The thermodynamic meaning of negative entropy

Lídia del Rio,[1] Johan Åberg,[1] Renato Renner,[1] Oscar Dahlsten,[1,2,3] and Vlatko Vedral[2,3]

[1] *Institute for Theoretical Physics, ETH Zurich, Switzerland.*
[2] *Clarendon Laboratory, University of Oxford, United Kingdom.*
[3] *Centre for Quantum Technologies, National University of Singapore, Singapore.*

Landauer's erasure principle exposes an intrinsic relation between thermodynamics and information theory: the erasure of information stored in a system, S, requires an amount of work proportional to the entropy of that system. This entropy, $H(S|O)$, depends on the information that a given observer, O, has about S, and the work necessary to erase a system may therefore vary for different observers. Here, we consider a general setting where the information held by the observer may be quantum-mechanical, and show that an amount of work proportional to $H(S|O)$ is still sufficient to erase S. Since the entropy $H(S|O)$ can now become negative, erasing a system can result in a net gain of work (and a corresponding cooling of the environment).

I. PRELIMINARIES

Statistical mechanics and information theory have a long standing and intricate relation. A famous example of this connection is Landauer's erasure principle [1], used to exorcise Maxwell's demon [2]. According to this principle, in order to perform irreversible operations on a system, like the erasure of a bit of information, we need to perform work on the system, which is dissipated as heat to the environment. The necessary amount of work is determined by our uncertainty about the system — the more we know about the system, the less it costs to 'erase' it. This result suggests that the seemingly elusive concept of 'information' is directly linked to a very concrete quantity, 'work'. Here, we analyse the relation between thermodynamics and information in a world that is fundamentally quantum mechanical.

Quantum information theory has peculiar properties that cannot be found in its classical counterpart. One example is that one's uncertainty about a system, as measured by an entropy, can become negative [3]. This motivates the following question: when our uncertainty about a system is negative, can we *gain* work by erasing the information stored in that system? Our results show that this is indeed possible; inherently non-classical aspects of quantum information theory, like negative uncertainty, are at a fundamental level part of thermodynamics.

A. Physics from an information-theoretic viewpoint

Our knowledge about the state of physical systems is usually limited, because the number of parameters that we can measure and store, as well as our precision, are finite. A typical example is a gas: we cannot keep track of the state of each particle, but only of a few macroscopic parameters, such as the volume or pressure of the gas. Despite this restricted information, it is possible to make accurate predictions about the behavior of systems using tools of statistical mechanics [4–6].

Information constraints can also result in different observers having considerably different knowledge about the same physical reality. To illustrate this *subjectivity of information*, consider an n-qubit system, S (e.g., n spin-1/2 particles). An observer, Alice, prepares the system in a known pure state. A second observer, Bob, does not know which state that is, but applies an energy measurement to the system. If S is degenerate, Bob remains ignorant about the exact state of the system.

A natural way to quantify the knowledge of these observers is to use entropy measures. The *entropy* of a system, S, given all the information available to a given observer, O, denoted by $H(S|O)$, increases with the uncertainty of the observer about the exact state of the system.[1] In the case where S is fully degenerate, the entropy of the system from the point of view of Alice is zero, $H(S|A) = 0$, as she has complete knowledge of the state of the system. On the other hand, Bob has maximal entropy, $H(S|B) = n$, because he does not know in which of the 2^n possible states the system is.[2]

This observer-dependence of entropy seems to contradict the traditional thermodynamics view, where entropy appears as a property of the system rather than of the observer. However, the two views can be reconciled by introducing a *standard observer* who has access to a well-defined set of macroscopic parameters, but whose uncertainty about the state of the system is otherwise maximal [4]. The idea is that the knowledge of this standard observer corresponds, to good approximation, to the

[1] For concreteness, one may think of the von Neumann entropy, which for a system, S, in state ρ, is defined by $H(S)_\rho := -\mathrm{Tr}(\rho \log_2 \rho)$. However, most of this section is valid for any reasonable entropy measure, and our technical statements will use smooth min- and max-entropies [7]. These are generalizations of the von Neumann entropy, and reduce to the latter for certain 'nicely behaved' distributions, e.g., in the thermodynamic limit (see Appendix B for details). The subscript in $H(S)_\rho$ can be dropped if the state is clear from the context.

[2] The entropy of S conditioned on the classical memory O, $H(S|O)$, can be defined as the expectation, taken over all states of the memory, m_O, of the entropy of ρ^m, the state of S conditioned on m_O: $H(S|O) := \mathbb{E}_m [H(S)_{\rho^m}]$.

knowledge we typically have about large systems in realistic situations: in general, we do not know microscopic details such as the spin direction of individual particles, but only parameters like the energy of a system (in the above example, it would make sense to take Bob as the standard observer). One may nevertheless ask whether the difference between the entropies $H(S|A)$ and $H(S|B)$ has any physical significance. As we shall see, this is indeed the case.

B. Quantum knowledge

The observers we described require an internal memory to store the information they have about the system S (for Alice this memory needs to be large enough to include a full description of the state of S, while Bob only stores the value of the energy). It often is implicitly assumed that this memory is *classical*. We go beyond this classical scenario and consider observers who may have access to information about S that is itself represented as the state of a *quantum* system — a quantum memory.

To illustrate the effects of a quantum memory, let us consider a third observer, Charlie, who has one. Charlie prepares each of the n particles of S such that it is maximally entangled with a corresponding qubit of his memory. Note that this quantum memory is at least as useful as the classical data held by Alice. In fact, the latter may be recovered by applying a measurement on Charlie's memory.

In order to quantify the uncertainty that Charlie has about S, we need entropy measures that account for the quantum-mechanical nature of the information he holds. In the field of quantum information, such measures are known as *conditional entropies* and generalize classical conditional entropies. The conditional von Neumann entropy can be written as a difference, $H(S|C) = H(SC) - H(C)$.[3] Here, $H(SC)$ denotes the von Neumann entropy of the joint state of the system, S, and the quantum memory, C. Since this joint state is pure, its entropy is zero. On the other hand, the reduced state of the memory, ρ_C, is fully mixed, which corresponds to the maximal entropy $H(C) = n$. We therefore find that, for Charlie, the conditional entropy is negative, $H(S|C) = -n$. Such negative entropies cannot occur for purely classical observers like Alice and Bob.

This raises the question of whether these 'negative uncertainties' have any operational meaning. The answer is yes. They can be used to quantify, for instance, the amount of entanglement needed to send a state to a receiver with side information, a task commonly referred to as 'state merging' [3]. Another example where negative conditional entropies play a crucial role was given recently in the context of Heisenberg's uncertainty principle. The principle bounds the minimum uncertainty one has about the outcome of a measurement on a system, S, chosen from two complementary observables, e.g., a spin measured in the X or Z basis.[4] This bound is, however, violated if quantum information about the initial state of S is available. It was shown that this violation can be quantified by the negativity of the entropy of S conditioned on the memory [10].[5]

In this work, we go one step further and establish a relation between a *physical* quantity (namely the work necessary to 'erase' the state of a system) and the conditional entropy. Remarkably, the validity of this relation extends to the quantum regime and, in particular, yields a direct thermodynamical interpretation of negative conditional entropies.

