

Number of trials required to estimate a free-energy difference, using fluctuation relations

Nicole Yunger Halpern^{1,*} and Christopher Jarzynski^{2,3}

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

²*Department of Chemistry and Biochemistry, University of Maryland, College Park, Maryland 20742, USA*

³*Institute for Physical Science and Technology, University of Maryland, College Park, Maryland 20742, USA*

(Received 25 January 2016; revised manuscript received 12 March 2016; published 26 May 2016)

The difference ΔF between free energies has applications in biology, chemistry, and pharmacology. The value of ΔF can be estimated from experiments or simulations, via fluctuation theorems developed in statistical mechanics. Calculating the error in a ΔF estimate is difficult. Worse, atypical trials dominate estimates. How many trials one should perform was estimated roughly by Jarzynski [Phys. Rev. E **73**, 046105 (2006)]. We enhance the approximation with the following information-theoretic strategies. We quantify “dominance” with a tolerance parameter chosen by the experimenter or simulator. We bound the number of trials one should expect to perform, using the order- ∞ Rényi entropy. The bound can be estimated if one implements the “good practice” of bidirectionality, known to improve estimates of ΔF . Estimating ΔF from this number of trials leads to an error that we bound approximately. Numerical experiments on a weakly interacting dilute classical gas support our analytical calculations.

DOI: [10.1103/PhysRevE.93.052144](https://doi.org/10.1103/PhysRevE.93.052144)

I. INTRODUCTION

The numerical estimation of free-energy differences is an active area of research, having applications to chemistry, microbiology, pharmacology, and other fields. Fluctuation relations can be used to estimate equilibrium free-energy differences ΔF from nonequilibrium experimental and simulation data. One repeatedly measures the amount W of work extracted from, or performed on, a system during an experiment or simulation. Fluctuation relations express the value of ΔF in terms of averages over infinitely many trials. Finitely many trials are performed in practice, introducing errors into estimates of ΔF . Efforts to quantify these errors, and to promote “good practices” in estimating ΔF , have been initiated (e.g., [1–12]).

How many trials should one perform to estimate ΔF reliably? The work W extracted from a system is a random variable that assumes different values in different trials. Typical trials involve W values that contribute little to the averages being estimated. *Dominant* W values, which largely determine the averages, characterize few trials [13]. Until observing a dominant W value, one cannot estimate ΔF with reasonable accuracy. The probability that some trial will involve a dominant W value determines the number N of trials one should expect to perform.

A rough estimate of N was provided in [13]. In this paper, we enhance the estimate’s precision. First, we introduce fluctuation relations and *one-shot information theory*, a mathematical toolkit for quantifying efficiencies at small scales. Next, we quantify dominance in terms of a tolerance parameter w^δ . We bound the number N_δ of trials expected to be required to observe a dominant work value. This bound depends on the thermal order- ∞ Rényi entropy H_∞^β , a quantity inspired by one-shot information theory [14]. The bound can be estimated during an implementation of the “bidirectionality good practice” recommended in [1]. Finally, we approximately

bound the error in a ΔF estimate inferred from N_δ trials. A weakly interacting dilute classical gas [15] illustrates our analytical results.

II. TECHNICAL INTRODUCTION

Let us introduce nonequilibrium fluctuation relations and the thermal order- ∞ Rényi entropy H_∞^β .

A. Nonequilibrium fluctuation relations

Nonequilibrium fluctuation relations govern statistical mechanical systems arbitrarily far from equilibrium. Consider a system in thermal equilibrium with a heat bath at inverse temperature $\beta \equiv \frac{1}{k_B T}$, wherein k_B denotes Boltzmann’s constant. We focus on classical systems for simplicity, though fluctuation relations have been extended to quantum systems [16]. Suppose that a time-dependent external parameter λ_t determines the system’s Hamiltonian: $H = H(\lambda_t, \mathbf{z})$, wherein \mathbf{z} denotes a phase-space point. If the system consists of an ideal gas in a box, λ_t may denote the height of the piston that caps the gas. Suppose that, at time $t = -\tau$, the system begins with the equilibrium phase-space density $e^{-\beta H(\lambda_{-\tau}, \mathbf{z})} / Z_{-\tau}$, wherein the partition function $Z_{-\tau}$ normalizes the state. The external parameter is then varied according to a predetermined schedule λ_t , from $t = -\tau$ to $t = \tau$. The system evolves away from equilibrium if τ is finite. In the gas example, the piston is lowered, compressing the gas. We call this process the *forward protocol*.

The *reverse protocol* begins with the system at equilibrium relative to $H(\lambda_\tau, \mathbf{z})$. The external parameter is changed to $\lambda_{-\tau}$ along the time reverse of the path followed during the forward protocol. In the gas example, the piston is raised, and the gas expands.

Changing the external parameter requires or outputs some amount of work. We use the following sign convention: the forward process tends to require an investment of a positive amount $W > 0$ of work, and the reverse process tends to output $W > 0$. The value of W varies from trial to trial.

*nicoleyh@caltech.edu

After performing many trials, one can estimate the probability $P_{\text{fwd}}(W)$ that any particular forward trial will cost an amount W of work and the probability $P_{\text{rev}}(-W)$ that any particular reverse trial will output an amount W .

