Near-optimal protocols in complex nonequilibrium transformations

Todd R. Gingrich[a,b,1], Grant M. Rotskoff[c], Gavin E. Crooks[d,e], and Phillip L. Geissler[b,c,d,f]

[a]Physics of Living Systems Group, Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139; [b]Department of Chemistry, University of California, Berkeley, CA 94720; [c]Biophysics Graduate Group, University of California, Berkeley, CA 94720; [d]Physical Biosciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720; [e]Kavli Energy NanoSciences Institute, Berkeley, CA 94720; and [f]Chemical Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720

Edited by Ken A. Dill, Stony Brook University, Stony Brook, NY, and approved July 19, 2016 (received for review April 19, 2016)

The development of sophisticated experimental means to control nanoscale systems has motivated efforts to design driving protocols that minimize the energy dissipated to the environment. Computational models are a crucial tool in this practical challenge. We describe a general method for sampling an ensemble of finite-time, nonequilibrium protocols biased toward a low average dissipation. We show that this scheme can be carried out very efficiently in several limiting cases. As an application, we sample the ensemble of low-dissipation protocols that invert the magnetization of a 2D Ising model and explore how the diversity of the protocols varies in response to constraints on the average dissipation. In this example, we find that there is a large set of protocols with average dissipation close to the optimal value, which we argue is a general phenomenon.

nonequilibrium | irreversible thermodynamics | entropic sampling methods | Metropolis algorithm | Ising model

When a system is guided gradually from one equilibrium state to another, the amount of heat dissipated into its surroundings is insensitive to the manner of driving. In the more realistic case of an irreversible transformation in finite time, however, the dissipation can vary greatly from one driving protocol to another. These basic tenets of thermodynamics have received renewed attention in recent years due to improved capabilities for manipulating systems at small scales (1–7) and advances in the theoretical understanding of nonequilibrium fluctuations (8–10). In particular, many studies have sought to identify which finite-time protocols transform a system with the minimum amount of dissipation (11–18). Protocols that are optimal in this sense provide the most efficient route to measuring equilibrium free-energy differences—in simulations and in experiments (19)—via the Jarzynski relation (14, 20). More generally, low-dissipation protocols provide insight into the optimal design of nanoscale machines, both synthetic (4, 5) and natural (21).

However, it remains challenging to identify the minimum-dissipation protocol for complex, many-body systems driven far from equilibrium, despite recent progress (22–24). The difficulty of computing strictly optimal protocols motivates a pragmatic question: how large is the set of nearly optimal protocols? In this paper, we develop a framework to characterize that set. We introduce an entropy that indicates how many different protocols realize the same value of dissipation. For low values of dissipation, this protocol entropy quantifies how prevalent the near-optimal protocols are, highlighting when the system may be efficiently driven in many different ways. In analogy with common techniques of statistical physics, we present Monte Carlo methods to numerically compute the entropy by sampling protocols with a preference for low dissipation. The samples generated by this procedure demonstrate the distinct ways in which the system can be driven while maintaining the expectation of low dissipation. Variation among the sampled protocols accentuates features that are unimportant for ensuring low dissipation; similarly, the lack of variation highlights features that are essential for this goal. These ideas and capabilities complement previous approaches to determining optimal control procedures, using tools with many similar features (22, 24, 25). We elaborate on the connections in *Discussion*.

We illustrate our protocol-sampling framework with a numerical study of spin inversion of a ferromagnet, an essential process for copying information encoded in magnetic storage devices. Reducing dissipation in this context is of practical interest because thermodynamic costs of copying and erasing bits is projected to account for a significant fraction of future computational energy expenditures (26, 27). We examine a simple microscopic model of this process, based on the 2D Ising model (Fig. 1). At low temperature and in the presence of an external magnetic field, spins align strongly in the direction of the field. By adjusting the magnetic field and the temperature as functions of time, the magnetization may be rapidly inverted with a dissipation that depends on the manner in which the field and temperature are changed. Analysis of the protocol entropy in this model indicates that a large set of nonoptimal protocols can be used to control the system with a dissipation comparable to that of the optimal protocol.

Protocol Entropy

We first consider an ensemble of protocols $\Lambda(t)$ sharing the same value ω of average dissipation. The protocol entropy $S(\omega)$ of this ensemble measures the density of protocols that have mean

Significance

Classical thermodynamics was developed to help design the best protocols for operating heat engines that remain close to equilibrium at all times. Modern experimental techniques for manipulating microscopic and mesoscopic systems routinely access far-from-equilibrium states, demanding new theoretical tools to describe the optimal protocols in this more complicated regime. Prior studies have sought, in simple models, the protocol that minimizes dissipation. We use computational tools to investigate the diversity of low-dissipation protocols. We show that optimal protocols can be accompanied by a vast set of near-optimal protocols, which still offer the substantive benefits of the optimal protocol. Although solving for the optimal protocol is typically difficult, computationally identifying a near-optimal protocol can be comparatively easy.

Author contributions: T.R.G., G.M.R., G.E.C., and P.L.G. designed research; T.R.G. and G.M.R. performed research; T.R.G., G.M.R., G.E.C., and P.L.G. analyzed data; and T.R.G., G.M.R., G.E.C., and P.L.G. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence should be addressed. Email: toddging@mit.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1606273113/-/DCSupplemental.

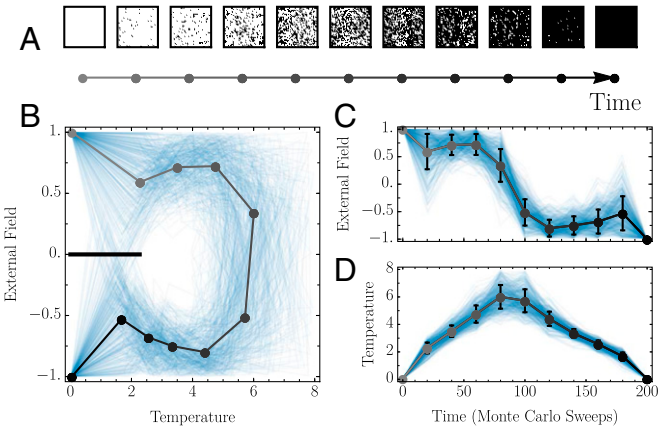

Fig. 1. Low-dissipation protocols that invert a 2D Ising magnet in finite time. (*A*) Snapshots of the 40×40 periodically replicated magnet during field inversion. (*B*) The 450 representative samples of low-dissipation protocols (blue), collected from the $\lambda=0.5$, $N=5$ protocol ensemble (Eq. **6**). Protocols are discretized into 10 intervals in the temperature vs. field plane. Each interval represents a fixed amount of time (20 Monte Carlo sweeps), and the temperature and field values are linearly interpolated between the endpoints. (*C*) The external field as a function of time. (*D*) The temperature as a function of time. In *B*–*D*, lines ranging from gray to black indicate averages over the 450 protocols, with shading corresponding to the times of the snapshots in *A*.

dissipation ω. In analogy to the standard microcanonical ensemble of statistical mechanics, we define the entropy

$$S(\omega)=\ln\left[\Omega_0\int\mathcal{D}\Lambda(t)\;\delta(\omega-\langle\omega\rangle_\Lambda)\right], \qquad [1]$$

where the integral runs over the space of time-dependent protocols $\Lambda(t)$ and Ω_0 is a constant that sets the arbitrary zero of entropy. The delta function picks out protocols whose average dissipation $\langle\omega\rangle_\Lambda$ lies within an infinitesimal interval around the specified value of ω. The Λ subscript denotes an average taken over the probability distribution $P_{\text{traj}}[x(t)|\Lambda(t)]$ of stochastic trajectories evolving under the fixed protocol $\Lambda(t)$:

$$\langle\omega\rangle_\Lambda=\int\mathcal{D}x(t)\;P_{\text{traj}}[x(t)|\Lambda(t)]\omega[x(t),\Lambda(t)]. \qquad [2]$$

For a single trajectory $x(t)$, the dissipation ω can be cast in terms of an imbalance between forward and reversed dynamics (10) giving

$$\omega[x(t),\Lambda(t)]=\ln\frac{P_{\text{traj}}[x(t)|\Lambda(t)]}{P_{\text{traj}}\left[\tilde{x}(t)|\tilde{\Lambda}(t)\right]}. \qquad [3]$$

Tildes signify time reversal, so the numerator and denominator are probabilities of forward and reverse trajectories, respectively.

