

Thermodynamics of Computations with Absolute Irreversibility, Unidirectional Transitions, and Stochastic Computation Times


Gonzalo Manzano¹, Gülce Kardeş,^{2,3} Édgar Roldán,⁴ and David H. Wolpert^{3,4}

¹*Institute for Cross-Disciplinary Physics and Complex Systems (IFISC) UIB-CSIC, Mallorca, Spain*

²*University of Colorado, Boulder, Colorado 80309, USA*

³*Santa Fe Institute, Santa Fe, New Mexico 87501, USA*

⁴*ICTP—The Abdus Salam International Centre for Theoretical Physics, Strada Costiera 11, 34151 Trieste, Italy*

 (Received 25 July 2023; revised 16 February 2024; accepted 14 March 2024; published 13 May 2024)

Developing a thermodynamic theory of computation is a challenging task at the interface of nonequilibrium thermodynamics and computer science. In particular, this task requires dealing with difficulties such as stochastic halting times, unidirectional (possibly deterministic) transitions, and restricted initial conditions, features common in real-world computers. Here, we present a framework which tackles all such difficulties by extending the martingale theory of nonequilibrium thermodynamics to generic nonstationary Markovian processes, including those with broken detailed balance and/or absolute irreversibility. We derive several universal fluctuation relations and second-law-like inequalities that provide both lower and upper bounds on the intrinsic dissipation (mismatch cost) associated with any periodic process—in particular, the periodic processes underlying all current digital computation. Crucially, these bounds apply even if the process has stochastic stopping times, as it does in many computational machines. We illustrate our results with exhaustive numerical simulations of deterministic finite automata processing bit strings, one of the fundamental models of computation from theoretical computer science. We also provide universal equalities and inequalities for the acceptance probability of words of a given length by a deterministic finite automaton in terms of thermodynamic quantities, and outline connections between computer science and stochastic resetting. Our results, while motivated from the computational context, are applicable far more broadly.

DOI: [10.1103/PhysRevX.14.021026](https://doi.org/10.1103/PhysRevX.14.021026)

Subject Areas: Computational Physics,
Statistical Physics

I. INTRODUCTION

A. Background and motivation

In the past three decades there has been major progress in formulating far from equilibrium systems and processes. Using stochastic thermodynamics, we can now rigorously formulate the thermodynamic behavior of systems ranging from biological molecular machines to electronic circuits, evolving arbitrarily away from equilibrium. Celebrated results of stochastic thermodynamics include fluctuation relations that generalize the second law of thermodynamics [1–3], speed limit theorems [4–6], thermodynamic uncertainty relations [7–10], large deviation approaches [11,12], martingale fluctuation relations for extrema and stopping

times [13–19], and universal bounds on various kinetic and frenetic properties [20–22].

The past decade also witnessed progress in thermodynamics of computation. Although many initial studies on energetic costs of computation have mostly concerned unit operations such as bit erasure [23–26], which is too primitive to be pertinent to the formal models of computation in theoretical computer science (TCS), very recent work started to investigate energetic costs of implementing computational machines central to TCS, which perform tasks such as string matching algorithms (which are justifiably more complex than bit erasure) [27–29]. Figure 1(a) shows the general model of a computational machine (henceforth called a computer) which implements a basic algorithm presented in Fig. 1(b).

An algorithm is a finite procedure for implementing a given task, which can be executed in various physical ways, e.g., while modifying the current on electrical wires or the structure of a DNA origami. Formally, an algorithm consists of the instructions to be performed (which is implemented by the dynamics of the computer), the local variables and the memory arrays (stored by the computer),

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

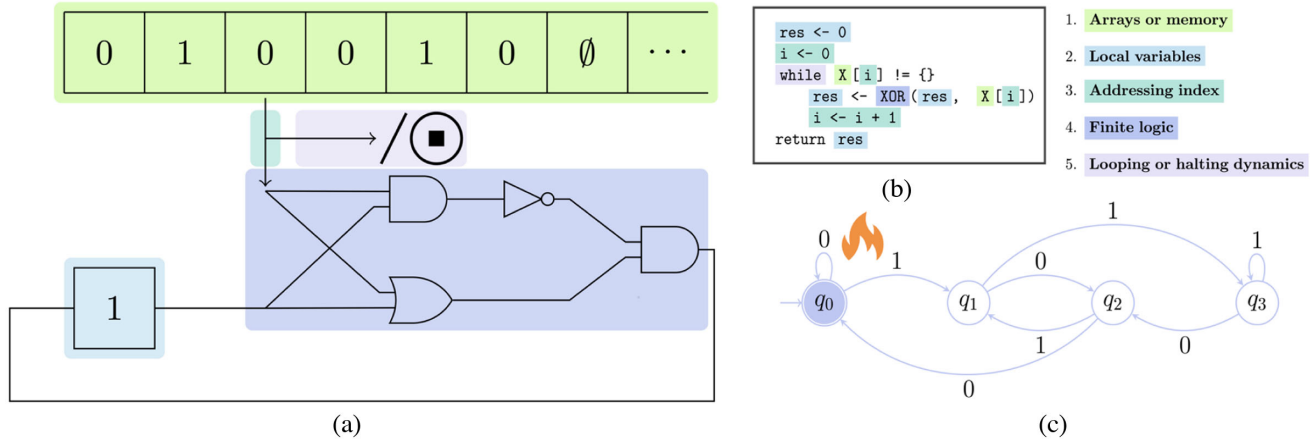


FIG. 1. Illustrations of computations with absolute irreversibility, unidirectional transitions, and stochastic computation times. An algorithm is executed by using a set of instructions, a finite control unit or local variables, a working memory to store the input or intermediate execution values, an address index, and mechanisms which specify when to loop and when to halt (the latter is of interest to us in thinking of stopping times). Both in computer science models and physical computers, the finite control unit generally corresponds to a circuit [as in (a), implementing the algorithm in (b)] or a DFA as in (c), here deciding whether input bit strings are divisible by four. In physical implementation of such devices which solve a myriad of computational tasks, energetic costs are inevitable.

as well as mechanisms to decide when to repeat steps and when to halt. A computer executes an algorithm on a given set of inputs, starting from a certain initial state, potentially following unidirectional transitions in its state space, and halting at an arbitrary stochastic time that depends on the computation. Hence, a general thermodynamic model of computers which implement arbitrary algorithms should be able to account for the energetic costs of implementing computational processes (i) at arbitrary stopping times, with (ii) unidirectional (possibly deterministic) transitions, and (iii) “absolute irreversibility” due to the computer being initialized to a designated start state.

However, most of the central results in stochastic thermodynamics do not directly apply to processes having the aforementioned three key ingredients of computational processes (stopping times, unidirectional transitions, and absolute irreversibility). In fact, a central assumption in much of stochastic thermodynamics is the condition of local detailed balance (LDB), which requires the system to have only bidirectional transitions; i.e., all transitions between any two states $i \rightarrow j$ with their reverse $j \rightarrow i$ have a finite, nonzero probability to occur in a finite time. On the other hand, LDB can be formally avoided by taking an inclusive Hamiltonian approach [30–33], which has been applied recently to computational machines [29]. However, in general there may be hidden nonequilibrium (driven) degrees of freedom which need to be included in the thermodynamic framework beyond a surrounding thermal bath. Moreover, even when assuming LDB, the ratio of transition probabilities between $i \rightarrow j$ and their reverse $j \rightarrow i$ can grow exponentially with the entropy exchanged with the environment, so that the reverse transition will not be effectively seen in relevant timescales, leading to an effective unidirectionality. Until now, little

has been known about the stochastic thermodynamics of systems with broken LDB reflecting unidirectional transitions [21,34–38] or athermal and nonequilibrium environments [39–43]. In addition, the recently established martingale theory of thermodynamics (see Ref. [19] for a review)—which formulates fluctuation theorems and second-law-like inequalities at generic stopping times—has not yet been extended to apply to systems with either unidirectional transitions or absolute irreversibility.

In this paper, we develop a nonequilibrium thermodynamics theory for computations with stochastic computational times that may have absolute irreversibility and unidirectional transitions [44]. Throughout this work, we focus on the intrinsic thermodynamic costs of computations, modeled by generic Markovian dynamics, with minimal or no details about their physical implementation. We derive fluctuation relations for key thermodynamic quantities applicable to all computational processes which can be modeled as discrete-time Markov chains (DTMC), with both unidirectional and bidirectional transitions, and restricted initial probability distributions. These relations hold at both fixed and stopping times, and so simultaneously extend the martingale theory for stochastic thermodynamics to the case of DTMCs with absolute irreversibility and unidirectional transitions.

The thermodynamic meaning of our results is established by introducing a generic physical implementation of the computer as a periodically driven process operated over a set of hidden degrees of freedom. This allows us to link dissipation from underlying (physical) to visible (computational) levels and obtain quantifiers for the energetic costs of computations up to their halting time or between halting times of consecutive computations and their statistics. Our results, while emphasized here for computational

processes, can also describe a wide range of systems, including, e.g., biochemical processes with irreversible release of molecules.

We illustrate our results in deterministic finite automata (DFA). Loosely speaking, a DFA is a system with a finite state space, initialized to a special start state q_0 at $t = 0$, and a logical computer by itself, which can solve basic computational tasks such as string matching. More importantly, it constitutes the “finite logic” part [see Figs. 1(b) and 1(c)] of the engineered computers at use today, and formally corresponds to the finite logic component of Turing machines (TMs).

B. Summary of our contributions and road map

Consider a discrete-time computational task which is implemented by iteratively processing an input sequence of symbols w through a maximum of τ iterations. We refer to τ as the *limit time* of the computation. As an example, τ could be the length $|w|$ of a bit string w , and the computation could involve iteratively processing each of those bits in sequence.

During such a computation, the state of the computer evolves in a stochastic manner, tracing a stochastic trajectory on a set of computational variables $\mathbf{x}_{[0,\tau]} = x_0, \dots, x_\tau$. The computation finishes either at the limit time or, if it is earlier, at a stochastic computation time \mathcal{T} , which in general will be a function of the precise input w . Formally, computational times \mathcal{T} are specific examples of *stopping times*. A stopping time is the first time that a stochastic trajectory meets a specific predefined criterion [45]. In this work, we deal with stopping times which associate to each specific trajectory $\mathbf{x}_{[0,\tau]} = x_0, \dots, x_\tau$ a stochastic time $\mathcal{T} \leq \tau$ that is always smaller than or equal to the limit time. As an example of stopping time in computation, \mathcal{T} could be the first time a DFA reaches a prescribed computational state A . \mathcal{T} could also be defined as the first time a DFA reaches a state $B \neq A$ once having left state A . Stopping times thus provide a flexible yet rigorous mathematical toolbox to tackle computations that last a stochastic amount of time.

In Sec. II, we provide the elementary concepts of our framework, including a formal definition of DFAs: the description of computational processes (such as the running of a DFA) as Markov chains and the physical implementation of those Markov chains. We also review the relevant thermodynamic quantities in this context.

In this paper, we investigate the thermodynamic costs of computations, paying particular attention to those with such stochastic durations. As described below, this will lead us to concentrate on a specific thermodynamic quantity, which we call the *intrinsic mismatch cost*. The intrinsic mismatch cost associated with a stochastic trajectory $\mathbf{x}_{[0,\mathcal{T}]}$ that takes place in the interval $[0, \mathcal{T}]$ of stochastic duration \mathcal{T} is defined as

$$\Sigma(\mathcal{T}) = \sum_{t=0}^{\mathcal{T}-1} \left[\ln \frac{\rho_t(x_t)}{r(x_t)} - \ln \frac{\rho_{t+1}(x_{t+1})}{r'(x_{t+1})} \right]. \quad (1)$$

In Eq. (1), $\rho_t(x)$ is the probability for the computer to be in state x at time t during the computation, whereas $r(x)$ is an arbitrary *reference* probability distribution. On the other hand, $\rho_{t+1}(x)$ and $r'(x)$ correspond to the distributions retrieved applying one iteration of the computer to $\rho_t(x)$ and $r(x)$, respectively. In particular, when the reference distribution minimizes the dissipation in the computer, it is called the *prior distribution*, and the associated quantity Eq. (1) is known as the *mismatch cost* of the computation up to time \mathcal{T} . This cost provides a lower bound on the entropy production incurred by any digital synchronous computer that implements this dynamics over values of the computational variables x , without any precise assumptions about the continuous-time process implementing each successive iteration of the computation [46]. Note that while the distributions ρ_t and ρ_{t+1} are indexed by the iteration number, changing as the computation proceeds, the distributions r and r' are the same for every iteration of the computation.

Some of our most important results are fluctuation relations and inequalities for the statistics of $\Sigma(\mathcal{T})$. These are threefold universal, in that they are valid for all (i) reference distributions $r(x)$ [although we will make specific choices for it to give a concrete thermodynamic meaning to Σ], (ii) stopping rules \mathcal{T} that halt before $\tau + 1$, (iii) discrete-time Markov chains (DTMCs), even those with restricted initial conditions, and/or with unidirectional transitions, and even DTMCs that are implemented by underlying non-Markovian continuous-time processes.

To illustrate these results we need to first introduce more definitions. Here and throughout this paper, we use $\langle A(\mathcal{T}) \rangle \equiv \mathbb{E}[A(\mathcal{T})]$ for the expectation of any functional A of $x_{[0,\mathcal{T}]}$ over many realizations of the computation, each ending at the possibly different time \mathcal{T} . In addition, we introduce the *stochastic distinguishability* at stopping times,

$$\delta_\tau(\mathcal{T}) = \ln \frac{\rho_t(x_t)}{\bar{\rho}_{\tau-t}(x_t)} \Big|_{t=\mathcal{T}}, \quad (2)$$

which quantifies the asymmetry of the computational process under time reversal through its distribution and plays a key role for identifying possible reductions in thermodynamic costs along processes with stochastic duration [16]. The distribution $\bar{\rho}_t(x)$ above is the probability for an *auxiliary computation* running backward in time to be in state x at time t during the computation. Such auxiliary computation evolves with a transition matrix that is the Bayesian inverse of the original computation with respect to r , such that it leads to the perfect retrodiction in time of the distribution r [51]. In Sec. III, we provide a formal definition for the auxiliary computational process, that allows us to address the thermodynamics of computations ending at a fixed time at the fluctuating level, and discuss how to incorporate explicitly the role of absolute irreversibility and unidirectional transitions, which is crucial because

conventionally formulated computational machines have precisely those features.

In Sec. IV, we present our main results for general computations starting and ending at stochastic halting times, which include fluctuation theorems and second-law inequalities that provide lower bounds to the average dissipation in a computation. A central result of our work is the derivation of an integral fluctuation relation at stopping times which applies to arbitrary computations:

$$\langle e^{-\Sigma(\mathcal{T}) - \delta_\tau(\mathcal{T})} \rangle = 1 - \Gamma_\tau. \quad (3)$$

In this equation, Γ_τ quantifies absolute irreversibility at stopping times. It equals the functional $e^{-\delta_\tau(\mathcal{T})}$ averaged over the restricted set of trajectories that can take place in the auxiliary dynamics but have zero probability in the original dynamics; see Eq. (36). This contribution introduces an unavoidable source of irreversibility that limits possible reductions in the dissipation of the computer.

Among other things, Eq. (3) provides a second-law inequality at stopping times for the intrinsic mismatch cost:

$$\langle \Sigma(\mathcal{T}) \rangle \geq -\langle \delta_\tau(\mathcal{T}) \rangle - \ln[1 - \Gamma_\tau]. \quad (4)$$

As we show in this paper, the right-hand side of Eq. (4), which sums up the net effects of time-asymmetry and absolute irreversibility, gives a universal lower bound not only for the intrinsic mismatch cost of the computation but also for the underlying average entropy production incurred by the computer. Moreover, for the particular case of a stationary reference distribution ($r = \pi$, with $\pi' = \pi$) $\langle \Sigma(\mathcal{T}) \rangle$ equals the average (discrete-time) nonadiabatic entropy production [52,53].

We remark that Eq. (3) follows from a stronger result that we derive here. In particular, $e^{-\Sigma(t) - \delta_\tau(t)}$ is a *supermartingale* process, i.e., $e^{-\Sigma(t) - \delta_\tau(t)}$ decreases with time when conditioned on an earlier part of the trajectory of states: $\langle e^{-\Sigma(t) - \delta_\tau(t)} | x_{[0,s]} \rangle \leq e^{-\Sigma(s) - \delta_\tau(s)}$, where $t \geq s \geq 0$; see Eq. (31) [54].

Another contribution in this paper arises when we apply the general martingale theory for thermodynamics [19], to extend our results to multiple, ordered stopping times. In particular, for the case of two stopping times \mathcal{T}_1 and \mathcal{T}_2 with $\mathcal{T}_2 \geq \mathcal{T}_1$, we obtain another central result:

$$\langle \Sigma(\mathcal{T}_2) + \delta_\tau(\mathcal{T}_2) \rangle \geq \langle \Sigma(\mathcal{T}_1) + \delta_\tau(\mathcal{T}_1) \rangle, \quad (5)$$

which provides a powerful second-law inequality applicable to both starting and ending stochastic times of computations. As an example, in the case of a DFA, \mathcal{T}_1 could be the first time that some particular state A is reached, and \mathcal{T}_2 could be the first time state B is reached after state A has been reached. Alternatively, \mathcal{T}_2 could be the *second* time state A is reached after the system has first reached state A , then left it, and then returned. See Sec. V

for numerical illustrations of these ideas in a specific minimal model of a DFA processing binary strings.

Moreover, from Eq. (5) a sandwich inequality for $\langle \Sigma(\mathcal{T}) \rangle$ can be derived [55],

$$D(\rho_0 || \bar{\rho}_\tau) - \langle \delta_\tau(\mathcal{T}) \rangle \leq \langle \Sigma(\mathcal{T}) \rangle \leq \langle \Sigma(\tau) \rangle - \langle \delta_\tau(\mathcal{T}) \rangle, \quad (6)$$

extending recent research in upper bounds and inverse thermodynamic uncertainty relations in stochastic thermodynamics given in Refs. [56–58].

In addition to the aforementioned fundamental results, in Sec. VI we also combine the supermartingale property of $e^{-\Sigma(t) - \delta_\tau(t)}$ with the fluctuation relation Eq. (3) to derive universal equalities and inequalities for the probability that a computation is completed within a certain amount of time. Such an idea can be applied, e.g., to compute the probability that a sequence of τ ordered computational states visited by a DFA during its evolution reaches an accept state.

Section VII is then devoted to sketch how our theory can be applied to investigate the thermodynamics of multiple concatenations of runs of a DFA, where after each run ends the system is reset to an initial start state and the next run begins. We conclude with Sec. VIII, where we present our main conclusions and further discuss future research directions motivated by our findings. Mathematical details of the derivations, proofs, and extra discussions are left to corresponding appendixes.

It is worth remarking that all our contributions, while originally motivated from problems arising in the context of the kinds of computational machines central to computer science theory, are applicable to any periodic process implementing a time-homogeneous DTMC, one that results in trajectories $\mathbf{x}_{[0,\mathcal{T}]}$.

II. MARKOVIAN COMPUTATIONS

In the following, we make the assumption that the implementation of a task on a given computer is realized through a physical process which induces Markovian (discrete) dynamics over a set of relevant computational states. The actual physical process being modeled will be a generic physical, chemical, or biological system, whose dynamics can be described at a microscopic level over a set of hidden degrees of freedom [47,48], here assumed to be not directly accessible. In particular, it is customary to model a computation as a continuous-time Markov chain (CTMC) [48,59–62].