These probabilities satisfy *Crooks' Theorem* [17],

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

Here, $\Delta F := F_\tau - F_{-\tau}$ denotes the difference between the free energy $F_\tau = -\beta^{-1} \log(Z_\tau)$ of the Gibbs distribution $e^{-\beta H(\lambda_\tau, \mathbf{z})}/Z_\tau$ corresponding to the final Hamiltonian and the free energy $F_{-\tau} = -\beta^{-1} \log(Z_{-\tau})$ of the Gibbs distribution corresponding to $H(\lambda_{-\tau}, \mathbf{z})$. Multiplying each side of Crooks' Theorem by $P_{\text{rev}}(-W)e^{\beta\Delta F}$, then integrating over W , yields a version of the *nonequilibrium work relation* [18]:

$$e^{\beta\Delta F} = \langle e^{\beta W} \rangle_{\text{rev}} \quad (2)$$

$$:= \int_{-\infty}^{\infty} dW e^{\beta W} P_{\text{rev}}(-W). \quad (3)$$

The angle brackets denote an average over infinitely many trials. To calculate ΔF , one performs many trials, estimates the average, and substitutes into Eq. (2).

B. Thermal order- ∞ Rényi entropy (H_∞^β)

Entropies quantify uncertainties in statistical mechanics and in information theory. Let $P := \{p_i\}$ denote a probability distribution over a discrete random variable X . The *Shannon entropy* $H_S(P) := -\sum_i p_i \log(p_i)$ quantifies an average, over infinitely many trials, of the information one gains upon learning the value assumed by X in one trial [19].

H_S has been generalized to a family of *Rényi entropies* H_α . The parameter $\alpha \in [0, \infty)$ is called the *order*. The H_α 's quantify uncertainties related to finitely many trials. In the limit as $\alpha \rightarrow \infty$, H_α approaches

$$H_\infty(P) = -\log(p_{\max}), \quad (4)$$

wherein p_{\max} denotes the greatest p_i . This maximal entropy has applications to randomness extraction: the efficiency with which finitely many copies of P can be converted into a uniformly random distribution $(\underbrace{\frac{1}{d}, \dots, \frac{1}{d}}_d)$ is quantified with

$H_\infty(P)$ [20].

The distributions P_{fwd} and P_{rev} in Crooks' Theorem are continuous. Hence we need a continuous analog of H_∞ . The definition

$$H_\infty^\beta(P) := -\log(p_{\max}/\beta) \quad (5)$$

has been shown to be useful in contexts that involve heat baths [14]. p_{\max} denotes the greatest value of the probability density P . p_{\max} can diverge, e.g., if P represents a Dirac delta function. But delta functions characterize the work distributions of quasistatic protocols, whose work $W = \Delta F$ in every trial. We focus on more-realistic, quick protocols. P_{fwd} and P_{rev} are short and broad, so p_{\max} is finite.

The density p_{\max} has dimensions of inverse energy, which are canceled by the β in Eq. (5). Hence the logarithm's argument is dimensionless. For further discussion about H_∞^β , see [14].

III. QUANTIFICATION OF DOMINANCE

Let us return to the nonequilibrium work relation (3). The exponential enlarges already-high W values, which dominate the integral. To estimate the integral accurately, one must perform trials that output large amounts of work. Few trials do; dominant W values are *atypical* [13]. How many trials should one expect to need to perform to achieve reasonable convergence of the exponential average in Eq. (3)?

An approximate answer was provided in [13]:

$$N \sim e^{\beta \langle W \rangle_{\text{fwd}} - \Delta F}, \quad (6)$$

wherein $\langle \cdot \rangle_{\text{fwd}}$ denotes an average with respect to $P_{\text{fwd}}(W)$. The *average dissipated work* $\langle W \rangle_{\text{fwd}} - \Delta F$ represents the mean amount of work wasted as heat. Switching λ_t quasistatically (infinitely slowly) would cost an amount ΔF of work. Switching at a finite speed costs more: work is dissipated into the bath as heat when the system is driven away from equilibrium. The dissipated work $W - \Delta F$ signifies the extra work paid to switch λ_t in a finite amount of time.

How large must a W value be to qualify as dominant? This question remained open in [13]. We propose a definition inspired by information-theoretic protocols in which an agent specifies an error tolerance. The experimenter who switches λ_t , or the programmer who simulates trials, chooses a threshold value of w^δ used to lower bound the W values considered large.

Definition 1. A work value W extracted from a reverse-protocol trial is called *w^δ -dominant* if $W \geq w^\delta$ for the fixed value w^δ chosen by the agent.

A similar quantity is defined in [3]. Lu and Kofke assess the accuracy of free-energy-perturbation (FEP) calculations. FEP is used to estimate free-energy differences ΔF . FEP results from a limit of nonequilibrium-fluctuation theory [18]. In [3], a fixed-length simulation is assumed to be performed. A difference u between potential energies is measured. u , in FEP, plays the role of W in nonequilibrium fluctuation relations. Lu and Kofke denote by $p(u)$ the probability that a fixed-length simulation yields the potential-energy difference u . *Limit energies* u_1 and u_2 are defined as the extreme realizable u values. Lu and Kofke fix the simulation length, then calculate the most likely limit energy, u^* . In contrast, the agent in the present work fixes a tolerance w^δ . The number N_δ of required trials (similar to the simulation length) is then bounded. Lu and Kofke also use the mode of W^* to calculate the error in ΔF . The *neglected-tail model* of [3] was extended from FEP to nonequilibrium fluctuation relations in [6]. When calculating the error in ΔF , Wu and Kofke average over possible values of the limit energy W^* . The framework in [3,6] accommodates arbitrary W^* values. Yet statistical properties, such as the mean and mode, are emphasized. That emphasis is complemented by the present paper's information-theory-inspired choice of w^δ by the agent. Additionally, the choice $w^\delta = \langle W \rangle_{\text{fwd}} - \Delta F$ of the dissipated work is analyzed below.

