Quantum Thermodynamics

In this chapter we wish to give some insights to quantum thermodynamics. It can be seen as an introduction to the topic, however not a broad one but rather an introduction by examples. More specifically, we will pick particular topics and concepts through which it is possible to illustrate how information theoretic viewpoints can illuminate problems in thermodynamics.

The topics treated in the following include Landauer's principle, Szilard engines and different ways of interpreting them thermodynamically, the subjective definition of thermodynamic entropy and the emergence of thermalization in closed systems. It is important to note that for the topics concerning Landauer's principle, Szilard engines and entropy the word 'quantum' in the title is not so important. Only when it comes to thermalization we will see purely quantum features in our considerations. Nevertheless, we will use the formalism of density operators to describe all systems, classical and quantum.[1]

1 Landauer's Principle

Erasing a two-level system is the task of resetting the current state to a reference state, say $|0\rangle$. The mechanism implementing this state change must work independently of the initial state of the system. Landauer's principle states that no matter how you implement an erasure process, if it happens in an environment at temperature T then it has a minimal work cost of $k_B T \ln 2$, where k_B is the Boltzmann constant. This principle can be seen as a consequence of the second law of thermodynamics, as pointed out below in Section 2.2.

1.1 Work cost of erasure

Before going into more detail what an erasure process might look like in Section 1.3 we give a more quantitative version of Landauer's principle for more general cases than just two-level systems. For illustrative purposes though we restrict ourselves to n-bit strings, i.e. multiple two-level systems that can be processed globally. Let X be a random variable describing this n-bit string and P_X be the distribution according to which X is distributed. Furthermore, let $000\cdots0$ be the reference state to which we want to bring the system by erasing it. What is the best way of doing so? Doing it naïvely by erasing each bit individually would have a work cost of $n\, k_B T \ln 2$. However, this is not optimal as we will see now.

[1] Because we can.

In fact, we can use a basic result from source coding (see Exercise Sheet 4) stating that the information contained in n bits with distribution P_X can be reversibly compressed to $m = H_{\max}(X)_{P_X}$ bits while transforming the remaining $n-m$ bits to the state $00\cdots0$. Now it is of course more favourable to erase only m out of n bits with a work cost of $m\,k_BT\,\ln 2$. We thus end up with the more general result that erasure of a bit string described by $X \sim P_X$ can be done with a work cost of

$$H_{\max}(X)_{P_X}\, k_BT\, \ln 2. \tag{1}$$

A few remarks are in order. First, we assumed that the reversible compression step has no work cost, which is indeed the case in an optimal implementation. After all, compression boils down to the implementation of a bijective function from $\{0,1\}^n$ to $\{0,1\}^n$ which, by a simple argument given in the lecture, can be done without investing work. Second, although not proved here, one can show that this work cost is in fact optimal and extends directly to any (classical) system with finitely many states.[2] Finally, it might be the case that the erasure process is allowed to fail with probability ε. Failure here means that the final state is not $000\cdots0$. Then again we can make use of results from source coding, telling us that the number of bits needed to compress the content of X is $H^{\varepsilon}_{\max}(X)_{P_X}$, the smoothed max-entropy. Consequently, the work cost can be reduced to $H^{\varepsilon}_{\max}(X)_{P_X}\, k_BT\, \ln 2$ in this setting.

Interestingly we find that $H_{\max}$, initially a quantity used only in *information theory*, is the relevant *thermodynamic* entropy to describe the work cost of erasure. Thus the information theoretic viewpoint on thermodynamics is not just another way of looking at it, but is actually adding new insights.

1.2 I.i.d. limit

When having access to many, say k, copies of the same system $X \sim P_X$, applying the above result directly yields a work cost of erasure of $H_{\max}(X^{\times k})_{P_X^{\times k}}\, k_BT\, \ln 2$. By the asymptotic equipartition theorem this relaxes to $k\,H(X)_{P_X}\, k_BT\, \ln 2$ in the limit $k \to \infty$, where H is now the Shannon entropy. Hence, on average the relevant thermodynamic entropy is the Shannon entropy in this setting.

1.3 Work extraction from a pure qubit

The opposite of erasure is work extraction. In this task one starts with a well known (pure) state and transforms it into a mixed state while extracting work. To make the rather abstract discussion above more explicit we present a setting below in which work extraction can be described. Letting the process described below run backwards is one way of implementing erasure with an optimal work cost.