In “synchronous” physical computers—such as all real-world digital computers—this CTMC is driven externally following a periodic protocol induced, e.g., by the electric current powering a computer. Such underlying periodic driving might be ignored in modeling the computational process, by describing its evolution by coarse graining it in time, and this results in an effective model given by a time-homogeneous DTMC. Throughout this

paper, we will work at such a coarse-grained level, and consider computational processes as generic DTMCs with time-independent transition probabilities. In doing so, we will map the underlying physical process to the DTMC dynamics of the (symbolic) computational states to formulate the actual physical dissipation in a thermodynamically consistent manner.

A. Stochastic computational processes

We consider computational processes described by a DTMC that can take values over a discrete set of $N \geq 1$ computational states $x_t \in \mathcal{X}$, with $t = 0, 1, 2, \dots$. For simplicity, we assume that the transition probabilities between the computational states are time independent (however, our results can be extended to time-dependent transition probabilities).

We write $P(x_{t+1}|x_t)$ for the conditional probability of jumping to state x_{t+1} given that the previous state was x_t in a single time step or iteration of the computational process. (Note that in a DTMC x_t can be the same as x_{t+1} , allowing for time instances where the system dwells in a given state.) We write $\rho_t(x)$ for the probability of being in state x at time t , given an ensemble of realizations of the Markovian process. The associated discrete-time master equation $\rho_{t+1} = \mathbf{W}\rho_t$, where ρ_t is an $N \times 1$ column vector and $[\mathbf{W}]_{i,j} = P(x_{t+1} = i|x_t = j)$ is the transition probability matrix. The transition matrix \mathbf{W} has at least one fixed point with distribution $\pi(x)$ such that $\mathbf{W}\pi = \pi$, and if aperiodic and irreducible, π becomes the unique stationary distribution in the long time run, that is, $\lim_{t \rightarrow \infty} \rho_t = \pi$. However, what follows does not require π to be unique.

Throughout the paper we will write τ for the *limit time* of a computation, i.e., the maximum time that can be spent to execute a computation, and assume it to be fixed. The probability of a sequence $\mathbf{x}_{[0,\tau]} = x_0, x_1, \dots, x_\tau$ is

$$P(\mathbf{x}_{[0,\tau]}) = \rho_0(x_0) \prod_{t=0}^{\tau-1} P(x_{t+1}|x_t). \quad (7)$$

Here we allow for arbitrary initial distributions $\rho_0(x_0)$ and transition probabilities $P(x_{t+1} = j|x_t = i) = P(j|i)$. In particular, some of the transitions might be bidirectional ($i \leftrightarrow j$) and others unidirectional ($i \rightarrow j$). Bidirectional transitions are characterized by conditional probabilities verifying $P(i|j) > 0$ whenever $P(j|i) > 0$, while for unidirectional transitions we can have $P(i|j) = 0$ with $P(j|i) > 0$. We notice that exactly because of the existence of unidirectional transitions, it is mandatory to relax the condition of local detailed balance, which is arguably among the most common assumptions adopted in the formulation of stochastic thermodynamics [63].

One of the main quantities of interest in stochastic thermodynamics is stochastic entropy production (EP) which equals the logarithm of the ratio between forward

and time-reversed path probabilities of a thermodynamic process [3,64]. This quantity, however, generically depends on the details of the underlying physical process implementing the computation, hence is not directly accessible unless certain simplifying assumptions, such as the condition of local detailed balance. Nevertheless, here we aim to obtain a thermodynamic description of the computational processes as deduced solely from the (discrete-time) dynamics of the visible variables defining the computation $x_t \in \mathcal{X}$. While our analysis holds for arbitrary DTMCs, we focus on digital synchronous computers, which undergo a time-homogeneous dynamics over discrete time, and which we connect to the underlying physical process generating it in a simple manner. This allows us to express and bound the entropy production of the computational task implemented by the DTMC, alongside the work and heat dissipated into the environment. For simplicity, we take the *continuous-time* physical process that implements the time-stationary DTMC to be periodic and choose units so that the period of the physical process is 1.

As an example (and to help ground the reader's intuition), suppose that our time-homogeneous DTMC is implemented by a time-inhomogeneous CTMC. It is well known that, in general, this requires that the CTMC evolves over an enlarged version of the DTMC's state space $\mathcal{Y} \supseteq \mathcal{X}$, which includes "hidden states" in addition to the "visible" states of the DTMC [48,65]. In particular, this is true when the DTMC is the update function of a computational machine. Therefore, our assumption that the continuous-time physical process is periodic implies that the time-inhomogeneous CTMC is periodic. As a result, the thermodynamics arising in any single iteration of the physical system [implementing the computational machine that starts at discrete time t in some state $y(t) \in \mathcal{Y}$] is independent of t . In the following, for further simplicity (and to ensure a time-homogeneous DTMC) we also assume that noncomputational degrees of freedom in \mathcal{Y} are reinitialized within every single iteration to their (possibly nonequilibrium) initial states.

B. Mismatch cost

Enlarging our original description over \mathcal{Y} to include all relevant physical variables of the computer is crucial to define the associated entropy production and other relevant thermodynamic costs of computation. Here we show that this can be done in a standard way. In particular, suppose that we are interested in some generic thermodynamic *average cost function* that can be written as

$$\begin{aligned} \mathcal{C}(\tau) &= S(q_\tau) - S(q_0) + F \\ &:= S(Gq_0) - S(q_0) + F, \end{aligned} \quad (8)$$

where $S(q) = -\sum_i q(i) \ln q(i)$ is Shannon entropy, q_0 is any initial distribution over the (extended set of) states of

the system, G is the linear map that transforms that distribution to an associated ending distribution q_τ , and F is an arbitrary linear functional of the initial state. Note that $\mathcal{C}(\tau)$ is an implicit function of q_0 .

As a canonical example, in CTMC-based stochastic thermodynamics obeying local detailed balance, the EP generated during a process is given by Eq. (8) by setting F equal to the *average entropy flow* to the environment:

$$F = \int_0^\tau dt \sum_v \sum_{i,j} q_\tau(j) K_{ij}^v(t) \ln \left[\frac{K_{ji}^v(t)}{K_{ij}^v(t)} \right], \quad (9)$$

where $K_{ij}^v(t)$ is the rate matrix associated to thermal reservoir v , and the rate matrix of the CTMC is $\sum_v K_{ij}^v(t)$ [66]. For different choices of F , $\mathcal{C}(\tau)$ gives different thermodynamic quantities besides EP, such as the drop in nonequilibrium free energy of the system during the process [49], among many others.

For any cost \mathcal{C} of the form in Eq. (8), and *any* physical process represented by G , the *prior* distribution is defined to be the initial distribution q_0 that minimizes $\mathcal{C}(\tau)$. (It is called the prior because it is, formally speaking, a prior distribution for calculating the posterior probability of an initial state of a thermodynamic process given its final state [67].) We write the prior as q_{\min} . The associated *average mismatch cost* is

$$\mathcal{M}(\tau) := D(q_0 \| q_{\min}) - D(Gq_0 \| Gq_{\min}), \quad (10)$$

where $D(q_1 \| q_2) = \sum_y q_1(y) \ln[q_1(y)/q_2(y)]$ denotes the Kullback-Leibler (KL) divergence between the distributions q_1 and q_2 for the case of a discrete random variable. Note that $\mathcal{M}(\tau)$ is implicitly a function of q_0 , q_{\min} and the linear function G —but nothing else. It depends on *no* other property of the process besides q_{\min} and G .

As shown in Refs. [27,47–49] we have, for all q_0 ,

$$\mathcal{C}(\tau) = \mathcal{M}(\tau) + \mathcal{R}(\tau), \quad (11)$$

where $\mathcal{R}(\tau)$ is an additive non-negative contribution independent of q_0 called “residual cost.” In the specific case in which F is identified with the entropy flow, residual cost is often called “residual EP.” See Appendix A for a discussion of residual cost and why we ignore it in this paper.

Expressions analogous to Eqs. (8), (10), and (11) hold for other state spaces, e.g., real-valued states, density matrices, etc. Moreover, there are no assumptions of detailed balance or the like in the derivation of Eq. (10); it holds purely for mathematical reasons. For the trajectory-level version of mismatch cost in Eq. (10) see Appendix B.

By the data-processing inequality for KL divergence [68], $\mathcal{M}(\tau)$ is never negative. Moreover, it can be shown that the prior q_{\min} in Eq. (10) has full support (see Appendix A in Ref. [27]), which ensures that the mismatch cost is finite.

Note also that the mismatch cost formula (10) is based on evaluating q at both the beginning and the end of the time interval $[0, \tau]$. This means this general formula applies to *any* physical process that maps q_0 to q_τ , for *any* choice of \mathcal{C} , i.e., any choice of the linear functional F . All the (messy) physical details of the process and the precise choice of F are buried in the prior and the residual cost.

In the following, unless explicitly stated otherwise, we will focus on the case in which the cost function $\mathcal{C}(\tau)$ in Eq. (11) is the average entropy production of the continuous-time periodic process implementing the computation. If such a process is moreover Markovian, the entropy flow F will be given by the canonical CTMC expression in Eq. (9). However, Markovianity of the continuous-time periodic process is not a necessary assumption for our results (see also Appendix C).

C. Strictly positive lower bounds for dissipation in periodic processes

As described above, in real-world (synchronous, digital) physical computers, the underlying physical process implementing each iteration of the computer is identical. This is true whether that physical process is a CTMC, a quantum operation, and so on. As noted in Ref. [28], this means that the prior q_{\min} for each iteration of the computer is the same [69]. Using also the fact that noncomputational variables are reinitialized in every single iteration (period) of the computational process, we can write the overall mismatch cost for any computation that takes exactly τ iterations in terms of computational variables as

$$\sum_{t=0}^{\tau-1} \mathcal{M}(\rho_t) = \sum_{t=0}^{\tau-1} [D(\rho_t \| \mu) - D(\rho_{t+1} \| \mu')], \quad (12)$$

where $\rho_t(x) = \sum_{y \notin \mathcal{X}} q_t(y)$ is the (marginal) distribution over computational states $x \in \mathcal{X}$ only, $\mu(x) = \sum_{y \notin \mathcal{X}} q_{\min}(y)$ is the prior over computational states at the beginning of (every) iteration, and $\mu' = \mathbf{W}\mu$ the prior at the end of every iteration (i.e., it is μ evolved to the end of the iteration). More details about the derivation of Eq. (12) for the DTMC are provided in Appendix C, where we explicitly elaborate its relation to the underlying EP in the continuous-time periodic process.

We note that if $\rho_0 = \mu$ in Eq. (12), then the first difference of KL divergences being summed equals 0. However, unless \mathbf{W} is degenerate (e.g., the identity matrix), $\mathbf{W}\rho_0 \neq \rho_0$, and therefore $\mathbf{W}\rho_0 \neq \mu$. This in turn means that the *second* difference of KL divergences being summed in Eq. (12) does *not* equal 0 (so long as \mathbf{W} is not logically invertible, i.e., not a permutation matrix). Therefore in this case, the overall sum will be strictly positive. This argument can be extended to prove that so long as \mathbf{W} is not logically invertible (and ρ_0 is not a fixed

point of the dynamics), the mismatch cost sum in Eq. (12) is not zero (see Appendix D).

Since the above reasoning is true for all actual μ , we can lower bound the sum Eq. (12) by minimizing over all distributions λ in the unit simplex Δ_X , whether or not they are a valid prior in some physical scenario:

$$\sum_{t=0}^{\tau-1} \mathcal{M}(\rho_t) \geq \inf_{\lambda \in \Delta_X} \sum_{t=0}^{\tau-1} [D(\rho_t \|\lambda) - D(\rho_{t+1} \|\lambda')] > 0, \quad (13)$$

with again $\lambda' = \mathbf{W}\lambda$ (see also Ref. [28]). The precise prior μ in Eq. (12) for the EP cost function will depend on the details of the precise physical process under consideration. On the other hand, the sum Eq. (13) is independent of those details. We therefore obtain a strictly positive lower bound on EP, given *in toto* by ρ_0 and \mathbf{W} . This strengthened second law arises solely from the fact that we have a periodic process with a nonlogically invertible \mathbf{W} . Moreover, because minimization in Eq. (13) is over all possible priors, it provides a lower bound on all costs that can be written as in Eq. (8). We therefore refer to it as the *minimal dissipation*.

As an example, suppose that our DTMC is the dynamics of a noise-free, synchronous, digital computer, with update function $f: X \rightarrow X$. Plugging in Eq. (13), the minimal possible EP is

$$\min_{\lambda \in \Delta_X} \left(\sum_{t=0}^{\tau-1} \sum_{x \in \Omega_t} \rho_0[f^{-t}(x)] \ln \left[\frac{\rho_0[f^{-t}(x)]}{\lambda(x)} \right] - \rho_0[f^{-t-1}(x)] \ln \left[\frac{\rho_0[f^{-t-1}(x)]}{\lambda[f^{-1}(x)]} \right] \right). \quad (14)$$

The term Ω_t in Eq. (14) is the set of all states that have nonzero probability if the update function is applied to the actual distribution ρ_0 a total of t times. The term $\rho_0[f^{-t}(x)]$ in Eq. (14) is the probability, under ρ , of the entire set of those states in X which, after t iterations of (the periodic processes underlying) the update function f of the digital computer, are in state x (and similarly for $\rho_0[f^{-t-1}(x)]$). Suppose that f is not just a permutation of the states of the computational machine that lie in the support of ρ_0 . Then Eq. (14) provides a strictly positive lower bound on the dissipation incurred by any physical device that implements that computation, f .

In the sense that it depends only on the conditional distribution \mathbf{W} and the initial distribution ρ_0 , the bound for periodic processes in Eq. (13) is similar to the generalized Landauer's bound. In particular, the thermodynamic uncertainty relations and speed limit theorems are also lower bounds on EP that depend on the initial distribution over states and the discrete-time conditional distribution of the dynamics. However, unlike the lower bound above, those other bounds depend on other properties of the process besides the initial distribution and the conditional distribution giving the dynamics (for example, current

precisions or expected activities). In this sense, the minimal dissipation given in Eq. (13) is more powerful than those other lower bounds on EP (a closed form of this result in terms of Jensen-Shannon divergence has also been reported very recently in Ref. [70]).

In this paper, we calculate mismatch costs by summing the cost over single iterations of a computational machine operating periodically, as in Eq. (12). In general this does *not* equal the standard mismatch cost for the entire computation, with an overall prior and a single drop in KL divergence between initial and final time τ . We remark that, to our knowledge, the necessary and sufficient conditions for this quantity to be larger than the one we use in this paper are not known. However, there is a particularly interesting case in which these two expressions become the same, namely, when EP is minimized at the stationary state of the DTMC; i.e., the prior μ coincides with π . In such case we recover from Eq. (11) the well-known decomposition of EP into adiabatic and nonadiabatic contributions [52,53,71,72], where mismatch cost reduces to nonadiabatic EP (also called excess EP [19]) and the residual cost becomes adiabatic EP (housekeeping heat [53,71]).

D. Deterministic finite automata


An important class of computational machines that can be described within our framework are the DFA. There are several different, very similar definitions of DFA, some of which overlap with common definitions of “finite state machines.” To fix the discussion, here we adopt the following definition. A *deterministic finite automaton* is a 5-tuple (Q, θ, q_0, A, f) , where

- (1) Q is a finite set of (*logical*) *states*,
- (2) θ is a finite (input) *alphabet*,
- (3) $q_0 \in Q$ is the *start state*,
- (4) $A \subseteq Q$ is the set of *accept states*, and
- (5) $f: Q \times \theta \rightarrow Q$ is the *update function*, mapping a current input symbol and the current logical state to a next logical state.

A finite string of successive input symbols, i.e., an input string $\omega \in \theta$, is sometimes called an (*input*) *word*. To operate a finite automaton on a particular input word, one begins with the automaton in its start state, and feeds that state together with the first symbol in the input word into the update function, to produce a new logical state. Then one feeds in the next symbol in the input word (if any), to produce a next logical state. Note that one can represent any given DFA's update function as a directed graph, where each edge (q_1, q_2) taking logical state q_1 to state q_2 is labeled by the input symbols that would cause that transition [see Figs. 1(c) and 2 for illustrations].

Our analysis of stochastic computational processes (as introduced above) in DFAs requires assigning probabilities to the input words (or to the symbols inside them) that are fed into the automaton, as well as identifying the computational states of the DTMC \mathcal{X} , which may coincide

(a) Original dynamics



(b) Auxiliary dynamics




FIG. 2. (a) Discrete-time Markov chain (DTMC) associated with the DFA recognizing binary i.i.d. sequences that are multiple of four [see Fig. 1(c)]. The transition matrix of such DTMC is given by Eq. (54) where p_0 and $p_1 = 1 - p_0$ denote, respectively, the probability for a 0 and a 1 in the input string. (b) DTMC associated with the auxiliary dynamics associated with the stationary prior, with transition probability matrix obtained from Eq. (15) and given by Eq. (57).

or not with the set Q of logical states of the DFA (typically, \mathcal{X} may contain more variables as, e.g., previously processed symbols). An important contribution of our work will be to show how one can do this analysis even though the dynamics of a DFA—its update function—is deterministic and often noninvertible (i.e., unidirectional), and given that the initial distribution over states of the DFA (though not over the input words) is a delta function, centered on the start state (i.e., leading to absolute irreversibility).

Physically, the (probabilistically generated) input word ω may be encoded in a tape whose symbols are read by the DFA “head” one by one in each cycle of the computation, but are not modified by the automaton operation [73]. In this way the input tape behaves as an (energyless) information reservoir [74], whose Shannon entropy is kept constant during the computation. More formally, we can consider that tape as forming part of the physical states of the computation in the extended state space \mathcal{Y} (a Cartesian-product factor of Q , ω , and other physical variables depending on the implementation), but we will not generically include it within the computational states in \mathcal{X} [75]. On the other hand, we will eventually incorporate some already processed symbols explicitly into \mathcal{X} , which are then assumed to be stored (and modified) in extra physical variables acting as a memory for the computer.

A typical question of interest in computer science is whether the DFA is in an accept state of the set A after the

last symbol from the input word is processed. If that is the case, one says that the automaton *accepts* that input word. In this way any given automaton uniquely specifies a *language* of all input words that automaton accepts, which is called a *regular* language. Importantly, any particular DFA can process input words of arbitrary length [76], and in general may enter and exit its set of accepting states multiple times, before the end of the input word. While the definition of whether an input word is accepted depends only on whether the ending logical state is an accepting state, the statistics of whether, how often, and precisely when a given DFA enters an accept state (when fed words generated by some given distribution) can be of independent interest.

III. INTRINSIC THERMODYNAMICS OF COMPUTATIONS AT FIXED TIMES

The mismatch cost sum introduced in Eq. (12) depends only on the computational degrees of freedom involved in the original DTMC dynamics and provides a lower bound on the average entropy production generated by the machine implementing the computation. It is hence a particularly useful candidate to assess the intrinsic (minimal) thermodynamic costs of computations. The prior $\mu(x)$ encodes the specific details of the physical implementation of the computational process. Concern for such details can even be avoided by considering the distribution $\nu(x)$ given by the infimum of Eq. (13), which still provides a useful (positive) bound on EP.

To begin, we construct a stochastic description based on thermodynamic quantities that can be computed by introducing an *auxiliary process*. This process is defined in terms of the “forward” discrete-time dynamics $P(j|i)$, the initial distribution of that dynamics $\rho_0(x)$, and a *reference distribution* $r(x)$ over computational states. The reference $r(x)$ is arbitrary, and in particular could be chosen to obtain stochastic versions of the mismatch cost sum in Eq. (12) [$r(x) = \mu(x)$] and the minimum dissipation in Eq. (13) [$r(x) = \nu(x)$].