Definition 1 enables us to bound the number N_δ of trials expected to be performed before one trial outputs a w^δ -dominant amount of work.

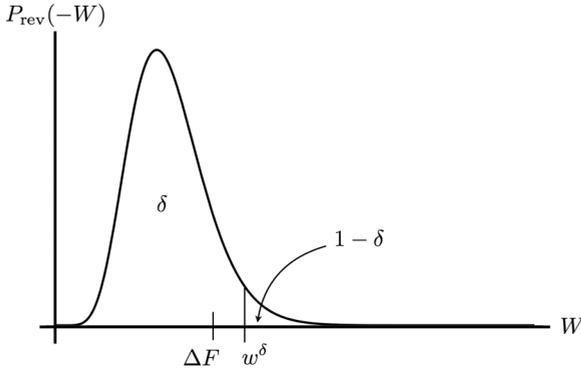


FIG. 1. Dominant values of work extractable from reverse-protocol trials. Large values W of work contribute the most to the integral in the nonequilibrium fluctuation relation (2). An amount W of extracted work is called w^δ -dominant if it is at least as great as the threshold w^δ specified by the experimenter: $W \geq w^\delta$. The probability that any particular reverse trial will output a w^δ -dominant amount of work is $\int_{w^\delta}^{\infty} dW P_{\text{rev}}(-W) = 1 - \delta$. This probability equals the area of the region under the distribution's right-hand tail.

IV. BOUND ON EXPECTED NUMBER N_δ OF TRIALS REQUIRED

Imagine implementing reverse trials until extracting a w^δ -dominant amount of work from one trial. One might have luck and extract $W \geq w^\delta$ on the first try. But one would not expect to. One would expect the number of trials to equal the inverse $1/\int_{w^\delta}^{\infty} dW P_{\text{rev}}(-W)$ of the probability that any particular reverse trial will output $W \geq w^\delta$. In the notation of [14], $\int_{w^\delta}^{\infty} dW P_{\text{rev}}(-W) = 1 - \delta$ (see Fig. 1),

$$N_\delta = \frac{1}{1 - \delta}. \quad (7)$$

Let us clarify what “expect to perform N_δ trials” means. Imagine performing M sets of reverse trials. In each set, one performs trials until extracting $W \geq w^\delta$ from one trial. Let N_δ^i denote the number of trials performed during the i th set. Consider averaging N_δ^i over the M sets of trials: $\frac{1}{M} \sum_{i=1}^M N_\delta^i$. As the number of sets grows large, the average of the number of required trials in a set approaches the “expected” value N_δ :

$$\lim_{M \rightarrow \infty} \frac{1}{M} \sum_{i=1}^M N_\delta^i = N_\delta. \quad (8)$$

This interpretation will facilitate our bounding of N_δ .

Theorem 1 (Bound on expected number of trials). The number N_δ of reverse trials expected to be performed before one trial outputs a w^δ -dominant amount $W \geq w^\delta$ of work is bounded as

$$N_\delta \geq e^{\beta(w^\delta - \Delta F) + H_\infty^\beta(P_{\text{fwd}})}. \quad (9)$$

Proof. The inequality

$$w^\delta \leq \Delta F - \frac{1}{\beta} [H_\infty^\beta(P_{\text{fwd}}) + \log(1 - \delta)] \quad (10)$$

was derived in [14]. The derivation relies on the definitions of $1 - \delta$ and H_∞^β , on Crooks' Theorem, and on the bound

$P_{\text{fwd}}(W) \leq p_{\text{max}} \forall W$. Solving for $1 - \delta$, then inverting the probability [Eq. (7)], yields inequality (9). ■

Inequality (9) implies that the bound on N_δ increases with w^δ , which makes sense. As we raise the threshold w^δ , fewer work values qualify as w^δ -dominant. Hence more trials are expected to be required before a w^δ -dominant work value is observed.

A. Improvement over relation (6)

Inequality (9) resembles its inspiration, relation (6), which states that the number N of trials required to achieve convergence of the average in Eq. (3) increases exponentially with the average dissipated work $\langle W \rangle_{\text{fwd}} - \Delta F$. Similarly, the bound on N_δ increases exponentially with the “one-shot dissipated work” $w^\delta - \Delta F$. This $w^\delta - \Delta F$ represents the work sacrificed for time in a forward trial that costs an amount w^δ of work.

Moreover, N_δ is defined in terms of the reverse process. Yet the bound on N_δ given by inequality (9) depends on the forward work distribution, via $H_\infty^\beta(P_{\text{fwd}})$. Similarly, in relation (6), the number N of repetitions of the reverse process required for the convergence of Eq. (3) depends on the forward work distribution $P_{\text{fwd}}(W)$, via $\langle W \rangle_{\text{fwd}}$.

Despite its similarity to relation (6), inequality (9) offers three advantages. First, inequality (9) quantifies dominance with δ , reflecting the agent's accuracy tolerance. Next, relation (6) is a rough estimate. Inequality (9) is a strict bound on the number of trials expected to be performed before a w^δ -dominant amount of work is extracted. Finally, inequality (9) contains an entropy that has no analog in relation (6). The entropy tightens the bound when

$$p_{\text{max}} < \beta. \quad (11)$$

This inequality is satisfied, for instance, in RNA-hairpin experiments used to test fluctuation theorems [21].

