

What's the worst that could happen? One-shot dissipated work from Rényi divergences

Nicole Yunger Halpern,¹ Andrew J. P. Garner,^{2,3} Oscar C. O. Dahlsten,² and Vlatko Vedral^{2,3}

¹*Institute for Quantum Information and Matter, Caltech, Pasadena, CA 91125, USA*

²*Atomic and Laser Physics, Clarendon Laboratory,
University of Oxford, Parks Road, Oxford OX13PU, United Kingdom*

³*Center for Quantum Technologies, National University of Singapore, Republic of Singapore*

(Dated: May 26, 2015)

Thermodynamics describes large-scale, slowly evolving systems. Two modern approaches generalize thermodynamics: *fluctuation theorems*, which concern finite-time nonequilibrium processes, and *one-shot statistical mechanics*, which concerns small scales and finite numbers of trials. Combining these approaches, we calculate a one-shot analog of the average dissipated work defined in fluctuation contexts: the cost of performing a protocol in finite time instead of quasistatically. The average dissipated work has been shown to be proportional to a relative entropy between phase-space densities, one between quantum states, and one between probability distributions over possible values of work. We derive one-shot analogs of all three equations, demonstrating that the order- ∞ Rényi divergence is proportional to the *maximum dissipated work* in each case. These one-shot analogs of fluctuation-theorem results contribute to the unification of these two toolkits for small-scale, nonequilibrium statistical physics.

Thermodynamics concerns large scales and infinitesimally slow evolutions. In the thermodynamic limit, a system's size approaches infinity and is typified by mean behaviors. Infinitesimally slow, *quasistatic*, processes are described with the free energy F , with temperature, and with other equilibrium quantities.

Two recently developed frameworks generalize thermodynamic concepts, such as work and heat, beyond slow processes and infinite sizes. *Fluctuation relations* interrelate equilibrium quantities such as F with nonequilibrium processes (e.g., [1–6]). *One-shot statistical mechanics* quantifies the efficiency with which work can be invested or extracted, not only on average as the number of trials approaches infinity, but also if few trials are performed (e.g., [7–11]). One-shot statistical mechanics grew from *one-shot information theory* (e.g., [13–16]), the study of entropies apart from Shannon's and von Neumann's [12], to describe protocols whose trials are not necessarily independent and identically distributed according to the same probability distribution or quantum state. A combination of fluctuation relations and one-shot statistical mechanics describes quite general thermodynamic systems [17].

Transforming one equilibrium state quasistatically into another requires an amount W of work equal to the difference between the states' free energies: $W = \Delta F$. Implementing a protocol in finite time yields a nonequilibrium state and costs extra work, some dissipated as heat. This penalty of irreversibility is called the *dissipated work*, or *irreversible work*. The average $\langle W_{\text{diss}} \rangle := \langle W \rangle - \Delta F$ over many trials has been studied in fluctuation contexts (e.g. [18–20]).¹ We define the *one-shot dissipated work* $W_{\text{diss}} := W - \Delta F$ as the penalty paid in one trial [21].

$\langle W_{\text{diss}} \rangle$ has been shown to be proportional to three instances of the *Kullback-Leibler (KL) divergence*, or *average relative entropy*, D . D quantifies how much two probability distributions, or two quantum states, differ. $\langle W_{\text{diss}} \rangle$ has been related to a D between phase-space densities $\rho(p, q, t)$ and $\tilde{\rho}(p, -q, t)$ [4], a D between quantum states $\rho(t)$ and $\tilde{\rho}(t)$ [22], and a D between probability distributions $P_{\text{fwd}}(W)$ and $P_{\text{rev}}(-W)$. We derive one-shot analogs of all three relationships.

Rényi divergences have recently appeared in fluctuation-relation contexts [23]. The latter work pertains specifically to resource theories, which we will not use. We follow the approach of [17], building on assumptions used to derive Crooks' Theorem.

We begin by reviewing fluctuation theorems and Rényi divergences, focusing on the one-shot *order- ∞ Rényi divergence* D_{∞} . We recall each $\langle W_{\text{diss}} \rangle$ proportionality and derive its one-shot analog. Our main results relate the maximum possible penalty $W_{\text{diss}}^{\text{worst}}$ of investing work in finite time to three instances of D_{∞} . Our one-shot analogs of fluctuation-relation results illustrate the insights offered by merging fluctuation relations with one-shot statistical mechanics.