Optimal protocols, which carry some minimal dissipation, represent a small fraction of the possible protocols. Consequently, the protocol entropy evaluated at this minimal dissipation is low. As the dissipation increases, the entropy increases, and the growth rate reflects how flexibly low-dissipation protocols may be constructed. In particular, rapid entropic growth near the minimum dissipation suggests that targeting an exact optimal protocol is both challenging and gratuitous. Many other protocols, in practice, will perform comparably to the optimum.

The limiting behavior of $S(\omega)$ near the minimum dissipation ω^* can be calculated exactly, assuming only that $\langle\omega\rangle_\Lambda$ depends smoothly on variations in the low-dissipation protocols. As shown in *Supporting Information*, protocol entropy grows logarithmically in this regime:

$$S(\omega)=\text{const}+\left(\frac{n}{2}-1\right)\ln(\omega-\omega^*), \qquad [4]$$

where n is the total number of protocol degrees of freedom, that is, the total number of parameter values defining any given protocol. Quite generally, therefore, protocol entropy increases sharply over a narrow range of dissipation values just above the minimum ω^*.

A Canonical Protocol Ensemble

To numerically compute the protocol entropy, it is useful to introduce a canonical protocol ensemble with distribution

$$P_{\text{canon}}[\Lambda(t)]\propto e^{-\gamma\langle\omega\rangle_\Lambda}. \qquad [5]$$

In correspondence with the canonical ensemble of statistical mechanics, $\langle\omega\rangle_\Lambda$ acts as an effective energy for each protocol and γ plays the role of inverse temperature, tuning the mean dissipation, that is, the average of ω over the distribution P_{canon}. Searching for a protocol with strictly minimum dissipation amounts to a zero "temperature" ($\gamma\to\infty$) quench, whereas near-optimal protocols are identified by large values of γ.

When using a sufficiently large γ, the samples reveal representative low-dissipation protocols. By using different biasing strengths γ in Eq. **5**, we can learn about the characteristics of protocols with distinct values of average dissipation. In addition, the protocols sampled with various choices of γ can be combined to calculate $S(\omega)$ over a broad range.

Sampling Protocols and Trajectories. In principle, the ensemble defined by Eq. **5** may be directly sampled with a Monte Carlo procedure that conditionally accepts protocol changes based on the corresponding changes in $\langle\omega\rangle_\Lambda$. For complex systems, however, values of $\langle\omega\rangle_\Lambda$ are typically not known exactly. They can be estimated from the sample mean $\overline{\omega}_\Lambda=N^{-1}\sum_i\omega[x_i(t)|\Lambda(t)]$ of a collection of N trajectories drawn from $P_{\text{traj}}[x(t)|\Lambda(t)]$. However, for finite N, replacing $\langle\omega\rangle_\Lambda$ by $\overline{\omega}_\Lambda$ in Eq. **5** yields a distribution of protocols that differs from $P_{\text{canon}}[\Lambda(t)]$. Strategies to correct for the finite-N bias have been formulated to enable conventional Boltzmann sampling when configurational energies cannot be calculated with certainty (28–32). Here, we consider an analogous strategy in the context of sampling protocols.

To sample the canonical protocol distribution, we construct a Monte Carlo procedure that performs a random walk through the joint space of protocols and N independent trajectories $x_1(t),x_2(t),\ldots x_N(t)$. A trial move amounts to an attempt to make changes in both $\Lambda(t)$ and in $\{x_i(t)\}$. Operationally, this proposed change can be achieved by first perturbing the protocol (with a symmetric generation probability) before regenerating the trajectories using the new protocol. For simplicity, we consider the case that the trajectory generation probabilities are also symmetric, as in the noise-guided shooting procedures of transition path sampling (33–36). Accepting a trial move with the Metropolis probability $\min[1,\exp(-\lambda\Delta\overline{\omega})]$, where $\Delta\overline{\omega}$ is the difference between the sample means under the original and trial protocols, yields the stationary distribution

$$P_{\text{sampled}}[x_1(t),\ldots,x_N(t),\Lambda(t)]\propto\exp\left(-\frac{\lambda}{N}\sum_{i=1}^{N}\omega[x_i(t)|\Lambda(t)]\right). \qquad [6]$$

The resulting marginal distribution of protocols,

$$P_{\text{sampled}}[\Lambda(t)]\propto\left\langle e^{-\lambda\omega/N}\right\rangle_\Lambda^N, \qquad [7]$$

is determined by the dissipation statistics of each protocol, but in a more complicated way than P_{canon}. Nevertheless, we show that

the sampled protocols are drawn from a canonical protocol distribution in two special situations: the case of Gaussian dissipation distributions and the large N limit.

Cumulant Expansion. The expression for the marginal distribution Eq. **7** may be recast as

$$P_{\text{sampled}}[\Lambda(t)] \propto e^{N\psi_\Lambda(-\lambda/N)}, \quad [8]$$

where $\psi_\Lambda(k) = \ln\langle e^{k\omega}\rangle_\Lambda$ is the cumulant-generating function for the dissipation. A cumulant expansion,

$$\psi_\Lambda(-\lambda/N) = -\frac{\lambda\langle\omega\rangle_\Lambda}{N} + \frac{\lambda^2\langle\delta\omega^2\rangle_\Lambda}{2N^2} - \mathcal{O}\left(\frac{\lambda^3}{N^3}\right), \quad [9]$$

relates the probability of sampled protocols to the cumulants of the dissipation distribution $P(\omega|\Lambda(t))$, where $\delta\omega = \omega - \langle\omega\rangle_\Lambda$. As N grows, the contribution from higher-order cumulants vanishes, consistent with the central limit theorem. In the limit $N\to\infty$, sampling Eq. **5** becomes equivalent to sampling Eq. **6** because the sample mean converges to the average dissipation.

In the special case that $P(\omega|\Lambda(t))$ is Gaussian for each protocol $\Lambda(t)$, a powerful simplification arises, averting the need to use large values of N. Gaussian dissipation distributions occur in many contexts—as a defining feature of linear response (37), in the limit of slow adiabatic driving (37), and when Brownian particles evolve in driven harmonic potentials (38). In all these cases, the cumulants of $P(\omega|\Lambda(t))$ beyond the variance vanish, allowing us to exactly truncate Eq. **9** at second order. If we further take $\Lambda(t)$ to be symmetric under time reversal, then the fluctuation theorem provides an exact relationship between the mean and variance: $\langle\delta\omega^2\rangle_\Lambda = 2\langle\omega\rangle_\Lambda$. As a result, the biased protocol distribution can be expressed in terms of mean dissipation alone:

$$P_{\text{Gaussian}}[\Lambda(t)] \propto e^{-\lambda(1-\lambda/N)\langle\omega\rangle_\Lambda}. \quad [10]$$

Eq. **10** has precisely the form of the canonical protocol distribution (Eq. **5**), with an effective bias $\gamma = \lambda(1-\lambda/N)$. This result offers tremendous flexibility. An exact bias toward low average dissipation can be achieved with any N, for example, by sampling a small number of trajectories for each proposed change in protocol. Because generating trajectories dominates the computational expense of our sampling scheme, the freedom to choose small N is very attractive.

The limitation with using small N when sampling protocols is that the achievable bias strength γ cannot exceed $\gamma_{\max} = N/4$. This constraint arises because the λ bias in Eq. **5** directly favors low-dissipation trajectories, and not necessarily low-dissipation protocols. For small values of λ, the trajectories sampled for a given protocol are typical of the unbiased trajectory distribution $P_{\text{traj}}[x(t)|\Lambda(t)]$. In this case, there is thus a strong correlation between sampled low-dissipation trajectories and protocols that yield low dissipation on average. This correspondence degrades for large λ. In fact, for $\lambda > N/2$, sampled trajectories have negative dissipation on average,[†] which cannot be typical of any protocol according to the second law of thermodynamics. Thus, as λ is increased toward $N/2$, the joint ensemble of trajectories and protocols switches from highlighting low-dissipation protocols to emphasizing rare negative-dissipation trajectories. Moreover, sampling with large values of λ requires generation of increasingly rare trajectories, complicating efficient path sampling as discussed in *Supporting Information*.

[†]The cumulant-generating function $\psi_\Lambda(k)$ is symmetric about $k=-1/2$ (39), and $P_{\text{sample}}[\Lambda(t)]$ is correspondingly symmetric in λ about $\lambda = N/2$. This symmetry implies negative average dissipation when $\lambda > N/2$.