To appreciate these advantages over relation (6), we can define w^δ -dominant work values by choosing $w^\delta = \langle W \rangle_{\text{fwd}}$, as in [13]. The bound becomes

$$N_\delta \geq e^{\beta(\langle W \rangle_{\text{fwd}} - \Delta F) + H_\infty^\beta(P_{\text{fwd}})}. \quad (12)$$

When $p_{\text{max}} < \beta$ (such that $H_\infty^\beta > 1$), the number of trials required for Eq. (3) to converge exceeds the prediction in relation (6).

We can gain further insight by rewriting inequality (9) as

$$N_\delta \geq \frac{\beta}{p_{\text{max}}} e^{\beta(\langle W \rangle_{\text{fwd}} - \Delta F)}, \quad (13)$$

using the definition of H_∞^β [Eq. (5)]. The fraction β/p_{max} represents approximately the number of forward trials performed before one trial's W value falls within a width- $(k_B T)$ window about the most probable work value W_{max} : $W \in [W_{\text{max}} - \frac{k_B T}{2}, W_{\text{max}} + \frac{k_B T}{2}]$. The value of β/p_{max} generically increases with the width of the distribution $P_{\text{fwd}}(W)$. Hence the bound on N_δ , as written in inequality (13), is a product of two factors. The first depends on the forward work distribution's width, and the second, on its mean. In contrast, relation (6) depends only on the mean.

The area under distributions' tails is evoked also in [6]. Wu and Kofke use their *neglected tail model* to estimate the bias in ΔF .

B. Classical vs quantum applications

Classical mechanics describes most experiments and numerical simulations for which N_δ needs calculating. Nonetheless, quantum experiments merit consideration.

We have assumed that the work distributions $P_{\text{fwd}}(W)$ and $P_{\text{rev}}(-W)$ are continuous. Classical systems have continuous work distributions: a classical system's possible energies form a continuous set. So do the differences between possible energy values—the possible work values. Continuousness leads to inequality (10), from which Theorem 1 is derived. How to extend inequality (10) to discrete sets of possible work values is unclear.

Quasiclassical systems can have continuous work distributions. By *quasiclassical*, we mean systems whose energies form a discrete set but whose states (density operators) commute with the Hamiltonian. Consider a quasiclassical system that exchanges heat with a bath throughout the work extraction. The system always occupies an energy eigenstate if the energy is measured frequently [14,22]. The work performance lowers the system's energy levels. Suppose that two levels fall at different rates. The system can hop from level to level at any time. Hopping at time t can output infinitesimally more work than hopping at time $t + dt$ [14]. Such quasiclassical systems obey Theorem 1.

Discrete work distributions characterize quantum systems that undergo the *two-time-measurement protocol* [23,24]. A quantum system undergoes an energy measurement, is isolated from the bath, performs work unitarily, and suffers another energy measurement. The differences between the possible measurement outcomes form a discrete set. Extending Theorem 1 to such protocols could merit investigation. One might incorporate the bin width of the histograms used to approximate $P_{\text{fwd}}(W)$ and $P_{\text{rev}}(-W)$. On the other hand, bin widths are artificial approximation tools, chosen by the experimenter. One might prefer a theory independent of such an approximation [14]. Extensions may be galvanized by the evolution of quantum experiments to a point that requires N_δ estimations.

C. Fail safety

Fail safety is a property of certain estimates calculated from incomplete data. The bound on N_δ depends on the free-energy difference ΔF . ΔF is estimated from forward-trial data. Finitely many forward trials are performed. Hence the ΔF estimate is biased. This bias skews one's estimate of the N_δ bound. Suppose that the estimate lay above the true value of N_δ . The N_δ -bound estimate would lead the agent to perform enough trials to estimate ΔF with reasonable accuracy. The N_δ -bound estimate would be *fail safe* [8,9]. Fail safety is often desirable. Surprisingly, a lack of fail safety benefits Theorem 1, because inequality (9) lower-bound N_δ .

The bias in the ΔF estimate lowers estimates of the N_δ bound below the bound's true value: the nonequilibrium fluctuation relation can be expressed as $e^{-\beta\Delta F} = \langle e^{-\beta W} \rangle_{\text{fwd}}$

[18]. Solving for ΔF yields

$$\Delta F = -\frac{1}{\beta} \log \langle e^{-\beta W} \rangle_{\text{fwd}}. \quad (14)$$

Forward trials tend to cost large amounts of work: typical W values are high. High W values lower the estimate of $\langle e^{-\beta W} \rangle_{\text{fwd}}$ below the average's true value. This low estimate raises the ΔF estimate above the true ΔF value, by Eq. (14). This overestimate of ΔF lowers the estimate of the N_δ bound below the bound's true value, by inequality (9).

In summary, inequality (9) lower bounds N_δ . Estimating this lower bound with biased data generates an even lower bound on N_δ :

$$N_\delta \geq (\text{true lower bound}) \geq (\text{estimated lower bound}). \quad (15)$$

This second lower bounding renders Theorem 1 robust against the bias in the ΔF estimate.

