}}^\dagger = \varrho_{\mathcal{SCH}} + \Delta p |0, 1, 0\rangle\langle 0, 1, 0|_{\mathcal{SCH}} - \Delta p |1, 0, 1\rangle\langle 1, 0, 1|_{\mathcal{SCH}}, \quad (\text{F4})$$

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, $\Delta E_{\mathcal{X}} = \text{tr} [H_{\mathcal{X}}(\varrho'_{\mathcal{X}} - \varrho_{\mathcal{X}})]$, can be calculated via

$$\Delta E_{\mathcal{S}} = -\omega_{\mathcal{S}} \Delta p, \quad \Delta E_{\mathcal{C}} = \omega_{\mathcal{C}} \Delta p, \quad \Delta E_{\mathcal{H}} = -\omega_{\mathcal{H}} \Delta p. \quad (\text{F5})$$

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.

F1. Diverging Energy: Proof of Theorem 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 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_S with associated Hamiltonian with finite but otherwise arbitrary energy gaps $H_S = \sum_{i=0}^{d_S-1} \omega_S^i |i\rangle\langle i|_S$, and let d_C and d_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 $p_{ijk} := p_S^i p_C^j p_H^k$.

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 $\frac{1}{d_C d_H}$ 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 $|0jk\rangle_{SC\mathcal{H}}$. This intuition is formalised via Lemma 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 $|i00\rangle_{SC\mathcal{H}}$ must be degenerate with some of the form $|0jk\rangle_{SC\mathcal{H}}$. This degeneracy condition, in turn, also implies that every energy eigenstate of the form $|0jk\rangle_{SC\mathcal{H}}$ 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 $|0jk\rangle_{SC\mathcal{H}}$, 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 $|i00\rangle_{SC\mathcal{H}}$, which does not contribute to the final ground state population of the target. We formalise this intuition in the discussion following Lemma 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 $SC\mathcal{H}$ 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 2. *Given $\mathcal{S}$, d_C and d_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 $|i00\rangle_{SC\mathcal{H}}$, where $i \in \{1, \dots, d_S - 1\}$, energy eigenstate is degenerate with at least one $|0jk\rangle_{SC\mathcal{H}}$ energy eigenstate, where $j \in \{0, \dots, d_C - 1\}, k \in \{0, \dots, d_H - 1\}$.*

Proof. Suppose that there exists an $i^* \in \{1, \dots, d_S - 1\}$ such that $|i^*00\rangle_{SC\mathcal{H}}$ is not degenerate with any $|0jk\rangle_{SC\mathcal{H}}$, where $j \in \{0, \dots, d_C - 1\}, k \in \{0, \dots, d_H - 1\}$. We show that, then, one cannot cool $\mathcal{S}$ arbitrarily close to zero.

Let B_i denote the degenerate subspace of the total Hamiltonian $H_S + H_C + H_H$, where $H_{\mathcal{X}}$ denotes the Hamiltonian of system $\mathcal{X} \in \{\mathcal{S}, \mathcal{C}, \mathcal{H}\}$, that contains the eigenvector $|i00\rangle_{SC\mathcal{H}}$. 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 $[U_{\text{EC}}, H_S + H_C + H_H] = 0$, see, e.g., Lemma 5 of Ref. [30]). This means in particular that the initial population of $|i^*00\rangle_{SC\mathcal{H}}$ can only be distributed within B_{i^*} , and as no eigenvector of the form $|0jk\rangle_{SC\mathcal{H}}$ is contained in B_{i^*} by assumption, it can never contribute to the final ground state population of $\mathcal{S}$, which we denote $\tilde{p}_S^0$. So we have

$$\tilde{p}_S^0 \leq 1 - p_{i^*00}. \quad (\text{F6})$$

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 $\mathcal{Z}_{\mathcal{X}}$ that

$$\mathcal{Z}_{\mathcal{X}} = 1 + e^{-\beta\omega_{\mathcal{X}}^1} + \dots + e^{-\beta\omega_{\mathcal{X}}^{d_{\mathcal{X}}-1}} \leq d_{\mathcal{X}}, \quad (\text{F7})$$

and so we have the following bound on the initial populations associated to each eigenvector $|i00\rangle_{SC\mathcal{H}}$

$$p_{i00} = \frac{e^{-\beta\omega_S^i}}{\mathcal{Z}_S \mathcal{Z}_C \mathcal{Z}_\mathcal{H}} \geq \frac{e^{-\beta\omega_S^i}}{\mathcal{Z}_S d_C d_\mathcal{H}} > 0 \quad \forall i \in \{1, \dots, d_S - 1\}. \quad (\text{F8})$$

Combining the above, we have that

$$\hat{p}_S^0 \leq 1 - \frac{e^{-\beta\omega_S^{i^*}}}{\mathcal{Z}_S d_C d_\mathcal{H}} < 1. \quad (\text{F9})$$

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 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 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

$$p_{0j^*k^*} = \frac{e^{-\beta\omega_C^{j^*}} e^{-\beta\omega_\mathcal{H}^{k^*}}}{\mathcal{Z}_S \mathcal{Z}_C \mathcal{Z}_\mathcal{H}} \geq \frac{e^{-2\beta\omega_S^{i^*}}}{\mathcal{Z}_S \mathcal{Z}_C \mathcal{Z}_\mathcal{H}} \geq \frac{e^{-2\beta\omega_S^{i^*}}}{\mathcal{Z}_S d_C d_\mathcal{H}} =: q_{i^*}. \quad (\text{F10})$$

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

$$B_{i^*} = \text{span}\{|0jk\rangle, |0j_2k_2\rangle, \dots, |0j_\mu k_\mu\rangle, |i^*00\rangle, |i_2\ell_2 m_2\rangle, \dots, |i_\nu \ell_\nu m_\nu\rangle\}. \quad (\text{F11})$$

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 [54], 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 $\mathbf{v}_\alpha$, we have

$$\tilde{p}_{i^*00} + \sum_{\alpha=2}^{\nu} \tilde{p}_{i_\alpha \ell_\alpha m_\alpha} \geq \sum_{\alpha=1}^{\nu} \mathbf{v}_\alpha^\uparrow. \quad (\text{F12})$$

We now claim that $\sum_{\alpha=1}^{\nu} \mathbf{v}_\alpha^\uparrow \geq q_{i^*}$ from Eq. (F10). 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 $\sum_{\alpha=1}^{\nu} \mathbf{v}_\alpha^\uparrow$. Let x be that element. As $\mathbf{v}_\alpha^\uparrow \geq 0$ for all $\alpha = 1, \dots, \pi = \mu + \nu$, we have

$$\sum_{\alpha=1}^{\nu} \mathbf{v}_\alpha^\uparrow \geq x. \quad (\text{F13})$$

Now as $p_{0j_\gamma k_\gamma} \geq q_{i^*}$ for all $\gamma = 2, \dots, \mu$, we have

$$x \geq \min(q_{i^*}, p_{i^*00}) = q_{i^*}, \quad (\text{F14})$$

where $p_{i^*00} \geq q_{i^*}$ can be seen from Eq. (F10), as claimed.

As the l.h.s. of Eq. (F12) 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

$$\hat{p}_S^0 \leq 1 - \left(\tilde{p}_{i^*00} + \sum_{\alpha=2}^{\nu} \tilde{p}_{i_\alpha \ell_\alpha m_\alpha} \right) \leq 1 - q_{i^*} = 1 - \frac{e^{-2\beta\omega_S^{i^*}}}{\mathcal{Z}_S d_C d_\mathcal{H}}. \quad (\text{F15})$$

So, for any finite-dimensional machine, one cannot cool the system $\mathcal{S}$ beyond $1 - \frac{e^{-\beta\omega_S^{i^*}}}{\mathcal{Z}_S d_C d_\mathcal{H}}$, 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

$$\varrho_{sc\mathcal{H}}^{(0)} = \frac{(|0\rangle\langle 0| + e^{-\beta\omega_S}|1\rangle\langle 1|)_S \otimes (|0\rangle\langle 0| + e^{-\beta\omega_C}|1\rangle\langle 1|)_C \otimes (|0\rangle\langle 0| + e^{-\beta_H\omega_{\mathcal{H}}}|1\rangle\langle 1|)_{\mathcal{H}}}{\mathcal{Z}_S(\beta, \omega_S)\mathcal{Z}_C(\beta, \omega_C)\mathcal{Z}_{\mathcal{H}}(\beta_H, \omega_{\mathcal{H}})}. \quad (\text{F16})$$

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

$$p'_S(0) = \langle 0 | \text{tr}_{c\mathcal{H}} [U\varrho_{sc\mathcal{H}}^{(0)}U^\dagger] | 0 \rangle_S = \frac{1 + e^{-\beta_H\omega_{\mathcal{H}}}(1 + e^{-\beta\omega_S} + e^{-\beta\omega_C})}{\mathcal{Z}_S(\beta, \omega_S)\mathcal{Z}_C(\beta, \omega_C)\mathcal{Z}_{\mathcal{H}}(\beta_H, \omega_{\mathcal{H}})} < 1. \quad (\text{F17})$$

Indeed, using $e^{-\beta\omega_S}e^{-\beta_H\omega_{\mathcal{H}}} \geq e^{-\beta\omega_C}$,

$$p'_S(0) \leq \frac{1 + e^{-\beta_H\omega_{\mathcal{H}}}e^{-\beta\omega_S}}{\mathcal{Z}_S(\beta, \omega_S)\mathcal{Z}_C(\beta, \omega_C)} \leq \frac{1}{\mathcal{Z}_C(\beta, \omega_C)} \leq 1. \quad (\text{F18})$$

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 $p'_S(0) = p_S(0) < 1$.