Fluctuation theorems—Consider a system governed by a time-dependent Hamiltonian $H(\lambda_t)$. The external parameter λ_t changes in time: $t \in [-\tau, \tau]$. Suppose the system begins in the thermal state $\gamma_{-\tau} := e^{-\beta H(\lambda_{-\tau})}/Z_{-\tau}$, wherein β denotes a heat bath's inverse temperature and $Z_{-\tau}$ normalizes the state. Suppose an agent switches λ_t from $\lambda_{-\tau}$ to λ_{τ} while the system interacts with the bath. The switching costs work, the amount of which varies from trial to trial. A probability distribution $P_{\text{fwd}}(W)$ represents the probability that a given trial costs work W . By $P_{\text{rev}}(-W)$, we denote the probability that initializing the Hamiltonian to $H(\lambda_{\tau})$ and initializing the system in $\gamma_{\tau} := e^{-\beta H(\lambda_{\tau})}/Z_{\tau}$, then reversing the

¹ Our discussion of work can be phrased alternatively in terms of entropy production (e.g., [19]).

drive according to λ_{-t} , outputs work W .

Fluctuation relations such as Crooks' Theorem govern these distributions [18]. Let $\Delta F := F(\gamma_\tau) - F(\gamma_{-\tau})$ denote the difference between the free energy of γ_τ and that of $\gamma_{-\tau}$. (Throughout this letter, we shall assume ΔF is finite.) Assuming the system is classical; coupled to a bath; and undergoing a Markovian, microscopically reversible evolution, Crooks proved that

$$\frac{P_{\text{fwd}}(W)}{P_{\text{rev}}(-W)} = e^{\beta(W - \Delta F)} \quad (1)$$

[18]. Identical theorems have been shown to govern quantum systems isolated from [3], or interacting with the bath while work is performed (e.g., [5]).

Rényi divergences—The *order- α Rényi divergence* quantifies the distinctness of probability distributions $P(x)$ and $Q(x)$ [13, 24],

$$D_\alpha(P||Q) := \frac{1}{\alpha - 1} \ln \left(\int dx p^\alpha(x) q^{1-\alpha}(x) \right), \quad (2)$$

or of quantum states ρ and σ [25]:

$$D_\alpha(\rho||\sigma) := \frac{1}{\alpha - 1} \ln (\text{Tr}(\rho^\alpha \sigma^{1-\alpha})), \quad (3)$$

wherein Tr denotes the trace, for $\alpha \in [0, 1) \cup (1, \infty)$. The order-1 Rényi divergence, known also as the KL divergence and the average relative entropy, follows from the limit as $\alpha \rightarrow 1$:

$$D(P||Q) = \int dx P(x) \ln (P(x)/Q(x)) \quad (4)$$

for classical distributions, and $D(\rho||\sigma) = \text{Tr}(\rho[\ln(\rho) - \ln(\sigma)])$ for quantum states. We will focus on the order- ∞ divergences:

$$D_\infty(P||Q) = \ln \left(\min\{\lambda \in \mathbb{R} : P(x) \leq \lambda Q(x) \ \forall x\} \right) \quad (5)$$

for classical distributions, and

$$D_\infty(\rho||\sigma) = \ln \left(\max_{i,j} \left\{ \frac{r_i}{s_j} : \langle r_i | s_j \rangle \neq 0 \right\} \right) \quad (6)$$

for quantum states $\rho = \sum_i r_i |r_i\rangle\langle r_i|$ and $\sigma = \sum_j s_j |s_j\rangle\langle s_j|$ [26].

Divergences between phase-space densities—Kawai *et al.* consider a classical system that remains isolated from the bath while work is performed [4]. Governed by Hamiltonian dynamics, the system follows a deterministic trajectory through phase space. Specifying a phase-space point (q, p) at any time t uniquely specifies a trajectory and a work cost $W(q, p, t)$.