A. Thermodynamic costs of periodic computations at the fluctuating level

We start by introducing the discrete-time auxiliary dynamics of the auxiliary process $\bar{W}_{i,j}$, with transition probabilities defined from the transition probabilities in W by

$$\bar{P}(i|j) \equiv \frac{P(j|i)r(i)}{r'(j)}, \quad (15)$$

where $r' = Wr$ is the reference distribution r evolved for one iteration; i.e., $r'(j) = \sum_i P(j|i)r(i)$ [77]. This auxiliary process is a bona fide Markov chain with $\sum_i \bar{P}(i|j) = \sum_i [P(j|i)r(i)]/r'(j) = 1$ and $0 \leq \bar{P}(i|j) \leq 1$ [78].

Moreover, r' transforms back into r in a single iteration under \bar{W} . That is, \bar{W} corresponds to the Bayesian inverse of W with respect to the reference distribution r , leading to perfect retrodiction for the distribution r [62,79–81].

To fully specify the auxiliary dynamics we must specify its initial distribution; here we will always set it to distribution of the original actual dynamics at its limit time, i.e., $\bar{\rho}_0(x) = \rho_\tau(x)$. So the joint distribution of a trajectory $\mathbf{x}_{[0,\tau]}$ under the auxiliary dynamics is

$$\bar{P}(\mathbf{x}_{[0,\tau]}) = \bar{\rho}_0(x) \prod_{t=0}^{\tau-1} \bar{P}(x_{t+1}|x_t). \quad (16)$$

Note that this choice of the initial distribution of the auxiliary process is not restricted by the choice of r in any way. Note as well that $\bar{P}(i|j)$ does not necessarily coincide with the transition probabilities induced by the time-reversed implementation of the underlying physical process, but it is solely defined from the distribution $r(x)$ and the original Markov chain transition probabilities.

Using Eqs. (15) and (16), we can write the probability of a time-reversed discrete-time trajectory, $\Theta\mathbf{x}_{[0,\tau]} = x_\tau, x_{\tau-1}, \dots, x_0$, under the auxiliary dynamics as

$$\begin{aligned} \bar{P}(\Theta\mathbf{x}_{[0,\tau]}) &= \bar{\rho}_0(x_\tau) \bar{P}(x_{\tau-1}|x_\tau) \dots \bar{P}(x_0|x_1) \\ &= \rho_\tau(x_\tau) \prod_{t=0}^{\tau-1} P(x_{t+1}|x_t) \frac{r(x_t)}{r'(x_{t+1})}. \end{aligned} \quad (17)$$

The ratio between the path probability to observe a given trajectory of states and the path probability to observe its time reversal under the auxiliary dynamics is

$$\begin{aligned} \Sigma(\mathbf{x}_{[0,\tau]}) &\equiv \ln[P(\mathbf{x}_{[0,\tau]})/\bar{P}(\Theta\mathbf{x}_{[0,\tau]})] \\ &= \sum_{t=0}^{\tau-1} \left[\ln \frac{\rho_t(x_t)}{r(x_t)} - \ln \frac{\rho_{t+1}(x_{t+1})}{r'(x_{t+1})} \right], \end{aligned} \quad (18)$$

providing us, for $r = \mu$, a stochastic version of the mismatch cost sum in Eq. (12), and for $r = \nu$, the minimal dissipation in Eq. (13). The functional $\Sigma(\mathbf{x}_{[0,\tau]})$ is an example of a “ Σ -entropic functional,” as introduced in Ref. [19].

The specific choice for the transition probability of the auxiliary dynamics introduced in Eq. (15) is crucial for avoiding divergences that would be induced by unidirectional links if we evaluate expressions like $\ln[P(i|j)/P(j|i)]$ —expressions that appear in most functionals associated with entropy production. This makes the functional Σ given by Eq. (18) suitable to tackle fluctuations of Markovian processes with unidirectional transitions, which are precisely the (idealized) dynamics of many computational processes.

Here and in the following, as shorthand, we will often write trajectory-level quantities such as $\Sigma(\mathbf{x}_{[0,\tau]})$ simply as $\Sigma(\tau)$, with the precise trajectory left implicit. Following such shorthand notation, Eq. (18) can be decomposed as

$$\Sigma(\tau) = \Delta S_{\text{sys}}(\tau) - \Delta\phi(\tau), \quad (19)$$

where we write the change in stochastic Shannon entropy of the computer as

$$\begin{aligned} \Delta S_{\text{sys}}(\mathbf{x}_{[0,\tau]}) &:= - \sum_{t=0}^{\tau-1} [\ln \rho_t(x_t) - \ln \rho_{t+1}(x_{t+1})] \\ &= - \ln \rho_\tau(x_\tau) + \ln \rho_0(x_0), \end{aligned} \quad (20)$$

and write the change in the *nonequilibrium potential* as

$$\Delta\phi(\mathbf{x}_{[0,\tau]}) := \sum_{t=0}^{\tau-1} [-\ln r'(x_{t+1}) + \ln r(x_t)]. \quad (21)$$

Such nonequilibrium potentials have been fruitfully employed in steady-state thermodynamics [41,53,82], and account for the excess of entropy absorbed from the environment during the computation $\mathbf{x}_{[0,\tau]}$ whenever the state of the system ρ_t differs from the distribution r along its time evolution.

Suppose that the initial distribution $\rho_0(x)$ has full support. Then if we average Eq. (18) over $P(\mathbf{x}_{[0,\tau]})$ we get $\langle \bar{P}(\Theta\mathbf{x}_{[0,\tau]})/P(\mathbf{x}_{[0,\tau]}) \rangle = 1$, which is an integral fluctuation relation [3], $\langle e^{-\Sigma(\tau)} \rangle = 1$. Moreover, $\langle \ln[P(\mathbf{x}_{[0,\tau]})/\bar{P}(\Theta\mathbf{x}_{[0,\tau]})] \rangle \geq 0$ is a KL divergence, which can be rewritten in an appealing form as

$$\langle \Sigma(\tau) \rangle = \sum_{t=0}^{\tau-1} [D(\rho_t||r) - D(\rho_{t+1}||r')] \geq 0. \quad (22)$$

We notice that for the choice $r = \mu$ we recover the expression for mismatch cost sum in Eq. (12), while for $r = \nu$ we obtain Eq. (13), as expected. Crucially, for the two choices $r = \mu$ and $r = \nu$, the quantity $\langle \Sigma(\tau) \rangle$ provides a lower bound on the total average entropy production incurred in the physical implementation of the computational process, and therefore we may refer to it as the intrinsic mismatch cost associated to a given computation. We remark that here and above averages are over trajectories of fixed length τ ; that is, $\langle \Sigma(\tau) \rangle \equiv \mathbb{E}(\Sigma(\tau))$.

For more general choices of r , the quantity $\Sigma(\tau)$ can still be defined (as long as the distribution r has full support over \mathcal{X}); however, it cannot be guaranteed in general that $\langle \Sigma(\tau) \rangle$ would provide a lower bound on the underlying entropy production anymore. In particular, by taking $r = \pi$, the stationary state of the DTMC, $\Sigma(\tau)$ becomes the discrete-time nonadiabatic entropy production for a relaxation process, whose average reads

$$\langle \Sigma(\tau) \rangle = D(\rho_0||\pi) - D(\rho_\tau||\pi) \geq 0; \quad (23)$$

thus we recover the expression for EP proposed by Spohn [83] (see also Ref. [41]). Remarkably, in this case $\langle \Sigma(\tau) \rangle$ becomes nonextensive in time, contrary to the general case [cf. Eq. (22)]. As a consequence, the steady state π of the DTMC (whenever aperiodic and irreducible) becomes the natural candidate for the prior ν providing the infimum in Eq. (13) in the large time limit. Therefore we expect the nonadiabatic entropy production in Eq. (23) to provide the minimum dissipation of the computation in many cases of interest. However, it is worth remarking that for ensuring the average nonadiabatic entropy production to be a lower bound on the EP would require π to share support with the initial distribution ρ_0 —which would often not be the case in the computational context—and π being also invariant state in the time-reversed (underlying) physical dynamics of the computer [41].

B. Role of absolute irreversibility

In many models of computation in TCS, the initial distribution $\rho_0(x)$ over the states of the computational machine is restricted to a subset of computational states in \mathcal{X} . For instance, almost any automaton—in particular, not just a DFA but also a TM—starts in a single, predetermined state x_0 . Such a system may have a delta function initial distribution, $\rho_0(x) = \delta_{x,x_0}$. For such an initial distribution the quantity $e^{-\Sigma(\tau)} = \bar{P}(\Theta \mathbf{x}_{[0,\tau]}) / P(\mathbf{x}_{[0,\tau]})$ may become ill defined as there might be trajectories for which $P(\mathbf{x}_{[0,\tau]}) = 0$, but $\bar{P}(\Theta \mathbf{x}_{[0,\tau]}) > 0$, e.g., trajectories in the auxiliary dynamics that do only reach states different from x_0 . This phenomenon has been often referred to as *absolute irreversibility* [84,85].

Following the techniques in Refs. [84–86] one can circumvent the divergence associated with absolute irreversibility by restricting the averages over sets of trajectories for which the functional of the intrinsic mismatch cost $e^{-\Sigma(\tau)}$ is well defined. Adopting the language of modern probability theory [87], we call such sets *filtrations* (see also Ref. [19]). In particular, we denote \mathcal{F} the filtration containing all possible trajectories $\mathbf{x}_{[0,\tau]}$ taking place in $[0, \tau]$. Similarly, we call \mathcal{F}_{AI} the filtration containing all “absolutely irreversible” trajectories, that is, trajectories for which $P(\mathbf{x}_{[0,\tau]}) = 0$, but $\bar{P}(\Theta \mathbf{x}_{[0,\tau]}) > 0$. On the other hand, we denote the complementary set of “absolutely continuous” trajectories as \mathcal{F}_{AC} , such that $\mathcal{F} = \mathcal{F}_{\text{AC}} \cup \mathcal{F}_{\text{AI}}$. Using these definitions, an extended version of the integral fluctuation theorem (IFT) for the intrinsic mismatch cost follows

$$\langle e^{-\Sigma(\tau)} \rangle = 1 - \gamma_\tau, \quad (24)$$

where $0 \leq \gamma_\tau \leq 1$ the total probability that the time-reversed picture of any absolutely irreversible trajectory (i.e., belonging to \mathcal{F}_{AI}) occurs in the auxiliary dynamics

$$\gamma_\tau = \sum_{\mathbf{x}_{[0,\tau]} \in \mathcal{F}_{\text{AI}}} \bar{P}(\Theta \mathbf{x}_{[0,\tau]}) \leq 1. \quad (25)$$

Applying Jensen’s inequality $\langle e^x \rangle \geq e^{\langle x \rangle}$ to the IFT Eq. (24) we obtain a lower bound on the intrinsic mismatch cost, implying a minimum dissipation due to the restricted initial condition:

$$\langle \Sigma(\tau) \rangle \geq -\ln[1 - \gamma_\tau] \geq 0, \quad (26)$$

where the second inequality follows from $\gamma_\tau \geq 0$ and hence extends the applicability of Eq. (22) to systems showing absolute irreversibility. We remark that here absolute irreversibility arises because of the restricted initial distribution, but not because of the unidirectional transitions, since they have been flipped in the auxiliary dynamics according to Eq. (15). Fluctuation theorems similar to Eq. (24) have been previously derived within the canonical framework of stochastic thermodynamics for entropy production [84] and standard mismatch cost [49], as well as in the inclusive Hamiltonian framework for entropy production [29].

IV. THERMODYNAMICS OF COMPUTATIONS AT STOCHASTIC STOPPING TIMES

We now extend our analysis to investigate the thermodynamics of computations which first reach a computational state of interest at a time that varies depending on the random input provided to the computer. In doing so, we extend the martingale theory for stochastic thermodynamics [19] to accommodate unidirectional transitions and arbitrary initial distributions leading to absolute irreversibility.

Consider a random sequence of τ bits sequentially fed into a computer (e.g., a DFA) (see also Fig. 1):

$$\underbrace{0\ 0\ 0\ 1\ 0\ 1\ \dots\ 0\ 1\ 1\ 1}_{\tau \text{ bits}}, \quad (27)$$

with $\tau \geq 1$ being the *word length* processed by the machine. While processing a specific sequence, the computer jumps between its computational states, as described in Sec. II D. We are interested in the thermodynamics of the (physical implementation of the) computer during the time from when it starts to a *stopping time* \mathcal{T} , that is until when a *stopping condition* is met. For example, we will often consider that the stopping condition is simply that the computer has for the first time reached an accept state. Note that this stopping time generally takes a different value when processing different words. Since the words are generated by sampling a distribution, this means that the stopping time is a random variable.

Generalizing from this case to give a fully formal definition, a stopping time is the earliest instance when a particular condition concerning the entire trajectory generated by a stochastic process is met:

$$\mathcal{T}(\mathbf{x}_{[0,t]}) := \inf\{t \in [0, \tau] | \mathbf{x}_{[0,t]} \in \Omega\}, \quad (28)$$

where $\Omega \subseteq \mathcal{F}$ denotes the set of trajectories satisfying the stopping condition. For example, Ω might be the set of trajectories of a given DFA that have reached an accept state at least once.

Note that its definition in Eq. (28) involves a limit time τ . So the stopping time associated with each stochastic trajectory is a bounded random variable that obeys $0 \leq \mathcal{T} \leq \tau$. As shorthand, from now on we will typically just write “ \mathcal{T} ,” leaving the precise trajectory $\mathbf{x}_{[0,\mathcal{T}]}$ implicit. It is also worth remarking that the computational machine does not necessarily stop functioning at \mathcal{T} , but this variable can just signal to us the time at which a specific computation is processed (e.g., accepting a word). We will therefore sometimes refer to \mathcal{T} in this context as the *computation time*, which is a particular instance of a (bounded) stopping time.

A. Martingale theory with absolute irreversibility

Inspired by Ref. [16], we now introduce the stochastic distinguishability between the computational process and the auxiliary process. Stochastic distinguishability (with respect to time τ) evaluated at time $t \leq \tau$ is defined as

$$\delta_\tau(t) := \ln \frac{\rho_t(x_t)}{\bar{\rho}_{\tau-t}(x_t)}, \quad (29)$$

where $\bar{\rho}_{\tau-t}(x)$ is the probability distribution of the auxiliary process defined in Eq. (15), evaluated at the conjugate time $\tau - t$ for the state x_t . [Recall that the auxiliary dynamics has initial distribution $\bar{\rho}_0(x) = \rho_\tau(x)$; i.e., it is the distribution of the original dynamics at the limit time τ .] Stochastic distinguishability is a measure of the asymmetry between the original and the auxiliary dynamics and plays a crucial role in martingale theory for stochastic thermodynamics of nonstationary processes [19].

It will be useful to consider an associated process involving the stochastic distinguishability,

$$M_\tau(t) := e^{-\Sigma(t) - \delta_\tau(t)} = \frac{\bar{\rho}_{\tau-t}(x_t)}{\rho_t(x_t)} \left[\prod_{s=0}^{t-1} \frac{r(x_s) \rho_{s+1}(x_{s+1})}{\rho_s(x_s) r'(x_{s+1})} \right]. \quad (30)$$

In the second line of Eq. (30) we used the first equality in Eq. (18) together with Eq. (29) and the second line in Eq. (18). Note that $M_\tau(\tau) = e^{-\Sigma(\tau)}$ because $\delta_\tau(\tau) = 0$. In general though $\delta_\tau(t) \neq 0$ for $t < \tau$, and so $M_\tau(t) \neq e^{-\Sigma(t)}$ for such t .

An important property of M_τ is that the expectation of $M_\tau(\tau)$ conditioned on a fixed trajectory ending at a time $0 \leq t \leq \tau$ satisfies

$$\langle M_\tau(\tau) | \mathbf{x}_{[0,t]} \rangle = M_\tau(t) [1 - \alpha_\tau(t)], \quad (31)$$

where we introduced the quantity defined by

$$\alpha_\tau(t) := \sum_{\mathbf{x}_{[t+1,\tau]} \in \mathcal{F}_{\text{AI}}} \frac{\bar{P}(\Theta \mathbf{x}_{[t,\tau]})}{\bar{\rho}_{\tau-t}(x_t)} \leq 1, \quad t < \tau, \quad (32)$$

and $\alpha_\tau(\tau) := 0$ (see Appendix E for details). Combining Eq. (31) with the fact that $\alpha_\tau(t) \leq 1$, we establish that $M_\tau(t) = e^{-\Sigma(t) - \delta_\tau(t)}$ is a supermartingale,

$$\langle M_\tau(\tau) | \mathbf{x}_{[0,t]} \rangle \leq M_\tau(t); \quad (33)$$

i.e., its conditional expectation given a fixed trajectory of length $t < \tau$ monotonically decreases over time. Note that for $t = 0$ one has $\Sigma(0) = 0$ and, hence, Eq. (31) yields the IFT with absolute irreversibility [cf. Eq. (24)]:

$$\begin{aligned} \langle e^{-\Sigma(\tau)} \rangle &= \langle M_\tau(\tau) \rangle = \sum_{x_0} \rho_0(x_0) \langle M_\tau(\tau) | x_0 \rangle \\ &= \sum_{x_0} \rho_0(x_0) M_\tau(0) [1 - \alpha_\tau(0)] = 1 - \gamma_\tau, \end{aligned} \quad (34)$$

where we have used Eq. (25) in the last equality. In addition, in the absence of absolute irreversibility, \mathcal{F}_{AI} is the empty set and $\alpha_\tau(t) = 0$ for all $t \in [0, \tau]$. In such a case $M_\tau(t)$ in Eq. (31) becomes a martingale. Therefore, in that limit we would be able to use the analysis in Refs. [16,19] on the thermodynamics of systems with stochastic stopping times. However, that analysis does not directly apply for generic initial states $\rho_0(x)$ without full support.

B. Integral fluctuation relations with absolute irreversibility at stopping times

Fortunately, the fact that $M_\tau(t)$ is a supermartingale rather than a martingale when our system has absolute irreversibility does not prevent us from analyzing its thermodynamics at stopping times. To carry out such analysis, here we closely follow the derivation of Doob’s optimal stopping theorem for martingales, generalizing it to apply to supermartingales that are written as in Eq. (31).

As elaborated in Appendix F, this generalized form of the optimal stopping theorem provides a fluctuation theorem at stopping times, which is valid even in the presence of absolute irreversibility:

$$\langle e^{-\Sigma(\mathcal{T}) - \delta_\tau(\mathcal{T})} \rangle = \langle M_\tau(\mathcal{T}) \rangle = 1 - \Gamma_\tau, \quad (35)$$

where $\mathcal{T} \leq \tau$ is the (stochastic) stopping time, $\Gamma_\tau \in [0, 1]$ is a contribution from absolute irreversibility, and therefore $\langle e^{-\Sigma(\mathcal{T}) - \delta_\tau(\mathcal{T})} \rangle \leq 1$. Since $\langle \cdot \rangle$ is an average over trajectories, and different trajectories have different stopping times, $\langle e^{-\Sigma(\mathcal{T}) - \delta_\tau(\mathcal{T})} \rangle$ involves averaging over (stochastic) values of \mathcal{T} . This introduces statistical coupling between the time \mathcal{T} and the value $\Sigma(\mathcal{T})$.