F2. 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 C, 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

$$\varrho_{sc\mathcal{H}}^{(n)} = \tau_S(\beta_n, \omega_S) \otimes \tau_C(\beta, \omega_n + \epsilon_n) \otimes \tau_{\mathcal{H}}(\beta_H, \omega_n + \epsilon_n - \omega_S). \quad (\text{F19})$$

In order to cool the target system via said unitary, we must have that $p_{101} \geq p_{010}$ for the state in Eq. (F19), which implies that ϵ_n must satisfy the following condition:

$$e^{-\beta\omega_n - \beta_H(\omega_n + \epsilon_n - \omega_S)} \geq e^{-\beta(\omega_n + \epsilon_n)} \Rightarrow \epsilon_n \geq \gamma(\omega_n - \omega_S) \quad \text{where} \quad \gamma := \frac{\beta_H}{\beta - \beta_H}. \quad (\text{F20})$$

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 6.

F3. 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. [38], it is shown that by repeating the incoherent cooling process (i.e., implementing the unitary in Eq. (F3)) 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

$$e^{-\beta\omega_n^{\text{eff}}} := \frac{p_{1_{\mathcal{C}}0_{\mathcal{H}}}}{p_{0_{\mathcal{C}}1_{\mathcal{H}}}} = e^{-\beta(\omega_n + \epsilon_n)} e^{\beta_H(\omega_n + \epsilon_n - \omega_S)} \Rightarrow \omega_n^{\text{eff}} = \omega_n + \epsilon_n - \frac{\beta_H}{\beta}(\omega_n + \epsilon_n - \omega_S). \quad (\text{F21})$$

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 C, i.e., considering a diverging number of machines, each of which having the part connected to the cold bath with energy gap $\omega_{\mathcal{C}_n} = \omega_n + \epsilon_n$ 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 C: 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. (F5), 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

$$\begin{aligned} \Delta E_S^{(n)} &= \text{tr} [H_S(\omega_S)(\varrho_S^*(\beta, \omega_n) - \varrho_S^*(\beta, \omega_{n-1}))] \\ \Delta E_C^{(n)} &= -\text{tr} [H_C(\omega_n)(\varrho_S^*(\beta, \omega_n) - \varrho_S^*(\beta, \omega_{n-1}))] \\ \Delta E_{\mathcal{H}}^{(n)} &= \text{tr} [H_{\mathcal{H}}(\omega_n - \omega_S)(\varrho_S^*(\beta, \omega_n) - \varrho_S^*(\beta, \omega_{n-1}))] \end{aligned} \quad (\text{F22})$$

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. (F20)]. 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 $\omega_N = \omega_{\text{max}}$ for the maximum energy gap of the cold machine subsystems, the energy exchanged throughout the entire cooling protocol here is given by

$$\begin{aligned} \Delta E_S &= \lim_{N \rightarrow \infty} \sum_{n=1}^N \Delta E_S^{(n)} = \text{tr} [H_S(\omega_S)(\varrho_S^*(\beta, \omega_{\text{max}}) - \varrho_S^*(\beta, \omega_S))] \\ \Delta E_C &= \lim_{N \rightarrow \infty} \sum_{n=1}^N \Delta E_C^{(n)} = \frac{1}{\beta} \{S[\varrho_S^*(\beta, \omega_S)] - S[\varrho_S^*(\beta, \omega_{\text{max}})]\} = \frac{1}{\beta} \tilde{\Delta} S_S \\ \Delta E_{\mathcal{H}} &= \lim_{N \rightarrow \infty} \sum_{n=1}^N \Delta E_{\mathcal{H}}^{(n)} = -\Delta E_S - \Delta E_C. \end{aligned} \quad (\text{F23})$$

Here, the expression for ΔE_C can be derived using the same arguments as presented in Appendix C1. 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.

F4. 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. (F21), 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 4; in Appendix G, 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. (F19), where $\epsilon_n = \gamma(\omega_n - \omega_S) + \epsilon$ and $\omega_n = \omega_S + n\epsilon$. Here, γ is as in Eq. (F20). 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. (F20)]. Based on Eq. (F5), 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. (F22), we have

$$\begin{aligned} \Delta E_c^{(n+1)} &\geq -\text{tr}\{H_c(\omega_n + \epsilon_n)[\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\} \\ &= -\text{tr}\{H_c[(\gamma + 1)\omega_n - \gamma\omega_S + \epsilon][\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\} \\ &= -\text{tr}\{H_c[(\gamma + 1)\omega_n - \gamma\omega_S + \epsilon + \gamma\epsilon - \gamma\epsilon][\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\} \\ &= -(\gamma + 1)\text{tr}\{H_c(\omega_n + \epsilon)[\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\} + \gamma\text{tr}\{H_c(\omega_S + \epsilon)[\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\} \\ &= (\gamma + 1)\Delta E_c^{*(n+1)} + \gamma\Delta E_s^{*(n+1)} + \gamma\text{tr}\{H_c(\epsilon)[\varrho_S^*(\beta, \omega_n + \epsilon) - \varrho_S^*(\beta, \omega_n)]\}, \end{aligned} \quad (\text{F24})$$

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 ΔE_c^* to denote the idealised energy cost [i.e., that corresponding to what would be achievable in the infinite-temperature setting; see Eq. (F22)] 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 $\tau_S(\beta_{\text{max}}, \omega_S)$ is given by

$$\begin{aligned} \Delta E_c &= \lim_{N \rightarrow \infty} \sum_{n=1}^N \Delta E_c^{(n+1)} \\ &\geq (\gamma + 1) \frac{1}{\beta} \tilde{\Delta} S_s + \gamma \Delta E_s \\ &= \frac{1}{\beta} \tilde{\Delta} S_s + \gamma \left(\Delta E_s + \frac{1}{\beta} \tilde{\Delta} S_s \right). \end{aligned} \quad (\text{F25})$$

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. (A14), 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.

Appendix G: 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 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.,*

$$|\beta' - \beta^*| < \epsilon_1, \quad (\text{G1})$$

with an energy cost, measured as heat drawn from the hot bath, that is arbitrarily close to the ideal Carnot-Landauer limit, i.e.,

$$\left| \Delta E_{\mathcal{H}} - \eta^{-1} \tilde{\Delta} F_s^{(\beta)} \right| < \epsilon_2, \quad (\text{G2})$$

where $\eta = 1 - \beta_H/\beta$ 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.

G1. 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:

- 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 $\tau_C(\beta, \omega_S + n\theta) \otimes \tau_H(\beta_H, n\theta)$.
- The energy-preserving three qubit unitary cycle in the $\{010, 101\}_{\mathcal{SCH}}$ subspace is performed [see Eq. (F3)], after which the cold and hot qubits are rethermalised to their respective initial temperatures.
- The above steps are repeated m_n times.

The energy increment θ is defined as

$$\theta := \frac{\omega_S}{N} \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right), \quad (\text{G3})$$

while the number of repetitions within each stage is given by

$$m_n = \left\lceil \frac{\log(\delta)}{\log(1 - N_V^{(n)})} \right\rceil. \quad (\text{G4})$$

$\lceil \cdot \rceil$ 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.,

$$N_V^{(n)} := \langle 01 | \tau_C(\beta, \omega_S + n\theta) \otimes \tau_H(\beta_H, n\theta) | 01 \rangle + \langle 10 | \tau_C(\beta, \omega_S + n\theta) \otimes \tau_H(\beta_H, n\theta) | 10 \rangle. \quad (\text{G5})$$

The parameter δ will be chosen appropriately to complete the proof ($\delta = 1/N^2$ 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 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 $|1\rangle\langle 1|_{\mathcal{S}}$. 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. [38]). Thus q_n is defined via the Gibbs ratio

$$\frac{q_n}{1 - q_n} = e^{-\beta(\omega_S + n\theta)} e^{+\beta_H n\theta} = e^{-\beta\omega_S} e^{-(\beta - \beta_H)n\theta}. \quad (\text{G6})$$

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 $\{010, 101\}_{SC\mathcal{H}}$ 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

$$\tilde{\Delta}E_{\mathcal{H}}^{(n)} = \omega_{\mathcal{H}}^{(n)}(p_{n-1} - p_n) = n\theta(p_{n-1} - p_n). \quad (\text{G7})$$

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

$$\frac{q_N}{1 - q_N} = e^{-\tilde{\beta}\omega_S} \Rightarrow e^{-\beta\omega_S} e^{-(\beta - \beta_H)N\theta} = e^{-\tilde{\beta}\omega_S} \Rightarrow \tilde{\beta} = \beta^*, \quad (\text{G8})$$

where we have made use of the definition of θ in Eq. (G3). 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

$$\tilde{\Delta}E_{\mathcal{H}}^* = \sum_{n=1}^N n\theta(q_{n-1} - q_n), \quad (\text{G9})$$

which can alternatively be expressed as

$$\tilde{\Delta}E_{\mathcal{H}}^* = \sum_{n=1}^N [(n-1)\theta(q_{n-1} - q_n)] + \theta(q_0 - q_N) \quad (\text{G10})$$

