Maximum one-shot dissipated work from Rényi divergences

Nicole Yunger Halpern, Andrew J. P. Garner, Oscar C. O. Dahlsten, and Vlatko Vedral

Phys. Rev. E **97**, 052135 — Published 25 May 2018

DOI: 10.1103/PhysRevE.97.052135

One-shot dissipated work from Rényi divergences

Nicole Yunger Halpern,[1] Andrew J. P. Garner,[2,3] Oscar C. O. Dahlsten,[4,2,5] and Vlatko Vedral[2,3,6]

[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 OX1 3PU, United Kingdom*
[3] *Center for Quantum Technologies, National University of Singapore, Republic of Singapore*
[4] *Institute for Quantum Science and Engineering, and Department of Physics, Southern University of Science and Technology, Shenzhen 518055, China*
[5] *London Institute for Mathematical Sciences, 35a South Street, Mayfair, London, W1K 2XF, United Kingdom*
[6] *Department of Physics, National University of Singapore, 2 Science Drive 3, Singapore 117542*

(Dated: May 10, 2018)

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-infinity Rényi divergence is proportional to the maximum possible 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.

I. INTRODUCTION

Thermodynamics concerns large scales and infinitesimally slow evolutions. In the thermodynamic limit, a system's size approaches infinity and is typified by mean behaviors. *Quasistatic* processes proceed slowly enough that the system remains in equilibrium. Equilibrium quantities describe quasistatic processes—for example, the temperature T and the free energy F such as Helmholtz's, $-k_{\text{B}} T \ln Z$ (wherein k_{B} denotes Boltzmann's constant and Z denotes a partition function).

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 and nonequilibrium processes (e.g., [1–6]). One-shot statistical mechanics is used to quantify the efficiency with which work can be invested or extracted, including outside the assumptions of conventional statistical mechanics (e.g., [7–11]): First, the system may be small, violating the thermodynamic limit. Second, the work performed in any given trial—rather than just the work averaged over trials—may be reasoned about. Third, the system may occupy a quantum state coherent relative to the energy eigenbasis.

One-shot statistical mechanics relies on the mathematical toolkit of *one-shot information theory*, or *information theory beyond i.i.d.* (independent and identically distributed variables and quantum states) (e.g., [13–16]). Conventional information theory concerns information-processing tasks such as data compression. One assumes that n random variables X, or quantum states ρ, are processed. The variables and states are assumed to be i.i.d. For example, the probability p_x that X evaluates to x is the same for all instances of X. One calculates the optimal efficiency with which the task can be performed, on average over n, in the limit as $n \to \infty$. *Asymptotic* entropies, such as the Shannon and von Neumann entropies [12], quantify these efficiencies. These entropies are generalized in one-shot information theory. Examples include the Rényi divergences D_α, discussed below. The generalized entropies quantify the efficiencies with which more-general information-processing tasks can be performed. For example, few copies of X or ρ may be processed. The variables or states may not be i.i.d.

One-shot information theory generalizes conventional information theory, as one-shot statistical mechanics extends conventional statistical mechanics. 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.

$\langle W_{\text{diss}} \rangle$ has been shown to be proportional to three instances of the *Kullback-Leibler divergence*, or *average relative entropy*, D_1. D_1 quantifies how much two probability distributions, or two quantum states, differ. (See the "Rényi divergences" section and [12] for

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

reviews.) $\langle W_{\rm diss}\rangle$ has been related to three average relative entropies: (i) a D_1 between phase-space densities $\rho(p,q,t)$ and $\tilde{\rho}(p,-q,t)$, associated with forward and time-reversed processes [4]; (ii) a D_1 between quantum states $\rho(t)$ and $\tilde{\rho}(t)$, associated with forward and reverse processes [21]; and (iii) a D_1 between probability distributions $P_{\rm fwd}(W)$ and $P_{\rm rev}(-W)$ over the work performed during the forward and reverse processes.

This D_1 belongs to a family of *Rényi divergences* D_α that quantify the discrepancies between distributions or between states. D_1 quantifies a discrepancy in terms of an average over many copies of a distribution or state. The *order-∞ Rényi divergence* D_∞ quantifies the distinguishability apparent, in a worst case, from just one copy.