This robustness precludes fail safety. Suppose that the protocol were fail safe. The estimate of the N_δ bound would lie above the true bound:

$$N_\delta \geq (\text{true lower bound}) \leq (\text{estimated lower bound}). \quad (16)$$

One's estimate of the lower bound on N_δ would not necessarily lower-bound N_δ . An experimentalist could not use Theorem 1. The theorem benefits, unusually, from a lack of fail safety.

V. EVALUATING THE N_δ BOUND

Not only does inequality (9) have a theoretically satisfying form, but it can also be estimated in practice. We will discuss how to estimate the $H_\infty^\beta(P_{\text{fwd}})$ and the ΔF in the bound. The bound can be estimated reasonably, we argue, from not too many trials.

The experimental setup determines β , and the agent chooses w^δ . $H_\infty^\beta(P_{\text{fwd}})$ and ΔF can be estimated if one implements the "good practice" of bidirectionality. To mitigate errors in (ΔF) estimates, one should perform forward trials, perform reverse trials, and combine all the data [1]. Upon performing several forward trials, one can estimate $H_\infty^\beta(P_{\text{fwd}})$ and ΔF . One can estimate the N_δ bound, then perform (probably at least N_δ) reverse trials until observing a w^δ -dominant work value, and improve the ΔF estimate.¹

H_∞^β depends on p_{max} , the greatest probability (per unit energy) of any possible forward-trial outcome. This outcome will likely appear in many trials. Hence one expects to estimate H_∞^β well from finitely many forward trials.

¹ N_δ can be estimated from reverse trials alone, less reliably. One could perform a few reverse trials, estimate $P_{\text{rev}}(-W)$, and estimate ΔF . From these estimates and from Crooks' Theorem, one could estimate $P_{\text{fwd}}(W)$. From $P_{\text{fwd}}(W)$, one could estimate $H_\infty^\beta(P_{\text{fwd}})$, then estimate the N_δ bound. One could repeat this process, improving one's estimate of the bound, until observing a w^δ -dominant work value. But the estimate of ΔF is expected to jump repeatedly [13]. This sawtooth behavior, as well as the piling of estimate upon estimate, may taint the estimates of the bound.

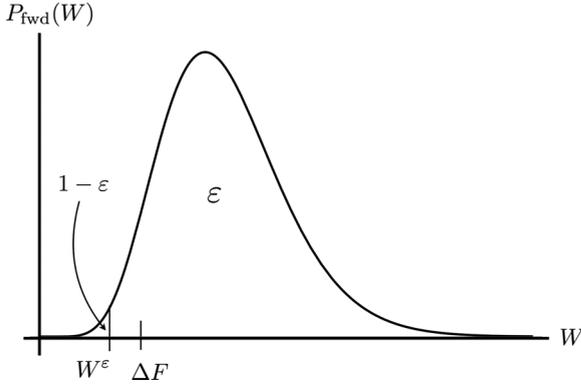


FIG. 2. Dominant values of work invested in forward-protocol trials. Small values W of work dominate the nonequilibrium work relation (17). An amount W of invested work is called W^ϵ -dominant if it lies below or on the threshold W^ϵ chosen by the experimenter: $W \leq W^\epsilon$. The probability that any particular forward trial will require a W^ϵ -dominant amount of work is $1 - \epsilon = \int_{-\infty}^{W^\epsilon} dW P_{\text{fwd}}(W)$. This probability equals the area under the distribution's left-hand tail.

VI. FORWARD-PROTOCOL BOUND

Trials or computations performed in one direction can cost more time than trials or computations performed in the opposite direction [3]. We have bounded a number N_δ of reverse trials. Similarly, we should bound the number N_ϵ of forward trials expected to be performed before a W^ϵ -dominant amount of work is invested. The analysis is analogous to that of N_δ .

The nonequilibrium work relation for the forward process is

$$\langle e^{-\beta W} \rangle_{\text{fwd}} = e^{-\beta \Delta F}. \quad (17)$$

The forward trials that dominate the average in Eq. (17) cost unusually small amounts of work. In the notation of [14], W^ϵ -dominant work values satisfy $W \leq W^\epsilon$, for a tolerance W^ϵ chosen by the agent. Each forward trial has a probability $1 - \epsilon$ of costing a W^ϵ -dominant amount of work (see Fig. 2). Theorem 4 of [14] bounds W^ϵ in terms of $1 - \epsilon$. Solving for $1 - \epsilon$, then inverting, bounds the number $N_\epsilon = 1/(1 - \epsilon)$ of forward trials expected to be performed before any trial costs a W^ϵ -dominant amount of work:

$$N_\epsilon \geq e^{-\beta(W^\epsilon - \Delta F) + H_\infty^\beta(P_{\text{rev}})}. \quad (18)$$

VII. ERROR ESTIMATE

Calculating the error in a ΔF estimate is crucial but difficult. Whenever one infers a value from data, the inference's reliability must be reported. Common error analyses do not suit estimates of ΔF values, for two reasons. First, ΔF depends on the random variable W logarithmically [see Eq. (2)]. Second, W tends not to be Gaussian. Approaches such as an uncontrolled approximation, in the form of a truncation of a series expansion, have been proposed [1]. Our approach centers on the agent's choice of w^δ .

Consider choosing a w^δ value and performing N_δ trials. With what accuracy can one estimate ΔF ? We will bound the

percent error

$$\epsilon := \left| \frac{\Delta F - (\Delta F)_{\text{est}}}{\Delta F} \right| \quad (19)$$

roughly. To render the problem tractable, we assume that one knows the exact form of $P_{\text{rev}}(-W)$ for all $W \leq w^\delta$.