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.,

$$I := - \int_{q_0}^{q^*} y \, dq, \quad \text{where} \quad \frac{q(y)}{1 - q(y)} = e^{-\beta\omega_S} e^{-(\beta - \beta_H)y}, \quad (\text{G11})$$

from Eq. (G6). 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

$$\sum_{n=1}^N (n-1)\theta(q_{n-1} - q_n) \leq \int_{q_0}^{q^*} y \, dq \leq \sum_{n=1}^N n\theta(q_{n-1} - q_n), \quad (\text{G12})$$

from which we can deduce that the value of $\Delta E_{\mathcal{H}}^*$ is itself is bounded both ways from Eqs. (G9) and (G10):

$$\int_{q_0}^{q^*} y \, dq \leq \tilde{\Delta}E_{\mathcal{H}}^* \leq \int_{q_0}^{q^*} y \, dq + \theta(q_0 - q^*). \quad (\text{G13})$$

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

$$F(q) = \langle E \rangle(q) - \frac{S(q)}{\beta} = q\omega_S + \frac{1}{\beta} [q \log(q) + (1-q) \log(1-q)]. \quad (\text{G14})$$

$$\frac{\partial F}{\partial q} = \omega_S + \frac{1}{\beta} \log\left(\frac{q}{1-q}\right) = \left(\omega_S + \frac{1}{\beta} (-\beta\omega_S - (\beta - \beta_H)y)\right) = -\frac{\beta - \beta_H}{\beta} y. \quad (\text{G15})$$

Using the above expression, the definite integral in Eq. (G11) amounts to

$$I = \frac{1}{\eta} [F(q^*) - F(q_0)] =: \frac{1}{\eta} (F^* - F_0), \quad (\text{G16})$$

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

$$\frac{1}{\eta} (F^* - F_0) \leq \tilde{\Delta}E_{\mathcal{H}}^* \leq \frac{1}{\eta} (F^* - F_0) + \theta(q_0 - q^*) \leq \frac{1}{\eta} (F^* - F_0) + \frac{\omega_S}{N} \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right), \quad (\text{G17})$$

where the inequality in the second line 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 [38, 58] (in particular, see Appendix G in Ref. [38]), we have that in each stage

$$\frac{p_n - q_n}{p_{n-1} - q_n} = \left(1 - N_v^{(n)}\right)^{m_n}. \quad (\text{G18})$$

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

$$\frac{p_n - q_n}{p_{n-1} - q_n} \leq \delta. \quad (\text{G19})$$

From this, we can prove that

$$p_n - q_n \leq \delta^n q_0 - \delta q_n + (1 - \delta) \delta \sum_{j=1}^{n-1} \delta^{n-j-1} q_j. \quad (\text{G20})$$

The proof is by induction. For $n = 0$, $p_0 = q_0$ (initial state), and for $n = 1$, using Eq. (G19)

$$\begin{aligned} p_1 - q_1 &\leq \delta(p_0 - q_1) \\ &= \delta(q_0 - q_1). \end{aligned} \quad (\text{G21})$$

Suppose that the above statement holds true for p_k . Then from Eq. (G19)

$$\begin{aligned} p_{k+1} - q_{k+1} &\leq \delta(p_k - q_{k+1}) \\ &= \delta(p_k - q_k + q_k - q_{k+1}) \\ &\vdots \\ &\leq \delta^{k+1} q_0 - \delta q_{k+1} (1 - \delta) \delta + \sum_{j=1}^{(k+1)-1} \delta^{(k+1)-j-1} q_j. \end{aligned} \quad (\text{G22})$$

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

$$\tilde{\Delta}E_{\mathcal{H}} - \tilde{\Delta}E_{\mathcal{H}}^* = \sum_{n=1}^N n\theta(p_{n-1} - p_n) - \sum_{n=1}^N n\theta(q_{n-1} - q_n) \leq \theta \left(q_0 \sum_{j=1}^{N-1} \delta^{N-j} - \sum_{j=1}^{N-1} \delta^{N-j} q_j \right). \quad (\text{G23})$$

We again use proof by induction. First note that we can rewrite

$$\sum_{n=1}^N n\theta(f_{n-1} - f_n) = \theta \left(\sum_{n=1}^N f_{n-1} \right) - N\theta f_N, \quad (\text{G24})$$

for $f_n \in \{p_n, q_n\}$. Therefore, we can rewrite the difference

$$\tilde{\Delta}E_{\mathcal{H}} - \tilde{\Delta}E_{\mathcal{H}}^* = \theta \sum_{n=1}^N (p_{n-1} - q_{n-1}) - N\theta(p_N - q_N) \leq \theta \left(\sum_{n=1}^N (p_{n-1} - q_{n-1}) \right), \quad (\text{G25})$$

since the last subtracted term is always strictly positive. Consider now the partial sum

$$\mathcal{E}_k = \sum_{n=1}^k (p_{n-1} - q_{n-1}). \quad (\text{G26})$$

For $k = 1$, $\mathcal{E}_1 = 0$, since $p_0 = q_0$. For $k = 2$, we have

$$\mathcal{E}_1 = (p_1 - q_1) \leq \delta(q_0 - q_1) = \left(q_0 \sum_{j=1}^1 \delta^{2-j} - \sum_{j=1}^1 \delta^{2-j} q_j \right), \quad (\text{G27})$$

which matches the hypothesis of Eq. (G23). Assuming that the same holds true for $\mathcal{E}_k$, then for $\mathcal{E}_{k+1}$, we have

$$\begin{aligned} \mathcal{E}_{k+1} &= \mathcal{E}_k + (p_k - q_k) \\ &\leq \left(q_0 \sum_{j=1}^{k-1} \delta^{k-j} - \sum_{j=1}^{k-1} \delta^{k-j} q_j \right) + \left(\delta^k q_0 + (1 - \delta) \delta \sum_{j=1}^{k-1} \delta^{k-j-1} q_j - \delta q_k \right) \\ &\vdots \\ &= q_0 \sum_{j=1}^k \delta^{k+1-j} - \sum_{j=1}^k \delta^{k+1-j} q_j. \end{aligned} \quad (\text{G28})$$

Then, by dropping the second sum, which is a strictly positive quantity, the difference in Eq. (G23) can be further simplified to

$$\tilde{\Delta} E_{\mathcal{H}} - \tilde{\Delta} E_{\mathcal{H}}^* \leq \theta q_0 \sum_{j=1}^{N-1} \delta^{N-j} = \theta q_0 \delta \sum_{k=0}^{N-2} \delta^k < \theta q_0 \delta (N-1) < \theta q_0 \delta N < \omega_S \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right) \delta, \quad (\text{G29})$$

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:

$$\begin{aligned} N_V^{(n)} &= \langle 01 | \tau_c(\beta, \omega_S + n\theta) \otimes \tau_{\mathcal{H}}(\beta_H, n\theta) | 01 \rangle + \langle 10 | \tau_c(\beta, \omega_S + n\theta) \otimes \tau_{\mathcal{H}}(\beta_H, n\theta) | 10 \rangle \\ &= \frac{e^{-\beta_H n\theta} + e^{-\beta(\omega_S + n\theta)}}{(1 + e^{-\beta_H n\theta})(1 + e^{-\beta(\omega_S + n\theta)})} \\ &> \frac{e^{-\beta(\omega_S + n\theta)}}{4}. \\ \Rightarrow \log \left[1 - N_V^{(n)} \right] &< \log \left[1 - \frac{e^{-\beta(\omega_S + n\theta)}}{4} \right] \\ &< -\frac{e^{-\beta(\omega_S + n\theta)}}{4} \quad \text{if } x \in (0, 1) \Rightarrow \log(1 - x) < -x. \\ \Rightarrow -\frac{1}{\log \left[1 - N_V^{(n)} \right]} &< 4e^{+\beta(\omega_S + n\theta)} \end{aligned} \quad (\text{G30})$$

Thus we can bound the number of repetitions in each stage from Eq. (G4). Noting that $\log(\delta) < 0$, we have

$$m_n < 4 \log(1/\delta) e^{+\beta(\omega_S + n\theta)} + 1. \quad (\text{G31})$$

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

$$M = \sum_{n=1}^N m_n < N \left[4 \log(1/\delta) e^{+\beta(\omega_S + N\theta)} + 1 \right] = N \left[4 \log(1/\delta) e^{\omega_S(\beta^* - \beta_H)/\eta} + 1 \right]. \quad (\text{G32})$$

Also, note that $\lim_{\delta \rightarrow 0} p_N = q_N = q^*$. More precisely, using Eq. (G20), we have

$$p_N - q^* < \delta \left(\delta^{N-1} q_0 + (1 - \delta) \sum_{j=1}^{N-1} \delta^{N-j-1} q_j - q_N \right)$$

$$< \delta (1 + (1 - \delta)(N - 1)) < \delta N. \quad (\text{G33})$$

In summary, we have the following bounds on the protocol in which each stage consists of a finite number of steps

$$\begin{aligned} p_N - q^* &< \delta N \\ \tilde{\Delta}E_{\mathcal{H}} &< \frac{1}{\eta} (F^* - F_0) + \omega_s \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right) \left(\frac{1}{N} + \delta \right), \end{aligned} \quad (\text{G34})$$

where we have combined Eqs. (G17) and (G29) for the second expression. For simplicity, we choose $\delta = 1/N^2$, so that

$$\begin{aligned} p_N - q^* &< \frac{1}{N} \\ \tilde{\Delta}E_{\mathcal{H}} &< \frac{1}{\eta} (F^* - F_0) + \omega_s \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right) \left(\frac{2}{N} \right). \end{aligned} \quad (\text{G35})$$