Results

Spin Inversion Protocols. To illustrate the use of the low-dissipation protocol ensemble, we consider the inversion of spins in a ferromagnet. Specifically, we imagine initializing a system of interacting spins at low temperature, where its equilibrium state has long-range "up" or "down" order. We then ask how best to flip this "bit." That is, how should we vary the temperature and external field as functions of time to flip the state of the magnet without excess dissipation? This problem, relevant to the design of low-power magnetic hard drives, has been investigated as an optimal control problem elsewhere (23, 40). Here, we also consider the near-optimal drivings.

We represent the ferromagnet as a 2D Ising model with periodic boundary conditions and dynamics generated by a succession of individual spin flips. With an external magnetic field h, the energy of a configuration is given by

$$E = -h\sum_i \sigma_i - \sum_{\langle ij\rangle}\sigma_i\sigma_j, \quad [11]$$

where $\sigma_i = \pm 1$, and $\langle ij\rangle$ indicates a sum over nearest neighbor sites i and j. An attempted spin flip that alters the energy by ΔE is accepted with Glauber probability $e^{-\Delta E/T}/(1+e^{-\Delta E/T})$, where T is the temperature of the bath. Unlike equilibrium Ising model dynamics, the temperature and magnetic field are time dependent as prescribed by a nonequilibrium protocol $\Lambda(t) = \{T(t), h(t)\}$. In a finite amount of time t_{obs}, we aim to switch from the macroscopic up state to the down state. We therefore consider only protocols that begin and end at low temperature $[T(0) = T(t_{\text{obs}}) = 0.05]$ and that switch from a positive to a negative field $[h(0) = -h(t_{\text{obs}}) = 1]$.

One consequence of the nonequilibrium driving is that the dynamics is not microscopically reversible. For ordinary Glauber dynamics, the equilibrium probability of a trajectory segment is equal to its time-reversed counterpart, but our time-dependent driving breaks this equality. By tracking the random numbers that generate each spin flip, we explicitly compute the forward and time-reversed probabilities of each trajectory, thereby computing the stochastic thermodynamic dissipation via Eq. **3**. Physically, the dissipation of each microscopic step multiplied by T is the heat transferred from the thermal bath into the system.

We use Monte Carlo techniques, discussed further in *Materials and Methods*, to sample low-dissipation protocols from Eq. **6** with $\lambda = 0.5, N = 5$. Fig. 1 shows 450 representative protocols, all of which avoid the region of parameter space near the Ising critical point. Control is particularly costly (23) in this vicinity due to critical slowing down, which causes the spin system to lag behind changes in the control parameters. There is a natural connection between dissipation and lag (41, 42): the farther the system falls out of equilibrium with control parameters' instantaneous values, the more heat is dissipated to the reservoir during the relaxation.

Roughly, the optimal strategy requires that we first heat the magnet, next invert the field, and then cool the magnet. However, the varied protocols in Fig. 1 demonstrate significant leeway in how these steps are carried out. Most notably, while the system is held at low temperature, the magnetic field need not be precisely tuned, as evidenced by large variations both early and late in the protocol. Some low-dissipation protocols even transiently invert the field at low temperature, thereby crossing the equilibrium coexistence curve, only to restore the field's original sign a short time later. This seemingly wasteful procedure in fact incurs little dissipative cost, because it is highly ineffectual. The low-temperature field inversion is too brief for nucleation of the new phase to occur with significant probability, so the extent of relaxation is negligible. Absent relaxation, no heat is dissipated

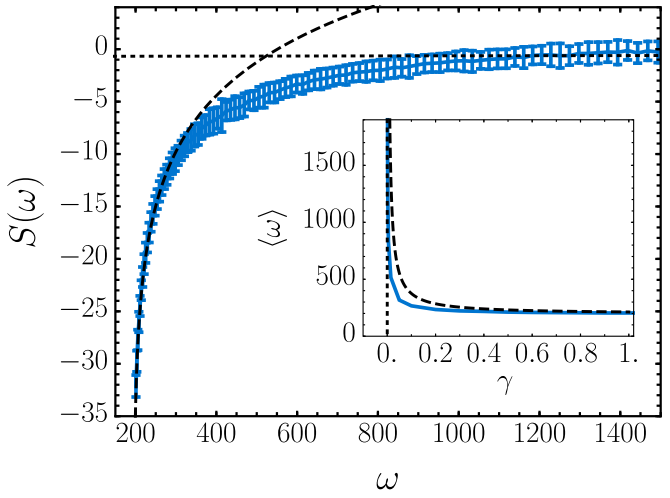

Fig. 2. Protocol entropy quantifies the diversity of protocols with average dissipation ω. By sampling the distribution equation (Eq. **5**) with $t_{obs} = 200$ Monte Carlo sweeps and using various biasing strengths γ, the protocol entropy $S(\omega)$ was computed using the MBAR method (43). The slope of the protocol entropy at $\omega = \langle\omega\rangle$ gives the strength of the bias necessary to yield mean dissipation $\langle\omega\rangle$. Asymptotic limits for small and large γ are plotted as the dotted and dashed lines, respectively. (*Inset*) The average dissipation for protocols as a function of the bias γ. Note that a small bias produces most of the achievable dissipation reduction.

to the bath. Perhaps counterintuitively, the lag is so severe in this case as to be irrelevant.

Protocol Entropy. The protocol entropy determined by sampling magnetization inversion dynamics is shown in Fig. 2. Near the apparent minimum dissipation, a small increase in mean dissipation is accompanied by a steep rise in $S(\omega)$, that is, the number of protocols grows rapidly as we permit modest excess dissipation. In fact, the density of protocols increases by several orders of magnitude for a change in dissipation that is very small relative to dissipation fluctuations for a fixed protocol. This rapid initial growth is captured well by the asymptotic form of Eq. **4**, which depends on the number of degrees of freedom in the protocol. Farther from the minimal value of mean dissipation, the protocol entropy climbs much more gradually.

The slope of $S(\omega)$ reflects the strength of bias γ needed to depress the average dissipation. The *Inset* of Fig. 2 illustrates a crossover between two regimes: small biases greatly reduce the mean dissipation but further reduction requires very large biases. Thus, weak biases on $\langle\omega\rangle_\Lambda$ can be greatly effective at directing the protocol sampling toward the optimum. We anticipate that these limiting behaviors are generic, and the corresponding asymptotic forms are derived in *Supporting Information*. The reduction in dissipation due to small values of γ is governed by the variability of $\langle\omega\rangle_\Lambda$ in an unbiased protocol ensemble. Because complex systems typically depend sensitively on one or more of their control parameters, this variability should be substantial in general. Large values of γ favor nearly optimal protocols, whose diversity is well described by Eq. **4**. Correspondingly, the mean dissipation in the large γ limit decays slowly as $\langle\omega\rangle = \omega^* + n/(2\gamma)$.

Gaussian Fluctuations. We have computed $S(\omega)$ for the spin inversion process both with and without the simplifying assumption that the dissipation is Gaussian distributed. We find that the Gaussian approximation provides an estimate that is accurate within statistical error despite requiring a significantly reduced number of trajectories. To more explicitly demonstrate the validity of the approximation, we selected three protocols from our sampling, which are shown in Fig. 3*A*.

For each protocol, we computed the dissipation distribution $P(\omega|\Lambda)$. Empirically, we find that these distributions are strikingly Gaussian over a large range of ω that includes $\omega = 0$. At large positive values of dissipation, we observe "fat" exponential tails, consistent with the structure of generic current large deviation functions that has recently been demonstrated for the case of time-independent driving (44, 45). This fat tail, associated with clusters of spins that resist reorientation, only weakly restricts our use of the Gaussian dissipation assumption. The positive λ bias, useful to study low-dissipation behavior, focuses the sampling toward the Gaussian region of the distribution, rendering the exponential tails insignificant.

Discussion

Low dissipation, the focus of our exploration of driving protocols, is one of many possible objectives for nonequilibrium control. Indeed, minimizing dissipation can be viewed as an instance of the extensively studied problem of stochastic optimal control (46–50).