This assumption features also in the neglected-tail model of [3,4,6]. The percent error in $e^{-\beta \Delta F}$ is calculated, with free-energy perturbation theory (FEP), in [3]. This percent error, if small, approximates the absolute error $\Delta F - (\Delta F)_{\text{est}}$ in the free-energy difference [4]. Bias calculations are extended from FEP to nonequilibrium work fluctuation relations in [6].

Theorem 2 (Approximate error bound). Let the work tolerance be $w^\delta \in (-\infty, \infty)$. Let $(\Delta F)_{\text{est}}$ denote the estimate of the free-energy difference ΔF inferred from data taken during N_δ trials. If $(\Delta F)_{\text{est}}$ is calculated from the exact form of $P_{\text{rev}}(-W) \quad \forall W \leq w^\delta$, the estimate has a percent error of

$$\epsilon \geq \frac{1}{\beta(\Delta F)} [\eta + O(\eta^2)], \quad (20)$$

wherein

$$\eta := \frac{e^{\beta w^\delta}}{N_\delta \langle e^{\beta W} \rangle_{\text{rev}}}. \quad (21)$$

Proof. Let us solve the nonequilibrium work relation (2) for ΔF :

$$\Delta F = \frac{1}{\beta} \log(\langle e^{\beta W} \rangle_{\text{rev}}) \quad (22)$$

$$= \frac{1}{\beta} \log \left(\int_{-\infty}^{\infty} dW e^{\beta W} P_{\text{rev}}(-W) \right). \quad (23)$$

The estimate has a similar form:

$$(\Delta F)_{\text{est}} = \frac{1}{\beta} \log \left(\int_{-\infty}^{w^\delta} dW e^{\beta W} P_{\text{rev}}(-W) \right) \quad (24)$$

$$= \frac{1}{\beta} \log \left(\int_{-\infty}^{\infty} dW e^{\beta W} P_{\text{rev}}(-W) - \int_{w^\delta}^{\infty} dW e^{\beta W} P_{\text{rev}}(-W) \right). \quad (25)$$

We replace the first integral with $\langle e^{\beta W} \rangle_{\text{rev}}$, using Eq. (22). The second term, representing the error, is expected to be much smaller than the first term. This second term will serve as a small parameter in a Taylor expansion:

$$(\Delta F)_{\text{est}} = \frac{1}{\beta} \left[\log(\langle e^{\beta W} \rangle_{\text{rev}}) + \log \left(1 - \frac{\int_{w^\delta}^{\infty} dW e^{\beta W} P_{\text{rev}}(-W)}{\langle e^{\beta W} \rangle_{\text{rev}}} \right) \right] \quad (26)$$

$$= \Delta F - \frac{1}{\beta} \{ \eta' + O([\eta']^2) \}, \quad (27)$$

wherein

$$\eta' := \frac{\int_{w^\delta}^{\infty} dW e^{\beta W} P_{\text{rev}}(-W)}{\langle e^{\beta W} \rangle_{\text{rev}}}. \quad (28)$$

We can bound the numerator, using Fig. 1:

$$\int_{w^\delta}^{\infty} dW e^{\beta W} P_{\text{rev}}(-W) \quad (29)$$

$$\geq e^{\beta w^\delta} \int_{w^\delta}^{\infty} dW P_{\text{rev}}(-W) \quad (30)$$

$$= e^{\beta w^\delta} (1 - \delta) = \frac{e^{\beta w^\delta}}{N_\delta}. \quad (31)$$

Substituting into Eq. (28) yields $\eta' \geq \eta$. Hence Eq. (27) reduces to

$$(\Delta F)_{\text{est}} \leq \Delta F - \frac{1}{\beta} [\eta + O(\eta^2)]. \quad (32)$$

Substituting into the percent error's definition [Eq. (19)] yields inequality (20).

The approximate error bound can be estimated from agent-chosen parameters and from data: the experiment's setup determines the value of β . The agent chooses the value of w^δ . For N_δ , one can substitute the number of trials performed [or can substitute from inequality (9)]. ΔF and $\langle e^{\beta W} \rangle_{\text{rev}}$ can be estimated from data.

VIII. NUMERICAL EXPERIMENTS

To illustrate our analytical results, we considered the weakly interacting dilute classical gas. This system's forward and reverse work distributions can be calculated exactly [15]. The gas begins in equilibrium with a heat bath at inverse temperature $\beta \equiv \frac{1}{k_B T}$. During the forward protocol, the gas is isolated from the bath at $t = -\tau$. The gas is quasistatically compressed, its temperature rising from T . During the reverse protocol, the gas expands and cools. When discussing either direction, we denote the initial volume by V_0 and the final volume by V_1 .

The probability densities over the possible work values were calculated in [15]:

$$P(W) = \frac{\beta}{|\alpha| \Gamma(k)} \left(\frac{\beta W}{\alpha} \right)^{k-1} e^{-\beta W/\alpha} \theta(\alpha W). \quad (33)$$

During the forward protocol, $\alpha := (V_0/V_1)^{2/3} - 1 > 0$; during the reverse, $\alpha < 0$. The gamma function is denoted by $\Gamma(k)$, and its argument by $k := \frac{3}{2}n$, wherein n denotes the number of particles. The theta function $\theta(\alpha W)$ ensures that $W \geq 0$ is invested in forward trials (for which $P = P_{\text{fwd}}$), and $W \leq 0$, in reverse trials (for which $P = P_{\text{rev}}$).