C. Information-work relation

In this section we illustrate Landauer's erasure principle and express it in terms of conditional entropies. The process of *erasing a system* is defined as taking the system to a pre-defined pure state, $|0\rangle$. Note that while erasing a system leads to the loss of information that could be encoded there, it may also reduce our uncertainty about the system (if we did not know the previous state of the system, now we are sure that it is $|0\rangle$).

For a concrete example of how to erase a bit, consider a spin-1/2 particle exposed to a tunable magnetic field that can be adjusted to manipulate the energy of states $|\downarrow\rangle$ and $|\uparrow\rangle$, according to a Hamiltonian like $\mathcal{H}_B = j\vec{B}\cdot\vec{s}$. Initially, the magnetic field is turned off, so the system is degenerate. We define 'erasing' as taking the spin to the pure state $|0\rangle := |\downarrow\rangle$. Let us see how two different observers could do this.

Our first observer, Alice, knows that the particle is in a pure state, for instance $|\uparrow\rangle$. In order to take the particle to $|\downarrow\rangle$, she may apply a unitary operation, in this case a NOT gate. This operation is reversible and has no energy cost.

The second observer, Bob, has no information about the initial state of the system, describing it as a fully mixed state, $\frac{1}{2}$. One strategy he can follow to erase the bit is to couple the particle to a heat bath and slowly increase the magnetic field, raising the energy of state $|\uparrow\rangle$ until its occupation decays, as shown in Fig. 1. This erasure process has an energy cost of $kT\ln 2$, where T is the temperature of the bath and k the Boltzman constant.

---

[3] If C was classical, this expression would be equivalent to $H(S|C) := \mathbb{E}_m\left[H(S)_{\rho^m}\right]$, as before.

[4] More precisely, in its formulation proposed by Deutsch [8] and Maassen and Uffink [9], the principle asserts that $H(X|O) + H(Z|O) \geq \log_2 \frac{1}{c}$, where O is any *classical* description of the initial state of S, and where $\log_2 \frac{1}{c} \geq 0$ is a measure for the non-commutativity of the observables X and Z.

[5] In the generalized form where O may be non-classical, the relation reads $H(X|O) + H(Z|O) \geq \log_2 \frac{1}{c} + H(S|O)$.

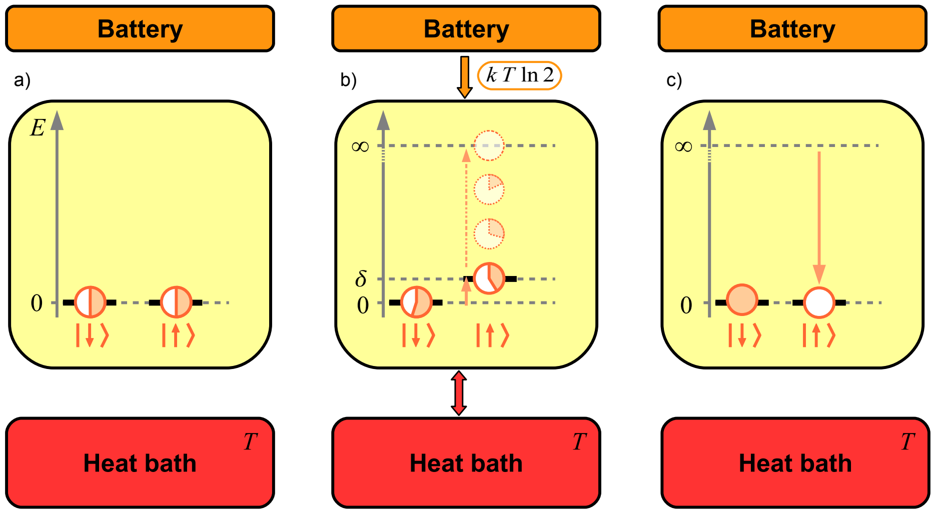

FIG. 1: Erasing a fully mixed qubit. *a*) We start from a fully mixed state in a degenerate system. The filling of each circle represents the probability, $\langle n_{\downarrow/\uparrow}\rangle$, that the system is in the respective state. *b*) We couple the system to a heat bath at temperature T and slowly raise the energy of state $|{\uparrow}\rangle$. Thermalized by the bath, the system equilibrates in a Gibbs state of temperature T. As the energy of $|{\uparrow}\rangle$ increases, it becomes less occupied, according to $\langle n_\uparrow\rangle(E) = [1 + e^{E/kT}]^{-1}$. We continue raising that level until it is empty. The total cost of this operation is $\int_0^\infty \langle n_\uparrow\rangle(E)\ dE = kT\ln 2$. *c*) Finally, we isolate the system and lower the energy of state $|{\uparrow}\rangle$. Since the state is empty, this operation is energy neutral.

More generally, in a hybrid setting where the system, S, may be quantum mechanical but the information about it is classical, the work, $W(S)$, required to erase S is given by [11]

$$W(S) = H(S)\ kT\ln 2\ . \tag{1}$$

Crucially, Eq. 1 relates work to a quantity that is, according to our discussion above, dependent on an observer. This apparent contradiction is resolved by reconsidering the meaning of $W(S)$. Note that in order to erase a system, we need to design an experimental setup that can, and in general *must*, depend on the knowledge we have about it. Hence, rather than describing $W(S)$ simply as the 'amount of work one needs to perform to erase system S', one may interpret it as the 'amount of work that an observer with memory O needs to erase S', and denote it by $W(S|O)$. For an observer with a classical memory, O_C,[6] we have in general

$$W(S|O_C) = H(S|O_C)\ kT\ln 2\ . \tag{2}$$

We emphasize that this formula does not condradict Eq. 1. Instead, it makes it explicit that the relevant quantities may depend on the knowledge of the observer and, in particular, may differ for different observers (in our example, Alice had zero entropy and consequently erased the bit at zero cost, while Bob had $H(S|B) = 1$ and had to perform work $kT\ln 2$; see also [12] for a discussion).

Our contribution is to generalize this relation to the fully quantum case. We will be able to analyse what observers with quantum memories can do to erase a system, and how much that costs them.

II. THE GENERAL RELATION BETWEEN INFORMATION AND WORK

In this section we state and explain our main result, a general relation between the work necessary to erase a system and the information one has about this system.

Several approaches have been proposed in the past to formalize the idea of a thermal process and to study erasure, work extraction and their relation to Maxwell's demon [1, 12–21]. This has spurred a rather extensive literature (for overviews see [22–25]) as well as debates (see, e.g., [11, 26–28]). Correlations and entanglement can affect erasure and work extraction, as has been noted by several authors. For instance, in [29] the system to be erased is bipartite and the observer is restricted to local operations and classical communication (LOCC); the difference between quantum and classical 'demons' is addressed in [30]; see also [31] for a discussion on 'local' and 'global' demons in the context of the thermodynamic arrow of time.

Here, we consider a setting as depicted in Fig. 2, where an observer, who has a quantum memory, O, tries to erase a system, S, using a heat bath at temperature T and performing operations on S and O (which are not restricted to LOCC). We assume that the inital Hamil-

[6] In the literature on Landauer's erasure principle the system to be erased is sometimes referred to as a 'memory'. However, for the sake of clarity we reserve the term 'memory' exclusively for the observer's memory resources.

