Entanglement and thermodynamics in general probabilistic theories

Giulio Chiribella*

Department of Computer Science, University of Hong Kong, Pokfulam Road, Hong Kong

Carlo Maria Scandolo†

Department of Computer Science, University of Oxford, Oxford, UK

Entanglement is one of the most striking features of quantum mechanics, and yet it is not specifically quantum. More specific to quantum mechanics is the connection between entanglement and thermodynamics, which leads to an identification between entropies and measures of pure state entanglement. Here we search for the roots of this connection, investigating the relation between entanglement and thermodynamics in the framework of general probabilistic theories. We first address the question whether an entangled state can be transformed into another by means of local operations and classical communication. Under two operational requirements, we prove a general version of the Lo-Popescu theorem, which lies at the foundations of the theory of pure-state entanglement. We then consider a resource theory of purity where free operations are random reversible transformations, modelling the scenario where an agent has limited control over the dynamics of a closed system. Our key result is a duality between the resource theory of entanglement and the resource theory of purity, valid for every physical theory where all processes arise from pure states and reversible interactions at the fundamental level. As an application of the main result, we establish a one-to-one correspondence between entropies and measures of pure bipartite entanglement. The correspondence is then used to define entanglement measures in the general probabilistic framework. Finally, we show a duality between the task of information erasure and the task of entanglement generation, whereby the existence of entropy sinks (systems that can absorb arbitrary amounts of information) becomes equivalent to the existence of entanglement sources (correlated systems from which arbitrary amounts of entanglement can be extracted).

I. INTRODUCTION

The discovery of quantum entanglement [1, 2] introduced the revolutionary idea that a composite system can be in a pure state while its components are mixed. In Schrödinger's words: *"maximal knowledge of a total system does not necessarily imply maximal knowledge of all its parts"* [2]. This new possibility, in radical contrast with the paradigm of classical physics, is at the root of quantum non-locality [3–6] in all its counterintuitive manifestations [7–13]. With the advent of quantum information, it quickly became clear that entanglement was not only a source of foundational puzzles, but also a resource [14]. Harnessing this resource has been the key to the invention of groundbreaking protocols such as quantum teleportation [15], dense coding [16], and secure quantum key distribution [17, 18], whose implications deeply impacted physics and computer science [19, 20].

The key to understand entanglement as a resource is to consider distributed scenarios where spatially-separated parties perform local operations (LO) in their laboratories and exchange classical communication (CC) from one laboratory to another [21–23]. The protocols that can be implemented in this scenario, known as LOCC protocols, provide a means to characterize entangled states and to compare their degree of entanglement. Precisely, a state

is *i)* entangled if it cannot be generated by an LOCC protocol, and *ii)* more entangled than another if there exists an LOCC protocol that transforms the former into the latter.

Comparing the degree of entanglement of two quantum states is generally a hard problem [24–30]. Nevertheless, the solution is simple for pure bipartite states, where the majorization criterion [31] provides a necessary and sufficient condition for LOCC convertibility. The criterion identifies the degree of entanglement of a bipartite system with the degree of mixedness of its parts: the more entangled a pure bipartite state is, the more mixed its marginals are. Mixed states are compared here according to their spectra, with a state being more mixed than another if the spectrum of the latter majorizes the spectrum of the former [32–35].

The majorization criterion shows that for pure bipartite states the notion of entanglement as a resource beautifully matches Schrödinger's notion of entanglement as non-maximal knowledge about the parts of a pure composite system. Moreover, majorization establishes an intriguing duality between entanglement and thermodynamics [36–40], whereby the reduction of entanglement caused by LOCC protocols becomes dual to the increase of mixedness (and therefore entropy [41]) caused by thermodynamic transformations. This duality has far-reaching consequences, such as the existence of a unique measure of pure state entanglement in the asymptotic limit [36, 37, 42]. In addition, it has provided guidance for the development of entanglement theory beyond the case of pure bipartite states [39].

---

* giulio@cs.hku.hk
† carlomaria.scandolo@st-annes.ox.ac.uk

The duality between entanglement and thermodynamics is a profound and fundamental fact. As such, one might expect it to follow directly from basic principles. However, what these principle are is far from clear: up to now, the relation between entanglement and thermodynamics has been addressed in a way that depends heavily on the Hilbert space framework, using technical results that lack an operational interpretation (such as, e.g. the singular value decomposition). It is then natural to search for a derivation of the entanglement-thermodynamics duality that uses only high-level quantum features, such as the impossibility of instantaneous signalling or the no-cloning theorem. In the same spirit, one can ask whether the duality holds for physical theories other than quantum mechanics, adopting the broad framework of *general probabilistic theories* [43–52]. In the landscape of general probabilistic theories, entanglement is a generic feature [44, 53], which provides powerful advantages for a variety of information-theoretic tasks [54–59]. But what about its relation with thermodynamics? Is it also generic, or rather constitutes a specific feature of quantum theory?

In this paper we explore the relation between entanglement and thermodynamics in an operational, theory-independent way. Our work is part of a larger project that aims at establishing a common axiomatic foundation to quantum information theory and quantum thermodynamics. Within this broad scope, we start our investigation from the resource theory of entanglement, asking which conversions are possible under LOCC protocols. Our first result is a generalization of the Lo-Popescu theorem [23]: we show that under suitable assumptions every LOCC protocol acting on a pure bipartite state can be simulated by a protocol using only one round of classical communication. Our assumptions are satisfied by quantum theory on both real and complex Hilbert spaces, and also by all bipartite extreme no-signalling boxes studied in the literature [60–63].

In order to establish the connection with thermodynamics, we then move our attention to mixed states. We consider the scenario where an agent has limited control over the dynamics of a closed system, thus causing it to undergo a random mixture of reversible transformations and degrading it to a more disordered state. This notion of degradation coincides with the notion of "adding noise" put forward by Müller and Masanes for the problem of encoding spatial directions into physical systems [64], and represents a natural generalization of the notion of majorization [65]. Provided that that every pure state can be reached from any other pure state through some reversible dynamics, we show that the relevant resource in this scenario can be identified with the purity of the state. This observation leads to an operational theory of purity, which in the quantum case turns out to be equivalent to the theory of purity defined by Horodecki and Oppenheim [66].

Once the resource theories of entanglement and purity are put into place, we set out to establish a dual-ity between them. To this purpose, we consider physical theories that admit a fundamental level of description where all states are pure and all interactions are reversible. Such theories are identified by the Purification Principle [47], which expresses a strengthened version of the conservation of information [67, 68]. The possibility of a pure and reversilble description is particularly appealing for the foundations of thermodynamics, as it reconciles the mixedness of thermodynamic ensembles with the pure and reversible picture provided by fundamental physics. In the quantum case, Purification is the starting point for all recent proposals to derive thermodynamic ensembles from the typicality of pure entangled states [69–75], an idea that has been recently explored also in the broader framework of general probabilistic theories [76, 77]. Building on the Purification Principle, we establish the desired duality between entanglement and thermodynamics, showing that the degree of entanglement of a pure bipartite system coincides with the degree of mixedness of its parts. As a consequence, every measure of single-system mixedness becomes equivalent to a measure of pure bipartite entanglement. Exploiting this result, we define a class of measures of entanglement, which can be extended to from pure to mixed state via the the convex roof construction [22, 78, 79], exactly in the same way as in the quantum case.

Finally, we apply the duality to the task of information erasure [80], namely the task of converting a mixed state into a fixed pure state via a set of allowed operations (in our case, the set of random reversible operations). As a result, erasing information becomes equivalent to generating entanglement. Quite surprisingly, we find out that the impossibility of erasing information with the assistance of a catalyst implies the existence of a special type of purification, where the purifying system is a twin of the purified system. This observation completes the physical picture of the entanglement-thermodynamics duality, which appears to be a consequence of the possibility to describe every physical process in terms of pure states, reversible interactions, and pure measurements—in particular, modelling physical systems in mixed states through the introduction of a mirror image that completes the description.

The paper is organized as follows. In section II we introduce the framework. The resource theory of entanglement is discussed in section III. In section IV we prove an operational version of Lo-Popescu theorem, which provides the starting point for the entanglement-thermodynamics duality. In section V we formulate an operational resource theory of dynamical control, which gives rise to a resource theory of purity under the condition that all pure state are equivalent under the allowed reversible dynamics. In section VI we prove the duality between entanglement and thermodynamics, focussing our attention to theories that admit a fundamental level where all processes are pure and reversible. The consequences of the duality are examined in section VII: specifically, we discuss the equivalence between measures

of mixedness and measures of entanglement for pure bipartite states, and we establish the relation between information erasure and entanglement generation. In section VIII we show that the requirement that information cannot be erased for free leads to the requirement of Symmetric Purification. Finally, section IX draws the conclusions and highlights the implications of our results.

II. FRAMEWORK

In this paper we adopt the framework of *operational-probabilistic theories (OPTs)* [47, 48, 52], which combines the toolbox of probability theory with the graphical language of symmetric monoidal categories [81–84]. Here we give a quick recap, referring the reader to Refs. [47, 48, 52] and to Hardy's works [50, 51] for a more extended presentation.

The OPT framework describes circuits that can be built up by combining physical processes in sequence and in parallel, as in the following example
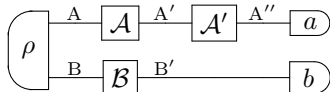
Here, A, A′, A″, B, B′ label physical systems, ρ is a bipartite state, $\mathcal{A}$, $\mathcal{A}'$ and $\mathcal{B}$ are transformations, a and b are effects. The two transformations $\mathcal{A}$ and $\mathcal{A}'$ are composed in sequence, while the transformations $\mathcal{A}$ and $\mathcal{B}$ and the effects a and b are composed in parallel. The circuit has no external wires—circuits of this form are associated with probabilities. Two transformations that give the same probabilities in all circuits are identified. The short-hand notation $(a|\rho)$ is used to indicate the probability that the effect a takes place on the state ρ, diagrammatically represented as

$$(a|\rho) := \boxed{\rho} \!-\!\!^{\text{A}}\!\!-\! \boxed{a} \, .$$

The set of all possible physical systems, denoted by Sys, is closed under composition: given two systems A and B one can form the composite system A⊗B. We denote the trivial system as I, which represents "nothing" (or, more precisely, nothing that the theory cares to describe). The trivial system satisfies the obvious conditions $A \otimes I = I \otimes A = A$, $\forall A \in$ Sys. For generic systems A and B, we denote as

- St (A) the set of states of system A

- Transf (A, B) the set of transformations from system A to system B

- Eff (A) the set of effects on system A

The sets of states, transformations, and effects span vector spaces over the real numbers, denoted by $\mathsf{St}_{\mathbb{R}}$ (A), $\mathsf{Transf}_{\mathbb{R}}$ (A, B), and $\mathsf{Eff}_{\mathbb{R}}$ (A), respectively. We denote by D_A the dimension of the vector space $\mathsf{St}_{\mathbb{R}}$ (A) and say

that system A is *finite* iff $D_A < +\infty$. Transformations and effects act linearly on the vector space of states. For every system A, we assume the existence of an identity transformation $\mathcal{I}_A$, which does nothing on the states of the system.

A *test* is a collection of transformations that can occur as alternatives in an experiment. Specifically, a test of type A to B is a collection of transformations $\{\mathcal{C}_i\}_{i \in \mathsf{X}}$ with input A and output B. A transformation is called *deterministic* if it belongs to a test with a single outcome. We will often refer to deterministic transformations as *channels*, following the standard terminology of quantum information. A channel $\mathcal{U}$ from A to B is called *reversible* if there exists a channel $\mathcal{U}^{-1}$ from B to A such that $\mathcal{U}^{-1}\mathcal{U} = \mathcal{I}_A$ and $\mathcal{U}\mathcal{U}^{-1} = \mathcal{I}_B$. We denote by RevTransf (A, B) the set of reversible transformations from A to B. If there exists a reversible channel transforming A into B we say that A and B are *operationally equivalent*, denoted by A $\simeq$ B.

The composition of systems is required to be *symmetric* [81–84], meaning that A $\otimes$ B $\simeq$ B $\otimes$ A. The reversible channel that implements the equivalence is the *swap channel*, SWAP, and satisfies the condition

(1)

for every pair of transformations $\mathcal{A}$ and $\mathcal{B}$ and for generic systems A, A′, B, B′, as well as the conditions

the wires in the r.h.s. representing identity transformations, and

In this paper we restrict our attention to *causal theories* [47], namely theories where the choice of future measurement settings does not influence the outcome probability of present experiments. Mathematically, causality is equivalent to the fact that for every system A there is only one deterministic effect, which we denote here by Tr_A, in analogy with the trace in quantum mechanics. The uniqueness of the deterministic effect provides a canonical way to define marginal states:

Definition 1 *The* marginal state *of a bipartite state ρ_{AB} on system A is the state $\rho_A := \mathrm{Tr}_B \rho_{AB}$ obtained by applying the deterministic effect on B.*

Moreover, one can define the *norm* of a state ρ as

$$\|\rho\| := \mathrm{Tr}\,\rho.$$

The set of normalized states of A will be denoted by

$$\mathsf{St}_1(A) := \{\rho \in \mathsf{St}(A) \mid \|\rho\| = 1\} \, .$$

In a causal theory, every state is proportional to a normalized state [47]. In quantum mechanics, $\mathsf{St}_1(A)$ is the set of normalized density matrices of system A, while $\mathsf{St}(A)$ is the set of all sub-normalized density matrices.

In a causal theory channels admit a simple characterization, which will be useful later in the paper:

Proposition 1 *Let* $\mathcal{C} \in \mathsf{Transf}(A, B)$. *$\mathcal{C}$ is a channel if and only if* $\mathrm{Tr}_B \mathcal{C} = \mathrm{Tr}_A$.

The proof can be found in lemma 5 of Ref. [47].

A. Pure states and transformations

In every probabilistic theory one can define pure states, and, more generally, pure transformations. Both concepts are based on the notion of *coarse-graining*, i.e. the operation of joining two or more outcomes of a test. More precisely, a test $\{\mathcal{C}_i\}_{i \in \mathsf{X}}$ is a *coarse-graining* of the test $\{\mathcal{D}_j\}_{j \in \mathsf{Y}}$ if there is a partition $\{\mathsf{Y}_i\}_{i \in \mathsf{X}}$ of Y such that $\mathcal{C}_i = \sum_{j \in \mathsf{Y}_i} \mathcal{D}_j$ for every $i \in \mathsf{X}$. In this case, we say that $\{\mathcal{D}_j\}_{j \in \mathsf{Y}}$ is a *refinement* of $\{\mathcal{C}_i\}_{i \in \mathsf{X}}$. The refinement of a given transformation is defined via the refinement of a test: if $\{\mathcal{D}_j\}_{j \in \mathsf{Y}}$ is a refinement of $\{\mathcal{C}_i\}_{i \in \mathsf{X}}$, then the transformations $\{\mathcal{D}_j\}_{j \in \mathsf{Y}_i}$ are a refinement of the transformation $\mathcal{C}_i$.

