

Quantum Horn's lemma, finite heat baths, and the third law of thermodynamics

Jakob Scharlau¹ and Markus P. Müller^{2,3,4,5,1}

¹Department of Theoretical Physics, University of Heidelberg, Heidelberg, Germany

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

³Department of Applied Mathematics, University of Western Ontario, London, ON N6A 5BY, Canada

⁴Department of Philosophy, University of Western Ontario, London, ON N6A 5BY, Canada

⁵Perimeter Institute for Theoretical Physics, Waterloo, ON N2L 2Y5, Canada

February 2, 2018

Interactions of quantum systems with their environment play a crucial role in resource-theoretic approaches to thermodynamics in the microscopic regime. Here, we analyze the possible state transitions in the presence of “small” heat baths of bounded dimension and energy. We show that for operations on quantum systems with fully degenerate Hamiltonian (noisy operations), all possible state transitions can be realized exactly with a bath that is of the same size as the system or smaller, which proves a quantum version of Horn’s lemma as conjectured by Bengtsson and Życzkowski. On the other hand, if the system’s Hamiltonian is not fully degenerate (thermal operations), we show that some possible transitions can only be performed with a heat bath that is unbounded in size and energy, which is an instance of the third law of thermodynamics. In both cases, we prove that quantum operations yield an advantage over classical ones for any given finite heat bath, by allowing a larger and more physically realistic set of state transitions.

I Introduction

Thermodynamics is traditionally viewed as a statistical theory of a large number of particles. In the last few years, however, there has been a wave of activity in extending the realm of thermodynamics to microscopic quantum systems, motivated by ideas from quantum information theory [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 20, 21, 22, 23, 24]. One key idea is to formulate thermodynamics as a *resource theory* [6, 25, 26, 27, 28]: by specifying and restricting to a set of experimental operations that an agent is allowed

to implement “for free”, one obtains a rigorous framework which allows to address typical thermodynamics questions even in physical scenarios that are far from the usual thermodynamic limit.

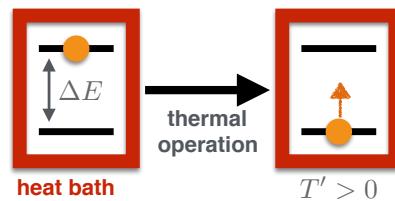


Figure 1: Theorem 9 can be seen as a simple instance of the third law of thermodynamics. We start with a qubit in its excited state, which serves on the one hand as a work reservoir, and on the other hand represents the system that we would like to cool. For any finite heat bath, thermal operations can never achieve the ground state exactly; there will always remain finite fluctuations, indicated by the arrow. The final temperature T' will always be strictly positive, with a lower bound that depends inversely-extensively on the size and energy of the heat bath. The physical intuition of this will be elaborated in more detail in the conclusions.

One way to formulate a suitable resource theory is by declaring the set of *thermal operations* as allowed. A thermal operation consists of an arbitrary global energy-preserving unitary transformation on the given quantum system A and an arbitrary heat bath B (at fixed background temperature), and the subsequent disposal of the heat bath. By making A itself consist of several parts (say, a single particle and a work-storage device), we can study the interplay of heat and work in microscopic systems. For example, we can ask how much work can be extracted with high probability from a given quantum state by using only thermal operations. Many of these questions have been answered explicitly,

and have been shown to refine – and converge to – the usual thermodynamic results in the limit of a large number of weakly interacting particles.

Even though this approach has been tremendously successful in providing insights into thermodynamics in the microscopic regime, it has also raised questions about the way that thermal operations should be modelled and understood [43, 44]. While the quantum system A is allowed to be microscopic (say, a single particle), the heat bath B is often assumed to be macroscopic, or even explicitly assumed to be unbounded. However, in many physical situations, this is not a realistic assumption. In particular, we may think of a situation where a thermodynamic protocol has to be completed in a finite amount of time, allowing A to effectively interact only with a restricted number of degrees of freedom in the bath. The third law of thermodynamics [16] can be seen as an example of this, as we will argue further below. This leads us to study thermal operations with bounded heat baths B in this paper.

The thermodynamical questions notwithstanding, there is also a purely mathematical motivation to study finite baths, which comes from more traditional quantum information problems. For example, consider the special case that all Hamiltonians (of system and bath) are fully degenerate, i.e. proportional to the identity. Then thermal operations are maps of the form

$$\rho'_A = \text{Tr}_B \left[U \left(\rho_A \otimes \frac{\mathbf{1}_B}{m} \right) U^\dagger \right], \quad (1)$$

with $m = \dim B$, and U an arbitrary unitary on AB . These maps have been called *noisy operations* [2, 32] in quantum information theory. They have been studied first without explicit reference to thermodynamics, often under different names (like “ K -unistochastic maps” [35, 36, 37] or “exactly factorizable maps” [41, 42]), formalizing the idea of a map that dilutes quantum systems with random noise.

One important question in this context is whether for two given quantum states ρ_A and ρ'_A , there exists a noisy operation that maps ρ_A to ρ'_A . This question has been partially answered in [2]: if the spectrum (vector of eigenvalues) of ρ_A majorizes the spectrum of ρ'_A , denoted

$$\rho_A \succ \rho'_A, \quad (2)$$

then one can obtain ρ'_A to arbitrary accuracy by noisy operations from ρ_A . Since unitary basis changes are allowed operations, we may assume that both ρ_A and ρ'_A are diagonal matrices. Then (2) can be proven by reduction to the case of discrete classical probability distributions, where the unitary U acts as a suitable permutation of basis vectors.

This strategy of “proof by reduction to classical” has one important drawback, however: it only works if the

heat bath is assumed to be unbounded. To see this, suppose that ρ_A is a pure state, and ρ'_A a mixed state with irrational eigenvalues. In this case, eq. (2) holds. But if U is a classical permutation, then the right-hand side of (1) will always be a density matrix with entries that are all rational numbers, and never be equal to ρ'_A which has irrational diagonal elements. In other words, to approximate ρ'_A to arbitrary accuracy, the dimension m of the bath has to become arbitrarily large.

It has been conjectured [36] that one can in fact perform all transitions $\rho_A \rightarrow \rho'_A$ exactly with a finite bath of dimension $\dim B = \dim A$, by using more general unitaries U which are not permutations. In this paper, we give a simple proof of this conjecture (see Lemma 4 below), which can be seen as a quantum version of Horn’s lemma. This shows that a small bath is always sufficient to implement any noisy state transition $\rho_A \rightarrow \rho'_A$ exactly.

A “proof by reduction to classical” has also been the standard approach in the thermodynamic setting, where the Hamiltonians of system and bath are allowed to be non-trivial: the possibility of a transition $\rho_A \rightarrow \rho'_A$ by thermal operations has standardly been proven by resorting to the classical case, where U is an (energy-preserving) permutation. This leads to a setup where one is forced to allow a heat bath B which is unbounded in dimension and energy, in order to approximate all possible state transitions. As mentioned above, this assumption does not always seem to be physically realistic, in particular since it presupposes the fine control to implement arbitrary energy-preserving unitaries on the joint quantum system AB , where B is very large.

One way that this problem has recently been addressed is by considering a small subset of “coarse” operations that are still sufficient to realize all possible thermal state transitions [44]. Similarly, Ref. [45] has studied the role of catalysts (auxiliary systems that ease a state transformation) that are bounded in size or energy. In this paper, we take a different route. We hold on to thermal operations, and ask whether there is an analogue of quantum Horn’s lemma for them, in particular in the quasiclassical case where both initial and final state are blockdiagonal in energy. If this was the case, then every thermal state transition between quasiclassical states could be realized exactly with a finite heat bath.

Here we give two results that answer this question in a complementary way. First, we show that “quantum” thermal operations (that is, ones with arbitrary energy-preserving unitaries U) allow indeed many transitions exactly which could only be realized approximately by “classical” thermal operations (where U is an energy-preserving permutation). For quantum systems with non-degenerate Hamiltonian H_A , Figure 4

illustrates this: for a fixed initial state and fixed finite heat bath, the possible final states obtained by classical thermal operations form a discrete finite set. In contrast, the final states obtainable by quantum thermal operations correspond to the convex hull of this set (up to some minor amendments in the case of degeneracies). This removes the unphysical discreteness originating from the finite set of classical permutations, and paints a physically more realistic and also mathematically more convenient picture. It is formalized in our first main result, Theorem 10, which generalizes quantum Horn’s lemma to thermal operations.

Figure 5 shows a simple consequence of this fact, which is the content of Corollary 12: the set of all states that can be obtained to arbitrary accuracy from an initial state by thermal operations (that is, the states “thermomajorized” by the given initial state) is a compact convex set. All states in its relative interior can be obtained exactly by quantum thermal operation with a suitable finite heat bath.

However, we show that this is not in general true for the states in the relative boundary, which represents our second main result. In particular, for every non-degenerate two-level system A , we show explicitly that there is a state that can be obtained to arbitrary accuracy, but never exactly with any finite heat bath.

This result has an interesting physical interpretation, illustrated in Figure 1. Consider the special case of a two-level system with some energy gap ΔE , which is initially in its (pure) excited state. Suppose we would like to couple this qubit system to a heat bath, in such a way that it is in its ground state after the interaction. Physically, this is arguably the simplest possible instance of *cooling*: we draw energy from the qubit system itself (coming from the excited state) in order to cool it down to zero temperature, which is less than the ambient (heat bath) temperature. However, we show in Theorem 9 that the ground state can never be obtained exactly with any finite heat bath: it can only be approximated, such that the final temperature T' can never be exactly zero. We give a strictly positive lower bound on T' which depends on the energy and dimension of the heat bath, which can be seen as an instance of the third law of thermodynamics. While the third law has been examined in more detail elsewhere [16], our result suggests an intuitive picture in terms of the geometry of the state space: states of zero temperature lie on the relative boundary of the set of thermomajorized states (boundary of the light-gray region in Figure 5), and not all of these states can be obtained with a finite heat bath.

Note that our results do not prove that zero temperature can never be attained with any finite heat bath. In the geometric picture of Figure 5, there are states

on the boundary which can be obtained exactly (for example, the initial state itself). Thus, we here only consider a very special unattainability result. On the other hand, our formulation suggests a generalization of the third law beyond cooling to zero temperature. For example, in the scenario of Figure 1, we can also look at initial states which are only partially excited, and obtain a potential target temperature $T'_0 > 0$ that can only be achieved approximately, but never exactly, by interactions with any finite heat bath.

Our paper is organized as follows. In Section II, we provide the mathematical framework and notation. We prove our main results for noisy operations in Section III, and for thermal operations in Section IV. In the thermal case, we first consider the case of a two-level system in Subsection IV.1, and then the general case in Subsection IV.2, before discussing the physical intuition and comparison to other approaches in more detail in the Conclusions, Section V.

II Mathematical preliminaries

We will assume that all quantum systems (labelled by uppercase letters like A) are finite-dimensional, and carry a Hamiltonian, i.e. a self-adjoint operator $H_A = H_A^\dagger$. In the context of noisy operations (in particular Section III), we assume that all Hamiltonians are zero. We also assume that every quantum system comes with a choice of orthonormal basis, in such a way that its Hamiltonian is diagonal in that basis.

A *classical state* is a probability vector $p = (p_1 \dots, p_n) \in \mathbb{R}^n$, with $p_i \geq 0$, $\sum_i p_i = 1$. We will treat probability vectors as column vectors in matrix calculations, but for notational simplicity, we denote them by row vectors. We identify classical states $p \in \mathbb{R}^n$ with the corresponding quantum states (diagonal density matrices) \hat{p} on \mathbb{C}^n , where

$$p = (p_1, \dots, p_n) \quad \Rightarrow \quad \hat{p} := \begin{pmatrix} p_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & p_n \end{pmatrix}.$$

Given two classical states $p, q \in \mathbb{R}^n$, we say that p *majorizes* q [30], and write $p \succ q$, if and only if

$$\sum_{i=1}^k p_i^\downarrow \geq \sum_{i=1}^k q_i^\downarrow \quad \text{for all } k = 1, \dots, n,$$

where $p^\downarrow = (p_1^\downarrow, \dots, p_n^\downarrow)$ denotes the entries of p in non-increasing order, i.e. $p_1^\downarrow \geq p_2^\downarrow \geq \dots \geq p_n^\downarrow$. Furthermore, we write $\rho \succ \sigma$ for quantum states ρ, σ if and only if $\text{spec}(\rho) \succ \text{spec}(\sigma)$, where $\text{spec}(\rho)$ denotes the vector of eigenvalues of ρ . According to the Hardy-Littlewood-Polyá theorem [30], $p \succ q$ if and only if there exists

a bistochastic matrix $D \in \mathbb{R}^{n \times n}$, i.e. a matrix with $D_{ij} \geq 0$, $\sum_i D_{ij} = \sum_j D_{ij} = 1$, such that $q = Dp$. The following lemma is easily checked by direct calculation, but will be very useful in the following. It has also been observed in [38].

