

The power of being explicit: demystifying work, heat, and free energy in the physics of computation

Thomas. E. Ouldridge,^{*} Rory A. Brittain,[†] Pieter Rein ten Wolde[‡]

^{*}To whom correspondence should be addressed.

1 *Introduction*

Interest in the thermodynamics of computation has revived in recent years, driven by developments in science, economics and technology. Given the consequences of the growing demand for computational power, the idea of reducing the energy cost of computations has gained new importance (DeBenedictis, Mee, and Frank 2017; Frank 2017). Simultaneously, many biological networks are now interpreted as information-processing or computational systems constrained by their underlying thermodynamics (Andrieux and Gaspard 2008; Lan et al. 2012; Mehta and Schwab 2012; Ito and Sagawa 2015; Govern and ten Wolde 2014; Barato, Hartich,

^{*}Department of Bioengineering and Centre for Synthetic Biology, Imperial College London, London, SW7 2AZ, UK; t.ouldridge@imperial.ac.uk

[†]Department of Mathematics, Imperial College London, London, SW7 2AZ, UK; r.brittain15@imperial.ac.uk

[‡]FOM Institute AMOLF, Science Park 104, 1098 XE Amsterdam, The Netherlands; ten-wolde@amolf.nl

and Seifert 2014; Mehta, Lang, and Schwab 2016; ten Wolde et al. 2016; Ouldrige and ten Wolde 2017; Ouldrige, Govern, and ten Wolde 2017; Barato and Seifert 2017). Indeed, some suggest (Bennett 1982; Adleman 1994; Organick et al. 2018) that low-cost, high-density biological systems may help to mitigate the rising demand for computational power and the “end” of Moore’s law of exponential growth in the density of transistors (Frank 2017).

In this chapter we address widespread misconceptions about thermodynamics and the thermodynamics of computation. In particular, we will argue against the general perception that a measurement or copy operation can be performed at no cost; against the emphasis placed on the significance of erasure operations; and against the careless discussion of heat and work. While not universal, these misconceptions are sufficiently prevalent (particularly within interdisciplinary contexts) to warrant a detailed discussion. In the process, we will argue that explicitly representing fundamental processes is a useful tool, serving to demystify key concepts.

We first give a brief overview of thermodynamics, then the history of the thermodynamics of computation - particularly in terms of copy and measurement operations inherent to classic thought experiments. Subsequently, we analyse these ideas via an explicit biochemical representation of the entire cycle of Szilard’s engine. In doing so we show that molecular computation is both a promising engineering paradigm, and a valuable tool in providing fundamental understanding.

2 Basic thermodynamics

Here we briefly recap the key ideas of work, heat, entropy and free energy. In effect, we introduce several consistent ways to think about the second law of thermodynamics. Experienced readers may wish to skip to Section 3.1; those wishing for more detail on the background of equilibrium thermodynamics and statistical mechanics may refer to Atkins

2010 and Huang 1987. Note that the quantities we consider here are implicitly statistical averages, rather than quantities defined for individual fluctuating trajectories (Jarzynski 1997; Crooks 1999; Esposito and Van den Broeck 2011; Seifert 2005), and that we will restrict ourselves to the classical regime.

Work, heat, the second law and entropy

Energy is conserved, but not all energy is equivalent. Certain forms of energy, like that stored when a mass is raised in a gravitational field, are more useful than others, such as that stored in a hot object. The energy of an ideal mass is stored in one accessible (collective) degree of freedom: the position of the centre of mass in a gravitational field. The energy of a hot object such as a volume of gas is stored in many degrees of freedom (the momenta of the gas particles). Unlike the mass, it is impossible for an experimenter to couple to these degrees of freedom individually and use the stored energy. This distinction underlies the second law.

Energy is transferred as heat when the transfer occurs through many degrees of freedom that are individually inaccessible. By contrast, work is done when the transfer of energy occurs via the accessible degrees of freedom. We define a heat reservoir as an ideal, arbitrarily large system that can only accept energy transfer via heat (think of a large gas at fixed volume). The heat reservoir is characterised by a temperature T , which quantifies the average energy per degree of freedom. We define an ideal work reservoir as an ideal system, like a mass in a gravitational field, that only stores energy in an accessible degree of freedom. We note that this definition of work – in terms of the transfer of energy via the manipulation of accessible coordinates – is consistent with the definition used in modern stochastic thermodynamics (Horowitz and Jarzynski 2008). The definition of heat is then, necessarily, also consistent.

One of the traditional statements of the second law is: "No process is possible whose sole result is the conversion of heat into work." This statement makes precise the idea that "not all energies are equal". No

engine exists that cyclically takes energy from a single heat reservoir and deposits all of it into a work reservoir (Fig. 1 (a)). The cyclical operation is important here – if the engine underwent a sustained change, then the transfer of energy would not be the "sole result" of the process.

A more typical modern statement of the second law would be: "In a closed system, the entropy is non-decreasing." The entropy quantifies the uncertainty in the precise state of a system. Specifically (Jaynes 1957),

$$\mathcal{S}[p(x)] = -k_B \sum_x p(x) \ln p(x). \quad (1)$$

Here, x denotes the possible states of a system, and $p(x)$ the probability that the system will be found in state x . The flatter the distribution $p(x)$, the larger the entropy. If we are completely certain of the state of a system, then all terms in the sum in Eq. 1 are identically zero.

Entropy incorporates the traditional statement of the second law. When energy is transferred to a work reservoir, there is no change in the work reservoir's entropy since it only has one degree of freedom – its energy uniquely specifies its state. By contrast, when energy is transferred to a heat reservoir, it is shared between the many degrees of freedom. More energy increases the uncertainty, since there are more ways to divide the energy up. Thus the transfer of energy to a heat reservoir is associated with an entropy increase, and converting heat directly into work is forbidden because it would be an overall decrease in the entropy of a closed system. To be precise, if heat dQ is transferred to a heat reservoir, the reservoir's entropy increase is $d\mathcal{S} = TdQ$. Importantly, entropic definition of the second law allows us to analyse processes in which systems that are not ideal reservoirs undergo change, using Eq. 1; this framework allows us to analyse non-cyclic systems, and steps of cyclic operations.

Equilibrium, the Boltzmann distribution and free energies

An isolated system with sufficiently many interacting degrees of freedom will evolve towards an equilibrium steady state in which there is no systematic variation of observables with time, or net currents of matter




Figure 1. (a) Graphical representation of the second law of thermodynamics: a cyclic engine cannot solely take in heat from a reservoir T and convert it to work to raise a mass m . (b) Illustration of the canonical ensemble: a small system of interest is embedded within a larger system setting the temperature T . Heat Q can be exchanged between system and reservoir. (c) The input of external work, or the coupling to another subsystem, can be used to drive up the free energy of a subsystem A so that $\mathcal{F}'_A > \mathcal{F}_A$. In the absence of the second subsystem, the second law dictates that $W_{in} \geq \mathcal{F}'_A - \mathcal{F}_A$. In the other case of no external work input, $\mathcal{F}_B - \mathcal{F}'_B \geq \mathcal{F}'_A - \mathcal{F}_A$.

or energy. Since the entropy of an isolated system cannot decrease, the equilibrium distribution $p_{eq}(x)$ maximises the entropy subject to whatever constraints are relevant (e.g., the overall conservation of energy). It is often convenient to consider a small system of interest coupled to a heat reservoir at temperature T , as illustrated in Fig. 1 (b). In this case, representing the states of the system by x , it is possible to show that

$$p_{eq}(x) \propto \exp(-\epsilon(x)/k_B T), \quad (2)$$

where $\epsilon(x)$ is the energy of configuration x (Huang 1987). The distribution $p_{eq}(x)$ shares a fixed amount of energy between the system and the heat reservoir so as to maximise the overall entropy of the combined distribution of system and reservoir. For simplicity, we restrict ourselves to this "canonical ensemble".

Prior to reaching equilibrium, systems have a non-equilibrium distribution, $p(x) \neq p_{eq}(x)$. The distance from equilibrium can be characterised by the non-equilibrium generalized free energy (Esposito and Van

den Broeck 2011; Parrondo, Horowitz, and Sagawa 2015)

$$\begin{aligned}\mathcal{F}[p(x)] &= \sum_x p(x) \epsilon(x) + k_B T \sum_x p(x) \ln p(x) = \langle \epsilon \rangle - T \mathcal{S}[p(x)] \\ &= \mathcal{F}[p_{\text{eq}}(x)] + k_B T \sum_x p(x) \ln \frac{p(x)}{p_{\text{eq}}(x)} \geq \mathcal{F}[p_{\text{eq}}(x)].\end{aligned}\tag{3}$$

Following a process in which $p(x)$ evolves over time, the change in the total entropy of the system and reservoir is given by $-\Delta \mathcal{F}[p(x)]/T$; $\mathcal{F}[p(x)]$ then *decreases* over time, and is *minimal* at equilibrium. This minimum represents a balance between minimisation of energy $\langle \epsilon \rangle$ and maximisation of entropy $\mathcal{S}[p(x)]$ within the system.

Although non-equilibrium free energies cannot increase spontaneously, they can be driven upwards as in Fig. 1 (c). Specifically, the change in free energy determines the minimal work that must be applied to move a system A coupled to a heat reservoir between $p(x)$ and $p'(x)$ (Esposito and Van den Broeck 2011): $\mathcal{W}_{\text{in}} \geq \mathcal{F}_A[p'(x)] - \mathcal{F}_A[p(x)]$. As a consequence, the maximum work that can be extracted is $\mathcal{W}_{\text{out}} \leq \mathcal{F}_A[p(x)] - \mathcal{F}_A[p'(x)]$. The system's free energy must change to obtain positive work since work cannot be extracted from the heat reservoir alone.