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

$$N = \left\lceil \max \left\{ \epsilon_1^{-1}, 2\omega_s \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right) \epsilon_2^{-1} \right\} \right\rceil, \quad (\text{G36})$$

we automatically have that $p_N - q^* < \epsilon_1$ and $\Delta E_{\mathcal{H}} < (F^* - F_0)/\eta + \epsilon_2$. The total number of unitary operations (each of which is followed by rethermalisation of the machine) is then bounded by Eq. (G32)

$$M < N \left(8 \log[N] e^{\omega_s(\beta^* - \beta_H)/\eta} + 1 \right). \quad (\text{G37})$$

We can see from Theorem 9 that the protocol is asymptotically optimal with respect to the energy extracted from the hot bath. $\square$

G2. 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. (F3)] 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.,

$$\beta_n = \beta + n\theta (\beta - \beta_{\mathcal{H}}), \quad (\text{G38})$$

$$\theta = \frac{1}{N} \left(\frac{\beta^* - \beta}{\beta - \beta_{\mathcal{H}}} \right), \quad (\text{G39})$$

so that $\beta_N = \beta^*$ by construction. This corresponds to taking the system through the following sequence of thermal states

$$\varrho_s^{(n)} = \frac{e^{-\beta_n H_S}}{\mathcal{Z}_S(H_S, \beta_n)}. \quad (\text{G40})$$

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:

$$G^{(n)} = -n\theta \Delta E_S^{(n)} = -n\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n)} - \varrho_S^{(n-1)} \right) \right]. \quad (\text{G41})$$

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

$$\begin{aligned} \tilde{\Delta}E_{\mathcal{H}}^* &= \sum_{n=1}^N G^{(n)} \\ &= \sum_{n=1}^N n\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n-1)} - \varrho_S^{(n)} \right) \right] \end{aligned} \quad (\text{G42})$$

$$= \left\{ \sum_{n=1}^N (n-1)\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n-1)} - \varrho_S^{(n)} \right) \right] \right\} + \theta \operatorname{tr} \left[H_S \left(\varrho_S^{(0)} - \varrho_S^{(N)} \right) \right]. \quad (\text{G43})$$

The sums on the second and third lines above, (G42) and (G43) respectively, are the *right* and *left* Riemann sums corresponding to the following integral:

$$\begin{aligned} I &= \int_{q_i}^{q_f} q (-dx) = \int_{q_f}^{q_i} q dx, \\ \text{where } n\theta &\rightarrow q, \\ x &= \operatorname{tr} [H_S \varrho_S(q)], \\ \varrho_S(q) &= \frac{e^{-[\beta+q(\beta-\beta_{\mathcal{H}})]H_S}}{\operatorname{tr} [e^{-[\beta+q(\beta-\beta_{\mathcal{H}})]H_S}]}. \end{aligned} \quad (\text{G44})$$

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

$$\sum_{n=1}^N (n-1)\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n-1)} - \varrho_S^{(n)} \right) \right] \leq \int_{q_f}^{q_i} q dx \leq \sum_{n=1}^N n\theta \operatorname{tr} \left[H_S \left(\varrho_S^{(n-1)} - \varrho_S^{(n)} \right) \right]. \quad (\text{G45})$$

We can thus bound the idealised heat cost in both directions via

$$I \leq \tilde{\Delta}E_{\mathcal{H}}^* \leq I + \theta \operatorname{tr} \left[H_S \left(\varrho_S^{(0)} - \varrho_S^{(N)} \right) \right]. \quad (\text{G46})$$

The integral in Eq. (G44) can be shown to be equal to the change in free energy of the target system (w.r.t. inverse temperature β)

$$\begin{aligned} F_{\beta}[\varrho_S(q)] &= \operatorname{tr} [H_S \varrho_S(q)] + \frac{1}{\beta} \operatorname{tr} [\varrho_S(q) \log \varrho_S(q)], \\ \frac{d}{dq} F_{\beta}[\varrho_S(q)] &= \operatorname{tr} \left[\left(H_S + \frac{\mathbb{1}_S + \log \varrho_S(q)}{\beta} \right) \frac{d\varrho_S(q)}{dq} \right]. \end{aligned} \quad (\text{G47})$$

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 $\operatorname{tr} [d\varrho_S(q)] = d \operatorname{tr} [\varrho_S(q)] = 0$ and label the partition function $\mathcal{Z}(q) := \operatorname{tr} [e^{-[\beta+q(\beta-\beta_{\mathcal{H}})]H_S}]$ to obtain

$$\begin{aligned} \frac{d}{dq} F_{\beta}[\varrho_S(q)] &= \operatorname{tr} \left[\left(H_S + \frac{\log \varrho_S(q)}{\beta} \right) \frac{d\varrho_S(q)}{dq} \right] \\ &= \operatorname{tr} \left[\left(H_S - \frac{\beta + q(\beta - \beta_{\mathcal{H}})}{\beta} H_S - \mathbb{1}_S \frac{\log \mathcal{Z}(q)}{\beta} \right) \frac{d\varrho_S(q)}{dq} \right] \end{aligned}$$

$$= -q \left(1 - \frac{\beta_{\mathcal{H}}}{\beta}\right) \frac{d}{dq} \text{tr} [H_s \varrho_s(q)] = -q\eta \frac{dx}{dq}, \quad (\text{G48})$$

where we have identified the Carnot efficiency η for an engine operating between β and $\beta_{\mathcal{H}}$. The integral thus simplifies to

$$I = \eta^{-1} (F_{\beta}[\varrho_s(q_f)] - F_{\beta}[\varrho_s(q_i)]) =: \eta^{-1} \Delta F_s^{(\beta)}. \quad (\text{G49})$$

The idealised heat cost is thus bounded by

$$\eta^{-1} \Delta F_s^{(\beta)} \leq \tilde{\Delta} E_{\mathcal{H}}^* \leq \eta^{-1} \Delta F_s^{(\beta)} + \theta \text{tr} \left[\mathcal{H}_s \left(\varrho_s^{(0)} - \varrho_s^{(N)} \right) \right]. \quad (\text{G50})$$

The left inequality is Landauer's bound applied to cooling a target system with Hamiltonian H_s (see Theorem 4), and the error term on the right can be bounded quite easily; for instance, for $\beta > 0$, we have

$$\begin{aligned} \text{tr} \left[\mathcal{H}_s \left(\varrho_s^{(0)} - \varrho_s^{(N)} \right) \right] &= \text{tr} \left[\left(\mathcal{H}_s - E_s^{\min} \mathbb{1}_s \right) \left(\varrho_s^{(0)} - \varrho_s^{(N)} \right) \right] \\ &\leq \text{tr} \left[\left(\mathcal{H}_s - E_s^{\min} \mathbb{1}_s \right) \varrho_s^{(0)} \right] && \text{since } \mathcal{H}_s - E_s^{\min} \mathbb{1}_s \text{ is a positive operator,} \\ &\leq \text{tr} \left[\left(\mathcal{H}_s - E_s^{\min} \mathbb{1}_s \right) \frac{\mathbb{1}_s}{d_s} \right] \leq \frac{\omega_s^{\max}}{d_s}, \end{aligned} \quad (\text{G51})$$

where $\omega_s^{\max} := E_s^{\max} - E_s^{\min}$ is the largest energy gap in the target system Hamiltonian and d_s is the system dimension. 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, $\mathbb{1}_s/d_s$. Since $\theta \propto 1/N$, 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

$$H_c = (1 + n\theta) H_s \quad (\text{G52})$$

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 $d_s(d_s - 1)/2$ 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)

$$H_{\mathcal{H}}^{(ij)} = [\omega_i + (1 + n\theta)\omega_j - (\omega_j + (1 + n\theta)\omega_i)] |1\rangle\langle 1|_{\mathcal{H}}^{(ij)} = n\theta(\omega_j - \omega_i) |1\rangle\langle 1|_{\mathcal{H}}^{(ij)}, \quad (\text{G53})$$

where we have labelled the energy eigenvalues of H_s by $\{\omega_i\}$. Each of these hot qubits begins at inverse temperature $\beta_{\mathcal{H}}$. 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

$$\Delta^{(ij)} = \omega_j(1 + n\theta) - \omega_i(1 + n\theta) - n\theta(\omega_j - \omega_i) = \omega_j - \omega_i, \quad (\text{G54})$$

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

$$g^{(ij)} = \frac{e^{-\beta\omega_j(1+n\theta)}}{e^{-\beta\omega_i(1+n\theta)} e^{-\beta_{\mathcal{H}}n\theta(\omega_j-\omega_i)}} = e^{-(\omega_j-\omega_i)(\beta+n\theta(\beta-\beta_{\mathcal{H}}))}. \quad (\text{G55})$$