Lemma 1 (Ruch and Mead [38]). *Let $p \in \mathbb{R}^n$ be a classical state, and let $U \in \mathrm{U}(n)$ be a unitary. Then the vector of diagonal elements of $U\hat{p}U^\dagger$ can be written*

$$(U\hat{p}U^\dagger)_{11}, \dots, (U\hat{p}U^\dagger)_{nn}) = Dp,$$

where D is a bistochastic matrix with entries $D_{ij} = |U_{ij}|^2$, i.e. $D = U \circ U^*$ for the Hadamard [29] product \circ , i.e. $(X \circ Y)_{ij} = X_{ij}Y_{ij}$.

Given any quantum system A , a *noisy operation* is a map on the quantum states of A which can be written in the form (1), for some suitable choice of quantum system B of dimension $\dim B = m$ and unitary transformation U on AB . (We will not consider more general noisy or thermal operations which map states on one Hilbert space to states on another Hilbert space in this paper.) The known relation between majorization and noisy operations is as follows.

Lemma 2 ([2, 36]). *If $\rho \rightarrow \rho'$ by some noisy operation, then $\rho \succ \rho'$. Conversely, if $\rho \succ \rho'$, then for every $\varepsilon > 0$, there exists some noisy operation Φ with $\Phi(\rho) = \rho'_\varepsilon$, where $\|\rho'_\varepsilon - \rho'\| < \varepsilon$. If both ρ and ρ' are diagonal in the canonical eigenbasis, then Φ can be (and is standardly) chosen in such a way that the unitary U in (1) is a (“classical”) permutation of basis vectors.*

In the case of non-trivial Hamiltonians, all these notions are generalized. We always assume that there is a fixed background inverse temperature $\beta = 1/(k_B T)$, with k_B the Boltzmann constant. The classical Gibbs state of a system B with Hamiltonian H_B is $\gamma_B = (e^{-\beta E_1}, \dots, e^{-\beta E_n})/Z$, where E_1, \dots, E_n are the eigenvalues (and thus diagonal elements) of H_B , and $Z = \sum_{i=1}^n e^{-\beta E_i}$ is the partition function. The corresponding quantum Gibbs state is $\hat{\gamma}_B = \exp(-\beta H_B)/Z$. We will sometimes write γ^B for γ_B if we need space for indices.

A *thermal operation* on a quantum system A is a map of the form

$$\rho'_A = \mathrm{Tr}_B [U(\rho_A \otimes \hat{\gamma}_B)U^\dagger], \quad (3)$$

where U is an energy-preserving unitary, i.e. $[U, H_A + H_B] = 0$ (this constraint becomes void in the case that $H_A = H_B = 0$, which recovers the noisy operations). Any choice of bath B is allowed, but the bath must be in its thermal state $\hat{\gamma}_B$. An important property of noisy operations Φ is that they cannot generate coherences between different energies: that is, if $[\rho, H_A] = 0$, then

$[\Phi(\rho), H_A] = 0$. That is, blockdiagonal states (in the energy eigenbasis) are mapped to blockdiagonal states.

Thermal operations are closely related to a generalized notion of majorization, named *thermomajorization* (this also corresponds to “majorization relative to d ” for $d = \gamma$ in [30]). For any two classical states p, p' on a system A with Gibbs state $\gamma = \gamma_A$, we write

$$p \succ_\gamma p'$$

and say that “ p thermomajorizes p' ” if and only if there exists a stochastic γ -preserving matrix D such that $p' = Dp$. Stochasticity of D means that D maps states to states, i.e. $D_{ij} \geq 0$ and $\sum_i D_{ij} = 1$ for all j , and γ -preservation is the property that $D\gamma = \gamma$.

The following was shown (in different notation) by Janzing et al. [1], see also [39, 40].

Lemma 3 (Janzing et al. [1]). *If $\hat{p} \rightarrow \hat{p}'$ by some thermal operation, where p and p' are classical states on the same system A , then $p \succ_\gamma p'$, where $\gamma = \gamma_A$ is the Gibbs state of A .*

Conversely, if $p \succ_\gamma p'$, then for every $\varepsilon > 0$, there exists some thermal operation Φ with $\Phi(\hat{p}) = \hat{p}'_\varepsilon$, where $\|\hat{p}'_\varepsilon - \hat{p}'\| < \varepsilon$. This map Φ can be chosen in such a way that the unitary U in (1) is a permutation of basis vectors; the bath Hamiltonian H_B can be chosen to be composed of a sufficient number of non-interacting copies of systems with Hamiltonian H_A .

In the following, we are only interested in transitions between classical states, i.e. instances of (3) where input and output state are both diagonal in the canonical (energy eigen)basis. Note that this does not mean that we restrict to operations Φ that map *all* diagonal input states to diagonal output states.

We are thus not considering the role of quantum coherence in this paper, which is currently an active area of research [21, 24, 23]. However, we are going beyond the earlier quasiclassical setups by explicitly considering unitaries U in (3) which are not just classical permutations of basis vectors.

III Noisy operations with small auxiliary systems

As mentioned in the introduction, in most of the literature on single-shot quantum thermodynamics (e.g. [2, 5, 6, 11, 32]), noisy operations are constructed by implementing a classical permutation in a high-dimensional degenerate eigenspace of system and bath. To approximate target states to arbitrary accuracy, this needs in general arbitrarily large auxiliary systems. (We also use the word “bath” for these auxiliary systems for brevity,

even though they do not carry a non-trivial Hamiltonian and are thus not heat baths in a strict sense.)

We will now show that this is not necessary if we consider more general unitaries. In this case, all transitions can be implemented exactly with a bath that is of the same dimension as the system.

Lemma 4. *If $p, p' \in \mathbb{R}^n$ are states such that $p \succ p'$, then the transition $p \rightarrow p'$ can be accomplished by a noisy operation with an auxiliary system of size n . That is, there is a unitary U on AB , $A = B = \mathbb{C}^n$, such that*

$$\hat{p}' = \text{Tr}_B \left[U \left(\hat{p} \otimes \frac{\mathbf{1}}{n} \right) U^\dagger \right]. \quad (4)$$

Proof. We will apply the *Schur-Horn Theorem* [33], which tells us that there is a unitary V on A such that $V\hat{p}V^\dagger$ has diagonal elements p'_1, \dots, p'_n . Therefore, $\hat{p}' = \Phi(V\hat{p}V^\dagger)$, where Φ is the “decoherence map”, satisfying

$$\langle i | \Phi(\rho) | j \rangle = \begin{cases} \langle i | \rho | j \rangle & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

It remains to show that Φ can be implemented as a noisy operation with a bath of size n . To this end, choose an arbitrary orthonormal basis $|1\rangle, \dots, |n\rangle$ on $B = \mathbb{C}^n$, and let π be the cyclic permutation with $\pi|i\rangle = |i+1\rangle$ if $i < n$ and $\pi|n\rangle = |1\rangle$. Define the unitary W by

$$W = \sum_{k=1}^n |k\rangle\langle k| \otimes \pi^k.$$

Then, if $\rho' = \text{Tr}_B [W (\rho \otimes \frac{\mathbf{1}}{n}) W^\dagger]$, we have $\langle i | \rho' | j \rangle = \frac{1}{n} \langle i | \rho | j \rangle \text{tr}(\pi^{i-j})$, which is zero if $i \neq j$ and $\langle i | \rho | j \rangle$ otherwise. In other words, the unitary W implements Φ as a noisy operation, and we can set $U := W(V \otimes \mathbf{1}_B)$. \square

This result confirms a conjecture by Bengtsson and Życzkowski [36], who however use slightly different terminology. They call a noisy operation on a system of size n , mediated by a bath of size n , “unistochastic”, regarded as a map between quantum states (note that this conflicts with our use of this term in Lemma 13 in the appendix). For two quantum states ρ and σ , write $\rho \succ \sigma$ if and only if the spectrum of ρ majorizes the spectrum of σ . Then, in the terminology of [35, 36, 37]), Lemma 4 implies “Quantum Horn’s Lemma”: *If $\rho \succ \sigma$, then there exists a unistochastic quantum operation that maps ρ to σ exactly.* Clearly, if p and p' denote the probability vectors of eigenvalues of ρ and σ , then there is a unitary U such that (4) holds true, with a bath of size n . Since there are unitaries V and W such that $\hat{p} = V\rho V^\dagger$ and $\hat{p}' = W\sigma W^\dagger$, we can set $U' := (W^\dagger \otimes \mathbf{1})U(V \otimes \mathbf{1})$, and then U' will generate a “unistochastic quantum operation” that maps ρ to σ .

It is interesting to see that Lemma 4 above is closely related to (and, in fact, has an alternative proof in terms of) a variant of the quantum marginal problem as formulated, for example, in [46]: given the spectrum of some bipartite quantum state ρ_{AB} , then what are the possible spectra of the local reduced state ρ_A ?

The following lemma gives a complete answer to this question in the case $\dim A \leq \dim B$. In a nutshell, it says that the possible spectra of ρ_A are those that are majorized by the partial trace of the (non-increasingly ordered) spectrum of ρ_{AB} .

Lemma 5. *Let ρ be a quantum state on AB , and σ a quantum state on A . If $\dim A \leq \dim B$, there exists a unitary U on AB with*

$$\text{Tr}_B (U\rho U^\dagger) = \sigma$$

if and only if

$$\text{Tr}_B \hat{\lambda}(\rho) \succ \sigma, \quad (5)$$

where $\lambda(\rho)$ is the vector of eigenvalues of ρ , ordered in non-increasing order, and $\hat{\lambda}(\rho)$ is the corresponding diagonal density matrix.

This lemma was conjectured in [46] in 2005, where it was also shown that (5) is *necessary* for the existence of a suitable unitary U which satisfies the statement of the lemma. After Lemma 5 was proven by the first author in his Master thesis [47], it turned out (cf. Acknowledgments) that the statement had just recently been proven independently in [48]. The lemma can also be inferred from the results in [49]. We give our proof in the appendix for completeness.

Lemma 4 follows as a corollary from Lemma 5 by setting $\dim A = \dim B = n$, $\rho = \rho_A \otimes \mathbf{1}/n$, $\sigma = \rho'_A$, and noting that

$$\lambda(\rho) = \frac{1}{n} (\underbrace{\lambda_1, \dots, \lambda_1}_n, \underbrace{\lambda_2, \dots, \lambda_2}_n, \dots, \underbrace{\lambda_n, \dots, \lambda_n}_n),$$

where $\lambda_1 \geq \lambda_2 \geq \dots$ are the eigenvalues of ρ_A . The partial trace corresponds to summing over n -blocks, thus $\text{Tr}_B \hat{\lambda}(\rho) = \hat{\lambda}$, where $\lambda = (\lambda_1, \dots, \lambda_n)$. These are the eigenvalues of ρ_A , which must therefore majorize the eigenvalues of ρ'_A .

Given that all possible transitions on a system A of size n can be achieved with a bath B of size n , it is natural to ask whether one can more generally classify the set of all *maps* on the classical probability distributions that can be implemented in this setting. We give a partial answer in Lemma 13 in the appendix, where we show that the corresponding set of maps lies somewhere in between the unistochastic and the bistochastic maps.

The relation to the quantum marginal problem allows us to explore what happens in the case that the heat

bath dimension $m = \dim B$ is smaller than the system dimension $n = \dim A$. The case $m = 2, n = 3$ has been considered in [46] in Sec. 7.3, where the authors have worked out the full set of inequalities relating the eigenvalues of ρ_{AB} to the eigenvalues of possible marginals ρ_A . Their result implies the following:

Lemma 6 (Case $n = 3, m = 2$). *If $p, p' \in \mathbb{R}^3$ are states such that $p \succ p'$, then the transition $p \rightarrow p'$ can be accomplished by a noisy operation with an auxiliary system of size $m = 2$. That is, there is a unitary U on AB , $A = \mathbb{C}^3$, $B = \mathbb{C}^2$, such that*

$$\hat{p}' = \text{Tr}_B \left[U \left(\hat{p} \otimes \frac{\mathbf{1}}{2} \right) U^\dagger \right].$$

That is, in this case, even a bath that is smaller than the system is sufficient to implement all noisy state transitions.