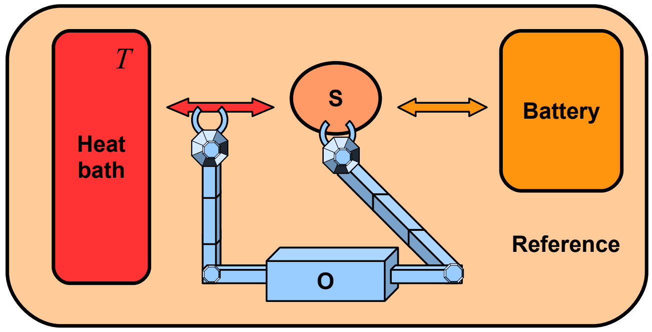

FIG. 2: Our setting: an observer, here represented by a machine with a quantum memory (O), will erase a system, S, using a heat bath at temperature T. The observer can store and withdraw energy from a battery. The rest of the universe is represented by the reference system.

tonian of S and O is fully degenerate. Details on the setting can be found in Appendix A.

Since the memory$O$ is quantum mechanical, accessing it may in general change its content. Also, there is no reason why the memory would only contain information about S; it could also carry information about other systems. Here we take a cautious position and require that those memory contents are kept intact in the erasure process. Note that this requirement is crucial, since the contents may generally be needed for other purposes, e.g., if the erasure of S is part of a larger procedure. As a simple example, suppose we erase system S, and later possibly would like to erase another system Z. If the erasure of S removed the information about Z, the subsequent erasure of Z could become unnecessarily costly.

In order to specify this memory preservation condition on a formal level, it is convenient to introduce a 'reference system' R, which models all systems other than S that the memory can have information about. To guarantee that the information about R is unaltered, we assume that the joint state of the memory and the reference, ρ_{OR}, is preserved by the erasure process and that system R is not touched.

A. A special case

The general idea of what an observer with a quantum memory can do to erase a system and *gain* work in the process can be illustrated with a simple example. Consider a single qubit system S, and an observer, Charlie, who has a memory formed by two qubits, $C = C_1 \otimes C_2$. The first qubit is maximally entangled with S, in state $|C_1S\rangle$, while the second is maximally entangled with a qubit of the reference system, R, in state $|C_2R\rangle$. Charlie will try to erase S but keep his memory about R intact, preserving the joint state $\rho_{CR} = \frac{\mathbb{1}_{C_1}}{2} \otimes |C_2R\rangle\langle C_2R|$.[7]

In a first step, Charlie uses the two-qubit pure state $|C_1S\rangle$ and a heat bath at temperature T to extract work $2kT \ln 2$, as described in Fig. 3. The system formed by C_1 and S is left in a fully mixed state. In particular, the reduced state of C_1 is fully mixed, which implies that the joint state of the memory and the reference is still ρ_{CR}. Charlie then erases the fully mixed qubit S, like Bob did in Section I C, performing work $kT \ln 2$. The net work gain of the whole procedure is $kT \ln 2$. Note that if Charlie had not preserved his memory and later wanted to erase R, he would have to perform unecessary work.

This case illustrates how the relation between entropy and the work necessary to erase a system applies in a quantum scenario: Charlie had negative conditional entropy about S, $H(S|C) = -1$, which resulted in negative work cost for erasure, $W(S|C) = -kT \ln 2$.

Naturally, the energy 'gained' in this process comes from the heat bath. As Charlie not only extracted work but also took S to a pure state, while leaving ρ_{CR} intact, one may at first sight fear that he has violated the second law of thermodynamics. This is, however, not the case, since those gains are balanced by the reduction in correlations between S and C. In fact, the entropy of the global state, $H(CSR)$, increased, and erasing S made Charlie lose all the entanglement between his memory and S. His knowledge about the final state of S is only classical — it can be expressed by a non-negative conditional entropy, $H(S|C) = 0$. This prevents him from gaining more work if he erases S again, using this process in a perpetual motion scheme. The same observation also explains why a negative cost of erasure would not enable Maxwell's demon to violate the second law.

[7] The reduced state of C_1 is fully mixed, because $|C_1S\rangle$ is maximally entangled.

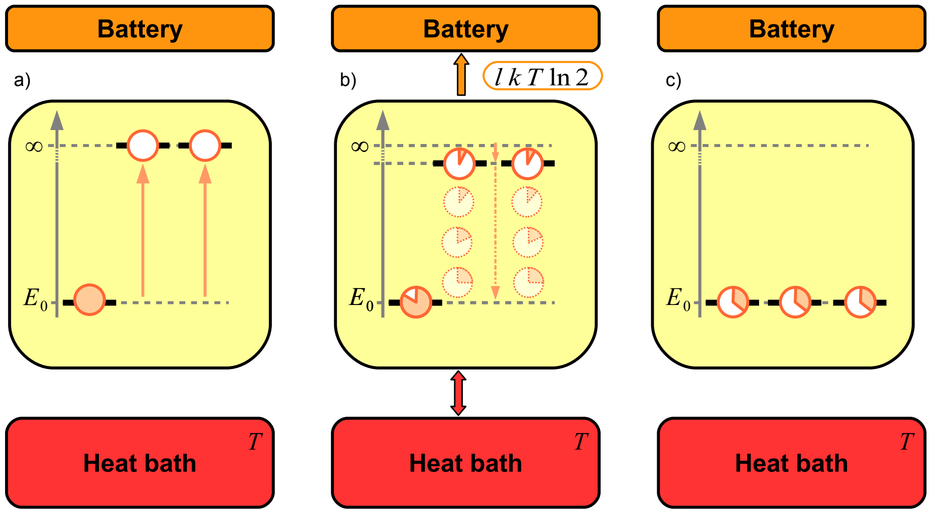

FIG. 3: Extracting work from a ℓ-qubit system in a pure state. This process can be seen as the reverse of erasure (Fig. 1). *a*) Only one state is occupied, at energy E_0; the energy of the empty levels is raised to a very high value at zero cost. *b*) We couple the system to the bath and slowly decrease the energy of the empty states. These will become gradually populated according to the Gibbs distribution. Lowering the partially occupied states results in energy gain of $\ell\ kT\ln 2$ in total. This energy is stored in the battery. *c*) In the end of the procedure, the system is degenerate and fully mixed.

B. Single-shot erasure

In general, the work required to erase a system is a random variable, i.e., the cost of erasure may fluctuate each time it is performed. Here we characterize a *single instance* of erasure with a probabilistic statement, and in Section II C we will consider the average work cost of erasure in a thermodynamic limit.

Theorem 1 guarantees that the cost of erasing a system does not exceed a bound given in terms of the entropy of S conditioned on O, except with a small probability.

Theorem 1. *There exists a process to erase a system S, conditioned on a memory, O, and acting at temperature T, whose work cost satisfies*

$$W(S|O) \leq [H^\varepsilon_{\max}(S|O) + \Delta]\ k\ T\ln 2, \tag{3}$$

except with probability less than $\delta = \sqrt{2^{-\frac{\Delta}{2}} + 12\varepsilon}$, $\forall \delta, \varepsilon > 0$.

The quantity $H^\varepsilon_{\max}(S|O)$ denotes the ε-smooth max-entropy of system S conditioned on the quantum memory O, a single-shot generalization of the von Neumann entropy [7]. In particular, as we shall see, this quantity reduces to the von Neumann entropy in a thermodynamic limit (we refer to Appendix B for definition and properties of smooth entropies).