This model illustrates accuracies also in [10]. Kofke synthesizes theoretical results about ΔF estimates. Relevant results include the neglected-tail model [6]. Numerical experiments on the gas illustrate those results.

We sampled 10^5 values of W from the forward (compression) work distribution and 10^5 values from the reverse (expansion) work distribution. Figure 3 shows the probability densities and the sampled data. We chose $V_0/V_1 = 2$ and $n = 6$, following [15], and $\beta = 10$. Dividing a histogram of the forward-protocol data into 50 bins yielded $p_{\text{max}} = 1.577$. Satisfying inequality (11), this p_{max} enables $H_\infty^\beta(P_{\text{fwd}})$ to tighten the N_δ bound.

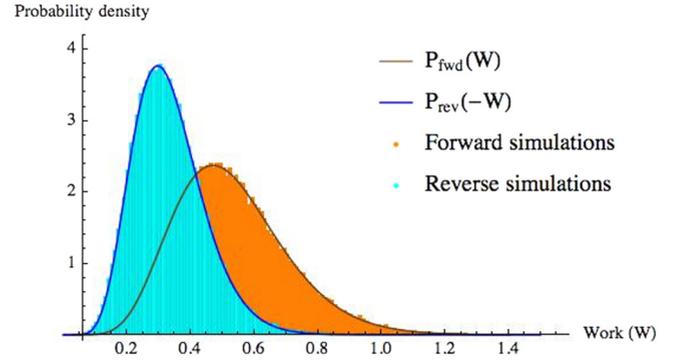


FIG. 3. Probability densities and numerical data for a weakly interacting dilute gas. We considered a gas undergoing compression (a forward protocol) and expansion (a reverse protocol). The probability per unit energy that any particular trial will involve an amount W of work [Eq. (33)] was calculated in [15]. The short, right-hand, brown curve represents $P_{\text{fwd}}(W)$. The tall, left-hand, dark-blue curve represents $P_{\text{rev}}(-W)$. By sampling work values from these distributions, we effectively simulated each protocol 10^5 times. The cyan bars (under the left-hand curve) depict the data gathered from the forward-protocol samples. The orange bars (under the right-hand curve) depict the data from the reverse-protocol samples.

Figure 4 illustrates our results. Possible values of w^δ appear along the abscissa. The blue (gently sloping) curve shows the N_δ bound, calculated from forward-trial samples, in Theorem 1. The red (staggered) curve, calculated from reverse-trial samples, shows after how many reverse trials (N_{true}) $W \geq w^\delta$ was extracted during one trial. N_{true} has a jagged, steplike shape, as one might expect.

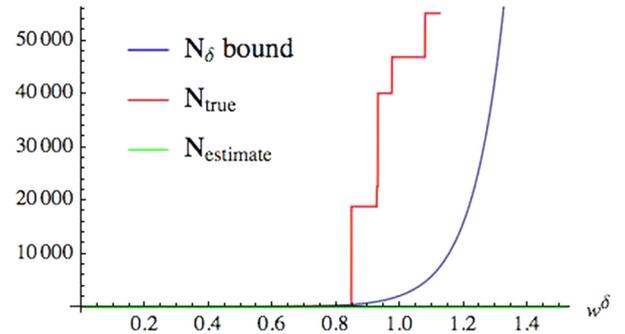


FIG. 4. Three number-of-required-trial measures. The abscissa shows possible choices of the threshold w^δ for w^δ -dominant work values. The blue (gently sloping) curve, calculated from 10^5 forward-trial samples, represents the bound on the number N_δ of reverse trials expected to be performed before any trial outputs a w^δ -dominant amount $W \geq w^\delta$ of work (Theorem 1). The red (staggered) curve, calculated from 10^5 reverse-trial samples, depicts the actual number N_{true} of trials performed before $W \geq w^\delta$ is extracted. The green curve (flat, nearly coincident with the abscissa) was calculated from forward-trial samples. This green curve represents relation (6): an estimate N_{est} of the number of trials required to extract a dominant amount of work, wherein the meaning of “dominant” is unspecified. The blue (gently sloping) curve follows the red (staggered) curve's shape more faithfully than the green (flat) does, illustrating the precision of Theorem 1. As expected, the blue (gently sloping) curve lower bounds the red (staggered) at most w^δ values.

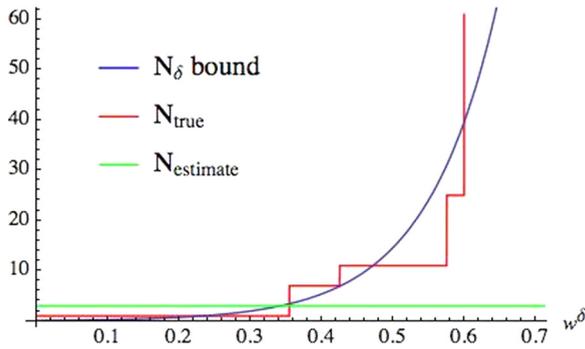


FIG. 5. Three number-of-required-trial measures at low threshold work values w^δ . At most threshold values w^δ , the N_δ bound (blue, gently sloping) lower bounds the actual number N_{true} (red, staggered) of reverse trials performed before any trial outputs a w^δ -dominant amount $W \geq w^\delta$ of work. At low w^δ values, the red curve zigzags across the blue (gently sloping). This zigzagging stems from the technical definition of N_δ .