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. (G38)]. In summary, we have constructed a machine featuring $d_s(d_s - 1)/2$ 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. Within each stage of the protocol, a single step consists of a unitary operation on $SC\mathcal{H}$, 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

$$\Delta E_s = \text{tr} [H_s (\varrho'_s - \varrho_s)] = -\delta p (\omega_j - \omega_i). \quad (\text{G56})$$

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

$$\tilde{\Delta} E_{\mathcal{H}} = \delta p n \theta (\omega_j - \omega_i) = -n \theta \Delta E_s. \quad (\text{G57})$$

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

$$D(\varrho||\tau) = \sum_k p_k \log \left(\frac{p_k}{q_k} \right). \quad (\text{G58})$$

We now focus on a single qubit subspace labelled by $\{i, j\}$, which leads to

$$\begin{aligned} D(\varrho||\tau) &= p_i \log \left(\frac{p_i}{q_i} \right) + p_j \log \left(\frac{p_j}{q_j} \right) + \sum_{k \notin \{i, j\}} p_k \log \left(\frac{p_k}{q_k} \right) \\ &= (p_i + p_j) \left[\frac{p_i}{p_i + p_j} \log \left(\frac{\frac{p_i}{q_i}}{\frac{p_i + p_j}{q_i + q_j}} \right) + \frac{p_j}{p_i + p_j} \log \left(\frac{\frac{p_j}{q_j}}{\frac{p_i + p_j}{q_i + q_j}} \right) \right] + \sum_{k \notin \{i, j\}} p_k \log \left(\frac{p_k}{q_k} \right) \\ &= N \left(\bar{p}_i \log \frac{\bar{p}_i}{\bar{q}_i} + \bar{p}_j \log \frac{\bar{p}_j}{\bar{q}_j} + \log \frac{N}{N_v} \right) + \sum_{k \notin \{i, j\}} p_k \log \frac{p_k}{q_k}. \end{aligned} \quad (\text{G59})$$

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

$$D(\varrho||\tau) = N \left[D(\varrho_v||\tau_v) + \log \left(\frac{N}{N_v} \right) \right] + \sum_{k \notin \{i, j\}} p_k \log \left(\frac{p_k}{q_k} \right). \quad (\text{G60})$$

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 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 [38, 58], 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

$$\begin{aligned} \left| \tilde{\Delta}E_{\mathcal{H}} - \tilde{\Delta}E_{\mathcal{H}}^* \right| &= \left| \sum_{n=1}^N \left\{ -n\theta \operatorname{tr} \left[H_s \left(\tilde{\varrho}_s^{(n)} - \tilde{\varrho}_s^{(n-1)} \right) \right] \right\} - \sum_{n=1}^N \left\{ -n\theta \operatorname{tr} \left[H_s \left(\varrho_s^{(n)} - \varrho_s^{(n-1)} \right) \right] \right\} \right| \\ &= \left| \sum_{n=0}^{N-1} \theta \operatorname{tr} \left[H_s \left(\tilde{\varrho}_s^{(n)} - \varrho_s^{(n)} \right) \right] - N\theta \left(\tilde{\varrho}_s^{(N)} - \varrho_s^{(N)} \right) \right| \\ &\leq 2N\theta\delta_E = 2 \left(\frac{\beta^* - \beta}{\beta - \beta_{\mathcal{H}}} \right) \delta_E. \end{aligned} \quad (\text{G61})$$

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

$$|\beta' - \beta^*| \leq \epsilon_\beta, \quad (\text{G62})$$

$$\left| \tilde{\Delta}E_{\mathcal{H}} - \eta^{-1} \Delta F_s^{(\beta)} \right| \leq \epsilon_E. \quad (\text{G63})$$

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. $\square$

Appendix H: 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.

H1. 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 G1, we derived the deviation from the idealised heat dissipation in the incoherent

control setting for a sequence of N machines [see Eq. (G17)], which we repeat below:

$$\frac{1}{\eta} (F^* - F_0) \leq \tilde{\Delta} E_{\mathcal{H}}^* \leq \frac{1}{\eta} (F^* - F_0) + \frac{\omega_s}{N} \left(\frac{\beta^* - \beta}{\beta - \beta_H} \right). \quad (\text{H1})$$

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

$$\Delta F_s \leq W \leq \Delta F_s + \frac{\omega_s}{N} \left(\frac{\beta^*}{\beta} - 1 \right). \quad (\text{H2})$$

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

$$\left[\frac{W - \Delta F}{\omega_s} \right] N \sim \frac{1}{2} \left(\frac{\beta^*}{\beta} - 1 \right). \quad (\text{H3})$$

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 $\beta^*/\beta - 1$, which, for cold enough final temperatures, is approximately the ratio β^*/β .

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:

$$\left[\frac{\eta \Delta E_{\mathcal{H}} - \Delta F}{\omega_s} \right] N \sim \frac{1}{2} \left(\frac{\beta^*}{\beta} - 1 \right), \quad (\text{H4})$$

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.

H2. 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., $|1\rangle_s|0\rangle_c|0\rangle_h \leftrightarrow |0\rangle_s|1\rangle_c|1\rangle_h$. 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 G, 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 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 N/m different stages with m steps (repetitions) within each 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. 4.

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. (H3) and (H4)], 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.




FIG. 4. *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 $W - \beta \Delta S_s$ (with $\beta = 1$) is plotted (in units of the smallest machine energy gap, $\omega_{\mathcal{M}}^{\min}$) 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.

Chapter V

Conclusion

This thesis represents a collection of the scientific contributions I made during my time as a PhD student in the form of a cumulative thesis. In my research I investigated the role that control plays in the achievability of a variety of tasks in physics, as well as elucidating the notion of control itself. As control is a notion that is used in a wide array of areas in physics, and science in general, investigating it with too narrow a focus can only lead to highly specialised notions and understanding that bares no fundamental relevance. To avoid such pitfalls in my investigations I incorporated a diverse set of tasks and settings that elucidate different aspects of control which together paint a richer picture than each of them alone. To that end I distinguishing between two fundamentally different notions of control, *i.e.* *temporal control* and *structural control*, and analyse their interplay. In Chapter III I establish the intricate connection between these two notions of control by analysing the fundamental limitations of clocks as constrained by the structural control that can be exerted over them, thus establishing the resources that are fundamentally required to maintain temporal control. Going beyond this, I establish the fundamental limitations clocks have to abide by, irrespective of the amount of structural control, thus presenting limitations to the amount of temporal control that are in principle achievable. Restrictions on our ability to keep track of time furthermore also impact our ability to carry out other tasks, as the implementation of many operations requires precise timing [86]. While it is hard to derive a general quantifier for control, which is why we have to rely on proxies for the most part, narrowing down the number of potential candidates which might qualify as good complexity quantifiers is a challenge that becomes more tractable if we focus on specific tasks. In Chapter IV we focus on the paradigmatic task of ground-state cooling. Despite this narrow focus on only one specific task, this setting allows us to draw rather general conclusions about specific properties any notion of control has to fulfill due to the thermodynamic and resource-theoretic value that comes with zero-temperature states. By identifying *control*

complexity as a quantity whose divergence enables ground-state cooling in finite time and with finite energy cost we not only expand and make more precise expand the third law of thermodynamics, but also highlight the long-neglected role of control as a resource that fundamentally alters the achievability of certain tasks. While the achievability of many tasks relies on our ability to control physical systems to a high degree, some physical processes, such as spontaneous processes, rely on descriptions that fundamentally exclude the ability of an agent to exert control. Finally, in Chapter II we analyse the role that control plays and is allowed to play when modelling the process of measurements. By formulating the Measurement-Equilibration Hypothesis, stating that measurements need to feature an equilibration process, we are able to develop a first version of a thermodynamically consistent model of measurements, while making realistic assumptions on the amount of control that can be exerted throughout this process.

While this thesis represents a first step towards a deeper understanding of the intricate relationship between the laws of thermodynamics, control and the ability of agents to obtain information about the world, many questions remain. As science is an endeavour that is fundamentally social and relies on discourse and exchange of ideas of a diverse set of individuals it is my hope that this thesis, and in particular the individual research topics that it touches on and develops in each of its constituting publications further this endeavour by inspiring fruitful discussion, providing new perspectives and sparking unique ideas.