The quantity Γ_τ appearing in Eq. (35) is an average of the functional $e^{-\delta_\tau(\mathcal{T})}$ evaluated at stopping times \mathcal{T} for trajectories leading to absolute irreversibility:

$$\Gamma_\tau := \sum_{\mathcal{T}=0}^{\tau} \sum_{\mathbf{x}_{[0,\mathcal{T}]} \in \mathcal{F}_{\text{AI}}^{(\mathcal{T})}} \bar{P}(\Theta \mathbf{x}_{[0,\mathcal{T}]}) \frac{\bar{\rho}_{\tau-\mathcal{T}}(x_{\mathcal{T}})}{\rho_{\mathcal{T}}(x_{\mathcal{T}})}. \quad (36)$$

To understand its meaning intuitively, first note that the second summation in Γ_τ is done over trajectories $\mathbf{x}_{[0,\mathcal{T}]}$ that belong to $\mathcal{F}_{\text{AI}}^{(\mathcal{T})}$, that is, trajectories verifying the stopping condition for the first time at \mathcal{T} , but that have zero probability to occur in the original process $P(\mathbf{x}_{[0,\mathcal{T}]}) = 0$, due to the restricted shape of the initial distribution $\rho_0(x)$. We notice also the presence of the distribution $\bar{\rho}_{\tau-\mathcal{T}}(x)$, which is due to $\delta_\tau(\mathcal{T})$. That is, Γ_τ consists of the total probability of trajectories starting at the stopped point $x_\mathcal{T}$ according to distribution $\bar{\rho}_{\tau-t}(x)$, and not turning back to the set of states with $\rho_0(x) > 0$ under the auxiliary dynamics. Recall also that the reference distribution r determining the precise meaning of $\Sigma(\mathcal{T})$ appears in Eq. (36) only implicitly, due to the definitions of \bar{P} and $\bar{\rho}_t$.

The inequality $\Gamma_\tau \leq 1$ is saturated when all trajectories are in the set $\mathcal{F}_{\text{AI}} = \mathcal{F}$, for which the sum over all trajectories in Eq. (36) is obtained; that is, $\Gamma_\tau = \sum_{t=0}^{\tau} \sum_{\mathbf{x}_{[0,t]} \in \mathcal{F}^{(t)}} \bar{P}(\Theta \mathbf{x}_{[0,t]}) \bar{\rho}_{\tau-t}(x_t) / \rho_t(x_t) = 1$. Moreover, we also have $\Gamma_\tau \geq 0$, since it is a sum of probabilities. Whenever the initial distribution $\rho_0(x)$ is not restricted in the state space, we obtain $\Gamma_\tau = 0$, and recover the standard form of the fluctuation theorem at stopping times for nonstationary processes [16,19].

It is worth remarking here that our previous results for fixed times [Eqs. (24) and (25)] can be directly obtained from Eqs. (35) and (39) by letting $\mathcal{T} = \tau$, i.e., when all trajectories are stopped at the final time τ , as we also discuss below in more detail. Our results thus provide an extension of martingale theory to cover different versions of mismatch costs in physical scenarios with absolute irreversibility, where martingales can be transformed into supermartingales via the correction term $\alpha_\tau(t)$ in Eq. (32), and stopping-time fluctuation relations can be derived from them.

Moreover, using the fact that $M_\tau(t)$ is a supermartingale [cf. Eq. (33)], we can also readily apply Doob's optional sampling theorem [88] for supermartingales to obtain (see Appendix G)

$$\langle e^{-\Sigma(\mathcal{T}_2) - \delta_\tau(\mathcal{T}_2)} \rangle \leq \langle e^{-\Sigma(\mathcal{T}_1) - \delta_\tau(\mathcal{T}_1)} \rangle, \quad (37)$$

where \mathcal{T}_1 and \mathcal{T}_2 are two stopping times, ordered such that $P(\mathcal{T}_2 \geq \mathcal{T}_1) = 1$, but otherwise arbitrary. Taking $\mathcal{T}_1 = \mathcal{T}$ and $\mathcal{T}_2 = \tau$, the above Eq. (37), together with the fluctuation theorem for stopping times [Eq. (35)] and fixed times [Eq. (34)] implies

$$\begin{aligned} \Gamma_\tau &= 1 - \langle e^{-\Sigma(\mathcal{T}_2) - \delta_\tau(\mathcal{T}_2)} \rangle \\ &\leq 1 - \langle e^{-\Sigma(\tau)} \rangle = \gamma_\tau, \end{aligned} \quad (38)$$

where we have used $\delta_\tau(\tau) = 0$. The above inequality implies that the absolute irreversibility term at stopping times Γ_τ is always smaller than its fixed-time counterpart γ_τ ; that is, absolute irreversibility implies always greater dissipation at fixed times than at stopping times.

C. Second-law inequalities at stopping times: Universal lower and upper bounds

If we apply Jensen's inequality $\langle e^x \rangle \geq e^{\langle x \rangle}$ to the fluctuation theorem of Eq. (35) we derive a second-law inequality at stopping times:

$$\langle \Sigma(\mathcal{T}) \rangle \geq -\langle \delta_\tau(\mathcal{T}) \rangle - \ln[1 - \Gamma_\tau]. \quad (39)$$

This sets a strict lower bound on the average dissipation incurred by a given computation up to an arbitrary stopping time \mathcal{T} , from its time-reversal symmetry breaking [as quantified by $\langle \delta_\tau(\mathcal{T}) \rangle$] and the absolute irreversibility (as quantified by Γ_τ).

Moreover, $\Gamma_\tau \geq 0$ implies that $-\ln[1 - \Gamma_\tau] \geq 0$. Therefore, Eq. (39) also implies the simpler bound:

$$\langle \Sigma(\mathcal{T}) \rangle \geq -\langle \delta_\tau(\mathcal{T}) \rangle. \quad (40)$$

These inequalities suggest that $\langle \Sigma(\mathcal{T}) \rangle$ might be negative whenever $\langle \delta_\tau(\mathcal{T}) \rangle \geq -\ln[1 - \Gamma_\tau] \geq 0$, as we discuss in detail below.

Any concave function [such as $\ln(x)$] of a supermartingale yields another supermartingale by Jensen's inequality. Therefore, the supermartingale property of $M_\tau(t)$ also implies that $\ln[M_\tau(t)] = -\Sigma(t) - \delta_\tau(t)$ is supermartingale. So $\Sigma(t) + \delta_\tau(t)$ is a submartingale; i.e., it conditionally increases with time. If we now invoke Doob's optional sampling theorem for submartingales, we get the inequality

$$\langle \Sigma(\mathcal{T}_2) + \delta_\tau(\mathcal{T}_2) \rangle \geq \langle \Sigma(\mathcal{T}_1) + \delta_\tau(\mathcal{T}_1) \rangle, \quad (41)$$

where again \mathcal{T}_1 and \mathcal{T}_2 are two ordered stopping times with $P(\mathcal{T}_2 \geq \mathcal{T}_1) = 1$. This inequality has several implications, the most immediate one being a second law for intervals between two ordered stopping times \mathcal{T}_1 and \mathcal{T}_2 :

$$\langle \Delta \Sigma(\mathcal{T}_1, \mathcal{T}_2) \rangle \geq -[\langle \delta_\tau(\mathcal{T}_2) \rangle - \langle \delta_\tau(\mathcal{T}_1) \rangle], \quad (42)$$

where $\langle \Delta \Sigma(\mathcal{T}_1, \mathcal{T}_2) \rangle := \langle \Sigma(\mathcal{T}_2) \rangle - \langle \Sigma(\mathcal{T}_1) \rangle$. This inequality provides a result applicable to both stochastic stopping and starting times, bounding the entropy production incurred for computations that both start and end at stochastic times.

As an example, inequality (42) provides a bound concerning the stochastic interval between the first time that a DFA enters an accept state and the earliest *subsequent* time that it again enters an accept state, after having left the set of accept states in between. Then the time up to \mathcal{T}_1 can be interpreted as the time it took for the DFA to accept a first substring of the full input word, and the time between \mathcal{T}_1 and \mathcal{T}_2 can be interpreted as the time it took for the DFA to accept a *second* substring of the full input word, a substring which follows the first one. Again, the inequality in Eq. (42) suggests that $\langle \Delta \Sigma(\mathcal{T}_1, \mathcal{T}_2) \rangle$ might eventually become negative for such a case, whenever

there is an increasing time-reversal asymmetry, i.e., for $\langle \delta_\tau(\mathcal{T}_2) \rangle > \langle \delta_\tau(\mathcal{T}_1) \rangle$.

Moreover, for the choice $\mathcal{T}_1 = \mathcal{T}$ and $\mathcal{T}_2 = \tau$, the inequality (41) gives us the following *upper* bound for the intrinsic mismatch cost at stopping times:

$$\langle \Sigma(\mathcal{T}) \rangle \leq \langle \Sigma(\tau) \rangle - \langle \delta_\tau(\mathcal{T}) \rangle. \quad (43)$$

The inequality (43) implies that whenever $\langle \delta_\tau(\mathcal{T}) \rangle \geq 0$, the intrinsic mismatch cost at stopping times will be upper bounded by its fixed-time counterpart, suggesting a drop in the thermodynamic costs of the computation at stopping times. On the other hand by taking $\mathcal{T}_1 = 0$ and $\mathcal{T}_2 = \mathcal{T}$ in Eq. (41), we obtain an alternative second law at stopping times, namely,

$$\langle \Sigma(\mathcal{T}) \rangle \geq D(\rho_0 \| \bar{\rho}_\tau) - \langle \delta_\tau(\mathcal{T}) \rangle, \quad (44)$$

to be compared with Eqs. (39) and (40). Here we have used that $\Sigma(0) = 0$ and

$$\langle \delta_\tau(0) \rangle = \sum_{x_0} \rho_0(x_0) \ln \frac{\rho_0(x_0)}{\bar{\rho}_\tau(x_0)} = D(\rho_0 \| \bar{\rho}_\tau). \quad (45)$$

This inequality provides us an alternative lower bound on the intrinsic cost of the computation. We notice that, while we expect it to be less tight in general than Eq. (39), it has the advantage of relying on the KL divergence between initial distribution ρ_0 and the final distribution in the auxiliary dynamics $\bar{\rho}_\tau$, which we expect to be more easily computable than Γ_τ in Eq. (36). Remarkably, combining Eqs. (43) and (44) we find a sandwich inequality for the intrinsic mismatch cost at stochastic times,

$$D(\rho_0 \| \bar{\rho}_\tau) - \langle \delta_\tau(\mathcal{T}) \rangle \leq \langle \Sigma(\mathcal{T}) \rangle \leq \langle \Sigma(\tau) \rangle - \langle \delta_\tau(\mathcal{T}) \rangle, \quad (46)$$

which provides both upper and lower bounds on $\langle \Sigma(\mathcal{T}) \rangle$.

The stopping-time fluctuation relation in Eq. (35) and the inequalities (39)–(44) for the intrinsic thermodynamic costs in computational processes with stochastic stopping times provide our main results. In the following we further discuss their interpretation and some of their implications, while in Sec. V we investigate their applications to computer science setups with some illustrative examples.

D. Thermodynamic interpretation and implications

The second-law inequality (40), $\langle \Sigma(\mathcal{T}) \rangle \geq -\langle \delta_\tau(\mathcal{T}) \rangle$ [as well the stronger versions Eqs. (39) and (44)], suggests that both the intrinsic mismatch cost and the underlying entropy production incurred in a given computation may be negative on average when evaluated at stopping times. To understand how this is possible in light of the data-processing inequality, we write $\langle \Sigma(\mathcal{T}) \rangle$ explicitly as the functional Eq. (18) averaged over many trajectories that are stopped each at a stochastic time \mathcal{T} :

$$\begin{aligned} \langle \Sigma(\mathcal{T}) \rangle = & \sum_{\mathcal{T}=0}^{\tau} p(\mathcal{T}) \sum_{t=0}^{\mathcal{T}-1} \left[\sum_{x_t} \rho_t(x_t | \mathcal{T}) \ln \frac{\rho_t(x_t)}{\mu(x_t)} \right. \\ & \left. - \sum_{x_{t+1}} \rho_{t+1}(x_{t+1} | \mathcal{T}) \ln \frac{\rho_{t+1}(x_{t+1})}{\mu'(x_{t+1})} \right]. \end{aligned} \quad (47)$$

Here, $p(\mathcal{T})$ denotes the probability that the stopping time takes value \mathcal{T} . Similarly, $\rho_t(x | \mathcal{T})$ denotes the conditional probability that the process takes the value x at time t given that the stopping condition is met at time \mathcal{T} . Because $\rho_t(x | \mathcal{T}) \leq \rho_t(x)$ in general, the terms $\sum_{x_t} \rho_t(x_t | \mathcal{T}) \ln[\rho_t(x_t)/\mu(x_t)]$ and $\sum_{x_{t+1}} \rho_{t+1}(x_{t+1} | \mathcal{T}) \times \ln[\rho_{t+1}(x_{t+1})/\mu'(x_{t+1})]$ are *not* KL divergences in general, and thus not necessarily greater than or equal to zero (see also Chap. 8.3 in Ref. [19]). This implies that $\langle \Sigma(\mathcal{T}) \rangle$ can in principle be negative. The second law at stopping times Eq. (40) permits $\langle \Sigma(\mathcal{T}) \rangle \leq 0$ whenever $\langle \delta_\tau(\mathcal{T}) \rangle \geq 0$; yet it is not clear when this would actually be the case.

The explicit expression for the stochastic distinguishability at stopping times reads

$$\langle \delta_\tau(\mathcal{T}) \rangle = \sum_{\mathcal{T}=0}^{\tau} \sum_{x_{\mathcal{T}}} p(\mathcal{T}) \rho_{\mathcal{T}}(x_{\mathcal{T}} | \mathcal{T}) \ln \frac{\rho_{\mathcal{T}}(x_{\mathcal{T}})}{\bar{\rho}_{\tau-\mathcal{T}}(x_{\mathcal{T}})}. \quad (48)$$

Equation (48) also reveals that $\langle \delta_\tau(\mathcal{T}) \rangle$ is not a KL divergence in general, and thus can in principle take any sign, yet so far only examples where $\langle \delta_\tau(\mathcal{T}) \rangle \geq 0$ have been reported in the literature. We remark that $\langle \delta_\tau(\mathcal{T}) \rangle$ is not a KL divergence unless $\mathcal{T} = \tau$, for which the process “stops” at the deterministic limit time τ , and one has that the joint stopping-time probability distribution

$$p(\mathcal{T}) \rho_t(x_t | \mathcal{T}) = \begin{cases} 0 & \text{if } \mathcal{T} < \tau \\ \rho_\tau(x_\tau) & \text{if } \mathcal{T} = \tau; \end{cases} \quad (49)$$

i.e., it takes the value, at time τ , of the solution of the master equation. Plugging in Eq. (49) in Eq. (48) one gets $\langle \delta_\tau(\mathcal{T} = \tau) \rangle = D(\rho_\tau \| \bar{\rho}_0) = 0$ because $\bar{\rho}_0 = \rho_\tau$. Analogously for $\mathcal{T} = \tau$, intrinsic mismatch cost $\langle \Sigma(\mathcal{T} = \tau) \rangle$ takes the expression Eq. (22), thus retrieving non-negativity $\langle \Sigma(\tau) \rangle \geq 0$. Note that other examples of negative entropy production at stopping times based on threshold criteria for work were first reported in Ref. [16] and for free energy more recently [89]. Such *gambling demon* [16] effect is allowed whenever $\langle \delta(\mathcal{T}) \rangle > 0$, which is not guaranteed for arbitrary stopping conditions but possible for wise stopping strategies, as shown experimentally in Refs. [16,89].

We can obtain further insight on this effect by decomposing the intrinsic mismatch cost at fixed times τ in two terms, one associated to intervals $[0, \mathcal{T}]$ up to the stopping time \mathcal{T} and $[\mathcal{T}, \tau]$ from the stopping time to the limit time τ ; that is,

$$\langle \Sigma(\tau) \rangle = \langle \Sigma(\mathcal{T}) \rangle + \langle \Delta \Sigma(\tau, \mathcal{T}) \rangle \geq 0, \quad (50)$$

which follows from the fact that \mathcal{T} is a single-valued function of the trajectory. Since $\langle \Sigma(\tau) \rangle \geq 0$, the above decomposition implies that, whenever $\langle \Sigma(\mathcal{T}) \rangle < 0$, such a negative value must be compensated by an incremented mismatch cost $\langle \Delta \Sigma(\tau, \mathcal{T}) \rangle \geq \langle \Sigma(\tau) \rangle$ incurred in the interval $[\mathcal{T}, \tau]$, if no external action is taken on the system at time \mathcal{T} to physically stop the dynamics. These considerations will be valid also in cases where the stopping condition is structurally imposed through the dynamical evolution of the computational process, e.g., using absorbing accept states to “stop” the computation, as it is the case in some models of DFAs.

The role of absolute irreversibility as captured in the stronger inequality (39) with $-\ln[1 - \Gamma_\tau] \geq 0$ makes more difficult the observation of negative average intrinsic mismatch cost, since it would require a higher time-reversal asymmetry in the dynamical evolution leading to large distinguishabilities $\langle \delta(\mathcal{T}) \rangle > -\ln[1 - \Gamma_\tau]$ [and similarly for inequality (44)]. Remarkably, however, the examples explored in Sec. V show how still dissipation can be reduced at stopping times thanks to a positive time-reversal asymmetry $\langle \delta(\mathcal{T}) \rangle > 0$, in agreement with Eq. (43) above. This reduction might be linked to the information needed to execute the stopping condition \mathcal{T} , similarly to what happens in feedback control scenarios [90,91]. However, a general relation between these two quantities remains unknown.

The second-law inequality at stopping times Eq. (40) can be further rewritten using Eq. (19) in a form reminiscent of Landauer’s principle:

$$-\langle \Delta \phi(\mathcal{T}) \rangle \geq -\langle \Delta S_{\text{sys}}(\mathcal{T}) \rangle - \langle \delta_\tau(\mathcal{T}) \rangle, \quad (51)$$

where the lhs accounts for the excess entropy flow dissipated into the environment as a consequence of a drop in Shannon entropy of the computational states, $-\langle \Delta S_{\text{sys}}(\mathcal{T}) \rangle$. Again, whenever $\langle \delta_\tau(\mathcal{T}) \rangle > 0$, the above inequality suggests that the entropy flow to the environment may be eventually reduced. Here it is also worth noticing that even in the case in which trajectories are stopped when returned to the initial state (as in the DFA example in Sec. V), the average system entropy change at stopping times, namely, $\langle \Delta S_{\text{sys}}(\mathcal{T}) \rangle = \langle S_{\text{sys}}(\mathcal{T}) \rangle - S(\rho_0)$, with

$$\langle S_{\text{sys}}(\mathcal{T}) \rangle = -\sum_{\mathcal{T}} p(\mathcal{T}) \sum_{x_{\mathcal{T}}} \rho_{\mathcal{T}}(x_{\mathcal{T}} | \mathcal{T}) \ln \rho_{\mathcal{T}}(x_{\mathcal{T}}), \quad (52)$$

is nonzero even when $x_{\mathcal{T}} = x_0$ for all \mathcal{T} , since in general the distribution $\rho_{\mathcal{T}}(x) \neq \rho_0(x)$, as corresponds to a relaxation process.

The second-law inequalities derived above not only can be applied to assess stochastic stopping times of a computation, but also to stochastic starting times; see Eqs. (41) and (42). This extension allow us to apply our theory to

computations that may “stop” at multiple consecutive times $\mathcal{T}_1 < \mathcal{T}_2 < \dots < \mathcal{T}_n$ (see Sec. V for a particular example in a DFA) or to the concatenations of simpler computations that start at a stochastic time, after the previous one is accomplished. We will further elaborate on the application of starting times to the computation of concatenated words with stochastic resetting in Sec. VII.

V. APPLICATION TO DETERMINISTIC FINITE AUTOMATA

In this section, we analyze minimal yet insightful examples of computations executed by deterministic finite automata. A computational task for a DFA starts by it receiving a sequence of exogenously generated symbols, an input string or an input word ω . As the DFA iteratively processes the symbols of the input string, it makes associated transitions among its possible states. Here we first assume that the sequence of symbols to the DFA are produced in an independent identically distributed (i.i.d.) manner and so the time evolution over the DFA states while processing those strings can be modeled using a DTMC. Then we will move to the case of input symbols that are not produced in an i.i.d. manner, but from a Markovian source. In the following examples, we consider two minimal DFA models that processes binary strings. In the first example involving i.i.d. symbol sources, the DFA under consideration accepts strings which encode binary numbers divisible by four, e.g., 0 (zero), 100 (four), 1100 (twelve), etc. In the second example, involving non-i.i.d. sources, we use a DFA that accepts strings which encode binary numbers divisible by three. In all cases, we assume that the input string behaves as an information reservoir [74] whose symbols are not modified by the computation, hence not leading to further energy or entropy changes (see Sec. II D).