We shortly describe a setting with which optimal average work extraction, the 'opposite' of erasure, from a pure qubit can be made explicit. This example can easily be generalized to pure d-dimensional systems. Starting with a qubit in state $\rho = |0\rangle\langle 0|$ with degenerate Hamiltonian $H = E_0$, we make the assumptions:

[2] It also extends to quantum systems of finite dimension, as can be shown e.g. by combining results from del Rio *et al.*, Nature 10123 (2011) and Faist *et al.*, Nature Commun. 6 7669 (2015).

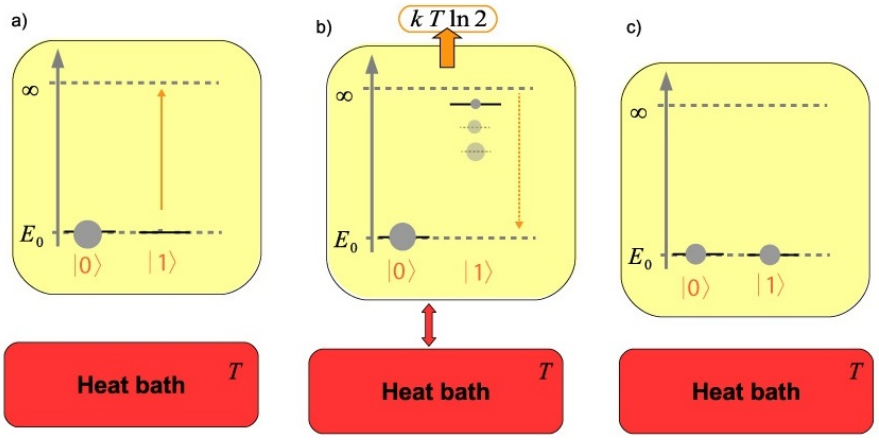

Figure 1: The setting consists of a system, here a qubit, initially in a pure state, and a heat bath. The bath itself is not further specified. Instead, its action in the system is described: when brought into contact with the bath the system thermalizes its state to the Gibbs state of the current Hamiltonian. The different panels are mentioned and described in the text.

- The energy difference between the energy eigenstates of the system (here denoted by $|0\rangle$ and $|1\rangle$) can be changed arbitrarily by changing the Hamiltonian of the system.

- Raising an energy level that is populated with probability p by energy $\mathrm{d}E$ costs work $p\,\mathrm{d}E$. If $\mathrm{d}E$ is negative this corresponds to lowering the energy of that state and work can be extracted.

- There is a heat bath at temperature T. Letting our system (here qubit) interact with it changes its current state to the Gibbs state $\tau = \frac{e^{-\beta H}}{\mathrm{tr}[e^{-\beta H}]}$, where $\beta = \frac{1}{k_B T}$ is the inverse temperature, k_B is the Boltzmann constant and H is the current Hamiltonian of the system.

Working in this paradigm we now carry out two steps depicted in Fig. 1 as a) and b). In step a we raise the unoccupied level $|1\rangle$ to a very high energy E_1. Later we let $E_1 \to \infty$. This has no work cost nor a work gain, because the initial state is $\rho = |0\rangle\langle 0|$. Then, in step b, we alternately connect the qubit to the thermal bath at temperature T and lower the energy of state $|1\rangle$ infintesimally by $\mathrm{d}E$. If the current eigenenergy of $|1\rangle$ is E, then the probability for being in this state is $p(E) = e^{-\beta(E-E_0)}/(1+e^{-\beta(E-E_0)})$. Hence the total work extracted

in this step if we start at E_1 and go back to the initial configuration at E_0 is

$$\begin{aligned} W_{\text{ext}} &= -\int_{E_1}^{E_0} p(E)\,\mathrm{d}E \\ &= \int_0^{E_1-E_0} \frac{e^{-\beta E}}{1+e^{-\beta E}}\,\mathrm{d}E \\ &= k_B T\left[\ln 2 - \ln\left(1+e^{-\beta(E_1-E_0)}\right)\right] \\ &\overset{E_1\to\infty}{\longrightarrow} k_B T\,\ln 2\,. \end{aligned} \tag{2}$$

We find that if the eigenenergy of $|1\rangle$ approaches ∞ after step a, then the extracted work can be arbitrarily close to the optimal value, $k_B T\,\ln 2$. Furthermore, the final state, depicted in c) of Fig. 1, is maximally mixed, $\rho = \frac{\mathbb{1}}{2}$. This makes sense, as the last thermalization step happens when the Hamiltonian is degenerate again. The Gibbs state of a system with degenerate Hamiltonian is fully mixed.