We do not currently know the smallest possible value of m that would allow all noisy state transitions on a system of dimension n . However, we have the following elementary bound:

Lemma 7. *If all noisy state transitions on a system of size n can be implemented with a bath of size m , then $m \geq \sqrt{n}$.*

Proof. If all noisy state transitions can be implemented, then in particular a pure state can be mapped to a maximally mixed state. That is, there is a pure state $|\psi\rangle$ on $A = \mathbb{C}^n$ and a unitary U on AB such that

$$\begin{aligned} \frac{1}{n} &= \text{Tr}_B \left[U \left(|\psi\rangle\langle\psi|_A \otimes \frac{\mathbf{1}}{m} \right) U^\dagger \right] \\ &= \frac{1}{m} \sum_{i=1}^m \text{Tr}_B \left[U (|\psi\rangle\langle\psi|_A \otimes |i\rangle\langle i|_B) U^\dagger \right], \end{aligned}$$

where $\{|i\rangle\}_{i=1}^m$ is any orthonormal basis of $B = \mathbb{C}^m$. The Schmidt rank of the pure state $U|\psi\rangle_A|i\rangle_B$ is upper-bounded by m , and so the rank of every density operator $\text{Tr}_B[\dots]$ on the right-hand side is upper-bounded by m . Due to the subadditivity of the rank, this means that the rank of the right-hand side is less than or equal to m^2 . Hence $\text{rank}(\mathbf{1}/n) = n \leq m^2$. \square

We do not know whether this bound is tight. If it was, then this would say that any noisy operation on k qubits can be achieved with a bath of $k/2$ qubits only (rounding up accordingly). This would be a significant tightening of quantum Horn's lemma.

IV Thermal operations with small heat baths

We now turn to the more general case of thermal operations. We will start by analyzing two-level systems in

Subsection IV.1, which illustrates most of the features of arbitrary-dimensional systems that will be proven in Subsection IV.2.

IV.1 The qubit case with a finite heat bath

Consider a two-dimensional Hilbert space with orthonormal basis $\{|1\rangle, |2\rangle\}$ and Hamiltonian

$$H_A = \Delta E |2\rangle\langle 2|,$$

and assume that $\Delta E > 0$. The corresponding Gibbs state at inverse temperature $\beta > 0$ is given by

$$\gamma_A = (|1\rangle\langle 1| + e^{-\beta\Delta E} |2\rangle\langle 2|) / (1 + e^{-\beta\Delta E}).$$

Now we would like to see which (classical) states p' are thermomajorized by a given state p . According to Lemma 3, we need a stochastic matrix D which satisfies

$$Dp = p' \quad \text{and} \quad D\gamma = \gamma.$$

In two dimensions the stochastic matrices that preserve γ are parametrized by a single parameter α :

$$D_\alpha = \begin{pmatrix} 1 - \alpha & \alpha \gamma_1 / \gamma_2 \\ \alpha & 1 - \alpha \gamma_1 / \gamma_2 \end{pmatrix} \quad (6)$$

with the requirement $0 \leq \alpha \leq \gamma_2/\gamma_1$, where γ_1, γ_2 are the components of $\gamma = \gamma_A$. The set of states thermomajorized by p is thus also parametrized by α , and we can use the notation $p'_\alpha := D_\alpha p$. Three special cases are as follows:

- $\alpha = 0$ gives $D_\alpha = \mathbf{1}$ and $p'_\alpha = p$.
- $\alpha = \gamma_2$ gives $D_\alpha = \begin{pmatrix} \gamma_1 & \gamma_1 \\ \gamma_2 & \gamma_2 \end{pmatrix}$, which maps all states p to the thermal state $p'_\alpha = \gamma$.
- $\alpha = \gamma_2/\gamma_1 = e^{-\beta\Delta E}$ gives the final state

$$p'_\alpha = p^* := \left(1 - \frac{\gamma_2}{\gamma_1} p_1, \frac{\gamma_2}{\gamma_1} p_1 \right). \quad (7)$$

The set of states thermomajorized by p , including these special cases, is shown in Figure 2. In the case $\Delta E = 0$, the D_α are bistochastic, and we recover the case of standard majorization from Section III.

In order to implement the D_α as thermal operations, it is necessary to choose H_B in a way that $H_A + H_B$ exhibits some degeneracies, since the unitaries in (3) have to commute with the total Hamiltonian. So the spectrum of H_B has to contain the energy gap ΔE , and furthermore it seems plausible that carrying the energy gap between many different energy levels should allow for a greater variety of transitions.



Figure 2: The line shows the diagonal states of a qubit system with $(1, 0)$ and $(0, 1)$ corresponding to the eigenstates $|1\rangle$ and $|2\rangle$. The states thermomajorized by p correspond to the points between p and p^* .

Let us therefore consider a truncated harmonic oscillator

$$H_B = \Delta E \sum_{i=1}^m (i-1) |i\rangle\langle i|, \quad (8)$$

that is, an m -level system that has ΔE as energy difference between all neighbouring eigenstates. In the case $m = 2$, we have one copy of the original system, and for $m \rightarrow \infty$ we obtain a harmonic oscillator. The joint Hamiltonian of system and bath $H_{AB} = H_A + H_B$ becomes

$$H_{AB} = \Delta E \sum_{i=1}^2 \sum_{j=1}^m (i+j-2) |i,j\rangle\langle i,j|$$

where $|i,j\rangle_{AB} = |i\rangle_A \otimes |j\rangle_B$. All energy levels of H_{AB} except for the lowest and highest one are doubly degenerate. As explained above, this degeneracy is necessary to implement non-trivial thermal operations.

In the case $m = 2$, we can implement one of the special cases above (apart from the trivial case), namely D_α for $\alpha = \gamma_2$. This is achieved by the energy-preserving permutation π with $\pi|12\rangle = |21\rangle$ and $\pi|21\rangle = |12\rangle$ which swaps the two systems, leading to $\hat{p} \otimes \hat{\gamma} \mapsto \hat{\gamma} \otimes \hat{p}$. Tracing out the second system, we end up with the Gibbs state of the qubit, independent of the state we started with.

By implementing other unitaries $U_{2 \times 2}$ on the subspace $\text{span}\{|12\rangle, |21\rangle\}$, we can in fact implement all maps D_α with $0 \leq \alpha \leq \gamma_2$, and thus obtain all corresponding target state p'_α . This follows from the simple observation that $SU(2)$ is connected, and thus by varying over all $U_{2 \times 2} \in SU(2)$, we obtain target states p' that vary continuously between $p' = p$ and $p' = \gamma$.

Can we obtain more target states than this? To answer this in the negative, consider the parametrization

$$U_{2 \times 2} = \begin{pmatrix} \kappa & -\lambda^* \\ \lambda & \kappa^* \end{pmatrix}$$

with $\kappa, \lambda \in \mathbb{C}$ such that $|\kappa|^2 + |\lambda|^2 = 1$. Implementing this unitary on the degenerate subspace, we obtain a target state $p'(U_{2 \times 2})$. A simple calculation shows that

$$p'(U_{2 \times 2}) = |\kappa|^2 p + |\lambda|^2 \gamma = |\kappa|^2 p'(\mathbf{1}_{2 \times 2}) + |\lambda|^2 p'(\pi_{2 \times 2}),$$

where π is the swap introduced above. This shows that the extremal cases are generated by classical permutations (swap or identity on the degenerate subspace) and choosing general energy-preserving unitaries give us all states that are in the convex hull or ‘in between’ these two states, but not more.

Now we would like to see which states we can achieve for $m > 2$. Consider the permutation

$$\pi_{\max} : \quad |1, j\rangle \longleftrightarrow |2, j-1\rangle \quad \text{for } j = 2, \dots, m$$

which swaps the basis vectors of all two-dimensional energy eigenspaces of H_{AB} . Intuitively, π_{\max} is the classical permutation for which we expect to obtain the D_α with the largest possible value of α (for this specific heat bath). Calculating the thermal operation for $p = (1, 0)$ and reading off $1 - \alpha$ from $p' = (1 - \alpha, \alpha)$ shows that π_{\max} yields

$$1 - \alpha = \gamma_1^B = \frac{1}{\text{tr}(e^{-\beta H_B})}.$$

which is indeed the minimal value for $1 - \alpha$ that is achievable by choosing any energy-preserving permutation. This becomes clear from the the fact that all energy-preserving permutations can be written as π_J , where $J \subset \{2, \dots, m\}$, and

$$\pi_J |1, j\rangle := \begin{cases} |2, j-1\rangle & \text{if } j \in J, \\ |1, j\rangle & \text{otherwise.} \end{cases}$$

That is, J indicates in which subspaces we choose to swap the two basis vectors, and $\pi_{\max} = \pi_{\{2, \dots, m\}}$. Choosing some permutation π_J leads a corresponding value α of $1 - \alpha = \sum_{j \in \{1, \dots, m\} \setminus J} \gamma_j^B$ which is indeed minimal for $J = \{2, \dots, m\}$. We thus define

$$\alpha_{\max}^{(m)} = 1 - \gamma_1^B = 1 - \frac{1 - e^{-\beta \Delta E}}{1 - e^{-m \beta \Delta E}}, \quad (9)$$

which is the maximal value of the parameter α that we can achieve with this kind of heat bath of dimension m , even if we choose not just a permutation but a more general unitary for the thermal operation. This can be seen similarly as in the case $m = 2$. The energy spectrum of H_{AB} is $0, \Delta E, \dots, m \Delta E$, and all energies except for the highest and lowest are two-fold degenerate. Thus, all energy-preserving unitaries consist of $(m-1)$ independently chosen unitaries from $SU(2)$ acting in the respective energy eigenspaces. For every single one of these eigenspaces, choosing a general unitary instead of a swap or the identity (as in π_J) will give a state that is a mixture of the two permutation cases, as above. Thus we can conclude that *all target states p' obtainable by energy preserving unitaries are in the convex hull of those that we obtain by energy preserving classical permutations*.

So the set of states p' that we can obtain with a truncated harmonic oscillator as a heat bath is

$$\text{conv} \left\{ p, p'_{\alpha_{\max}^{(m)}} \right\}.$$

This means that the set of achievable states p' is growing with increasing m , but never exhausts the set of all thermomajorized states: p^* , corresponding to $\alpha = e^{-\beta\Delta E}$, can never be achieved exactly for finite m . However, we can get arbitrarily close to p^* by increasing the dimension m since $\lim_{m \rightarrow \infty} \alpha_{\max}^{(m)} = e^{-\beta\Delta E}$. In particular, all states $p' \neq p^*$ which are thermomajorized by p can be obtained *exactly* with a finite heat bath of suitable size. This is also sketched in Figure 3.

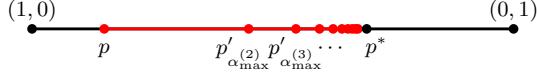


Figure 3: The states $p'_{\alpha_{\max}^{(m)}}$ for $m = 2, \dots, 10$. They converge to p^* in the limit $m \rightarrow \infty$, but they never reach it. That is, no finitely truncated harmonic oscillator can allow all possible state transitions on the qubit.

What if we choose any other (finite) heat bath Hamiltonian H_B which is *not* of the truncated harmonic oscillator form – could it be that we achieve the state p^* exactly? The following result answers this question in the negative, and gives a bound on how close we can get to p^* . It also says that truncated harmonic oscillators are in a certain sense optimal.

Lemma 8. *Let A be a qubit system with energy gap $\Delta E > 0$, and B a finite-dimensional heat bath with arbitrary Hamiltonian H_B . Then the stochastic maps D_α of the form (6) which can be implemented by thermal operations on this system and heat bath satisfy*

$$\alpha \leq e^{-\beta\Delta E} (1 - g_{\max}^B \gamma_{\max}^B), \quad (10)$$

where $\gamma_{\max}^B = e^{-\beta E_{\max}^B} / \text{tr}(e^{-\beta H_B})$ is the thermal occupation of the highest energy level of H_B , and g_{\max}^B is its degeneracy. Moreover, inequality (10) is tight, i.e. there is some thermal operation that achieves this value of α , if and only if

- for every energy eigenvalue E of H_B , except for the highest, also $E + \Delta E$ is an energy eigenvalue, and
- the degeneracies satisfy $g_E^B \geq g_{E+\Delta E}^B$.

In particular, given some classical initial state p on A , the state p^* (corresponding to $\alpha = e^{-\beta\Delta E}$ as in (7)) can never be created exactly from p with any finite heat bath. However, all other states that are thermomajorized by p can be obtained exactly with a finite heat bath.

Proof. Denote the orthonormal basis vectors of B by $|i\rangle$ and the corresponding energies by E_i (there may be $i \neq j$ with $E_i = E_j$). Furthermore, denote the set of “matching” energies of H_B by M_B , where an energy E is matching if there is some i such that $E_i = E$, and if there is some j such that $E_j = E - \Delta E$. That is, the matching energies $E \in M_B$ are those energy values which have a suitable energy gap downwards in order to allow for thermal operation in conjunction with the qubit on A .