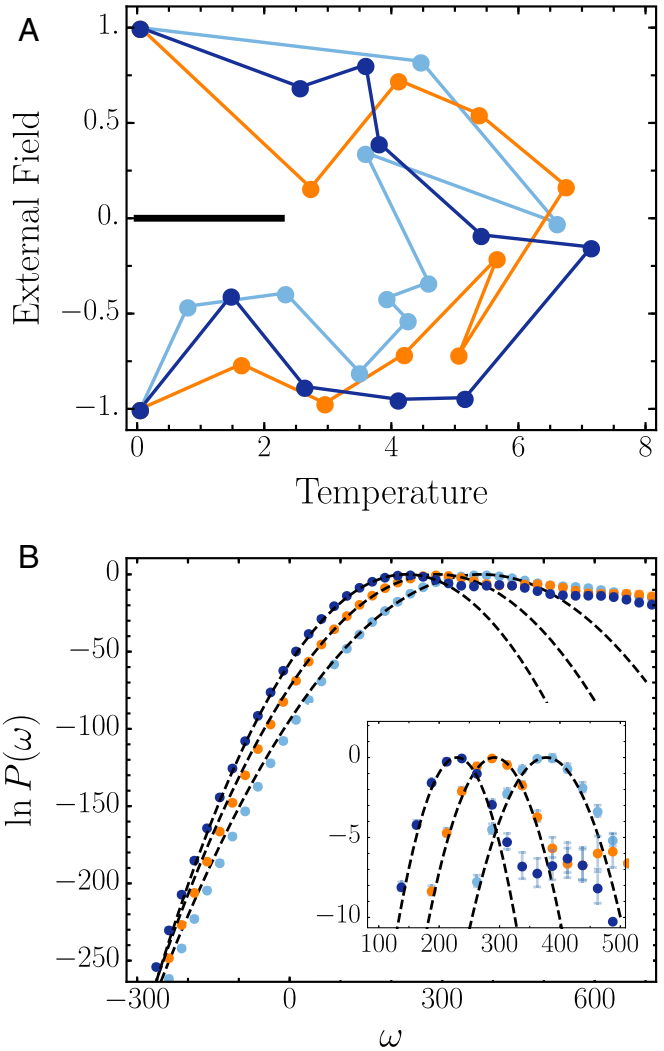

Fig. 3. Dissipation distributions for representative protocols. (*A*) Three randomly selected protocols from the ensemble equation (Eq. **6**) are plotted on the temperature, external field plane as in Fig. 1. (*B*) The distribution of dissipation values $P(\omega|\Lambda)$ for the three protocols, displayed with corresponding colors. Dashed black lines show Gaussian distributions with the same means $\langle\omega\rangle_\Lambda$ as the sampled distributions and with variances $2\langle\omega\rangle_\Lambda$. (*Inset*) The neighborhood around the average dissipation values is shown in greater detail.

The formulation of stochastic optimal control as a path integral control problem (25, 51–53) has a particularly close connection to our work, in that importance sampling of trajectories can be used to iteratively refine a protocol toward the optimum (22, 24). A distinct feature of our approach is the systematic sampling of protocols, as opposed to the more limited goal of strictly optimizing them. By simultaneously sampling trajectories and protocols with a well-defined bias, we identify low-dissipation protocols without the challenging task of converging trajectory-space averages for any particular protocol, as is required in iterative optimization methods. Protocols harvested by our procedure, in contrast to those encountered in an optimization, quantitatively reflect the diversity of low-dissipation protocols.

Scrutinizing near-optimal protocols complements the search for optimality in several ways. In simple model systems that can be optimized exactly, minimum-dissipation protocols are known to involve features that are singular or may be impractical to implement (14). In such cases, the collection of near-optimal protocols becomes a natural target for design. Even when the optimal protocol may be physically achieved, its form does not directly indicate which features of a driving history are essential to its success, and which are irrelevant. In our approach, relevant features can be readily identified through their limited variability in the protocol ensemble, as illustrated by the cloud of protocols in Fig. 1*B*.

Finally, we note that efficient but suboptimal nonequilibrium transformations are almost certainly the norm in biology at many scales. Indeed, the evolutionary dynamics of biological adaptation might be viewed as an importance sampling on the space of protocols, roughly akin to the sampling methods developed in this paper. The surprising, often eccentric strategies used to perform simple tasks in biology are, perhaps, indicative of the myriad options provided by an ensemble of protocols evolving under a complex set of constraints.

Materials and Methods

We sample the joint space of trajectories and protocols using Markov chain Monte Carlo methods. Each point in this space consists of a protocol and N independent trajectories subject to that protocol. With tunable bias λ, the Markov chain samples the distribution in Eq. **6**. We restrict the space of protocols with a view toward experimental practicality. For the 2D Ising example, we impose two such restrictions. First, to allow only slowly varying protocols, we parameterize our protocol space by the values of temperature and external field at 11 evenly spaced times. We call the values at these special times the control points. Between any two neighboring control points, the temperature and field strength are linearly interpolated. Second, we require that $-1 \le h \le 1$ and $0 \le T \le 8$ for all control points to mimic that physical apparatuses can tune controls over bounded ranges. Limiting the protocol space in these ways can be viewed as a regularization scheme that simplifies the representation problem of optimal control theory (24). We note, however, that the protocol entropy depends on our choice of regularization. If, for example, we were to use many more control points with the same t_{obs}, then we would introduce many additional protocols with high-frequency features. Consequently, $S(\omega)$ would grow more rapidly, following Eq. **4**. Fig. S1 highlights this dependence on the number of protocol degrees of freedom.

Each Monte Carlo move first attempts to adjust the protocol by moving a single control point by a random displacement in the temperature–field plane. The move is constructed to be symmetric, meaning the probability of selecting any displacement vector equals the probability of a displacement of the opposite sign. Using this trial protocol, N new trajectories are simulated using a sequence of Glauber single spin flips. Conventionally, each step of the Glauber Ising dynamics chooses a random spin, which is flipped to generate a trial configuration. To enable more efficient noise-guided trajectory sampling, we use a modified Glauber dynamics: the trial configuration is given by setting the randomly selected spin to either the up or the down state without reference to its prior state (36). The move is futile when the selected spin is already in the trial configuration. Because one-half of the moves are futile on average, the Monte Carlo time is rescaled by a factor of 2 compared with ordinary single-spin flip Glauber dynamics. Following each move, the probability of running that step backward is computed, enabling an explicit calculation of the dissipation of each trajectory. The new protocol and trajectories are conditionally accepted with probability $\min[1, \exp(-\lambda\Delta\overline{\omega})]$, where $\Delta\overline{\omega}$ is the difference between the sample means under the original and trial protocols.

The protocol entropy is calculated, up to a constant offset $\ln\Omega_0$, using a weighted average over the protocols collected by the Monte Carlo procedure:

$$S(\omega) = \ln\left[\Omega_0 \int d\Lambda(t) \frac{P_{\text{sampled}}[\Lambda(t)]}{\langle e^{-\lambda\omega/N}\rangle_\Lambda^N} \delta(\omega - \langle\omega\rangle_\Lambda)\right]. \quad [12]$$

From a set of M sampled protocols, $\{\Lambda_1, \Lambda_2, \ldots, \Lambda_\alpha, \ldots, \Lambda_M\}$, we therefore estimate the following:

$$S(\omega) = \ln\left[\frac{\Omega_0}{M} \sum_{\alpha=1}^{M} \frac{\delta(\omega - \langle\omega\rangle_{\Lambda_\alpha})}{\langle e^{-\lambda\omega/N}\rangle_{\Lambda_\alpha}^N}\right]. \quad [13]$$

Operationally, this amounts to collecting a histogram of values of $\langle\omega\rangle_\Lambda$ with each entry weighted by the corresponding value of $\langle e^{-\lambda\omega/N}\rangle_\Lambda^N$. To generate Fig. 2, each of these weights is computed by estimating the exponential average from 1,000 independent trajectories. In practice, $S(\omega)$ is constructed using the multistate Bennett acceptance ratio (MBAR) method to combine samples collected with $N = 20$ and with several different values of the bias ranging from $\lambda = 0$ to $\lambda = 1$. The offset $\ln\Omega_0$ is chosen such that $S(\omega)$ is zero at its maximum.

The protocol entropy can be computed much more efficiently when the Gaussian approximation may be used to evaluate the exponential average. To evaluate the validity of this approximation for the Ising dynamics, we compute the actual dissipation distributions by sampling trajectories with a fixed protocol. These trajectories are importance sampled using harmonic biases, which restrain the dissipation to fluctuate around a specified value. By choosing several different harmonic biases, trajectories were biased into both tails of the dissipation distribution, which was reconstructing using the MBAR method (43).