A transformation is called *pure* if it has only trivial refinements:

Definition 2 *The transformation* $\mathcal{C} \in \mathsf{Transf}(A, B)$ *is* pure *if for every refinement* $\{\mathcal{D}_j\}$ *one has* $\mathcal{D}_j = p_j \mathcal{C}$, *where* $\{p_j\}$ *is a probability distribution.*

Pure transformations are those for which the experimenter has maximal information about the evolution of the system. We assume as part of the framework that tests satisfy a *pure decomposition property*:

Definition 3 *A test satisfies the* pure decomposition property *if it admits a refinement consisting only of pure transformations.*

Later in the paper, we will assume one axiom—Purification—that implies the validity of the pure decomposition property for every possible test [47].

The set of pure transformations from A to B will be denoted as $\mathsf{PurTransf}(A, B)$. In the special case of states (transformations with no input), the above definition coincides with the usual definition of pure state. We denote the set of pure states of system A as $\mathsf{PurSt}(A)$. As usual, non-pure states will be called *mixed*.

Pure states will play a key role in this paper. An elementary property of pure states is that they are preserved by reversible transformations.

Proposition 2 *Let* $\mathcal{U} \in \mathsf{Transf}(A, B)$ *be a reversible channel. Then a state* $\psi \in \mathsf{St}(A)$ *is pure if and only if the state* $\mathcal{U}\psi \in \mathsf{St}(B)$ *is pure.*

The proof is standard and is reported in Appendix A for convenience of the reader.

III. THE RESOURCE THEORY OF ENTANGLEMENT

The resource theory of quantum entanglement [14] is based on the notion of LOCC protocols, that is, protocols in which distant parties are allowed to communicate classically to one another and to perform local operations in their laboratories [21, 22]. Being operational, the notion of LOCC protocol can be directly exported to arbitrary theories.

In this paper we consider protocols involving only two parties, Alice and Bob. A generic LOCC protocol consists of a sequence of tests, performed by Alice and Bob, with the property that the choice of the test at a given step can depend on all the outcomes produced at the previous steps. For example, consider a two-way protocol where

1. Alice performs a test $\{\mathcal{A}_{i_1}\}$ and communicates the outcome to Bob

2. Bob performs a test $\left\{\mathcal{B}_{i_2}^{(i_1)}\right\}$ and communicates the outcome to Alice

3. Alice performs a test $\left\{\mathcal{A}_{i_3}^{(i_1, i_2)}\right\}$.

An instance of the protocol is identified by the sequence of outcomes (i_1, i_2, i_3) and can be represented by a circuit of the form
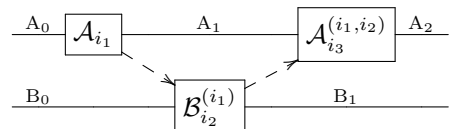
where the dashed arrows represent classical communication. By coarse-graining over all possible outcomes, one obtains a channel, given by

$$\mathcal{L} = \sum_{i_1, i_2, i_3} \left[\mathcal{A}_{i_3}^{(i_1, i_2)} \mathcal{A}_{i_1} \otimes \mathcal{B}_{i_2}^{(i_1)} \right] \, .$$

Entangled states are those states that cannot be generated using an LOCC protocol. Equivalently, they can be characterized as the states that are not *separable*, i.e. not of the form

$$\rho = \sum_i p_i \, \alpha^{(i)} \otimes \beta^{(i)} \, ,$$

where $\{p_i\}$ is a probability distribution allowed by the theory, $\alpha^{(i)}$ is a state of A, and $\beta^{(i)}$ is a state of B.

Like in quantum theory, LOCC protocols can be used to compare entangled states.

Definition 4 *Given two states $\rho \in \mathsf{St}\,(\mathrm{A} \otimes \mathrm{B})$ and $\rho' \in \mathsf{St}\,(\mathrm{A}' \otimes \mathrm{B}')$, we say that ρ is* more entangled *than ρ', denoted by $\rho \succeq_{\mathrm{ent}} \rho'$, if there exists an LOCC protocol that transforms ρ into ρ', i.e. if $\rho' = \mathcal{L}\rho$ for some LOCC channel $\mathcal{L}$.*

Mathematically, the relation $\succeq_{\mathrm{ent}}$ is a preorder, i.e. it is reflexive and transitive. Moreover, it is stable under tensor products, namely $\rho \otimes \sigma \succeq_{\mathrm{ent}} \rho' \otimes \sigma'$ whenever $\rho \succeq_{\mathrm{ent}} \rho'$ and $\sigma \succeq_{\mathrm{ent}} \sigma'$. In other words, the relation $\succeq_{\mathrm{ent}}$ turns the set of all bipartite states into a preordered monoid, the typical mathematical structure arising in all resource theories [85]. The resource theory of entanglement here fits completely into the framework of Ref. [85], with the LOCC channels as free operations. The states which can be prepared by LOCC (i.e. the separable states) are free by definition, and all the other states represent resources. If $\rho \succeq_{\mathrm{ent}} \rho'$ and $\rho' \succeq_{\mathrm{ent}} \rho$, then we say that ρ and ρ' are *equally entangled*, denoted by $\rho \simeq_{\mathrm{ent}} \rho'$. Note that $\rho \simeq_{\mathrm{ent}} \rho'$ does not imply that ρ and ρ' are equal: for example, every two separable states are equally (un)entangled.

IV. AN OPERATIONAL LO-POPESCU THEOREM

Given two bipartite states, it is natural to ask whether one is more entangled than the other. A priori, answering the question requires one to check all possible LOCC protocols. However, the situation is much simpler when the initial state is pure: here we prove that in this case every LOCC protocol can be replaced without loss of generality by a protocol involving only one round of classical communication—i.e. a *one-way LOCC protocol*. Our argument is based on two basic operational requirements and provides a generalized version of the Lo-Popescu theorem [23], the key result at the foundation of the quantum theory of pure-state entanglement.

A. Two operational requirements

Our derivation of the operational Lo-Popescu theorem is based on two requirements, the first being

Axiom 1 (Purity Preservation [45, 48, 68]) *The composition of two pure transformations yields a pure transformation, namely*
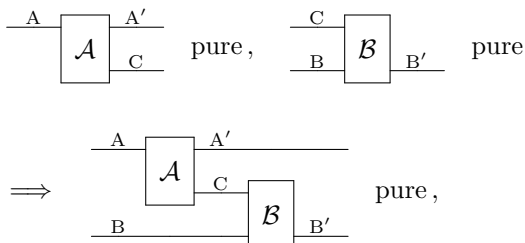
for every choice of systems $\mathrm{A}, \mathrm{A}', \mathrm{B}, \mathrm{B}', \mathrm{C}$.

As a special case, Purity Preservation implies that the product of two pure states is a pure state. This conclusion could also be obtained from the Local Tomography axiom [43–45, 47]. Nevertheless, counterexamples exist of theories that satisfy Purity Preservation and violate Local Tomography. An example is quantum theory on real vector spaces [86–88]. In general, we regard Purity Preservation as more fundamental than Local Tomography. Considering the theory as an algorithm to make deductions about physical processes, Purity Preservation ensures that, when presented with maximal information about two processes, the algorithm outputs maximal information about their composition [68].

Our second requirement imposes a symmetry of pure bipartite states:

Axiom 2 (Local Exchangeability) *For every pure bipartite state $\Psi \in \mathsf{PurSt}\,(\mathrm{A} \otimes \mathrm{B})$, there exist two channels $\mathcal{C} \in \mathsf{Transf}\,(\mathrm{A}, \mathrm{B})$ and $\mathcal{D} \in \mathsf{Transf}\,(\mathrm{B}, \mathrm{A})$ such that*

where SWAP *is the swap operation [cf. Eq. (1)].*

Note that, in general, the two channels depend on the specific pure state Ψ.

Local Exchangeability is trivially satisfied by classical probability theory, where all pure states are of the product form. Less trivially, it is satisfied by quantum theory, both on complex and on real Hilbert spaces. This fact is illustrated in the following

Example 1 *Suppose that A and B are quantum systems, and let $\mathcal{H}_{\mathrm{A}}$ and $\mathcal{H}_{\mathrm{B}}$ be the corresponding Hilbert spaces. By the Schmidt decomposition, every pure state in the tensor product Hilbert space can be written as*

$$|\Psi\rangle = \sum_{i=1}^{r} \sqrt{p_i}\, |\alpha_i\rangle \otimes |\beta_i\rangle \,,$$

where $\{|\alpha_i\rangle\}_{i=1}^{r} \subset \mathcal{H}_{\mathrm{A}}$ and $\{|\beta_i\rangle\}_{i=1}^{r} \subset \mathcal{H}_{\mathrm{B}}$ are orthonormal vectors. The Schmidt decomposition implies the relation

$$\mathrm{SWAP}|\Psi\rangle = (C \otimes D)\,|\Psi\rangle \,,$$

where C and D are the partial isometries $C := \sum_{i=1}^{r} |\beta_i\rangle\langle\alpha_i|$ and $D := \sum_{i=1}^{r} |\alpha_i\rangle\langle\beta_i|$. From the partial isometries C and D it is immediate to construct the desired channels $\mathcal{C}$ and $\mathcal{D}$, which can be defined as

$$\mathcal{C}\,(\rho) := C\rho C^{\dagger} + \sqrt{I_{\mathrm{A}} - C^{\dagger}C}\,\rho\,\sqrt{I_{\mathrm{A}} - C^{\dagger}C}$$
$$\mathcal{D}\,(\sigma) := D\sigma D^{\dagger} + \sqrt{I_{\mathrm{B}} - D^{\dagger}D}\,\sigma\,\sqrt{I_{\mathrm{B}} - D^{\dagger}D} \,,$$

where ρ and σ are generic input states of systems A and B, respectively. With this definition, one has

$$(\mathcal{C} \otimes \mathcal{D})\,(|\Psi\rangle\langle\Psi|) = \mathrm{SWAP}\,|\Psi\rangle\langle\Psi|\,\mathrm{SWAP} \,,$$

which is the Hilbert space version of the local exchangeability condition of Eq. (2).

Local Exchangeability is also satisfied by all the extreme bipartite non-local correlations characterized in the literature [60–63]:

Example 2 *Consider a scenario where two space-like separated parties, Alice and Bob, perform measurements on a pair of systems, A and B, respectively. Let x (y) be the index labelling Alice's (Bob's) measurement setting and let a (b) the index labelling the outcome of the measurement done by Alice (Bob). In the theory known as box world [44, 89] all no-signalling probability distributions $p_{ab|xy}$ are physically realizable and represent states of the composite system A ⊗ B. Such probability distributions form a convex polytope [62], whose extreme points are the pure states of the theory.*

For $x, y, a, b \in \{0, 1\}$ the systems A and B are operationally equivalent. We will denote by $\mathcal{I}$ the reversible transformation that converts A into B. The extreme non-local correlations have been characterized in [60] and are known to be equal to the standard PR-box correlation [61]

$$p_{ab|xy} = \begin{cases} \frac{1}{2} & a + b \equiv xy \mod 2 \\ 0 & otherwise \end{cases} \tag{3}$$

up to exchange of 0 with 1 in the local settings of Alice and Bob and in the outcomes of their measurements. In the circuit picture, these operations are described by local reversible transformations: denoting by Φ the standard PR-box state, one has that every other pure entangled state $\Psi \in \mathsf{PurSt}(A \otimes B)$ is of the form

where $\mathcal{U}$ and $\mathcal{V}$ are reversible transformations.

To see that Local Exchangeability holds, note that swapping systems A and B is equivalent to exchanging x with y and a with b. Now, the standard PR-box correlation of Eq. (3) is invariant under exchange $x \leftrightarrow y$, $a \leftrightarrow b$, meaning that one has

Then it is clear that every pure state of A ⊗ B can be swapped by local operations: indeed, one has

$$\tag{4}$$

where $\mathcal{C} := \mathcal{V}\mathcal{I}\mathcal{U}^{-1}$ and $\mathcal{D} := \mathcal{U}\mathcal{I}^{-1}\mathcal{V}^{-1}$. *This proves the Local Exchangeability property for all pure bipartite states in the 2-setting/2-outcome scenario.*

The situation is analogous in the case of 2 settings and arbitrary number of outcomes. Let us denote the setting by $x, y \in \{0, 1\}$ and assume that a can take d_{A} values, whereas b can take d_{B} values. Then, all extreme non-local correlations are characterized in Ref. [62]. Up to local reversible transformations, they are labelled by a parameter $k \in \{2, \ldots, \min\{d_{\mathrm{A}}, d_{\mathrm{B}}\}\}$ and they are such that

$$p_{ab|xy} = \begin{cases} \frac{1}{k} & b - a \equiv xy \mod k \\ 0 & otherwise \end{cases}. \tag{5}$$

Thanks to the local equivalence, it is enough to prove the validity of Local Exchangeability for correlations in the standard form of Eq. (5). We distinguish between the two cases $xy = 0$ and $xy = 1$. For $xy = 0$, swapping x with y and a with b has no effect on $p_{ab|xy}$. For $xy = 1$, by swapping x with y and a with b, one obtains the probability distribution

$$p'_{ab|xy} = \begin{cases} \frac{1}{k} & a - b \equiv 1 \mod k \\ 0 & otherwise \end{cases}$$

This probability distribution can be obtained from the original one by relabelling the outputs as $a' := k - a$ and $b' := k - b$. Such a relabelling corresponds to local reversible operations on A and B. In other words, Local Exchangeability holds.

Finally, the last category of extreme non-local correlations characterized in the literature corresponds to the case of arbitrary number of settings and to 2-outcome measurements. In this case, the extreme correlations are characterized explicitly in Ref. [63]. Up to local reversible transformations, the pure states are invariant under swap. Hence, the same argument used in Eq. (4) shows that Local Exchangeability holds.

B. Inverting the direction of classical communication

Purity Preservation and Local Exchangeability have an important consequence. For one-way protocols acting on a pure input state, the direction of classical communication is irrelevant: every one-way LOCC protocol with communication from Alice to Bob can be replaced by a one-way LOCC protocol with communication from Bob to Alice, as shown by the following.