We could also couple to another system that is not an idealised work reservoir (Fig. 1 (c)). In this case, the second law dictates that the total free energy of the combined system must decrease, but a decrease in the free energy of one subsystem can compensate for an increase in the other. A non-equilibrium system is then fundamentally an exploitable resource. If a system A is out of equilibrium, we can in principle couple it to a second system B and use the relaxation of A to drive a change in B . If A is in equilibrium, however, it is not exploitable in this way.

It is important to note that, although the free energy has the dimensions of an energy, and although changes in free energy are bounded by work input, free energies are not energies in a very fundamental sense. Energies obey the first law of thermodynamics; energy can never be destroyed, only converted between different forms. If work is done on a system, either its internal energy must change, or energy must be transferred to the environment as heat. Free energies, however, also have an

entropic component (Eq. 3). They are therefore not conserved; it is perfectly possible for the free energy of a system to decrease with no other consequences. We shall explore this fact in more detail in Section 5.2.

Thermodynamic reversibility

Consider applying a time-varying manipulation of a system, driving it from $p(x)$ to $p(x')$. If and only if this process is thermodynamically reversible, applying the manipulation in a time-reversed manner would convert $p(x')$ back to $p(x)$, and the intermediate probability distributions obtained during the forward manipulation would be reproduced in the opposite order during the reverse protocol (see Sagawa 2014). Thermodynamically reversible processes necessarily involve no change in the overall entropy of the system and reservoirs to which it is coupled, since a positive entropy increase on the forwards manipulation would correspond to a forbidden decrease on the reverse. Reversible processes therefore require work input of exactly $\mathcal{W}_{\text{in}} = F[p'(x)] - F[p(x)]$ (Esposito and Van den Broeck 2011), and constitute an important idealised limit of practical processes.

3 *The physics of computation*

To perform computation, a physical system – be it silicon-based or biological – must be able to change in response to a range of inputs. Indeed, data and information must have a physical realisation simply to exist, let alone to be manipulated during a computation (Landauer 1991). If computation involves changes in the states of physical systems, then these processes can be analysed thermodynamically, and the resultant thermodynamic costs and requirements can be assessed. This approach is becoming increasingly fashionable, driven by a desire to reduce the consumption of electricity by our increasing demands on computational power (DeBenedictis, Mee, and Frank 2017; Frank 2017).

Although recent years have seen a wider focus (see e.g. Owen, Kolchinsky, and Wolpert 2017), the study of the thermodynamics of computation has been dominated by the discussion of two processes – erasure and measurement (or copying) of single two-state memories. Partly, this focus follows from the fact that these processes are prototypical computational operations that provide much of the framework for understanding more complex systems. Equally importantly, however, erasure and measurement lie at the heart of a fundamental debate in physics relating to the validity and meaning of the second law of thermodynamics. In the following paragraphs, we summarise this debate.

A history of Maxwell's demon and Szilard's engine

Physicists have long been concerned by the second law. Most famously, James Clerk Maxwell wondered whether an intelligent being (a "demon") could violate the second law by extracting work from random fluctuations of an equilibrium system (Dougal 2016). Maxwell's demon operated a trap door separating two volumes of gas at initially equal temperature. If the demon could determine the velocity of gas particles, it could selectively open the door to allow high energy particles to pass in one direction, and low energy particles in the other. Over time a temperature difference would develop, to which a work-extracting heat engine could be coupled. There is no fundamental lower bound on the work needed to open and close a trap door. Therefore it appeared to Maxwell that the demon could use its intelligence to perform measurement and feedback to exploit fluctuations and thereby violate the second law of thermodynamics by extracting work from an initially equilibrated system.

Many authors have attempted to justify why the second law is not violated by Maxwell's demon. Typically, they construct a demon-like system, analyze its behaviour and demonstrate that violation of the second law does not occur. This approach has been criticised (Norton 2005), on the grounds that the accuracy of the second law is generally assumed in setting up the analysis. We, however, believe that it is valid both to test

the self-consistency of the second law in this way, and to ask the conceptual question of *how* the self-consistency is manifested, as the understanding gained applies beyond the original context.

Whilst the term "Maxwell's demon" is used quite liberally, we will focus on systems in which a measurement is made, stored, and then used to implement a measurement-dependent feedback. In 1929, Leo Szilard described an analytically tractable "engine" with measurement and feedback (Szilard 1964). A version is schematically illustrated in Fig. 2. Here, the demon simply determines whether a single particle is located on the left or right hand side of a piston in a hollow cylinder, and configures the system to exploit this fact. Given knowledge of the particle's position, the piston can be coupled to a weight so that the pressure of the particle can do work to raise the weight. It is clear that measurement and feedback are required to couple the weight correctly. Assuming that this procedure is performed with 100% accuracy, the maximum work extracted per expansion can be calculated by integrating $p dV$. Here, p is the average pressure applied by the particle and V is the volume in which it is contained, which grows from $V_{\text{cylinder}}/2$ to V_{cylinder} . The result is $k_B T \ln 2$.

Szilard sought to identify a cost that compensated for the apparent ability to extract work from the initial equilibrium system. He realised that the essence of his system could be analysed without reference to an intelligent being. Instead, he argued that the demon's measurement generates a correlation between two physical degrees of freedom. In Fig. 2, these degrees of freedom are the position of the particle (the data) and the weight's connection to the partition (the memory). Crucially, Szilard emphasized that post-measurement, the memory must retain its state even if the data subsequently changes. In Fig. 2, for example, the weight must stay attached to the same side of the piston even as the particle's accessible volume expands to fill the whole cylinder. Szilard stated that "if the measurement could take place without compensation", then the second law could be violated. He claimed that there must be a minimum "production of entropy" during measurement that compensated for the

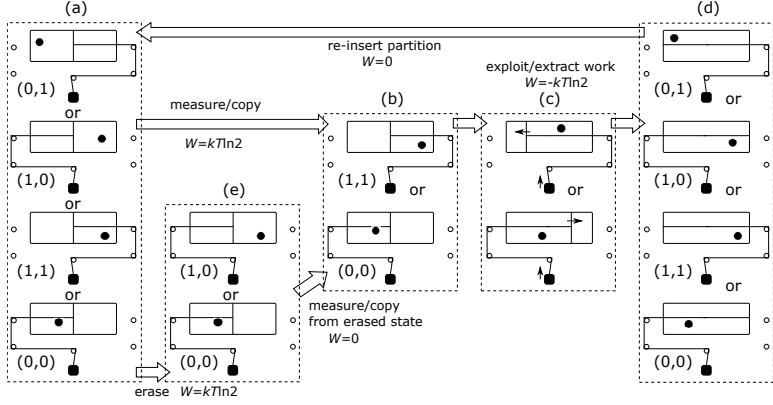


Figure 2. A cycle of Szilard's engine, with possible operating states illustrated. Steps are labelled with the minimal work input required for reversible operation (negative values correspond to work extraction). The engine consists of a single particle in a volume, a partition, and a weight that can be attached to either side of the partition via pulleys. (a) Initially the particle and the weight can both be associated with either side of the volume, and four states are possible (labelled (d,m) , where $d = 0(1)$ corresponds to the particle on the left(right) hand side of the system, and $m = 0(1)$ corresponds to the apparatus on the left(right) side of the partition). (b) A measurement correlates the weight's connection with the particle's position. Two states are possible after a perfectly accurate measurement. (c) The particle is allowed to expand against the partition, raising the weight. Note that the weight's attachment state does not change. Eventually state (d) is reached, when the particle's position is uncorrelated with the weight's attachment. The system is then restored to its initial condition by re-inserting the partition. (e) Alternatively, one might first "erase", setting the pulleys into a guaranteed configuration and leaving two possible states overall, prior to performing the measurement.

$k_B T \ln 2$ of work extracted during the expansion. In the computational sense, he was arguing that there is a cost to copying of data.

In 1951, Brillouin reopened the question. He too believed that the demon needed to pay a cost when measuring (Brillouin 1951). He associated this cost with the generation of a non-equilibrium photon that could be used to see the particle. Brillouin did not consider correlating a memory with the measured system, but merely the interaction mechanism. His proposed measurement costs are therefore of a fundamentally different character from Szilard's.

A decade later, Landauer asked whether heat generation was necessary during computation (Landauer 1961). Although he referred to Brillouin's work, he did not deal with demons. Rather, he argued that computational operations without a single-valued inverse (many-to-one functions) necessarily lead to an increase in the entropy of the environment. Landauer pointed out that such "logically irreversible" functions – for example, "erasing" a bit of unknown value to zero – decrease the entropy of the bit itself. He argued that there must be an entropy increase in the environment to preserve the second law. For a symmetric bit initially in state 0 with 50% probability, the initial entropy is $k_B \ln 2$ (from Eq. 1). After erasure is complete, the entropy of the bit is 0. Thus the entropy of the environment must increase by at least $k_B \ln 2$ to compensate. In a simple system with single heat and work reservoirs, this operation would involve a transfer of at least $k_B T \ln 2$ of energy from the work reservoir to the heat reservoir via the bit.

Eventually, Bennett drew upon Landauer's work to argue that measurement, or copying of a bit, could be performed with no work input, provided that the memory was in a well-defined initial state (say, 0) (Bennett 1982, 2003). In this case, the measurement is logically reversible, since two initial states of the memory and data are mapped to two final states, as shown in Fig. 2. Such a measurement requires no net increase in the entropy of the environment, and hence no work input. However, to operate again the memory must be reset. If the "data" bit has already

been exploited at this point, for example by a demon extracting work, it no longer retains its previously measured value. The reset cannot then be performed in a logically reversible manner – we must perform a logically irreversible "erase" (Fig. 2) to reset the memory. This reset must increase the entropy of the environment by Landauer’s principle, and hence the work extracted ($\mathcal{W}_{\text{out}} \leq k_{\text{B}} T \ln 2$) is paid for by the resetting of the device ($\mathcal{W}_{\text{in}} \geq k_{\text{B}} T \ln 2$). Bennett argued that this erasure cost explains why the second law is not violated.