When the Gaussian approximation is appropriate, it is wasteful to use a large value of N. Low-dissipation protocols may be sampled with an effective biasing strength $\gamma = \lambda(1 - \lambda/N)$ using various combinations of N and λ, and a small N reduces the computational expense, as demonstrated in Fig. S2. However, when N is too small or λ too large, the Monte Carlo acceptance probability drops precipitously, a fact elaborated upon in *Supporting Information*, particularly in Fig. S3. Sampling efficiency is poor under these conditions because the Markov chain favors a collection of rare trajectories with dissipation below the mean (and often below zero). This issue can be partially alleviated by introducing replica exchange moves—random swaps exchanging replicas with different biasing strengths λ. The implementation of this procedure naturally mirrors the use of replica exchange to surmount kinetic traps when sampling low-temperature molecular configurations. Further performance enhancements are obtained when trial trajectories are generated with random numbers (noises) that correlate with the noises of the previous collection of trajectories. An implementation of this noise-guided sampling is described in detail elsewhere (36, 54). The noise guidance technique is not strictly required to perform the protocol sampling, but in Fig. S4 we show that it can provide significant practical benefits. Fig. S5 shows further practical improvements.

ACKNOWLEDGMENTS. T.R.G. acknowledges support from the National Science Foundation (NSF) Graduate Research Fellowship, the Fannie and John Hertz Foundation, and the Gordon and Betty Moore Foundation as a Massachusetts Institute of Technology Physics of Living Systems Fellow. G.M.R. acknowledges support from the NSF Graduate Research Fellowship. P.L.G. was supported by the US Department of Energy, Office of Basic Energy Sciences, through the Chemical Sciences Division of the Lawrence Berkeley National Laboratory, under Contract DE-AC02-05CH11231. G.E.C. acknowledges support from the US Army Research Laboratory and the US Army Research Office under Contract W911NF-13-1-0390.

1. Smith SB, Finzi L, Bustamante C (1992) Direct mechanical measurements of the elasticity of single DNA molecules by using magnetic beads. *Science* 258(5085):1122–1126.
2. Liphardt J, Dumont S, Smith SB, Tinoco I, Jr, Bustamante C (2002) Equilibrium information from nonequilibrium measurements in an experimental test of Jarzynski's equality. *Science* 296(5574):1832–1835.
3. Collin D, et al. (2005) Verification of the Crooks fluctuation theorem and recovery of RNA folding free energies. *Nature* 437(7056):231–234.
4. Martínez IA, et al. (2016) Brownian Carnot engine. *Nat Phys* 12(1):67–70.
5. Blickle V, Bechinger C (2012) Realization of a micrometre-sized stochastic heat engine. *Nat Phys* 8(2):143–146.

6. Toyabe S, Sagawa T, Ueda M, Muneyuki E, Sano M (2010) Experimental demonstration of information-to-energy conversion and validation of the generalized Jarzynski equality. *Nat Phys* 6(12):988–992.
7. Jun Y, Gavrilov M, Bechhoefer J (2014) High-precision test of Landauer's principle in a feedback trap. *Phys Rev Lett* 113(19):190601.
8. Jarzynski C (1997) Nonequilibrium equality for free energy differences. *Phys Rev Lett* 78(14):2690.
9. Crooks GE (1999) Excursions in statistical dynamics. PhD thesis (University of California, Berkeley).
10. Spinney R, Ford I (2013) Fluctuation relations: A pedagogical overview. *Nonequilibrium Statistical Physics of Small Systems: Fluctuation Relations and Beyond*, eds Klages R, Just W, Jarzynski C (Wiley-VCH, Weinheim, Germany), pp 3–56.
11. Weinhold F (1975) Metric geometry of equilibrium thermodynamics. *J Chem Phys* 63(6):2479–2483.
12. Ruppeiner G (1979) Thermodynamics: A Riemannian geometric model. *Phys Rev A* 20(4):1608–1613.
13. Salamon P, Nitzan A, Andresen B, Berry RS (1980) Minimum entropy production and the optimization of heat engines. *Phys Rev A* 21(6):2115–2129.
14. Schmiedl T, Seifert U (2007) Optimal finite-time processes in stochastic thermodynamics. *Phys Rev Lett* 98(10):108301.
15. Gomez-Marin A, Schmiedl T, Seifert U (2008) Optimal protocols for minimal work processes in underdamped stochastic thermodynamics. *J Chem Phys* 129(2):024114.
16. Then H, Engel A (2008) Computing the optimal protocol for finite-time processes in stochastic thermodynamics. *Phys Rev E Stat Nonlin Soft Matter Phys* 77(4):041105.
17. Zulkowski PR, Sivak DA, Crooks GE, DeWeese MR (2012) Geometry of thermodynamic control. *Phys Rev E Stat Nonlin Soft Matter Phys* 86(4):041148.
18. Zulkowski PR, DeWeese MR (2015) Optimal control of overdamped systems. *Phys Rev E Stat Nonlin Soft Matter Phys* 92(3):032117.
19. Maragakis P, Ritort F, Bustamante C, Karplus M, Crooks GE (2008) Bayesian estimates of free energies from nonequilibrium work data in the presence of instrument noise. *J Chem Phys* 129(2):024102.
20. Dellago C, Hummer G (2013) Computing equilibrium free energies using non-equilibrium molecular dynamics. *Entropy (Basel)* 16(1):41–61.
21. Oster G, Wang H (2000) Why is the mechanical efficiency of F_1-ATPase so high? *J Bioenerg Biomembr* 32(5):459–469.
22. Zhang W, Wang H, Hartmann C, Weber M, Schütte C (2014) Applications of the cross-entropy method to importance sampling and optimal control of diffusions. *SIAM J Sci Comput* 36(6):A2654–A2672.
23. Rotskoff GM, Crooks GE (2015) Optimal control in nonequilibrium systems: Dynamic Riemannian geometry of the Ising model. *Phys Rev E Stat Nonlin Soft Matter Phys* 92(6):060102.
24. Kappen HJ, Ruiz HC (2016) Adaptive importance sampling for control and inference. *J Stat Phys* 162(5):1244–1266.
25. Kappen HJ (2005) Linear theory for control of nonlinear stochastic systems. *Phys Rev Lett* 95(20):200201.
26. Lambson B, Carlton D, Bokor J (2011) Exploring the thermodynamic limits of computation in integrated systems: Magnetic memory, nanomagnetic logic, and the Landauer limit. *Phys Rev Lett* 107(1):010604.
27. Hong J, Lambson B, Dhuey S, Bokor J (2016) Experimental test of Landauer's principle in single-bit operations on nanomagnetic memory bits. *Sci Adv* 2(3):e1501492.
28. Andrieu C, Roberts GO (2009) The pseudo-marginal approach for efficient Monte Carlo computations. *Ann Stat* 37(2):697–725.
29. Ball RC, Fink TMA, Bowler NE (2003) Stochastic annealing. *Phys Rev Lett* 91(3):030201.
30. Beaumont MA (2003) Estimation of population growth or decline in genetically monitored populations. *Genetics* 164(3):1139–1160.
31. Ceperley DM, Dewing M (1999) The penalty method for random walks with uncertain energies. *J Chem Phys* 110(20):9812–9820.
32. Lin L, Liu KF, Sloan J (2000) A noisy Monte Carlo algorithm. *Phys Rev D Part Fields* 61(7):074505.
33. Dellago C, Bolhuis PG, Chandler D (1998) Efficient transition path sampling: Application to Lennard-Jones cluster rearrangements. *J Chem Phys* 108(22):9236–9245.
34. Crooks GE, Chandler D (2001) Efficient transition path sampling for nonequilibrium stochastic dynamics. *Phys Rev E Stat Nonlin Soft Matter Phys* 64(2):026109.
35. Hartmann AK (2014) High-precision work distributions for extreme nonequilibrium processes in large systems. *Phys Rev E Stat Nonlin Soft Matter Phys* 89(5):052103.
36. Gingrich TR, Geissler PL (2015) Preserving correlations between trajectories for efficient path sampling. *J Chem Phys* 142(23):234104.
37. Speck T, Seifert U (2004) Distribution of work in isothermal nonequilibrium processes. *Phys Rev E Stat Nonlin Soft Matter Phys* 70(6):066112.
38. Mazonka O, Jarzynski C (1999) Exactly solvable model illustrating far-from-equilibrium predictions. arXiv:cond-mat/9912121.
39. Lebowitz JL, Spohn H (1999) A Gallavotti–Cohen-type symmetry in the large deviation functional for stochastic dynamics. *J Stat Phys* 95(1-2):333–365.
40. Venturoli M, Vanden-Eijnden E, Ciccotti G (2009) Kinetics of phase transitions in two dimensional Ising models studied with the string method. *J Math Chem* 45(1): 188–222.
41. Kawai R, Parrondo JMR, Van den Broeck C (2007) Dissipation: The phase-space perspective. *Phys Rev Lett* 98(8):080602.
42. Vaikuntanathan S, Jarzynski C (2009) Dissipation and lag in irreversible processes. *EPL* 87(6):60005.
43. Shirts MR, Chodera JD (2008) Statistically optimal analysis of samples from multiple equilibrium states. *J Chem Phys* 129(12):124105.
44. Gingrich TR, Horowitz JM, Perunov N, England JL (2016) Dissipation bounds all steady-state current fluctuations. *Phys Rev Lett* 116(12):120601.
45. Pietzonka P, Barato AC, Seifert U (2016) Universal bounds on current fluctuations. *Phys Rev E Stat Nonlin Soft Matter Phys* 93(5):052145.
46. Fleming WH (1971) Stochastic control for small noise intensities. *SIAM J Control* 9(3): 473–517.
47. Fleming WH (1977) Exit probabilities and optimal stochastic control. *Appl Math Optim* 4(1):329–346.
48. Dupuis P, Wang H (2004) Importance sampling, large deviations, and differential games. *Stochastics* 76(6):481–508.
49. Hartmann C, Schütte C (2012) Efficient rare event simulation by optimal nonequilibrium forcing. *J Stat Mech* 2012(11):P11004.
50. Chetrite R, Touchette H (2015) Variational and optimal control representations of conditioned and driven processes. *J Stat Mech* 2015(12):P12001.
51. Theodorou E, Buchli J, Schaal S (2010) A generalized path integral control approach to reinforcement learning. *J Mach Learn Res* 11:3137–3181.
52. Chernyak VY, Chertkov M, Bierkens J, Kappen HJ (2014) Stochastic optimal control as non-equilibrium statistical mechanics: Calculus of variations over density and current. *J Phys A* 47(2):022001.
53. Thijssen S, Kappen HJ (2015) Path integral control and state-dependent feedback. *Phys Rev E Stat Nonlin Soft Matter Phys* 91(3):032104.
54. Gingrich TR (2015) Two paths diverged: Exploring trajectories, protocols, and dynamic phases. PhD thesis (University of California, Berkeley).