The term Δ can be chosen to be small, and in the limit of large systems could be negleted. For instance, to allow a maximum probability of failure of only $\delta = 3\%$, one pays a price of approximately $20\ kT\ln 2$ in the work consumption of the process (in addiction to the one dictated by the entropy).

Theorem 1 implies that an observer with a quantum memory entangled with S (i.e., with $H^\varepsilon_{\max}(S|O) < 0$) can erase the system with negative work cost, actually *extracting* work in the process. Note that this is more general than the example of Section II A, where S was, conveniently, maximally entangled with a part of the memory: Theorem 1 implies that observers can make full use of the correlations between S and O, even if those are not present in the neat form of maximally entangled qubits.

As a byproduct of the proof of Theorem 1 we find an analogous result for work extraction. The goal of this process is to extract work from a system, S, under the assumption that the memory is kept intact(while the final state of S is arbitrary).

Corollary 1. *Given an n-qubit system S and a memory O, there exists a work extraction process acting at temperature T, such that the extracted work satisfies*

$$W_e(S|O) \geq [n - H^\varepsilon_{\max}(S|O) - \Delta]\ k\ T\ln 2,$$

except with a probability of at most $\delta = \sqrt{2^{-\frac{\Delta}{2}} + 12\varepsilon}$, $\forall \delta, \varepsilon > 0$.

C. Thermodynamic limit

We typically expect thermal fluctuations to disappear in macroscopic systems. Theoretically, this is usually handled by taking a thermodynamic limit, where we in some sense increase the size of the system such that fluctuations are averaged away. In order to define a thermodynamic limit in our scenario, we imagine to perform the erasure on a large collection of independent systems.

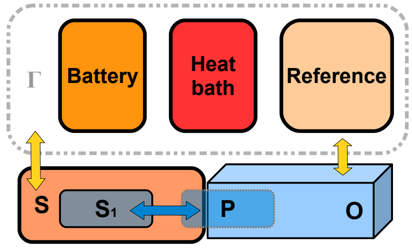

FIG. 4: Information compression, as used in the first step of our proof: a subsystem S_1 is decoupled from Γ. The size of S_1 decreases with the strength of the correlations between S and Γ, and therefore increases with correlations between S and the memory, O (see Appendix B). Since the global state is pure, S_1 is purified by a system P of equal size that belongs to the remaining systems, S and O. The state of $S_1 \otimes P$ is fully entangled. The arrows symbolize correlations between the different systems.

We define the *work cost rate* of an erasure process as the average work cost of the process in this limit,

$$\bar{w}(S|O) = \lim_{n\to\infty} \frac{1}{n} W(S^{\otimes n}|O^{\otimes n}).$$

This quantity can be evaluated if we perform the erasure of many copies of a system. To understand the implications of our claim in such a situation, we use a well-known statement from information theory, the Asymptotic Equipartition Property (AEP) [32]. The quantum version of this result essentially asserts that, for n-partite states that consitst of many identical copies of the same single subsystem state, the smooth max-entropy converges towards the von Neumann entropy (see Appendix B).

The work cost rate can now be evaluated using Theorem 1 combined with AEP, leading to the following result.

Corollary 2. *There exists a process to erase a system S, conditioned on a memory, O, and acting at temperature T, with work cost rate*

$$\bar{w}(S|O) \leq H(S|O)\ kT \ln 2.$$

III. OUTLINE OF THE PROOF

We prove our result by providing an explicit process that satisfies the bound of Theorem 1. We assume (without loss of generality) that S is an n-qubit system. The erasure process consists of three main steps:

1. We manipulate S in order to *compress the correlations* between the memory and S into a pure state of a subsystem of $S \otimes O$ that has approximately $n - H_{\max}(S|O)$ qubits. This state is maximally entangled between two subsystems of $S \otimes O$, like in the case of Charlie, from the example of Section II A.

2. We use that pure state to extract roughly $[n - H_{\max}(S|O)]\ kT \ln 2$ work ($kT \ln 2$ per qubit).

3. Finally, we erase system S, performing work $n\ kT \ln 2$ (again, $kT \ln 2$ per qubit).

We now describe these three steps in more detail, referring to technical proofs that can be found in the appendices when necessary.

In the first step, we show, using decoupling results [3, 33] that, after an appropriate transformation, the first $\ell/2$ qubits of S are almost (up to a probability determined by δ) uncorrelated to the collection, Γ, of systems outside S and O (see Appendix C 1 for details), with

$$\ell \geq n - H_{\max}^{\varepsilon}(S|O) + 2\log_2(\delta^2 - 12\varepsilon). \tag{4}$$

These $\ell/2$ qubits form the subsystem S_1. As illustrated in Fig. 4, the fact that S_1 is decoupled from Γ implies that there is an $(\ell/2)$-qubit subsystem, P, of $S \otimes O$ such that the state of $S_1 \otimes P$ is δ-close to a pure, fully entangled state (details in Appendix C 2).

In a second step, the observer extract work $\ell\ kT \ln 2$ from the state of $S_1 \otimes P$ using a heat bath at temperature T, as described in Fig. 3 and Appendix D. The system $S_1 \otimes P$ is left in a fully mixed state. Note that the state used was maximally entangled, so the reduced states of S_1 and P were already fully mixed before this step. In particular, the part of the memory involved in work extraction is not changed. The observer did not touch the memory before this second step and will not use it again, which implies that the reduced state of memory and reference, ρ_{OR}, is preserved by the erasure process. It is shown in Appendix D that the probability of failure of work extraction is upper bounded by δ. The work extraction process of Corollary 1 ends here.

In the last step of the erasure process, the observer uses energy from the battery to erase system S, as described in Fig. 1, performing work $n\ kT \ln 2$. The work balance of whole process is $(\ell - n)\ kT \ln 2$. The logarithmic term in Eq. 4 is usually negative, because we choose δ and ε to be small, so we can write the work consumption of the process as $W(S|O) \leq [H_{\max}^{\varepsilon}(S|O) + \Delta] kT \ln 2$.

IV. CONCLUSIONS

We have shown that conditional entropies, as measures of the uncertainty that an observer has about a system, have a direct physical significance in statistical mechanics. These results complement previous findings that conditional entropies have an operational meaning within information theory [3, 10]. More specifically, we have introduced an erasure process that uses the quantum information that an observer has about a system to erase the latter. The work cost of this erasure process depends on conditional entropies, and a curious implication of our findings is that negative entropies correspond to a negative work cost of erasure. We have also seen that an observer with a quantum memory can extract twice as much work from a system as one with a classical memory.

The strengthened connection between information theory and statistical mechanics may allow us to interchange concepts between the two areas. An example is the proof of our results, as an essential part is played by *decoupling*, which has shown to be a very powerful information theoretic primitive [33, 34]. The following observation suggests that we may also transfer ideas in the other direction. Intuitively, it appears rather clear that observers cannot extract more work by locally processing data in their memory. Combined with our bounds for work extraction, this gives an alternative 'thermodynamic' derivation, as well as interpretation, of the data processing inequality (also known as strong subadditivity) which, in information theory, is a crucial and nontrivial result.