We derive one-shot analogs of all three thermodynamic equalities. The averages $\langle W_{\rm diss}\rangle$ and D_1 are replaced with the one-shot $W_{\rm diss}^{\rm worst}$ and D_∞. The trio reveals the generality of the proportionality between the worst-case dissipated work and a one-shot entropy.

We begin by reviewing fluctuation theorems and Rényi divergences, focusing on D_∞. We recall each $\langle W_{\rm diss}\rangle$ proportionality and derive its one-shot analog. Our main results relate the maximum possible penalty $W_{\rm diss}^{\rm worst}$ of investing work in finite time to three instances of D_∞. We apply our results to a quantum quench, whose work distribution has been studied in several settings [20, 22–27]. Our one-shot analogs of fluctuation-relation results illustrate the insights offered by merging fluctuation relations with one-shot statistical mechanics.

II. BACKGROUND

We review fluctuation theorems, then Rényi divergences.

A. 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 $P_{\rm fwd}(W)$ represents the probability that a given trial costs work W. By $P_{\rm rev}(-W)$, 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 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 article, 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

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

[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]).

B. Rényi divergences

Let P and Q denote probability distributions over the set of values $\{x\}$. The *order-α Rényi divergence* quantifies the distinctness of P and Q [13, 28],

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

or of quantum states ρ and σ [29]:

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

wherein Tr denotes the trace, for $\alpha\in[0,1)\cup(1,\infty)$.

The order-1 Rényi divergence, known also as the *Kullback-Leibler divergence* and the *average relative entropy*, follows from the limit as $\alpha\to1$:

$$D_1(P||Q) = \int dx\; P(x)\ln\big(P(x)/Q(x)\big) \tag{4}$$

for classical distributions, and $D_1(\rho||\sigma) = \mathrm{Tr}\big(\rho[\ln(\rho)-\ln(\sigma)]\big)$ for quantum states. D_1 quantifies an average of the information learned when one mistakes Q for P, or σ for ρ, then is corrected [30, 31].

We focus on the order-∞ divergences: For classical distributions,

$$D_\infty(P||Q) = \ln\big(\min\{\lambda\in\mathbb{R} : P(x)\le\lambda Q(x)\;\;\forall x\}\big) \tag{5}$$

if the support $\mathrm{supp}(Q)\subseteq\mathrm{supp}(P)$, and $D_\infty(P||Q)=\infty$ otherwise. For quantum states,

$$D_\infty(\rho||\sigma) = \ln\left(\max_{i,j}\left\{\frac{r_i}{s_j} : \langle r_i|s_j\rangle\neq0\right\}\right) \tag{6}$$

for quantum states $\rho = \sum_i r_i|r_i\rangle\langle r_i|$ and $\sigma=\sum_j s_j|s_j\rangle\langle s_j|$ [32]. Imagine receiving just one copy of a state that is ρ or σ. Suppose, for simplicity, that the states share the eigenbasis $\{|r_i\rangle\langle r_i|\}$, which you measure. In the worst case, two events occur: (1) The outcome, i_0, maximizes the ratio r_{i_0}/s_{i_0}. Since r_{i_0} is enormous, while s_{i_0} is tiny, you guess that you received ρ. (2) You then learn that you received σ. The information gained from event (2), after (1), equals $D_\infty(\rho||\sigma)$.

III. RESULTS

We now derive equalities between the worst-case work and (i) phase-space densities, (ii) quantum states, and (iii) work distributions.

A. 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 $\tilde{t} = 2\tau - t$ 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"

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

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

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

The right-hand side (RHS) is well-defined if the support of ρ lies in the support of $\tilde{\rho}$: $\mathrm{supp}\big(\rho(q, p, t)\big) \subseteq \mathrm{supp}\big(\tilde{\rho}(q, -p, t)\big)$ [21].

The nonnegativity of D_1 implies that, on average, performing a protocol quickly dissipates positive work. The work penalty's nonnegativity has been interpreted as the Second Law of Thermodynamics [4, 33]. According to Stein's Lemma, $D_1(P||Q)$ quantifies the average probability that an attempt to distinguish between P and Q will fail [30, 34]. $D_1\big(\rho(q, p, t) || \tilde{\rho}(q, -p, t)\big)$ quantifies the distinguishability of the forward-process density from its time-reverse. $D_1(P||Q)$ vanishes if and only if $P = Q$ [30]. 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_1 quantifies roughly how far from equilibrium the system evolves.