An experimenter does not know which trajectory the system follows in any given forward trial, because the experimenter ascribes to the system the initial state $e^{-\beta H(\lambda_{-\tau})}/Z_{-\tau}$. The probability that the system occupies an area- $(dq dp)$ region centered on (q, p) at time t is

$\rho(q, p, t) dq dp$, wherein $\rho(q, p, t)$ denotes the phase-space density. $\tilde{\rho}(q, p, t)$ denotes the phase-space density after an amount $t = 2\tau - t$ of time has passed during the reverse protocol.

Kawai *et al.* proceed as follows. As the system loses no heat while work is performed, the work required to evolve the system along some trajectory equals the difference between the final and initial Hamiltonians: $W(p, q, t) = H(q_\tau, p_\tau, \tau) - H(q_{-\tau}, p_{-\tau}, -\tau)$. The forward process's initial ρ and the reverse process's initial $\tilde{\rho}$ are equated with thermal states. The Hamiltonian is assumed to have time-reversal invariance (TRI): $H(q, p, t) = H(q, -p, t)$. From TRI, the preservation of phase-space densities by Hamiltonian dynamics, and the correspondence of $\rho(q, p, t)$ and $\tilde{\rho}(q, -p, t)$ to the same Hamiltonian follows the "generalized Crooks relation"

$$e^{\beta[W(q,p,t) - \Delta F]} = \frac{\rho(q, p, t)}{\tilde{\rho}(q, -p, t)}. \quad (7)$$

By taking logs, multiplying each side by $\tilde{\rho}(q, -p, t)$, and integrating over phase space, Kawai *et al.* derive

$$\langle W_{\text{diss}} \rangle = \frac{1}{\beta} D(\rho(q, p, t) || \tilde{\rho}(q, -p, t)). \quad (8)$$

The right-hand side (RHS) is well-defined if the support of ρ lies in the support of $\tilde{\rho}$: $\text{supp}(\rho(q, p, t)) \subseteq \text{supp}(\tilde{\rho}(q, -p, t))$ [22].

The nonnegativity of D implies that, on average, performing a protocol quickly costs positive work. The work penalty's nonnegativity has been interpreted as the Second Law of Thermodynamics [4, 27]. According to Stein's Lemma, $D(P||Q)$ quantifies the average probability that an attempt to distinguish between P and Q will fail [28, 29]. $D(\rho(q, p, t) || \tilde{\rho}(q, -p, t))$ quantifies the distinguishability of the forward-process density from its time-reverse. $D(P||Q)$ vanishes if and only if $P = Q$ [28]. Equation (18) shows that reversing the trajectory followed during the forward protocol yields the trajectory followed during the reverse protocol if and only if the system dissipates no work on average. No work is dissipated if the process proceeds quasistatically, such that the system remains in equilibrium. Hence D quantifies roughly how far from equilibrium the system evolves.

Let us turn from averages over infinitely many trials to single trials.

Theorem 1. *The worst-case dissipated work of the foregoing protocol is proportional to an order- ∞ Rényi divergence between phase-space distributions:*

$$W_{\text{diss}}^{\text{worst}} = \frac{1}{\beta} D_\infty(\rho(q, p, t) || \tilde{\rho}(q, -p, t)), \quad (9)$$

if $\text{supp}(\rho(q, p, t)) \subseteq \text{supp}(\tilde{\rho}(q, -p, t))$.

Proof. First, we take the logarithm of each side of the generalized Crooks relation [Eq. (7)]:

$$W - \Delta F = \frac{1}{\beta} \ln \left(\frac{\rho(q, p, t)}{\tilde{\rho}(q, -p, t)} \right). \quad (10)$$

We maximize each side of the equation, invoking the logarithm's monotonicity to shift the maximum into the argument:

$$W_{\max} - \Delta F = \frac{1}{\beta} \ln \left(\max \left\{ \frac{\rho(q, p, t)}{\tilde{\rho}(q, -p, t)} \right\} \right). \quad (11)$$

Comparing the left-hand side (LHS) with the definition of $W_{\text{diss}}^{\text{worst}}$ and the RHS with the definition of D_{∞} yields Eq. (9). \square

Like Eq. (8), Theorem 1 relates dissipated work to a measure of the difference between $\rho(p, q, t)$ and $\tilde{\rho}(p, -q, t)$. The more work is dissipated during the most expensive possible trial, the less the forward-process density can resemble its time-reversed cousin. The lesser the resemblance, the farther the system is expected to depart from equilibrium. As in Eq. (8), the LHS of Eq. (9) is time-independent, so the RHS remains constant for all $t \in [-\tau, \tau]$.