To every $E \in M_B$, there is a corresponding global energy eigenspace

$$S_E := \text{span}\{|1i\rangle, |2j\rangle \mid E_i = E \text{ and } E_j = E - \Delta E\},$$

and clearly $E, E' \in M_B$ with $E \neq E'$ implies $S_E \perp S_{E'}$. Let P_E be the orthogonal projector onto S_E , and $P_c := \mathbf{1} - \sum_{E \in M_B} P_E$ the projector onto the orthogonal complement S_c of the span of all these (matching) energy eigenspaces.

Any energy-preserving unitary must be a direct sum of unitaries U_E acting on S_E and U_c acting on S_c . Define

$$(\hat{p} \otimes \hat{\gamma})_E := P_E (\hat{p} \otimes \hat{\gamma}) P_E$$

and similarly $(\hat{p} \otimes \hat{\gamma})_c$. Any thermal operation D_α acting on a probability vector p which can be implemented with the given heat bath can thus be written

$$(D_\alpha(p))^\wedge = \text{Tr}_B \left(\sum_{E \in M_B} U_E (\hat{p} \otimes \hat{\gamma})_E U_E^\dagger + U_c (\hat{p} \otimes \hat{\gamma})_c U_c^\dagger \right).$$

In particular, since thermal operations do not generate off-diagonal elements [5], the right-hand side will yield a diagonal density matrix on A . Due to (6), we can determine the corresponding parameter α by setting $p = (1, 0)$ and computing $\alpha = \langle 2 | (D_\alpha(p))^\wedge | 2 \rangle$. Since

$$(\hat{p} \otimes \hat{\gamma})_c = |1\rangle\langle 1| \otimes \sum_{i: E_i \notin M_B} \gamma_i^B |i\rangle\langle i|$$

with $\gamma_i^B := \exp(-\beta E_i)/Z$ and $Z = \text{tr}(e^{-\beta H_B})$, the energy-preserving unitary U_c cannot map any $|1i\rangle$ in this sum to any state of non-zero overlap with $|2j\rangle$ for any j . Thus, $\text{Tr}_B U_c (\hat{p} \otimes \hat{\gamma})_c U_c^\dagger = |1\rangle\langle 1|$, and so

$$\begin{aligned} \alpha &= \langle 2 | \left(\text{Tr}_B \sum_{E \in M_B} U_E (\hat{p} \otimes \hat{\gamma})_E U_E^\dagger \right) | 2 \rangle \\ &= \sum_{E \in M_B} \text{tr} \left(|2\rangle\langle 2| \otimes \mathbf{1}_B P_E U_E (\hat{p} \otimes \hat{\gamma})_E U_E^\dagger P_E \right) \\ &\stackrel{(*)}{\leq} \sum_{E \in M_B} \|P_E |2\rangle\langle 2| \otimes \mathbf{1}_B P_E\|_1 \cdot \|(\hat{p} \otimes \hat{\gamma})_E\|_\infty \\ &= \sum_{E \in M_B} g_{E-\Delta E}^B e^{-\beta E} / Z, \end{aligned}$$

where $g_{E'}^B$ is the degeneracy of the energy E' in H_B . Thus

$$\begin{aligned}\alpha &\leq \frac{1}{Z} \sum_{j: E_j + \Delta E \in M_B} e^{-\beta(E_j + \Delta E)} \\ &= \frac{1}{Z} e^{-\beta \Delta E} \sum_{j: E_j + \Delta E \in M_B} e^{-\beta E_j} \\ &\stackrel{(**)}{\leq} \frac{1}{Z} e^{-\beta \Delta E} \sum_{j: E_j \neq E_{\max}^B} e^{-\beta E_j} \\ &= \frac{1}{Z} e^{-\beta \Delta E} \left(Z - g_{\max}^B \cdot e^{-\beta E_{\max}^B} \right),\end{aligned}$$

where $g_{\max}^B \in \mathbb{N}$ is the degeneracy of the highest energy level of H_B . Inequality $(**)$ holds with equality if and only if for all eigenvalues $E \neq E_{\max}^B$ of H_B also $E + \Delta E$ is an eigenvalue. Furthermore, $(*)$ becomes an equality for some suitable choice of unitary if and only if every global eigenstate $|1i\rangle$ with $E_i = E$ can be unitarily mapped to another global energy eigenstate $|2j\rangle$ with $E_j = E - \Delta E$, which is possible by some suitable U_E if there are enough of the latter. \square

Lemma 8 has a compelling physical interpretation in terms of the *third law of thermodynamics*. Think of the task to cool a physical system. In general, this needs two ingredients: first, a heat bath, and second, a work reservoir. We can model both by considering a qubit system in a state $p = (0, 1)$, that is, an excited eigenstate of energy ΔE . Since p thermomajorizes all other classical qubit states p' , it is clear that this state contains enough free energy to allow a transition to any other state p' to arbitrary accuracy by coupling it to a suitable heat bath.

Suppose we would like to obtain another state

$$p' = (1, e^{-\beta' \Delta E}) / (1 + e^{-\beta' \Delta E}) \quad (11)$$

of inverse temperature $\beta' = 1/(k_B T')$. The third law of thermodynamics tells us that we cannot achieve temperature $T' = 0$ in finite time. More in detail, Nernst's formulation [16], first stated in 1912 [17], can be phrased as follows [18, 19]: *It is impossible by any procedure, no matter how idealized, to reduce any assembly to the absolute zero in a finite number of operations.* The thermal state of temperature zero is the ground state $p' = (1, 0)$ which is nothing but the state p^* from (7) (for our choice of initial state p). And, in fact, Lemma 8 says that we cannot obtain $p' = p^*$ exactly with any finite heat bath. But we can say more.

Theorem 9 (Third law of thermodynamics for qubits). *Let A be a qubit system with energy gap ΔE which is in its excited energy eigenstate, and B an n -dimensional heat bath with smallest and largest energy eigenvalues*

E_{\min}^B and E_{\max}^B , which is in a thermal state at temperature T . If we apply any thermal operation that yields a new qubit state p' with non-negative temperature $T' \geq 0$ (i.e. without population inversion), then

$$T' \geq \frac{T \Delta E}{E_{\max}^B - F_B} \geq \frac{T \Delta E}{E_{\max}^B - E_{\min}^B + k_B T \log n},$$

where $F_B = -k_B T \log \text{tr}(\exp(-H_B/(k_B T)))$ is the free energy of the heat bath.

Proof. If we apply a map D_α of the form (6) to the excited state $p = (0, 1)$, then $D_\alpha p = p'$ with p' as in (11) yields

$$\beta' = -\frac{1}{\Delta E} \log \left(\frac{e^{-\beta \Delta E}}{\alpha} - 1 \right),$$

and the argument inside the log is strictly positive due to $\alpha < e^{-\beta \Delta E}$. Furthermore, $\beta' \in [0, +\infty]$ implies that $\alpha \geq \frac{1}{2} e^{-\beta \Delta E}$. From eq. (10), it follows in particular that $\alpha \leq e^{-\beta \Delta E} (1 - \gamma_{\max}^B)$ (since $g_{\max}^B \geq 1$), and so

$$\begin{aligned}\beta' &\leq -\frac{1}{\Delta E} \log \left[(1 - \gamma_{\max}^B)^{-1} - 1 \right] \\ &= \frac{1}{\Delta E} \log \left((\gamma_{\max}^B)^{-1} - 1 \right) \leq -\frac{1}{\Delta E} \log \gamma_{\max}^B \\ &= \frac{1}{\Delta E} (\beta E_{\max}^B + \log Z_B),\end{aligned}$$

where $Z_B = \text{tr}(\exp(-\beta H_B))$, with $\beta = 1/(k_B T)$ the inverse temperature. Substituting the definition of F_B proves the first claimed inequality. The second one follows from the additional estimate

$$\begin{aligned}\log Z_B &= \log \sum_{i=1}^n e^{-\beta E_i^B} \leq \log \sum_{i=1}^n e^{-\beta E_{\min}^B} \\ &= \log n - \beta E_{\min}^B,\end{aligned}$$

where E_i^B denotes the energy eigenvalues of H_B . \square

In the special case of the truncated harmonic oscillator, with H_B as in (8) and thus $E_{\max}^B = (m-1)\Delta E$, we get

$$T' = \frac{\Delta E}{k_B} \left(\log \frac{e^{-\beta \Delta E} + e^{m\beta \Delta E}}{e^{\beta \Delta E} - 1} \right)^{-1} \approx \frac{T}{m} \quad (12)$$

for large m (to first order in $1/m$).

For weakly interacting multi-particle baths, the denominators in the lower bounds for T' in Theorem 9 will be extensive, that is, proportional to the particle number N . In physically realistic scenarios, only a polynomial number $N = \text{poly}(t)$ of particles can interact with the qubit system within a finite time interval $[0, t]$. Supplemented by these plausible physical assumptions, Theorem 9 tells us that

$$T' \gtrsim 1/t^r,$$

where t is the time in which the given thermodynamical protocol is supposed to be completed, and $r > 0$ is some model-dependent exponent. This result is compatible with the findings in [16], where the authors specify the exponent r in more detail, dependent on the spectrum of the heat bath. For example, they argue that $r = 7$ for a radiation bath in three spatial dimensions, which is in some sense the largest physically plausible r . However, the approach of [16] is somewhat different from the one in this paper. We will give a more detailed comparison, and some more physical intuition on our result, in the conclusions.

It is also interesting to see that the specific form of our lower bound in Theorem 9 is very close to the temperature obtained via a cooling method described in [16]. It relies on a different scenario that has been explored, for example, in [34], which is a variation of a traditional bit reset protocol. Given any two-level system with energy gap ΔE , raise the energy of the higher level isothermally while in contact with a heat bath of temperature T . Thermalization will bring the two-level system closer to its ground state. Traditionally, one would raise the energy level up to infinity, but suppose we stop at some E_{\max} , and subsequently reduce the higher energy level to its initial value. The resulting temperature of the qubit turns out to be $(T\Delta E)/(\Delta E + E_{\max})$, which is comparable to the bound in Theorem 9 (and in particular (12)).

IV.2 General systems with a finite heat bath

In the previous subsection, we have seen by example of a qubit that the classical target states which can be obtained with a fixed finite heat bath (in particular, a truncated harmonic oscillator) are the *convex hull* of those that can be achieved by classical permutation of energy levels on the same heat bath. We will now prove that this is true in general, up to some minor amendments.

To this end, let us introduce some notation. For any given initial classical state p (with corresponding diagonal quantum state \hat{p}) and Hamiltonians H_A, H_B , we denote by $\mathcal{T}_Q(p, H_A, H_B)$ the set of all classical final states p' that can be obtained by a suitable choice of energy-preserving unitary U on AB , i.e. $[U, H_A + H_B] = 0$, such that

$$\hat{p}' = \text{Tr}_B(U(\hat{p} \otimes \hat{\gamma}_B)U^\dagger), \quad (13)$$

where $\hat{\gamma}_B$ is the Gibbs state on B for inverse temperature β . In the case of non-degenerate H_A , the final state \hat{p}' will automatically be diagonal, because thermal operations cannot create coherences. If H_A exhibits degeneracies, we can always make \hat{p}' by implementing suitable unitary transformations in the degenerate subspaces. Furthermore, by $\mathcal{T}_C(p, H_A, H_B)$, we denote the

set of final states p' which can be obtained by a suitable choice of energy-preserving *permutation* π on AB such that

$$\hat{p}' = \text{Tr}_B(\pi(\hat{p} \otimes \hat{\gamma}_B)\pi^\dagger) \quad (14)$$

(which will always result in diagonal states). In other words, $\mathcal{T}_C(p, H_A, H_B)$ denotes the set of states that can be created from p **classically** (which is a finite set), whereas $\mathcal{T}_Q(p, H_A, H_B)$ denotes those states that can be obtained **quantum-mechanically**. (These sets also depend on the inverse temperature β , which we regard as fixed in this section.) We have the following result.

Theorem 10. *We have*

$$\mathcal{T}_Q(p, H_A, H_B) \subseteq \text{conv } \mathcal{T}_C(p, H_A, H_B). \quad (15)$$

If furthermore H_A is non-degenerate, then we have

$$\mathcal{T}_Q(p, H_A, H_B) = \text{conv } \mathcal{T}_C(p, H_A, H_B). \quad (16)$$

If H_A is degenerate, convex combinations of classical transitions can still be implemented exactly quantumly by adding an extra heat bath C with trivial Hamiltonian $H_C = 0$:

$$\text{conv } \mathcal{T}_C(p, H_A, H_B) \subseteq \mathcal{T}_Q(p, H_A, H_B + 0_C),$$

where $\dim C = 1 + \sum_i(d_i - 1) \leq \dim A$, and d_i denotes the dimensions of the (energy) eigenspaces of H_A .