Within the wider scientific community, Bennett’s argument is generally accepted as the “true” explanation of why the second law survives (e.g. Feynman 1998; Plenio and Vitelli 2001; Bub 2001; Dillenschneider and Lutz 2009; Maruyama, Nori, and Vedral 2009; Bérut et al. 2012; Jun, Gavrilov, and Bechhoefer 2014). Although there are exceptions to this perspective (Fahn 1996; Maroney 2005; Sagawa and Ueda 2009; Sagawa 2014; Parrondo, Horowitz, and Sagawa 2015), Szilard’s argument that “measurement cannot take place without compensation” has largely been discounted, with measurement viewed as potentially costless, and the cost of erasure viewed as fundamental to resolving the paradox. Here, we argue against this perspective, leveraging a concrete molecular mechanism for implementing the entire measurement and feedback cycle of Szilard’s engine. We use this mechanism to demonstrate that, in fact, Szilard’s emphasis on the cost of correlating degrees of freedom, which corresponds to creating a non-equilibrium state, is essentially correct. We will simultaneously highlight fundamental issues that follow from an over-emphasis on the significance of erasure, and the stages at which work and heat are exchanged with the environment.

Our explicit molecular mechanism overcomes the inherent inscrutability of Maxwell’s demon. Previous work has included both physical and thought experiments on systems that perform erasure, measurement and exploitation (Landauer 1961; Szilard 1964; Bennett 1982; Sagawa and Ueda 2009; Lambson, Carlton, and Bokor 2011; Bérut et al. 2012; Jun, Gavrilov, and Bechhoefer 2014; Koski et al. 2014; Hong et al. 2016; Ouldridge,

Govern, and ten Wolde 2017). However, it is extremely rare to study a full cycle of Szilard’s engine in which all components and information-processing mechanisms are concrete and explicitly accounted for (we leverage our recent work in Brittain, Jones, and Ouldrige 2018). Based on our analysis of a molecular Szilard engine, we will make the case for explicit molecular systems as a natural environment in which to embed thermodynamic thought experiments more generally.

4 Implementing Szilard’s engine with molecular reactions

Well-mixed, dilute molecular systems fit naturally into the paradigm of Section 2 (Ouldrige 2018). The state of such systems is given by the number of each molecular species present. Changes in state are chemical reactions such as $A + B \rightarrow C + D$, which happen at rates determined by reactant concentrations. Some of these reactions may correspond to large-scale conformational changes of a single molecule. “Chemical Reaction Networks” (CRNs) of this kind have been extensively studied in the biophysics, mathematics and computer science literature (Krishnamurthy and Smith 2017), and provide a realistic model of many biological processes. Furthermore, it has been shown that arbitrary CRNs can be approximated in the laboratory through nucleic acid-based designs (Qian, Soloviechick, and Winfree 2011; Chen et al. 2013; Srinivas et al. 2017).

The discrete nature of the state space at the CRN level introduces a minor modification to the theory of Section 2. These “macrostates” each contain an enormous number of microscopic configurations. Macrostates are then stabilised not only by having low energy, but also by having many microscopic configurations (an “entropic” contribution). Thankfully, we can simply work at the level of macrostates, and replace the energies of individual states $\epsilon(x)$ appearing in Eqs. 2-3 with chemical free energies of macrostates $f_{\text{chem}}(m)$, which incorporate both energetic and




Figure 3. *Molecular Buffers. (a) A single molecule in a large bath of fuel molecules of two types, A and B, at a temperature T . (b) A system of interest coupled to one of a series of buffers at different concentrations of fuel molecules.*

entropic contributions.

$$\begin{aligned}
 f_{\text{chem}}(m) &= k_B T \ln \sum_{x \in m} \exp(-\epsilon(x)/k_B T), & p_{\text{eq}}(m) &\propto \exp(-f_{\text{chem}}(m)/k_B T), \\
 \mathcal{F}[p(m)] &= \sum_m p(m) f_{\text{chem}}(m) + k_B T \sum_m p(m) \ln p(m) = \langle f_{\text{chem}} \rangle - T S[p(m)] \\
 &= \mathcal{F}[p_{\text{eq}}(m)] + k_B T \sum_m p(m) \ln \frac{p(m)}{p_{\text{eq}}(m)} \geq \mathcal{F}[p_{\text{eq}}(m)].
 \end{aligned} \tag{4}$$

In practice, $f_{\text{chem}}(m)$ replaces the energies of the individual configurations. For a detailed discussion of the macrostate perspective, see Ouldrige 2018.

We will consider large baths of "fuel" molecules, or buffers; our systems will inter-convert these fuel molecules (Seifert 2011). The change in chemical free energy of a system when an additional molecule of A is added is called the *chemical potential* of A, μ_A . For dilute systems (Ouldrige 2018; Nelson 2004),

$$\mu_A = \mu_A^0 + k_B T \ln([A]/[1\text{M}]), \tag{5}$$

where μ_A^0 is the (theoretical) value of the chemical potential at a reference concentration of 1M, and $[A]$ is its actual concentration. If a reaction consumes a single A, and produces a single B, then there is a contribution

to the overall chemical free energy change of

$$\Delta f_{\text{chem}}(A \rightarrow B) = \mu_B - \mu_A = \mu_B^0 - \mu_A^0 + k_B T \ln([A]/[B]). \quad (6)$$

For example, consider the reaction $A + C \rightleftharpoons B + D$. A large, negative value of $\Delta f_{\text{chem}}(A \rightarrow B)$ would tend to favour the conversion of C into D , since macrostates with lower chemical free energy are more favourable.

Buffers as a source of chemical work

We will assume that the buffers are sufficiently large that any reactions have a negligible effect on the probability distribution of bath macrostates. In this limit, the baths become reservoirs of "chemical work", analogous to the (less concrete) work reservoirs introduced in Section 2.1. In this case, we now use m to label the macrostate of only the molecules that are not part of the buffers (the "system of interest"), with $p(m)$ describing the distribution of these macrostates. The change in free energy due to an arbitrary process is then

$$\begin{aligned} \Delta \mathcal{F} &= \Delta \left(\sum_m p(m) f_{\text{chem}}^{\text{sys}}(m) \right) + k_B T \Delta \left(\sum_m p(m) \ln p(m) \right) + \langle \sum_n \Delta f_{\text{chem}}^n \rangle \\ &= \Delta \langle f_{\text{chem}}^{\text{sys}} \rangle - T \Delta \mathcal{S}^{\text{sys}}[p(m)] + \langle \Delta f_{\text{chem}}^{\text{buffers}} \rangle = \Delta \mathcal{F}^{\text{sys}} - \mathcal{W}_{\text{chem}}, \end{aligned} \quad (7)$$

where $f_{\text{chem}}^{\text{sys}}(m)$ is the chemical free energy of the system of interest only, and $\langle \Delta f_{\text{chem}}^n \rangle$ is the expected change in chemical free energy of buffer n . The chemical work is defined as, $-\langle \Delta f_{\text{chem}}^{\text{buffers}} \rangle = -\langle \sum_n \Delta f_{\text{chem}}^n \rangle = \mathcal{W}_{\text{chem}}$; a positive value allows an increase in the free energy of the system of interest whilst satisfying $\Delta \mathcal{F} \leq 0$ overall.


By varying the chemical potentials of fuel molecules, we can implement a time-varying protocol on the system of interest (Ouldrige, Govern, and ten Wolde 2017; Rao and Esposito 2016; Schmiedl and Seifert 2007). A device for implementing such a protocol is illustrated in Fig. 3 (b). We have a few molecules of interest that are trapped inside a small reaction volume. These molecules could be tethered to the surface of the volume, or trapped by a membrane. Outside of the volume is a large bath

of fuel molecules, which can diffuse into the volume, setting the chemical potential and biasing reactions as desired. We implement a time-dependent protocol by connecting a series of buffers with different fuel concentrations to the system of interest - either in a linear fashion as in Fig. 3 (b) or in a cycle in order to implement a cyclic protocol.

Despite the explicit apparatus, it is still possible to have implicit decision-making; a demon could move the buffers in a particular fashion following a measurement, for example. We wish to eliminate such behaviour, and make all information-processing and decision-making explicit in the system of interest, thereby avoiding conceptual pitfalls. We therefore demand that any applied protocol incorporates no external feedback based on the state of the system of interest. One way to view the "no external feedback" criterion is that *a protocol applied to a single system can be trivially applied in parallel to many systems*. This idea marginalizes all costs of applying even complex protocols. In principle, if the buffers are moved slowly enough, there is no lower bound on the mechanical cost of moving them, given pistons of arbitrarily low friction. If we act on many systems in parallel, then any external cost associated with implementing the protocol can be taken as an edge effect. Further, if the differences in fuel concentration in consecutive buffers tends to zero, so does any thermodynamic cost associated with fuel molecules being exchanged between adjacent buffers via the volume of interest. In this limit, therefore, all thermodynamic costs arise from actual chemical reactions involving the system of interest. We now illustrate systems and protocols that implement erasing, copying and Szilard's engine.

Erasing a bit

A molecule with two long-lived states is a natural analog of a bit (Ouldrige, Govern, and ten Wolde 2017). Let X and X^* be two states of a molecule, and let us assume that it is possible to switch between X and X^* by coupling to fuel molecules F_1 and F_1^* , and a catalyst Y :



The catalyst Y must be present for the reaction to proceed at an appreciable rate, but is not consumed by the reaction. Reactions of this type are extremely common in natural systems (for example, Y could be a kinase, X its substrate, F_1^* , ATP and F_1 , ADP). They can also be engineered as nucleic acid strand displacement reactions (Qian, Soloviechick, and Winfree 2011; Chen et al. 2013; Srinivas et al. 2017). The chemical free energy change associated with the forwards reaction in Eq. 8 is

$$\begin{aligned}\Delta f_{\text{chem}}(X + F_1^* + Y \rightarrow X^* + F_1 + Y) &= \Delta\mu_{F_1^* \rightarrow F_1} + \Delta\mu_{X \rightarrow X^*}^0 \\ &= \mu_{F_1}^0 - \mu_{F_1^*}^0 + k_B T \ln \frac{[F_1]}{[F_1^*]} + \mu_{X^*}^0 - \mu_X^0.\end{aligned}\quad (9)$$

Here, $\Delta\mu_{X \rightarrow X^*}^0$ is the intrinsic stability difference between the two states (the $\ln[X]/[X^*]$ term is absent because we have a single molecule of X/X^*). For simplicity, we shall assume that $\Delta\mu_{X \rightarrow X^*}^0 = 0$, equivalent to assuming the bit is symmetric, as is typical. In this case, all changes in chemical free energy arise from the fuel molecules. The fundamental conclusions of this chapter are not affected by this assumption.