Equation (9) has the correct quasistatic limit: If work is invested infinitesimally slowly, the worst amount of work that can be dissipated—the only amount that can be dissipated—vanishes: $W_{\max} - \Delta F = \Delta F - \Delta F = 0$. Because the system remains in equilibrium, $H(\lambda_t)$ and β determine the state completely. The RHS of Ineq. (9) becomes $D(\rho(q, p, t) || \tilde{\rho}(q, -p, t)) = 0$.

Theorem 1 can aid an agent who has imperfect information about phase-space densities. Kawai *et al.* recommend using Eq. (8) to predict $\langle W_{\text{diss}} \rangle$ from ρ and $\tilde{\rho}$. Phase-space densities, they acknowledge, can be difficult to learn about. So they bound $\langle W_{\text{diss}} \rangle$ with a D between coarse-grained densities. Theorem 1 offers an alternative to coarse-graining. One can use the theorem upon learning just the maximum of $\rho/\tilde{\rho}$, rather than the densities' precise forms. Instead of bounding $\langle W_{\text{diss}} \rangle$, one can calculate a one-shot dissipated work exactly.

Interchanging the arguments of D_{∞} yields the worst-case forfeited work. One can extract less work by implementing the reverse protocol at finite speed than by implementing the protocol quasistatically, due to dissipation. The *worst-case forfeited work*

$$W_{\text{forfeit}}^{\text{worst}} := \Delta F - W_{\max} \quad (12)$$

is the most work an agent might sacrifice for time in any finite-speed reverse trial:

$$W_{\text{forfeit}}^{\text{worst}} = \frac{1}{\beta} D_{\infty}(\tilde{\rho}(q, -p, t) || \rho(q, p, t)), \quad (13)$$

if $\text{supp}(\tilde{\rho}(q, -p, t)) \subseteq \text{supp}(\rho(q, p, t))$.

Divergences between quantum states—Parrondo *et al.* have quantized Eq. (8) [22]. They consider a quantum system governed by a quantum Hamiltonian $H(\lambda_t)$ specified by an external parameter λ_t . Let $\rho(t)$ denote the state occupied by the system at time t . In the forward protocol, the system begins in thermal equilibrium: $\rho(-\tau) = e^{-\beta H_{-\tau}}/Z_{-\tau}$. During $t \in (-\tau, \tau)$, the system

is isolated from the bath, and an agent invests work to switch λ_t from $\lambda_{-\tau}$ to λ_{τ} . The state changes unitarily. During the reverse protocol, the system is prepared in the state $\tilde{\rho}(\tau) = e^{-\beta H_{\tau}}/Z_{\tau}$; time runs from $t = \tau$ to $t = -\tau$; and work is extracted via the time-reversed schedule λ_{-t} .

Assuming that $\text{supp}(\rho(t)) \subseteq \text{supp}(\tilde{\rho}(t))$, Parrondo *et al.* derive

$$\langle W_{\text{diss}} \rangle = \frac{1}{\beta} D(\rho(t) || \tilde{\rho}(t)). \quad (14)$$

Recycling their set-up, we will prove a proportionality between the worst-case dissipated work and an order- ∞ Rényi divergence. We must define “work” explicitly. In some quantum fluctuation-relation contexts, work is defined in terms of two energy measurements [3, 30]: The system begins in the thermal state $\gamma_{-\tau}$. An energy measurement at $t = -\tau$ yields some eigenvalue E_i of $H_{-\tau}$. The system is isolated from the bath, and the state evolves unitarily. An energy measurement at $t = \tau$ yields some eigenvalue \tilde{E}_j of H_{τ} . As the system exchanges no heat during the unitary evolution, the difference between the measurement outcomes equals the work performed: $W = \tilde{E}_j - E_i$.

We assume that the agent does not learn the initial measurement's outcome until the end of the protocol. Because the state begins block-diagonal relative to the initial Hamiltonian, this measure-and-forget operation preserves the initial state.