Supporting Information

Gingrich et al. 10.1073/pnas.1606273113

General Features of the Protocol Entropy

In the limit of large values of the bias parameter γ, the difference between a sampled protocol $\Lambda(t)$ and the optimal protocol $\Lambda^*(t)$ should be small. If the average dissipation $\langle\omega\rangle_\Lambda$ is a smooth functional of the protocol, we can approximate the deviation from the minimum average dissipation in terms of the small protocol variations $\delta\Lambda = \Lambda - \Lambda^*$:

$$\langle\omega\rangle_\Lambda = \omega^* + \frac{1}{2}\sum_{i,j=1} \delta\Lambda_i \cdot \mathbf{K}_{ij} \cdot \delta\Lambda_j, \qquad \text{[S1]}$$

where the indices i and j label times at which the protocol is manipulated. In this limit, we can explicitly calculate the moment generating function for the average dissipation:

$$Z(\gamma) = \left\langle e^{-\gamma\langle\omega\rangle_\Lambda} \right\rangle \underset{\gamma\to\infty}{\sim} \frac{1}{\sqrt{\det(\gamma\mathbf{K})}} e^{-\gamma\omega^*} \propto \gamma^{-n/2} e^{-\gamma\omega^*}, \qquad \text{[S2]}$$

and n denotes the total number of degrees of freedom in the protocol.

The unbiased distribution of average dissipations $P(\omega) \propto \exp[S(\omega)]$ is an inverse Laplace transform of $Z(\gamma)$,

$$P(\omega) \propto (\omega - \omega^*)^{n/2-1}, \qquad \text{[S3]}$$

and hence the entropy in the large γ limit is

$$S(\omega) = \text{const} + \left(\frac{n}{2} - 1\right)\ln(\omega - \omega^*). \qquad \text{[S4]}$$

Fig. S1 compares the asymptotic expression (Eq. **S4**) with the entropy computed for the Ising system.

From the asymptotic expression for $Z(\gamma)$, we can also compute the average dissipation associated with protocols sampled under the bias γ:

$$\langle\omega\rangle = -\frac{d}{d\gamma}\ln Z(\gamma), \qquad \text{[S5]}$$

$$= \omega^* + \frac{n}{2\gamma}. \qquad \text{[S6]}$$

In the absence of protocol bias, $\gamma = 0$, typical values of dissipation are quite large. Provided that $\Lambda(t)$ is bounded, the mean dissipation $\langle\omega\rangle_0$ at $\gamma = 0$ should nonetheless be finite, as is the corresponding variance $\langle(\omega - \langle\omega\rangle_0)^2\rangle_0$. Sufficiently close to the maximum of $P(\omega)$, we therefore have the following:

$$S(\omega) = S(\langle\omega\rangle_0) - \frac{(\omega - \langle\omega\rangle_0)^2}{2\left\langle(\omega - \langle\omega\rangle_0)^2\right\rangle_0}. \qquad \text{[S7]}$$

The parameters $\langle\omega\rangle_0$ and $\langle(\omega - \langle\omega\rangle_0)^2\rangle_0$ in this expression are determined numerically by computing average dissipation for protocols generated in the $\gamma = 0$ ensemble. The corresponding curve is plotted in Fig. 2 as a dotted line.

In the small-bias regime, the moment generating function,

$$Z(\gamma) \propto \exp\left[-\gamma\langle\omega\rangle_0 + \frac{1}{2}\left\langle(\omega - \langle\omega\rangle_0)^2\right\rangle_0\gamma^2\right], \qquad \text{[S8]}$$

gives a γ-biased average dissipation:

$$\langle\omega\rangle = \langle\omega\rangle_0 - \gamma\left\langle(\omega - \langle\omega\rangle_0)^2\right\rangle_0, \qquad \text{[S9]}$$

which is shown as a dotted line in the *Inset* of Fig. 2.

Sampling Efficiency

Ergodic sampling of $P_{\text{canon}}[\Lambda(t)]$ requires that decorrelated protocols be generated, and the efficiency of the sampling depends on how many Monte Carlo moves are required to produce each decorrelated sample. As discussed in the main text, the canonical distribution with bias γ may be accessed with various choices of N and λ, yet the sampling efficiency depends on these choices. A large value of N bears a clear computational cost because N new trajectories must be simulated for each trial protocol. It is not, however, always optimal to choose the minimal N capable of generating bias γ, $N^* = \text{Ceil}(4\gamma)$. The problem with using a small value of N is that rare low-dissipation trajectories (including those with negative dissipation) can depress the probability of accepting a protocol move, even when the trial protocol has a lower average dissipation. The smaller the choice of N, the more rare trajectories influence the acceptance of protocol moves. Consequently, the optimal choice of N often exceeds N^*.

In practice, the best choice of N depends on the Monte Carlo moves used to update trajectories and protocols. For example, there is a trade-off in choosing the best protocol space moves. Large changes to the protocol have a low acceptance probability, but small steps require many moves before sampling a protocol with a decorrelated value of mean dissipation. To quantify the trade-offs between possible choices of N and of Monte Carlo moves, we construct a correlation function:

$$C_\omega(j) = \left\langle \delta\langle\omega\rangle_{\Lambda_i} \delta\langle\omega\rangle_{\Lambda_{i+j}} \right\rangle, \qquad \text{[S10]}$$

where $\delta\langle\omega\rangle_{\Lambda_i}$ is the difference between the ith sampled protocol's average dissipation and the mean dissipation, found by averaging over all protocols in the ensemble. Assuming exponentially decaying correlations ($C_\omega(j)/C_\omega(0) = e^{-j/\tau_\omega}$), we set the correlation time to the decay constant τ_ω. This decorrelation time is the typical number of protocols that must be sampled before generating a protocol with a decorrelated value of the mean dissipation. Because each new protocol must be accompanied by the simulation of N new trajectories, the computational cost for obtaining each decorrelated sample is given by $N\tau_\omega$.