2 Szilard engine

Another paradigmatic example, proposed by Leo Szilard in 1929, that links information and thermodynamics is nowadays called Szilard engine. Szilard proposed a machine that seemingly produces work in a cyclic process while having access to only one heat bath. If this was possible it would be a violation of the second law of thermodynamics. In the following, we present the machine, describe its cycle, and analyse it from different perspectives to show that the second law is not violated by this machine. Also, the coming analysis emphasizes the subjectivity of thermodynamic statements.

A Szilard engine consists of a box with one particle inside, see the black boxes in Fig. 2. The box can be partitioned into two equal halves and it is in an environment (heat bath) at temperature T that thermalizes the box at all times. If the wall partitioning the box is moved by the particle, a weight can be attached to extract work during that process.[3] If work is extracted in this way in an isothermal process at temperature T, classical thermodynamics tells us that the extracted work for an N-particle ideal gas amounts to

$$W_{\text{ext}} = -\int_{V/2}^{V} -p\,\mathrm{d}V \overset{pV=Nk_BT}{=} \int_{V/2}^{V} \frac{Nk_BT}{V} = Nk_BT\,\ln 2\,. \tag{3}$$

For our ideal one-particle gas this is $k_B T\,\ln 2$.

A Szilard engine now operates in the following way: first a wall is inserted to partition the box into two halves. Since initially the particle could have been everywhere. It is then measured where the particle is, the outcome being 'L' (left) or 'R' (right). This information can be used to either turn the box around (if the outcome was 'R') or not (if it was 'L'). Attaching the weight on the left side as shown in Fig. 2 consequently allows us to extract $k_B T\,\ln 2$ work. The final state of the box will be the same as the initial one and everything happened at one temperature T.

[3] Of course, if the particle pushes the wall from the other side in this setting, the weight will be lowered and work is lost.

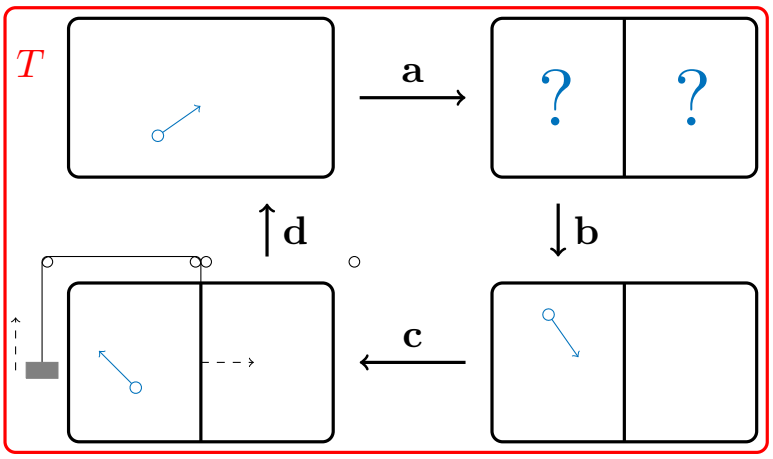

Figure 2: The working cycle of a Szilard engine. **a)** After inserting the partition it is not known on which side of the wall the particle is. **b)** Measuring allows us to find out about the position of the particle. **c)** Using the information gained in the measurement we can attach a weight on the correct side of the box such that letting the one-particle gas expand isothermally allows us to extract $k_B T \ln 2$ work. **d)** After the work extraction step the initial state of the box is restored. Here, every step happens at one and the same temperature T, all processes are assumed to be isothermal.

Since a measurement can be carried out at no work cost the work balance of this seemingly cyclic process tells us that in each cycle $k_B T \ln 2$ is extracted and no work is invested. So where is the problem?

2.1 Resolving the paradox

In fact, when we say that it is measured on which side of the wall the particle is, and that this information is used later on, we imply that there is a memory that stores this bit (information). Taking this memory into account, the analysis changes fundamentally because resetting (erasing) a random bit has a work cost. To see this, consider an agent A with a device storing this information. The box itself, essentially a two-level system, shall be denoted by B from now on. Initially the state of the Agent's device is idle, denoted by '$\perp$'. After the measurement it can either be 'L' or 'R'. We consider the states of the agent and the device in the three situations after step a (upper right), after step b (lower right) and after step d (upper left), see Fig. 2. We do so from three different views that will be analysed in detail.

View 1. Consider only B while ignoring A, in particular ignoring the microscopic knowledge accessible to A. The entropy of (uncertainty about) B in each of the relevant situations is then:

after a: The particle is equally likely to be on the left or on the right, hence $H(B) = 1$.

after b: The agent's knowledge is explicitly *not* taken into account. Thus the entropy of B has not changed and stays $H(B) = 1$.

after d: As argued above after this step the box is again in its initial state and the entropy is $H(B) = 1$.