Bibliography

- [1] Emanuel Schwarzhans, Maximilian P. E. Lock, Paul Erker, Nicolai Friis, and Marcus Huber. “Autonomous Temporal Probability Concentration: Clockworks and the Second Law of Thermodynamics”. In: *arXiv:2007.01307 [cond-mat, physics:quant-ph]* (2020-07). DOI: [10.1103/PhysRevX.11.011046](https://doi.org/10.1103/PhysRevX.11.011046). arXiv: [2007.01307 \[cond-mat, physics:quant-ph\]](https://arxiv.org/abs/2007.01307).
- [2] Flavio Del Santo and Emanuel Schwarzhans. ““Philosophysics” at the University of Vienna: The (Pre-)History of Foundations of Quantum Physics in the Viennese Cultural Context”. In: *Physics in Perspective* 24.2-3 (2022-10), pp. 125–153. ISSN: 1422-6944, 1422-6960. DOI: [10.1007/s00016-022-00290-y](https://doi.org/10.1007/s00016-022-00290-y). arXiv: [2110.05217 \[physics, physics:quant-ph\]](https://arxiv.org/abs/2110.05217).
- [3] Philip Taranto et al. “Landauer vs. Nernst: What Is the True Cost of Cooling a Quantum System?” In: *arXiv:2106.05151 [quant-ph]* (2021-06). arXiv: [2106.05151 \[quant-ph\]](https://arxiv.org/abs/2106.05151).
- [4] Florian Meier, Emanuel Schwarzhans, Paul Erker, and Marcus Huber. “Fundamental Accuracy-Resolution Trade-off for Timekeeping Devices”. In: *arXiv:2301.05173* (2023-01). arXiv: [2301.05173 \[quant-ph\]](https://arxiv.org/abs/2301.05173).
- [5] Emanuel Schwarzhans, Felix Binder, Marcus Huber, and Maximilian P. E. Lock. “Quantum Measurements and Equilibration: The Emergence of Objective Reality via Entropy Maximisation”. In preperation, manuscript not yet submitted for publication.
- [6] Adán Cabello. “Interpretations of Quantum Theory: A Map of Madness”. In: (2015-09). DOI: [10.1017/9781316494233.009](https://doi.org/10.1017/9781316494233.009). URL: <http://arxiv.org/abs/1509.04711><http://dx.doi.org/10.1017/9781316494233.009>.
- [7] J. R. Partington. *A Short History of Chemistry*. 3rd ed., rev. and enl. Dover Science Books. New York: Dover Publications, 1989. ISBN: 978-0-486-65977-0.
- [8] Sadi Carnot, Robert Henry Thurston, Robert Henry Carnot, and William Thomson Kelvin. *Reflections on the Motive Power of Heat [Microform] : And on Machines Fitted to Develop That Power*. New York : J. Wiley, 1890.

- [9] John Goold, Marcus Huber, Arnau Riera, L dia del Rio, and Paul Skrzypczyk. “The role of quantum information in thermodynamics — a topical review”. In: *J. Phys. A* 49 (2016), p. 143001. DOI: [10.1088/1751-8113/49/14/143001](https://doi.org/10.1088/1751-8113/49/14/143001). arXiv: [1505.07835](https://arxiv.org/abs/1505.07835).
- [10] Ronnie Kosloff. “Quantum Thermodynamics: A Dynamical Viewpoint”. In: *Entropy* 15.6 (2013-06), pp. 2100–2128. ISSN: 1099-4300. DOI: [10.3390/e15062100](https://doi.org/10.3390/e15062100).
- [11] Ronnie Kosloff and Amikam Levy. “Quantum Heat Engines and Refrigerators: Continuous Devices”. In: *Annual Review of Physical Chemistry* 65.1 (2014), pp. 365–393. DOI: [10.1146/annurev-physchem-040513-103724](https://doi.org/10.1146/annurev-physchem-040513-103724).
- [12] David Gelbwaser-Klimovsky, Wolfgang Niedenzu, and Gershon Kurizki. “Chapter Twelve - Thermodynamics of Quantum Systems Under Dynamical Control”. In: *Advances In Atomic, Molecular, and Optical Physics*. Ed. by Ennio Arimondo, Chun C. Lin, and Susanne F. Yelin. Vol. 64. Academic Press, 2015-01, pp. 329–407. DOI: [10.1016/bs.aamop.2015.07.002](https://doi.org/10.1016/bs.aamop.2015.07.002).
- [13] Luigi Amico, Rosario Fazio, Andreas Osterloh, and Vlatko Vedral. “Entanglement in Many-Body Systems”. In: *Reviews of Modern Physics* 80.2 (2008-05), pp. 517–576. DOI: [10.1103/RevModPhys.80.517](https://doi.org/10.1103/RevModPhys.80.517).
- [14] Tobias Fritz. “Resource Convertibility and Ordered Commutative Monoids”. In: *Mathematical Structures in Computer Science* 27.6 (2017-09), pp. 850–938. ISSN: 0960-1295, 1469-8072. DOI: [10.1017/S0960129515000444](https://doi.org/10.1017/S0960129515000444).
- [15] Wassim M. Haddad. “Thermodynamics: The Unique Universal Science”. In: *Entropy* 19.11 (2017-11), p. 621. ISSN: 1099-4300. DOI: [10.3390/e19110621](https://doi.org/10.3390/e19110621).
- [16] Jochen Gemmer. *Quantum Thermodynamics : Emergence of Thermodynamic Behavior within Composite Quantum Systems*. Springer, 2012. ISBN: 978-3-642-26099-5.
- [17] C. Gogolin and J. Eisert. “Equilibration, Thermalisation, and the Emergence of Statistical Mechanics in Closed Quantum Systems”. In: *Reports on Progress in Physics* 79.5 (2016-05), p. 056001. ISSN: 0034-4885, 1361-6633. DOI: [10.1088/0034-4885/79/5/056001](https://doi.org/10.1088/0034-4885/79/5/056001).
- [18] Anatoli Polkovnikov, Krishnendu Sengupta, Alessandro Silva, and Mukund Vengalattore. “Colloquium: Nonequilibrium Dynamics of Closed Interacting Quantum Systems”. In: *Reviews of Modern Physics* 83.3 (2011-08), pp. 863–883. DOI: [10.1103/RevModPhys.83.863](https://doi.org/10.1103/RevModPhys.83.863).

- [19] A. J. Daley, M. Rigol, and D. S. Weiss. “Focus on Out-of-Equilibrium Dynamics in Strongly Interacting One-Dimensional Systems”. In: *New Journal of Physics* 16.9 (2014-09), p. 095006. ISSN: 1367-2630. DOI: [10.1088/1367-2630/16/9/095006](https://doi.org/10.1088/1367-2630/16/9/095006).
- [20] H. E. D. Scovil and E. O. Schulz-DuBois. “Three-Level Masers as Heat Engines”. In: *Physical Review Letters* 2.6 (1959-03), pp. 262–263. DOI: [10.1103/PhysRevLett.2.262](https://doi.org/10.1103/PhysRevLett.2.262).
- [21] J. E. Geusic, E. O. Schulz-DuBios, and H. E. D. Scovil. “Quantum Equivalent of the Carnot Cycle”. In: *Physical Review* 156.2 (1967-04), pp. 343–351. DOI: [10.1103/PhysRev.156.343](https://doi.org/10.1103/PhysRev.156.343).
- [22] R. Alicki. “The Quantum Open System as a Model of the Heat Engine”. In: *Journal of Physics A: Mathematical and General* 12.5 (1979-05), p. L103. ISSN: 0305-4470. DOI: [10.1088/0305-4470/12/5/007](https://doi.org/10.1088/0305-4470/12/5/007).
- [23] Engelbert Broda. *The interaction of Boltzmann with Mach, Ostwald and Planck, and his influence on Nernst and Einstein*. Tech. rep. Bucharest: Institute of Physical Chemistry Vienna University, 1981.
- [24] Harvey Leff and Andrew F. Rex. *Maxwell’s Demon 2 Entropy, Classical and Quantum Information, Computing*. 2002. ISBN: 9780750307598.
- [25] L. Szilard. “über die Entropieverminderung in einem thermodynamischen System bei Eingriffen intelligenter Wesen”. In: *Zeitschrift für Physik* 53.11 (1929-11), pp. 840–856. ISSN: 0044-3328. DOI: [10.1007/BF01341281](https://doi.org/10.1007/BF01341281).
- [26] R. Landauer. “Irreversibility and Heat Generation in the Computing Process”. In: *IBM Journal of Research and Development* 5.3 (1961-07), pp. 183–191. ISSN: 0018-8646. DOI: [10.1147/rd.53.0183](https://doi.org/10.1147/rd.53.0183).
- [27] Meir Hemmo and Orly R. Shenker. *Maxwell’s Demon in Quantum Mechanics*. Entropy. 2020-02-05. URL: <http://philsci-archive.pitt.edu/19566/> (visited on 2022-12-30).
- [28] Charles H. Bennett. “The Thermodynamics of Computation—a Review”. In: *International Journal of Theoretical Physics* 21.12 (1982-12), pp. 905–940. ISSN: 1572-9575. DOI: [10.1007/BF02084158](https://doi.org/10.1007/BF02084158).