We now implement erasure on X/X^* using the apparatus outlined in Fig. 4(a) (Ouldrige, Govern, and ten Wolde 2017). Let a single Y molecule and a single X/X^* be confined in the volume. Initially, we start with buffers with both $[F_1], [F_1^*] \rightarrow 0$, so both X and X^* are essentially stable. We then slowly increase both $[F_1]$ and $[F_1^*]$ to a point where the forwards and backwards reactions occur at appreciable rates, maintaining a ratio of $[F_1]/[F_1^*]$ that ensures that $\Delta\mu_{F_1^* \rightarrow F_1} = 0$. We now see interchanges of X and X^* , but do not yet bias the bit's state. In the energy landscape representation shown in Fig. 4, this initial manipulation corresponds to lowering the barrier between macrostates prior to quasi-statically destabilizing one with respect to the other. Next, we slowly increase $[F_1]$ whilst keeping $[F_1^*]$ fixed, thereby increasing $\Delta\mu_{F_1^* \rightarrow F_1}$. A large and positive value of $\Delta\mu_{F_1^* \rightarrow F_1}$ pushes X^* towards X . When we have sufficiently large $\mu_{F_1^*} - \mu_{F_1} \gg k_B T$ that the probability of finding X^* is essentially zero, we start to decrease both $[F_1]$ and $[F_1^*]$ towards zero, whilst maintaining a constant $\mu_{F_1^*} - \mu_{F_1} \gg k_B T$, raising the barrier and "freezing"

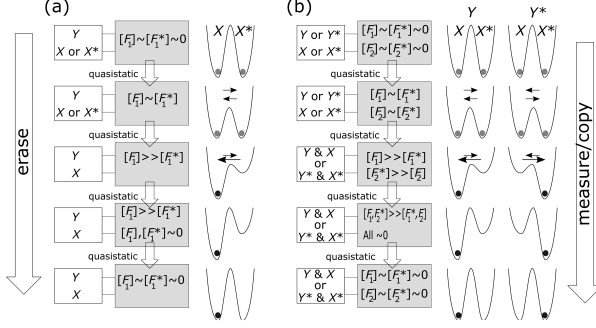


Figure 4. Reversible erase and copy protocols can be implemented in molecular systems. (a) Erasure of a molecular bit X/X^* to the X state. The system of interest (far left), containing X/X^* and catalyst Y , is coupled to a series of buffers with varying concentrations of fuel molecules; we indicate key stages in this process. $[F_1] \sim [F_1^*]$ means that $\Delta\mu_{F_1 \rightarrow F_1^*} = 0$; $[F] \sim 0$ means that the concentration is so low reactions requiring $[F]$ are negligible. The landscapes on the right hand represent the effective free energy landscapes for the interconversion of X and X^* at a given buffer and coupled to the catalyst Y , using an ideal reaction coordinate. Arrows indicate the dominant reaction, if any, and shaded balls indicate the probability of occupying a given state. (b) As in (a), but for copying of a Y/Y^* bit using two pairs of fuel molecules. Landscapes are now shown for both possibilities of the data bit Y/Y^* .

our molecular bit in the X state. Finally, we reduce $[F_1]$ further so that $[F_1], [F_1^*] \rightarrow 0$ and $\mu_X = \mu_{X^*}$. The original buffer is now re-connected, but the initially uncertain bit has been set to a guaranteed X state.

Overall, the free energy change of the reaction volume itself is $\Delta\mathcal{F}^{\text{sys}} = k_B T \ln 2$, since both X and X^* have the same chemical free energy, but the entropy $\mathcal{S}^{\text{sys}}[p(m)] = -k_B \sum_m p(m) \ln p(m)$ decreases from $k_B T \ln 2$ to 0.

The cost of erasure is supplied by the buffers. Initially, when $[F_1]$ and $[F_1^*]$ are increased simultaneously whilst maintaining $\mu_{F_1} = \mu_{F_1^*}$, no reactions occur on average. Subsequently, $[F_1]$ is increased relative to $[F_1^*]$, and net reactions do occur, consuming chemical free energy and performing chemical work. If the buffers are updated sufficiently slowly, then the X/X^* system reaches a new buffer-dependent equilibrium at each point

(Ouldrige, Govern, and ten Wolde 2017). From Eqs. 4 and 9, the equilibrium probability of observing X in each buffer is

$$p(X) = \frac{1}{1 + \exp(-\Delta\mu_{F_1^* \rightarrow F_1}/k_B T)}. \quad (10)$$

When the system is coupled to the next buffer, there is an infinitesimal change $d\Delta\mu_{F_1^* \rightarrow F_1}$. Consequently, $p(X)$ changes by

$$\begin{aligned} dp(X) &= \frac{1}{1 + \exp(-(\Delta\mu_{F_1^* \rightarrow F_1} + d\Delta\mu_{F_1^* \rightarrow F_1})/k_B T)} - \frac{1}{1 + \exp(-\Delta\mu_{F_1^* \rightarrow F_1}/k_B T)} \\ &= \frac{\exp(-\Delta\mu_{F_1^* \rightarrow F_1}/k_B T)}{\left(1 + \exp(-\Delta\mu_{F_1^* \rightarrow F_1}/k_B T)\right)^2} \frac{d\Delta\mu_{F_1^* \rightarrow F_1}}{k_B T}. \end{aligned} \quad (11)$$

When X^* changes to X , F_1 changes to F_1^* in the buffer, with an associated $\Delta f_{\text{chem}}^{\text{buffers}} = -\Delta\mu_{F_1^* \rightarrow F_1}$. The expected drop in free energy, or chemical work done, by the change $dp(X)$ is

$$d\mathcal{W}_{\text{chem}}^{\text{in}} = \Delta\mu_{F_1^* \rightarrow F_1} dp(X) = \Delta\mu_{F_1^* \rightarrow F_1} \frac{\exp(-\Delta\mu_{F_1^* \rightarrow F_1}/k_B T)}{\left(1 + \exp(-\Delta\mu_{F_1^* \rightarrow F_1}/k_B T)\right)^2} \frac{d\Delta\mu_{F_1^* \rightarrow F_1}}{k_B T}. \quad (12)$$

To find the total expected work done by the all of the buffers, we simply integrate over the whole protocol: $\Delta\mu_{F_1^* \rightarrow F_1} = 0$ to $\Delta\mu_{F_1^* \rightarrow F_1} \rightarrow \infty$. Using the substitution $u = \Delta\mu_{F_1^* \rightarrow F_1}/k_B T$, we obtain


$$\frac{\mathcal{W}_{\text{chem}}^{\text{in}}}{k_B T} = \int_0^\infty \frac{u \exp(-u)}{(1 + \exp(-u))^2} du = \int_0^\infty -\frac{\exp(u)}{1 + \exp(u)} + \frac{d}{du} \left(\frac{u \exp(u)}{1 + \exp(u)} \right) du = \ln 2. \quad (13)$$

The chemical work done to erase our molecular bit perfectly is $k_B T \ln 2$, as expected. The process is thermodynamically reversible; applying the buffers in reverse would return the system to the initial state, and the overall change in the non-equilibrium free energy is zero ($\mathcal{W}_{\text{chem}}^{\text{in}} = \Delta\mathcal{F}^{\text{sys}}$).

Copying or measuring a bit

To allow for copying, we consider two states of the catalyst, Y and Y^* . We shall label the catalyst as the "data" molecule or data bit. Moreover,

we shall require that Y and Y^* separately catalyse reactions that couple the reactions of the measurement molecule or memory bit $X \rightleftharpoons X^*$ to two distinct fuels (Ouldrige, Govern, and ten Wolde 2017):



Such a scheme approximates natural "bifunctional kinases" (Stock, Robinson, and Goudreau 2000), and could easily be engineered with nucleic acid circuitry (Qian, Soloviechick, and Winfree 2011; Chen et al. 2013; Srinivas et al. 2017).

Performing a perfect copy corresponds to correlating the data and measurement molecules. If we knew for certain that the data molecule was in the Y state, we could simply repeat the protocol outlined in Section 4.2. However, we also need the system to automatically handle the possibility that the data molecule is in the Y^* state via a single protocol. If we follow the protocol of Section 4.2 and Fig. 4(a), but with the data molecule in the Y^* state, nothing would happen, because reactions require the presence of either F_1, F_1^* and Y or F_2, F_2^* and Y^* . We can therefore implement two erasures simultaneously as a single protocol: driving the system to X^* if Y^* is present, and to X if Y is present. Such a protocol is illustrated in Fig. 4(b). The buffer concentrations of $[F_2]$ and $[F_2^*]$ mirror those of $[F_1]$ and $[F_1^*]$, in that $[F_2]$ and $[F_2^*]$ are initially increased at $\Delta\mu_{F_2^* \rightarrow F_2} = 0$, but then $[F_2^*]$ is increased further so that $-\Delta\mu_{F_2^* \rightarrow F_2} \gg k_B T$, and the system is driven to towards the X^* state if Y^* is present.