So what exactly happened in between these situations? In step b the position of the particle is measured. But when restricting our view to B only, this is measuring without storing the outcome. Then in steps c and d work should be extracted. But in this view this is not possible because the information about the outcome of the measurement is needed to do so. In other words, the work extraction process acts on both A and B, not only on B. Hence, in this view one cannot talk about the work extraction step. We emphasize that this does not mean that view 1 is bad in any way. On the contrary, it shows that, as long as one respects that work extraction is impossible within this view, no contradiction with the second law arises.

View 2. We now consider B from the agent's point of view, in particular taking microscopic knowledge of A about B into account. The uncertainty about B in the relevant situations is:

after a: At this point A has no more information about B than we had before, so $H(B|A) = 1$.

after b: Now A stores the outcome of the position measurement, thus reducing the uncertainty about B to the minimum: $H(B|A) = 0$.

after d: After this step A has used her knowledge about the state of B to extract work, leaving B in the mixed initial state. The agent now no longer knows whether the particle is to the left or to the right of the wall, $H(B|A) = 1$.

In steps c and d the usual work extraction procedure happens, which can be described from A's perspective. Thereby, the agent's knowledge is used to extract $k_B T \ln 2$ work while increasing the entropy of B. However, now the problem occurs in step b. When the measurement happens this affects not only B but also A. But from the viewpoint of the agent it is impossible to describe changes to herself, this can only be done from outside.[4] Opposite to the previous view, in this view the analysis can be done for the work extraction step but not for the measurement step. Again, no contradiction with the second law arises.

View 3. Now we describe A and B from outside together.

after a: Here $H(AB) = H(A) + H(B|A) = 0 + 1 = 1$ because A is initially in the idle state '$\perp$' and has no knowledge of B's state.

after b: Now A stores the outcome of the position measurement, thus reducing the uncertainty about B . But at the same time looking at it from outside, the reduced state of A is fully mixed because both outcomes are equally likely, so $H(AB) = H(A) + H(B|A) = 1 + 0 = 1$.

[4]For instance, a global correlated states like $\frac{1}{2}|L\rangle\langle L|_A \otimes |L\rangle\langle L|_B + \frac{1}{2}|R\rangle\langle R|_A \otimes |R\rangle\langle R|_B$ makes sense from outside, but not when describing things from the point of view of A.

after d: After extracting work the state of the memory, i.e. of A, has not changed. On the other hand, A has no longer an accurate description of B, so we have $H(AB) = H(A) + H(B|A) = 1 + 1 = 2$.

From an outside view we can describe both steps accurately. During the measurement[5] correlations between A and B are built up. During work extraction information in A about B is used, but the state of A is not changed. Afterwards, B is again in the initial state and its entropy increased. Since our description now involves both systems, also this step can be described.

We emphasize that neither in the third view is the second law violated. The entropy of the total system, on which work extraction is performed, increases during the process by one bit.

2.2 Discussion

In neither of the three views inconsistencies arise, nor is any of them violating the second law. Only the third view can describe both steps which is why we use it to resolve the apparent paradox raised by the Szilard engine. Obviously, in this view, the entropy increased throughout the cycle of steps a–d. Hence the overall final state cannot be the same as the initial one. Instead of

$$|\perp\rangle\langle\perp|_A \otimes (\tfrac{1}{2}|L\rangle\langle L|_B + \tfrac{1}{2}|R\rangle\langle R|_B)$$

it is

$$(\tfrac{1}{2}|L\rangle\langle L|_A + \tfrac{1}{2}|R\rangle\langle R|_A) \otimes (\tfrac{1}{2}|L\rangle\langle L|_B + \tfrac{1}{2}|R\rangle\langle R|_B)\,.$$

To make it a cyclic process a fifth step, say e, resetting A would be necessary. This step is erasure and, if the second law holds, must have a minimal work cost of $k_B T \ln 2$ if it happens at temperature T. Thus Landauer's principle is a necessary consequence of the second law of thermodynamics. The net work gain of the cyclic process is therefore

$$W_{\text{ext}}^{\text{tot}} = k_B T \ln 2 - k_B T \ln 2 = 0\,, \tag{4}$$

in accordance with the second law.