The green curve (flat) lies close to the abscissa. This curve depicts the estimate, in [13], of the number of reverse trials expected to be performed before one trial outputs a dominant work value, for an unspecified meaning of “dominant.” We calculated $N_{\text{est}} = 3$ by simulating forward trials, calculating the average dissipated work, and substituting into relation (6).

The curves’ shapes and locations illustrate the N_δ bound’s advantages. The bound (the blue, gently sloping curve) hugs the actual number N_{true} of trials required (the red, staggered curve) more closely than N_{est} (the green, flat curve) does. N_{est} remains flat, whereas the N_δ bound rises as N_{true} rises. The N_δ bound often lower bounds N_{true} , as expected. When w^δ is small, the N_δ bound weaves above and below N_{true} , as shown in Fig. 5. The reason was explained above Theorem 1: N_δ

denotes the number of trials *expected*, in a sense defined by probability and frequency, to be required. One might get lucky and extract $W \geq w^\delta$ before performing N_δ trials. The dropping of the N_{true} curve below the N_δ bound represents such luck. But one expects to perform N_δ trials, and the N_δ bound lower bounds N_{true} for most w^δ values.

IX. CONCLUSIONS

We have sharpened predictions about the number of experimental trials required to estimate ΔF from fluctuation relations. We improved the approximation in [13] to an inequality, tightened the bound (in scenarios of interest) with an entropy H_∞^β , freed the experimenter to choose a tolerance w^δ for dominance, and approximately bounded the error in an estimate of ΔF . How to choose w^δ merits further investigation. We wish to be able to specify the greatest error ϵ acceptable in an estimate of ΔF . From ϵ , we wish to infer the number N^ϵ of trials we should expect to perform. This entire investigation improves the rigor with which free-energy differences ΔF can be estimated from experimental and numerical-simulation data.

ACKNOWLEDGMENTS

N.Y.H. thanks Yi-Kai Liu for conversations about error probability and thanks Alexey Gorshkov for hospitality at QuICS. Part of this research was conducted while N.Y.H. was visiting the QuICS and the UMD Department of Chemistry and Biochemistry. N.Y.H. was supported by an IQIM Fellowship and NSF Grant No. PHY-0803371. The Institute for Quantum Information and Matter (IQIM) is an NSF Physics Frontiers Center supported by the Gordon and Betty Moore Foundation. C.J. was supported by NSF Grant No. DMR-1506969.

-
- [1] A. Pohorille, C. Jarzynski, and C. Chipot, *J. Phys. Chem. B* **114**, 10235 (2010).
 - [2] C. M. Rohwer, F. Angeletti, and H. Touchette, *Phys. Rev. E* **92**, 052104 (2015).
 - [3] N. Lu and D. A. Kofke, *J. Chem. Phys.* **114**, 7303 (2001).
 - [4] N. Lu and D. A. Kofke, *J. Chem. Phys.* **115**, 6866 (2001).
 - [5] J. Gore, F. Ritort, and C. Bustamante, *Proc. Natl. Acad. Sci. USA* **100**, 12564 (2003).
 - [6] D. Wu and D. A. Kofke, *J. Chem. Phys.* **121**, 8742 (2004).
 - [7] D. Wu and D. A. Kofke, *Phys. Rev. E* **70**, 066702 (2004).
 - [8] D. Wu and D. A. Kofke, *J. Chem. Phys.* **123**, 054103 (2005).
 - [9] D. Wu and D. A. Kofke, *J. Chem. Phys.* **123**, 084109 (2005).
 - [10] D. A. Kofke, *Mol. Phys.* **104**, 3701 (2006).
 - [11] A. M. Hahn and H. Then, *Phys. Rev. E* **80**, 031111 (2009).
 - [12] S. Kim, Y. W. Kim, P. Talkner, and J. Yi, *Phys. Rev. E* **86**, 041130 (2012).
 - [13] C. Jarzynski, *Phys. Rev. E* **73**, 046105 (2006).
 - [14] N. Yunger Halpern, A. J. P. Garner, O. C. O. Dahlsten, and V. Vedral, *New J. Phys.* **17**, 095003 (2015).
 - [15] G. E. Crooks and C. Jarzynski, *Phys. Rev. E* **75**, 021116 (2007).
 - [16] M. Campisi, P. Hänggi, and P. Talkner, *Rev. Mod. Phys.* **83**, 771 (2011).
 - [17] G. E. Crooks, *Phys. Rev. E* **60**, 2721 (1999).
 - [18] C. Jarzynski, *Phys. Rev. Lett.* **78**, 2690 (1997).
 - [19] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (John Wiley & Sons, New York, 2012).
 - [20] R. Renner and S. Wolf, in *Proceedings of International Symposium on Information Theory, ISIT 2004* (IEEE, New York, 2004), p. 232.
 - [21] D. Collin *et al.*, *Nature (London)* **437**, 231 (2005).
 - [22] H. T. Quan and H. Dong, [arXiv:0812.4955](https://arxiv.org/abs/0812.4955).
 - [23] H. Tasaki, [arXiv:cond-mat/0009244](https://arxiv.org/abs/cond-mat/0009244).
 - [24] J. Kurchan, [arXiv:cond-mat/0007360](https://arxiv.org/abs/cond-mat/0007360).