The state of the DFA when a stopping condition is reached (e.g., when the DFA enters a designated accept state) defines a computation that the DFA performs on that string. However, this computation can be followed by further processing of input symbols up to a limit time τ (e.g., the DFA may exit the accept state in forthcoming iterations). In this sense our results for stopping times can be applied to various situations, for example, (i) computations generated by input words of fixed length τ where we ask about the value of thermodynamic quantities when visiting the accept state for the first (or the n th) time and (ii) computations that may actually end when visiting the accept state for some reason (e.g., the accept state is an absorbing state of the DFA or there exists an external mechanism that activates when the accept state is reached to stop the dynamics). In particular, we can always modify a given DFA by removing all edges of the associated directed graph that leave an accepting state. This turns the accept state into an absorbing state (or set of states, if there is more than one accepting state).

A. Processing symbols from i.i.d. sources

As mentioned above, consider the DFA from Fig. 1, initialized to state q_0 with certainty, and that its computation starts by processing a stream of binary letters generated as an i.i.d. sequence of 0's and 1's, with $p_0 \leq 1$ the probability to observe a 0 and $p_1 = 1 - p_0$ the probability to observe a 1. Under this assumption, the time evolution of the DFA's states follows a DTMC over four computational states $q_0, q_1, q_2,$ and q_3 , with transition probabilities as indicated in Fig. 2(a). All together, the Markov chain associated with the DFA's dynamics is characterized by its initial state,

$$\rho_0 = [1 \ 0 \ 0 \ 0]^\dagger, \quad (53)$$

with \dagger denoting here matrix transposition, and the transition matrix

$$\mathbf{W} = \begin{bmatrix} p_0 & 0 & p_0 & 0 \\ p_1 & 0 & p_1 & 0 \\ 0 & p_0 & 0 & p_0 \\ 0 & p_1 & 0 & p_1 \end{bmatrix}. \quad (54)$$

It follows that for $t = 1$ we have

$$\rho_1 = \mathbf{W}\rho_0 = [p_0 \ p_1 \ 0 \ 0]^\dagger, \quad (55)$$

whereas for larger times $t \geq 2$,

$$\rho_t = \mathbf{W}^t \rho_0 = [p_0^2 \ p_0 p_1 \ p_0 p_1 \ p_1^2]^\dagger \equiv \pi; \quad (56)$$

i.e., the dynamics already reaches the stationary state at the second iteration.

For computing the auxiliary dynamics for this DFA's DTMC, we would need to identify the reference distribution $r(x)$ appearing in Eq. (15) as the prior $\mu(x)$ minimizing the mismatch cost sum in Eq. (8) or $\nu(x)$ leading to its minimum in Eq. (13). For simplicity, here we assume $r(x) = \pi(x)$, the stationary state of the DFA dynamics. This is a reasonable assumption as long as the induced DTMC is aperiodic, irreducible, and π has full support over the computational states. As discussed before, since Σ becomes nonextensive in time in this case, there are reasons to expect minimal dissipation in the steady state (see also Refs. [92–94]).

The auxiliary dynamics starts in $\bar{\rho}_0 = \rho_\tau$, which can take two possible values depending on the value of the final maximum time of the computation τ : If $\tau = 1$, we have $\bar{\rho}_0 = \rho_1$, whereas for $\tau \geq 2$, we have $\bar{\rho}_0 = \pi$. Following Eq. (15) for the transition probability, with $r = \pi$, the stationary distribution given in Eq. (56), we obtain the transition matrix associated with the auxiliary dynamics:

$$\bar{\mathbf{W}} = \begin{bmatrix} p_0 & p_0 & 0 & 0 \\ 0 & 0 & p_0 & p_0 \\ p_1 & p_1 & 0 & 0 \\ 0 & 0 & p_1 & p_1 \end{bmatrix}, \quad (57)$$

as illustrated in Fig. 2(b). It then follows that for computations ending at $\tau = 1$ we have $\bar{\rho}_1 = \bar{\mathbf{W}}\bar{\rho}_0 = \bar{\mathbf{W}}\rho_1 = [p_0 \ 0 \ p_1 \ 0]^\dagger$. On the other hand, since by construction, the auxiliary dynamics Eq. (15) will always preserve the steady state for $r = r' = \pi$, it follows that in the case $\tau \geq 2$, the auxiliary dynamics is stationary at all times t , that is, $\bar{\rho}_t^{(\tau)} = \bar{\mathbf{W}}^t \bar{\rho}_0 = \bar{\mathbf{W}}^t \pi = \pi$, with π given by Eq. (56).

The intrinsic mismatch cost in Eq. (18), evaluated over a trajectory $\mathbf{x}_{[0,\tau]}$, reduces in this case to the (discrete-time) stochastic nonadiabatic EP:

$$\begin{aligned} \Sigma(\mathbf{x}_{[0,\tau]}) &= \sum_{t=0}^{\tau-1} \left[\ln \frac{\rho_t(x_t)}{\pi(x_t)} - \ln \frac{\rho_{t+1}(x_{t+1})}{\pi(x_{t+1})} \right] \\ &= \ln \frac{\rho_0(x_0)}{\pi(x_0)} - \ln \frac{\rho_\tau(x_\tau)}{\pi(x_\tau)}, \end{aligned} \quad (58)$$

with $x_t \in \mathcal{X} = \{q_0, q_1, q_2, q_3\}$ for all t , which depends only on the initial and final states.

Having obtained the system probability distribution at all times for the original and auxiliary dynamics, we are now ready to compute thermodynamic quantities at stopping times. In particular, we consider the family of stopping times

$$\mathcal{T} = \min(\mathcal{T}_1, \tau), \quad (59)$$

with τ fixing a time horizon and $\mathcal{T}_1 \geq 1$ the first time the DFA returns to the accept state q_0 , hence accepting a word as a multiple of four (including “0”). From numerical simulations, we obtained sample histograms for the stopping time \mathcal{T} given by Eq. (59) for three different choices of the limit time τ ; see Fig. 3. There we observe the first peak at $\mathcal{T} = 1$ in the three plots, corresponding to the cases where the first incoming symbol is “0” and the word is then accepted. In order to allow longer accepted words we need $\tau > 2$, such that $\mathcal{T}_1 = 3$ (accepting four “100”) or $\mathcal{T}_1 = 4$ (accepting twelve “1100”), etc. Note, however, that with the stopping condition given in Eq. (59), we do not capture the acceptance of some of the multiples of four like, e.g., eighth “1000,” since the stopping condition would already be verified at previous symbol of the string, “100,” corresponding to four. The same happens for any other accepted number to which an arbitrary number of zeros are attached at the end. For assessing the acceptance of such numbers, extra stopping conditions such as \mathcal{T}_n , i.e., the n th time the DFA returns to the accept state q_0 , are needed (see example in Sec. V C).

For all trajectories in which $\mathcal{T} = \mathcal{T}_1 < \tau$, i.e., the word is accepted before the limit time τ is reached, we have $x_{\mathcal{T}_1} = q_0$, the accept state, and thus




FIG. 3. Stopping-time statistics for the DFA recognizing binary numbers divisible by four, obtained from 10^4 numerical simulations of the DTMC sketched in Fig. 2, with initial state q_0 , and an absorbing accept state set also at q_0 . Simulations are done by feeding the DFA with i.i.d. binary sequences with probability of observing the letter 0 given by $p_0 = 0.9$, obtained from Monte Carlo simulations of the discrete-time Markov chain sketched in Fig. 2(c).

$$\Sigma(\mathcal{T}_1) = -\ln \rho_{\mathcal{T}_1}(q_0) = \begin{cases} -\ln p_0 & \text{if } \mathcal{T}_1 = 1 \\ -2 \ln p_0 & \text{if } \mathcal{T}_1 \geq 2, \end{cases} \quad (60)$$

and the stochastic distinguishability in Eq. (29) is

$$\delta_\tau(\mathcal{T}_1) = \ln \frac{\rho_{\mathcal{T}_1}(q_0)}{\bar{\rho}_{\tau-\mathcal{T}_1}(q_0)} = \begin{cases} -\ln p_0 & \text{if } \mathcal{T}_1 = 1 \\ 0 & \text{if } \mathcal{T}_1 \geq 2, \end{cases} \quad (61)$$

where we used the fact that (by construction) $\tau > \mathcal{T}_1 \geq 1$ and, hence, $\bar{\rho}_{\tau-\mathcal{T}_1} = \pi$.

If however $\mathcal{T}_1 \geq \tau$, the dynamics stops at the maximum time $\mathcal{T} = \tau$, independently of the state x_τ , and we obtain

$$\Sigma(\tau) = -\ln \frac{\rho_\tau(x_\tau)\pi(q_0)}{\pi(x_\tau)} = \begin{cases} -\ln p_0 & \text{if } \tau = 1 \\ -2 \ln p_0 & \text{if } \tau \geq 2. \end{cases} \quad (62)$$

Note that the case $\tau \geq 2$ is independent of x_τ because the system has already reached its stationary state, and thus $\Sigma(\tau) = -\ln \pi(q_0) = -2 \ln p_0$ for all $x_\tau \neq q_0$. On the other hand, the stochastic distinguishability verifies $\delta_\tau(\tau) = 0$ since $\bar{\rho}_0 = \rho_\tau$ always.

Using the above calculations we obtain the average intrinsic mismatch cost at the stopping time Eq. (59) for all $\tau \geq 2$ as

$$\langle \Sigma(\mathcal{T}) \rangle = (p_0 - 2) \ln p_0 \geq 0, \quad (63)$$

which follows from

$$\begin{aligned} \langle \Sigma(\mathcal{T}) \rangle &= -P(\mathcal{T}_1 = 1) \ln p_0 - \sum_{t=2}^{\tau} P(\mathcal{T}_1 = t) 2 \ln p_0 \\ &\quad - P(\mathcal{T}_1 > \tau) 2 \ln p_0 \\ &= -p_0 \ln p_0 - (1 - p_0) 2 \ln p_0 \\ &= (p_0 - 2) \ln p_0, \end{aligned} \quad (64)$$

where we have used $P(\mathcal{T}_1 = 1) = p_0$ and, thus, $P(\mathcal{T}_1 > 1) = 1 - p_0$. In addition, using Eq. (61), we obtain the average stochastic distinguishability at the stopping time Eq. (59) for all $\tau \geq 2$:

$$\langle \delta_\tau(\mathcal{T}) \rangle = -p_0 \ln p_0 \geq 0, \quad (65)$$

where again we have used $P(\mathcal{T}_1 = 1) = p_0$. Note that the above expressions also remain valid in the limit of large input word lengths, $\tau \rightarrow \infty$.

To tackle the contribution from absolute irreversibility at stopping times, it is convenient to first identify which trajectories contribute to Γ_τ in Eq. (36). These are trajectories that are stopped at $\mathcal{T} \leq \tau$ (either with or without reaching q_0) and have zero probability to occur in the original dynamics. Note that the original dynamics is a Markov chain with initial state $\rho_0(x) = \delta_{x,q_0}$. The set of absolutely irreversible trajectories at stopping times consists of two sets: (i) trajectories that do not start in q_0 and reach q_0 with $\mathcal{T} \leq \tau$ in the original dynamics and (ii) trajectories of length τ that do not start at q_0 and do not reach q_0 in the original dynamics.

Let us now flesh out the list of such trajectories $\mathbf{x}_{[0,\mathcal{T}]}$ classified by the value of \mathcal{T} for the special case $\tau = 2$.

- (i) $q_2 q_0$ reaches the accept state at $\mathcal{T} = 1$, yet it has zero probability to occur in the original dynamics with $\rho_0(q_2) = 0$.
- (ii) $q_1 q_2 q_0$ and $q_3 q_2 q_0$ reach the accept state at $\mathcal{T} = 2$, yet they have zero probability to occur in the original dynamics since $\rho_0(q_1) = \rho_0(q_3) = 0$.
- (iii) $q_1 q_2 q_1$, $q_1 q_3 q_2$, $q_1 q_3 q_3$, $q_2 q_1 q_2$, $q_2 q_1 q_3$, $q_3 q_2 q_1$, $q_3 q_3 q_2$, and $q_3 q_3 q_3$ are stopped at $\mathcal{T} = 2$ without reaching the accept state. They have zero probability in the original dynamics because their initial state is different from q_0 .

All the sequences listed above are such that they would halt the computation at the stopping time $\mathcal{T} = \min(\mathcal{T}_1, 2)$; they have nonzero probability in the auxiliary dynamics but zero probability in the original dynamics. In order to calculate the absolute irreversibility correction term Γ_τ in Eq. (36), we thus need the probability of the above trajectories to occur in time-reversed order in the auxiliary dynamics. More precisely, one needs to compute $\bar{P}(\Theta \mathbf{x}_{[0,\mathcal{T}]})$ multiplied by $\bar{\rho}_{\tau-\mathcal{T}}(x_\tau)/\rho_\tau(x_\tau) = \pi(x_\tau)/\rho_\tau(x_\tau)$, which in this

case is equivalent to modifying their initial condition to $\pi(x_T)$, i.e., to compute the following path probabilities:

$$\begin{aligned}
\bar{P}(q_0, q_2|q_0)\pi(q_0) &= p_0^2 p_1, \\
\bar{P}(q_0, q_2, q_1|q_0)\pi(q_0) &= p_0^2 p_1 p_0, \\
\bar{P}(q_0, q_2, q_3|q_0)\pi(q_0) &= p_0^2 p_1 p_1, \\
\bar{P}(q_1, q_2, q_1|q_1)\pi(q_1) &= p_0 p_1 p_1 p_0, \\
\bar{P}(q_1, q_2, q_3|q_1)\pi(q_1) &= p_0 p_1 p_1 p_1, \\
\bar{P}(q_2, q_1, q_2|q_2)\pi(q_2) &= p_0 p_1 p_0 p_1, \\
\bar{P}(q_2, q_3, q_3|q_2)\pi(q_2) &= p_0 p_1 p_1 p_1, \\
\bar{P}(q_2, q_3, q_1|q_2)\pi(q_2) &= p_0 p_1 p_1 p_0, \\
\bar{P}(q_3, q_3, q_3|q_3)\pi(q_3) &= p_1^2 p_1 p_1, \\
\bar{P}(q_3, q_3, q_1|q_3)\pi(q_3) &= p_1^2 p_1 p_0, \\
\bar{P}(q_3, q_1, q_2|q_3)\pi(q_3) &= p_1^2 p_0 p_1.
\end{aligned} \tag{66}$$

Summing up all the contributions in Eq. (66) leads us to the absolute irreversibility contribution [cf. Eq. (36) for the general formula]:

$$\begin{aligned}
\Gamma_2 &= p_0 p_1 (p_0 + p_0^2 + 4p_1 p_0 + 4p_1^2) + p_1^4 \\
&= 1 - p_0^2.
\end{aligned} \tag{67}$$

Combining all the terms above, we observe that for the stopping time $\mathcal{T} = \min(\mathcal{T}_1, 2)$:

$$\langle \Sigma(\mathcal{T}) \rangle = -\langle \delta_2(\mathcal{T}) \rangle - \ln[1 - \Gamma_2]. \tag{68}$$

In other words, the second law at stopping times given by Eq. (39) is saturated over the stopping time given by Eq. (59) for $\tau = 2$, as it is illustrated in Fig. 4 for different values of the probability of incoming zeros p_0 . As can be appreciated in this figure, the positive sign of the term $\langle \delta_2(\mathcal{T}) \rangle > 0$ implies that the intrinsic mismatch cost at stopping times $\langle \Sigma(\mathcal{T}) \rangle = -2 \ln p_0 + p_0 \ln p_0$, see Eq. (65), is smaller than its value at fixed times,

$$\langle \Sigma(\tau) \rangle = -2 \ln p_0 \geq \langle \Sigma(\mathcal{T}) \rangle, \tag{69}$$

in spite of the presence of the absolute irreversibility contribution with Γ_2 . For $\tau > 2$, we have $\Gamma_\tau \leq \Gamma_2$, which follows by combining the equality in Eq. (68) with the generic bound in Eq. (39). In any case, the inequality $\langle \Sigma(\tau) \rangle \geq \langle \Sigma(\mathcal{T}) \rangle$ holds for any limit time τ for this example.

When p_0 approaches 1 (words with a high number of zeros) the dynamics cannot escape from the initial state q_0 and the steady state π becomes equal to the initial distribution ρ_0 . In this limit, the DTMC dynamics becomes fully stationary and, hence, the intrinsic mismatch cost becomes zero for every trajectory, the time-reversal




FIG. 4. Illustration of analytical results for the second law at stopping times applied to the discrete-time Markov chain model of the DFA recognizing binary strings whose length is a multiple of four [see Fig. 2(a)]. Here the computation stops at $\mathcal{T} = \min(\mathcal{T}_1, \tau)$ for the limit time $\tau = 2$, or earlier if the accept state is reached in one iteration. Symbols represent analytical results for the averages at the stopping time for the relevant thermodynamic quantities: intrinsic mismatch cost (nonadiabatic entropy production) $\langle \Sigma(\mathcal{T}) \rangle$, given by Eq. (63) for prior equal to the stationary probability (blue filled squares); minus the stochastic distinguishability $-\langle \delta_2(\mathcal{T}) \rangle$, given by Eq. (65) (blue dotted line); fixed-time nonadiabatic entropy production $\langle \Sigma(\tau) \rangle$ evaluated over trajectories of the same length $\tau = 2$ (open symbols); and the absolute irreversibility contribution $-\ln[1 - \Gamma_2]$, given by Eq. (36) (blue dashed line). The blue solid line is given by the sum $-\langle \delta_2(\mathcal{T}) \rangle - \ln[1 - \Gamma_2]$ which in this example equals $\langle \Sigma(\mathcal{T}) \rangle$, thus saturating the second law Eq. (39). The horizontal black thick line is set to zero as a reference value.

asymmetry is lost, and the absolute irreversibility is no longer present, leading to a drop in the three quantities on the rhs of Fig. 4. As we move away from that limit, the mismatch cost increases (both at stopping and fixed times), signaling the energetic costs incurred by the computational task, which grow as p_0 decreases. This can be justified by the fact that the dynamics on the DTMC spreads more easily over all computational states as p_1 increases [see Fig. 4(a)], leading to a greater distinction between initial and steady-state distributions. In this case we also observe nonzero stochastic distinguishability and an increasingly large absolute irreversibility term. In the limit $p_0 \rightarrow 0$ (words with a high number of ones) accepting a word becomes almost impossible, and hence the stopping occurs most probably at the maximum time $\mathcal{T} \simeq \tau$, leading again to zero stochastic distinguishability. We notice that in this limit π tends to localize at state q_3 and, hence, it would lead to $\langle \Sigma(\tau) \rangle \rightarrow \infty$, which is not physically meaningful. The catch point is that in this limit the fixed point π would not be equal to the prior μ or ν anymore.