Let us turn from averages over infinitely many trials to single trials, starting with our first theorem.

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

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

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

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

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

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

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

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

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, as measured by D_∞. 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 $D_\infty\big(\rho(q, p, t) || \tilde{\rho}(q, -p, t)\big) = 0$.

Theorem 1 can aid an agent who has imperfect information about phase-space densities. Kawai *et al.* recommend using Eq. (8) to predict $\langle W_{\mathrm{diss}} \rangle$ from ρ and $\tilde{\rho}$. Phase-space densities, they acknowledge, can be difficult to learn about. So they bound $\langle W_{\mathrm{diss}} \rangle$ with a D_1 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 $\langle W_{\rm diss}\rangle$, one can calculate a one-shot dissipated work exactly.

One might worry that the RHS of Eq. (9) diverges. For instance, a point particle has a Dirac-delta-function ρ, if the particle has a particular momentum. Evaluating D_∞ on a divergent ρ would yield infinity. In reality, however, finite precision limits measurements of a particle's position and momentum. This practicality regulates the divergence, rendering Theorem 1 applicable to realistic particles.

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*

$$W_{\rm forfeit}^{\rm worst} := \Delta F - W_{\rm max} \tag{12}$$

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

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

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

B. Divergences between quantum states

Parrondo *et al.* have quantized Eq. (8) [21]. 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 $\mathrm{supp}(\rho(t)) \subseteq \mathrm{supp}(\tilde\rho(t))$, Parrondo *et al.* derive

$$\langle W_{\rm diss}\rangle = \frac{1}{\beta} D_1\big(\rho(t)||\tilde\rho(t)\big). \tag{14}$$

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, 35]: 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*

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

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 $t=\tau$:

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

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

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

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). □

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. $W_{\rm diss}^{\rm worst}$ is bounded when $H_{-\tau}$ and H_τ have bounded spectra. Bounded spectra characterize many realistic systems, including one-shot problems (e.g., [11]).

An unbounded example seemingly curtails the theorem's applicability: the classical harmonic oscillator (HO). Specifically, take a positively charged classical particle that moves in one dimension (the x-axis), in a potential well centered at $x=0$. Consider turning on and off an electric field. In the worst case, *prima facie,* the field pushes the particle to the top of the well—infinitely high up, costing $W_{\rm diss}^{\rm worst} = \infty$. However, an HO accurately models a realistic particle only near $x=0$. Farther away, a realistic potential likely flattens, or turns over into a deeper potential well, or becomes well-modeled by an infinitely hard wall, etc. In real-world situations, therefore, $W_{\rm diss}^{\rm worst}$ is finite.[2]

Let us apply Theorem 2 to a sudden quench. Quantum quenches' work distributions have been studied in the context of the transverse-field Ising model [22, 23, 26], trapped ions [20], randomly quenched finite-dimensional systems [24], Fermi gases [25], and semiclassical approximations [27]. Consider a finite-dimensional quantum system $\mathcal{S}$, e.g., a set of N qubits (two-level systems). Let $H(\lambda_t)$ denote the Hamiltonian. The parameter λ_t is *quenched* (changed instantaneously) from λ to $\tilde{\lambda}$ during the forward protocol and from $\tilde{\lambda}$ to λ during the reverse protocol [4]. $\mathcal{S}$ begins the forward protocol in the state $\rho(-\tau) = e^{-\beta H(\lambda)}/Z_{-\tau}$, wherein $H(\lambda) = \sum_j E_j |E_j\rangle\langle E_j|$. $\mathcal{S}$ begins the reverse protocol in $\tilde{\rho}(\tau) = e^{-\beta H(\tilde{\lambda})}/Z_\tau$, wherein $H(\tilde{\lambda}) = \sum_j \tilde{E}_j |\tilde{E}_j\rangle\langle \tilde{E}_j|$. By the density operator's statistical interpretation, $\mathcal{S}$ can be regarded as starting each trial in an energy eigenstate chosen according to a Gibbs distribution. In the worst case, $\mathcal{S}$ begins in the forward process in the lowest-energy eigenstate of $H(\lambda)$, $|E_{\rm min}\rangle$, which is the highest-energy eigenstate of $H(\tilde{\lambda})$, $|\tilde{E}_{\rm max}\rangle$. The work dissipated is $W_{\rm diss}^{\rm worst} = \tilde{E}_{\rm max} - E_{\rm min} - \Delta F$. Now, we calculate D_∞. The state's form has no time to change during the quench. Hence $\rho(t) = \rho(0)$ and $\tilde{\rho}(t) = \rho(\tau)$ $\forall t$. Therefore, $D_\infty(\rho(t)||\tilde{\rho}(t)) = \log\left(\min_{j,k}\left\{\lambda \in \mathbb{R} : \frac{e^{-\beta E_j}}{Z_{-\tau}} \leq \lambda \frac{e^{-\beta \tilde{E}_k}}{Z_\tau}\right\}\right) = \log\left(e^{-\beta E_{\rm min} + \beta \tilde{E}_{\rm max}}\right) + \log(Z_{-\tau}/Z_\tau)$. Equation (15) is satisfied.