- [29] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. 1st ed. Wiley, 2005-09-16. ISBN: 978-0-471-24195-9 978-0-471-74882-3. DOI: [10.1002/047174882X](https://doi.org/10.1002/047174882X). URL: <https://onlinelibrary.wiley.com/doi/book/10.1002/047174882X> (visited on 2022-12-31).
- [30] Rolf Landauer. “Information Is Physical”. In: *Physics Today* 44.5 (1991-05), pp. 23–29. ISSN: 0031-9228. DOI: [10.1063/1.881299](https://doi.org/10.1063/1.881299).
- [31] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. 2010. ISBN: 9781107002173.
- [32] Anthony J Short. “Equilibration of Quantum Systems and Subsystems”. In: *New Journal of Physics* 13.5 (2011-05), p. 053009. ISSN: 1367-2630. DOI: [10.1088/1367-2630/13/5/053009](https://doi.org/10.1088/1367-2630/13/5/053009).
- [33] Anthony J Short and Terence C Farrelly. “Quantum Equilibration in Finite Time”. In: *New Journal of Physics* 14.1 (2012-01), p. 013063. ISSN: 1367-2630. DOI: [10.1088/1367-2630/14/1/013063](https://doi.org/10.1088/1367-2630/14/1/013063).
- [34] Pedro Figueroa-Romero. “Equilibration and Typicality in Quantum Processes”. In: *arXiv:2102.02289 [quant-ph]* (2021), 1773554 Bytes. DOI: [10.26180/13684795.v1](https://doi.org/10.26180/13684795.v1). arXiv: 2102.02289 [quant-ph].
- [35] Yelena Guryanova, Nicolai Friis, and Marcus Huber. “Ideal Projective Measurements Have Infinite Resource Costs”. In: *Quantum* 4 (2020), p. 222. DOI: [10.22331/q-2020-01-13-222](https://doi.org/10.22331/q-2020-01-13-222). arXiv: 1805.11899. URL: <https://doi.org/10.22331/q-2020-01-13-222>.
- [36] P. W. Bridgman. *The Logic of Modern Physics*. Three Centuries of Science in America. New York: Arno Press, 1980. ISBN: 978-0-405-12594-2.
- [37] Anna N. Pearson, Yelena Guryanova, Paul Erker, Edward A. Laird, G. Andrew Davidson Briggs, Marcus Huber, and Natalia Ares. *Measuring the thermodynamic cost of time-keeping*. 2020. arXiv: [2006.08670](https://arxiv.org/abs/2006.08670) [quant-ph].
- [38] Samuel L. Braunstein, Carlton M. Caves, and G. J. Milburn. “Generalized Uncertainty Relations: Theory, Examples, and Lorentz Invariance”. In: *Annals of Physics* 247.1 (1996-04), pp. 135–173. ISSN: 00034916. DOI: [10.1006/aphy.1996.0040](https://doi.org/10.1006/aphy.1996.0040). arXiv: [quant-ph/9507004](https://arxiv.org/abs/quant-ph/9507004).
- [39] G. J. Milburn. “The thermodynamics of clocks”. In: *Contemp. Phys.* 61.2 (2020), pp. 69–95. DOI: [10.1080/00107514.2020.1837471](https://doi.org/10.1080/00107514.2020.1837471). arXiv: [2007.02217](https://arxiv.org/abs/2007.02217).

- [40] Paul Erker. “The Quantum Hourglass - Approaching time measurement with quantum information theory”. MA thesis. ETH Zurich, 2014. DOI: [10.3929/ethz-a-010514644](https://doi.org/10.3929/ethz-a-010514644). URL: <https://doi.org/10.3929/ethz-a-010514644>.
- [41] Carlo Rovelli, Erica Segre, and Simon Carnell. *The Order of Time*. First American edition. New York: Riverhead Books, 2018. ISBN: 978-0-7352-1610-5.
- [42] Massimiliano Esposito, Upendra Harbola, and Shaul Mukamel. “Nonequilibrium Fluctuations, Fluctuation Theorems, and Counting Statistics in Quantum Systems”. In: *Reviews of Modern Physics* 81.4 (2009-12), pp. 1665–1702. DOI: [10.1103/RevModPhys.81.1665](https://doi.org/10.1103/RevModPhys.81.1665).
- [43] Michele Campisi, Peter Hänggi, and Peter Talkner. “Colloquium. Quantum Fluctuation Relations: Foundations and Applications”. In: *Reviews of Modern Physics* 83.3 (2011-07), pp. 771–791. ISSN: 0034-6861, 1539-0756. DOI: [10.1103/RevModPhys.83.771](https://doi.org/10.1103/RevModPhys.83.771). arXiv: [1012.2268](https://arxiv.org/abs/1012.2268).
- [44] Udo Seifert. “Stochastic Thermodynamics, Fluctuation Theorems and Molecular Machines”. In: *Reports on Progress in Physics* 75.12 (2012-11), p. 126001. ISSN: 0034-4885. DOI: [10.1088/0034-4885/75/12/126001](https://doi.org/10.1088/0034-4885/75/12/126001).
- [45] D.W. Allan. “Statistics of Atomic Frequency Standards”. In: *Proceedings of the IEEE* 54.2 (1966-02), pp. 221–230. ISSN: 1558-2256. DOI: [10.1109/PROC.1966.4634](https://doi.org/10.1109/PROC.1966.4634).
- [46] David W. Allan, Neil Ashby, and Clifford C. Hodge. *The Science of Timekeeping Application Note 1289*. Packard Application Note 1289. Hewlett Packard. 1997. URL: http://www.allanstime.com/Publications/DWA/Science_Timekeeping/index.html.
- [47] William Riley and David Howe. *Handbook of Frequency Stability Analysis*. 2008-07. URL: https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=50505.
- [48] Norman Margolus and Lev B. Levitin. “The maximum speed of dynamical evolution”. In: *Physica D* 120 (1998), pp. 188–195. DOI: [10.1016/S0167-2789\(98\)00054-2](https://doi.org/10.1016/S0167-2789(98)00054-2). arXiv: [quant-ph/9710043](https://arxiv.org/abs/quant-ph/9710043). URL: [https://doi.org/10.1016/S0167-2789\(98\)00054-2](https://doi.org/10.1016/S0167-2789(98)00054-2).
- [49] L. Mandelstam and Ig. Tamm. “The Uncertainty Relation Between Energy and Time in Non-relativistic Quantum Mechanics”. In: *Selected Papers*. Ed. by Boris M. Bolotovskii, Victor Ya. Frenkel, and Rudolf Peierls. Berlin, Heidelberg: Springer Berlin Heidelberg,

- 1991, pp. 115–123. ISBN: 978-3-642-74628-4 978-3-642-74626-0. DOI: [10.1007/978-3-642-74626-0_8](https://doi.org/10.1007/978-3-642-74626-0_8).
- [50] Sebastian Deffner and Steve Campbell. “Quantum Speed Limits: From Heisenberg’s Uncertainty Principle to Optimal Quantum Control”. In: *Journal of Physics A: Mathematical and Theoretical* 50.45 (2017-11), p. 453001. ISSN: 1751-8113, 1751-8121. DOI: [10.1088/1751-8121/aa86c6](https://doi.org/10.1088/1751-8121/aa86c6). arXiv: [1705.08023](https://arxiv.org/abs/1705.08023).
- [51] Sandra Ranković, Yeong-Cherng Liang, and Renato Renner. “Quantum Clocks and Their Synchronisation - the Alternate Ticks Game”. In: *arXiv:1506.01373 [quant-ph]* (2015-06). arXiv: [1506.01373 \[quant-ph\]](https://arxiv.org/abs/1506.01373).
- [52] Paul Erker, Mark T. Mitchison, Ralph Silva, Mischa P. Woods, Nicolas Brunner, and Marcus Huber. “Autonomous Quantum Clocks: Does Thermodynamics Limit Our Ability to Measure Time?” In: *Physical Review X* 7.3 (2017-08), p. 031022. ISSN: 2160-3308. DOI: [10.1103/PhysRevX.7.031022](https://doi.org/10.1103/PhysRevX.7.031022). arXiv: [1609.06704](https://arxiv.org/abs/1609.06704).
- [53] Nicolas Brunner, Noah Linden, Sandu Popescu, and Paul Skrzypczyk. “Virtual Qubits, Virtual Temperatures, and the Foundations of Thermodynamics”. In: *Phys Rev E* 85.5 (2012), pp. 1–16. ISSN: 15393755. DOI: [10.1103/PhysRevE.85.051117](https://doi.org/10.1103/PhysRevE.85.051117).
- [54] Yuxiang Yang and Renato Renner. “Ultimate Limit on Time Signal Generation”. In: *arXiv:2004.07857 [quant-ph]* (2020-04). arXiv: [2004.07857 \[quant-ph\]](https://arxiv.org/abs/2004.07857).
- [55] Mischa P. Woods, Ralph Silva, Gilles Pütz, Sandra Stupar, and Renato Renner. *Quantum clocks are more accurate than classical ones*. 2018. arXiv: [1806.00491 \[quant-ph\]](https://arxiv.org/abs/1806.00491).
- [56] Bob Coecke, Tobias Fritz, and Robert W. Spekkens. “A Mathematical Theory of Resources”. In: *Information and Computation. Quantum Physics and Logic* 250 (2016-10), pp. 59–86. ISSN: 0890-5401. DOI: [10.1016/j.ic.2016.02.008](https://doi.org/10.1016/j.ic.2016.02.008).
- [57] E T Jaynes. “Information Theory and Statistical Mechanics”. In: *Physical Review* 106.4 (1957), pp. 620–630. ISSN: 0031899X. DOI: [10.1103/PhysRev.106.620](https://doi.org/10.1103/PhysRev.106.620).
- [58] Nicolas Gisin, Grégoire Ribordy, Wolfgang Tittel, and Hugo Zbinden. “Quantum Cryptography”. In: *Reviews of Modern Physics* 74.1 (2002-03), pp. 145–195. DOI: [10.1103/RevModPhys.74.145](https://doi.org/10.1103/RevModPhys.74.145).
- [59] S. Pirandola et al. “Advances in Quantum Cryptography”. In: *Advances in Optics and Photonics* 12.4 (2020-12), pp. 1012–1236. ISSN: 1943-8206. DOI: [10.1364/AOP.361502](https://doi.org/10.1364/AOP.361502).