Why did the paradox show up in the first place even though all of the above views, if treated properly, are consistent? It is because often naïvely one jumps between view 1 and 2 without noticing. When doing so, the step describable in view 1 is seen as measuring *and* storing the outcome, even though there is no memory. Then one switches to view 2, where work extraction has a proper description, and concludes that in total work was extracted and all systems are returned in their initial state. So changing views can lead to inconsistencies even though each view in itself is consistent.

An analogous explanation can be used to resolve Maxwell's demon paradox or Gibbs paradox (see exercises).

[5]The measurement is essentially a CNOT operation, where B is the control bit and A the target bit. In our case it changes the initial state $|\perp\rangle\langle\perp|_A \otimes (\frac{1}{2}|L\rangle\langle L|_B + \frac{1}{2}|R\rangle\langle R|_B)$ to $\frac{1}{2}|L\rangle\langle L|_A \otimes |L\rangle\langle L|_B + \frac{1}{2}|R\rangle\langle R|_A \otimes |R\rangle\langle R|_B$.

2.3 Further remarks

So far we assumed that both the Szilard engine and the agent are in the same environment and thus operate at the same temperature. What if the box B is held at temperature T_B and the memory of the agent at T_A? If $T_A < T_B$ the work needed to erase the memory would be smaller than the work extracted during the expansion of the one-particle gas and the cycle would yield a net work gain. However, this does not contradict the second law either because now there are two heat baths instead of just one and a net heat flow from the hot to the cold bath. This net heat flow is analogous to the heat flow from hot to cold in thermal machines.

More generally, in this case one could define an efficiency of the 'information driven' thermal machine as

$$\eta = \frac{\text{extracted work} - \text{invested work}}{\text{invested work}} = \frac{k_B T_B \ln 2 - k_B T_A \ln 2}{k_B T_A \ln 2} = \frac{T_B}{T_A} - 1 \,. \tag{5}$$

There are three regimes in which such a machine can work:

$\underline{T_A = T_B}$: This case has been treated above and we effectively found $\eta = 0$. Hence the process neither produces work nor uses work to do some other task.

$\underline{T_A < T_B}$: Here, the work cost for erasure is smaller than the work extracted, which corresponds to the situation described in the above paragraph. Therefore we have $\eta > 0$ and the 'information driven' engine essentially acts as a heat engine, producing work while letting heat flow from a hot to a cold bath.

$\underline{T_A > T_B}$: If the work cost for erasure surmounts the amount of extracted work it follows that $\eta < 0$. Now the machine acts a refrigerator (or, equivalently, a heat pump) pumping heat from the cold bath to the warmer one.

All in all, through the Szilard engine we found an information theoretic description of thermal engines which can operate in different regimes as heat engines or refrigerators / heat pumps.

2.4 Remark on the objectivity of thermodynamic entropy

In traditional thermodynamics it is usually assumed that entropy is an objective quantity. This seemingly follows from the operational definition of thermodynamic entropy as

$$\delta S = \frac{\mathrm{d}Q_{\mathrm{rev}}}{T} \,, \tag{6}$$

the quotient of the reversibly exchanged heat $\mathrm{d}Q_{\mathrm{rev}}$ and the temperature T. Both heat and temperature are measurable quantities and thus two agents doing the same experiment should get the same outcomes.

On the other hand we argued above that one can have different views on thermodynamic processes and, depending on which view we take on, different entropies are relevant. However, the above analysis was conducted from an information theoretic perspective and it remains open whether one could also arrive at a subjective definition of entropy in a thermodynamic analysis.

We here argue that, opposite to the argument given above, the definition given in Eq. 6 is a subjective one. The reason for this lies in the fact that $\mathrm{d}Q_{\mathrm{rev}}$ is not an objective quantity. To see this we have to go back to the definition of reversible heat, which goes via the definition of reversible work. Work is usually seen as the energy that is exchanged with a system in a *controlled* way in thermodynamic processes, meaning that energy that is called work could also be used to drive other processes like lifting a weight, for instance. What is seen as controlled energy sources largely depends on the observer, as well as the reversibility of such processes. For instance, taking the Gibbs paradox as an example, whether there is a thermodynamic process that separates two types of particles in a gas depends on whether the experimenter knows about the two types or not. Two observers with different knowledge will therefore call the process of mixing two gases to one gas reversible or not depending on their knowledge. But if it can subjective whether a process is reversible or not then also the quantity $\mathrm{d}Q_{\mathrm{rev}}$ is a subjective one.

Having argued that $\mathrm{d}Q_{\mathrm{rev}}$ is not the same for all observers it directly follows that δS defined through Eq. 6 is subjective. It depends on what operations / processes the observer considers reversible.