Calculating the chemical work done during the copy is trivial. We assume for simplicity that there is a 50% chance that Y is present, and a 50% chance that Y^* (similar arguments can be constructed for any initial bias). In either case, the decrease in chemical free energies of the buffers is $k_B T \ln 2$, and so $\mathcal{W}_{\text{chem}}^{\text{in}} = k_B T \ln 2$. This decrease in free energy exactly matches the increase of \mathcal{F}^{sys} . As in Section 4.2, the chemical free energy of the system of interest is unchanged, but $S[p(m)]$ has decreased from $k_B \ln 4$ to $k_B \ln 2$ – we have gone from four equally probable macrostates to just two due to the correlations. We have paid chemical work to decrease


the entropy of the data/measurement molecule system, thereby storing free energy.

Crucially, the protocol outlined above is true measurement in the sense discussed by Szilard. The two molecules' states are correlated without direct interaction, and were the catalyst (the data bit) to undergo a state change, there would be no tendency for the measurement molecule (memory bit) to change. We have previously described this type of copying as *persistent* (Ouldrige, Govern, and ten Wolde 2017).

A molecular Szilard engine

Having implemented a molecular copy (or measurement), we now exploit the measurement through feedback. The essential challenge is to allow the data Y/Y^* to evolve so that it decorrelates with the fixed memory X/X^* , releasing the free energy stored in correlations in a controlled way so that chemical work can be extracted. Decorrelation could also be performed by undoing the measurement: here, Y/Y^* does not change but X/X^* is randomized. This process, which doesn't involve feedback, can be achieved reversibly by reversing the copy protocol.

Feedback is normally implicit in Szilard's engine; it is difficult to imagine concrete systems in which one degree of freedom influences the other during an initial period, and then this influence is reversed. Fortunately, our chemical system allows this behaviour. On top of the reactions used during the copy (Eqs. 8 and 14), we consider catalytic interconversion of Y/Y^* by X/X^* using the fuel molecules. G_1 , G_2 , G_1^* and G_2^* :



The system is now a mutual bi-functional catalyst network. Such a system is unusual, but by no means unimaginable.

The overall protocol of our molecular Szilard engine is illustrated in Fig. 5 (a). The initial measurement, or correlation stage, proceeds exactly as in Section 4.3. The concentration of fuel molecules $[G_1], [G_1^*], [G_2], [G_2^*] \rightarrow$

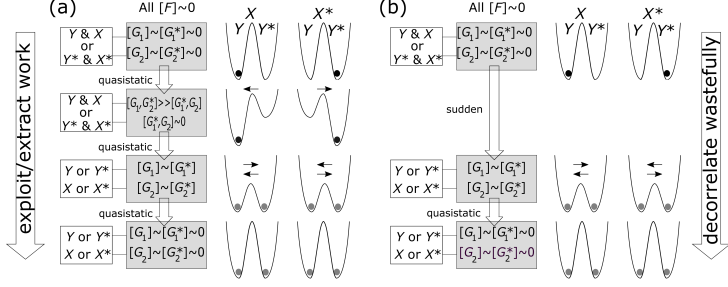


Figure 5. The work extraction step of a Szilard engine can be implemented in a molecular system. (a) Work extraction step following the copy/measurement step in Fig 4 (b). The final state is equivalent to the initial state of the measurement protocol, so combining the two produces a full cyclic operation. The combined data and measurement molecules are exposed to a series of buffers of varying fuel concentrations; key points are illustrated. $[G_1] \sim [G_1^*]$ means that $\Delta\mu_{G_1^* \rightarrow G_1} = 0$; $[G] \sim 0$ means that the concentration is so low reactions requiring $[G]$ are negligible. The landscapes on the right hand represent the effective free energy landscapes for the interconversion of Y and Y^* at a given buffer and coupled to either X of X^* acting as a catalyst. Arrows indicate the dominant reaction, if any, and shaded balls indicate the probability of occupying a given state. (b) Relaxation of the correlated state without careful extraction of chemical work during the decorrelation.

0, preventing the state of the data molecule from changing during this period. This part of the protocol requires $\mathcal{W}_{\text{chem}}^{\text{in}} = k_B T \ln 2$.

The subsequent exploitation stage proceeds as follows. First, $[G_1]$ and $[G_2^*]$ are increased whilst maintaining $[F_1], [F_1^*], [F_2], [F_2^*], [G_1^*], [G_2] \rightarrow 0$, until $\Delta\mu_{G_1^* \rightarrow G_1}, -\Delta\mu_{G_2^* \rightarrow G_2} \gg k_B T$. Since only (Y, X) and (Y^*, X^*) combinations are possible at this point, no reactions occur and no chemical work is done. Subsequently, $[G_1^*]$ and $[G_2]$ are slowly increased until $\Delta\mu_{G_1^* \rightarrow G_1} = \Delta\mu_{G_2^* \rightarrow G_2} = 0$. During this period, the data is decorrelated from the measurement, and free energy is transferred back to the buffers. Given a measurement X , the data molecule is in state Y^* with probability

$$p(Y^*|X) = \frac{\exp(-\Delta\mu_{G_1^* \rightarrow G_1}/k_B T)}{1 + \exp(-\Delta\mu_{G_1^* \rightarrow G_1}/k_B T)} \quad (16)$$

when coupled to a buffer at $\Delta\mu_{G_1^* \rightarrow G_1}$. Proceeding as in Section 4.2, we multiply the increase in $p(Y^*|X)$ associated with a buffer increment of $d\Delta\mu_{G_1^* \rightarrow G_1}$ by the resulting free energy change of the buffer for a conversion of Y into Y^* , $\Delta\mu_{G_1^* \rightarrow G_1}$, and integrate:

$$\Delta f_{\text{chem}}^{\text{buffers}} = - \int_{\infty}^0 \Delta\mu_{G_1^* \rightarrow G_1} \frac{\exp(-\Delta\mu_{G_1^* \rightarrow G_1}/k_B T)}{(1 + \exp(-\Delta\mu_{G_1^* \rightarrow G_1}/k_B T))^2} \frac{d\Delta\mu_{G_1^* \rightarrow G_1}}{k_B T}. \quad (17)$$

The integral is identical to Eq. 13, and thus $\Delta f_{\text{chem}}^{\text{buffers}} = k_B T \ln 2$. The same $\Delta f_{\text{chem}}^{\text{buffers}}$ is obtained for a measurement result of X^* , although it is the combined free energy of the G_2/G_2^* fuels that increases. Thus the average change in the total chemical free energies of the buffers is $\mathcal{W}_{\text{chem}}^{\text{in}} = -k_B T \ln 2$, and so the chemical work done in performing the measurement is extracted. Finally, no chemical work is done as the concentrations of all fuels are taken to zero at fixed $\Delta\mu_{G_1^* \rightarrow G_1} = \Delta\mu_{G_2^* \rightarrow G_2} = 0$. This final operation corresponds to "raising the barrier" between the metastable states of the bit in the energy landscape analog, as shown in Fig. 5.

Our molecular device is an exact analogue of the engine proposed by Szilard, with all components explicit. We perform a binary measurement by creating a record of the data molecule with the measurement molecule. Chemical work of $k_B T \ln 2$ is required by the measurement step, in order to create a low-entropy, high-free energy correlated state. This work is then extracted from these correlations during exploitation, when the measurement molecule itself couples the data molecule to the appropriate fuel reservoir. Importantly, the measurement is persistent in that the X/X^* state does not change during exploitation.

5 What we learn from explicit biochemical protocols

Erasure is not the most compelling explanation for the self-consistency of the second law

There is no erasure in the protocol above. In principle, we could add one by including a third molecule Z , which couples the transition of $X \rightleftharpoons X^*$

to another pair of fuel molecules H and H^* . By implementing a protocol equivalent to that in Section 4.2 after the exploitation step, we could reversibly set the measurement molecule to state X , at a cost of $k_B T \ln 2$ of chemical work. Having paid for the erasure, the minimal chemical work required to go from the erased state to the post-measurement, correlated state would now be zero since $\Delta \mathcal{F}^{\text{sys}} = 0$; in both cases there are two equally likely macrostates for the measurement/data system ($(d, m) = (0, 0)$ OR $(1, 0)$ and $(0, 0)$ OR $(1, 1)$, respectively).

Although cycles involving erasure are possible, it is not required, and so it seems slightly perverse to attribute the self-consistency of the second law in Szilard's engine to erasure (Maroney 2005). Indeed, in the simplest implementation of measurement with an erased bit, one first "un-erases" the measurement molecule to an equal probability of X and X^* , prior to even coupling to Y/Y^* via the F fuels. During this un-erase, $k_B T \ln 2$ of chemical work is extracted, giving zero overall chemical work for measurement from an erased state when combined with the $k_B T \ln 2$ required to subsequently create correlations. Erasing then appears to be an accounting trick; we add a process to the cycle (erase and un-erase) that achieves nothing, but then incorporate the un-erase into a "measurement from an erased state" process (Fig. 2)) that overall has no cost. We note that Bennett's analysis also includes an "un-erase" prior to correlation (Bennett 1982, 2003). It is possible to incorporate erasure in a less trivial manner, but we still find the emphasis on any one step to be misleading. The second law is not violated at *any* stage of the protocol, whether erasure is performed or not (Fahn 1996; Maroney 2005; Sagawa 2014; Parrondo, Horowitz, and Sagawa 2015).

Instead, it is more illuminating to focus on the non-equilibrium free energy. As originally pointed out by Szilard, the measured state contains correlations between degrees of freedom that must persist despite an absence of direct interactions. The existence of correlations without direct interactions implies a low entropy state without a compensating low chemical free energy (Fahn 1996; Sagawa 2014; Parrondo, Horowitz,

and Sagawa 2015). In this case, the correlations mean only two of a total of four equally stable macrostates are possible; the entropy of the post-measurement state is $k_B \ln 2$ rather than $k_B \ln 4$ and an excess free energy of $\mathcal{F}^{\text{sys}} - \mathcal{F}_{\text{eq}}^{\text{sys}} = k_B T \ln 2$ is available. The stored free energy in the measured state makes the exploitation step consistent with the second law, regardless of the details of the cycle. The only difference between protocols with and without erasure is whether we chose to compensate for the creation of the measurement by initially creating an out-of-equilibrium memory bit.