Our work can be related to *discord*, a quantity originally introduced in the context of open systems theory and decoherence [35, 36], and also intensively studied in quantum information theory [37]. Discord quantifies the difference between the uncertainty about a system, S, for an observer that possesses a quantum memory, O_Q, and one that has only a classical memory, O_C, obtained by performing a measurement on O_Q, $\delta(S|O) = H(S|O_Q) - H(S|O_C)$. Similarly to [30], our results suggest that $\delta(S|O)kT \ln 2$ can be interpreted as the difference between the work cost of an erasure procedure that makes full use of the quantum nature of the memory and a process that is restricted to the classical properties of that memory. In fact, since our relation between work and entropy is valid for a single instance of an erasure process, one may consider a generalized definition of discord based on the smooth max-entropy, which retains its operational meaning in the single-shot case.

A. Applications

Our result can also have implications on the fundamental limits of computation. Today, one of the major challenges to the miniaturization of circuitry for high-performance computing lies in the heat generation. With the increased compactification, the heat generated per square unit of circuitry is rapidly becoming difficult to handle. Although our investigation certainly cannot help with the practical issues, it might nevertheless be extended to a theory that provides the ultimate bounds on dissipation. As is well known, computation per se can be made reversible [38, 39]. However, this comes at the expense of keeping extra information about the computation in a memory. Whenever we wish to erase a part of this memory, Landauer's erasure principle dictates that this unavoidably comes at the cost of generationg heat.

A very common scenario in a computation is that we wish to erase a part of a memory, but keep the rest of the memory intact. How much work do we need to dissipate in order to do this? The naive answer would be that the cost is given by the entropy solely of the part of the memory to be erased. However, our analysis shows that one can do better, namely that the required work is upper bounded by a condicional entropy, which in general can be much smaller.

Note that our result requires almost perfect control of the quantum systems involved, and one may wonder why we should consider such a theoretical idealization. As an analogue one can think of the Carnot cycle. Although the ideal performance of the Carnot engine in many cases can be a practically unattainable ideal limit, it nevertheless provides the theoretical foundation in terms of which the performance of heat engines can be gauged. Reversible computation together with the erasure principle provides a similar ideal limit for minimally heat generating computation.

Acknowledgments

We would like to thank Roger Colbeck, Sania Jevtic, Terry Rudolph, Tony Short and Will Matthews for discussions that helped shape the final form of this note. We acknowledge support from the Swiss National Science Foundation (grant no. 200021-119868) and Fundação para a Ciência e Tecnologia (grant no. SFRH/BD/43263/2008).

Appendix A: Formal setting

In this appendix we formalize the setting and the conditions for an erasure process that we use to derive Theorem 1.

Setting: our setting consists of a system S, a quantum memory, O, a heat bath at temperature T, a battery and a reference system, R (Fig. 2), so that the initial global state is pure, and the composed system $S \otimes O$ is fully degenerate.

Allowed actions: the following physical processes on any subsystem, X, of $S \otimes O$ are allowed: unitary transformations on X; manipulation of the energy levels of X; coupling between X and the heat bath or battery. One may not perform any operations on the reference system.

Erasure process: in the setting described, a successful erasure process is one that erases system S and preserves the joint state of the memory and the reference, ρ_{OR}. A system is said to be *erased* when it is in a pre-defined pure state. The *work cost* of the process is defined as the difference between the initial and final charge of the battery.

Altering the energy of a state from E_0 to $E_0 + \Delta E$ has an average energy cost of $\langle n \rangle \Delta E$, where $\langle n \rangle$ is the probability that the system is in that state. This energy can be withdrawn from a battery, modeled as follows.

Battery. A battery is a system characterized by an energy value, E, called *charge*, and the following operations:

- *Withdrawing energy (performing work).* If performing an operation on a system requires energy ΔE, *coupling* between the system and the battery is modeled by performing that operation and decreasing the charge of the battery by ΔE.

- *Storing energy (extracting work).* Conversely, if an operation on a system has a negative energy cost ΔE, coupling the battery to the system and performing the operation results in an increment of ΔE of the charge of the battery.

Heat bath. We assume that the heat bath is large enough to thermalize a system like S without altering its own temperature. We model contact between a system and the heat bath by replacing the state of the system with a thermal Gibbs state of temperature T. Physically, this corresponds to letting the system be in contact with heat bath for long enough to thermalize. This condition does not imply that the state of the heat bath does not change — it does, losing or gaining the energy required to thermalize the system, but not enough to affect the temperature of the bath.

Appendix B: Smooth entropies

The main result, Theorem 1, relies on the smooth max-entropy, $H^{\varepsilon}_{\max}$, as a measure to quantify uncertainty [7]. Smooth entropies have, so far, mainly been used in information theory, where they proved to be the relevant quantities to characterize information-processing tasks such as randomness or entanglement distillation, channel coding, data compression, or key distribution.

The formulation of the entropy-work relation in terms of the smooth max-entropy — rather than the more standard von Neumann entropy — has the advantage that the relation is valid independently of the structure of the underlying quantum states. A work-entropy relation involving the von Neumann entropy (Corollary 2) is obtained from this general result by introducing appropriate assumptions, as explained below.

In the following, we briefly review the definition of smooth entropies and show how they are related to the von Neumann entropy. For a more detailed discussion of smooth entropies, their properties, and their information-theoretic significance, we refer to [7, 40–42].

1. Definition and properties

Let $\rho = \rho_{SO}$ be the state of a bipartite system, consisting of subsystems S and O. The *ε-smooth max-entropy of S conditioned on O* can be expressed in terms of the fidelity,[8] F, as

$$H^{\varepsilon}_{\max}(S|O)_{\rho} := \inf_{\rho'_{SO}} \sup_{\sigma_O} \log_2 F(\rho'_{SO}, \mathbb{1}_S \otimes \sigma_O)2 \ .$$

The supremum ranges over all density operators σ_O on O. The infimum is taken over all (subnormalized) density operators ρ'_{SO} that are ε-close[9] to ρ_{SO}, where $\varepsilon \geq 0$ is the *smoothness parameter*, which is usually chosen to be small but nonzero.

The proof of Theorem 1 also involves the smooth min-entropy, which can be seen as the *dual* of the smooth max-entropy, in the following sense. Consider a purification $\rho_{SO\Gamma}$ of the given bipartite state ρ_{SO}, with a purifying system Γ. The *ε-smooth min-entropy of S conditioned on* Γ then corresponds to the negative smooth max-entropy conditioned on O [41, 43],

$$H^{\varepsilon}_{\min}(S|\Gamma)_{\rho} = -H^{\varepsilon}_{\max}(S|O)_{\rho} \ . \tag{B1}$$

Smooth entropies have properties analogous to those of the von Neumann entropy. For example, for $\varepsilon \to 0$, both $H^{\varepsilon}_{\min}(S|O)_{\rho}$ and $H^{\varepsilon}_{\min}(S|O)_{\rho}$ are 0 if the reduced state on S is pure, 1 for a qubit S that is fully mixed and uncorrelated to O, and -1 for a qubit S that is maximally entangled with O. Furthermore, they satisfy a data-processing inequality. It asserts that the entropy of S conditioned on O can only increase if information is processed locally at O. Formally,

$$H^{\varepsilon}_{\max}(S|O')_{\bar{\rho}} \geq H^{\varepsilon}_{\max}(S|O)_{\rho},$$

where $\bar{\rho} = \bar{\rho}_{SO'}$ is the state obtained from ρ_{SO} when a completely positive map $\mathcal{M}$ is applied on system O.