C. 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 $W_{\rm diss}^{\rm worst}$ to a D_∞ between distributions over possible values of work.

[2] Even extreme settings lead to finite $W_{\rm diss}^{\rm worst}$ values. Consider, as an example, a work protocol $\mathcal{P}$ that preserves the system's volume, V. The greatest amount W of work that could be performed would turn the system into a volume-V black hole. Adding more energy would raise the black hole's mass, M. The mass varies directly with the radius, R: $M \propto R$. Hence adding more energy would violate the protocol's finite-volume constraint and so would not be work performable during $\mathcal{P}$. Though this example might appear contrived, it has relevance to contemporary physics: The intersection of general relativity and quantum thermodynamics, especially together with high-energy physics, forms a frontier being explored now. Initial steps in this direction include [36–38] and many works inspired by the black-hole-information paradox. We leave a detailed unification of these fields with one-shot statistical mechanics as an opportunity for further study.

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

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

[39, 40]. The first equality follows from the substitution from Crooks' Theorem [Eq. (1)] for $P_{\rm fwd}(W)/P_{\rm rev}(-W)$ in the definition of $D(P_{\rm fwd}(W)||P_{\rm rev}(-W))$. 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 $P_{\rm fwd}(W)$ and $P_{\rm rev}(-W)$:*

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

if the set of possible work-values is bounded.

Proof. By the definition of D_∞,

$$\begin{aligned} & D_\infty(P_{\rm fwd}(W)||P_{\rm rev}(-W)) \\ & = \ln(\min\{\lambda \in \mathbb{R} : P_{\rm fwd}(W) \leq \lambda P_{\rm rev}(-W) \ \forall W\}). \end{aligned} \tag{20}$$

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

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

Substituting into the RHS from Crooks' Theorem yields $\lambda_{\rm min} \geq e^{\beta(W - \Delta F)}$. The bound saturates when W assumes its maximal value $W_{\rm max}$: $\lambda_{\rm min} = e^{\beta(W_{\rm max} - \Delta F)} = e^{\beta W_{\rm diss}^{\rm worst}}$. Substituting into Eq. (20) yields Eq. (19). $\square$

Just as $\frac{1}{\beta} D_1(P_{\rm fwd}(W)||P_{\rm rev}(-W))$ equals the average, over many trials, of dissipated work, $\frac{1}{\beta} D_\infty(P_{\rm fwd}(W)||P_{\rm rev}(-W))$ equals the most work that could be dissipated in any trial. An agent can calculate this dissipated work upon inferring $P_{\rm fwd}$ and $P_{\rm rev}$ 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)]:

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

IV. OUTLOOK

We have developed one-shot analogs of three relationships between the average dissipated work $\langle W_{\text{diss}} \rangle$ and the average Rényi divergence D_1. 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. The triptych of theorems demonstrates an unexpected generality of the equality $W_{\text{diss}}^{\text{worst}} = \frac{1}{\beta} D_\infty(.||.)$.

Beyond this theoretical contribution, our results may have applications to experiments and simulations. We applied Theorem 2 to a quantum quench, whose work distribution has been studied in diverse settings [20, 22–27]. Work distributions have been studied also for trapped ions [42], single-electron boxes [43], and classical gases [41].