An illustration of Theorem 10 (the non-degenerate case) is given in Figure 4. Note that the dimension of C could be chosen even smaller if it turned out that a generalization of quantum Horn's lemma along the lines of Lemma 6 was possible in general (for example, if Lemma 7 turned out to be tight).

Proof. Let $p' \in \mathcal{T}_Q(p, H_A, H_B)$, and let U be a unitary such that (13) holds. Since $[U, H_A + H_B] = 0$, we can decompose the total Hilbert space AB into a direct sum $\bigoplus_i S_i$, where the S_i are the energy eigenspaces of $H_A + H_B$. Then $U = \bigoplus_i U_i$, where the U_i are block matrices acting on the S_i .

Let $\lambda := p \otimes \gamma_B$. Since $\hat{\lambda}$ is diagonal in the energy eigenbasis, we can decompose $\hat{\lambda} = \bigoplus_i \hat{\lambda}_i$ accordingly, which means that $\lambda = \bigoplus_i \lambda_i$. Due to Lemma 1, we have

$$U\hat{\lambda}U^\dagger = \bigoplus_i U_i \hat{\lambda}_i U_i^\dagger = \bigoplus_i ((D_i \lambda_i)^\wedge + \Omega_i), \quad (17)$$

where Ω_i are matrices with all diagonal elements zero, and $D_i := U_i \circ U_i^*$ are bistochastic matrices. Let $\Omega := \bigoplus_i \Omega_i$, then all diagonal elements of Ω are zero, i.e. $\langle jk|\Omega|jk\rangle = 0$ for all j, k , and thus

$$\langle j|\text{Tr}_B \Omega |j\rangle = 0 \text{ for all } j.$$

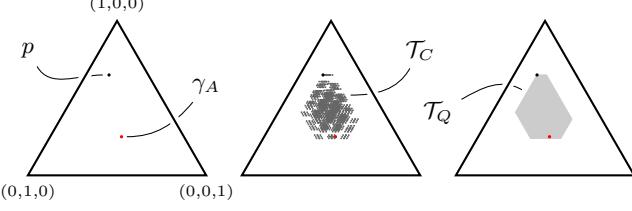


Figure 4: For a given initial state p in the classical probability simplex, given system Hamiltonian H_A , and bath Hamiltonian H_B , the set of states $\mathcal{T}_C(p, H_A, H_B)$ that can be obtained by classical thermal operations is a discrete set of points. In this picture we see an example calculation of this set for a three-level system where H_A is non-degenerate and $H_B = H_A + H_A$ (middle triangle). However, according to Theorem 10, since H_A is non-degenerate, the set of states that can be obtained via thermal quantum operations $\mathcal{T}_Q(p, H_A, H_B)$ is the convex hull of this set of states (right triangle), which is arguably a simpler and physically more realistic description. (Data used: $p = (65, 22, 13)/100$, $\gamma_A = (25, 35, 40)/100$, $\gamma_B = \gamma_A \otimes \gamma_A$.)

Every D_i is a bistochastic matrix acting on S_i , and thus, according to Birkhoff's Theorem [29], a convex combination of permutations $\pi_i^{(j)}$ acting on S_i . Therefore, according to Lemma 14 in the appendix, $D := \bigoplus_i D_i$ is a convex combinations of direct sums of permutations $\pi_i^{(k)}$ that act on the S_i , i.e.

$$D = \sum_k \mu_k \bigoplus_i \pi_i^{(k)}, \quad \mu_k \geq 0, \quad \sum_k \mu_k = 1.$$

Let $\pi^{(k)} := \bigoplus_i \pi_i^{(k)}$, which is an energy-preserving permutation due to the block matrix form. We have

$$\begin{aligned} p'_j &= \langle j | \text{Tr}_B U \hat{\lambda} U^\dagger | j \rangle = \langle j | \text{Tr}_B [(D\lambda)^\wedge] | j \rangle \\ &= \sum_k \mu_k \langle j | \text{Tr}_B \left[(\pi^{(k)} \lambda)^\wedge \right] | j \rangle \\ &= \sum_k \mu_k \langle j | \text{Tr}_B \left[\pi^{(k)} \hat{\lambda} (\pi^{(k)})^\dagger \right] | j \rangle, \end{aligned}$$

and since \hat{p}' is diagonal, this leads to

$$\hat{p}' = \sum_k \mu_k \text{Tr}_B \left[\pi^{(k)} (\hat{p} \otimes \hat{\gamma}) (\pi^{(k)})^\dagger \right] \quad (18)$$

which is a convex combination of outputs of maps of the form (14). Thus $p' \in \text{conv } \mathcal{T}_C(p, H_A, H_B)$, which proves (15).

Now we would like to show the second half of the theorem, that every $p' \in \text{conv } \mathcal{T}_C(p, H_A, H_B)$ can be obtained quantum-mechanically. So suppose there are energy-preserving permutations $\pi^{(k)}$ and coefficients μ_k such that (18) holds. Energy preservation implies that

$\pi^{(k)} = \bigoplus_i \pi_i^{(k)}$, where the $\pi_i^{(k)}$ act on the global energy eigenspaces S_i . Let us define

$$\hat{\lambda}' = \sum_k \mu_k \pi^{(k)} \hat{p} \otimes \hat{\gamma} (\pi^{(k)})^\dagger = \sum_k \mu_k \pi^{(k)} \hat{\lambda} (\pi^{(k)})^\dagger$$

which gives $\hat{p}' = \text{Tr}_B \hat{\lambda}'$. As before, we can decompose $\lambda = \bigoplus_i \lambda_i$ and $\lambda' = \bigoplus_i \lambda'_i$ into blocks corresponding to the global energy eigenspaces, and it follows that $\lambda'_i = \sum_k \mu_k \pi_i^{(k)} \lambda_i$, and so $\lambda_i \succ \lambda'_i$, i.e. λ_i majorizes λ'_i by a general result on majorization [30]. Thus, according to [30, Thm. B.6], there are unistochastic maps D_i such that $\lambda'_i = D_i \lambda_i$ (note that D_i need not be equal to $\sum_k \mu_k \pi_i^{(k)}$), and therefore unitary matrices U_i such that $D_i = U_i \circ U_i^*$. We set $U := \bigoplus_i U_i$ and by Lemma 1, as in (17), we will have

$$U \hat{\lambda} U^\dagger = \hat{\lambda}' + \Omega,$$

where $\Omega = \bigoplus_i \Omega_i$ is a matrix with all diagonal elements zero. It follows that

$$\langle j | \text{Tr}_B (U \hat{\lambda} U^\dagger) | j \rangle = \langle j | \text{Tr}_B \hat{\lambda}' | j \rangle + \sum_k \langle j k | \Omega | j k \rangle = p'_j.$$

For thermal operations, we know that $\text{Tr}_B (U \hat{\lambda} U^\dagger)$ must be block-diagonal [5] in the energy eigenbasis due to energy conservation. Thus, if H_A is non-degenerate, it follows that this matrix is diagonal in the unique energy eigenbasis, and we have

$$\hat{p}' = \text{Tr}_B [U (\hat{p} \otimes \hat{\gamma}) U^\dagger],$$

which proves (16). However, if H_A is degenerate, we do not necessarily obtain a diagonal matrix. We deal with this by introducing an auxiliary system used to implement a decoherence map that sets the off-diagonal elements of $\rho'_A := \text{Tr}_B (U \hat{\lambda} U^\dagger)$ to zero. Since ρ'_A is block-diagonal in the energy eigenbasis, with block sizes (at most) d_i , Lemma 11 below shows that we can implement the decoherence map by adding a fully degenerate auxiliary system of dimension $1 + \sum_i (d_i - 1)$ and performing a suitable thermal operation. \square

Note that this theorem contains Lemma 4 from Section III as a special case: in the case of a trivial $n \times n$ Hamiltonian $H_A = 0$, we have a single degeneracy $d_1 = n$. Due to Birkhoff's Theorem [29], all bistochastic maps can be implemented as convex combinations of classical permutations on system A alone. Therefore, denoting a trivial heat bath B of dimension $\dim B = 1$ (and arbitrary Hamiltonian $H_B \in \mathbb{R}$) simply by a bullet, we have

$$\text{conv } \mathcal{T}_C(p, H_A, \bullet) = \{p' \mid p \succ p'\}.$$

Then Theorem 10 tells us that we can implement all these transitions exactly by noisy quantum operations with a heat bath C of size

$$\dim C = 1 + (d_1 - 1) = n.$$

The following lemma has been used in the proof of Theorem 10.

Lemma 11 (Thermal decoherence). *Let A be an n -dimensional quantum system with Hamiltonian H_A , and $\{|i\rangle\}_{i=1}^n$ be an arbitrary orthonormal eigenbasis of H_A . Furthermore, suppose that ρ_A is a block-diagonal density matrix on A , i.e. $[\rho_A, H_A] = 0$. Then there is a heat bath B with trivial Hamiltonian $H_B = 0$ of dimension*

$$\dim B = 1 + \sum_i (d_i - 1) \leq n,$$

where d_i denotes the dimensions of the energy eigenspaces of H_A , such that there is a thermal operation on A with this heat bath B that maps ρ_A to ρ'_A , where $\langle i|\rho'_A|i\rangle = \langle i|\rho_A|i\rangle$, and $\langle i|\rho'_A|j\rangle = 0$ for $i \neq j$.

In other words: small heat baths allow to decohere block-diagonal quantum states. If H_A is non-degenerate, i.e. all $d_i = 1$, then ρ_A is already diagonal and we do not need any heat bath at all, i.e. $\dim B = 1$.

Proof. We construct a thermal operation Φ on A which sets suitable matrix elements of the input density matrix ρ to zero. Suppose we are given any subset $S \subset \{1, 2, \dots, n\}$, then we will construct Φ such that it sets all elements in rows and columns appearing in S to zero (except for the diagonal elements) and leaves all other matrix elements invariant. For example, suppose that $n = 4$, and that $S = \{2, 4\}$. Then we will construct a thermal operation Φ such that

$$\begin{pmatrix} \rho_{11} & \dots & \rho_{14} \\ \vdots & & \vdots \\ \rho_{41} & \dots & \rho_{44} \end{pmatrix} \xrightarrow{\Phi} \begin{pmatrix} \rho_{11} & 0 & \rho_{13} & 0 \\ 0 & \rho_{22} & 0 & 0 \\ \rho_{31} & 0 & \rho_{33} & 0 \\ 0 & 0 & 0 & \rho_{44} \end{pmatrix}.$$

Clearly, if ρ is already block-diagonal with block sizes d_i (e.g. in this case $d_1 = 2$ and $d_2 = 2$), then this map fully decoheres ρ , as long as we decohere $(d_i - 1)$ rows and columns in every i -th block. In other words, we will choose the subset S such that

$$|S| = \sum_i (d_i - 1),$$

and if Φ acts as stated, we achieve the decoherence as claimed in the statement of the lemma. Let $B = \mathbb{C}^d$, $d = |S| + 1$, and $|1\rangle, \dots, |d\rangle$ some orthonormal basis on B . Define the permutation matrix π by $\pi|i\rangle = |i+1\rangle$ if

$1 \leq i \leq d-1$ and $\pi|d\rangle = |1\rangle$. We have $\pi^T = \pi^\dagger = \pi^{-1}$, and, for $j \in \mathbb{Z}$,

$$\text{tr}(\pi^j) = \begin{cases} d & \text{if } j \in \{\dots, -2d, -d, 0, d, 2d, \dots\} \\ 0 & \text{otherwise.} \end{cases}$$

Label the elements of S (in increasing order) by $S = \{s_1, s_2, \dots, s_{d-1}\}$, and define the unitary U on AB by

$$U = \sum_{i \notin S} |i\rangle\langle i| \otimes \mathbf{1} + \sum_{j=1}^{|S|} |s_j\rangle\langle s_j| \otimes \pi^j.$$

Clearly $[U, H_A \otimes \mathbf{1}] = 0$, so if we regard B as a quantum system with trivial Hamiltonian $H_B = 0$, the following map $\Phi : \rho \mapsto \rho'$ is a thermal operation:

$$\rho' = \text{Tr}_B \left[U \left(\rho \otimes \frac{\mathbf{1}}{d} \right) U^\dagger \right].$$

It is easy to check by direct calculation that

$$\rho'_{ij} = \begin{cases} \rho_{ij} & \text{if } i = j \text{ or } (i \notin S \text{ and } j \notin S), \\ 0 & \text{otherwise,} \end{cases}$$

where $\rho_{ij} = \langle i|\rho|j\rangle$, and similarly for ρ'_{ij} . For example, if $i \notin S$ and $j \in S$, then $\langle ik|U = \langle ik|$, and

$$U^\dagger|jk\rangle = \sum_{m=1}^{|S|} |s_m\rangle\langle s_m|j\rangle \otimes \pi^{-m}|k\rangle = |j\rangle \otimes \pi^{-m}|k\rangle,$$

where m is such that $s_m = j$. Thus, we get

$$\begin{aligned} \rho'_{ij} &= \frac{1}{d} \sum_k \langle ik|\rho \otimes \mathbf{1}|j\rangle \otimes \pi^{-m}|k\rangle \\ &= \frac{1}{d} \rho_{ij} \sum_k \langle k|\pi^{-m}|k\rangle = \rho_{ij} \frac{1}{d} \text{tr}(\pi^{-m}) = 0. \end{aligned}$$

All other cases of i and j can be checked similarly. \square

Theorem 10 implies in particular the following result, which is illustrated in Figure 5.