Theorem 2. *The worst-case work dissipated during any such quantum forward trial is*

$$W_{\text{diss}}^{\text{worst}} = \frac{1}{\beta} D_{\infty}(\rho(t) || \tilde{\rho}(t)). \quad (15)$$

Proof. Let $\rho(t) = \sum_i p_i |i(t)\rangle\langle i(t)|$ and $\tilde{\rho} = \sum_j \tilde{p}_j |\tilde{j}(t)\rangle\langle \tilde{j}(t)|$ denote the states' eigenvalue decompositions. The eigenvalues, and the inner products $\langle i(t) | \tilde{j}(t) \rangle$, remain constant throughout the unitary evolution. $D_{\infty}(\rho(t) || \tilde{\rho}(t))$ therefore remains constant. Without loss of generality, we can evaluate the definition [Eq. (6)] at $t = \tau$:

$$D_{\infty}(\rho(t) || \tilde{\rho}(t)) = \ln \left(\max_{i,j} \left\{ \frac{p_i}{\tilde{p}_j} : \langle i(\tau) | \tilde{j}(\tau) \rangle \neq 0 \right\} \right). \quad (16)$$

Let U denote the unitary that evolves the initial state to the final in the forward process: $\rho(\tau) = U\rho(-\tau)U^{\dagger}$. We can express the inner product as $\langle i(-\tau) | U^{\dagger} | \tilde{j}(\tau) \rangle$. The thermal natures of $\rho(-\tau)$ and $\tilde{\rho}(\tau)$ imply that $p_i = e^{-\beta E_i}/Z_{-\tau}$ and $\tilde{p}_j = e^{-\beta \tilde{E}_j}/Z_{\tau}$. Since $Z_{\tau}/Z_{-\tau} = e^{-\beta \Delta F}$, Eq. (16) is equivalent to

$$D_{\infty}(\rho(t) || \tilde{\rho}(t)) = \ln \left(\max_{i,j} \left\{ e^{\beta(\tilde{E}_j - E_i - \Delta F)} : \langle i(-\tau) | U^{\dagger} | \tilde{j}(\tau) \rangle \neq 0 \right\} \right). \quad (17)$$

The work dissipated in some forward trial is proportional to the exponential's argument. The forward protocol is unable to map $|i(-\tau)\rangle$ to $|\tilde{j}(\tau)\rangle$ if and only if

$\langle i(-\tau)|U^\dagger|\tilde{j}(\tau)\rangle = 0$, i.e., if and only if the condition in Eq. (17) is violated. Hence the worst-case work that can be dissipated during any forward trial is proportional to exponential’s argument, maximized under the condition in Eq. (17). Rearranging Eq. (17) yields Eq. (15). \square

The discussion of irreversibility, distinguishability, t -dependence, the quasistatic limit, and coarse-graining that characterizes the classical Theorem 1 characterizes also the quantum Theorem 2. $W_{\text{diss}}^{\text{worst}}$ is bounded when $H_{-\tau}$ and H_τ have bounded spectra, as in many problems in one-shot statistical mechanics (e.g., [11]).

Divergences between work distributions—We have related dissipated work to a divergence D_∞ between phase-space densities and to a D_∞ between quantum states. We now relate $W_{\text{diss}}^{\text{worst}}$ to a D_∞ between distributions over possible values of work.

The Kullback-Leiber divergence between $P_{\text{fwd}}(W)$ and $P_{\text{rev}}(-W)$ is proportional to the average dissipated work:

$$\frac{1}{\beta}D(P_{\text{fwd}}(W)||P_{\text{rev}}(-W)) = \langle W \rangle_{\text{fwd}} - \Delta F = \langle W_{\text{diss}} \rangle. \quad (18)$$

The first equality follows from the substitution from Crooks’ Theorem [Eq. (1)] for $P_{\text{fwd}}(W)/P_{\text{rev}}(-W)$ in the definition of $D(P_{\text{fwd}}(W)||P_{\text{rev}}(-W))$. We will derive a one-shot analog of Eq. (18).

Theorem 3. *The worst-case work that can be dissipated in any forward trial is proportional to the order- ∞ Rényi divergence between $P_{\text{fwd}}(W)$ and $P_{\text{rev}}(-W)$:*

$$W_{\text{diss}}^{\text{worst}} = \frac{1}{\beta}D_\infty(P_{\text{fwd}}(W)||P_{\text{rev}}(-W)), \quad (19)$$

if the set of possible work-values is bounded.