3 Information theoretic views on thermalization

We now come to the last part of the excursion to thermodynamics and information theoretic views on it. Also, this is the only section in which purely quantum phenomena show up while the previous analysis classical.

The main question with which we will be concerned in this section is: how can one justify that an initially pure state of a system S will turn to a mixed state after some time?

Thinking of quantum mechanics a closed system evolves unitarily. Hence an initially pure state will stay pure for all times. More generally, (von Neumann) entropy is constant under unitary evolutions and so in this case no mixedness can be introduced. If S can interact with an environment, say E, it is no longer closed and its entropy can increase during a joint evolution with E.

The same argument holds for a classical system that evolves according to the laws of classical mechanics, which are deterministic. Clearly for a closed system S entropy will stay constant, but if the system interacts with an environment E, an entropy increase of S during the evolution is possible.

In the following we will investigate how exactly this can happen, and how typical it is.

3.1 Analogy: a die

The reason why we say that systems thermalize after some time is the same as the reason why we say that the outcome of a die is random. see Fig. 3. Also for a classical die D starting in one of six possible states $d \in \{1, \ldots, 6\}$, say $d = 4$, the initial entropy is zero, $H(D) = 0$. After throwing it, we usually say that each outcome is equally likely, hence $H(D) = \log 6$ after that. Clearly D is not a closed system in this consideration, so entropy can increase. Nevertheless the question remains how.

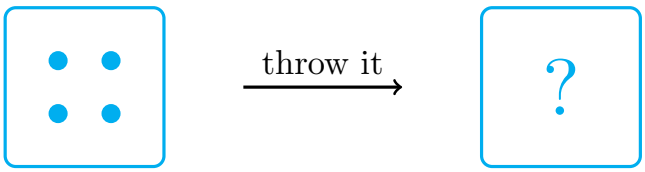

$H(D) = 0$ $\qquad\qquad$ $H(D) = \log 6$

Figure 3: Even if we know the initial state of a die D, after throwing it we usually say that its outcome is random. The (Shannon) entropy of D is initially $H(D) = 0$, but finally it is maximal, $H(D) = \log 6$, as the actual state $d \in \{1, \ldots, 6\}$ is unknown.

To understand this let us model the environment too. For simplicity we assume that the environment consists only of an agent throwing the die and looking at the outcome. Agent and die together form a closed system. The agent has many different ways to throw D and initially it is random which one is chosen. Therefore there is some uncertainty about A initially, which we quantify as $H(A) = c$ for some $c \geq 0$. For the total system in the initial state we have $H(AD) = H(A) + H(D) = c + 0 = c$. Since AD together are assumed to be closed the entropy does not change throughout the evolution, so $H(AD) = c$ will stay constant. However, $H(D)$ may change.

This becomes possible because initially A and D are uncorrelated, $I(A : D) = 0$, but finally $I(A : D) = \log 6$ is possible.[6] To see this, consider the simplified case when there are exactly six ways of throwing the die, $\{a_1, \ldots, a_6\}$, and where a_i leads to outcome i for D. Assume that initially, all 6 ways of throwing the die are equally likely, which means that the initial state of AD can be described as

$$\rho_{AD}^{\text{in}} = \sum_{i=1}^{6} \frac{1}{6} |a_i\rangle\langle a_i|_A \otimes |4\rangle\langle 4|_D \,,$$

and $H(A) = \log 6$. After throwing D the final state may be[7]

$$\rho_{AD}^{\text{out}} = \sum_{i=1}^{6} \frac{1}{6} |a_i\rangle\langle a_i|_A \otimes |i\rangle\langle i|_D \,.$$

The total entropy of AD will not have changed, but the reduced state on D after the joint evolution is $\rho_D^{\text{out}} = \frac{1}{6}\sum_i |i\rangle\langle i|_D$, which is fully mixed and thus $H(D) = \log 6$. Also it is easy to check that in this case for the final state $I(A : D) = \log 6$.

To sum up, the reason why the outcome of the die is random is that randomness that was already in the agent throwing the die. This randomness

[6] If c was too small, $c < \log 6$, $I(A : D) = \log 6$ is not achievable. Above we assumed that there are many ways in which the agent can throw the die and it is unknown to us which one is going to be chosen. Hence it is fair to assume that $c \geq \log 6$ which allows the final mutual information to be $I(A : D) = \log 6$ for some evolutions.

[7] Check for yourself that closed evolution (i.e. unitary evolution in this notation) can lead to this state change.