Corollary 12. *Let p be any classical state on a quantum system with some Hamiltonian H_A , and $\mathcal{T}(p, H_A)$ the (compact convex) set of all classical states that are thermomajorized by p . Then all states $p' \in \text{relint } \mathcal{T}(p, H_A)$ (i.e. in the relative interior) can be obtained exactly by quantum thermal operations with finite heat baths, but not in general all states in the relative boundary $\text{rebd } \mathcal{T}(p, H_A)$.*

Proof. We use Lemma 3 which shows that all states $p' \in \mathcal{T}(p, H_A)$ can be obtained to arbitrary accuracy by classical thermal operations with suitable heat baths. Formally, if we denote the set of all states that can be

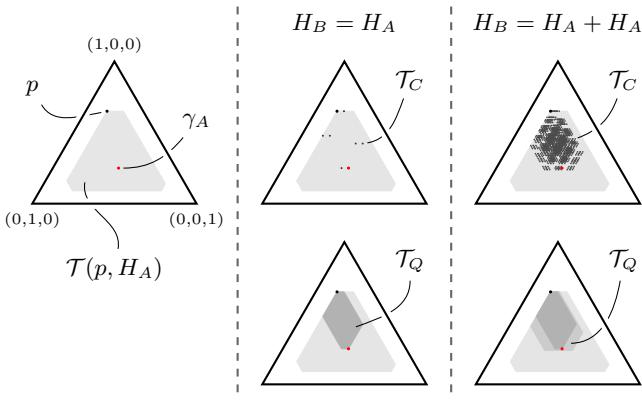


Figure 5: The set of (classical) states in the probability simplex which are thermomajorized by p , denoted $\mathcal{T}(p, H_A)$, is a compact convex set (light gray). If H_A is non-degenerate, then the set of states $\mathcal{T}_Q(p, H_A, H_B)$ that one can obtain with a fixed, finite heat bath, by quantum thermal operations, is a compact convex subset (dark gray, second row). With increasing resources (for example, by having several copies of the heat bath), this subset will grow (third column), and finally cover every point in the relative interior of $\mathcal{T}(p, H_A)$. However, some points on the boundary will never be covered (we have proven this for the qubit in Lemma 8, and conjecture that it is also true for higher-level systems). This can be interpreted as an instance of the third law of thermodynamics.

obtained from p classically with any possible heat bath by

$$\mathcal{T}_C(p, H_A) := \bigcup_{\dim B \in \mathbb{N}, H_B} \mathcal{T}_C(p, H_A, H_B),$$

then $\mathcal{T}_C(p, H_A)$ is dense in $\mathcal{T}(p, H_A)$. It follows with a bit of convex geometry (Lemma 15 in the appendix) that

$$\text{relint } \mathcal{T}(p, H_A) \subseteq \text{conv } \mathcal{T}_C(p, H_A).$$

By construction, if $p' \in \text{conv } \mathcal{T}_C(p, H_A)$, then there exist $m \in \mathbb{N}$, $\lambda_1, \dots, \lambda_m \geq 0$ with $\sum_i \lambda_i = 1$, heat baths B_i with Hamiltonians H_{B_i} , and states $q^{(i)}$ on A such that $p' = \sum_{i=1}^m \lambda_i q^{(i)}$ and $q^{(i)} \in \mathcal{T}_C(p, H_A, H_{B_i})$. If we have a composite heat bath on the tensor product $B_1 \otimes \dots \otimes B_m$ with Hamiltonian $H_{B_1} + \dots + H_{B_m}$, then

$$\mathcal{T}_C(p, H_A, H_{B_i}) \subseteq \mathcal{T}_C(p, H_A, H_{B_1} + \dots + H_{B_m}),$$

since we can always perform a classical permutation that ignores (acts as the identity on) some of the tensor factors. Thus, $q^{(i)} \in \mathcal{T}_C(p, H_A, H_{B_1} + \dots + H_{B_m})$, and so $p' \in \text{conv } \mathcal{T}_C(p, H_A, H_{B_1} + \dots + H_{B_m})$ which, according to Theorem 10, is contained in $\mathcal{T}_Q(p, H_A, H_{B_1} + \dots + H_{B_m} + 0_C)$, with C a trivial heat bath of suitable finite size. Therefore, all elements of $\text{relint } \mathcal{T}(p, H_A)$ can be obtained exactly by quantum thermal operations with finite heat baths.

On the other hand, states in the relative boundary cannot always be obtained exactly. This is shown by example in Lemma 8. \square

V Conclusions

The study of quantum operations with finite heat baths reveals some interesting mathematical structure and physical insights. In this paper, we have analyzed basic properties of noisy and thermal operations with finite heat baths. In order to study whether some state transitions are unattainable with finite baths (resembling the third law of thermodynamics), we first had to remove an unphysical apparent “unattainability” that is simply a mathematical artefact of a semiclassical treatment: the fact that infinite auxiliary systems are needed for some simple state transitions if one restricts to classical permutations in the energy eigenbasis.

We did this by proving Corollary 12, which is also sketched in Figure 5: the target states that can be obtained from a given (classical) initial state with unbounded heat baths is a compact convex set. All the states in its interior can be obtained exactly via thermal quantum operations with finite heat baths. It is then (some of) the states in the boundary of this set that become the subject of potential unattainability results.

In the case of a two-level quantum system, this boundary consists of two states (cf. Figure 3): the initial state p itself, and another state p^* that is in some sense “thermodynamically antipodal” to p . In fact, we have shown that p^* can never be obtained exactly from p with any finite heat bath. In the case of a pure excited initial state, this led us to an instance of the third law of thermodynamics (Theorem 9), giving a (physically meaningful) lower bound on the temperature to which the system can be cooled by using any finite-dimensional heat bath. While we have formally proven this result only for qubits, it seems plausible that an analogous statement holds for higher-dimensional quantum systems, too. This gives an intuitive structural picture of unattainability in terms of the geometry of the state space (cf. Figure 5), and it suggests a generalization of the third law to situations beyond cooling to absolute zero (since p^* is not in general the ground state).

We can elucidate the relation to other recent unattainability results by elaborating on the physical intuition of our result. Consider cooling with thermal operations, as in the scenario of Figure 1. Denoting the ground state of the qubit by $|1\rangle$ and the excited state by $|2\rangle$, we would like to apply an energy-preserving unitary to the global initial state $|2\rangle\langle 2| \otimes \hat{\gamma}_B$ in a way that moves as much probability weight as possible (in fact, all of it) into the local ground state, i.e. the subspace spanned by the global eigenstates $|1i\rangle$. However, if $|m\rangle$

denotes the maximal energy level of the bath (assuming for the moment that it is unique), then the (globally maximal) energy level $|2m\rangle$ is “trapped”: energy preservation forbids to move its probability weight into any other subspace. And since the bath starts in a thermal state, this “trapped” population will be non-zero, and it will not be transferred into the ground state. We expect that this simple intuition will remain robust also in more complicated cases beyond the simple qubit scenario, for example in cases where the quantum system A is itself composite, $A = A_1 A_2$. This encompasses in particular situations in which A_1 is the system of interest (say, a quantum system to be cooled), while A_2 is a work storage system, such as a weight [54] or work bit [5].

This “inability to move probability weight around” in high-lying energy levels of the heat bath is also a major ingredient of the version of the third law proven in [16]. In our setting, however, this inability comes from an assumption of finite-dimensionality of the bath, while in [16], it is the bath’s spectral density together with the assumption of a strict upper bound on the available work w_{\max} that limits the available energy levels. Our approach should not be understood as claiming that physical heat baths are generically finite-dimensional. Instead, the motivation is to give a very general formulation of the idea that we can only engineer interactions with a finite number of energy levels of the bath in a controlled way within a thermodynamic protocol. If understood in this way, the approach of [16] is conceptually related to ours (but [16] goes into much more physical detail).

Our unattainability statement (and the one in [16]) is also different from the scenario studied by Janzing et al. in [1], where they consider the amount of non-thermal *resource states* needed to cool a system (close) to absolute zero. In contrast, our formulation assumes that all non-thermal resources are in principle there (e.g. the energy of the excited state in Figure 1 is enough for cooling to zero), but they can only be harvested by thermodynamic protocols in the presence of an unbounded heat bath.

Several further recent publications have considered limitations for cooling that arise from finite heat baths, cf. [50] and [51] (see also [52] and [53]), but within approaches that are not formulated within the resource-theoretic framework of thermal operations. Instead, arbitrary (in general energy-non-preserving) unitaries are allowed, and the energy cost or yield is taken into account in retrospect, by bookkeeping the system’s average energy content. That is, the fundamental microscopic reversibility of the physical laws is taken into account directly by restricting to unitary transformations, but the principle of total energy conservation is

only considered indirectly.

In the context of the third law, this leads to a subtle restriction that comes from the resulting focus on entropic properties. Typically, in approaches that allow arbitrary unitaries, the impossibility of cooling to absolute zero is ultimately attributed to a simple property of linear algebra: namely, that *unitary operations cannot create zero eigenvalues from non-zero eigenvalues*. In contrast, our setup yields more severe restrictions that rely on microscopic reversibility *in conjunction with* total energy preservation. The simplest example is the one depicted in Figure 1: even though we start with a pure state of the system, the pure ground state is still unattainable.

Our paper leaves several interesting open questions. First, we have shown by example in Lemma 6 that quantum Horn’s lemma is true (Lemma 4), but not optimal: for example, all noisy state transitions on a three-dimensional quantum system can be accomplished by means of a two-dimensional auxiliary system. What is the smallest possible bath dimension in general? Is the bound in Lemma 7 tight? If so, then this would also improve our more general result for non-trivial Hamiltonians, Theorem 10. As we have shown, this question is closely related to a version of the quantum marginal problem.

From the point of view of the physical interpretation, the most interesting open problem is to analyze in more detail whether, and if so in what way, our third law for qubits (Theorem 9) generalizes to higher-dimensional quantum systems and scenarios with a work storage device in a physically meaningful way. Restricting the dimension of the heat bath, as we have done here, is only one of several possibilities to study unattainability results in thermodynamics. But the picture of some states, not necessarily of temperature zero, that lie “on the edge” of the set of attainable states (as in Figure 5) is a quite suggestive hint towards possible generalizations beyond perfect cooling.

Acknowledgments

We are grateful to Manfred Salmhofer for discussions, and we would like to thank Michael Walter and Matthias Christandl for drawing our attention to [48] and [49].