The decorrelation time τ_ω implicitly depends on N as well as on the details of the sampling moves. To gain intuition about the most computationally efficient choice of N, we first consider a simplified Gaussian model. For this model, the optimal N slightly exceeds N^*. The Ising dynamics is more complicated, requiring τ_ω to be computed from simulations. We show that the computational expense depends on N in a similar manner as in the Gaussian model, although the computational expense may be reduced by using noise-guided path sampling.

Gaussian Model. We sample a scalar protocol Λ in the vicinity of the minimum dissipation protocol Λ^*. The advantage of sampling near Λ^* is that we can Taylor expand to arrive at an expression for the average dissipation associated with a perturbation $\delta\Lambda = \Lambda - \Lambda^*$:

$$\langle\omega\rangle_{\Lambda} = \langle\omega\rangle_{\Lambda^*} + \frac{1}{2}K\delta\Lambda^2. \quad \text{[S11]}$$

We assume that the dissipation distribution associated with a protocol Λ is Gaussian:

$$p(\omega|\Lambda) \propto \exp\left(-\frac{(\omega - \langle\omega\rangle_{\Lambda})^2}{4\langle\omega\rangle_{\Lambda}}\right). \quad \text{[S12]}$$

We furthermore use a Gaussian distribution to propose a trial protocol:

$$P(\Lambda_i \to \Lambda_{i+1}) \propto \exp\left(-\frac{(\Lambda_i - \Lambda_{i+1})^2}{2D^2}\right). \quad \text{[S13]}$$

Using the new trial protocol, the dissipation for N independent trajectories is drawn from Eq. **S12** to compute the sample mean dissipation for the new protocol $\overline{\omega}$. For a choice of bias γ, we choose λ and N such that $\gamma = \lambda(1 - \lambda/N)$ and accept the new protocol and trajectories with the following probability:

$$P_{\text{accept}}^{\Lambda\to\Lambda'} = \min\left[1, e^{-\lambda\Delta\overline{\omega}}\right], \quad \text{[S14]}$$

where $\Delta\overline{\omega}$ is the difference between the new protocol's sample mean dissipation and that of the old protocol and trajectories.

The correlation between subsequent samples is given by

$$C_{\omega}(1) = \int d\Lambda_i \int d\Lambda_{i+1} P(\Lambda_i) P(\Lambda_i \to \Lambda_{i+1}) \delta\langle\omega\rangle_{\Lambda_i} \times \left(P_{\text{accept}}^{\Lambda_i\to\Lambda_{i+1}} \delta\langle\omega\rangle_{\Lambda_{i+1}} + \left(1 - P_{\text{accept}}^{\Lambda_i\to\Lambda_{i+1}}\right)\delta\langle\omega\rangle_{\Lambda_i}\right), \quad \text{[S15]}$$

with

$$P_{\text{accept}}^{\Lambda_i\to\Lambda_{i+1}} = \int_{-\infty}^{0} d\Delta\overline{\omega}\, P(\Delta\overline{\omega}|\Lambda_i, \Lambda_{i+1}) + \int_{0}^{\infty} d\Delta\overline{\omega}\, P(\Delta\overline{\omega}|\Lambda_i, \Lambda_{i+1}) e^{-\lambda\Delta\overline{\omega}} \quad \text{[S16]}$$

giving the acceptance probability for the Monte Carlo protocol move. Because the sample mean dissipations (for both the old and trial protocol) are Gaussian distributed, $P(\Delta\overline{\omega}|\Lambda_i, \Lambda_{i+1})$ is also a Gaussian:

$$P(\Delta\overline{\omega}|\Lambda_i, \Lambda_{i+1}) \propto \exp\left(-\frac{N\left(\Delta\overline{\omega} - \left(\langle\omega\rangle_{\Lambda_{i+1}} - \left(1 - \frac{2\lambda}{N}\right)\langle\omega\rangle_{\Lambda_i}\right)\right)^2}{4\left(\langle\omega\rangle_{\Lambda_i} + \langle\omega\rangle_{\Lambda_{i+1}}\right)}\right). \quad \text{[S17]}$$

The acceptance rate may therefore be expressed in terms of error functions. The number of attempted protocol moves required to sample protocols with decorrelated mean dissipations is given by $\tau_{\omega} = -1/\ln(C_{\omega}(1)/C_{\omega}(0))$. A complicated integral expression for τ_{ω} in terms of $\langle\omega\rangle_{\Lambda^*}, D, \gamma$, and N may be derived in terms of error functions. Numerically evaluating the integral expression yields the N dependence of the computational cost, $N\tau_{\omega}$, for generating protocols with decorrelated average dissipation. Analogously, we may compute the time to find decorrelated values of the scalar protocol τ_{Λ} as the decay constant for the correlation function:

$$C_{\Lambda}(j) = \langle\delta\Lambda_i \delta\Lambda_{i+j}\rangle. \quad \text{[S18]}$$

Both τ_{ω} and τ_{Λ} yield the same heuristic: the computational expense is minimal when N slightly exceeds N^*, as plotted in Fig. S2.

Ising Dynamics. Using various choices of λ and N, the canonical protocol ensemble is sampled by a sequence of protocol moves that alter the control point at a single time. Either the control point's magnetic field strength is increased by a random uniform number between −0.5 and 0.5 or the temperature is increased by a random uniform number between −2.5 and 2.5. The protocol move is resampled if the magnitude of the magnetic field strength at any time exceeds 1 or if the dimensionless temperature exits the range [0,8]. N trajectories are generated by repropagating forward dynamics from the initial time or by running time-reversed dynamics from the final time slice and reweighting the trajectories into the forward-trajectory ensemble. When noise guidance path sampling is not being used, each trial trajectory is generated by drawing new random noises to carry out the steps of the spin flip dynamics. The noise guidance path sampling scheme recycles, with probability $1 - \epsilon$, each random number from the previous trajectory. Over time, each random number gets resampled uniformly, but the correlations between an old trajectory and a new trajectory are enhanced (36, 54).

To assess how rapidly the protocol space is sampled, protocols were stored and an accurate estimate of $\langle\omega\rangle_{\Lambda}$ was found by averaging over 1,000 trajectories with the fixed protocol. Fig. S3*A* shows the (Monte Carlo) time series of this average dissipation. Using this time series, the correlation function $C(j)$ is computed for each choice of N and λ. The correlation functions are fit to exponential decays, and the decay constant τ is extracted to yield the computational cost ($N\tau$) as a function of N. Fig. S4 illustrates that the optimal choice of N exceeds N^*, as in the Gaussian model. The noise guidance strategy aids small-N sampling, thereby lowering the optimal N. For these small choices of N, the noise guidance scheme offers a roughly order of magnitude speed up.

Sample Mean Fluctuations

The biasing strength γ necessary to sample protocols with dissipation $\langle\omega\rangle$ is given by the slope of $S(\omega)$ at $\omega = \langle\omega\rangle$. When $S(\omega)$ is especially steep, it therefore requires a very strong bias to access the low-dissipation entropy. Sampling efficiency worsens for large γ, so we note an alternative method for computing $S(\omega)$ that makes use of the fluctuations in the sample mean of N trajectories.

We define the entropy $\bar{S}(\omega)$ to give the density of protocols and trajectories that have a sample mean dissipation within an infinitesimal window of ω:

$$\bar{S}(\omega) = \ln\left[\overline{\Omega}_0 \int d\Lambda(t) dx_1(t) \ldots dx_N(t) \delta\left(\sum_i \frac{\omega[x_i(t), \Lambda]}{N} - \overline{\omega}_{\Lambda}\right)\right]. \quad \text{[S19]}$$

In the limit $N \to \infty$, this entropy must converge to the protocol entropy, $S(\omega)$, but $\bar{S}(\omega)$ falls off less rapidly around the minimal average dissipation. The tails of $\bar{S}(\omega)$ with low sample mean may therefore be importance sampled with a bias λ. Sampling $S(\omega)$ directly, on the other hand, requires the stronger bias γ. By making the Gaussian approximation, we can reconstruct $S(\omega)$ from N dependence of $\bar{S}(\omega)$ using the following:

$$\bar{S}(\omega) = \ln\left[\overline{\Omega}_0 \int d\omega \exp\left(-\frac{N(\overline{\omega} - \omega)^2}{4\omega} + S(\omega)\right)\right], \quad \text{[S20]}$$

a technique illustrated in Fig. S5.