This explanation of Szilard’s engine generalizes. In information theory, correlations between two degrees of freedom m_1 and m_2 are described using the mutual information (Shannon and Weaver 1949)

$$\mathcal{I}[p(m_1), p(m_2)] = \sum_{m_1, m_2} p(m_1, m_2) \ln \left(\frac{p(m_1, m_2)}{p(m_1)p(m_2)} \right) \geq 0, \quad (18)$$

with the equality if the two degrees of freedom are independent. It is relatively straightforward to show that if the energy/chemical free energy is additive in m_1 and m_2 (the two degrees of freedom do not interact), then (Esposito and Van den Broeck 2011; Horowitz, Sagawa, and Parrondo 2013; Parrondo, Horowitz, and Sagawa 2015)

$$\mathcal{F}[p(m_1), p(m_2)] = \mathcal{F}[p(m_1)] + \mathcal{F}[p(m_2)] + k_B T \mathcal{I}[p(m_1), p(m_2)]. \quad (19)$$

Since m_1 and m_2 must be independent in equilibrium, a non-zero $\mathcal{I}[p(m_1), p(m_2)]$ necessarily represents a store of free energy. This principle has been applied to systems that create and exploit correlations between non-interacting degrees of freedom (Horowitz, Sagawa, and Parrondo 2013; Ouldrige, Govern, and ten Wolde 2017; McGrath et al. 2017; Ouldrige and ten Wolde 2017; Brittain, Jones, and Ouldrige 2018; Stopnitzky et al. 2018; Chapman and Miyake 2015; Boyd, Mandal, and Crutchfield 2016).

Bennett’s argument that, in a setting with heat and work reservoirs coupled to a pair of bits, measurement can be performed with no work input using an erased memory, and that erasing memories requires work input, is not incorrect. We nonetheless believe that the emphasis on erasure

causes unnecessary confusion. Taken at face value, the Bennett explanation might seem to suggest that the second law is transiently violated, since erasure compensates for the work extraction after the fact, when a whole cycle is completed (Landauer 1991). Indeed, the finite size of the demon's memory is sometimes invoked (Feynman 1998; Deffner 2015), suggesting that the second law can be violated until the memory is full and needs to be erased. But if the second law can be violated until memory has to be erased, then what happens if erasure never happens? If, instead, the necessarily non-equilibrium nature of the post-measurement state is used to support the second law, these confusions evaporate; all sub-steps are consistent with the second law (Parrondo, Horowitz, and Sagawa 2015; Sagawa 2014). Furthermore, the perceived importance of erasure leads researchers to look for its presence in systems that could perhaps be more fruitfully analysed in terms of the generation and usage of non-equilibrium correlations (Mehta and Schwab 2012; Andrieux and Gaspard 2013).

Focussing on heat and work can be misleading

There is a tendency to use heat generation as a shorthand for irreversibility, or the increase in the total entropy. This shorthand, however, can lead to significant confusion. Heat generation does not imply thermodynamic irreversibility. It is possible, at least in principle, to transfer heat to a system's environment in a thermodynamically reversible way, provided that the increase in the entropy of the environment is compensated by a decrease in the entropy of the system. A common example would be quasistatic compression of a piston. Heat transfer to the environment only guarantees irreversibility if the system operates cyclically, in which case its entropy cannot change to compensate. Nonetheless, the transfer of heat is frequently described as "dissipation", without clarifying the degree to which the total entropy actually increases (Maroney 2005).

From a technical perspective, the distinction is between entropy transferred and entropy produced. The entropy production is the increase in

total entropy. It is the sum of the change in entropy of the system of interest and the entropy transferred to the rest of the environment due to its interaction with the system. In the case where the system of interest is in contact with a single heat reservoir the entropy transferred is the heat transferred to the reservoir divided by the temperature (Par-rondo, Horowitz, and Sagawa 2015) – the heat transferred is therefore not directly informative of the total entropy production, even in simple settings settings with a single heat reservoir.

In Landauer’s analysis of erasure, heat generation compensates for the entropy decrease of the bit. Despite the fact that the bound of $k_{\text{B}}T \ln 2$ is explicitly derived for a reversible protocol satisfying $\Delta S^{\text{tot}} = 0$, Landauer himself describes the heat transfer as a necessarily thermodynamically irreversible process (Landauer 1961). Many others have apparently repeated this logic (Plenio and Vitelli 2001; Bub 2001; Ladyman et al. 2007; Dillenschneider and Lutz 2009; Bérut et al. 2012; Jun, Gavrilov, and Bechhoefer 2014; Hong et al. 2016). This confusion contributes to the erroneous narrative that thermodynamic irreversibility of erasure pays for a transient violation of the second law during a cycle of Szilard’s engine, and leads to statements such as "logical irreversibility is associated with physical irreversibility and requires a minimal heat generation" (Landauer 1961).

We defined thermodynamic reversibility in Section 2.3; it is a statement about whether a particular protocol that converts a system’s distribution from $p(x)$ into $p'(x)$ would recreate $p(x)$ when applied in a time reversed fashion to an initial distribution $p'(x)$. Thermodynamic reversibility depends on the actual physical process undergone to move from $p(x)$ to $p(x')$. However, exact trajectories $x \rightarrow x'$ do not need to be reversed; $p(x)$ need only be recreated in the statistical sense. Whether or not a process is thermodynamically reversible depends upon both the protocol and the initial distribution $p(x)$ (Kolchinsky and Wolpert 2017); in general, a protocol that is thermodynamically reversible for one $p(x)$ will not be reversible for another.

By contrast, a process is logically reversible if and only if the initial logical state y can be inferred unambiguously from the final logical state y' – which is not true for erasure, for example (Sagawa 2014). Logical irreversibility is then solely a condition on the overall transition matrix between input and output logical states $T_{y'y}$, and is independent of the initial state to which $T_{y'y}$ is applied; the possibility of multiple physical states x within a logical state y , and the detailed physical mechanism by which $T_{y'y}$ is achieved. Logical reversibility is, however, crucially dependent on the specific input/output map represented by $T_{y'y}$, rather than just the overall statistics. Given these differences between logical and thermodynamic irreversibility, it is not surprising that there is in fact no causal connection between the two; both logically reversible and logically irreversible processes can in principle be implemented in thermodynamically reversible and irreversible ways. Sagawa has given specific examples to illustrate this fact in general settings (Sagawa 2014).

Moreover, if a process as a whole is thermodynamically irreversible, one cannot immediately associate the part of the process that transfers heat to the environment (or requires work input) as the cause of the irreversibility. For example, irreversibility in certain systems has been ascribed to the need to erase (Mehta and Schwab 2012; Mehta, Lang, and Schwab 2016; Andrieux and Gaspard 2013). But if the erasure considered by Landauer is not necessarily thermodynamically irreversible, then it is a logical non-sequitur to attribute thermodynamic irreversibility to the presence of erasure. Landauer’s principle in isolation cannot explain an increase in the total entropy of a closed system. For example, consider the molecular Szilard engine including erasure as outlined in Section 5.1. This device is totally reversible. We could, however, consider an alternative protocol in which, instead of carefully relaxing the correlated state after the measurement, we simply suddenly expose the system to buffers with high concentrations of $[G_1]$, $[G_1^*]$, $[G_2]$ and $[G_2^*]$ with $\Delta\mu_{G_1^* \rightarrow G_1} = \Delta\mu_{G_2^* \rightarrow G_2} = 0$, as in Fig. 5(b), prior to erasing. The measurement and data molecules would immediately decorrelate, and

no chemical work would be extracted. The overall cycle would then require $\mathcal{W}_{\text{chem}} = k_B T \ln 2$, the total free energy of the buffers would decrease in a cycle, and the total entropy would increase. The same uncontrolled decorrelation would happen if we tried to erase directly from the correlated state using Z and fuels H and H^* , without gradually undoing the correlations first.

It is tempting to associate the irreversibility in this cycle with a step at which we have to expend net chemical work due to Landauer’s principle – erasure. In the setting of magnetic bits, this process would require work input and generate heat. However, in either context, erasure is thermodynamically reversible. The (chemical) work input is exactly compensated by the increase in \mathcal{F}^{sys} . Rather, it is the decorrelation step – which requires no (chemical) work, and exchanges no heat with the environment in either the molecular or magnetic contexts – that is the *cause* of the irreversibility (Fahn 1996; Ouldrige, Govern, and ten Wolde 2017). Reversing this step would not restore the correlated state, and during this step \mathcal{F}^{sys} decreases by $k_B T \ln 2$ due to the loss of correlations, without a compensating extraction of (chemical) work, implying an increase in the combined entropy of the system and environment.

The above discussion highlights a further issue; the second law doesn’t only apply to the mechanical heat engines that drove the industrial revolution; it also applies to the molecular systems that underlie life. In these settings, and others, mechanical work and heat can play a less direct role, with processes relying on other forms of free energy transduction. Indeed, the erasure, copying and Szilard engine protocols described in Section 4 involve no mechanical work at all. This distinction matters because, as discussed in Section 2.2, free energy is not an energy; it is not conserved. If *chemical* work $W_{\text{chem}}^{\text{in}}$ is done on a system during a cyclic evolution (in which $\Delta U = 0$), physical heat $Q^{\text{out}} = W_{\text{chem}}^{\text{in}}$ need not be transferred to the environment to compensate. Nonetheless, heat generation is often used as a shorthand for thermodynamic irreversibility, even in molecular contexts (see e.g. Lan et al. 2012; Mehta, Lang, and Schwab

2016; Govern and ten Wolde 2014; Barato and Seifert 2017), as noted in Seifert 2011. Heat flow is a real energy flux that can be measured; it is misleading to describe entropy generation in a different form in this way. For example, one might incorrectly assume that a living system needs to develop machinery to rid itself of heat exactly equal to $T\dot{S}$, where \dot{S} is its entropy generation rate, to maintain a constant temperature.