Lemma 1 (Inverting CC) *Let Ψ be a pure state of A ⊗ B and let ρ' be a (possibly mixed) state of A' ⊗ B'. Under the validity of axioms 1 and 2, the following are equivalent:*

1. Ψ can be transformed into ρ' by a one-way LOCC protocol with communication from Alice to Bob

2. Ψ *can be transformed into* ρ' *by a one-way LOCC protocol with communication from Bob to Alice.*

Proof. Suppose that Ψ can be transformed into ρ' by a one-way LOCC protocol with communication from Alice to Bob, namely

$$\left(\rho'\begin{array}{l}\text{A}'\\\text{B}'\end{array}\right.=\sum_{i\in\mathsf{X}}\left(\Psi\begin{array}{l}\text{A}\\\text{B}\end{array}\boxed{\mathcal{A}_i}\,\text{A}'\searrow\atop\boxed{\mathcal{B}^{(i)}}\,\text{B}'\right.\qquad(6)$$

where $\{\mathcal{A}_i\}_{i\in\mathsf{X}}$ is a test, and, for every outcome $i\in\mathsf{X}$, $\mathcal{B}^{(i)}$ is a channel. Note that one can assume without loss of generality that all transformations $\{\mathcal{A}_i\}_{i\in\mathsf{X}}$ are pure: if the transformations are not pure, we can refine them by the pure decomposition property (cf. definition 2) and apply the argument to the refined test consisting of pure transformations.

For every fixed $i\in\mathsf{X}$, one has

$$\qquad(7)$$

By Local Exchangeability, the first swap can be realized by two local channels $\mathcal{C}:\text{A}\to\text{B}$ and $\mathcal{D}:\text{B}\to\text{A}$. Moreover, since $\mathcal{A}_i$ is pure, Purity Preservation implies that the (unnormalized) state $(\mathcal{A}_i\otimes\mathcal{I}_{\text{B}})\,\Psi$ is pure. Hence, also the second swap in Eq. (7) can be realized by two local channels $\mathcal{C}^{(i)}:\text{A}'\to\text{B}$ and $\mathcal{D}^{(i)}:\text{B}\to\text{A}'$. Substituting into Eq. (7) one obtains

and, therefore,

$$=:\qquad(8)$$

having defined $\widetilde{\mathcal{A}}^{(i)}:=\mathcal{D}^{(i)}\,\mathcal{C}$ and $\widetilde{\mathcal{B}}^{(i)}=\mathcal{B}^{(i)}\,\mathcal{C}^{(i)}\,\mathcal{A}_i\,\mathcal{D}$. By construction $\left\{\widetilde{\mathcal{B}}_i\right\}_{i\in\mathsf{X}}$ is a test, because it can be realized by performing the test $\{\mathcal{A}_i\}_{i\in\mathsf{X}}$ after the channel $\mathcal{D}$

and subsequently applying the channel $\mathcal{B}^{(i)}\,\mathcal{C}^{(i)}$, depending on the outcome (the ability to perform conditional operations is guaranteed by causality [47]). On the other hand, $\widetilde{\mathcal{A}}^{(i)}$ is a channel for every $i\in\mathsf{X}$. Hence, we have constructed a one-way LOCC protocol with communication from Bob to Alice. Combining Eqs. (6) and (8) we obtain

meaning that Ψ can be transformed into ρ' by a one-way LOCC protocol with communication from Bob to Alice. Clearly, the same argument can be applied to prove the converse direction. ■

Note that the target state ρ' need not be pure: the fact that the direction of classical communication can be exchanged relies only on the purity of the input state Ψ.

C. Reduction to one-way protocols

We are now ready to derive the operational version of the Lo-Popescu theorem. Our result shows that the action of an arbitrary LOCC protocol on a pure state can be simulated by a one-way LOCC protocol:

Theorem 1 (Operational Lo-Popescu theorem)
Let Ψ *be a pure state of* $\text{A}\otimes\text{B}$ *and* ρ' *be a (possibly mixed) state of* $\text{A}'\otimes\text{B}'$. *Under the validity of axioms 1 and 2, the following are equivalent*

1. Ψ *can be transformed into* ρ' *by an LOCC protocol*

2. Ψ *can be transformed into* ρ' *by a one-way LOCC protocol.*

Proof. The non-trivial implication is $1\implies2$. Suppose that Ψ can be transformed into ρ' by an LOCC protocol with N rounds of classical communication. Without loss of generality, we assume that Alice starts the protocol and that all transformations occurring in the first $N-1$ rounds are pure.

Let $s=(i_1,i_2,\ldots,i_{N-1})$ be the sequence of all classical outcomes obtained by Alice and Bob up to step $N-1$, p_s be the probability of the sequence s, and Ψ_s be the pure state after step $N-1$ conditional on the occurrence of s. For concreteness, suppose that the outcome i_{N-1} has been generated on Alice's side. Then, the rest of the protocol consists in a test $\left\{\mathcal{B}_{i_N}^{(s)}\right\}$, performed on Bob's side, followed by a channel $\mathcal{A}^{(s,i_N)}$ performed on Alice's

side. By hypothesis, one has

$$
\rho' \!\begin{array}{l} \text{A}' \\ \text{B}' \end{array} = \sum_s p_s \;\; \times
$$

$$
\times \sum_{i_N} \left(\Psi_s \;\; \begin{array}{c} A_{N-1} \\ B_{N-1} \end{array} \; \mathcal{A}^{(s,i_N)} \; A' \;\; \mathcal{B}^{(s)}_{i_N} \; B' \right) .
$$

Now, using lemma 1 one can invert the direction of the classical communication in the last round, obtaining

$$
\rho' \!\begin{array}{l} \text{A}' \\ \text{B}' \end{array} = \sum_s p_s \;\; \times
$$

$$
\times \sum_{i_N} \left(\Psi_s \;\; \begin{array}{c} A_{N-1} \\ B_{N-1} \end{array} \; \widetilde{\mathcal{A}}^{(s)}_{i_N} \; A' \;\; \widetilde{\mathcal{B}}^{(s,i_N)} \; B' \right)
$$

for a suitable test $\left\{ \widetilde{\mathcal{A}}^{(s)}_{i_N} \right\}$ and suitable channels $\widetilde{\mathcal{B}}^{(s,i_N)}$. Now, since both the $(N-1)$-th and the N-th tests are performed by Alice, they can be merged into a single test, thus reducing the original LOCC protocol to an LOCC protocol with $N-1$ rounds. Iterating this argument for $N-1$ times one finally obtains a one-way protocol. $\blacksquare$

In quantum theory, the Lo-Popescu theorem provides the foundation for the resource theory of pure-state entanglement. Having the operational version of this result will be crucial for our study of the relation between entanglement and thermodynamics. Before entering into that, however, we need to put into place a suitable resource theory of purity, which will provide the basis for our thermodynamic considerations.

V. THE RESOURCE THEORY OF PURITY

A. A resource theory of dynamical control

Consider the scenario where a closed system A undergoes a reversible dynamics governed by some parameters under the experimenter's control. For example, system A could be a charged particle moving in an electric field, whose intensity and direction can be tuned in order to obtain a desired trajectory. In general, the experimenter may not have full control and the actual values of the parameters may fluctuate randomly. As a result, the evolution of the system will be described by a *Random Reversible (RaRe) channel*, that is a channel $\mathcal{R}$ of the form

$$
\text{A} \; \boxed{\mathcal{R}} \; \text{A} = \sum_{i \in \mathsf{X}} p_i \; \text{A} \; \boxed{\mathcal{U}^{(i)}} \; \text{A} \;,
$$

where $\left\{ \mathcal{U}^{(i)} \mid i \in \mathsf{X} \right\}$ is a set of reversible transformations and $\{p_i\}_{i \in \mathsf{X}}$ is their probability distribution. Assuming

that the system remains closed during the whole evolution, RaRe channels are the most general transformations the experimenter can implement.

An important question in all problems of control is whether a given input state can be driven to a target state using the allowed dynamics. With respect to this task, an input state is more valuable than another if the set of target states that can be reached from the former contains the set of target states that can be reached from the latter. In our model, this idea leads to the following definition.

Definition 5 (More controllable states) *Given two states ρ and ρ' of system A, we say that ρ is more controllable than ρ', denoted by $\rho \succeq \rho'$, if ρ' can be obtained from ρ via a RaRe channel.*

This definition appeared independently in an earlier work by Müller and Masanes [64], where the authors explored the use of two-level systems as indicators of spatial directions (cf. definition 8 in the appendix). In this paper we propose to consider it as the starting point for an axiomatic theory of thermodynamics.

Definition 5 fits into the general framework of resource theories [85], with RaRe channels playing the role of free operations. Note that at this level of generality there are no free states: since the experimenter can only control the evolution, every state is regarded as a resource. Physically, this is in agreement with the fact that the input state in a control problem is not chosen by the experimenter—for example, it can be a thermal state or the ground state of an unperturbed Hamiltonian.

As it is always the case in resource theories, the relation $\succeq$ is reflexive and transitive, i.e. it is a preorder. Moreover, since the tensor product of two RaRe channels is a RaRe channel, the relation $\succeq$ is stable under tensor products, namely $\rho \otimes \sigma \succeq \rho' \otimes \sigma'$ whenever $\rho \succeq \rho'$ and $\sigma \succeq \sigma'$. For finite systems (i.e. systems with finite-dimensional state space), Müller and Masanes [64] showed the additional property

$$
\rho \succeq \sigma, \quad \sigma \succeq \rho \quad \Longrightarrow \quad \rho = \mathcal{U}\sigma, \tag{9}
$$

for some reversible transformation $\mathcal{U}$. In other words, two states that are equally controllable can only differ by a reversible transformation.

B. From dynamical control to purity

There is a close relation between the controllability of a state and its purity. For example, a state that is more controllable than a pure state must also be pure.

Proposition 3 *If $\psi \in \mathsf{St}(A)$ is a pure state and $\rho \in \mathsf{St}(A)$ is more controllable than ψ, then ρ must be pure. Specifically, $\rho = \mathcal{U}\psi$ for some reversible channel $\mathcal{U}$.*

Proof. Since ψ is pure, the condition $\sum_i p_i \mathcal{U}^{(i)}\rho = \psi$ implies that $\mathcal{U}^{(i)}\rho = \psi$ for every i, meaning that

$\rho = \mathcal{V}^{(i)}\psi$, where $\mathcal{V}^{(i)}$ is the inverse of $\mathcal{U}^{(i)}$. Proposition 2 then guarantees that ρ is pure. ∎

In other words, pure states can be reached only from pure states. A natural question is whether every state can be reached from some pure state. The answer is positive in quantum theory and in a large class of theories. Nevertheless, counterexamples exist that prevent an easy identification of the resource theory of dynamical control with a "resource theory of purity". This fact is illustrated in the following

Example 3 *Consider a system with the state space depicted in Fig. 1(a). In this case, there are only two reversible transformations, namely the identity and the reflection around the vertical symmetry axis. As a consequence, there is no way to obtain the mixed states on the two vertical sides by applying a RaRe channel to a pure state. These states represent a valuable resource, even though they are not pure. Since some mixed states are a resource, the resource theory of dynamical control cannot be viewed as a resource theory of purity.*

As a second example, consider instead a system whose state space is a half-disk, like in Fig. 1(b). Also in this case there are only two reversible transformations (the identity and the reflection around the vertical axis). However, now every mixed state can be generated from some pure state via a RaRe channel. The state space can be foliated into horizontal segments generated by pure states under the action of RaRe channels. As a result, the pure states are the most useful resources and one can interpret the relation $\succeq$ as a way to compare the degree of purity of different states. Nevertheless, pure states on different segments are inequivalent resources. In this case there are different, inequivalent classes of pure states: purity is not the only relevant resource into play.

Finally, consider a system with a square state space, like in Fig. 1(c) and suppose that all the symmetry transformations in the dihedral group D_4 are allowed reversible transformations. In this case, all the pure states are equivalent under reversible transformations and every mixed state can be obtained by applying a RaRe channel to a fixed pure state. Here, the resource theory of dynamical control becomes a full-fledged resource theory of purity.

The above examples show that not every operational theory supports a sensible resource theory of purity. Motivated by the examples, we put forward the following definition:

Definition 6 *A* theory of purity *is a resource theory of dynamical control where every state ρ can be compared with at least one pure state. The theory is called* canonical *if every pure state is comparable to any other pure state.*

In this paper we will focus on canonical theories of purity.
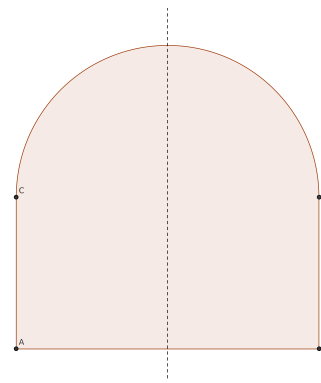
(a) Example of state space leading to a theory of dynamical control that cannot be interpreted as a theory of purity. Due to the shape of the state space, the reversible transformations can only be the identity and the reflection around the vertical axis. The states on the vertical sides are maximally controllable (and therefore, maximally resourceful) even though they are not pure.
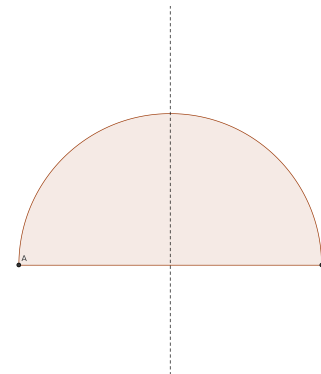
(b) Example of state space leading to a non-canonical theory of purity. In this case, maximal purity is equivalent to maximal resourcefulness: only the pure states are maximally controllable. However, some pure states are inequivalent resources, meaning that purity is not the only resource into play.
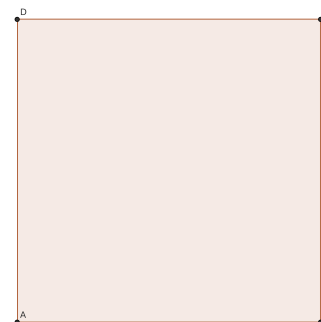
(c) Example of state space compatible with a canonical theory of purity. Here the set of maximally controllable states coincides with the set of pure states, and, in addition, all pure states are equivalent resources.

Figure 1. Three different examples of theories of dynamical control.

Proposition 4 *The following are equivalent:*

1. *the theory is a canonical theory of purity*

2. *for every system* A*, the group of reversible channels acts transitively on the set of pure states.*

3. *for every system* A*, there exists at least one state that is more controllable than every state.*

The proof is provided in Appendix B.

Combining all the statements of proposition 4, one can see that in a canonical theory of purity every pure state is more controllable than any other state.

Starting from Hardy's work [43], the transitivity of the action of reversible channels on pure states has featured in a number of axiomatizations of quantum theory, either directly as an axiom [90–92] or indirectly as a consequence of an axiom, as in the case of the Purification axiom [48]. Proposition 4 provides a new motivation for this axiom, now identified as a necessary and sufficient condition for a well-behaved theory of purity.