Proof. By the definition of D_∞ ,

$$D_\infty(P_{\text{fwd}}(W)||P_{\text{rev}}(-W)) = \ln \left(\min \{ \lambda \in \mathbb{R} : P_{\text{fwd}}(W) \leq \lambda P_{\text{rev}}(-W) \forall W \} \right). \quad (20)$$

Let us solve for the minimal λ -value λ_{min} that satisfies the inequality. First, we check that we can divide the inequality by $P_{\text{rev}}(-W)$. Crooks’ Theorem implies that $P_{\text{fwd}}(W) = e^{\beta(W-\Delta F)}P_{\text{rev}}(-W)$. By assumption, $P_{\text{fwd}}(W)$ and $P_{\text{rev}}(-W)$ are nonzero only if W is finite. Also, ΔF is finite. Hence Crooks’ Theorem implies that $P_{\text{rev}}(-W) = 0$ if and only if $P_{\text{fwd}}(W) = 0$. In this case, the inequality becomes $0 \leq \lambda \cdot 0$, which is satisfied by any finite λ and so does not determine λ_{min} . To solve for λ_{min} , we can restrict our focus to $P_{\text{rev}}(-W) \neq 0$, then divide each side of the inequality in Eq. (20) by $P_{\text{rev}}(-W)$:

$$\lambda_{\text{min}} \geq \frac{P_{\text{fwd}}(W)}{P_{\text{rev}}(-W)} \quad \forall W. \quad (21)$$

Substituting into the RHS from Crooks’ Theorem yields $\lambda_{\text{min}} \geq e^{\beta(W-\Delta F)}$. The bound saturates when W

assumes its maximal value W_{max} : $\lambda_{\text{min}} = e^{\beta(W_{\text{max}}-\Delta F)} = e^{\beta W_{\text{diss}}^{\text{worst}}}$. Substituting into Eq. (20) yields Eq. (19). \square

Just as $\frac{1}{\beta}D(P_{\text{fwd}}(W)||P_{\text{rev}}(-W))$ equals the average, over many trials, of dissipated work, $\frac{1}{\beta}D_\infty(P_{\text{fwd}}(W)||P_{\text{rev}}(-W))$ equals the most work that could be dissipated in any trial. An agent can calculate this dissipated work upon inferring P_{fwd} and P_{rev} from experimental or simulation statistics.

Theorem 3 contains a Rényi divergence between work distributions, rather than a D_∞ between phase-space distributions or a D_∞ between quantum states. Hence Theorem 3 governs more protocols than Theorems 1 and 2, as it describes all protocols—quantum or classical, regardless of whether the system exchanges heat while work is performed—that obey Crooks’ Theorem.

Interchanging the divergence’s arguments yields the worst-case forfeited work [Eq. (12)]:

$$W_{\text{forfeit}}^{\text{worst}} = \frac{1}{\beta}D_\infty(P_{\text{rev}}(-W)||P_{\text{fwd}}(W)). \quad (22)$$

Outlook and discussion—We have developed one-shot analogs of three relationships between the average dissipated work $\langle W_{\text{diss}} \rangle$ and an “average” Rényi divergence D . We related the worst-case dissipated work $W_{\text{diss}}^{\text{worst}}$ to an order- ∞ Rényi divergence D_∞ between classical phase-space distributions, between quantum states, and to a D_∞ between work distributions. In all three cases, the proportionality between the averages $\langle W_{\text{diss}} \rangle$ and D also characterizes the one-shot quantities $W_{\text{diss}}^{\text{worst}}$ and D_∞ .

The incorporation of risk tolerance into these results merits investigation. An agent can trade off the guarantee that each trial will accomplish its purpose with the possibility of paying less work (or extracting more work) than by exerting caution. Risk tolerance can be quantified with a parameter $\epsilon \in [0, 1]$. This *failure probability*, chosen by the agent, has been incorporated into Rényi divergences [15] and one-shot statistical mechanics (e.g., [11, 21]). The incorporation of ϵ into the results above, as well as the consideration of different-order Rényi divergences $D_\alpha \neq D_\infty$, should provide further insights into fluctuation relations via one-shot statistical mechanics.