Applications to such settings could assume many forms. For instance, experimentalists simulating their systems, before performing experiments, might infer the right-hand side of Eq. (9) or of Eq. (15). The $W_{\text{diss}}^{\text{worst}}$ estimate could inform the preparation of work resources (e.g., a sufficiently charged battery) sufficient to ensure that any implementation of the protocol succeeds. Also, dissipated work may manifest as heat. Equipment such as transistors can break if inundated with too much heat. Such equipment may be strengthened to withstand $W_{\text{diss}}^{\text{worst}}$. Additionally, high-precision measurements of small heat quantities are being developed (e.g., [44]). Our results could provide a "sanity check" on whether new instruments are working properly. If the measured heat exceeds $W_{\text{diss}}^{\text{worst}}$, the instrument is likely malfunctioning.

A few practicalities merit consideration in applications of our theorems. Consider applying Theorem 3 to experimental data. One measures W in each of several trials, and bins the outcomes to form a histogram. Only finitely many trials can be performed, so $P_{\text{fwd}}(W)$ and $P_{\text{rev}}(-W)$ are estimated with finite precision [45–47]. Some $W = W_0$ bin in the $P_{\text{rev}}(-W)$ histogram might have height zero, though $P_{\text{rev}}(-W_0) \neq 0$. The worst-case work would appear to diverge. Given physical expectations that $W_{\text{diss}}^{\text{worst}} \neq \infty$, one could vary the histograms' bin widths to model $P_{\text{rev}}(-W)$ better.

Relatedly, the histograms can be *smoothed*. An agent can trade off the guarantee that each trial will accomplish its purpose for the possibility of paying less work (or extracting more work). An agent's risk tolerance can be quantified with a parameter $\epsilon \in [0, 1]$. The agent ignores area-ϵ tails of the distributions, because they correspond to highly unlikely W-values [48]. This process, called *smoothing*, has been introduced into Rényi divergences [15] and into one-shot statistical mechanics (e.g., [9, 11, 49–51]). Smoothing offers a theoretical and practical opportunity to advance this article's results further into applications.

Note added: Theorem 3 appeared in a preprint of [17], not in the published article. Since the first preprint of this article appeared, [48, 52–54, 56] have addressed other aspects of the one-shot-and-fluctuation-relation overlap. Rényi divergences were applied to fluctuation relations within a resource-theory model in [49].

ACKNOWLEDGMENTS

This work was supported by a Virginia Gilloon Fellowship; an IQIM Fellowship; a Barbara Groce Fellowship; a KITP Graduate Fellowship; NSF grants PHY-0803371, PHY-1125565, and PHY-1125915; 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 (GBMF-2644). VV and OD acknowledge funding from the EU Collaborative Project TherMiQ (Grant Agreement 618074). NYH thanks Ning Bao for conversations about high-energy scenarios. We thank all our referees for feedback that enhanced this article.