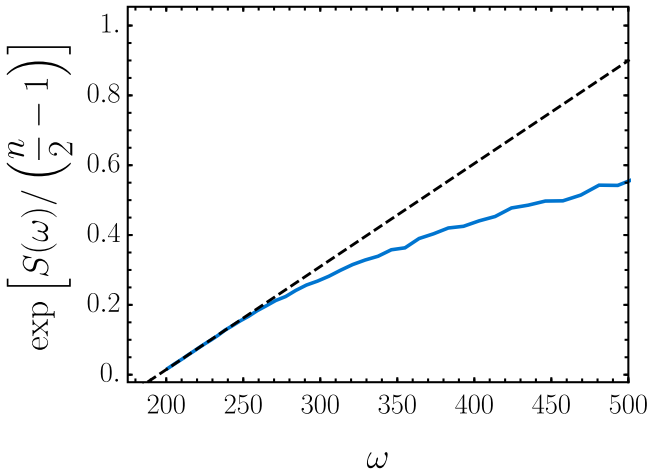

Fig. S1. Asymptotic analysis of the Ising inversion protocol entropy. The protocol entropy $S(\omega)$ from Fig. 2 is transformed and plotted in blue to highlight the large-γ–small-ω limiting behavior. As described in *Materials and Methods*, the protocol entropy calculations used nine control points, each with two degrees of freedom. The dashed black line is the asymptotic form given in Eq. **S4** with $n = 18$.

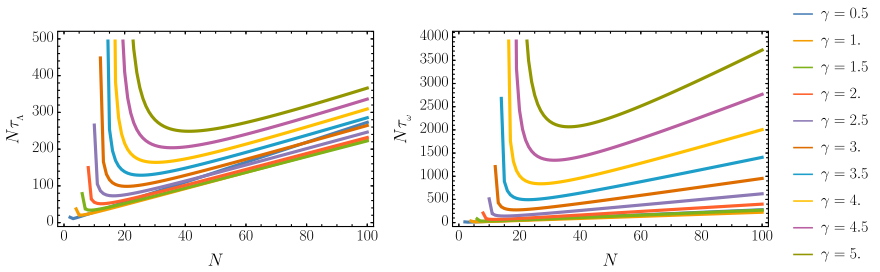

Fig. S2. Computational expense of sampling varies nonmonotonically with N. (*Left*) Computational cost, in number of trajectories needed to generate a decorrelated protocol, for the Gaussian model with various values of the bias γ. (*Right*) Cost to generate protocols with a decorrelated value of the mean dissipation.

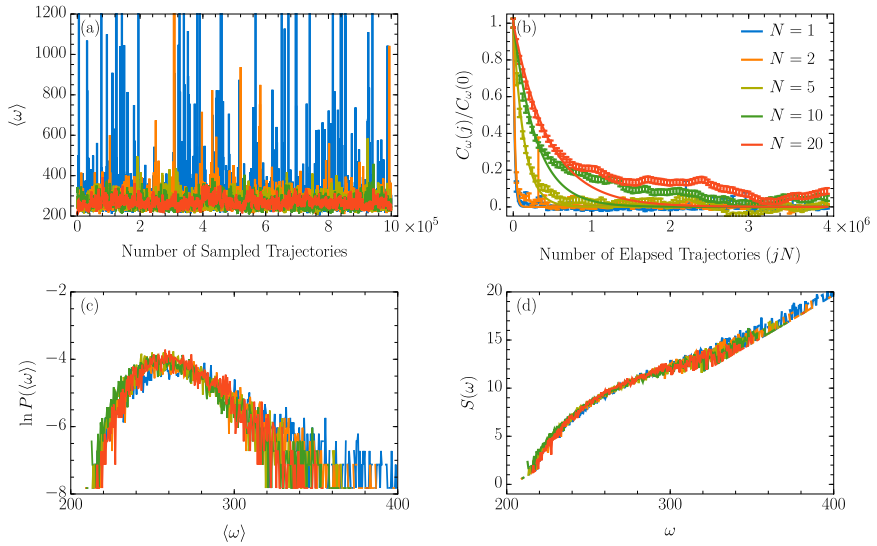

Fig. S3. Sampling $e^{-\gamma\langle\omega\rangle_\Lambda}$ with $\gamma = 0.1$ and various choices of N. To analyze the marginal protocol distributions generated from different choices of N (these marginal distributions agree with each other in the Gaussian dissipation distribution limit), we report on the statistics of $\langle\omega\rangle_\Lambda$, estimated as a sample mean of 1,000 independent trajectories. (*A*) Fluctuations in this mean dissipation for different sampled protocols. The x axis measures "time" for the protocol-sampling Monte Carlo procedure by counting the total number of simulated spin flip trajectories, which differs from the total number of sampled protocols by a factor of N. Counting the number of trajectories roughly reflects the computational expense. (*B*) Correlation functions showing how many trajectories must be simulated before a new protocol is produced with a statistically uncorrelated value of the mean dissipation. Using large values of N can be unnecessarily wasteful because each protocol requires sampling N trajectories. (*C*) Histogram of the sampled average dissipation values. Note that the different choices of N produce the same distributions except that extreme events can be overrepresented by small-N sampling. (*D*) The histogram of C is reweighted to yield the protocol entropy. In the neighborhood of the histogram's peak, the protocol entropy calculation is robust to errors from small-N sampling. By choosing various biasing strengths γ, we highlight different ranges of $\langle\omega\rangle$ and stitch them together to compute the protocol entropy curve reported in the main text.

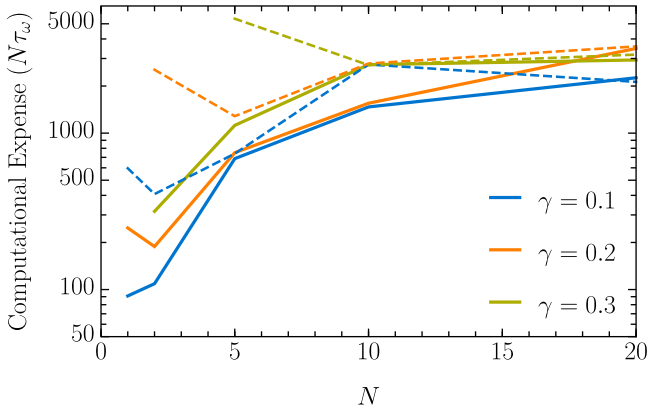

Fig. S4. Computational expense of generating decorrelated protocols. For various N and γ, an exponential time constant τ was extracted from the correlation functions plotted in Fig. S3*B* to obtain the number of trajectories that must be simulated to obtain a decorrelated protocol. This computational expense depends not only on N and γ but also on the manner in which new trajectories are generated. The dashed lines show the expense when each trajectory is drawn at random. Solid lines show the improvement that can be achieved by using noise guidance. When generating a new Ising trajectory, the noise guidance scheme reused each of the previous trajectory's random numbers with probability 0.999 ($\epsilon = 10^{-3}$).

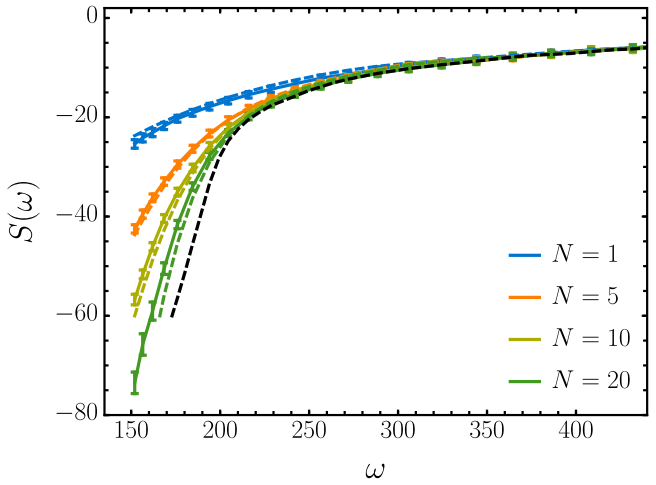

Fig. S5. Inferring protocol entropy from finite-*N* sampling. The colored lines with error bars show $\bar{S}(\omega)$ for four choices of *N*, computed using MBAR to combine sampling with various choices of γ. The form of $S(\omega)$ is inferred (black dashed line) by fitting to Eq. **S20** for $N = 1,\ 5,\ 10,\ 20$. The dashed colored lines show the expected $\bar{S}(\omega)$ given the inferred protocol entropy. It is typically harder to sample in regions where $\bar{S}(\omega)$ is steep. Because the entropy is less steep for small *N*, it can be productive to compute $\bar{S}(\omega)$ with modest *N*, and then infer the $N \to \infty$ limit $S(\omega)$.