The benefits of having an explicit supply of chemical work

Work reservoirs are not usually explicitly modelled like the chemical work buffers in Section 4. Using abstract (and interchangeable) work buffers is often conceptually useful, but it can help to make certain results seem overly remarkable. For example, measurement can be performed without net (chemical) work expenditure if the memory is initially in a well-defined state, suggesting that the measurement has no cost. Given that the correlated post-measurement state stores free energy that can be extracted as work, this lack of a cost seems remarkable. However, there is a cost – the initial non-equilibrium state of the memory is consumed to pay for the measurement. At a fundamental physical level, there is no real difference between using free energy stored in the memory bit or a (chemical) work buffer. Either way, a resource is consumed to create non-equilibrium correlations between memory and data. There is no reason to value resources in a (chemical) work reservoir over and above sources of free energy explicitly incorporated into that system of interest, which is particularly clear when the work reservoir is made explicit.

An explicit supply of (chemical) work also helps establish what is definitely possible – a so-called "positive result". Work reservoirs are often pictured as masses in a gravitational field, as in Fig. 2. It is often imagined, however, that work can be done by implementing an arbitrary change in the potential energy of a system as a function of its internal coordinates (eg. Schmiedl and Seifert 2007), despite the fact that a single weight would actually be insufficient even for the simple Szilard engine of Fig.2. Without an explicit mechanism to achieve the required work

coupling, it is arguable whether one can positively prove the possibility of certain operations – even if one can prove the impossibility of others. Perhaps even more challenging than implementing an arbitrary control of a potential is the idea of recovering work back from the system of interest, as is required by Szilard’s engine (Deshpande et al. 2017).

The operations that can be achieved with the molecular setup outlined in Section 4 are clearly limited. However, they are at least definitely possible, albeit in a highly-idealised limit. Furthermore, since the interconversion of fuel species directly determines the chemical work done, any chemical work recovered is indeed stored. By contrast, consider recent experiments in which devices such as optical feedback traps are used to implement an arbitrary force on a colloid (Bérut et al. 2012; Jun, Gavrilov, and Bechhoefer 2014). Although this force is designed to mimic a conservative, time-dependent potential, the thermodynamic cost of applying the control is orders of magnitude higher than calculated, and work is not actually recovered when the colloid moves against the pseudo-potential.

6 *Biochemical systems as a platform for studying fully autonomous systems*

Although conceptually useful, the apparatus in Section 4 is highly idealised. We assume that molecular states are stable unless enzymes are present, and that the fuels only couple to the intended enzymes (in practice, lifetimes need to be long and leak rates need to be slow relative to intended reactions). Furthermore, although we did not consider intelligent intervention or implicit decision-making, it is still somewhat unsatisfying to invoke an external protocol at all. The most intellectually-satisfying description would include the device (or demon) implementing any protocol as part of the system itself (Mehta, Lang, and Schwab 2016).

Such systems, which do not involve an externally-applied time-varying control, are *autonomous*. Molecular systems are even more suited to exploring autonomous contexts than those with externally-applied protocols. Fixed buffers of non-equilibrium fuel concentrations can drive active processes. Crucially, diffusion is automatically present in molecular systems, allowing an arbitrary complex network of interactions at no cost. We have previously shown that a simple, autonomously-operating enzymatic network can be mapped directly to a copy process in the sense of Szilard’s engine (Ouldrige, Govern, and ten Wolde 2017). The network could not reach the bound set by the non-autonomous, quasistatic protocols considered in Section 4. It could come remarkably close, however, despite operating at a finite rate, and requiring less stringent assumptions about the slowness of unintended leak reactions.

7 Conclusion

Basing our discussion around molecular protocols for manipulating bits, we have demonstrated the advantages of being more explicit in our analysis of the fundamental thermodynamic and computational systems. We have argued that doing so removes much of the mystery surrounding the physics of computation and measurement, and thermodynamics more generally. We have shown that molecular realisations are a particularly helpful way to understand basic thermodynamic ideas.

We argue that it is time to revisit the generally accepted explanation for the self-consistency of the second law of thermodynamics during the operation of Maxwell’s demon and Szilard’s engine. Attributing the survival of the second law to one part of the process or another is an accounting exercise, and misses the fundamental point: the second law is self-consistent throughout all stages of operation. Nonetheless, we believe that Szilard’s original contention that creating a correlated state is fundamental to the device, and that creating this correlated state requires

"compensation", is essentially correct. Whether this cost is paid by exploiting a reservoir, or a supply of free energy that is part of the explicitly-modelled system, is moot. We emphasize that Bennett's analysis is not wrong; merely the significance of erasure in the context of the second law has been overstated. The lack of a relationship between logical irreversibility and thermodynamic reversibility more generally calls into question the value of pursuing logically reversible computing approaches in an effort to reduce power consumption (Grochow and Wolpert 2018), although tracking the amount of erasure in a human-designed computation is potentially useful for some specific approaches to low-cost computation. We should also not impose a language of erasure on (biological) systems for which it is unnatural.

We are not claiming to advance this view for the first time in this chapter, and we have drawn upon a number of works published throughout the last few decades (including Fahn 1996; Maroney 2005; Sagawa and Ueda 2009; Sagawa 2014; Parrondo, Horowitz, and Sagawa 2015). One might ask whether the development of new theories and tools since the middle of the 20th Century has led to this alternative rationalisation. On the one hand, modern methodology is not really necessary to understand the problem. Non-equilibrium generalized free energies and the language of information theory are a convenient tool that make Szilard's argument more formal, but the essential insight was already provided in 1929, and Landauer and Bennett's calculations themselves were not incorrect. Really it is a question of interpretation, and the development and application of stochastic thermodynamics (Jarzynski 1997; Crooks 1999; Esposito and Van den Broeck 2011; Seifert 2005) has facilitated a much deeper general understanding of information and thermodynamics for small, fluctuating systems in recent years. This general understanding has in turn yielded a clearer approach to this particular problem. We note that for many (including ourselves), this insight has been developed whilst studying concrete molecular systems that accomplish information-processing or computational tasks (Ouldrige, Govern, and

ten Wolde 2017).

We additionally make the case for more careful use of the terms heat and work – heat is not a synonym for an increase in the total entropy, and free energy is not an energy. The distinction between free energies and work touches on a further important point: fundamentally, chemical work buffers are reservoirs of free energy due to an entropic component to their deviation from equilibrium. Much has been made of the fact that a low-entropy string of erased bits is a thermodynamic resource (Feynman 1998; Mandal and Jarzynski 2012). However, entropic contributions to stored free energies are natural in molecular contexts. The effective force driving the diffusion of uncharged molecules from a high concentration to a low concentration is entirely entropic. Once the physical object encoding the low entropy state is made explicit, "information reservoirs" typically seem much less mysterious.

We have also argued that molecular systems are a convenient basis in which to ground explicit (if idealised) representations of processes that we wish to analyse at a fundamental level. Explicit mechanisms can demystify the underlying concepts, and without a method of implementing a protocol, it is difficult to argue for positive rather than negative results. Furthermore, without a concrete mechanism, it is easy to design systems that are not even thermodynamically well-defined – as discussed in Stopnitzky et al. 2018. Needless to say, despite our enthusiasm with molecular contexts and their relevance to complex natural systems, alternative paradigms may be more relevant in other settings.

Full autonomy is the natural limit of explicitness. Currently, there is a major lack of understanding of the possibilities of autonomous systems (Mehta, Lang, and Schwab 2016). Can they approach the performance of externally manipulated systems, particularly in the limit of finite time operation? How challenging is it to develop systems in which one subsystem effectively generates time-varying conditions for another, and is this effective? Molecular systems are both an important conceptual tool and an experimental testing-ground for exploring these ideas, particularly as

we move beyond the classic questions of copying and erasing.

8 Acknowledgements

We greatly thank Dave Doty, Manoj Gopalkrishnan and Nick Jones for helpful conversations.

References

- Adleman, L. M. 1994. "Molecular computation of solutions to combinatorial problems." *Science* 266:1021–1024.
- Andrieux, D., and P. Gaspard. 2013. "Information erasure in copolymers." *Europhys. Lett.* 103:30004.
- . 2008. "Nonequilibrium generation of information in copolymerization processes." *Proc. Nat. Acad. Sci. USA* 105:9516–9521.
- Atkins, P. 2010. *The laws of thermodynamics: A very short introduction*. Oxford University Press.
- Barato, A. C., D. Hartich, and U. Seifert. 2014. "Efficiency of cellular information processing." *New. J. Phys.* 16:103024.
- Barato, A. C., and U. Seifert. 2017. "Coherence of biochemical oscillations is bounded by driving force and network topology γ ." *Phys. Rev. E* 95:062409.
- Bennett, C. H. 2003. "Notes on Landauer's principle, reversible computation, and Maxwell's Demon." *Stud. Hist. Philos. M. P.* 34:501–510.
- . 1982. "The thermodynamics of computation - a review." *Int. J. Theor. Phys.* 21:905–940.
- Bérut, A., A. Arakelyan, A. Petrosyan, S. Ciliberto, R. Dillenschneider, and E. Lutz. 2012. "Experimental verification of Landauer's principle linking information and thermodynamics." *Nature* 483:187–189.