[1] C. Jarzynski, Phys. Rev. Lett. **78**, 2690 (1997).
[2] G. E. Crooks, Phys. Rev. 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, Rep. Prog. Phys. Physical Society (Great Britain) **75**, 126001 (2012).
[7] O. C. O. Dahlsten, R. Renner, E. Rieper, and V. Vedral, New J. Phys. **13**, 053015 (2011).
[8] D. Egloff, O. C. O. Dahlsten, R. Renner, and V. Vedral, New J. Phys. **17**, 073001 (2015).
[9] J. Åberg, Nat. Comm. **4**, 1925 (2013).
[10] O. C. O. Dahlsten, Entropy **15**, 5346 (2013).
[11] M. Horodecki and J. Oppenheim, Nat. Comm. **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 Int. Congress on Math. Phys., 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, New J. Phys. **17**, 095003 (2015).
[18] G. E. Crooks, J. Stat. Phys. **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, Phys. Rev. Lett. **109**, 160601 (2012).
[21] J. M. R. Parrondo, C. V. den Broeck, and R. Kawai, New J. of Phys. **11**, 073008 (2009).
[22] Phys. Rev. Lett. **101**, 120603 (2008).
[23] N. O. Abeling and S. Kehrein, Phys. Rev. B **93**, 104302 (2016).
[24] M. Łobejko J. Łuczka and P. Talkner, Phys. Rev. E **95**, 052137 (2017).
[25] A. Sindona and J. Goold and N. Lo Gullo and F. Plastina, New J. Phys. **16** (2014).
[26] A. Gambassi and A. Silva, arXiv e-print (2011), cond-mat/1106.2671.
[27] I. García-Mata and A. J. Roncaglia and D. A. Wisniacki, Phys. Rev. E **95**, 050102 (2017).
[28] T. van Erven and P. Harremoës, ArXiv e-prints (2012), 1206.2459.
[29] S. Beigi, J. Math. Phys. **54**, 122202 (2013), 1306.5920.
[30] T. Cover and J. Thomas, *Elements of Information Theory* A Wiley-Interscience publication (Wiley, 2006).
[31] F. Hiai and D. Petz, Commun. Math. Phys. **143**, 99 (1991).
[32] M. Tomamichel, M. Berta, and M. Hayashi, J. Math. Phys. **55**, 082206 (2014).
[33] C. Jarzynski, Euro. Phys. J. B **64**, 331 (2008).
[34] S. Vaikuntanathan and C. Jarzynski, Europhys. Lett. **87**, 60005 (2009).
[35] J. Kurchan, arXiv e-print (2000), cond-mat/0007360.
[36] A. Bartolotta and S. Deffner, Phys. Rev. X **8**, 011033 (2018).
[37] C. Cirstoiu and D. Jennings, arXiv e-print (2017), quant-ph/1707.09826.
[38] T. Opatrný and L. Richterek, Am. J. Phys. **80**, 66-71 (2012).
[39] A. Gomez-Marin and J. M. R. Parrondo and C. Van den Broeck, Europhys. Lett. **82**, 50002 (2008).
[40] D. Wu and D. A. Kofke, J. Chem. Phys. **123**, 054103 (2005).
[41] G. E. Crooks and C. Jarzynski, Phys. Rev. E **75**, 021116 (2007).
[42] S. An and J.-N. Zhang and M. Um and D. Lv and Y. Lu and J. Zhang and Z.-Q. Yin and H. T. Quan and K. Kim, Nat. Phys. **11**, 193199 (2015).
[43] O.-P. Saira, Y. Yoon and T. Tanttu and M. Möttönen and D. V. Averin and J. P. Pekola Phys. Rev. Lett. **109**, 180601 (2012).
[44] J. P. Pekola, Nat. Phys. **11**, 118 (2015).
[45] A. Pohorille and C. Jarzynski and C. Chipot, J. Phys. Chem. B **114**, 32 (2010).
[46] C. Jarzynski, Phys. Rev. E **73**, 046105 (2006).
[47] N. Yunger Halpern and C. Jarzynski, Phys. Rev. E **93**, 052144 (2016).
[48] O. C. O. Dahlsten, M. S. Choi, D. Braun, A. J. P. Garner, N. Yunger Halpern, and V. Vedral, New J. Phys. **19**, 043013 (2017).
[49] S. Salek and K. Wiesner, arXiv e-print (2015), quant-ph/1504.05111.
[50] F. G. S. L. Brandão, M. Horodecki, N. H. Y. Ng, J. Oppenheim, and S. Wehner, Proc. Nat. Acad. Sci. **112**, 3275 (2015).
[51] R. van der Meer and N. H. Y. Ng and S. Wehner, arXiv e-print (2017) quant-ph/1706.03193.
[52] B.-B. Wei and M. B. Plenio, New J. Phys. **19**, 023002 (2017).
[53] Á. M. Alhambra, L. Masanes, J. Oppenheim, and C. Perry, Phys. Rev. X **6**, 041017 (2016).
[54] J. Åberg, arXiv e-print (2016), quant-ph/1601.01302.
[55] G. Gour and M. P. Müller and V. Narasimhachar and R. Spekkens and N. Yunger Halpern, Phys. Rep. **583**, 1-58 (2015).
[56] S.-Y. Guo and Y. Peng and C. Peng and H. Deng and Y.-R. Jin and C. Tang and X. Zhu and D. Zheng and H. Fan, arXiv e-print (2017), quant-ph/1710.10234.