B. Uniform prior

We now implement the analysis in Sec. VA for a different setting, where the strings are generated i.i.d. and we consider the same four-states DFA, now with a uniform prior distribution over its states:

$$r = [1/4 \quad 1/4 \quad 1/4 \quad 1/4]^\dagger. \quad (70)$$

The evolution of r under \mathbf{W} after one iteration yields

$$r' = \mathbf{W}r = [p_0/2 \quad p_1/2 \quad p_0/2 \quad p_1/2]^\dagger. \quad (71)$$

Because r changes after one iteration, we write Σ as in Eq. (19) for $\tau > 0$,

$$\Sigma(\mathbf{x}_{[0,\tau]}) = \ln \frac{\rho_0(x_0)}{\rho_\tau(x_\tau)} + \sum_{t=0}^{\tau-1} \ln \frac{r'(x_{t+1})}{r(x_t)}, \quad (72)$$

where the first term is the system entropy change $\Delta S_{\text{sys}}(\mathbf{x}_{[0,\tau]})$ and the second term is the nonequilibrium potential $\Delta\phi(\mathbf{x}_{[0,\tau]})$ in Eq. (21). Unlike for the stationary prior, now this term is extensive with time [cf. Eq. (58)]. Note that in this case Σ is no longer equal to the nonadiabatic EP associated with the stochastic trajectory $\mathbf{x}_{[0,\tau]}$.

The uniform distribution is not invariant under the map \mathbf{W} ; hence the intrinsic mismatch cost Σ associated with a stochastic trajectory $\mathbf{x}_{[0,\tau]}$ is extensive with time. This implies that, unlike for the case of stationary prior (see Sec. VA), the averages of Σ at fixed times τ as well as at stopping times with limit time τ [of the form of Eq. (59)] will crucially depend on τ . This is also the case for any other choice for the prior distribution which differs from the stationary distribution.

In Fig. 5 we show the intrinsic mismatch cost $\langle \Sigma(\mathcal{T}) \rangle$ at the stopping time $\mathcal{T} = \min(\mathcal{T}_1, \tau)$ for the DFA with the uniform prior for two different values of τ , and compare it with the case of stationary prior, Eq. (63). We observe that the uniform prior leads to higher values for the intrinsic mismatch cost for high values of p_0 , while for low p_0 values the tendency can be inverted. However, when increasing τ sufficiently we always obtain a lower cost for the stationary prior, as expected from its nonextensivity. Indeed we observe a tendency for the mismatch cost at stopping times $\langle \Sigma(\mathcal{T}) \rangle$ to saturate when increasing the limit time τ , in contrast with the linear scaling of $\langle \Sigma(\tau) \rangle$ with τ . In Appendix H we confirm this point by studying in more detail the scaling behavior of these two quantities as a function of τ .

We test the sandwich inequality in Eq. (46) comprising the upper and lower bounds to $\langle \Sigma(\mathcal{T}) \rangle$ in Eqs. (43) and (44), respectively. As can be appreciated in Fig. 5, both inequalities provide useful bounds that become tighter for small τ , and are simultaneously saturated at the point $p_0 = 1/2$. This example also reveals that again there is a




FIG. 5. Numerical results for the average intrinsic mismatch cost $\langle \Sigma(\mathcal{T}) \rangle$ (symbols) at stopping times $\mathcal{T} = \min(\mathcal{T}_1, \tau)$ with uniform prior, evaluated for the DFA in Fig. 2(a) processing i.i.d. binary input data, as a function of the probability of input symbol 0. We used different limit times: $\tau = 5$ (blue filled squares), and $\tau = 14$ (red filled circles). The solid lines correspond to the lower bound predicted by Eq. (44) and given by $D(\rho_0 \parallel \bar{\rho}_\tau) - \langle \delta_\tau(\mathcal{T}) \rangle$ for $\tau = 5$ (blue solid line), and $\tau = 14$ (red solid line), while the dashed lines are the corresponding upper bounds in Eq. (43), $\langle \Sigma(\tau) \rangle - \langle \delta_\tau(\mathcal{T}) \rangle$ for same values of τ . Averages are estimated from 10^4 numerical simulation for each parameter value. The thick gray line is the average cost at stopping times for stationary prior π , see Eq. (63) and Fig. 4, and the vertical dashed line is set to $p_0 = 1/2$ as a reference value. Inset: $\langle \Sigma(\tau) \rangle - \langle \Sigma(\mathcal{T}) \rangle$ (solid line) and $\langle \delta_\tau(\mathcal{T}) \rangle$ (dashed line) as a function of p_0 for $\tau = 14$. The horizontal dotted line is set to zero as a reference value.

reduction of intrinsic costs at stopping times with respect to fixed times; that is, $\langle \Sigma(\tau) \rangle \geq \langle \Sigma(\mathcal{T}) \rangle$ holds over the entire parameter range of probability of symbol 0, p_0 , and the limit time τ , as shown in the inset of Fig. 5. This reduction is guaranteed by a positive value of the stochastic distinguishability $\langle \delta(\mathcal{T}) \rangle > 0$ in the range $p_0 \geq 1/2$ [cf. Eq. (43)], but, interestingly, it is also verified even for $\langle \delta(\mathcal{T}) \rangle < 0$ as it happens for $p_0 \leq 1/2$.

C. Beyond i.i.d. sources

Thus far we have analyzed the statistics of a DFA processing inputs generated by a source of i.i.d. bits, which induces a Markovian dynamics for the time evolution of the computational states. This is, however, one of the simplest possible computational processes, as, e.g., regular languages recognized by DFAs are often composed of correlated words. To illustrate the applicability of our theory to computing thermodynamic costs of DFAs processing arbitrary strings from arbitrary languages, it is mandatory to consider DFAs processing non-i.i.d. sequences.

In processing a generic non-i.i.d. sequence, the dynamics over the computational states of a DFA is in general a




FIG. 6. (a) Minimal DFA that accepts binary multiples of three with three states $z_i = \{q_0, q_1, q_2\}$ with same start and accept state q_0 . (b) Associated DTMC, where the automaton processes input strings generated by a non-i.i.d. source of input symbols with probabilities depending only on the last processed symbol $b_i = \{0, 1\}$. See Eq. (73) for the corresponding transition probability matrix.

non-Markovian process. However, one can extend the computational state space such that our formalism can be applied. For the analysis in this section it is important to mention the distinction between the “computational states” of the DTMC computational state space \mathcal{X} and the states of the DFA. In particular, we refer to states of the DFA (as in the usual TCS definition) as logical states of the DFA, and remind that with computational states we refer to the sets of variables which describe the entire state space for a computational process of interest, as introduced in Sec. II.

Now consider that the (process generating the) input string itself is a DTMC characterized by time-independent transition probabilities $p(b_{i+1}|b_i)$ for the $(i+1)$ th bit to be equal to $b_{i+1} = \{0, 1\}$ given that the i th symbol (bit) of the string is $b_i = \{0, 1\}$. In this case, the logical state of, e.g., a three-state DFA $z_i = \{q_0, q_1, q_2\}$ processing this input string is not a DTMC, although by constructing the computational state space as the Cartesian product of $z_i = \{q_0, q_1, q_2\}$ and $b_i = \{0, 1\}$, we encode the current computational state $x_i = \{z_i, b_i\}$ as the logical state of the DFA z_i and the most recent input symbol fed to the DFA b_i . In this case, one is left with a DTMC with six possible computational states, for which our formalism can be readily applied to tackle the thermodynamic properties.

As an example, we consider the minimal DFA that accepts binary multiples of 3, shown in Fig. 6(a), leading to the DTMC represented in Fig. 6(b). The probabilities to obtain input bits 0 or 1 given the last input symbol are fixed

and denoted by $p(i|j) := p_{ij}$ for $i, j = \{0, 1\}$. They satisfy $p_{00} + p_{10} = 1$ and $p_{01} + p_{11} = 1$. Ordering the computational states’ distribution as $\rho_t = [\rho_t(q_0, 0), \rho_t(q_1, 0), \rho_t(q_2, 0), \rho_t(q_0, 1), \rho_t(q_1, 1), \rho_t(q_2, 1)]$, we obtain a 6×6 transition matrix given by

$$\mathbf{W} = \begin{bmatrix} p_{00} & 0 & 0 & p_{01} & 0 & 0 \\ 0 & 0 & p_{00} & 0 & 0 & p_{01} \\ 0 & p_{00} & 0 & 0 & p_{01} & 0 \\ 0 & p_{10} & 0 & 0 & p_{11} & 0 \\ p_{10} & 0 & 0 & p_{11} & 0 & 0 \\ 0 & 0 & p_{10} & 0 & 0 & p_{11} \end{bmatrix}. \quad (73)$$

We choose as the initial condition the probability distribution:

$$\rho_0 = [1/2 \ 0 \ 0 \ 1/2 \ 0 \ 0]^\dagger, \quad (74)$$

with initial equal probabilities over the DTMC states corresponding to q_0 as the DFA start state and zero otherwise. We assume that the underlying entropy production is minimized for the uniform prior,

$$r = [1/6 \ 1/6 \ 1/6 \ 1/6 \ 1/6 \ 1/6]^\dagger, \quad (75)$$

which is transformed, after one iteration, into $r' = \mathbf{W}r$:

$$r' = [a/6 \ a/6 \ a/6 \ b/6 \ b/6 \ b/6]^\dagger, \quad (76)$$

with $a := p_{00} + p_{01}$ and $b := p_{10} + p_{11}$.

Using the above definitions we can compute $\Sigma(\tau)$ in Eq. (18), at arbitrary fixed times. Moreover, in order to evaluate thermodynamic quantities at stopping times, we embrace again the family of stopping times $\mathcal{T} = \min(\mathcal{T}_1, \tau)$ with τ the fixed time horizon and $\mathcal{T}_1 \geq 1$ the first time the DFA returns to the accept state q_0 for either $b = \{0, 1\}$.

We show numerical results in Fig. 7, where $\langle \Sigma(\mathcal{T}) \rangle$, together with the corresponding upper and lower bounds given by Eq. (46), are plotted as a function of the probability $p_{01} = 1 - p_{11}$ to obtain symbol 0 after a symbol 1, for different values of $p_{00} = 1 - p_{10}$. Again we obtain relevant bounds on the intrinsic mismatch cost at stopping times, which, interestingly, become tightest when $p_{00} = 1 - p_{01}$, i.e., when $p_{00} = p_{11}$ and $p_{10} = p_{01}$. This corresponds to the situation in which the input sequence is a Markovian process with homogeneous stationary probabilities, $p_0^{\text{st}} = p_1^{\text{st}} = 1/2$. The fact that our bounds become tight for homogeneous input sequences was also observed for the i.i.d. example (see Fig. 5) and makes us conjecture that this phenomenon may be generic to correlated input sequence, maybe also non-Markovian.

As also commented for the previous examples, however, using a stopping time of the form $\mathcal{T} = \min(\mathcal{T}_1, \tau)$ allows us to describe computation times for the DFA to reach the




FIG. 7. Numerical results for the intrinsic mismatch cost $\langle \Sigma(\mathcal{T}) \rangle$ with uniform prior for the DFA in Fig. 6(a) processing Markovian input strings, as a function of the input symbol transition probability. Here $\mathcal{T} = \min(\mathcal{T}_1, \tau)$, with \mathcal{T}_1 the first return time to any of the states $(q_0, 0)$ or $(q_0, 1)$, and $\tau = 5$ a prescribed limit time. Symbols are numerical estimates for $\langle \Sigma(\mathcal{T}) \rangle$ obtained from 10^4 numerical simulations, for two different values of the transition probability p_{00} of the input string containing two consecutive zeros (see legend). Solid lines are estimates from the numerical simulations of the quantity $D(\rho_0 \| \bar{\rho}_\tau) - \langle \delta_\tau(\mathcal{T}) \rangle$, confirming the lower bound given by Eq. (44).

accept state for the first time. That corresponds to the acceptance of only some of the multiples of three, e.g., “0” (zero), “11” (three), “1001” (nine), but not other multiples like “110” (six) or any other word that already contains an acceptable prefix. In order to explore thermodynamic costs associated to these words, we now consider more general stopping times $\mathcal{T} = \min(\mathcal{T}_n, \tau)$, where \mathcal{T}_n is the n th time the DFA returns to the accept state. Therefore, \mathcal{T}_2 is related with the acceptance of words like “110” (six) or “10010” (eighteen), while \mathcal{T}_3 corresponds to accept words like “1100” (twelve), among many others.

In Fig. 8 we plot $\langle \Sigma(\mathcal{T}) \rangle$ with $\mathcal{T} = \min(\mathcal{T}_n, \tau)$ as a function of the return time to the accept state, $n = 1, 2, 3, 4, 5$. We notice that different behaviors are obtained depending on the choice of input symbols probabilities, p_{00} and p_{01} , leading to either increasing values of the intrinsic mismatch cost or a nonmonotonic behavior. Interestingly, considering different stopping times allows us to test inequality (42) for two stopping times, which is shown in the inset of Fig. 8 for $\mathcal{T}_1 = \min(\mathcal{T}_{n-1}, \tau)$ and $\mathcal{T}_2 = \min(\mathcal{T}_n, \tau)$ as a function of $n > 1$. In particular, we observe that the mismatch cost between consecutive returning times to the accept state can be eventually negative for specific choices of parameters (probabilities p_{00} and p_{01}), that is, $\langle \Delta \Sigma(\mathcal{T}_n, \mathcal{T}_{n-1}) \rangle < 0$ for $n = 3, 4, 5$, owing to a reduction in the associated stochastic distinguishability and despite having $\langle \Sigma(\tau) \rangle \geq 0$ at fixed times.




FIG. 8. Intrinsic mismatch cost up to the second, third, fourth, and fifth accepted word. Numerical results for $\langle \Sigma(\min(\mathcal{T}_n, \tau)) \rangle$ with uniform prior, with \mathcal{T}_n the n th return time to the accept states, and $\tau = 40$ a prescribed limit time. Results are obtained for the DFA in Fig. 6(a) processing Markovian input strings, as a function of the n input symbol transition probability, for parameter values $p_{00} = 0.25$, $p_{01} = 0.4$ (open symbols) and $p_{00} = 0.25$, $p_{01} = 0.75$ (filled symbols). The lines are estimates from the numerical simulations of the quantity $D(\rho_0 \| \bar{\rho}_\tau) - \langle \delta_\tau(\mathcal{T}) \rangle$ for $p_{00} = 0.25$, $p_{01} = 0.4$ (red dashed line) and for $p_{00} = 0.25$, $p_{01} = 0.75$ (red solid line). Inset: mismatch cost between returning times $\langle \Delta \Sigma(\mathcal{T}_n, \mathcal{T}_{n-1}) \rangle := \langle \Sigma(\min(\mathcal{T}_n, \tau)) \rangle - \langle \Sigma(\min(\mathcal{T}_{n-1}, \tau)) \rangle$ (red circles) and $-\langle \delta_\tau(\min(\mathcal{T}_n, \tau)) \rangle - \langle \delta_\tau(\min(\mathcal{T}_{n-1}, \tau)) \rangle$ (red solid line) as a function of n ; see Eq. (42). The horizontal dotted line is set to zero as a reference value.

VI. UNIVERSAL EQUALITIES AND INEQUALITIES FOR ACCEPTANCE PROBABILITIES

Our formalism can be further applied to address other issues in computer science theory, beyond automata literature, and besides second laws and fluctuation theorems at stopping times. Both in this section and in Sec. VII we develop further theoretical predictions for key statistical properties of interest for computer science that may inspire numerical and experimental illustrations of future work. An example that we develop in this section is using our formalism to establish universal equalities and inequalities concerning the probabilities of acceptance or rejection of sets of distinct bit sequences when a given DFA is implemented. In what follows we focus on a specific choice of such sets, namely, (1) the set of all strings or trajectories that end in an accept state before the limit time τ versus (2) the set of all trajectories that do not end in an accept state before the limit time τ . However, we emphasize that this formalism can be generalized to arbitrary pairs of sets of trajectories, by specifying suitable *filtrations* as done in martingale approaches.

Thus we will explore a class of simple examples of “acceptance” statistics for binary words of length $\tau \geq 2$ that are processed by a computer. We will use the notation *accept* to signify that a computer reaches a prescribed accept state before the limit time τ , and *reject* otherwise. The probabilities $P_a(\tau)$ denote the probability for the computer to have reached the accept state within $[0, \tau]$, and $P_r(\tau) = 1 - P_a(\tau)$ the probability for the complementary. Recall that for simple computer architectures (e.g., DFAs processing i.i.d. binary strings), $P_a(\tau)$ and $P_r(\tau)$ can often be evaluated analytically or with Monte Carlo simulations. The approach we reveal below is complementary to Monte Carlo approaches in such simple computations; however, we highlight its usefulness in revealing how such accept or reject statistics are related to thermodynamic quantities. Note that here those thermodynamic quantities can be used as a tool of calculation, determined completely by the computer update function and the distribution over input words. In particular, they need not correspond to any “real” thermodynamic quantities that one would measure in the laboratory. That is, our formalism provides a way to derive the relative probabilities of accepting or rejecting a string while sidestepping the conventional technical difficulties found in the traditional approaches to this issue [95,96]. On the other hand, one can also interpret the results presented below as a way for obtaining information about the intrinsic thermodynamic costs of computations by looking at the acceptance probabilities (of languages solved by machines), which might be calculated by other means, such as Monte Carlo approaches.

So we consider again a stopping time $T = \min(T_1, \tau)$, which signifies the first time that the computer reaches the accept state T_1 or the limit τ in the case that the accept state is not visited before τ . So $T < \tau$ if a word of length $\tau - 1$ is accepted by the computer. Otherwise, $T = \tau$, if the word is not accepted before τ . The probabilities that a word of length $\tau - 1$ is accepted or not are then given by

$$P_a(\tau) = P(T < \tau), \quad (77)$$

$$P_r(\tau) = 1 - P(T < \tau) = P(T = \tau), \quad (78)$$

respectively. We now make use of our formalism to derive bounds for $P_a(\tau)$ and $P_r(\tau)$ in terms of thermodynamic quantities.

Using our fluctuation theorem at stopping times with absolute irreversibility, Eq. (35), $\langle M_\tau(T) \rangle = 1 - \Gamma_\tau$, we expand its lhs into terms corresponding to accepted and rejected words as

$$P_a(\tau) \langle M_\tau(T) | T < \tau \rangle + P_r(\tau) \langle M_\tau(T) | T = \tau \rangle, \quad (79)$$

with $\langle A(T) | c(T) \rangle = \mathbb{E}[A(T) | c(T)]$ being the conditional average of functional A over trajectories $x_{[0, T]}$ given that the condition $c(T)$ is fulfilled over the stopping time T . Upon

using $P_r(\tau) = 1 - P_a(\tau)$, the decomposition Eq. (79) gives us the following relation between the acceptance probability and the averages of the supermartingale $M_\tau(T)$ at stopping times:

$$P_a(\tau) = \frac{1 - \Gamma_\tau - \langle M_\tau(T) | T = \tau \rangle}{\langle M_\tau(T) | T < \tau \rangle - \langle M_\tau(T) | T = \tau \rangle}. \quad (80)$$

Equality (80) generalizes analytical expressions obtained in previous works for absorption probabilities [13,19,97] by including the absolute irreversibility contribution Γ_τ . As can be appreciated in Eq. (80), since $\Gamma_\tau \geq 0$, the role of absolute irreversibility is to decrease the acceptance probability $P_a(\tau)$ of a word of length $\tau - 1$ by the DFA. This can be intuitively understood from the fact that starting computation from a restricted set of initial states can only decrease the velocity at which the computational state space is explored, and hence the probability to reach a generic stopping condition before time τ .

Since $P_a(\tau)$ is a well-defined probability [i.e., $0 \leq P_a(\tau) \leq 1$], we further obtain from Eq. (80) that one of the two following chain inequalities holds:

$$\langle M_\tau(T) | T < \tau \rangle \geq 1 - \Gamma_\tau \geq \langle M_\tau(T) | T = \tau \rangle, \quad (81)$$

$$\langle M_\tau(T) | T = \tau \rangle \geq \langle M_\tau(T) | T < \tau \rangle \geq 1 - \Gamma_\tau \geq 0, \quad (82)$$

which provide us constraints on the values of $M_\tau(T)$ for generic T of the form $T = \min(T_1, \tau)$.