- [60] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. “Advances in Quantum Metrology”. In: *Nature Photonics* 5.4 (2011-04), pp. 222–229. ISSN: 1749-4893. DOI: [10.1038/nphoton.2011.35](https://doi.org/10.1038/nphoton.2011.35).
- [61] Géza Tóth and Iagoba Apellaniz. “Quantum Metrology from a Quantum Information Science Perspective”. In: *Journal of Physics A: Mathematical and Theoretical* 47.42 (2014-10), p. 424006. ISSN: 1751-8121. DOI: [10.1088/1751-8123/47/42/424006](https://doi.org/10.1088/1751-8123/47/42/424006).
- [62] Rafal Demkowicz-Dobrzański, Marcin Jarzyna, and Jan Kołodyński. “Chapter Four - Quantum Limits in Optical Interferometry”. In: *Progress in Optics*. Ed. by E. Wolf. Vol. 60. Elsevier, 2015-01, pp. 345–435. DOI: [10.1016/bs.po.2015.02.003](https://doi.org/10.1016/bs.po.2015.02.003).
- [63] John Preskill. “Fault-Tolerant Quantum Computation”. In: *Introduction to Quantum Computation and Information*. WORLD SCIENTIFIC, 1998-10, pp. 213–269. ISBN: 978-981-02-3399-0. DOI: [10.1142/9789812385253_0008](https://doi.org/10.1142/9789812385253_0008).
- [64] John Preskill. “Quantum Computing in the NISQ Era and Beyond”. In: *Quantum* 2 (2018-08), p. 79. DOI: [10.22331/q-2018-08-06-79](https://doi.org/10.22331/q-2018-08-06-79).
- [65] Walter Nernst. “Über Die Beziehungen Zwischen Wärmeentwicklung Und Maximaler Arbeit Bei Kondensierten Systemen.” In: *Sitzungsberichte Der Königlich Preussischen Akademie Der Wissenschaften*. Berlin, 1906, pp. 933–940.
- [66] Johan Åberg. “Truly Work-like Work Extraction via a Single-Shot Analysis”. In: *Nature Communications* 4.1 (2013-06), p. 1925. ISSN: 2041-1723. DOI: [10.1038/ncomms2712](https://doi.org/10.1038/ncomms2712).
- [67] Paul Skrzypczyk, Anthony J. Short, and Sandu Popescu. “Work Extraction and Thermodynamics for Individual Quantum Systems”. In: *Nature Communications* 5.1 (2014-09), p. 4185. ISSN: 2041-1723. DOI: [10.1038/ncomms5185](https://doi.org/10.1038/ncomms5185).
- [68] Matteo Lostaglio, David Jennings, and Terry Rudolph. “Description of Quantum Coherence in Thermodynamic Processes Requires Constraints beyond Free Energy”. In: *Nature Communications* 6.1 (2015-12), p. 6383. ISSN: 2041-1723. DOI: [10.1038/ncomms7383](https://doi.org/10.1038/ncomms7383).
- [69] Francesco Campaioli, Felix A. Pollock, and Sai Vinjanampathy. *Quantum Batteries - Review Chapter*. 2018-05. arXiv: [1805.05507 \[quant-ph\]](https://arxiv.org/abs/1805.05507).
- [70] John von Neumann. *Mathematische Grundlagen der Quantenmechanik*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1996 (orig. 1930). ISBN: 978-3-642-64828-1 978-3-642-61409-5. DOI: [10.1007/978-3-642-61409-5](https://doi.org/10.1007/978-3-642-61409-5).

- [71] H. D. Zeh. “On the Interpretation of Measurement in Quantum Theory”. In: *Foundations of Physics* 1.1 (1970-03), pp. 69–76. ISSN: 1572-9516. DOI: [10.1007/BF00708656](https://doi.org/10.1007/BF00708656).
- [72] Maximilian Schlosshauer. “Decoherence, the Measurement Problem, and Interpretations of Quantum Mechanics”. In: *Reviews of Modern Physics* 76.4 (2005-02), pp. 1267–1305. DOI: [10.1103/RevModPhys.76.1267](https://doi.org/10.1103/RevModPhys.76.1267).
- [73] Wojciech Hubert Zurek. “Decoherence, Einselection, and the Quantum Origins of the Classical”. In: *Reviews of Modern Physics* 75.3 (2003-05), pp. 715–775. ISSN: 0034-6861, 1539-0756. DOI: [10.1103/RevModPhys.75.715](https://doi.org/10.1103/RevModPhys.75.715).
- [74] Wojciech Hubert Zurek. “Quantum Darwinism”. In: *Nature Physics* 5.3 (2009-03), pp. 181–188. ISSN: 1745-2473, 1745-2481. DOI: [10.1038/nphys1202](https://doi.org/10.1038/nphys1202).
- [75] Wojciech Hubert Zurek. “Decoherence and the Transition from Quantum to Classical — Revisited”. In: *Quantum Decoherence*. Ed. by Bertrand Duplantier, Jean-Michel Raimond, and Vincent Rivasseau. Basel: Birkhäuser Basel, 2006, pp. 1–31. ISBN: 978-3-7643-7807-3 978-3-7643-7808-0. DOI: [10.1007/978-3-7643-7808-0_1](https://doi.org/10.1007/978-3-7643-7808-0_1).
- [76] R. Horodecki, J. K. Korbicz, and P. Horodecki. “Quantum origins of objectivity”. In: *Physical Review A* 91.3 (2015-03). ISSN: 1094-1622. DOI: [10.1103/physreva.91.032122](https://doi.org/10.1103/physreva.91.032122). URL: <http://dx.doi.org/10.1103/PhysRevA.91.032122>.
- [77] J. K. Korbicz. “Roads to Objectivity: Quantum Darwinism, Spectrum Broadcast Structures, and Strong Quantum Darwinism – a Review”. In: *Quantum* 5 (2021-11), p. 571. ISSN: 2521-327X. DOI: [10.22331/q-2021-11-08-571](https://doi.org/10.22331/q-2021-11-08-571). arXiv: [2007.04276 \[quant-ph\]](https://arxiv.org/abs/2007.04276).
- [78] Thao P. Le and Alexandra Olaya-Castro. “Strong Quantum Darwinism and Strong Independence are Equivalent to Spectrum Broadcast Structure”. In: *Physical Review Letters* 122.1 (2019-01). ISSN: 1079-7114. DOI: [10.1103/physrevlett.122.010403](https://doi.org/10.1103/physrevlett.122.010403). URL: <http://dx.doi.org/10.1103/PhysRevLett.122.010403>.
- [79] Noah Linden, Sandu Popescu, Anthony J. Short, and Andreas Winter. “Quantum Mechanical Evolution towards Thermal Equilibrium”. In: *Physical Review E* 79.6 (2009-06), p. 061103. DOI: [10.1103/PhysRevE.79.061103](https://doi.org/10.1103/PhysRevE.79.061103).
- [80] Sheldon Goldstein, Joel L. Lebowitz, Roderich Tumulka, and Nino Zanghì. “Canonical Typicality”. In: *Physical Review Letters* 96.5 (2006-02), p. 050403. DOI: [10.1103/PhysRevLett.96.050403](https://doi.org/10.1103/PhysRevLett.96.050403).

- [81] Peter Reimann. “Foundation of Statistical Mechanics under Experimentally Realistic Conditions”. In: *Physical Review Letters* 101.19 (2008-11), p. 190403. DOI: [10.1103/PhysRevLett.101.190403](https://doi.org/10.1103/PhysRevLett.101.190403).
- [82] Peter Reimann. “Equilibration of Isolated Macroscopic Quantum Systems under Experimentally Realistic Conditions”. In: *Physica Scripta* 86.5 (2012-10), p. 058512. ISSN: 1402-4896. DOI: [10.1088/0031-8949/86/05/058512](https://doi.org/10.1088/0031-8949/86/05/058512).
- [83] I. Devetak and A. Winter. “Distilling Common Randomness from Bipartite Quantum States”. In: *IEEE Transactions on Information Theory* 50.12 (2004-12), pp. 3183–3196. ISSN: 0018-9448. DOI: [10.1109/TIT.2004.838115](https://doi.org/10.1109/TIT.2004.838115). arXiv: [quant-ph/0304196](https://arxiv.org/abs/quant-ph/0304196).
- [84] J. Tuziemski and J. K. Korbicz. “Analytical Studies of Spectrum Broadcast Structures in Quantum Brownian Motion”. In: *Journal of Physics A: Mathematical and Theoretical* 49.44 (2016-11), p. 445301. ISSN: 1751-8113, 1751-8121. DOI: [10.1088/1751-8113/49/44/445301](https://doi.org/10.1088/1751-8113/49/44/445301). arXiv: [1603.04217](https://arxiv.org/abs/1603.04217).
- [85] Paul Busch and Pekka Lahti. “The Standard Model of Quantum Measurement Theory: History and Applications”. In: *Foundations of Physics* 26.7 (1996-07), pp. 875–893. ISSN: 0015-9018, 1572-9516. DOI: [10.1007/BF02148831](https://doi.org/10.1007/BF02148831). arXiv: [quant-ph/9603020](https://arxiv.org/abs/quant-ph/9603020).
- [86] Jake Xuereb, Paul Erker, Florian Meier, Mark T. Mitchison, and Marcus Huber. “The Impact of Imperfect Timekeeping on Quantum Control”. In: arXiv:2301.10767 (2023-01). arXiv: [2301.10767](https://arxiv.org/abs/2301.10767) [quant-ph].