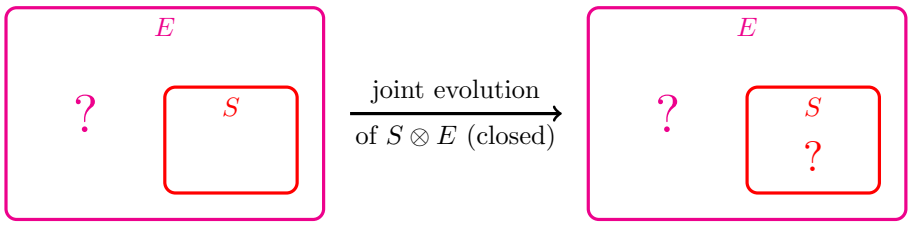

$H(S) = 0$, $H(E) = c$ $\qquad$ $H(S) = c'$, $H(E) = c$

Figure 4: Consider two systems, S and E, where E is much larger than S. The total system is $S \otimes E$. They are assumed to be uncorrelated initially, which implies $H(SE) = H(S) + H(E)$. Furthermore, S may be in a pure state with no entropy, whereas there may be some entropy c in E initially. Nevertheless, just like in the example of the die above, after joint evolution the entropy of E can be the same as the initial entropy while the entropy of S increased to $H(S) = c'$ due to correlations between S and E.

is extended to the correlations between A and D during the evolution, which makes the final state of D random, too. It is not the case that entropy flow from A to D. Also, it is important to note that no entropy 'flew' from A to D, as also in the final state $H(A) = \log 6$.

3.2 Qualitative analysis – classical and quantum

Let us now apply the above to understand why subsystems of larger systems thermalize. Thermalization is always accompanied with an entropy increase and in a first step we focus on this characteristic.

Let S be the system of interest and E be a much larger environment, see Fig. 4. Notice that S is *not* a subsystem of E but rather S and E together form a composite system $S \otimes E$. We have seen above in the example with the die that initial entropy in E can lead to an entropy increase in S during joint evolution, even though the entropy of E is not diminished. This shows that it would be wrong to think of it as entropy flowing from E to S. The effect relies on correlations built up during the process which allow S to 'inherit uncertainty' from E without making E more pure.

An obvious question is now whether such an effect could also occur if there was no initial randomness in E, i.e. $H(E) = 0$ initially. However, this view can only be consistent with our everyday observations if subsystems undergoing joint evolution with other systems can experience entropy increases.

Classically. In a classical world this is impossible, as we will show here. If S and E share no correlations initially and both contain no entropy, $H(S) = H(E) = 0$, then we have $H(SE) = H(S) + H(E) = 0$ initially, and also in the final state because $S \otimes E$ undergoes closed evolution. But then also in the final

state

$$0 = H(SE) = H(S) + \underbrace{H(E|S)}_{\geq 0}, \tag{7}$$

where $H(E|S) \geq 0$ because both E and S are classical systems. Since also $H(S)$ is non-negative, it follows that $H(S) = 0$ for all times. Thus, if the universe was classical and initially in a pure state, then subsystems cannot increase their entropy locally, which prohibits thermalization.[8]

Quantum. In a quantum world the phenomenon of having $H(S) > 0$ even though the universe is in a pure state is, surprisingly, possible. Even if there are no initial correlations between S and E, and both systems are in a pure state initially, joint evolution can lead to a local entropy increase on S. The reason for this is that in a quantum world $H(E|S)$ can become negative, giving space for $H(S)$ to be positive in Eq. 7.

This is a purely quantum feature because only entangled states on $S \otimes E$ can have negative conditional entropy.[9] We conclude that in a quantum world, even if $c = 0$ (the initial entropy in the universe), thermalization processes can occur.

3.3 Thermalization through entanglement

We now present a result that quantifies how typical it is to observe thermalization in quantum systems. This result is stated in very general terms and we will comment on its implications on thermalization processes after having stated it.

The setting is again the same: We have two systems, the systems of interest S, and the rest E is called environment. The Hilbert spaces of those systems are denoted by $\mathcal{H}_S$ and $\mathcal{H}_E$, respectively. Let $\mathcal{H}_R \subset \mathcal{H}_S \otimes \mathcal{H}_E$ be a subspace of the total Hilbert space, e.g. defined via energy constraints on the possible states. Define the fully mixed state on R to be $\mathcal{E}_R = \frac{\mathbb{1}_R}{d_R}$ and the *canonical* states on S and E as $\Omega_S = \mathrm{tr}_E[\mathcal{E}_R]$ and $\Omega_E = \mathrm{tr}_S[\mathcal{E}_R]$, respectively.