In a canonical theory of purity, we say that ρ is *purer* than ρ' if $\rho \succeq \rho'$ and we adopt the notation $\rho \succeq_{\mathrm{pur}} \rho'$. In this case, we also say that ρ' is *more mixed* than ρ, denoted by $\rho' \succeq_{\mathrm{mix}} \rho$. When $\rho \succeq_{\mathrm{mix}} \rho'$ and $\rho' \succeq_{\mathrm{mix}} \rho$ we say that ρ and ρ' are *equally mixed*, denoted by $\rho \simeq_{\mathrm{mix}} \rho'$. Clearly, every two states that differ by a reversible channel are equally mixed. The converse is true for finite dimensional systems, thanks to Eq. (9).

C. Maximally mixed states

We say that a state $\chi \in \mathsf{St}(A)$ is *maximally mixed* if it satisfies the property

$$\forall \rho \in \mathsf{St}(A): \quad \rho \succeq_{\mathrm{mix}} \chi \quad \Longrightarrow \quad \rho = \chi \,.$$

Maximally mixed states can be characterized as the states that are invariant under all reversible channels:

Proposition 5 *A state* $\chi \in \mathsf{St}(A)$ *is maximally mixed if and only if it is invariant, i.e. if and only if* $\chi = \mathcal{U}\chi$ *for every reversible channel* $\mathcal{U} : A \to A$.

We omit the proof, which is straightforward. Note that maximally mixed states do not exist in every theory: for example, infinite-dimensional quantum systems have no maximally mixed density operator, i.e. no trace-class operator that is invariant under the action of the full unitary group. For finite-dimensional canonical theories, however, the maximally mixed state exists and is unique under the standard assumption of compactness of the state space [47, 91]. In this case, the state χ is not only a maximal element, but also the *maximum* of the relation $\succeq_{\mathrm{mix}}$, namely,

$$\chi \succeq_{\mathrm{mix}} \rho \qquad \forall \rho \in \mathsf{St}(A) \,. \tag{10}$$
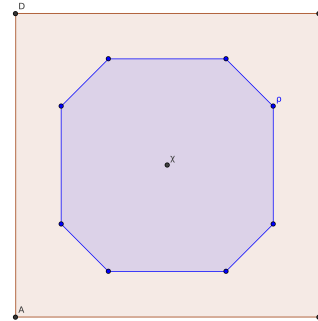
Figure 2. Mixedness relation for the state space of a square bit: the vertices of the octagon represent the states that can be reached from a given state ρ via reversible transformations. Their convex hull is the set of states that are more mixed than ρ. Note that it contains the invariant state χ, which can be characterized as the maximally mixed state.

This is in analogy with the quantum case, where the maximally mixed state is given by the density matrix $\chi = I/d$, where I is the identity operator on the system's Hilbert space and d is the Hilbert space dimension. Another example of finite-dimensional canonical theory is provided by the square bit:

Example 4 *Consider a system whose state space is a square, as in Fig. 1(c) and pick a generic (mixed) state* ρ*. The states that are more mixed than* ρ *are obtained by applying all possible reversible transformations to* ρ *(i.e. all the elements of the dihedral group* D_4*) and taking the convex hull of the orbit. The set of all states that are more mixed than* ρ *is an octagon, depicted in blue in Fig. 2. All the vertexes of the octagon are equally mixed. The centre of the square is the maximally state* χ*, the unique invariant state of the system.*

VI. ENTANGLEMENT-THERMODYNAMICS DUALITY

In quantum theory, it is well known that the ordering of pure bipartite states according to the degree of entanglement is equivalent to the ordering of their marginals according to the degree of mixedness [31–34, 93]. In this section we will prove the validity of this equivalence based only on first principles.

A. Purification

In order to establish the desired duality, we consider theories that satisfy the Purification Principle [47, 48]. Let us briefly summarize its content. We say that a state $\rho \in \mathsf{St}(A)$ has a *purification* if there exists a system B

and a pure state $\Psi \in \mathsf{PurSt}\,(A \otimes B)$ (the purification) such that

$$\Psi \begin{smallmatrix} A \\ B-\boxed{\mathrm{Tr}} \end{smallmatrix} = \boxed{\rho}\!-\!A$$

We say that the purification is *essentially unique* if every other purification Ψ' with the same purifying system B satisfies the condition

$$\Psi' \begin{smallmatrix} A \\ B \end{smallmatrix} = \Psi \begin{smallmatrix} A \\ B-\boxed{\mathcal{U}}-B \end{smallmatrix}, \qquad (11)$$

for some reversible transformation $\mathcal{U} : B \rightarrow B$. With these definitions, the Purification Principle can be phrased as

Axiom 3 (Purification [47, 48]) *Every state has a purification. Every purification is essentially unique.*

Purification has a number of important consequences. First of all, it implies that the group of reversible transformations acts transitively on the set of pure states:

Proposition 6 (Transitivity) *For every system* B *and every pair of pure states* $\psi, \psi' \in \mathsf{PurSt}\,(B)$ *there exists a reversible channel* $\mathcal{U} : B \to B$ *such that* $\psi' = \mathcal{U}\psi$.

The existence of a reversible transformation connecting ψ and ψ' is a consequence of the essential uniqueness of purification [Eq. (11)], in the special case A = I [47]. Since all pure states are equivalent under reversible transformations, every theory with purification gives rise to a *canonical* theory of purity, in the sense of definition 6. One could take this fact as a further indication that Purification is a good starting point for a well-behaved thermodynamics.

Another important consequence of Purification is the existence of entanglement:

Proposition 7 (Existence of entangled states) *For every pair of systems* A *and* B, *a pure state of* $A \otimes B$ *is entangled if and only if its marginal on system* A *is mixed.*

Proof. Let us denote the pure bipartite state by Ψ. If Ψ is not entangled, then it must be a product of two pure states, say $\Psi = \alpha \otimes \beta$. Clearly, this implies that the marginal on system A is pure.

Conversely, suppose that the marginal of Ψ on system A is pure and denote it by α. Then, for every pure state $\beta' \in \mathsf{PurSt}\,(B)$, the product state $\Psi' = \alpha \otimes \beta'$ is pure, thanks to Purity Preservation. Now, Ψ and Ψ' are two purifications of α. By the essential uniqueness of purification, one must have $\Psi = (\mathcal{I}_A \otimes \mathcal{U}_B)\,\Psi'$ for some reversible transformation $\mathcal{U}_B$ acting on system B. Hence, we have $\Psi = \alpha \otimes \beta$, with $\beta = \mathcal{U}_B \beta'$. ∎

Finally, Purification implies the *steering property* [2, 94], stating that every ensemble decomposition of a given state can be generated by a measurement on the purifying system:

Proposition 8 (Steering) *Let* ρ *be a state of system* A *and let* $\Psi \in \mathsf{PurSt}\,(A \otimes B)$ *be a purification of* ρ. *For every ensemble of states* $\{\rho_i\}_{i \in \mathsf{X}}$ *such that* $\sum_i \rho_i = \rho$, *there exists a measurement* $\{b_i\}_{i \in \mathsf{X}}$ *on the purifying system* B *such that the following relation holds*

$$\boxed{\rho_i}\!-\!A = \Psi \begin{smallmatrix} A \\ B-\boxed{b_i} \end{smallmatrix} \qquad \forall i \in \mathsf{X}.$$

See theorem 6 of Ref. [47] for the proof. The steering property will turn out to be essential in establishing the duality between entanglement and thermodynamics.

B. One-way protocols transforming pure states into pure states

The operational Lo-Popescu theorem guarantees that every LOCC protocol acting on a pure bipartite input state can be simulated by a one-way protocol. Purification buys us an extra bonus: not only is the protocol one-way, but also all the conditional operations are reversible.

Lemma 2 *Let* Ψ *and* Ψ' *be pure states of* $A \otimes B$. *Under the validity of Purification and Purity Preservation, every one-way protocol transforming* Ψ *into* Ψ' *can be simulated by a one-way protocol where all conditional operations are reversible.*

Proof. Suppose that Ψ can be transformed into Ψ' via a one-way protocol where Alice performs a test $\{\mathcal{A}_i\}_{i \in \mathsf{X}}$ and Bob performs a channel $\mathcal{B}^{(i)}$ conditional on the outcome i. By definition, we have

$$\sum_i \left(\Psi \begin{smallmatrix} A-\boxed{\mathcal{A}_i}-A \\ B \searrow \boxed{\mathcal{B}^{(i)}}-B \end{smallmatrix} \right) = \Psi' \begin{smallmatrix} A \\ B \end{smallmatrix}.$$

Since Ψ' is pure, this implies that there exists a probability distribution $\{p_i\}$ such that

$$\Psi \begin{smallmatrix} A-\boxed{\mathcal{A}_i}-A \\ B \searrow \boxed{\mathcal{B}^{(i)}}-B \end{smallmatrix} = p_i\, \Psi' \begin{smallmatrix} A \\ B \end{smallmatrix} \qquad (12)$$

for every outcome i. Now, without loss of generality each transformation $\mathcal{A}_i$ can be assumed to be pure (if not, one can always decompose it into pure transformations, thanks to the pure decomposition property). Then, Purity Preservation guarantees that the normalized state Ψ_i defined by

$$\Psi_i \begin{smallmatrix} A \\ B \end{smallmatrix} := p_i^{-1} \, \Psi \begin{smallmatrix} A-\boxed{\mathcal{A}_i}-A \\ B \end{smallmatrix} \qquad (13)$$

is pure. With this definition, Eq. (12) becomes

$$\Psi_i \begin{smallmatrix} A \\ B-\boxed{\mathcal{B}^{(i)}}-B \end{smallmatrix} = \Psi' \begin{smallmatrix} A \\ B \end{smallmatrix}.$$

Tracing out system B on both sides one obtains

$$\left(\Psi' \;\substack{A \\ B}\right)\!\!\!-\!\!\mathrm{Tr} \;=\; \left(\Psi_i \;\substack{A \\ B}\right)\!\!-\!\mathcal{B}^{(i)}\!-\!\mathrm{Tr}$$

$$=\; \left(\Psi_i \;\substack{A \\ B}\right)\!\!-\!\mathrm{Tr} \;,$$

the second equality coming from the normalization of the channel $\mathcal{B}^{(i)}$ (cf. proposition 1). Hence, the pure states Ψ_i and Ψ' have the same marginal on A. By the essential uniqueness of Purification, they must differ by a reversible channel $\mathcal{U}^{(i)}$ on the purifying system B, namely

$$\left(\Psi_i \;\substack{A \\ B}\right)\!\!-\!\mathcal{U}^{(i)}\!- \;=\; \left(\Psi' \;\substack{A \\ B}\right) \tag{14}$$

In conclusion, we obtained

$$\left(\Psi \;\substack{A-\mathcal{A}_i-A \\ B}\right)\!\!-\!\mathcal{B}^{(i)}\!- \;=$$
$$=\; p_i \left(\Psi' \;\substack{A \\ B}\right)$$
$$=\; p_i \left(\Psi_i \;\substack{A \\ B}\right)\!\!-\!\mathcal{U}^{(i)}\!-$$
$$=\; \left(\Psi \;\substack{A-\mathcal{A}_i-A \\ B}\right)\!\!-\!\mathcal{U}^{(i)}\!- \;,$$

where we have used Eqs. (12), (14), and (13). In other words, the initial protocol can be simulated by a protocol where Alice performs the test $\{\mathcal{A}_i\}$ and Bob performs the reversible transformation $\mathcal{U}^{(i)}$ conditionally on the outcome i. ∎

The reduction to one-way protocols with reversible operations is the key to connect the resource theory of entanglement with the resource theory of purity. The duality between these two resource theories will be established in the next subsections.

C. The more entangled a pure state, the more mixed its marginals

We start by proving one direction of the entanglement-thermodynamics duality: if a state is more entangled than another, then the marginals of the former are more mixed than the marginals of the latter:

Lemma 3 *Let Ψ and Ψ' be two pure states of system $A \otimes B$ and let ρ, ρ' and σ, σ' be their marginals on system A and B, respectively. Under the validity of Purification, Purity Preservation, and Local Exchangeability, if Ψ is more entangled than Ψ', then ρ (σ) is more mixed than ρ' (σ').*

Proof. By the operational Lo-Popescu theorem, we know that there exists a one-way protocol transforming Ψ into Ψ'. Moreover, thanks to Purification, the conditional operations in the protocol can be chosen to be reversible (lemma 2). Let us choose a protocol with classical communication from Alice to Bob, in which Alice performs the test $\{\mathcal{A}_i\}_{i \in \mathsf{X}}$ and Bob performs the reversible transformation $\mathcal{U}^{(i)}$ conditional on the outcome i. Since Ψ' is pure, we must have

$$\left(\Psi \;\substack{A-\mathcal{A}_i-A \\ B}\right)\!\!-\!\mathcal{U}^{(i)}\!- \;=\; p_i \left(\Psi' \;\substack{A \\ B}\right) \qquad \forall i \in \mathsf{X},$$

where $\{p_i\}$ is a suitable probability distribution. Denoting by $\mathcal{V}^{(i)}$ the inverse of $\mathcal{U}^{(i)}$ and applying it on both sides of the equation, we obtain

$$\left(\Psi \;\substack{A-\mathcal{A}_i-A \\ B}\right) \;=\; p_i \left(\Psi' \;\substack{A \\ B}\right)\!\!-\!\mathcal{V}^{(i)}\!- \;.$$

Summing over all outcomes the equality becomes

$$\left(\Psi \;\substack{A-\mathcal{A}-A \\ B}\right) \;=\; \left(\Psi' \;\substack{A \\ B}\right)\!\!-\!\mathcal{R}\!- \;, \tag{15}$$

with $\mathcal{A} := \sum_i \mathcal{A}_i$ and $\mathcal{R} := \sum_i p_i \mathcal{V}^{(i)}$. Finally, we obtain

$$\left(\sigma\right)\!-\!B \;=\; \left(\Psi \;\substack{A-\mathrm{Tr} \\ B}\right)$$
$$=\; \left(\Psi \;\substack{A-\mathcal{A}-A-\mathrm{Tr} \\ B}\right)$$
$$=\; \left(\Psi' \;\substack{A-\mathrm{Tr} \\ B}\right)\!\!-\!\mathcal{R}\!-$$
$$=\; \left(\sigma'\right)\!-\!\mathcal{R}\!- \;,$$

where we have used the normalization of channel $\mathcal{A}$ in the second equality and Eq. (15) in the third. Since $\mathcal{R}$ is a RaRe channel by construction, we have proved that σ is more mixed than σ'. The fact that ρ is more mixed than ρ' can be proved by the same argument, starting from a one-way protocol with classical communication from Bob to Alice and with reversible operations on Alice's side. ∎

The relation between degree of entanglement of a pure state and degree of mixedness of its marginals holds not only for bipartite states, but also for multipartite states. Indeed, suppose that Ψ and Ψ' are two pure states of system $A_1 \otimes A_2 \otimes \cdots \otimes A_N$ and that Ψ is more entangled

than Ψ', in the sense that there exists a (multipartite) LOCC protocol converting Ψ into Ψ'. For every subset $\mathsf{S} \subset \{1, \ldots, N\}$ one can define $\mathrm{A} := \bigotimes_{n \notin \mathsf{S}} \mathrm{A}_n$ and $\mathrm{B} := \bigotimes_{n \in \mathsf{S}} \mathrm{A}_n$ and apply lemma 3. As a result, one obtains that the marginals of Ψ are more mixed than the marginals of Ψ' *on every subsystem*.