Note added—Lemma 3 appeared previously in an early draft of [17] but was deleted from the manuscript. Theorems 1 and 2 have never, to our knowledge, appeared in the literature.

Acknowledgements—This work was supported by a Virginia Gilloon Fellowship, an IQIM Fellowship, NSF grant PHY-0803371, the FQXi Large Grant for “Time and the Structure of Quantum Theory,” the EPSRC, the John Templeton Foundation, the Leverhulme Trust, the Oxford Martin School, the NRF (Singapore), and the MoE (Singapore). The Institute for Quantum Information and Matter (IQIM) is an NSF Physics Frontiers Center with support from the Gordon and Betty Moore Foundation. VV and OD acknowledge funding from the EU Collaborative Project TherMiQ (Grant Agreement 618074).

-
- [1] C. Jarzynski, *Physical Review Letters* **78**, 2690 (1997).
- [2] G. E. Crooks, *Physical Review E* **60**, 2721 (1999).
- [3] H. Tasaki, arXiv e-print (2000), cond-mat/0009244.
- [4] R. Kawai, J. M. R. Parrondo, and C. Van den Broeck, *Phys. Rev. Lett.* **98**, 080602 (2007).
- [5] H. T. Quan and H. Dong, arXiv e-print , 6 (2008), 0812.4955.
- [6] U. Seifert, *Reports on progress in physics. Physical Society (Great Britain)* **75**, 126001 (2012).
- [7] O. C. O. Dahlsten, R. Renner, E. Rieper, and V. Vedral, *New Journal of Physics* **13**, 053015 (2011).
- [8] D. Egloff, O. C. O. Dahlsten, R. Renner, and V. Vedral, arXiv e-print (2012), 1207.0434.
- [9] J. Åberg, *Nature communications* **4**, 1925 (2013).
- [10] O. C. O. Dahlsten, *Entropy* **15**, 5346 (2013).
- [11] M. Horodecki and J. Oppenheim, *Nature Communications* **4**, 2059 (2013).
- [12] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, 2010).
- [13] A. Rényi, *Proc. Fourth Berkeley Symp. on Math. Statist. and Prob.* **1**, 547 (1961).
- [14] R. Renner, *Security of Quantum Key Distribution*, PhD thesis, ETH Zürich, 2005.
- [15] N. Datta, *IEEE T. Inform. Theory* **55**, 2816 (2009).
- [16] F. Dupuis, L. Krämer, P. Faist, J. M. Renes, and R. Renner, *XVIIth International Congress on Mathematical Physics*, in *XVIIth International Congress on Mathematical Physics*, Aalborg, Denmark, 2012, 1211.3141.
- [17] N. Yunger Halpern, A. J. P. Garner, O. C. O. Dahlsten, and V. Vedral, (2014), 1409.3878.
- [18] G. E. Crooks, *Journal of Statistical Physics* **90**, 1481 (1998).
- [19] S. Deffner and E. Lutz, *Phys. Rev. Lett.* **105**, 170402 (2010).
- [20] R. Dorner, J. Goold, C. Cormick, M. Paternostro, and V. Vedral, *Physical Review Letters* **109**, 160601 (2012).
- [21] O. C. O. Dahlsten *et al.*, arXiv e-print (2015), 1504.05152.
- [22] J. M. R. Parrondo, C. V. den Broeck, and R. Kawai, *New Journal of Physics* **11**, 073008 (2009).
- [23] S. Salek and K. Wiesner, ArXiv e-prints (2015), 1504.05111.
- [24] T. van Erven and P. Harremoës, ArXiv e-prints (2012), 1206.2459.
- [25] S. Beigi, *Journal of Mathematical Physics* **54**, 122202 (2013), 1306.5920.
- [26] M. Tomamichel, M. Berta, and M. Hayashi, *Journal of Mathematical Physics* **55**, 082206 (2014).
- [27] C. Jarzynski, *The European Physical Journal B* **64**, 331 (2008).
- [28] T. Cover and J. Thomas, *Elements of Information Theory* A Wiley-Interscience publication (Wiley, 2006).
- [29] S. Vaikuntanathan and C. Jarzynski, *EPL (Europhysics Letters)* **87**, 60005 (2009).
- [30] J. Kurchan, arXiv e-print (2000), cond-mat/0007360.