- Boyd, A. B., D. Mandal, and J. P. Crutchfield. 2016. "Identifying functional thermodynamics in autonomous Maxwellian ratchets." *New J. Phys.* 18:023049.
- Brillouin, L. 1951. "Maxwell's Demon cannot operate: Information and entropy. I." *J. Appl. Phys.* 22:334–337.
- Brittain, R. A., N. S. Jones, and T. E. Ouldridge. 2018. "Differing strategies for memory-limited inference and thermodynamic exploitation of structured molecular environments." *Submitted*.
- Bub, J. 2001. "Maxwell's Demon and the thermodynamics of computation." *Stud. Hist. Philos. M. P.* 32:569–579.
- Chapman, A., and A. Miyake. 2015. "How an autonomous quantum Maxwell demon can harness correlated information." *Phys. Rev. E* 92 (6): 062125.
- Chen, Y.-J., N. Dalchau, N. Srinivas, A. Phillips, L. Cardelli, D. Solveichik, and G. Seelig. 2013. "Programmable chemical controllers made from DNA." *Nat. Nanotechnol.* 8:755–762.
- Crooks, G. E. 1999. "Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences." *Phys. Rev. E* 60:2721–2726.
- DeBenedictis, E. P., J. K. Mee, and M. P. Frank. 2017. "The opportunities and controversies of reversible computing." *Computer* 50:76–80.
- Deffner, S. 2015. "Viewpoint: Exorcising Maxwell's Demon." *Physics* 8:127.
- Deshpande, A., M. Gopalkrishnan, T. E. Ouldridge, and N. S. Jones. 2017. "Designing the optimal bit: balancing energetic cost, speed and reliability." *Proc. Roy. Soc. A* 473 (2204).
- Dillenschneider, R., and E. Lutz. 2009. "Memory erasure in small systems." *Phys. Rev. Lett.* 102:210601.

- Dougal, R. C. 2016. "Kelvin, thermodynamics and the natural world." Chap. 5. Kelvin, Maxwell, Clausius and Tait: the correspondence of James Clerk Maxwell, edited by M. W. Collins, R. C. Dougal, C. Koenig, and I. S. Ruddock, 135–151. WIT Press.
- Esposito, M., and C. Van den Broeck. 2011. "Second law and Landauer principle far from equilibrium." *Europhys. Lett.* 95:40004.
- Fahn, P. N. 1996. "Maxwell's demon and the entropy cost of information." *Found. Phys.* 26:71–93.
- Feynman, R. P. 1998. *Feynman Lectures on Computation*. New York: Addison-Wesley Publishing Co., Inc.
- Frank, M. P. 2017. "Throwing computing into reverse." *IEEE Spectrum* 54:32–37.
- Govern, C. C., and P. R. ten Wolde. 2014. "Optimal resource allocation in cellular sensing systems." *Proc. Nat. Acad. Sci. USA* 111:17486–17491.
- Grochow, J. A., and D. H. Wolpert. 2018. "Beyond Number of Bit Erasures: New Complexity Questions Raised by Recently Discovered Thermodynamic Costs of Computation." *SIGACT News* 49:33–56.
- Hong, J., B. Lambson, S. Dhuey, and J. Bokor. 2016. "Experimental test of Landauer's principle in single-bit operations on nanomagnetic memory bits." *Sci. Adv.* 2:e1501492.
- Horowitz, J., and C. Jarzynski. 2008. "Comment on "Failure of the work-Hamiltonian connection for free-energy calculations". " *Phys. Rev. Lett.* 101:098901.
- Horowitz, J. M., T. Sagawa, and J. M. R. Parrondo. 2013. "Imitating chemical motors with optimal information motors." *Phys. Rev. Lett.* 111:010602.
- Huang, K. 1987. *Statistical Mechanics*. Second Edition. New York: John Wiley & Sons, Inc.

- Ito, S., and T. Sagawa. 2015. "Maxwell's demon in biochemical signal transduction with feedback loop." *Nat. Comm.* 6:7498.
- Jarzynski, C. 1997. "Nonequilibrium equality for free energy differences." *Phys. Rev. Lett.* 78:2690–2693.
- Jaynes, E. T. 1957. "Information theory and statistical mechanics." *Phys. Rev.* 106:620–630.
- Jun, Y., M. Gavrilov, and J. Bechhoefer. 2014. "High-precision test of Landauer's principle in a feedback trap." *Phys. Rev. Lett.* 113:190601.
- Kolchinsky, A., and D. H. Wolpert. 2017. "Dependence of dissipation on the initial distribution over states." *J. Stat. Mech.* 2017:083202.
- Koski, J. V., V. F. Maisi, T. Sagawa, and J. P. Pekola. 2014. "Experimental observation of the role of mutual information in the nonequilibrium dynamics of a Maxwell Demon." *Phys. Rev. Lett.* 113:030601.
- Krishnamurthy, S., and E. Smith. 2017. "Solving moment hierarchies for chemical reaction networks." *J. Phys. A* 50:425002.
- Ladyman, J., S. Presnell, A. J. Short, and B. Groisman. 2007. "The connection between logical and thermodynamic irreversibility." *Stud. Hist. Philos. M. P.* 38:58–79.
- Lambson, B., D. Carlton, and J. Bokor. 2011. "Exploring the thermodynamic limits of computation in integrated systems: Magnetic memory, nanomagnetic logic, and the Landauer limit." *Phys. Rev. Lett.* 107:010604.
- Lan, G., P. Sartori, S. Neumann, V. Sourjik, and Yv Tu. 2012. "The energy-speed-accuracy trade-off in sensory adaptation." *Nature Physics* 8:422–428.
- Landauer, R. 1991. "Information is physical." *Phys. Today* 44:23–29.
- . 1961. "Irreversibility and heat generation in the computing process." *IBM J. Res. Dev.* 5:183–191.

- Mandal, D., and C. Jarzynski. 2012. "Work and information processing in a solvable model of Maxwell's demon." *Proc. Nat. Acad. Sci. USA* 109:11641–11645.
- Maroney, O.J.E. 2005. "The (absence of a) relationship between thermodynamic and logical reversibility." *Stud. Hist. Philos. M. P.* 36:355 – 374.
- Maruyama, K., F. Nori, and V. Vedral. 2009. "Colloquium: The physics of Maxwell's demon and information." *Rev. Mod. Phys.* 81.
- McGrath, T., N. S. Jones, P. R. ten Wolde, and T. E. Ouldridge. 2017. "Biochemical machines for the interconversion of mutual information and work." *Phys. Rev. Lett.* 118:028101.
- Mehta, P., A. H. Lang, and D. J. Schwab. 2016. "Landauer in the age of synthetic biology: Energy consumption and information processing in biochemical networks." *J. Stat. Phys.* 162:1153–1166.
- Mehta, P., and D. J. Schwab. 2012. "Energetic costs of cellular computation." *Proc. Nat. Acad. Sci. USA* 109:17978–17982.
- Nelson, P. 2004. *Biological Physics: Energy, Information, Life*. New York, NY: W. H. Freeman / Co.
- Norton, J. D. 2005. "Eaters of the lotus: Landauer's principle and the re-turdemon Maxwell's demon." *Stud. Hist. Philos. M. P.* 36:375–411.
- Organick, L., S. D. Ang, Y.-J. Chen, R. Lopez, S. Yekhanin, K. Makarychev, M. Z. Racz, et al. 2018. "Random access in large-scale DNA data storage." *Nat. Biotechnol.* 36:242–248.
- Ouldridge, T. E. 2018. "The importance of thermodynamics for molecular systems and the importance of molecular systems for thermodynamics." *Nat. Comput.* 17:3–29.

- Ouldridge, T. E., C. C. Govern, and P. R. ten Wolde. 2017. “The thermodynamics of computational copying in biochemical systems.” *Phys. Rev. X* 7:021004.
- Ouldridge, T. E., and P. R. ten Wolde. 2017. “Fundamental costs in the production and destruction of persistent polymer copies.” *Phys. Rev. Lett.* 118:158103.
- Owen, J. A., A. Kolchinsky, and D. H. Wolpert. 2017. “Number of hidden states needed to physically implement a given conditional distribution.” *arXiv:1709.00765*.
- Parrondo, J. M., J. M. Horowitz, and T. Sagawa. 2015. “Thermodynamics of information.” *Nat. Phys.* 11:131–139.
- Plenio, M. B., and V. Vitelli. 2001. “The physics of forgetting: Landauer’s erasure principle and information theory.” *Contemp. Phys.* 42:25–60.
- Qian, L., D. Soloviechick, and E. Winfree. 2011. “Efficient Turing-universal computation with DNA polymers.” *DNA 16, LNCS* 6518:123–140.
- Rao, R., and M. Esposito. 2016. “Nonequilibrium thermodynamics of chemical reaction networks: Wisdom from stochastic thermodynamics.” *Phys. Rev. X* 6:041064.
- Sagawa, T. 2014. “Thermodynamic and logical reversibilities revisited.” *J. Stat. Mech.* P03025.
- Sagawa, T., and M. Ueda. 2009. “Minimal energy cost for thermodynamic information processing: Measurement and information erasure.” *Phys. Rev. Lett.* 102:250602.
- Schmiedl, T., and U. Seifert. 2007. “Stochastic thermodynamics of chemical reaction networks.” *J. Chem. Phys.* 126:044101.
- Seifert, U. 2005. “Entropy production along a stochastic trajectory and an integral fluctuation theorem.” *Phys. Rev. Lett.* 95:040602.

- Seifert, U. 2011. “Stochastic thermodynamics of single enzymes and molecular motors.” *Eur. Phys. J. E* 34:26.
- Shannon, C. E., and W. Weaver. 1949. *The Mathematical Theory of Communication*. University of Illinois Press.
- Srinivas, N., J. Parkin, G. Seelig, E. Winfree, and D. Soloveichik. 2017. “Enzyme-free nucleic acid dynamical systems.” *Science* 358:eaal2052.
- Stock, A. M., V. L. Robinson, and P. N. Goudreau. 2000. “Two-component signal transduction.” *Annu. Rev. Biochem.* 69:183–215.
- Stopnitzky, E., S. Still, T. E. Ouldridge, and L. Altenberg. 2018. “Physical limitations of work extraction from temporal correlations.” *Submitted*.
- Szilard, L. 1964. “On the decrease of entropy in a thermodynamic system by the intervention of intelligent beings.” *Behav. Sci.* 9:301–310.
- ten Wolde, P.R., N.B. Becker, T.E. Ouldridge, and Mugler A. 2016. “Fundamental limits to cellular sensing.” *J. Stat. Phys.* 162:1395.