D. The more mixed a state, the more entangled its purification

We now prove the converse direction of the entanglement-thermodynamics duality: if a state is more mixed than another, then its purification is more entangled. Remarkably, the proof of this fact requires only the validity of Purification.

Lemma 4 *Let ρ and ρ' be two states of system A and let Ψ (Ψ') be a purification of ρ (ρ'), with purifying system B. Under the validity of Purification, if ρ is more mixed than ρ', then Ψ is more entangled than Ψ'.*

Proof. By hypothesis, one has

$$\rho' - \mathrm{A} - \mathcal{R} - \mathrm{A} = \rho - \mathrm{A}$$

for some RaRe channel $\mathcal{R} := \sum_i p_i \mathcal{U}^{(i)}$. Let us define the bipartite state Θ as

$$\Theta \!\! \begin{array}{c} \mathrm{A} \\ \mathrm{B} \end{array} := \Psi' \!\! \begin{array}{c} \mathrm{A} - \mathcal{R} - \mathrm{A} \\ \mathrm{B} \end{array}$$

$$\equiv \sum_i p_i \, \Psi' \!\! \begin{array}{c} \mathrm{A} - \mathcal{U}^{(i)} - \mathrm{A} \\ \mathrm{B} \end{array} \tag{16}$$

By construction, Θ is an extension of ρ: indeed, one has

$$\Theta \!\! \begin{array}{c} \mathrm{A} \\ \mathrm{B} - \mathrm{Tr} \end{array} = \Psi' \!\! \begin{array}{c} \mathrm{A} - \mathcal{R} - \mathrm{A} \\ \mathrm{B} \qquad\quad \mathrm{Tr} \end{array}$$

$$= \rho' - \mathrm{A} - \mathcal{R} - \mathrm{A}$$

$$= \rho - \mathrm{A} \; .$$

Let us take a purification of Θ, say $\Gamma \in \mathsf{PurSt}(\mathrm{A} \otimes \mathrm{B} \otimes \mathrm{C})$. Clearly, Γ is a purification of ρ, since one has

$$\Gamma \!\! \begin{array}{c} \mathrm{A} \\ \mathrm{B} - \mathrm{Tr} \\ \mathrm{C} - \mathrm{Tr} \end{array} = \Theta \!\! \begin{array}{c} \mathrm{A} \\ \mathrm{B} - \mathrm{Tr} \end{array} = \rho - \mathrm{A} \; .$$

Then, the essential uniqueness of purification implies that Γ must be of the form

$$\Gamma \!\! \begin{array}{c} \mathrm{A} \\ \mathrm{B} \\ \mathrm{C} \end{array} = \Psi \!\! \begin{array}{c} \mathrm{A} \\ \mathrm{B} \\ \gamma \; \mathrm{C} \end{array} \!\! \mathcal{U} \!\! \begin{array}{c} \\ \mathrm{B} \\ \mathrm{C} \end{array} , \tag{17}$$

for some reversible transformation $\mathcal{U}$ and some pure state γ. In other words, Ψ can be transformed into Γ by local operations on Bob's side.

Now, Eq. (16) implies that the states $\left\{ p_i \left(\mathcal{U}^{(i)} \otimes \mathcal{I}_\mathrm{B} \right) \Psi \right\}_{i \in \mathsf{X}}$ are an ensemble decomposition of Θ. Hence, the steering property (proposition 8) implies that there exists a measurement $\{c_i\}_{i \in \mathsf{X}}$ on C such that

$$p_i \, \Psi' \!\! \begin{array}{c} \mathrm{A} - \mathcal{U}_i - \mathrm{A} \\ \mathrm{B} \end{array} = \Gamma \!\! \begin{array}{c} \mathrm{A} \\ \mathrm{B} \\ \mathrm{C} - c_i \end{array} \qquad \forall i \in \mathsf{X}. \tag{18}$$

Combining Eqs. (17) and (18), we obtain the desired result.

$$\Psi' \!\! \begin{array}{c} \mathrm{A} \\ \mathrm{B} \end{array} = \sum_{i \in \mathsf{X}} \Psi \!\! \begin{array}{c} \mathrm{A} \qquad\quad \mathcal{U}^{(i)} - \mathrm{A}' \\ \mathrm{B} - \mathcal{B}_i - \mathrm{B}' \end{array} ,$$

where $\{\mathcal{B}_i\}_{i \in \mathsf{X}}$ is the test defined by

$$\mathrm{B} - \mathcal{B}_i - \mathrm{B} := \begin{array}{c} \qquad \mathrm{B} \qquad\quad \mathrm{B} \\ \gamma \; \mathrm{C} \, \mathcal{U} \, \mathrm{C} - c_i \end{array}$$

In conclusion, if the marginal state of Ψ is more mixed than the marginal state of Ψ', then Ψ can be converted into Ψ' by a one-way LOCC protocol. ∎

E. The duality

Combining lemmas 3 and 4 we identify the degree of entanglement of a pure bipartite state with the degree of mixedness of its marginals:

Theorem 2 (Entanglement-thermodynamics duality) *Let Ψ and Ψ be two pure states of system $\mathrm{A} \otimes \mathrm{B}$ and let ρ, ρ' and σ, σ' be their marginals on system A and B, respectively. Under the validity of Purification, Purity Preservation, and Local Exchangeability, the following statements are equivalent:*

1. Ψ is more entangled than Ψ'

2. ρ is more mixed than ρ'

3. σ is more mixed than σ'.

Proof. The implications $1 \implies 2$ and $1 \implies 3$ follow from lemma 3 and require the validity of all the three axioms. The implications $2 \implies 1$ and $3 \implies 1$ follow from lemma 4 and require only the validity of Purification. ∎

The duality can be illustrated by the commutative diagrams

$$
\begin{array}{ccc}
\Psi & \xrightarrow{\text{LOCC}} & \Psi' \\
{\scriptstyle \text{Tr}_B} \downarrow & & \downarrow {\scriptstyle \text{Tr}_B} \\
\rho & \xleftarrow{\text{RaRe}} & \rho'
\end{array}
\qquad
\begin{array}{ccc}
\Psi & \xrightarrow{\text{LOCC}} & \Psi' \\
{\scriptstyle \text{Tr}_A} \downarrow & & \downarrow {\scriptstyle \text{Tr}_A} \\
\sigma & \xleftarrow{\text{RaRe}} & \sigma'
\end{array}
$$

and is implemented operationally by discarding one of the component systems. Another illustration of the duality is via the diagram

$$
\begin{array}{ccc}
\Psi & \xleftarrow{\text{LOCC}} & \Psi' \\
{\scriptstyle \text{purification}} \uparrow & & \uparrow {\scriptstyle \text{purification}} \\
\rho & \xrightarrow{\text{RaRe}} & \rho'
\end{array}
\qquad .
$$

Here the map implementing the duality is (a choice of) purification. Such a map cannot be realized as a physical operation [47]. Instead, it corresponds to the theoretical operation of modelling mixed states as marginals of pure states.

VII. CONSEQUENCES OF THE DUALITY

In this section we discuss the simplest consequences of the entanglement-thermodynamics duality, including the relation between maximally mixed and maximally entangled states, as well as a link between information erasure and generation of entanglement. From now on, the axioms used to derive the duality will be treated as standing assumptions and will not be written explicitly in the statement of the results.

A. Equivalence under local reversible transformations

Let us start from the easiest consequence of the duality:

Corollary 1 *Let Ψ and Ψ' be two states of system $A \otimes B$, with A finite-dimensional. Then, Ψ and Ψ' are equally entangled if and only if they are equivalent under local reversible transformations, namely*

$$
\Psi' = (\mathcal{U} \otimes \mathcal{V}) \Psi
$$

where $\mathcal{U}$ and $\mathcal{V}$ are reversible transformations acting on A and B, respectively.

This result, proved in appendix C, guarantees that the equivalence classes under the entanglement relation have a simple structure, inherited from the reversible dynamics allowed by the theory. For finite systems, pure bipartite entanglement is completely characterized by the quotient of the set of pure states under local reversible transformations.

B. Duality for states on different systems

Theorem 2 concerns the convertibility of states of the same system. To generalize it to arbitrary systems, it is enough to observe that the tensor product with local pure states does not change the degree of entanglement: for arbitrary pure states Ψ, α', and β' of systems $A \otimes B$, A', and B' one has

$$
\Psi \simeq_{\text{ent}} \alpha' \otimes \Psi \otimes \beta' , \tag{19}
$$

relative to the bipartition $(A' \otimes A) \otimes (B \otimes B')$. As a consequence, one has the equivalence

$$
\Psi \succeq_{\text{ent}} \Psi' \quad \Longleftrightarrow \quad \alpha' \otimes \Psi \otimes \beta' \succeq_{\text{ent}} \alpha \otimes \Psi' \otimes \beta
$$

for arbitrary pure states $\alpha, \alpha', \beta, \beta'$ of A, A', B, B', respectively. This fact leads directly to the generalization of the duality to states of different systems:

Corollary 2 *Let Ψ and Ψ' be two pure states of systems $A \otimes B$ and $A' \otimes B'$, respectively, and let ρ, ρ', σ and σ' be their marginals on system A, A', B and B' respectively. Under the validity of Purification, Purity Preservation, and Local Exchangeability, the following statements are equivalent:*

1. *Ψ is more entangled than Ψ'*

2. *$\rho \otimes \alpha'$ is more mixed than $\alpha \otimes \rho'$ for every pair of pure states $\alpha \in \mathsf{PurSt}(A)$ and $\alpha' \in \mathsf{PurSt}(A')$.*

3. *$\sigma \otimes \beta'$ is more mixed than $\beta \otimes \sigma'$ for every pair of pure states $\beta \in \mathsf{PurSt}(B)$ and $\beta' \in \mathsf{PurSt}(B')$.*

The duality is now implemented by the operation of discarding systems and preparing pure states, as illustrated by the commutative diagrams

$$
\begin{array}{ccc}
\Psi & \xrightarrow{\text{LOCC}} & \Psi' \\
{\scriptstyle \text{Tr}_B \otimes \alpha'} \downarrow & & \downarrow {\scriptstyle \alpha \otimes \text{Tr}_B} \\
\rho \otimes \alpha' & \xleftarrow{\text{RaRe}} & \alpha \otimes \rho'
\end{array}
\qquad
\begin{array}{ccc}
\Psi & \xrightarrow{\text{LOCC}} & \Psi' \\
{\scriptstyle \text{Tr}_A \otimes \beta'} \downarrow & & \downarrow {\scriptstyle \beta \otimes \text{Tr}_A} \\
\sigma \otimes \beta' & \xleftarrow{\text{RaRe}} & \beta \otimes \sigma'
\end{array}
$$

At this point, a cautionary remark is in order. Inspired by Eq. (19) one may be tempted compare the degree of mixedness of states of different systems, by postulating the relation

$$
\rho \simeq_{\text{mix}} \rho \otimes \alpha' \tag{20}
$$

for arbitrary states ρ and arbitrary pure states α'. The appeal of this choice is that the duality would maintain the simple form

$$\Psi \succeq_{\text{ent}} \Psi' \quad \Longleftrightarrow \quad \rho \succeq_{\text{mix}} \rho',$$

even for states of different systems. However, Eq. (20) would trivialize the resource theory of purity: as a special case, it would imply the relation $1 \simeq_{mix} \alpha'$ for a generic pure state α', meaning that pure states can be freely generated. Since in a canonical theory of purity pure states are the most resourceful, having pure states for free would mean having every state for free.

Another way to compare states of different systems according to their degree of purity would be to postulate the relation

$$\rho \simeq_{\text{pur}} \rho \otimes \chi_{\text{B}}, \tag{21}$$

where χ_{B} is the maximally mixed state of system B (assuming that such a state exists). The rationale for this choice would be that χ_{B} is the "minimum-resource state" in the resource theory of purity and therefore one may want to consider it as free. This choice would not trivialize the resource theory of purity, but would break the duality with the resource theory of entanglement. Indeed, Eq. (21) would imply as a special case $1 \simeq_{\text{pur}} \chi_{\text{B}}$, meaning that maximally mixed states can be freely generated from nothing. Clearly, this is not the case for their purifications, which are entangled and cannot be generated freely by LOCC. In summary, refraining from comparing mixed states on different systems seems to be the best way to approach the duality between the resource theory of entanglement and the resource theory of purity.

C.   Measures of mixedness and measures of entanglement

The duality provides the foundation for the definition of quantitative measures of entanglement. In every resource theory, one can define measures of "resourcefulness", by introducing functions that are non-increasing under the set of free operations [85]. In the resource theory of entanglement, this leads to the notion of entanglement monotones:

Definition 7 *An* entanglement monotone *for system* $A \otimes B$ *is a function* $E : \mathsf{St}(A \otimes B) \to \mathbb{R}$ *satisfying the condition*

$$E(\rho) \geq E(\rho') \qquad \forall \rho, \rho' \in \mathsf{St}(A \otimes B), \rho \succeq_{\text{ent}} \rho'.$$

More generally, one may want to compare entangled states on different systems. In this case, an entanglement monotone E is a *family* of functions $E = \{E_{\text{A}\otimes\text{B}} \mid A, B \in \mathsf{Sys}\}$ satisfying the condition

$$E_{\text{A}\otimes\text{B}}(\rho) \geq E_{\text{A}'\otimes\text{B}'}(\rho')$$

for every pair of states $\rho \in \mathsf{St}(A \otimes B)$ and $\rho' \in \mathsf{St}(A' \otimes B')$ satisfying $\rho \succeq_{\text{ent}} \rho'$.

Similarly, one can define monotones in the resource theory of purity:

Definition 8 *A* purity monotone *for system* A *is a function* $P : \mathsf{St}(A) \to \mathbb{R}$ *satisfying the condition*

$$P(\rho) \geq P(\rho') \qquad \forall \rho, \rho' \in \mathsf{St}(A), \rho \succeq_{\text{pur}} \rho'.$$

Recall that in our resource theory of purity we abstain from comparing states on different systems, for the reasons discussed in the end of the previous subsection. Purity monotones give a further indication that the definition of purity in terms of random reversible channels is a sensible one: indeed, if we restrict our attention to the classical case, the notion of purity monotone introduced here coincides with the canonical notion of Schur-convex function in the theory of majorization [65] (see appendix D). Schur-convex functions are the key tool to construct entropies and other measures of mixedness in classical statistical mechanics, and have applications in a number of diverse fields [95].