Analogously, we can exploit the second-law inequality at stopping times Eq. (44), namely, $\langle \Sigma(T) \rangle \geq -\langle \delta_\tau(T) \rangle + D(\rho_0 \| \tilde{\rho}_\tau)$, to derive universal bounds for the finite-time acceptance probability. Indeed, the average of the left-hand side of this equation at the stopping time $T = \min(T_1, \tau)$ can also be decomposed into two terms, accounting, respectively, for accepted and rejected words of maximum length $\tau - 1$:

$$\langle \Sigma(T) + \delta_\tau(T) \rangle = P_a(\tau) C_a(\tau) + P_r(\tau) C_r(\tau), \quad (83)$$

where we have introduced the conditional averages,

$$C_a(\tau) = \langle \Sigma(T) + \delta_\tau(T) | T < \tau \rangle, \quad (84)$$

$$C_r(\tau) = \langle \Sigma(T) | T = \tau \rangle. \quad (85)$$

Note that in Eq. (85) we have used the fact that $\delta_\tau(\tau) = 0$. We refer to these two conditional averages as the average thermodynamic *costs* associated with the acceptance and rejection of words of length $\tau - 1$, respectively.

Combining Eqs. (44) and (83) we obtain the two following lower and upper bounds for the acceptance probability,

$$P_a(\tau) \geq \frac{D(\rho_0 \| \tilde{\rho}_\tau) - C_r(\tau)}{C_a(\tau) - C_r(\tau)}, \quad (86)$$

valid whenever $C_a(\tau) > C_r(\tau)$, and similarly

$$P_a(\tau) \leq \frac{C_r(\tau) - D(\rho_0 \| \bar{\rho}_\tau)}{C_r(\tau) - C_a(\tau)}, \quad (87)$$

valid in the complementary case when $C_r(\tau) > C_a(\tau)$. These bounds express a constraint on the acceptance probability of a word with maximum length $\tau - 1$ in terms of the average costs associated with the accepted and rejected words as defined in Eqs. (84) and (85), and the KL divergence between the initial distribution of the computational state and the final distribution of the computational state under the auxiliary dynamics [see Eq. (45)].

Equation (86) provides a meaningful bound whenever its rhs is non-negative and smaller than one, i.e., when $C_a(\tau) \geq D(\rho_0 \| \bar{\rho}_\tau) \geq C_r(\tau)$. On the other hand, the bound Eq. (87) is meaningful when $C_a(\tau) \leq D(\rho_0 \| \bar{\rho}_\tau) \leq C_r(\tau)$. We expect the first condition to be satisfied if the probability of accepted words is large enough so that the associated cost $C_a(\tau)$ is larger than the cost of rejected words $C_r(\tau)$. So we expect the bound Eq. (86) to be helpful for parameter values of the DFA and distribution over input words in which the acceptance rate is high. On the contrary, when the probability of rejected words is large enough, we expect $C_r(\tau)$ to be larger than $C_a(\tau)$, and the bound Eq. (87) to be useful when the acceptance rate is low.

The above relations in Eqs. (80), (86), and (87) concerning the acceptance probability of a word can also be applied to any finite-horizon stopping time of the form $T = \min(\mathcal{T}_c, \tau)$, where \mathcal{T}_c represents the time at which a given arbitrary condition c is verified for the first time, e.g., the first time the accept state is reached twice, or the first time the accept state is reached after passing through any other arbitrary state (or sequence of them). Thus there is an ample flexibility in choosing the stopping condition \mathcal{T} , including the logical composition of any other set of conditions; e.g., $c = c_1 \cup c_2$ giving the first time either condition c_1 or condition c_2 is verified, or $c = c_1 \cap c_2$ for the first time both c_1 and c_2 are simultaneously verified.

VII. CONCATENATING RUNS OF A DFA WITH STOCHASTIC RESETTING

In this section we further elaborate on how our results would be applied to sequences of computations separated by a *reset* of the dynamics which implements concatenated computational rounds. This is an interesting avenue where our results might be fruitfully combined in the future with the powerful analytical tools from the framework of *stochastic resetting* [35,98–102]. Let us consider a random sequence of symbols fed into a computer,

$$0\ 0\ 0\ \sqcup\ 0\ 1\ 0\ 1\ 1\ 1\ \sqcup\ 0\ 1\ 0\ \dots, \quad (88)$$

where \sqcup is a blank symbol that flags the beginning of a new computation. For the example sequence Eq. (88), a computation starts at the random starting time $\mathcal{T}_{\text{start}} = 5$

and ends at the stochastic ending time $\mathcal{T}_{\text{end}} = 10$ just before the next blank symbol arrives, thus generating the input word “010111.” During this computation, the computer begins computing at $\mathcal{T}_{\text{start}} = 4$ from its start state, and ends the computation either in an accept state or in another logical state.

Now, stochastic starting times can be reformulated as stochastic stopping times [see also our results concerning multiple stopping times, Eq. (41)]. In particular, here the starting time $\mathcal{T}_{\text{start}}$ is the first appearance of a blank symbol \sqcup . Whenever the probability of a blank symbol $p_\sqcup > 0$ is greater than zero, then it is guaranteed that $P(\mathcal{T}_{\text{end}} < \infty) = 1$; i.e., there is a limit time τ that is a finite global upper limit to \mathcal{T}_{end} . This is the setting which would correspond to, e.g., stochastic starting times that are drawn from distributions with bounded support, say from Bernoulli or binomial distributions. Under such mild assumptions, it is then possible to establish thermodynamic constraints for computations starting at stochastic times.

Supposing $p_\sqcup > 0$, we outline how the stopping-time fluctuation relations derived in our work can be applied to a computation of the example sequence Eq. (88). First, we let the computer process the sequence “000,” which implies visiting the accept state at least once. At time $t = 3$, the computer may or may not be in the start state depending on its update rules. Next, at $t = 4$, the state of the computer is reset to its start state from whichever state x_3 it occupies at the previous time instance. Upon this, the computer processes the string 010111 before the arrival of the next blank symbol, during which the logical state may or may not have reached the accept state. This leaves us with an ordered sequence of stopping times: $T_0 = 0$, $T_1 = \min(\mathcal{T}_{\text{accept}}^{(1)}, \mathcal{T}_{\text{blank}}^{(1)})$, $T_2 = \mathcal{T}_{\text{blank}}^{(1)}$, $T_3 = \min(\mathcal{T}_{\text{accept}}^{(2)}, \mathcal{T}_{\text{blank}}^{(2)})$, $T_4 = \mathcal{T}_{\text{blank}}^{(2)}$, ..., $T_\infty = \tau$, which obey

$$T_0 \leq T_1 \leq T_2 \leq T_3 \leq T_4 \leq \dots \leq T_\infty. \quad (89)$$

Here above we have denoted by $\mathcal{T}_{\text{accept}}^{(i)}$ the first return time to the accept state during the computation of the i th word. Similarly, $\mathcal{T}_{\text{blank}}^{(i)}$ is the stochastic arrival time of the i th blank symbol. While the stochastic times $\mathcal{T}_{\text{accept}}^{(i)}$ have the same structure as the stopping times considered throughout our work, the times $\mathcal{T}_{\text{blank}}^{(i+1)}$ can be seen as stochastic starting times, which are also examples of stopping times for which our formalism applies.

Figure 9 provides an illustration of a DFA processing an i.i.d. sequence of bits interspersed by blank symbols. The DFA processing the symbols recognizes binary words multiples of four, as in the examples of Sec. V. Assuming time-independent probabilities p_0 , p_1 , and p_\sqcup for the occurrence of 0, 1, and blank \sqcup symbols, respectively (with $p_0 + p_1 + p_\sqcup = 1$), the DTMC associated with this computation can be represented by a discrete-time stochastic resetting process (see Fig. 9). In such processes,




FIG. 9. An application of stochastic resetting to computer science. Top: illustration of a random tape of binary symbols interspersed by blank symbols that are processed during a computation. Bottom: discrete-time Markov chain associated with the deterministic finite automaton recognizing binary words multiple of four [see Fig. 1(c)] with stochastic resetting to the start state. Along a stochastic computation, resetting takes place when a blank symbol is recognized by the DFA. For the model illustrated here, we have assumed that words are drawn from i.i.d. sequences with probabilities p_0 , p_1 , and p_\perp of 0,1 and blank symbols, respectively, with $p_0 + p_1 + p_\perp = 1$; however, more complex scenarios could be envisaged in future work.

resetting takes place from each logical state to the start state q_0 at a stochastic starting time. This requires a suitable description of computation whose transition matrices include resetting events. For the DFA example considered here, such transition matrix takes the form

$$\mathbf{W} = \begin{bmatrix} p_0 + p_\perp & p_\perp & p_0 + p_\perp & p_\perp \\ p_1 & 0 & p_1 & 0 \\ 0 & p_0 & 0 & p_0 \\ 0 & p_1 & 0 & p_1 \end{bmatrix}, \quad (90)$$

cf. Eq. (54) for the case where no resetting takes place, corresponding to $p_\perp = 0$. The DTMC described by the transition matrix Eq. (90) allows one to study multiple realistic computational scenarios where Σ at stochastic starting and stopping times can be efficiently tackled. For example, one may consider that the processing of the input string by the DFA is a nonequilibrium stationary process with resetting and apply results from the martingale theory for stationary processes (see Chap. 7 in Ref. [19]). Alternatively, one can apply the formalism in this work to establish bounds for the intrinsic mismatch costs of the computation between the first and the n th arrival of a blank symbol, from the n th blank symbol to the $(n+1)$ th one, etc. [see Eq. (42)], similarly to Sec. Vc.

VIII. DISCUSSION

In this work we have shown how to extend stochastic thermodynamics to describe the minimal costs associated

with a computer processing with a stochastic halting time, processing strings of arbitrary length. Our formalism applies to computations described by discrete-time Markov chains over a set of computational states that may have restricted initial conditions, unidirectional links, and start and/or stop at a stochastic time. We obtain quantifiers, which are collectively dubbed as the intrinsic mismatch cost of a computation, that lower bound the entropy production incurred by the computer and that can be formulated at the fluctuating level. A key insight here is that these quantifiers, which provide a tool to probe the entropy production associated with computations at stopping times, can be entirely obtained from the DTMC evolution and the prior, without further details about their physical implementation. Note that such an intrinsic cost is independent of the internal energy of the computational states, $x_t \in \mathcal{X}$, which can indeed be assumed to be equal for every computational state and constant over time. Still, nonzero entropy production through the irreversible dissipation of heat into the environment will be in general incurred for any physical computer which implements a given computation over such set of states.

Putting forward the modern martingale formalism of stochastic thermodynamics, we also unveiled a plethora of universal fluctuation relations and inequalities that are valid for the broad class of computations analyzed in this work. We obtained a main fluctuation theorem, Eq. (35), valid for settings which include arbitrary stopping times, unidirectional transitions, and absolute irreversibility. In doing that, we have extended the martingale theory for stochastic thermodynamics to account for this additional source of irreversibility in generic situations, which we expect to have broad applicability in nonequilibrium thermodynamics.

The rigor and flexibility of our theory for stopping times allowed us to formulate and interpret several second-law-like inequalities [Eqs. (39)–(44)] at stochastic stopping times, as well as relations for the probabilities of acceptance or rejection of input data by a computer in terms of thermodynamic quantities [equality (80) and inequalities (86) and (87)]. In particular, the second-law inequalities (39) and (44) provides us useful lower bounds on the minimum dissipation incurred by a generic computation stopping at an arbitrary stopping time, while Eq. (43) establishes formally how stopping times can be used to reduce the thermodynamic costs of a computation by means of time-reversal symmetry breaking. Moreover, we have also shown the relevance of accounting for absolutely irreversible sequences in providing accurate bounds for the intrinsic mismatch cost of the computation. In this sense, the bound we derived in Eq. (39) with the absolute irreversibility term Γ_τ is tighter, with respect to the alternative bound in Eq. (44). However, computing Γ_τ might be challenging depending on the setting considered, especially for large limit times τ . On the contrary, the alternative bound in Eq. (44) would be much easier to compute (as it depends only on two

probability distributions), while still providing a meaningful bound in all examples explored here.

The framework developed in this paper can be readily applied for assessing thermodynamic costs of computations in a broad range of models of relevance in computer science theory, including—but not being limited to—deterministic finite automata. Our results apply to every computation implemented by a synchronous digital computer and remains valid independently of how the computational variables are defined. In particular, they can include already processed input symbols (as in non-iid DFAs), stacks (as in pushdown automata models), or even entire words written on a random access tape (as in Turing machines). Hence our results provide a tool to classify abstract computational machines by their intrinsic (unavoidable) thermodynamic costs. Applying our framework to more complex models of computational machines such as pushdown automata or Turing machines halting at stochastic times is a natural step following the investigation initiated here.

Our results are amenable of experimental testing using state-of-the-art techniques in line with previous tests of Landauer’s principle [25,103,104] and other experimental platforms in stochastic thermodynamics [105] in setups ranging from colloidal particles [106] and nanoscale devices [107,108] to biopolymers [91]. Regarding the determination of the prior, in some cases the experimentalist will have designed the system in a sufficiently detailed manner such that it is possible to calculate the prior, or at least approximate it with reliable numerical estimates. In other cases, the experimentalist can estimate the prior by repeatedly running the system and observing the resultant behavior. In any case, our results have nonzero lower bounds that apply no matter what the prior is.

It would also be very interesting in the future to extend the framework developed here by combining analytical tools from stochastic resetting (e.g., renewal theory and first-passage-time ideas [109]) with computer science methods. This will allow us to obtain tight bounds for the statistics of starting time and entropy production bounds in specific models of DFAs and TMs processing regular languages, as follows from the ideas sketched in Sec. VII. Also, we note that even if current digital devices are *very close* to periodic, they are not exactly so. In other words, they are some first-order perturbation away from being periodic, which suggests other avenues for future work. In general, the prior is a function of the physical process implementing the computation. For example, it is a function of the time-dependent rate matrix in the case of a CTMC. Given this, we might be able to use the envelope theorem (often used in game theory) to calculate how much the prior can change under first-order perturbations away from an exactly periodic process. That in turn might allow us to modify Eq. (13) to involve some infinitesimal first-order perturbation parameter ϵ characterizing how much the process differs from being exactly periodic.

However, the results developed in this paper also provide new insights in the field of nonequilibrium thermodynamics. An important consequence of our work is the finding that the auxiliary dynamics introduced in Eq. (15) is suitable to treat processes at stopping times that may have unidirectional transitions and absolute irreversibility, hence making our framework applicable to generic situations where local detail balance is broken. Similar auxiliary dynamics has been invoked in the literature such as so-called “dual,” “dual-reversed, or “adjoint” dynamics, in the context of fluctuation theorems; see, e.g., Refs. [20,41,52,82,110,111]. In particular, as shown above, if the process admits a well-behaved stationary solution, we can obtain from Σ the so-called nonadiabatic (or excess) entropy production [52,72]. Within such scenario, our work is another brick in the wall of recent progress highlighting the role of nonadiabatic entropy and excess heat in characterizing the efficiency [112] and calorimetry [113,114] of active nonequilibrium systems.

We also expect our results to have potential applications outside statistical mechanics and computer science, e.g., in biological physics, for instance, within the field of biomolecular computation [115], enzyme kinetics [116], and information processing in biology [117]. As a minimal model, consider a minimal Michaelis-Menten scheme for enzyme kinetics in which an enzyme E transforms a substrate molecule S into a product molecule P . A typical assumption is that the conversion of the substrate into a product takes place through an irreversible chemical

reaction, $E + S \xrightleftharpoons[k_{-1}]{k_{+1}} ES \xrightarrow{k_2} E + P$, where k_i ’s here are suitable transition rates. Within this model, the enzyme’s state during enzymatic cycles follows a continuous-time Markov jump process with one irreversible transition. In previous works, the presence of the irreversible transition has been circumvented by considering virtual processes with a very slow transition rate. However, our formalism can be readily applied to describe the stochastic thermodynamics of such enzymatic reaction, inasmuch as that it does not require the presence of bidirectional transitions. Similarly, we expect our formalism to be suitable to describe fluctuations of biological populations in processes that include totally irreversible transitions such as cell-fate decisions [118,119], cell death, and apoptosis [120], among others. For such systems, our approach puts forward recent approaches to estimate dissipation developed within in the field of active matter [121–123] which did not contemplate the presence of unidirectional transitions which are commonplace in biophysical modeling [124,125].

ACKNOWLEDGMENTS

We thank Artemy Kolchinsky for useful comments on the first version of the manuscript. G.M. acknowledges “Ramón y Cajal” program (No. RYC2021-031121-I), the

CoQuSy project (No. PID2022-140506NB-C21), and support from the María de Maeztu project (No. CEX2021-001164-M) funded by MCIU/AEI/10.13039/501100011033 and European Union NextGenerationEU/PRTR. E. R. acknowledges fruitful discussions with Léa Bresque and financial support from PNRR MUR project No. PE0000023-NQSTI. D. H. W. acknowledges U.S. NSF Grant No. CHE-1648973. D. H. W. also thanks the Santa Fe Institute for support.

APPENDIX A: RESIDUAL COST

Suppose we are given a physical process which implements a single-valued function f over a space X . The *islands* of f are defined as the elements of the partition of X given by the preimages of f . Formally, if we write the image of f as $f(X)$, the islands of f are the sets $\{f^{-1}(x):x \in f(X)\}$. We write the set of islands of a function f as $L(f)$.

In this appendix we will expand the residual cost $\mathcal{R}(\varrho_0)$ arising in Eq. (11) in terms of islands, discuss some properties of the residual cost, and then justify why this cost can be ignored in our investigation in this paper. For simplicity we will restrict attention to the case where the linear term F in Eq. (8) is expected entropy flow generated in a process, so that $\mathcal{C}(\tau)$ is the EP generated in the process up to time τ . However, all of the discussion extends to other choices of F , with the obvious modifications.

Equation (11) provides entropy production as a sum of two terms, the mismatch cost and the residual cost, where $\mathcal{R}(\tau)$ corresponds to the residual cost. For any conditional distribution G , and any island $c \in L(G)$, we define the prior within that island as [126]

$$\varrho_{\min}^c \in \arg \min_{\varrho: \text{supp}(\varrho) \in \Delta_c} \mathcal{C}(\tau), \quad (\text{A1})$$

where the subscript means that ϱ is a distribution whose support is restricted to the unit simplex over the island c , Δ_c , viewed as a subset of the state space \mathcal{Y} . The associated minimum EP in that island is written as

$$\mathcal{C}_{\min}^c(\tau) := \min_{\varrho: \text{supp}(\varrho) \in \Delta_c} \mathcal{C}(\tau). \quad (\text{A2})$$

Now introduce an arbitrary distribution over islands $q(c)$ and define

$$\varrho_{\min} := \sum_{c \in L(G)} q(c) \varrho_{\min}^c. \quad (\text{A3})$$

Then the residual EP of the physical process that implements G is

$$\mathcal{R}(\tau) = \sum_{c \in L(G)} p_\tau(c) \mathcal{C}_{\min}^c(\tau), \quad (\text{A4})$$

as shown in Ref. [126]. Therefore, the EP for the process starting with distribution ϱ_0 is

$$\mathcal{C}(\tau) = D(\varrho_0 \| \varrho_{\min}) - D(G\varrho_0 \| G\varrho_{\min}) + \sum_{c \in L(G)} p_\tau(c) \mathcal{C}_{\min}^c. \quad (\text{A5})$$

Since expected EP $\mathcal{C}(\tau)$ is non-negative, by definition the residual cost of any island c , \mathcal{C}_{\min}^c , is non-negative. So the total residual cost given by an expectation over all islands, which is the residual cost for the entire thermodynamic process characterized by G , is also non-negative. Like priors, residual costs of islands in general will differ from one cost function \mathcal{C} to the next.