2. Specialization to the von Neumann entropy

For a bipartite quantum state ρ_{SO}, the *von Neumann entropy of S conditioned on O* is defined by $H(S|O)_{\rho} = H(\rho_{SO}) - H(\rho_O)$, where $H(\sigma)$ denotes the usual (non-conditional) von Neumann entropy of σ, i.e., $H(\sigma) = -\mathrm{Tr}(\sigma \log_2 \sigma)$. The conditional von Neumann entropy is always bounded by the smooth min- and max-entropies,

$$\begin{aligned}\lim_{\varepsilon \to 0} H^{\varepsilon}_{\min}(S|O)_{\rho} &\leq H(S|O)_{\rho} \\ &\leq \lim_{\varepsilon \to 0} H^{\varepsilon}_{\max}(S|O)_{\rho} \ .\end{aligned} \tag{B2}$$

[8] Note that the fidelity can be defined for arbitrary (not necessarily normalized) positive operators, R and S, by $F(R,S) := \|\sqrt{R}\sqrt{S}\|_1$, where $\|\cdot\|_1$ is the L_1-norm.

[9] Closeness is measured in terms of the *purified distance* [43].

In particular, if the smooth min- and max-entropies coincide, they are automatically equal to the von Neumann entropy. Hence, under this condition, the smooth max-entropy occurring in Theorem 1 can be replaced by the von Neumann entropy.

A typical situation where Eq. B2 holds (approximately) with equality is that of a large n-partite system with weakly correlated parts. In the limit when the correlations disappear, the state of the system is *independent and identically distributed* (i.i.d.), i.e., of the form $\sigma^{\otimes n}$. Such states are common in information theory and physics — they arise, for instance, naturally for systems with sufficiently high symmetries (e.g., when a system is invariant under permutations of its n parts [44]). One can show that the smooth min- and max-entropies converge for states of the form $\rho_{S^nO^n} = {\sigma_{SO}}^{\otimes n}$ [32]. Hence, by virtue of Eq. B2, and using the fact that the von Neumann entropy is additive, one has, for any $\varepsilon > 0$,

$$\begin{aligned}\lim_{n\to\infty}\frac{1}{n}&H^{\varepsilon}_{\max}(S^n|O^n)_{\sigma^{\otimes n}}\\ &=\lim_{n\to\infty}\frac{1}{n}H^{\varepsilon}_{\min}(S^n|O^n)_{\sigma^{\otimes n}}\\ &=H(S|O)_{\sigma}\ . \qquad \text{(B3)}\end{aligned}$$

In other words, for i.i.d. states, the work-entropy relation of Theorem 1 asymptotically also holds for the von Neumann entropy.

We note that Eq. B3 can be seen as a reformulation of the *Asymptotic Equipartition Property*, which plays a crucial role in the area of information theory. There, operational quantities (such as the compression rate of a random source or the amount of randomness that can be distilled from a given source) are usually related to either the smooth min- or the smooth max-entropy. The widespread use of the von Neumann entropy in (textbook) information theory is therefore mainly a consequence of the fact that one typically considers i.i.d. situations, such that Eq. B2 holds with equality.

Appendix C: Information Compression

Here we address information compression, used in the first step of the erasure process; in particular, we prove the bound of Eq. 4, of Section *III*.

Information compression uses correlations between two systems, S and O, as measured by an entropy measure, to create a pure state in a subsystem of $S\otimes O$, using only local reversible transformations on S. In this result, we consider a global system $S\otimes O\otimes\Gamma$. In the context of our work, S is the system the observer is trying to erase, O the memory of the observer, and Γ is formed by the battery, the heat bath and the reference system.

Theorem 2. *Given a system $\Omega = S\otimes O\otimes\Gamma$ in a pure state, where S is an n-qubit system, it is possible to create an ℓ-qubit state of a subsystem of $S\otimes O$, with*

$$\ell \geq n - H^{\varepsilon}_{\max}(S|O) + 2\log_2(\delta^2 - 12\varepsilon),$$

that is δ-close to a pure state, applying a local unitary transformation on S.

The last term is usually small. For instance, for $\delta = 0.003$ and $\varepsilon = 10^{-6}$, we have $2\log_2(\delta^2 - 12\varepsilon) \approx -20$. If the system S is large (say ≈ 1000 qubits), this logarithmic term can be negleted.

We will see later that the erasure process fails with maximum probability δ. This means that allowing a probability of failure of 3% has a cost of 10 qubits in the size of S_1, and results in an increase of $20kT\ln 2$ in the work consumption of the erasure process (see Section III).

The proof of Theorem 2 consists of two steps: first we will *decouple* a subsystem $S_1 \subseteq S$, of $\ell/2$-qubits, from Γ. Then we will see that, since the global state is pure, S_1 is purified by a subsystem of $S\otimes O$ of the same dimension. The pure state created has a total of ℓ qubits.

1. Decoupling

In this first step, we show that it is in general possible to identify a subsystem of S that can be decoupled from Γ, according to the following definition.

Definition 1 (Decoupling)**.** A system, X, is δ'-decoupled from another system, Y, if their joint state is δ'-close to a product state,

$$\delta\left(\rho_{XY}, \frac{\mathbb{1}_X}{|X|}\otimes\rho_Y\right) \leq \delta$$

where $\delta(\rho,\sigma) = \frac{1}{2}\|\rho-\sigma\|_1$ is the trace distance between two states.

Lemma 1 will show that the size of the decoupled system depends on the correlations between S and O, as measured by an entropy measure, the smooth conditional max-entropy, $H^{\varepsilon}_{\max}(S|O)$. This result uses the procedure of decoupling, first introduced by [3] and generalized by [33].

Lemma 1. *Given a system $\Omega = S\otimes O\otimes\Gamma$ in a pure state, where S is an n-qubit system, it is possible to δ'-decouple an m-qubit subsystem of S, S_1, from Γ. The maximum size of S_1 is given by*

$$m \geq \frac{n - H^{\varepsilon}_{\max}(S|O)}{2} + \log_2(2\delta' - 12\varepsilon).$$

Proof. The decoupling results [33, Cor. 6.2, p. 57] [10] implies that the average distance between the state actually obtained after applying a unitary on S and the desired,

[10] Here formulated in terms of the smooth min-entropy.

decoupled state, is given by

$$\int_{U_S} \delta\left(\mathrm{Tr}_{S_2}\left([U_S \otimes \mathbb{1}_\Gamma]\rho^0_{S\Gamma}\right), \frac{\mathbb{1}_{S_1}}{2^m} \otimes \rho^0_\Gamma\right) dU_S \leq 2^{-\frac{1}{2}(n-2m+2)} 2^{-\frac{1}{2}H^\varepsilon_{\min}(S|\Gamma)_{\rho^0}} + 6\varepsilon. \tag{C1}$$

Here, the integral is taken over all unitary operations on system S, and $H^\varepsilon_{\min}(S|\Gamma)_{\rho^0}$ is the smooth conditional min-entropy of S, given the information that Γ may provide about that system, before applying U_S. Since the bound of Eq. C1 applies to the average over all unitaries, there is at least one fixed unitary, U_S, that respects it. For an upper bound of δ' on the distance between the desired and the obtained states, we have

$$m = \frac{n + H^\varepsilon_{\min}(S|\Gamma)}{2} + \log_2(2\delta' - 12\varepsilon). \tag{C2}$$

The global state is pure, so one may use the duality relation between entropy measures, introduced in Eq. B1 of Appendix B, $H^\varepsilon_{\min}(S|\Gamma)_{\rho^0} = -H^\varepsilon_{\max}(S|O)_{\rho^0}$, where the latter is the smooth conditional max-entropy of system S given the memory. Inserting this to Eq. C2, we obtain

$$m = \frac{n - H^\varepsilon_{\max}(S|O)}{2} + \log_2(2\delta' - 12\varepsilon).$$

□

It can be proved that the bound of Lemma 1 is optimal, i.e., that there is no unitary U_S that allows us to decouple a system with more than m qubits from Γ [45].