Constructing purity monotones is fairly easy. For example, every function that is convex and invariant under reversible transformations is a purity monotone:

Proposition 9 *Let* $P : \mathsf{St}(A) \to \mathbb{R}$ *be a function satisfying*

1. *convexity:* $P\left(\sum_i p_i \rho_i\right) \leq \sum_i p_i P(\rho_i)$ *for every set of states* $\{\rho_i\}$ *and for every probability distribution* $\{p_i\}$, *and*

2. *invariance under reversible transformations:* $P(\mathcal{U}\rho) = P(\rho)$ *for every state* ρ *and for every reversible transformation* $\mathcal{U}$.

Then, P *is a purity monotone.*

The proof is elementary and is presented in the appendix D for the convenience of the reader. We highlight that the above proposition is the natural extension of a well-known result in majorization theory, namely that every convex function that is symmetric in its variables is automatically Schur-convex [65]. Again, it is worth highlighting the perfect match of the operational notions discussed here with the canonical results about majorization.

Using proposition 9, one can construct purity monotones aplenty: for every convex function $f : \mathbb{R} \to \mathbb{R}$ one can define the f-purity $P_f : \mathsf{St}(A) \to \mathbb{R}$ as

$$P_f(\rho) := \sup_{\boldsymbol{a} \text{ pure}} \sum_{x \in \mathsf{X}} f(p_x) \qquad p_x := (a_x|\rho),$$

where the supremum runs over all pure measurements $\boldsymbol{a} = \{a_x\}_{x \in \mathsf{X}}$ and over all outcome spaces X. It is easy to verify that every f-purity is convex and invariant under reversible transformations, and therefore is a purity monotone. In the special case of the function $f(x) = x \log x$, one has

$$P_f(\rho) = -H(\rho), \tag{22}$$

where H is the *measurement entropy* [96–98], namely the minimum over all pure measurements of the Shannon entropy of the probability distribution resulting from the measurement. In the case of $f(x) = x^2$ one obtains instead an generalized notion of "purity", which in the quantum case coincides with the usual notion $P(\rho) = \text{Tr}\left(\rho^2\right)$.

Another way to construct purity monotones is by using norms on the state space: thanks to proposition 9, every norm that is invariant under reversible transformations leads to a purity monotone. For systems that have an invariant state, an easy example is given by the *operational distance*

$$P(\rho) := \frac{1}{2} \left\| \rho - \chi \right\|,$$

where $\left\|\cdot\right\|$ is the *operational norm*, defined as $\left\|\delta\right\| := \sup_{a_0 \in \text{Eff}(A)} (a_0|\delta) - \inf_{a_1 \in \text{Eff}(A)} (a_1|\delta)$ [47], and χ is the invariant state. Another example of purity monotone induced by a norm is the notion of purity introduced in Refs. [76, 77], based on the Schatten 2-norm. In the quantum case, this notion of purity coincides with the ordinary notion $P(\rho) = \text{Tr}\left(\rho^2\right)$ and therefore coincides with the f-purity with $f(x) = x^2$. It is not a priori clear whether the 2-norm purity coincides with the x^2-purity for more general theories.

Now, thanks to the duality we can turn every purity monotone into an entanglement monotone. Given a purity monotone $P : \mathsf{St}(A) \to \mathbb{R}$, we can define the pure state entanglement monotone $E : \mathsf{PurSt}(A \otimes B) \to \mathbb{R}$ as

$$E(\Psi) := g\left[P(\rho)\right], \qquad \rho = \text{Tr}_B \Psi, \qquad (23)$$

where $g : \mathbb{R} \to \mathbb{R}$ is any monotonically decreasing function ($f(x) \leq f(y)$ for $x > y$). Here the monotonically decreasing behaviour of g implements the reversing of arrows in the duality. Furthermore, if the functions P and f have suitable convexity properties, the entanglement monotone can be extended from pure states to arbitrary states using the *convex roof construction* [22, 78, 79]. Specifically, one has the following

Corollary 3 *Let* $P : \mathsf{St}(A) \to \mathbb{R}$ *be a convex purity monotone,* $g : \mathbb{R} \to \mathbb{R}$ *be a concave, monotonically decreasing function, and* $E : \mathsf{PurSt}(A \otimes B) \to \mathbb{R}$ *be the pure state entanglement monotone defined in Eq. (23). Then, the convex roof extension* $E : \mathsf{St}(A \otimes B) \to \mathbb{R}$ *defined by*

$$E(\Sigma) := \inf_{\substack{\{p_i, \Psi_i\} \\ \sum_i p_i \Psi_i = \Sigma}} \sum_i p_i E(\Psi_i)$$

is a convex entanglement monotone.

The proof is the same as in the quantum case [78]. An easy way to generate entanglement measures is to pick an f-purity and take its negative, which corresponds to the choice $g(x) = -x$. For example, the choice $f(x) = x \log x$ leads to a generalization of the *entanglement of formation* [22] to all theories satisfying the duality.

D. Maximally entangled states

As a consequence of the duality, there exists a correspondence between maximally mixed and "maximally entangled" states, the latter being defined as follows

Definition 9 *A pure state* Φ *of system* $A \otimes B$ *is maximally entangled if no other pure state of* $A \otimes B$ *is more entangled than* Φ*, except for the states that are equivalent to* Φ *under local reversible transformations—i.e. if for every* $\Psi \in \mathsf{PurSt}(A \otimes B)$ *one has*

$$\Psi \succeq_{\text{ent}} \Phi \quad \implies \quad \Psi = (\mathcal{U} \otimes \mathcal{V}) \Phi$$

for some reversible transformations $\mathcal{U} : A \to A$ *and* $\mathcal{V} : B \to B$*.*

Theorem 2 directly implies the following.

Corollary 4 *The purification of a maximally mixed state is maximally entangled.*

Proof. Suppose that $\Psi \in \mathsf{PurSt}(A \otimes B)$ is more entangled than Φ, where Φ is a purification of the maximally mixed state of system A (assuming such a state exists for system A). By theorem 2, the marginal of Ψ on system A, denoted by ρ, must satisfy $\rho \succeq_{\text{mix}} \chi$. Since χ is maximally mixed, this implies $\rho = \chi$. The uniqueness of purification then implies the condition $\Psi = (\mathcal{I}_A \otimes \mathcal{V}_B) \Phi$ for some reversible transformation $\mathcal{V}_B$ on B. ∎

As noted earlier in the paper, under the standard assumptions of convexity and compactness of the state space, the maximally mixed state is not only a maximal element of the mixedness relation, but also the maximum [cf. Eq. (10)]. Similarly, under the same standard assumptions, it is immediate to obtain that the purification of a maximally mixed state is more entangled than every state, namely

$$\Phi \succeq_{\text{ent}} \Sigma \qquad \forall \Sigma \in \mathsf{St}(A \otimes B).$$

The relation follows directly from theorem 2 when Σ is a pure state and in the general case can be proved by convexity, using the fact that the set of LOCC channels is closed under convex combinations.

E. Duality between information erasure and entanglement generation

The entanglement-thermodynamics duality establishes a link between the two tasks of erasing information and generating entanglement. By *erasing information* we mean resetting a mixed state to a fixed pure state of the same system [80]. Clearly, erasure is a costly operation in the resource theory of purity: there is no way to transform a non-pure state into a pure state by using only RaRe channels (cf. proposition 3). The dual operation in the resource theory of entanglement is the generation

of entangled states from product states. By the duality, the impossibility of erasing information by RaRe channels and the impossibility of generating entanglement by LOCC are one and the same thing.

The relation between information erasure and entanglement generation suggests that the cost of erasing a mixed state ρ could be identified with the cost of generating the corresponding entangled state Ψ. For example, one may choose a fixed entangled state Φ as a reference "unit of entanglement" and ask how many copies of Φ are needed to generate Ψ through LOCC operations. The number of entanglement units needed to generate Φ could then be taken as a measure of the cost of erasing ρ. We now explore this idea at the heuristic level, discussing first a model of erasure and then connecting it with the generation of entanglement. Suppose that erasure is implemented by *i)* performing a pure measurement, *ii)* writing down the outcome on a classical register, *iii)* conditionally on the outcome, performing a reversible transformation that brings the system to a fixed pure state, and finally *iv)* erasing the classical register. Of course, this model assumes that some systems described by the theory can act as "classical registers", meaning that they have perfectly distinguishable pure states. Assuming the validity of Landauer's principle at the classical level, the cost of erasing the classical register is then equal to the Shannon entropy of the outcomes multiplied by $k_B T$, k_B and T being the Boltzmann constant and the temperature, respectively [80]. Minimizing the entropy over all possible measurements at step *i)*, one would then obtain the measurement entropy, as defined in Eq. (22). Hence, the minimum cost for erasing ρ is given by $k_B T H(\rho)$. Note that this heuristic conclusion implicitly assumes that the operations *i-iii)* can be performed for free. This is the case in quantum theory, where *i)* the measurement attaining minimum Shannon entropy is projective and the overall transformation associated to it is a random unitary channel, *ii)* the measurement outcome can be written down via a unitary operation on the system and the classical register, and *iii)* the state of the system can be reset via another joint unitary operation. In physical theories other than quantum and classical theories, however, the request that the operations *i-iii)* are free is non-trivial and would need to be further analysed in terms of physical axioms.

Suppose now that we want to erase an *unknown* state ρ. Since the state is unknown, the relevant quantity here is the *worst case cost of erasure*, defined as the supremum of the cost over all possible states. Since the measurement entropy is monotone under the mixedness relation, in finite dimensions the supremum is attained for the maximally mixed state χ, so that the worst case cost of erasure is given by $k_B T H(\chi)$. This result allows us to make an interesting connection with the work by Brunner *et al* [99], who considered the task of erasure in general probabilistic theories. Specifically, they considered the number of states that can be perfectly distinguished by a measurement and adopted the logarithm of this number as a measure of the cost of erasure. In their analysis they considered probabilistic quantum and classical theories, as well as alternative theories with hypercube state spaces, wherein measurements can distinguish at most two states. In all these theories the logarithm of the dimension is exactly equal to the measurement entropy of the maximally mixed state. Thanks to this fact, erasure cost defined in Ref. [99] coincides with the worst case erasure cost defined above. It is an open question whether the two definitions coincide in all canonical theories of purity, and, if not, which conditions are needed for the two definitions to coincide.

Let us now look at erasure from the dual point of view. Since the duality inverts the order, the dual of an erasure protocol consisting of operations *i),i),iii)* and *iv)* will be an entanglement generation protocol consisting of the dual operations in the opposite order *iv),iii),ii)* and *i)*. The dual of *iv)* is an operation that generates a purification of the classical register. The duals of the free operations *i-iii)* are LOCC operations that convert the initial entangled state into the state Ψ. Now, by the duality we can measure the cost of generating Ψ in terms of the measurement entropy $H(\rho)$. But what is the operational meaning of this choice? Again, the duality suggests an answer. Classically, the Shannon entropy can be interpreted as the asymptotic rate at which random bits can be extracted from a given probability distribution. Dually, the inverse relation must hold between the purifications: Referring to the purification of a random bit as an *ebit*, we have that the Shannon entropy is the number of ebits needed asymptotically to generate the purification of a given probability distribution by LOCC. Minimizing over all probability distributions one can characterize the measurement entropy as the minimum number of ebits needed to asymptotically generate the state Ψ by LOCC. Although partly based on heuristics, the argument provides already a good illustration of the far reaching consequences of the entanglement-thermodynamics duality, which allowed to identify the cost of erasing a state with the number of ebits needed to generate its purification.

F. Entropy sinks and entanglement reservoirs

Let us consider now the task of erasure assisted by a catalyst, namely a system C whose state remains unaffected by the erasure operation. In this case, the operation of erasure transforms the product state $\rho \otimes \gamma \in \mathsf{St}(A \otimes C)$ into the state $\alpha_0 \otimes \gamma$ for some pure state $\alpha_0 \in \mathsf{PurSt}(A)$. By duality, it is immediate to see that catalyst-assisted erasure is equivalent to catalyst-assisted entanglement generation:

Corollary 5 *Let Ψ and Γ be two pure states of systems $A \otimes B$ and $C \otimes D$, respectively, and let ρ and γ be their marginals on systems A and C, respectively. Then, the following are equivalent*

1. ρ can be erased by a RaRe channel using γ as a catalyst

2. Ψ can be generated by a LOCC channel using Γ as a catalyst.

If such catalysts existed, they would behave like "entropy sinks", which absorb mixed states without becoming more mixed, or like "entanglement reservoirs", from which entanglement can be borrowed indefinitely. For example, suppose that Ψ can be generated freely using Γ as a catalyst. Then, every measure of pure state entanglement E that it is additive on product states would have to satisfy the relation

$$E(\Gamma) \geq E(\Psi) + E(\Gamma) .$$

Assuming that the measure assigns a strictly positive value to every entangled state. the above relation can only be satisfied if $E(\Gamma) = +\infty$. In other words, the catalyst's state must be infinitely entangled. It is then natural to ask whether the impossibility of infinitely entangled/infinitely mixed states follow from our axioms. The answer is affirmative in the finite-dimensional case, but counterexamples exist in infinite dimensions. For the finite-dimensional case, we have the following

Proposition 10 *Let* $A \otimes C$ *be a finite system. Then, it is impossible to erase a mixed state of* A *using* C *as a catalyst.*

The proof is presented in appendix E.

In the infinite-dimensional case, a heuristic counterexample is as follows: imagine a scenario where system C consists of an infinite chain of identical systems, with each system in the chain equivalent to A, namely $C = \bigotimes_{i \in \mathbb{Z}} A_i$, $A_i \simeq A$. Loosely speaking, we may choose the state γ to be the product state $\gamma = \gamma_L \otimes \gamma_R$, where γ_L is a state on the left side of the chain, consisting of infinite copies of the pure state α_0, and γ_R is a state on the right of the chain, consisting of infinite copies of the mixed state ρ. It is then natural to expect that the state $\rho \otimes \gamma$ can be reversibly transformed into the state $\alpha_0 \otimes \gamma$, simply by swapping system A with the first system on the left of the chain and subsequently shifting the whole chain by one place to the right.

The above counterexample is heuristic, because the notion of infinite tensor product is not defined in our formalism. However, infinite tensor products can be treated rigorously, at least in the quantum case, and the intuition of our counterexample turns out to be correct. In the dual task of catalytic entanglement generation, the rigorous version of this argument was presented by Keyl, Matsui, Schlingemann, and Werner [100], who exhibited an example of infinite spin chain from which arbitrarily large amounts of entanglement can be generated for free.