Informally, the following theorem states that for a randomly chosen pure state on $\mathcal{H}_R$ the marginal of this state on S is very close to the canonical state on S, provided the dimension of S is much smaller than the one of E, $d_S \ll d_E$.

Theorem. [Popescu, Short, Winter (2006)]
Let $\mathcal{H}_S \otimes \mathcal{H}_E$ be a bipartite Hilbert space of dimension $d_S \cdot d_E$ and $\mathcal{H}_R \subset \mathcal{H}_S \otimes \mathcal{H}_E$ a subspace of dimension d_R. Define $\mathcal{E}_R = \frac{\mathbb{1}_R}{d_R}$ and the corresponding marginals $\Omega_S = tr_E[\mathcal{E}_R]$ and $\Omega_E = tr_S[\mathcal{E}_R]$. Then for a randomly chosen pure state on $\mathcal{H}_R$, $|\phi\rangle \in \mathcal{H}_R$, and arbitrary $\varepsilon > 0$, the trace distance between the actual reduced state on S, $\rho_S = tr_E[|\phi\rangle\langle\phi|]$, and the canonical state Ω_S is given probabilistically by

$$P[\, \|\rho_S - \Omega_S\|_1 \geq \eta \,] \leq \eta' \,, \tag{8}$$

[8] Equivalently this can be seen as a consequence of the fact that for a deterministic global probability distribution any marginal is deterministic, too.

[9] Notice that there is no 1-1 correspondence between entangled states and negative conditional entropy. However, if a conditional entropy is negative, then the corresponding state must be entangled.

where

$$\eta = \varepsilon + \sqrt{\frac{d_S}{d_E^{\text{eff}}}}\,, \quad \eta' = 2e^{-Cd_R\varepsilon^2}\,, \quad d_E^{\text{eff}} = \frac{1}{tr[\Omega_E^2]} \geq \frac{d_R}{d_S}\,, \quad C = \frac{1}{18\pi^3}\,. \tag{9}$$

In applications the environment will be much larger than the system, $d_E \gg d_S$, and $d_R \gg 1$ s.t. both η and η' will be small and the actual state ρ_S will be close to the canonical state Ω_S with high probability.

This theorem can be interpreted as a statement about typical states.[10] Namely, for typical states $|\phi\rangle \in \mathcal{H}_R$, even if they are pure on $S \otimes E$, the reduced state of S is mixed. Furthermore, this reduced state is mixed in a very specific way. It will be such that it is very close to the canonical state in trace distance.

What is the canonical state? Let us consider two examples. First, let R be the whole system $S \otimes E$, i.e. no restrictions. In this case the canonical states are fully mixed. Hence the theorem tells us that typical pure states on $S \otimes E$ are fully mixed on S with high probability.

The second example concerns thermodynamics. It can be shown (see exercises) that if R is restricted through energy constraints (e.g. 'the total energy of S and E must be equal to $\bar{E}$') then $\Omega_S \propto e^{-\beta H_S}$, where H_S is the Hamiltonian on S and β is the inverse temperature of the environment E. Then the theorem tells us that typical pure states on $S \otimes E$ with a energy $\bar{E}$ are very close to thermal on S with high probability.

Comment on typical states. The previous theorem made statements about *typical states* (also called *random*). How are typical / random states defined? A measure on a set is usually called uniform if it is invariant under some sort of translation. In the case of a complex Hilbert space $\mathcal{H}$ 'translations' are unitary operations on the vectors. Hence, a measure on $\mathcal{H}$ is called uniform if it does not change when first applying a unitary operation.

Such a measure can be derived from the Haar measure, $\mathrm{d}\mu$, over the unitaries on $\mathcal{H}$, $\mathcal{U}(\mathcal{H})$, simply by starting from an arbitrary reference state $|0\rangle_S|0\rangle_E$ and then applying a Haar-random unitary $U \in \mathcal{U}(\mathcal{H})$ to this state. The probability distribution is then invariant under the application of a unitary $V \in \mathcal{U}(\mathcal{H})$ because $\mathrm{d}\mu(VU) = \mathrm{d}\mu(UV) = \mathrm{d}\mu(U)$ for all $V \in \mathcal{U}(\mathcal{H})$.

Comment on typical evolution. When we talk about thermalization we usually speak of *typical evolution* that leads to thermalization, rather than typical states. In fact, the previous comment can be seen as defining typical states in terms of typical evolution. Thinking of typical evolution of a closed system as a Haar-random unitary, we ultimately find that typical evolution leads to typical states

[10] See comments below for a definition of what is referred to as typical states.