References

- [1] D. Janzing, P. Wocjan, R. Zeier, R. Geiss, and Th. Beth, *Thermodynamic Cost of Reliability and Low Temperatures: Tightening Landauer’s Principle*

- ciple and the Second Law*, Int. J. Theor. Phys. **39**(12), 2717–2753 (2000).
- [2] M. Horodecki, P. Horodecki, and J. Oppenheim, *Reversible transformations from pure to mixed states and the unique measure of information*, Phys. Rev. A **67**, 062104 (2003).
- [3] O. C. O. Dahlsten, R. Renner, E. Rieper, and V. Vedral, *Inadequacy of von Neumann entropy for characterizing extractable work*, New J. Phys. **13**, 053015 (2011).
- [4] L. del Rio, J. Åberg, R. Renner, O. C. O. Dahlsten, and V. Vedral, *The thermodynamic meaning of negative entropy*, Nature **474**, 61 (2011).
- [5] M. Horodecki and J. Oppenheim, *Fundamental limitations for quantum and nanoscale thermodynamics*, Nat. Comm. **4**, 2059 (2013).
- [6] F. G. S. L. Brandão, M. Horodecki, J. Oppenheim, J. M. Renes, and R. W. Spekkens, *Resource Theory of Quantum States Out of Thermal Equilibrium*, Phys. Rev. Lett. **111**, 250404 (2013).
- [7] J. Åberg, *Truly work-like work extraction via a single-shot analysis*, Nat. Comm. **4**, 1925 (2013).
- [8] P. Skrzypczyk, A. J. Short, and S. Popescu, *Work extraction and thermodynamics for individual quantum systems*, Nat. Comm. **5**, 4185 (2014).
- [9] J. M. Renes, *Work cost of thermal operations in quantum thermodynamics*, Eur. Phys. J. Plus **129**:153 (2014).
- [10] P. Faist, F. Dupuis, J. Oppenheim, and R. Renner, *The minimal work cost of information processing*, Nat. Comm. **6**, 7669 (2015).
- [11] F. G. S. L. Brandão, M. Horodecki, N. Ng, J. Oppenheim, and S. Wehner, *The second laws of quantum thermodynamics*, Proc. Natl. Acad. Sci. USA **112**, 3275 (2015).
- [12] N. Yunger Halpern, A. J. P. Garner, O. C. O. Dahlsten, and V. Vedral, *Introducing one-shot work into fluctuation relations*, New J. Phys. **17**, 095003 (2015).
- [13] D. Egloff, O. C. O. Dahlsten, R. Renner, and V. Vedral, *A measure of majorization emerging from single-shot statistical mechanics*, New J. Phys. **17**, 073001 (2015).
- [14] M. Lostaglio, M. P. Müller, and M. Pastena, *Stochastic independence as a resource in small-scale thermodynamics*, Phys. Rev. Lett. **115**, 150402 (2015).
- [15] J. Gemmer and J. Anders, *From single-shot towards general work extraction in a quantum thermodynamic framework*, New J. Phys. **17**, 085006 (2015).
- [16] Ll. Masanes and J. Oppenheim, *A general derivation and quantification of the third law of thermodynamics*, Nat. Comm. **8**, 14538 (2017).
- [17] W. Nernst, Sitzungsbericht der Königlich Preussischen Akademie der Wissenschaften, p. 134 (1912).
- [18] E. M. Loeb, *The Third Law of Thermodynamics, the Unattainability of Absolute Zero, and Quantum Mechanics*, J. Chem. Educ. **37**(7), 361–363 (1960).
- [19] R. Fowler and E. A. Guggenheim, *Statistical Thermodynamics*, Cambridge University Press, Cambridge, 1949.
- [20] Y. Guryanova, S. Popescu, A. J. Short, R. Silva, and P. Skrzypczyk, *Thermodynamics of quantum systems with multiple conserved quantities*, Nat. Comm. **7**, 12049 (2016).
- [21] M. Lostaglio, D. Jennings, and T. Rudolph, *Description of quantum coherence in thermodynamic processes requires constraints beyond free energy*, Nat. Comm. **6**, 6383 (2015).
- [22] M. Lostaglio, K. Korzekwa, D. Jennings, and T. Rudolph, *Quantum Coherence, Time-Translation Symmetry, and Thermodynamics*, Phys. Rev. X **5**, 021001 (2015).
- [23] K. Korzekwa, M. Lostaglio, J. Oppenheim, and D. Jennings, *The extraction of work from quantum coherence*, New J. Phys. **18**, 023045 (2016).
- [24] P. Ćwikliński, M. Studziński, M. Horodecki, and J. Oppenheim, *Limitations on the Evolution of Quantum Coherences: Towards Fully Quantum Second Laws of Thermodynamics*, Phys. Rev. Lett. **115**, 210403 (2015).
- [25] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, *Quantum entanglement*, Rev. Mod. Phys. **81**, 865 (2009).
- [26] B. Coecke, T. Fritz, and R. W. Spekkens, *A mathematical theory of resources*, Information and Computation **250**, 59–86 (2016).
- [27] F. G. S. L. Brandão and G. Gour, *Reversible Framework for Quantum Resource Theories*, Phys. Rev. Lett. **115**, 070503 (2015).
- [28] L. del Rio, L. Krämer, and R. Renner, *Resource theories of knowledge*, arXiv:1511.08818.
- [29] R. Bhatia, *Matrix Analysis*, Springer Verlag, New York, 1997.

- [30] A. W. Marshall, I. Olkin, and B. C. Arnold, *Inequalities: Theory of Majorization and Its Applications*, Springer Verlag, New York, 2011.
- [31] R. Webster, *Convexity*, Oxford University Press, Oxford, 1994.
- [32] G. Gour, M. P. Müller, V. Narasimhachar, R. W. Spekkens, and N. Yunger Halpern, *The resource theory of informational nonequilibrium in thermodynamics*, *Phys. Rep.* **583**, 1–58 (2015).
- [33] A. Horn, *Doubly stochastic matrices and the diagonal of a rotation matrix*, *Amer. J. Math.* **76**, 620–630 (1954).
- [34] C. Browne, A. J. P. Garner, O. C. O. Dahlsten, and V. Vedral, *Guaranteed Energy-Efficient Bit Reset in Finite Time*, *Phys. Rev. Lett.* **113**, 100603 (2014).
- [35] K. Życzkowski and I. Bengtsson, *On Duality between Quantum Maps and Quantum States*, *Open Syst. Inf. Dyn.* **11**, 3 (2004).
- [36] I. Bengtsson and K. Życzkowski, *Geometry of Quantum States*, Cambridge University Press, Cambridge, 2006.
- [37] M. Musz, M. Kuś, and K. Życzkowski, *Unitary quantum gates, perfect entanglers, and unistochastic maps*, *Phys. Rev. A* **87**, 022111 (2013).
- [38] E. Ruch and A. Mead, *The principle of increasing mixing character and some of its consequences*, *Theoret. Chim. Acta (Berl.)* **41**, 95–117 (1976).
- [39] R. Ruch, R. Schranner, and T. H. Seligman, *The mixing distance*, *J. Chem. Phys.* **69**, 386 (1978).
- [40] E. Ruch, R. Schranner, and T. H. Seligman, *Generalization of a Theorem by Hardy, Littlewood, and Pólya*, *J. Math. Anal. Appl.* **76**, 222–229 (1980).
- [41] P. Shor, *Structure of Unital Maps and the Asymptotic Quantum Birkhoff Conjecture*, presentation, slides at <http://science-visits.mccme.ru/doc/steklov-talk.pdf>, 2010.
- [42] U. Haagerup and M. Musat, *Factorization and dilation problems for completely positive maps on von Neumann algebras*, *Commun. Math. Phys.* **303**(2), 555–594 (2011).
- [43] P. Faist, J. Oppenheim, and R. Renner, *Gibbs-preserving maps outperform thermal operations in the quantum regime*, *New J. Phys.* **17**, 043003 (2015).
- [44] C. Perry, P. Ćwikliński, J. Anders, M. Horodecki, and J. Oppenheim, *A sufficient set of experimentally implementable thermal operations*, [arXiv:1511.06553](https://arxiv.org/abs/1511.06553).
- [45] N. H. Y. Ng, M. Mančinska, C. Cirstoiu, J. Eisert, and S. Wehner, *Limits to catalysis in quantum thermodynamics*, *New J. Phys.* **17**, 085004 (2015).
- [46] S. Daftuar and P. Hayden, *Quantum state transformations and the Schubert calculus*, *Ann. Phys.* **315**(1), 80–122 (2005).
- [47] J. Scharlau, *The resource theories of non-uniformity and athermality: State transitions with finite size environments*, Master thesis, Heidelberg University, 2015.
- [48] C. Li, Y. Poon, and X. Wang, *Ranks and eigenvalues of states with prescribed reduced states*, *Electron. J. Linear Algebra* **27**, 1074 (2014).
- [49] M. Vergne and M. Walter, *Inequalities for Moment Cones of Finite-Dimensional Representations*, *Journal of Symplectic Geometry* **15**(4), 1209–1250 (2017).
- [50] D. Reeb and M. M. Wolf, *An improved Landauer Principle with finite-size corrections*, *New J. Phys.* **16**, 103011 (2014).
- [51] A. E. Allahverdyan, K. V. Hovhannisan, D. Janzing, and G. Mahler, *Thermodynamic limits of dynamic cooling*, *Phys. Rev. E* **84**, 041109 (2011).
- [52] L.-A. Wu, D. Segal, and P. Brumer, *No-go theorem for ground state cooling given initial system-thermal bath factorization*, *Sci. Rep.* **3**, 1824 (2013).
- [53] F. Ticozzi and L. Viola, *Quantum resources for purification and cooling: fundamental limits and opportunities*, *Sci. Rep.* **4**, 5192 (2014).
- [54] P. Skrzypczyk, A. J. Short, and S. Popescu, *Extracting work from quantum systems*, [arXiv:1302.2811](https://arxiv.org/abs/1302.2811).

VI Appendix

The following lemma gives a partial classification of the set of *maps* that can be implemented via noisy operations on any n -dimensional system A with an n -dimensional bath B . Since $p \succ p'$ implies that there is a unistochastic map D with $Dp = p'$ [30, Thm. B.6], the lemma below implies Lemma 4 as a simple corollary. It can therefore be interpreted as yet another generalization of quantum Horn’s lemma.

Lemma 13. Let $\text{Uni}(n)$ and $\text{Bi}(n)$ be the sets of unistochastic resp. bistochastic maps on \mathbb{R}^n , i.e.

$$\begin{aligned}\text{Uni}(n) &:= \left\{ D \in \mathbb{R}^{n \times n} \mid \exists U : U^\dagger U = \mathbf{1}, D_{ij} = |U_{ij}|^2 \right\}, \\ \text{Bi}(n) &:= \left\{ D \in \mathbb{R}^{n \times n} \mid \sum_i D_{ij} = \sum_j D_{ij} = 1, D_{ij} \geq 0 \right\}.\end{aligned}$$

Moreover, let $\text{Noisy}(n, n)$ be the set of maps $p \mapsto p'$ on \mathbb{R}^n which can be written in the form

$$\hat{p}' = \text{Tr}_B \left[U \left(\hat{p} \otimes \frac{\mathbf{1}}{n} \right) U^\dagger \right]$$

with some unitary U on $\mathbb{C}^n \otimes \mathbb{C}^n$. Then we have

$$\begin{aligned}\text{Uni}(2) &= \text{Noisy}(2, 2) = \text{Bi}(2), \\ \text{Uni}(n) &\subsetneq \text{Noisy}(n, n) \subseteq \text{Bi}(n) \quad \text{for } n \geq 3.\end{aligned}$$

Moreover, the inclusion $\text{Noisy}(n, n) \subseteq \text{Bi}(n)$ is strict if and only if $\text{Noisy}(n, n)$ is not convex.

Remark. We do not currently know whether $\text{Noisy}(n, n)$ is convex for $n \geq 3$, which would imply that $\text{Noisy}(n, n) = \text{Bi}(n)$.

Proof. If $D \in \text{Uni}(n)$ then there is a unitary matrix V such that $D_{ij} = |V_{ij}|^2$. Thus, for any state $p \in \mathbb{R}^n$, we have

$$\langle i | V \hat{p} V^\dagger | i \rangle = \sum_j V_{ij} p_j V_{ij}^* = (Dp)_i;$$

in other words, $V \hat{p} V^\dagger$ has the vector Dp on its diagonal. We can compose V with the unitary W from the proof of Lemma 4 which implements the decoherence map to remove all off-diagonal elements with an n -dimensional bath. Then $U = W(V \otimes \mathbf{1}_B)$ will realize the map D as a noisy operation. This proves that $\text{Uni}(n) \subseteq \text{Noisy}(n, n)$. On the other hand, for any given operation in $\text{Noisy}(n, n)$ with associated unitary U , define

$$D_{ik} := \frac{1}{n} \langle i | \text{Tr}_B [U (|k\rangle\langle k| \otimes \mathbf{1}) U^\dagger] | i \rangle,$$

then it is straightforward to check that $\sum_i D_{ik} = \sum_k D_{ik} = 1$ and $D_{ik} \geq 0$. Moreover, if \hat{p}' is the output of the noisy operation on input \hat{p} , then $p' = Dp$. This shows that $\text{Noisy}(n, n) \subseteq \text{Bi}(n)$. For $n = 2$, we have $\text{Uni}(2) = \text{Bi}(2)$, which proves the first statement of the lemma.