VIII.   SYMMETRIC PURIFICATION

As we observed in the previous section, the ability to erase information/generate entanglement for free has un-

desirable consequences for the resource theories of purity and entanglement. These scenarios can be excluded at the level of first principles, by postulating the following

Axiom 4 (No Entropy Sinks) *Random reversible dynamics cannot achieve erasure, even with the assistance of a catalyst.*

In addition to being a requirement for a sensible resource theory of entanglement, Axiom 4 has a surprising twist: in the context of the other axioms, it implies that Local Exchangeability is equivalent to the existence of *symmetric* purifications, defined as follows

Definition 10 *Let* ρ *be a state of system* A *and let* Ψ *be a pure state of* $A \otimes A$. *We say that* Ψ *is a* symmetric purification *of* ρ *if*

and

This definition leads us to an upgraded version of the Purification axiom:

Axiom 5 (Symmetric Purification) *Every state has a symmetric purification. Every purification is essentially unique.*

The key result is then given by the following:

Theorem 3 *In a causal theory satisfying Purity Preservation and No Entropy Sinks, the following axioms are equivalent:*

1. *Local Exchangeability and Purification*

2. *Symmetric Purification.*

The proof is presented in appendix F.

This result identifies Purity Preservation and Symmetric Purification as the key axioms at the foundation of the entanglement-thermodynamics duality and, ultimately, as strong candidates for a reconstruction of quantum thermodynamics from first principles. We stress that the axiom No Entropy Sinks is needed only for infinite-dimensional systems, while for finite dimensional systems its validity can be proved (cf. proposition 10).

One of the bonuses of Symmetric Purification is that the marginals of a pure state are "equivalent", in the following sense:

Proposition 11 *Let* Ψ *be a pure state of system* $A \otimes B$ *and let* ρ_A *and* ρ_B *be its marginals on systems* A *and* B, *respectively. Then, one has*

where α and β are pure states of A *and* B*, respectively, and $\mathcal{U}$ is a reversible transformation.*

The proof can be found in appendix G. As a consequence of this result, we have that the states $\rho_A \otimes \beta$ and $\alpha \otimes \rho_B$ have the same purity, for every possible purity monotone. Equivalently, we can say that the two marginal states have the same mixedness, for every measure of mixedness.

IX. CONCLUSIONS

While entanglement is not a uniquely quantum feature, the remarkable ways in which it is intertwined with thermodynamics appear to be far more specific. Understanding these links at the level of basic principles is expected to reveal new clues to the foundations of quantum theory, as well as to the foundations of thermodynamics. With this motivation in mind, we set out to search for the roots of the relation between entanglement and entropy, adopting an operational, theory-independent approach. We attacked the problem from what is arguably the most primitive link: the duality between the resource theory of entanglement (where the free transformations are those achievable by two spatially separated agents via local operations and classical communication) and the resource theory of purity (where the free transformations are those achievable by an agent who has limited control on the dynamics of the system). By the duality, every free operation in the resource theory of purity admits an equivalent description as a free operation in the resource theory of entanglement. The duality leads to an identification between measures of mixedness (i.e. lack of purity) and measures of pure bipartite entanglement. Under suitable conditions, the latter can be extended to measures of mixed-state entanglement.

Let us elaborate on the implications of our results. Our reconstruction of the entanglement-thermodynamics duality hints at a simple, physically motivated idea: the idea that nature should admit a fundamental level of description where all states are pure, all dynamics reversible, and all measurements pure. Two of our axioms clearly express this requirement: *i)* Purification is equivalent to the existence of a pure and reversible level of description for states and channels, and *ii)* Purity Preservation ensures that such a description remains consistent when different, possibly non-deterministic processes are connected. The remaining axiom, Local Exchangeability, appeared to be slightly more mysterious at first sight. Nevertheless, the duality clarified its significance: for every finite system (and more generally, for every system where all states have finite entanglement), Local Exchangeability is equivalent to the existence of *symmetric* purifications—that is, purifications where the purifying system is a twin of the purified system. In summary, all the axioms used to derive the duality are requirements about the possibility to come up with an ideal description of the world, satisfying simple requirements of purity, reversibility, and symmetry.

A natural question is whether these axioms single out quantum theory. Strictly speaking, the answer cannot be affirmative, because all our axioms are satisfied by also by the variant of quantum theory based on real Hilbert spaces [86, 87]. Hence, the actual question is whether real and complex quantum theory are the only two examples of theories satisfying the axioms. While an affirmative answer is logically possible, we do not expect it to be the case. The reason is that our axioms do not place any restriction on measurements: for example, our proof of the duality does not require one to assume an operational analogue of Naimark's theorem, stating that every measurement can be implemented as an ideal measurement at the fundamental level. Overall, in the general purification philosophy of our work, it is natural to expect that full characterization of quantum theory will require at least one requirement about the existence of a class of ideal measurements that generalize projective quantum measurements.

Naimark-type axioms for measurements has been recently put forward by one of the authors [101, 102], for the purpose of deriving bounds on quantum nonlocality and contextuality. A natural development of our work is to investigate the consequences of these axioms on the entanglement-thermodynamics duality. From such development, we expect a solution to most of the outstanding questions arising from the present paper. Among them, an important one concerns the asymptotic limit of many identical copies: in quantum theory, it is well known that asymptotically there exists a unique measure of pure bipartite entanglement [36, 37, 42]—namely, the von Neumann entropy. Under which conditions does this result hold in the general probabilistic scenario? In order to address the question, the most promising route is to add an axiom about ideal measurements, which, combined with Purification and Purity Preservation, guarantees that mixed states can be "diagonalized", that is, decomposed as random mixtures of perfectly distinguishable pure states [103]. The consequences of this diagonalization result for the entanglement-thermodynamics duality will be discussed in a forthcoming paper [104].

Another open question concerns the physical interpretation of the duality. So far, the duality has been presented as a one-to-one correspondence between two operational scenarios, one involving a single agent with limited control and the other involving two spatially separated agents performing LOCC operations on a pure state. Inspired by the paradigm of the "fundamentally pure and reversible description", one may be tempted to regard the pure-state side of the duality as more fundamental. To push this idea further, one would have to consider a completely-coherent version of the LOCC operations, where Alice's and Bob's operations are replaced by control-reversible channels [105]. Restricting the global dynamics of composite systems to these completely-coherent evolutions appears as a promising

direction in the programme of deriving effective thermodynamic features from the reversible dynamics of a composite system [69–75, 77]. While it is early to predict all the applications of the completely-coherent paradigm, our work provides the basic theoretical framework and motivation to embark in this new exploration.

Appendix A: Proof of proposition 2

Suppose that $\mathcal{U}\psi$ can be written as a coarse-graining as follows

$$\mathcal{U}\psi = \sum_i \rho_i . \qquad (A1)$$

To prove that the state $\mathcal{U}\psi$ is pure, now we show that the refinement $\{\rho_i\}$ is trivial. Indeed, by applying $\mathcal{U}^{-1}$ to both sides of Eq. (A1), we obtain

$$\psi = \sum_i \mathcal{U}^{-1}\rho_i .$$

Since ψ is pure, this implies that $\mathcal{U}^{-1}\rho_i = p_i\,\psi$ for some probability distribution $\{p_i\}$. Hence, by applying $\mathcal{U}$ on both sides, we obtain $\rho_i = p_i\mathcal{U}\psi$. This concludes the proof that $\{\rho_i\}$ is a trivial refinement of $\mathcal{U}\psi$ and, therefore, that $\mathcal{U}\psi$ is pure. The converse can be proved in the same way by applying the reversible channel $\mathcal{U}^{-1}$ to $\mathcal{U}\psi$. ∎

Appendix B: Proof of proposition 4

$1 \implies 2$. If the theory is canonical, every pure state $\psi \in \mathsf{PurSt}\,(A)$ is comparable to every pure state $\varphi \in$ $\mathsf{PurSt}\,(A)$. Suppose, for instance that ψ is more controllable than φ. Then, by proposition 3, there exists a reversible channel $\mathcal{U}$ such that $\psi = \mathcal{U}\varphi$, thus showing that the group of reversible transformations acts transitively on the set of pure states.

$2 \implies 3$. Every state ρ can be expressed as a convex combination of the form $\rho = \sum_i p_i\varphi_i$, where $\{p_i\}$ is a probability distribution allowed by the theory and φ_i are pure states. Now, suppose that ψ is a pure state. For every i, by picking a reversible channel $\mathcal{U}^{(i)}$ such that $\mathcal{U}^{(i)}\psi = \varphi_i$, one obtains the relation $\rho = \sum_i p_i\mathcal{U}^{(i)}\psi$, meaning that ψ is more controllable than ρ. Since ρ is generic, we conclude that ψ is more controllable than every state.

$3 \implies 1$. Suppose there exists a state ρ that is more controllable than every state. Specifically, ρ must be more controllable than every pure state ψ. By proposition 3, ρ must be pure and there exists a reversible transformation $\mathcal{U}$ such that $\rho = \mathcal{U}\psi$. This shows that ψ is more controllable than ρ, which, in turn is more controllable than any state. Hence ψ is more controllable than every state, and, specifically, more controllable than every pure state. Since ψ is generic, the theory is canonical. ∎

Appendix C: Proof of corollary 1

Clearly, if Ψ and Ψ' are equivalent under local reversible transformations, then they are equally entangled. To prove the converse, note that, by the duality, the marginals of Ψ and Ψ' on system A, denoted by ρ and ρ', are equally mixed. Since A is finite-dimensional, this implies $\rho' = \mathcal{U}\rho$ for some reversible transformation $\mathcal{U}$. As a consequence, Ψ' and $(\mathcal{U} \otimes \mathcal{I}_{\mathrm{B}})\,\Psi$ are two purifications of ρ'. By the essential uniqueness of purification, we then have $\Psi' = (\mathcal{U} \otimes \mathcal{V})\,\Psi$ for some reversible transformation $\mathcal{V}$. ∎

Appendix D: Purity monotones and Schur-convex functions

In classical probability theory, states are probability distributions over finite sets and reversible transformations are permutation matrices, of the form $\Pi_{mn} = \delta_{m,\pi(n)}$ where π is a permutation. According to definition 5, a state $\boldsymbol{p} = \begin{pmatrix} p_1 & \ldots & p_n \end{pmatrix}^T$ is purer than another state $\boldsymbol{p}' = \begin{pmatrix} p'_1 & \ldots & p'_n \end{pmatrix}^T$ if

$$\boldsymbol{p}' = \sum_i q_i\Pi_i\boldsymbol{p} ,$$

where $\{q_i\}$ are probabilities and $\{\Pi_i\}$ are permutation matrices. This notion is equivalent to the classical notion of majorization: in short, $\boldsymbol{p}$ is purer than $\boldsymbol{p}'$ if and only

if the vector $\boldsymbol{p}'$ is majorized by the vector $\boldsymbol{p}$. Hence, a function $P : \mathbb{R}^n \to \mathbb{R}$ is a purity monotone iff it is a Schur-convex function.

The parallel between purity monotones and Schur-convex functions continues with proposition 9. In classical probability theory, a function $P : \mathbb{R}^n \to \mathbb{R}$ is *symmetric* if $P(\boldsymbol{x}) = P(\Pi \boldsymbol{x})$ for every permutation matrix Π. A well-known result is that every convex symmetric function is Schur-convex [65]. Our proposition 9 is the operational version of this statement: every convex function $P : \mathsf{St}(A) \to \mathbb{R}$ satisfying the condition $P(\rho) = P(\mathcal{U}\rho)$ for every reversible transformation $\mathcal{U}$ is a purity monotone. The proof is elementary. Suppose that ρ is purer than ρ', namely $\rho' = \sum_i p_i \mathcal{U}_i \rho$. Then, one has

$$
\begin{aligned}
P(\rho') &\leq \sum_i p_i P(\mathcal{U}_i \rho) \\
&= \sum_i p_i P(\rho) \\
&= P(\rho) ,
\end{aligned}
$$

having used convexity in the first inequality and invariance in the first equality. In the classical case, this (trivial) proof provides a simpler proof of the well-known result for convex symmetric functions (cf. C.2 of Ref. [65]).

Appendix E: Proof of proposition 10

Let us prove the contrapositive: if a state can be erased using system C as a catalyst, then the state must be pure. Specifically, suppose that $\rho \in \mathsf{St}(A)$ can be erased, with the catalyst in the state $\gamma \in \mathsf{St}(C)$. By definition, this means that $\rho \otimes \gamma \preceq_{\mathrm{mix}} \alpha_0 \otimes \gamma$ for some pure state α_0. On the other hand, one has $\rho \succeq_{\mathrm{mix}} \alpha_0$, which implies $\rho \otimes \gamma \succeq_{\mathrm{mix}} \alpha_0 \otimes \gamma$—hence $\rho \otimes \gamma$ and $\alpha_0 \otimes \gamma$ are equally mixed. Since $A \otimes C$ is a finite system, this means that there exists a reversible transformation $\mathcal{U}$ such that

$$
\mathcal{U}(\alpha_0 \otimes \gamma) = \rho \otimes \gamma , \tag{E1}
$$

[cf. Eq. (9)]. Now, let us choose a basis for $\mathsf{St}_{\mathbb{R}}(A \otimes C)$, such that the reversible transformations are represented by orthogonal matrices. Following Ref. [64], we consider the Schatten 2-norm associated with this basis, defined as

$$
\|v\|_2 := \sqrt{\sum_{i=1}^{D_{A \otimes C}} v_i^2} ,
$$

where v is a generic element of the vector space $\mathsf{St}_{\mathbb{R}}(A \otimes C)$ and $(v_i)_{i=1}^{D_{A \otimes C}}$ are the expansion coefficients of v. With this definition, we have the relation

$$
\begin{aligned}
\|\alpha_0 \otimes \gamma\|_2 &= \|\mathcal{U}(\alpha_0 \otimes \gamma)\|_2 \\
&= \|\rho \otimes \gamma\|_2 \\
&\leq \sum_i p_i \|\alpha_i \otimes \gamma\|_2 \\
&= \|\alpha_0 \otimes \gamma\|_2 ,
\end{aligned}
$$

the first and fourth lines following from the invariance of the 2-norm under orthogonal transformations, the second line following from Eq. (E1), and the third line following from the triangular inequality, having chosen a convex decomposition of ρ as $\rho = \sum_i p_i \alpha_i$ for suitable pure states $\{\alpha_i\}$. In conclusion, we must have the equality

$$
\left\| \sum_i p_i (\alpha_i \otimes \gamma) \right\|_2 = \sum_i p_i \|\alpha_i \otimes \gamma\|_2 .
$$

In order for this to be possible, all the terms $\alpha_i \otimes \gamma$ must be proportional to one another: in other words, ρ must be pure. ∎

Appendix F: Proof of theorem 3

Let us show that the first set of axioms (Purity Preservation, Local Exchangeability, Purification, and No Entropy Sinks) implies the second (Purity Preservation, Symmetric Purification, No Entropy Sinks). To this purpose, it is sufficient to show that every state has a symmetric purification. This can be done as follows:
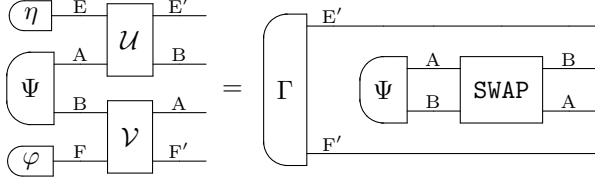
Let ρ be a state of system A, and let $\Psi \in \mathsf{PurSt}(A \otimes B)$ be one of its purifications. By Local Exchangeability there exist two channels $\mathcal{C}$ and $\mathcal{D}$ such that

Now, in a theory satisfying Purification, every channel can be realized through a reversible transformation acting on the system and on an environment, initially in a pure state and finally discarded [47]. Specifically, channel $\mathcal{C}$ can be realized as
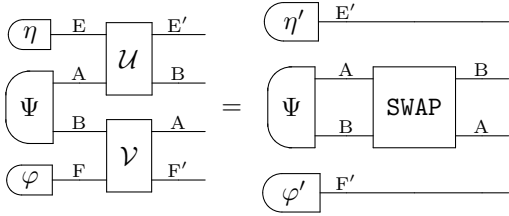

where E and E$'$ are suitable systems, $\mathcal{U}$ is a reversible transformation, and η is a pure state. Similarly, channel $\mathcal{D}$ can be realized as

Inserting the realizations of $\mathcal{C}$ and $\mathcal{D}$ in the local exchangeability condition, we obtain
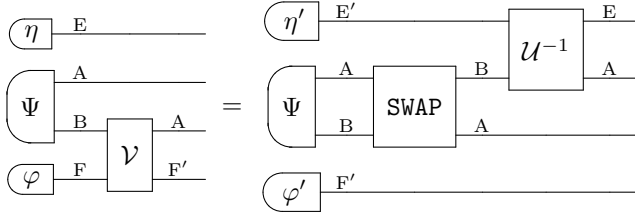
ß

[1] A. Einstein, B. Podolsky, and N. Rosen, Phys. Rev. **47**, 777 (1935).