2. Purification

To complete the proof of Theorem 2, it remains to show that, given an $\frac{\ell}{2}$-qubit system S_1 decoupled from Γ, it is possible to find an ℓ-qubit pure state in a subsystem of $S \otimes O$. Note that the global state of $S \otimes O \otimes \Gamma$ is still in a pure state, for we have only applied a local unitary transformation on S.

Lemma 2. *Consider a system $\Omega = (S_1 \otimes S_2) \otimes O \otimes \Gamma$ in a pure state, such that the m-qubit system S_1 is δ'-decoupled from Γ, in a fully mixed state.*

It is possible to find an m-qubit subsystem P of $S_2 \otimes O$ that purifies the state of S_1 such that the joint state of $S_1 \otimes P$ is $\sqrt{2\delta}$-close to a fully entangled state.

Proof. In a first step we assume that the state of S_1 and Γ if fully decoupled. We can expand it as

$$\frac{\mathbb{1}_{S_1}}{2^m} \otimes \rho_\Gamma = 2^{-m} \sum_k |k\rangle\langle k|_{S_1} \otimes \sum_i \lambda_i \; |i\rangle\langle i|_\Gamma.$$

We can find systems A_1 and A_2 that purify ρ_{S_1} and ρ_Γ. The composite system $A_1 \otimes A_2$ purifies $\rho_{S_1} \otimes \rho_\Gamma$,

$$|\phi\rangle = |\phi'\rangle_{S_1 A_1} \otimes |\phi''\rangle_{\Gamma A_2} = 2^{-\frac{m}{2}} \sum_k |k\rangle_{S_1} |k\rangle_{A_1} \otimes \sum_i \sqrt{\lambda_i} \; |i\rangle_\Gamma |i\rangle_{A_2}.$$

The statement for $\delta' = 0$ follows now from the fact that any two purifications of the same state are related by a unitary transformation on the purifying system. In particular, P is given as the image of A_1 under this unitary. The claim for strictly positive δ' follows similarly, using Uhlmann's theorem and properties of the trace distance [43, Lem. 6].

□

Appendix D: Work extraction

In this appendix we introduce in detail a process that allows us to extract energy from a heat bath and store it in a battery, using a pure state of a system X, as introduced in Fig. 3.

Theorem 3. *Given an ℓ-qubit subsystem of $S \otimes O$, X, in a pure state, a heat bath at temperature T, and a battery, it is possible to extract exactly $\ell kT \ln 2$ work. The system is left in a fully mixed state.*

Proof. Let E_0 be the energy of the inital state of X, $|\phi_0\rangle$, for a basis $\{|\phi_i\rangle\}_i, i = 0, 1, \ldots, N$. We start by lifting the energy of all unoccupied states $\{|\phi_1\rangle, \ldots, |\phi_N\rangle\}$ to a high value, E_1. This can be done with no energy cost, because those states are empty (Fig. 3 *a*)).

Now we couple X to the heat bath and let it thermalize; X is taken to a Gibbs state of temperature T. The probability that X is in each of the states $\{|\phi_1\rangle, \ldots, |\phi_N\rangle\}$ is given by $\langle n\rangle = \left[N + e^{\beta(E_1 - E_0)}\right]^{-1}$, where $\beta = (kT)^{-1}$. In total, the probability that the system is in one of the levels raised is $N\langle n\rangle = \left[1 + e^{\beta(E_1 - E_0)}/N\right]^{-1}$.

We then couple X to the battery and lower the energy of levels $\{|\phi_1\rangle, \ldots, |\phi_N\rangle\}$ by a small amount Δ. Since those states were partially occupied, this operation gives us a small amount of energy, $N\langle n\rangle\Delta$, that is stored in the battery (Fig. 3 *b*)).

We wait for the system to thermalize again. Because levels $\{|\phi_1\rangle, \ldots, |\phi_N\rangle\}$ have slightly lower energy than before, they will become a little more populated, so the machine can extract a little more energy when it decreases the energy of the levels by another Δ. The proccess is repeated until the energy of states $\{|\phi_1\rangle, \ldots, |\phi_N\rangle\}$ is lowered to E_0. At this point all $\{|\phi_i\rangle\}_i$ are degenerate again and the state of X is fully mixed (Fig. 3 c).

In the quasistatic limit of $\Delta \to 0$ and $E_1 \to \infty$, this proccess allows us to extract a total amount of work of

$$\lim_{E_1 \to \infty} \int_{E_0}^{E_1} \frac{1}{1 + \frac{\beta(E - E_0)}{N}} dE = \frac{\ln(N+1)}{\beta} = \log|X| \; kT \ln 2. \tag{D1}$$

□

The process described in Theorem 3 takes a system from a pure to a fully mixed state, extracting some work in the process. By inverting the process (Fig. 1), one can bring a system from a fully mixed to a pure state — in other words, *erase* the system.

Corollary 3. *To erase an ℓ-qubit system initially in a fully mixed state, using a heat bath at temperature T, it is sufficient to perform work $\ell kT \ln 2$.*

When compressing information between the system and the memory, we allowed the state created to be at most δ-distant from a pure state (Appendix C). The following lemma shows how that affects the probability of failure of the work extraction procedure.

Lemma 3. *If the process described in Theorem 3 is applied to a state δ-close to a pure state, it succeeds with probability at least $1-\delta$.*

Proof. The probability that two states, ρ and σ, of the same system can be distinguished in a one-shot approach using a physical process, such as a measurement after a reversible evolution, is given by $\Pr_{\max}(\rho,\sigma) = \frac{1}{2}[1+\delta(\rho,\sigma)]$, where $\delta(\rho,\sigma)$ is the trace distance between those states.

An example of a process to distinguish two states is the work extraction process described in Theorem 3. If the process is applied to the expected pure state, σ, the probability of error is zero and the quantity of work extracted is $\ell\ kT \ln 2$. We denote the probability of failure of the work extraction process for an arbitrary state, ρ, by p_ρ.

If we are given one of the two states, σ and ρ, at random, apply the work extraction process and obtain less than $\ell\ kT\ln 2$, we know that the state was ρ. This happens with probability $p_\rho/2$. In $(1-p_\rho)/2$ of the cases, we are given ρ and extract exactly $\ell\ kT \ln 2$, and with probability $1/2$ we had σ, extracting the same work, so our best guess if we obtain work $\ell\ kT \ln 2$ is to say we had state σ. In total, we will be right with probability $\frac{1}{2}[1+p_\rho]$.