For $n \geq 3$, we construct a matrix $D \in \text{Bi}(n) \setminus \text{Uni}(n)$ which can be implemented by a noisy operation. Let π be the cyclic permutation on \mathbb{C}^n from the proof of Lemma 4, and define

$$D := \left(1 - \frac{1}{n} \right) \mathbf{1} + \frac{1}{n} \pi.$$

Obviously D is bistochastic. To see that it is not unistochastic, suppose U is a unitary matrix with $D_{ij} = |U_{ij}|^2$, then orthogonality of the first two rows of U is $\sum_{j=1}^n U_{1j}^* U_{2j} = 0$. But for every $j \neq 1$, either $D_{1j} = 0$ or $D_{2j} = 0$, hence this sum equals $U_{11}^* U_{21} = e^{i\theta} \sqrt{D_{11} D_{21}} = e^{i\theta} \sqrt{(1 - \frac{1}{n}) \frac{1}{n}}$ for some $\theta \in \mathbb{R}$, and this is not zero. Define the unitary

$$U = \mathbf{1} \otimes \sum_{k=2}^n |k\rangle\langle k| + \pi \otimes |1\rangle\langle 1|.$$

A simple calculation shows that this unitary generates a noisy operation which maps \hat{p} to $(Dp)^\wedge$. This shows that $D \in \text{Noisy}(n, n)$, but $D \notin \text{Uni}(n)$.

Finally, since $\text{Bi}(n)$ is convex, non-convexity of $\text{Noisy}(n, n)$ would imply that $\text{Noisy}(n, n) \subsetneq \text{Bi}(n)$. Conversely, suppose that this strict inclusion was true (note that we do not know whether this is the case). Since every permutation is unistochastic, and the bistochastic maps are the convex hull of the permutations due to Birkhoff's Theorem [29], $\text{Noisy}(n, n)$ could then not be convex, because otherwise the permutations that it contains would generate all bistochastic maps. \square

We conjecture that $\text{Noisy}(n, n)$ is not convex (and thus a strict subset of the bistochastic maps), but we were not able to show this. This contrasts the simple behavior of the set of transitions that can be realized with an n -dimensional bath, characterized by majorization.

The following lemma has been used in the proof of Theorem 10.

Lemma 14. *Let B_1, \dots, B_n be matrices such that every B_i is a convex combination of other matrices $B_i^{(j)}$. Then*

$$B = B_1 \oplus \dots \oplus B_n$$

is a convex combination of the matrices

$$B_1^{(j_1)} \oplus \dots \oplus B_n^{(j_n)}.$$

Proof. The case $n = 1$ is trivial, and the cases $n \geq 3$ follow from repeated application of the statement for $n = 2$, so it is sufficient to prove the lemma for $n = 2$. Suppose that

$$B_1 = \sum_i \lambda_i B_1^{(i)}, \quad B_2 = \sum_j \mu_j B_2^{(j)}, \quad \lambda_i, \mu_j \geq 0, \quad \sum_i \lambda_i = \sum_j \mu_j = 1.$$

Then

$$B_1 \oplus B_2 = \left(\sum_{ij} \lambda_i \mu_j B_1^{(i)} \right) \oplus \left(\sum_{ij} \lambda_i \mu_j B_2^{(j)} \right) = \sum_{ij} \lambda_i \mu_j (B_1^{(i)} \oplus B_2^{(j)}).$$

Note that this statement can easily be illustrated by writing down the block matrices concretely. \square

The following lemma concerns the relative interior of a convex set [31]. It has been used in the proof of Corollary 12.

Lemma 15. *Let $C \subset \mathbb{R}^n$ be convex, and $X \subset C$ a dense subset of C . Then $\text{relint } C \subseteq \text{conv } X$.*

Proof. Denote the ε -ball around some point $c \in \mathbb{R}^n$ by $B_\varepsilon(c) = \{x \in \mathbb{R}^n \mid \|x - c\| \leq \varepsilon\}$. Let $c \in \text{relint } C$, then there is $\delta > 0$ small enough such that $B_\delta(c) \cap \text{aff}(C) \subset C$. Since $B_\delta(c) \cap \text{aff}(C)$ is an m -dimensional Euclidean unit ball, with $m = \dim C$, there is a regular m -simplex with vertices v_1, \dots, v_{m+1} on the surface of that ball, i.e. $\|v_i - c\| = \delta$, and center $c = \frac{1}{m+1} \sum_{i=1}^{m+1} v_i$. Since X is dense in C , we can find $v'_1, \dots, v'_{m+1} \in X$ that are arbitrarily close to the vertices v_1, \dots, v_{m+1} . If they are close enough, it is geometrically clear that their convex hull must contain c . \square

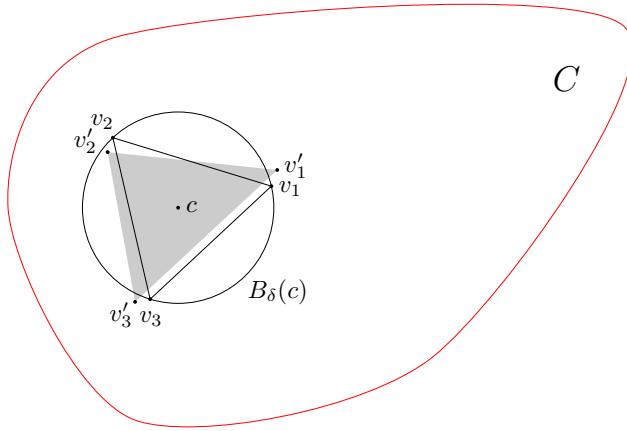


Figure 6: Illustration of the proof of Lemma 15.

As mentioned in the main text, Lemma 5 has been proven independently in [48], and it also follows from the results in [49]. Nevertheless, we give a self-contained proof for completeness.

Lemma 5. Let ρ be a quantum state on AB , and σ a quantum state on A . If $\dim A \leq \dim B$, there exists a unitary U on AB with

$$\mathrm{Tr}_B(U\rho U^\dagger) = \sigma$$

if and only if

$$\mathrm{Tr}_B \hat{\lambda}(\rho) \succ \sigma, \quad (19)$$

where $\lambda(\rho)$ is the vector of eigenvalues of ρ , ordered in non-increasing order, and $\hat{\lambda}(\rho)$ is the corresponding diagonal density matrix.

Proof. Suppose there is some unitary U such that $\mathrm{Tr}_B(U\rho U^\dagger) = \sigma$. Let us assume without loss of generality that ρ and σ are diagonal and their diagonal elements are non-increasingly ordered, so $\hat{\lambda}(\rho) = \rho$ and $\hat{\lambda}(\sigma) = \sigma$. With Lemma 1 we can write $U\hat{\lambda}(\rho)U^\dagger = (D\lambda(\rho))^\wedge + \Omega$ where D is a stochastic matrix with $D_{ij} = |U_{ij}|^2$ and Ω some matrix with only zeros on the diagonal. Thus, it holds $\lambda(\rho) \succ D\lambda(\rho)$, and $\mathrm{Tr}_B \Omega$ has also only zeros on its diagonal. Since σ is diagonal, the off-diagonal elements of $\mathrm{Tr}_B \Omega$ must vanish identically, and we get $\mathrm{Tr}_B \Omega = 0$ and $\lambda(\sigma) = \mathrm{Tr}_B(D\lambda(\rho))^\wedge$. If p, q are probability vectors on a joint system AB and p is non-increasingly ordered, then from $p \succ q$, it follows that $\Sigma_B p \succ \Sigma_B q$ where Σ_B is the marginal map with respect to the B system which satisfies $\mathrm{Tr}_B \hat{\lambda} = (\Sigma_B \lambda)^\wedge$. So we have

$$\mathrm{Tr}_B \hat{\lambda}(\rho) \succ \mathrm{Tr}_B(D\lambda(\rho))^\wedge = \hat{\lambda}(\sigma)$$

which proves the “only if” direction of the statement (and is independent of the dimensions of A and B).

To prove the converse direction, set $n = \dim A, m = \dim B$. Lemma 1 that we have just used is basically the one half of the Schur-Horn Theorem that states that the vector of diagonal elements of a Hermitian matrix is always majorized by the vector of its eigenvalues. For the second half of the proof one main ingredient will be the other half of the theorem which states: if λ is the vector of eigenvalues of a Hermitian matrix then for each vector p with $\lambda \succ p$ there exists a basis in which p is the vector of diagonal elements of this Hermitian matrix.

So assume that $\mathrm{Tr}_B \hat{\lambda}(\rho) \succ \sigma$ which means that $\Sigma_B \lambda(\rho) \succ \lambda(\sigma)$. Without loss of generality, we may assume that σ is diagonal (a local diagonalizing unitary can always be implemented via the global unitary U that we are constructing now), $\sigma = \hat{\lambda}(\sigma)$. By the Schur-Horn theorem we know there exist a unitary \bar{U} on the A system such that

$$\bar{U} \left(\mathrm{Tr}_B \hat{\lambda}(\rho) \right) \bar{U}^\dagger$$

has diagonal elements $\lambda(\sigma)$. We will use this to define a new unitary on the whole AB system such that $\mathrm{Tr}_B(U\rho U^\dagger) = \sigma$ is satisfied. So define

$$U = \sum_{i,j=1}^n u_{ij} |i\rangle\langle j| \otimes \pi^{j-i},$$

where u_{ij} are the components of \bar{U} , and π^{j-i} is a (later to be specified) permutation matrix π on the B system to the power $(j-i)$. Since $\pi^\dagger = \pi^{-1}$ and \bar{U} itself is unitary, this is a unitary matrix as well. Now the rest of the proof is merely showing by direct calculation that $\mathrm{Tr}_B(U\hat{\lambda}(\rho)U^\dagger) = \hat{\lambda}(\sigma)$.

For this we write the diagonal matrix $\hat{\lambda}(\rho)$ as

$$\hat{\lambda}(\rho) = \sum_{i=1}^n \sum_{\mu=1}^m \lambda_{i\mu} |i\rangle\langle i| \otimes |\mu\rangle\langle \mu| = \sum_{i=1}^n |i\rangle\langle i| \otimes \Lambda_i \quad \text{with} \quad \Lambda_i = \sum_{\mu=1}^m \lambda_{i\mu} |\mu\rangle\langle \mu|,$$

and with Lemma 1 we write again

$$U\hat{\lambda}(\rho)U^\dagger = (D\lambda(\rho))^\wedge + \Omega$$

where D is the unistochastic matrix $D = U \circ U^*$ and Ω some matrix with only zeroes on the diagonal. If we define $\bar{D} = \bar{U} \circ \bar{U}^*$ it is simple to show by direct calculation that

$$\Sigma_B(D\lambda(\rho)) = \bar{D}(\Sigma_B \lambda(\rho)) = \lambda(\sigma),$$

where the second equality sign holds by assumption. So we know that $\mathrm{Tr}_B(D\lambda(\rho))^\wedge = \hat{\lambda}(\sigma)$ and it remains to show that $\mathrm{Tr}_B \Omega = 0$.

By direct calculation, starting with the form of $\hat{\lambda}(\rho)$ and U as just given, we get the form

$$\Omega = \sum_{\substack{ij=1 \\ i \neq j}}^n |i\rangle\langle j| \otimes \sum_{k=1}^n u_{ik} u_{jk}^* \pi^{i-k} \Lambda_k \pi^{k-j} + \sum_{i=1}^n |i\rangle\langle i| \otimes \Omega^{(ii)},$$

where $\Omega^{(ii)} = \langle i|\Omega|i\rangle$ is a matrix on the B system with all diagonal elements equal to zero. Since $\text{Tr}_B \Omega^{(ii)} = 0$ the second term vanishes when calculating $\text{Tr}_B \Omega$. However we still have to look at the first term where we calculate

$$\text{tr}(\pi^{i-k} \Lambda_k \pi^{k-j}) = \text{tr}(\pi^{i-j} \Lambda_k)$$

and set $s = i - j$, with $s \in \{-n+1, \dots, -1, 1, \dots, n-1\}$ since $i \neq j$.

Now it is time to specify the permutation π . We choose π such that $\pi(i) = i+1$ for $i < m$ and $\pi(m) = 1$. This is a permutation with no fixed point in the set $\{1, \dots, m\}$ and thus $\text{tr}(\pi) = 0$ since the permutation matrix has only zeros on the diagonal. This is as well valid for all powers π^s of the matrix as long as $s \neq m$ and $s \neq 0$. This is true due to the restriction on s above and since we assumed $\dim A \leq \dim B$ i.e. $n \leq m$. Now, since π^s has no fixed point in $\{1, \dots, m\}$, it follows that

$$\text{Tr}(\pi^s \Lambda_k) = \sum_{\mu=1}^m \lambda_{i\mu} \langle \mu | \pi^s | \mu \rangle = 0,$$

finally showing that $\text{Tr}_B \Omega = 0$. □