[2] E. Schrödinger, Mathematical Proceedings of the Cambridge Philosophical Society **31**, 555 (1935).

[3] J. S. Bell, *Speakable and Unspeakable in Quantum Mechanics: Collected Papers on Quantum Philosophy*, 2nd ed. (Cambridge University Press, Cambridge, 2004).

[4] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, Phys. Rev. Lett. **23**, 880 (1969).

[5] H. Buhrman, R. Cleve, S. Massar, and R. de Wolf, Rev. Mod. Phys. **82**, 665 (2010).

[6] N. Brunner, D. Cavalcanti, S. Pironio, V. Scarani, and S. Wehner, Rev. Mod. Phys. **86**, 419 (2014).

[7] D. M. Greenberger, M. A. Horne, and A. Zeilinger, in *Bell's Theorem, Quantum Theory, and Conceptions of the Universe*, edited by M. Kafatos (Kluwer, Dordrecht, 1989) pp. 69–72.

[8] N. D. Mermin, Phys. Rev. Lett. **65**, 3373 (1990).

[9] A. Peres, Physics Letters A **151**, 107 (1990).

[10] L. Hardy, Phys. Rev. Lett. **68**, 2981 (1992).

[11] L. Hardy, Phys. Rev. Lett. **71**, 1665 (1993).

[12] S. Popescu, Nat Phys **6**, 151 (2010).

[13] S. Abramsky and L. Hardy, Phys. Rev. A **85**, 062114 (2012).

[14] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, Rev. Mod. Phys. **81**, 865 (2009).

[15] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, Phys. Rev. Lett. **70**, 1895 (1993).

[16] C. H. Bennett and S. J. Wiesner, Phys. Rev. Lett. **69**, 2881 (1992).

[17] C. H. Bennett and G. Brassard, in *Proceedings of the IEEE International Conference on Computers, Systems, and Signal Processing* (1984) pp. 175–179.

[18] A. K. Ekert, Phys. Rev. Lett. **67**, 661 (1991).

[19] D. Mayers and A. Yao, Quantum Info. Comput. **4**, 273 (2004).

[20] B. W. Reichardt, F. Unger, and U. Vazirani, arXiv:1209.0448 (2012).

[21] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters, Phys. Rev. Lett. **76**, 722 (1996).

[22] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, Phys. Rev. A **54**, 3824 (1996).

[23] H.-K. Lo and S. Popescu, Phys. Rev. A **63**, 022301 (2001).

[24] A. Peres, Phys. Rev. Lett. **77**, 1413 (1996).

[25] M. Horodecki, P. Horodecki, and R. Horodecki, Physics Letters A **223**, 1 (1996).

[26] P. Horodecki, Physics Letters A **232**, 333 (1997).

[27] M. Horodecki, P. Horodecki, and R. Horodecki, Phys. Rev. Lett. **80**, 5239 (1998).

[28] C. H. Bennett, D. P. DiVincenzo, T. Mor, P. W. Shor, J. A. Smolin, and B. M. Terhal, Phys. Rev. Lett. **82**, 5385 (1999).

[29] W. Dür, G. Vidal, and J. I. Cirac, Phys. Rev. A **62**, 062314 (2000).

[30] J. I. de Vicente, C. Spee, and B. Kraus, Phys. Rev. Lett. **111**, 110502 (2013).

[31] M. A. Nielsen, Phys. Rev. Lett. **83**, 436 (1999).

[32] A. Uhlmann, Wiss. Z. Karl-Marx-Univ. Leipzig **20**, 633 (1971).

[33] A. Uhlmann, Wiss. Z. Karl-Marx-Univ. Leipzig **21**, 421 (1972).

[34] A. Uhlmann, Wiss. Z. Karl-Marx-Univ. Leipzig **22**, 139 (1973).

[35] I. Bengtsson and K. Zyczkowski, *Geometry of Quantum States: An Introduction to Quantum Entanglement* (Cambridge University Press, Cambridge, 2006).

[36] S. Popescu and D. Rohrlich, Phys. Rev. A **56**, R3319 (1997).

[37] V. Vedral and M. B. Plenio, Phys. Rev. A **57**, 1619 (1998).

[38] M. Horodecki, J. Oppenheim, and R. Horodecki, Phys. Rev. Lett. **89**, 240403 (2002).

[39] F. G. S. L. Brandão and M. B. Plenio, Nat Phys **4**, 873 (2008).

[40] M. Horodecki, Nat Phys **4**, 833 (2008).

[41] W. Thirring, *Quantum Mathematical Physics: Atoms, Molecules and Large Systems* (Springer, Berlin Heidelberg, 2002).

[42] C. H. Bennett, H. J. Bernstein, S. Popescu, and B. Schumacher, Phys. Rev. A **53**, 2046 (1996).

[43] L. Hardy, arXiv quant-ph/0101012 (2001).

[44] J. Barrett, Phys. Rev. A **75**, 032304 (2007).

[45] G. M. D'Ariano, in *Philosophy of Quantum Information and Entanglement*, edited by A. Bokulich and G. Jaeger (Cambridge University Press, Cambridge, 2010) pp. 85–126.

[46] H. Barnum, J. Barrett, M. Leifer, and A. Wilce, Phys. Rev. Lett. **99**, 240501 (2007).

[47] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Phys. Rev. A **81**, 062348 (2010).

[48] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Phys. Rev. A **84**, 012311 (2011).

[49] H. Barnum and A. Wilce, Electronic Notes in Theoretical Computer Science **270**, 3 (2011), proceedings of the Joint 5th International Workshop on Quantum Physics and Logic and 4th Workshop on Developments in Computational Models (QPL/DCM 2008).

[50] L. Hardy, in *Deep Beauty: Understanding the Quantum World through Mathematical Innovation*, edited by H. Halvorson (Cambridge University Press, Cambridge, 2011) pp. 409–442.

[51] L. Hardy, arXiv:1104.2066 (2011).

[52] G. Chiribella, in Proceedings 11th workshop on *Quantum Physics and Logic*, Kyoto, Japan, 4-6th June 2014, Electronic Proceedings in Theoretical Computer Science, Vol. 172, edited by B. Coecke, I. Hasuo, and P. Panangaden (Open Publishing Association, 2014) pp. 1–14.

[53] H. Barnum, J. Barrett, M. Leifer, and A. Wilce, in *Proceedings of Symposia in Applied Mathematics*, Vol. 71 (2012) pp. 25–48.

[54] W. van Dam, *Nonlocality & communication complexity*, Ph.D. thesis, Faculty of Physical Sciences, University of Oxford (1999).

[55] G. Brassard, H. Buhrman, N. Linden, A. A. Méthot, A. Tapp, and F. Unger, Phys. Rev. Lett. **96**, 250401 (2006).

[56] N. Linden, S. Popescu, A. J. Short, and A. Winter, Phys. Rev. Lett. **99**, 180502 (2007).

[57] M. Pawłowski, T. Paterek, D. Kaszlikowski, V. Scarani, A. Winter, and M. Żukowski, Nature **461**, 1101 (2009).

[58] M. L. Almeida, J.-D. Bancal, N. Brunner, A. Acín, N. Gisin, and S. Pironio, Physical review letters **104**, 230404 (2010).

[59] T. Fritz, A. B. Sainz, R. Augusiak, J. B. Brask, R. Chaves, A. Leverrier, and A. Acín, Nature communications **4** (2013).

[60] B. Tsirelson, Hadronic J. Suppl. **8**, 329 (1993).

[61] S. Popescu and D. Rohrlich, Foundations of Physics **24**, 379 (1994).

[62] J. Barrett, N. Linden, S. Massar, S. Pironio, S. Popescu, and D. Roberts, Phys. Rev. A **71**, 022101 (2005).

[63] N. S. Jones and L. Masanes, Phys. Rev. A **72**, 052312 (2005).

[64] M. P. Müller and L. Masanes, New Journal of Physics **15**, 053040 (2013).

[65] A. W. Marshall, I. Olkin, and B. C. Arnold, *Inequalities: Theory of Majorization and Its Applications*, Springer Series in Statistics (Springer, New York, 2011).

[66] M. Horodecki, P. Horodecki, and J. Oppenheim, Phys. Rev. A **67**, 062104 (2003).

[67] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Entropy **14**, 1877 (2012).

[68] G. Chiribella and C. M. Scandolo, EPJ Web of Conferences **95**, 03003 (2015).

[69] S. Popescu, A. J. Short, and A. Winter, Nat Phys **2**, 754 (2006).

[70] E. Lubkin and T. Lubkin, International Journal of Theoretical Physics **32**, 933 (1993).

[71] J. Gemmer, A. Otte, and G. Mahler, Phys. Rev. Lett. **86**, 1927 (2001).

[72] S. Goldstein, J. L. Lebowitz, R. Tumulka, and N. Zanghì, Phys. Rev. Lett. **96**, 050403 (2006).

[73] J. Gemmer, M. Michel, and G. Mahler, *Quantum Thermodynamics: Emergence of Thermodynamic Behavior Within Composite Quantum Systems*, Lecture Notes in Physics, Vol. 784 (Springer Verlag, Heidelberg, 2009).

[74] M. P. Müller, D. Gross, and J. Eisert, Communications in Mathematical Physics **303**, 785 (2011).

[75] F. G. S. L. Brandão and M. Cramer, arXiv:1502.03263 (2015).

[76] M. P. Müller, O. C. Dahlsten, and V. Vedral, Communications in Mathematical Physics **316**, 441 (2012).

[77] M. P. Müller, J. Oppenheim, and O. C. O. Dahlsten, Journal of High Energy Physics **2012**, 116 (2012), 10.1007/JHEP09(2012)116.

[78] G. Vidal, Journal of Modern Optics **47**, 355 (2000).

[79] M. B. Plenio and S. S. Virmani, in *Quantum Information and Coherence*, Scottish Graduate Series, edited by E. Andersson and P. Öhberg (Springer International Publishing, 2014) pp. 173–209.

[80] R. Landauer, IBM J. Res. Dev. **5**, 183 (1961).

[81] S. Abramsky and B. Coecke, in *Handbook of Quantum Logic and Quantum Structures: Quantum Logic*, edited by K. Engesser, D. M. Gabbay, and D. Lehmann (Elsevier, 2008) pp. 261–324.

[82] B. Coecke, AIP Conference Proceedings **810**, 81 (2006).

[83] B. Coecke, Contemporary Physics **51**, 59 (2010).

[84] P. Selinger, in *New Structures for Physics*, Lecture Notes in Physics, Vol. 813, edited by B. Coecke (Springer, Berlin, Heidelberg, 2011) pp. 289–356.

[85] B. Coecke, T. Fritz, and R. W. Spekkens, arXiv:1409.5531 (2014).

[86] E. C. G. Stueckelberg, Helv. Phys. Acta **33**, 727 (1960).

[87] W. K. Wootters, in *Complexity, entropy and the physics of information*, edited by W. H. Zurek (Addison-Wesley, Boston, 1990) pp. 39–46.

[88] L. Hardy and W. K. Wootters, Foundations of Physics **42**, 454 (2012).

[89] A. J. Short and J. Barrett, New Journal of Physics **12**, 033034 (2010).

[90] B. Dakić and C. Brukner, in *Deep Beauty: Understanding the Quantum World through Mathematical Innovation*, edited by H. Halvorson (Cambridge University Press, Cambridge, 2011) pp. 365–392.

[91] L. Masanes and M. P. Müller, New Journal of Physics **13**, 063001 (2011).

[92] L. Masanes, M. P. Müller, R. Augusiak, and D. Pérez-García, Proceedings of the National Academy of Sciences **110**, 16373 (2013).

[93] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge University Press, Cambridge, 2010).

[94] H. Barnum, C. P. Gaebler, and A. Wilce, Foundations of Physics **43**, 1411 (2013).

[95] B. C. Arnold, Statist. Sci. **22**, 407 (2007).

[96] H. Barnum, J. Barrett, L. Orloff Clark, M. Leifer, R. Spekkens, N. Stepanik, A. Wilce, and R. Wilke, New Journal of Physics **12**, 033024 (2010).

[97] A. J. Short and S. Wehner, New Journal of Physics **12**, 033023 (2010).

[98] G. Kimura, K. Nuida, and H. Imai, Reports on Mathematical Physics **66**, 175 (2010).

[99] N. Brunner, M. Kaplan, A. Leverrier, and P. Skrzypczyk, New Journal of Physics **16**, 123050 (2014).

[100] M. Keyl, T. Matsui, D. Schlingemann, and R. F. Werner, Reviews in Mathematical Physics **18**, 935 (2006).

[101] G. Chiribella and X. Yuan, arXiv:1404.3348 (2014).

[102] G. Chiribella and X. Yuan, arXiv:1504.02395 (2015).

[103] G. Chiribella and C. M. Scandolo, arXiv:1506.00380 (2015).

[104] G. Chiribella and C. M. Scandolo, In preparation.

[105] A. Harrow, Phys. Rev. Lett. **92**, 097902 (2004).