This guessing probability is upper bounded by $\Pr_{\max}(\rho,\sigma)$, so $p_\rho \leq \delta(\sigma,\rho)$. Since we imposed a maximum distance δ between the pure state σ and ρ, the probability of failure of the process is at most δ.

□

[1] Landauer, R. Dissipation and heat generation in the computing process. *IBM Journal of Research and Development* **5**, 148–156 (1961).

[2] Bennett, C. H. The thermodynamics of computation — a review. *International Journal of Theoretical Physics* **12**, 905 (1982).

[3] Horodecki, M., Oppenheim, J. & Winter, A. Partial quantum information. *Nature* **436**, 673–676 (2005).

[4] Jaynes, E. T. Information theory and statistical mechanics. *The Physical Review* **106**, 620–630 (1957).

[5] Jaynes, E. T. Information theory and statistical mechanics, ii. *The Physical Review* **108**, 171–190 (1957).

[6] Penrose, O. Foundations of statistical mechanics. *Reports on Progress in Physics* **42** (1979).

[7] Renner, R. *Security of quantum key distribution.* Ph.D. thesis, Swiss Federal Institute of Technology, Zurich (2005). E-print `quant-ph/0512258`.

[8] Deutsch, D. Uncertainty in quantum measurements. *Physical Review Letters* **50**, 631–633 (1983).

[9] Maassen, H. & Uffink, J. B. Generalized entropic uncertainty relations. *Physical Review Letters* **60**, 1103–1106 (1988).

[10] Berta, M., Christandl, M., Colbeck, R., Renes, J. M. & Renner, R. The uncertainty principle in the presence of quantum memory. *Nature Physics* **6**, 659–662 (2010).

[11] Bennett, C. Notes on Landauer's principle, reversible computation and Maxwell's demon. *Studies in History and Philosophy of Modern Physics* **34**, 501–510 (2003).

[12] Dahlsten, O., Renner, R., Rieper, E. & Vedral, V. The work value of information. e-print `quant-ph/0908.0424` (2009).

[13] Szilard, L. Über die Entropieverminderung in einem thermodynamischen System bei Eingriffen intelligenter Wesen. *Zeitschrift für Physik* **53**, 840 (1929).

[14] Shizume, K. Heat generation required by information erasure. *Physical Review E* **52**, 3495 (1995).

[15] Piechocinska, B. Information erasure. *Physical Review A* **61**, 062314 (2000).

[16] Janzing, D., Wocjan, P., Zeier, R., Geiss, R. & Beth, T. Thermodynamic cost of reliability and low temperatures: Tightening Landauer's principle and the second law. *International Journal of Theoretical Physics* **39**, 2717 (2000).

[17] Parrondo, J. M. R. The Szilard engine revisited: Entropy, macroscopic randomness, and symmetry breaking phase transitions. *Chaos* **11**, 725 (2001).

[18] Allahverdyan, A. E., Balian, R. & Nieuwenhizen, T. M. Maximal work extractions from finite quantum systems. *Europhysics Letters* **67**, 565 (2004).

[19] Horodecki, M. *et al.* Local versus non-local information in quantum-information theory: Formalism and phenomena. *Physical Review A* **71**, 062307 (2005).

[20] Maroney, O. J. E. Generalising landauer's principle. *Physical Review E* **79**, 031105 (2009).

[21] Hilt, S., Shabbir, S., Anders, J. & Lutz, E. Validity of landauer's principle in the quantum regime. e-print `quant-ph/1004.1599` (2010).

[22] Leff, H. S. & Rex, A. F. *Maxwell's demon: Entropy, information, computing* (Taylor and Francis, 1990).

[23] Leff, H. S. & Rex, A. F. *Maxwell's demon 2: Entropy, classical and quantum information, computing* (Taylor and Francis, 2002).

[24] Plenio, M. B. & Vitelli, V. The physics of forgetting: Landauer's erasure principle and information theory. *Contemporary Physics* **42**, 25–60 (2001).

[25] Maruyma, K., Nori, F. & Vedral, V. Colloquium: The physics of Maxwell's demon and information. *Reviews of Modern Physics* **81**, 1 (2009).

[26] Earman, J. & Norton, J. D. Exorcist XIV: The wrath of Maxwell's demon. part i. from Maxwell to Szilard. *Studies in the History and Philosophy of Modern Physics* **29**, 435–471 (1998).

[27] Earman, J. & Norton, J. D. Exorcist XIV: The wrath of Maxwell's demon. Part II. from Szilard to Landauer and beyond. *Studies in the History and Philosophy of Modern Physics* **30**, 1–40 (1999).

[28] Bub, J. Maxwell's demon and the thermodynamics of computation. *Studies in History and Philosophy of Modern Physics* **32**, 569–579 (2001).

[29] Oppenheim, J., Horodecki, M., Horodecki, P. & Horodecki, R. Thermodynamical approach to quantifying quantum correlations. *Physical Review Letters* **89**, 180402 (2002).

[30] Zurek, W. H. Quantum discord and Maxwell's demons. *Physical Review A* **67**, 012320 (2003).

[31] Jennings, D. & Rudolph, T. Entanglement and the thermodynamic arrow of time. *Physical Review E* **81**, 61130 (2010).

[32] Tomamichel, M., Colbeck, R. & Renner, R. A fully quantum asymptotic equipartition property. *IEEE Transactions on information theory* **55**, 5840–5847 (2009).

[33] Dupuis, F. *The decoupling approach to quantum information theory*. Ph.D. thesis, Université de Montréal (2009). Also available as `quant-ph/1004.1641v1`.

[34] Hayden, P. & Preskill, J. Black holes as mirrors: quantum information in random subsystems. *Journal of High Energy Physics* **09**, 120 (2007).

[35] Zurek, W. H. Einselection and decoherence from an information theory perspective. *Annals of Physics* **9**, 853–862 (2000).

[36] Ollivier, H. & Zurek, W. H. Quantum discord: a measure of the quantumness of correlations. *Physical Review Letters* **88**, 017901 (2000).

[37] Cavalcanti, D., Aolita, L., S. Boixo, K. M., Piani, M. & Winter, A. Operational interpretations of quantum discord. e-print `quant-ph/1008.3205` (2010).

[38] Bennett, C. H. Logical reversibility of computation. *IBM Journal of Research and Development* **17**, 525 (1973).

[39] Fredkin, E. & Toffoli, T. Conservative logic. *International Journal of Theoretical Physics* **21**, 219 (1982).

[40] Renes, J. M. & Renner, R. One-shot classical data compression with quantum side information and the distillation of common randomness or secret keys. e-print `quant-ph/1008.0452` (2010).

[41] König, R., Renner, R. & Schaffner, C. The operational meaning of min- and max-entropy. *IEEE Transactions on Information Theory* **55**, 4337–4347 (2009).

[42] Datta, N. & Renner, R. Smooth rényi entropies and the quantum information spectrum. *IEEE Transactions on information theory* **55**, 2807–2815 (2009).

[43] Tomamichel, M., Colbeck, R. & Renner, R. Duality between smooth min- and max-entropies. *IEEE Transactions on information theory* **56**, 4674–4681 (2010).

[44] Renner, R. Symmetry of large physical systems implies independence of subsystems. *Nature Physics* **3**, 645–649 (2007).

[45] Berta, M., Dupuis, F., Renner, R. & Wullschleger, J. Optimal decoupling. in preparation.