In general, as the iteration t of a periodic process changes, the distribution $p_t(c)$ over the islands c will change. Therefore so will the associated total expected residual cost. However, since that total residual cost is always non-negative, all the lower bounds on EP in the main text that consider only mismatch cost apply. This is true even if the residual costs of the islands are strictly positive. In particular, Eq. (14) will still be a lower bound on the EP generated in the process.

On the other hand, if we write down the formula for minimal total residual cost which is analogous to Eq. (14), minimizing over the residual costs of each island, we just get zero, by taking those costs to all equal zero. So unless we fix the physical details underlying the process, and therefore fix the residual costs of the islands to be strictly positive, our analysis of lower bounds is not changed by the existence of residual costs. This is why such quantities are ignored in the main text.

As a final comment, note that in general both the prior and residual costs will vary with τ . However, the same mismatch cost formula bounds dissipation for any such choice of τ , once one plugs in the appropriate prior. This need not be true for residual cost, in the sense that the islands might change for different choices of τ .

APPENDIX B: TRAJECTORY-LEVEL MISMATCH COST

Define $[G_t \varrho](y)$ as the distribution ϱ evolved through the (linear) dynamics G up to time t , and then evaluated at state $y \in \mathcal{Y}$. Using this, we can define a trajectory-level version of mismatch cost (and associated instance-level version) as

$$\begin{aligned} m_{\varrho_0}(Y_{[0,\tau]}) &:= \ln \left[\frac{\varrho_0(y_0)}{\varrho_{\min}(y_0)} \right] - \ln \left[\frac{G\varrho_0(y_0)}{G\varrho_{\min}(y_0)} \right] \\ &= \sum_{t=0}^{\tau-1} \ln \left(\frac{[G_t \varrho_0](y_t)}{[G_t \varrho_{\min}](y_t)} \right) \\ &\quad - \ln \left(\frac{[G_{t+1} \varrho_0](y_{t+1})}{[G_{t+1} \varrho_{\min}](y_{t+1})} \right). \end{aligned} \quad (\text{B1})$$

By construction, the expected value of $m_{\rho_0}(Y_{[0,\tau]})$ over all trajectories $Y_{[0,\tau]}$ equals the (ensemble level) mismatch cost $\mathcal{M}(\rho_0)$ given in Eq. (10). In the context of inclusive thermodynamics, such definitions allowed the derivation of fluctuation theorems for mismatch cost [49].

APPENDIX C: ENTROPY PRODUCTION BOUNDS FROM DISCRETE-TIME EVOLUTION OVER COMPUTATIONAL STATES

We focus on the case in which the cost function $\mathcal{C}(\tau)$ introduced in Eq. (8) corresponds to the entropy production; that is, we identify F as the entropy flow due to heat exchange of the system with (one or several) heat baths, and under the eventual presence of nonconservative forces [3]. In the case in which the continuous-time periodic process is also Markovian, that is, it can be described with a set of rate matrices $K_{ij}^v(t)$ associated to the different reservoirs, the entropy flow F will take the form presented in Eq. (9). Otherwise, the explicit expression of the entropy flow F might be more involved [127–129].

In any case, the average entropy production in the continuous-time physical process can be written in terms of a Kullback-Leibler divergence between path probabilities of trajectories $\mathbf{y}_{[0,\tau]} = \{y_t | 0 \leq t \leq \tau\}$ in the extended state space, $y_t \in \mathcal{Y}$, and their time-reversed counterparts $\Theta\mathbf{y}_{[0,\tau]} = \{y_{\tau-t} | 0 \leq t \leq \tau\}$ as [64,130,131]

$$\mathcal{C}(\tau) = \langle S_{\text{tot}}(\tau) \rangle = D[\mathbb{P}(\mathbf{y}_{[0,\tau]}) \| \tilde{\mathbb{P}}(\Theta\mathbf{y}_{[0,\tau]})], \quad (\text{C1})$$

where \mathbb{P} and $\tilde{\mathbb{P}}$ stand for the probability measures in the forward and time-reversed driven processes, respectively.

By introducing the trajectories $\mathbf{x}_{[0,\tau]} = x_0, x_1, \dots, x_\tau$ on discrete-time, visible space dynamics with corresponding probability $P(\mathbf{x}_{[0,\tau]})$ as introduced in Eq. (7), we can bound the above Kullback-Leibler divergence as

$$\begin{aligned} \langle S_{\text{tot}}(\tau) \rangle &= D[\mathbb{P}(\mathbf{y}_{[0,\tau]}) \| \tilde{\mathbb{P}}(\Theta\mathbf{y}_{[0,\tau]})] \\ &= D[\mathbb{P}(\mathbf{y}_{[0,\tau]} | \mathbf{x}_{[0,\tau]}) \| \tilde{\mathbb{P}}(\Theta\mathbf{y}_{[0,\tau]} | \Theta\mathbf{x}_{[0,\tau]})] \\ &\quad + D[P(\mathbf{x}_{[0,\tau]}) \| \tilde{P}(\Theta\mathbf{x}_{[0,\tau]})] \\ &\geq D[P(\mathbf{x}_{[0,\tau]}) \| \tilde{P}(\Theta\mathbf{x}_{[0,\tau]})] =: \langle \check{S}_{\text{tot}}(\tau) \rangle, \end{aligned} \quad (\text{C2})$$

where we used chain rule for Kullback-Leibler divergence to write the second equality and denoted the coarse-grained path probabilities for forward and time-reversed driving protocols by $P(\mathbf{x}_{[0,\tau]}) = \int d\mathbf{y}_{[0,\tau]} \mathbb{P}(\mathbf{y}_{[0,\tau]}) \prod_{n=0}^{\tau-1} \delta(y_n - x_n)$ and $\tilde{P}(\Theta\mathbf{x}_{[0,\tau]}) = \int d\mathbf{y}_{[0,\tau]} \tilde{\mathbb{P}}(\Theta\mathbf{y}_{[0,\tau]}) \prod_{n=0}^{\tau-1} \delta(y_n - x_n)$, respectively.

The quantity $\langle \check{S}_{\text{tot}}(\tau) \rangle$ defined in the rhs of Eq. (C2) is a coarse-grained version of the average entropy production defined only over the DTMC dynamics and reduced state space \mathcal{X} verifying $\langle S_{\text{tot}}(\tau) \rangle \geq \langle \check{S}_{\text{tot}}(\tau) \rangle$. We can now decompose the coarse-grained entropy production $\langle \check{S}_{\text{tot}}(\tau) \rangle$

into the sum of mismatch and residual costs as in Eq. (11) at the DTMC level. In particular, by assuming that all non-computational degrees of freedom $y_t \notin \mathcal{X}$ are reinitialized to their starting values in every cycle of the periodic computational process, the prior probability minimizing \check{S}_{tot} in a single cycle will be the same for all cycles and can be easily obtained from marginalization of the fine-grained prior $q_{\min}(y)$ as $\mu(x) = \sum_{y \notin \mathcal{X}} q_{\min}(y)$.

As a consequence, the coarse-grained DTMC mismatch cost over a single cycle of the process reads:

$$\mathcal{M}(\rho_t) := D(\rho_t \| \mu) - D(\rho_{t+1} \| \mu') \geq 0, \quad (\text{C3})$$

with μ' the prior distribution evolved during a single iteration of the DTMC, $\mu' = \mathbf{W}\mu$. Analogously $\mathcal{R}(\rho_t) := \langle \check{S}_{\text{tot}}(\rho_t) \rangle - \mathcal{M}(\rho_t) = \min_{\rho_t} \langle \check{S}_{\text{tot}}(\rho_t) \rangle \geq 0$ represents the residual costs during a single iteration, which corresponds to the entropy production when the computational system starts the evolution in distribution μ . Hence for every single iteration $\langle \check{S}_{\text{tot}}(\rho_t) \rangle = \mathcal{M}(\rho_t) + \mathcal{R}(\rho_t)$, from which it follows that the above mismatch cost provides a lower bound to the average entropy production during a single cycle. Finally, by summing the mismatch costs per iteration Eq. (C3) over all iterations up to the final time τ , we recover Eq. (12) for the sum of mismatch costs during the entire computation (that we dub intrinsic mismatch cost), which verifies the chain inequality:

$$\sum_{t=0}^{\tau-1} \mathcal{M}(\rho_t) \leq \langle \check{S}_{\text{tot}}(\tau) \rangle \leq \langle S_{\text{tot}}(\tau) \rangle, \quad (\text{C4})$$

hence providing a lower bound on the average entropy production during the entire computation.

APPENDIX D: STRICT POSITIVITY OF THE MISMATCH COST SUM GIVEN BY Eq. (12)

Here we provide a proof of the strict positivity of the mismatch cost sum in Eq. (12); that is, $\sum_{t=0}^{\tau-1} \mathcal{M}(\rho_t) > 0$.

We know that, in general, $D(\rho_0 \| \mu) - D(G\rho_0 \| G\mu) = 0$ if and only if $\rho_0 = \mu$. However, if in fact G is not logically invertible, and yet $\rho_0 = \mu$, then $G\rho_0 \neq \mu$. This means that either $D(\rho_0 \| \mu) - D(G\rho_0 \| G\mu) = 0$ or $D(G\rho_0 \| \mu) - D(G^2\rho_0 \| G\mu) = 0$ —but not both. Therefore, so long as G is not logically invertible, the sum in Eq. (12) is not zero.

APPENDIX E: PROOF OF THE SUPERMARTINGALE PROPERTY Eq. (33)

Here we provide a detailed proof of the supermartingale property of the process $M_\tau(t)$ defined by Eq. (30):

$$\langle M_\tau(\tau) | X_{[0,i]} \rangle \equiv \sum_{X_{(t,\tau)} \in \mathcal{F}_{\text{AC}}} P(X_{[0,\tau]} | X_{[0,i]}) M_\tau(\tau) \quad (\text{E1})$$

$$\begin{aligned}
&= \sum_{X_{[t+1,\tau]} \in \mathcal{F}_{AC}} \frac{\bar{P}(\Theta X_{[0,\tau]})}{P(X_{[0,t]})} \quad (\text{E2}) \\
&= e^{-\Sigma(t)} \sum_{X_{[t+1,\tau]} \in \mathcal{F}_{AC}} \frac{\bar{P}(\Theta X_{[0,\tau]})}{\bar{P}(\Theta X_{[0,t]})} \\
&= e^{-\Sigma(t) - \delta_\tau(t)} \sum_{X_{[t+1,\tau]} \in \mathcal{F}_{AC}} \frac{\bar{P}(\Theta X_{[t,\tau]})}{\bar{\rho}_{\tau-t}(x_t)} \\
&= M_\tau(t) \left[1 - \sum_{X_{[t+1,\tau]} \in \mathcal{F}_{AI}} \frac{\bar{P}(\Theta X_{[t,\tau]})}{\bar{\rho}_{\tau-t}(x_t)} \right] \\
&= M_\tau(t) [1 - \alpha_\tau(t)] \leq M_\tau(t). \quad (\text{E3})
\end{aligned}$$

To obtain Eq. (E3) we first used the definition of conditional probability in Eq. (E1) to derive Eq. (E2). We then multiplied and divided by $\bar{P}(\Theta X_{[0,t]})$, in order to obtain the $e^{-\Sigma(t)}$ multiplicative factor. After that we expanded the remaining probabilities $\bar{P}(\Theta X_{[0,\tau]})$ in the numerator and $\bar{P}(\Theta X_{[0,t]})$ in the denominator, and multiplied the resulting expression with $\rho_t(x_t)/\bar{\rho}_{\tau-t}(x_t)$ to obtain the expression involving $M_\tau(t)$ [recall the definition of $M_\tau(t)$ in Eq. (29)]. Finally, we transformed the sum to involve \mathcal{F}_{AI} , the complementary filtration of \mathcal{F}_{AC} , and invoked the fact that $\sum_{X_{[t,\tau]} \in \mathcal{F}} \rho_\tau(x_\tau) \times \bar{P}(x_{\tau-1}|x_\tau) \dots \bar{P}(x_t|x_{t+1}) = \bar{\rho}_{\tau-t}(x_t)$ when the sum is taken over the complete set of trajectories \mathcal{F} .

APPENDIX F: STOPPING-TIMES FLUCTUATION THEOREM WITH ABSOLUTE IRREVERSIBILITY IN Eq. (35)

In this appendix we give a detailed proof of the main stopping-times fluctuation theorem with absolute irreversibility presented in Eq. (35). For the proof, it is convenient to split the filtration \mathcal{F} into subsets of filtrations $\mathcal{F}^{(t)}$ containing all trajectories that are stopped at time t , i.e., for which $M_\tau(t) = M_\tau(\mathcal{T})$. We take $t = 0, 1, \dots, \tau$, where τ is the maximum allowed time. Note that we enforce the dynamics to stop at τ if not previously done [however, one can later take $\tau \rightarrow \infty$ whenever $M_\tau(\tau)$ remains bounded]. Therefore, we have $\mathcal{F} = \mathcal{F}^{(1)} \cup \dots \cup \mathcal{F}^{(\tau)}$. Now, in analogy to the previous appendix, we define for each (discrete) instant of time t the sets $\mathcal{F}_{AC}^{(t)}$ and $\mathcal{F}_{AI}^{(t)}$ of trajectories that are both stopped at t and which are either allowed [$P(X_{[0,t]}) > 0$ or $P(X_{[0,t]}) = P(\Theta X_{[0,t]}) = 0$] or forbidden only in the original dynamics [$P(X_{[0,t]}) = 0$ with $\bar{P}(\Theta X_{[0,t]}) > 0$], respectively. This implies that $\mathcal{F}^{(t)} = \mathcal{F}_{AC}^{(t)} \cup \mathcal{F}_{AI}^{(t)}$ at any $t = 0, 1, \dots, \tau$. Then we have

$$\begin{aligned}
\langle M_\tau(\mathcal{T}) \rangle &\equiv \sum_{t=0}^{\tau} \sum_{X_{[0,t]} \in \mathcal{F}_{AC}^{(t)}} P(X_{[0,t]}) M_\tau(t) \\
&= \sum_{t=0}^{\tau} \sum_{X_{[0,t]} \in \mathcal{F}_{AC}^{(t)}} P(X_{[0,t]}) \\
&\quad \times [\langle M_\tau(\tau) | X_{[0,t]} \rangle + M_\tau(t) \alpha_\tau(t)] \\
&= \langle M_\tau(\tau) \rangle + \sum_{t=0}^{\tau} \sum_{X_{[0,t]} \in \mathcal{F}_{AC}^{(t)}} P(X_{[0,t]}) M_\tau(t) \alpha_\tau(t) \\
&= 1 - \gamma_\tau + \sum_{t=0}^{\tau} \sum_{X_{[0,t]} \in \mathcal{F}_{AC}^{(t)}} P(X_{[0,t]}) M_\tau(t) \alpha_\tau(t) \quad (\text{F1})
\end{aligned}$$

$$= 1 - \sum_{t=0}^{\tau} \sum_{X_{[0,t]} \in \mathcal{F}_{AI}^{(t)}} \bar{P}(\Theta X_{[0,t]}) = 1 - \Gamma_\tau, \quad (\text{F2})$$

where we used the fluctuation theorem at fixed times with absolute irreversibility, cf. Eq. (24), and in the last line we defined the correction term for stopping times:

$$\Gamma_\tau \equiv \sum_{t=0}^{\tau} \sum_{X_{[0,t]} \in \mathcal{F}_{AI}^{(t)}} \bar{P}(\Theta X_{[0,t]}) \frac{\bar{\rho}_{\tau-t}(X_t)}{\rho_t(X_t)} \geq 0, \quad (\text{F3})$$

summing up the probabilities for all the (reversed) stopped AI trajectories.

To reach Eq. (F2) from Eq. (F1) we used

$$\begin{aligned}
&\sum_{t=0}^{\tau} \sum_{X_{[0,t]} \in \mathcal{F}_{AC}^{(t)}} P(X_{[0,t]}) M_\tau(t) \alpha_\tau(t) \\
&= \sum_{t=0}^{\tau} \sum_{X_{[0,t]} \in \mathcal{F}_{AC}^{(t)}} \frac{\bar{P}(\Theta X_{[0,t]})}{\rho_t(X_t)} \sum_{X_{[t,\tau]} \in \mathcal{F}_{AI}} \bar{P}(\Theta X_{[t,\tau]}) \\
&= \sum_{t=0}^{\tau} \sum_{X_{[0,t]} \in \mathcal{F}_{AC}^{(t)}} \frac{\bar{P}(\Theta X_{[0,t]})}{\rho_t(X_t)} \left[\bar{\rho}_{\tau-t}(X_t) - \sum_{X_{[t,\tau]} \in \mathcal{F}_{AC}} \bar{P}(\Theta X_{[t,\tau]}) \right] \\
&= \sum_{t=0}^{\tau} \sum_{X_{[0,t]} \in \mathcal{F}_{AC}^{(t)}} \bar{P}(\Theta X_{[0,t]}) \frac{\bar{\rho}_{\tau-t}(X_t)}{\rho_t(X_t)} - \sum_{X_{[0,\tau]} \in \mathcal{F}_{AC}} \bar{P}(\Theta X_{[0,\tau]}) \\
&= \sum_{t=0}^{\tau} \sum_{X_{[0,t]} \in \mathcal{F}_{AC}^{(t)}} \bar{P}(\Theta X_{[0,t]}) \frac{\bar{\rho}_{\tau-t}(X_t)}{\rho_t(X_t)} - 1 + \gamma_\tau \\
&= - \sum_{t=0}^{\tau} \sum_{X_{[0,t]} \in \mathcal{F}_{AI}^{(t)}} \bar{P}(\Theta X_{[0,t]}) \frac{\bar{\rho}_{\tau-t}(X_t)}{\rho_t(X_t)} + \gamma_\tau \\
&= -\Gamma_\tau + \gamma_\tau, \quad (\text{F4})
\end{aligned}$$

where in the last line we used that $\mathcal{F}_{AC}^{(t)} \cup \mathcal{F}_{AI}^{(t)} = \mathcal{F}^{(t)}$ and $\sum_{i=0}^{\tau} \sum_{X_{[0,i]} \in \mathcal{F}^{(t)}} \bar{P}(\Theta X_{[0,t]}) \bar{\rho}_{\tau-t}(X_t) / \rho_t(X_t) = 1$, since the trajectories are stopped with certainty in the interval $[0, \tau]$ and $\bar{P}(\Theta X_{[0,t]}) \bar{\rho}_{\tau-t}(X_t) / \rho_t(X_t)$ is a normalized path probability.

APPENDIX G: ORDERED STOPPING-TIMES FLUCTUATION RELATION IN Eq. (37)

In this appendix we derive the stopping-times fluctuation relation for two ordered stopping times in Eq. (37) via Doob's optional sampling theorem. Let us start by considering a finite but otherwise generic stopping time \mathcal{T} of the form in Eq. (28), such that $0 \leq \mathcal{T} \leq \tau$. Doob's optional sampling theorem for supermartingales reads [88] (see also Ref. [19]):

$$\langle A(\mathcal{T}) | \mathbf{x}_{[0,t]} \rangle \leq A(\min\{\mathcal{T}, t\}), \quad (\text{G1})$$

where A is a supermartingale over the trajectories $\mathbf{x}_{[0,\tau]}$.

In the following we will apply Eq. (G1) to the supermartingale process $M_\tau(t) = e^{-\Sigma(t) - \delta_\tau(t)}$ introduced in Eq. (30), for two ordered stopping times \mathcal{T}_1 and \mathcal{T}_2 such that $P(\mathcal{T}_2 \geq \mathcal{T}_1) = 1$. In this case we have

$$\langle M_\tau(\mathcal{T}_2) | \mathbf{x}_{[0,\mathcal{T}_1]} \rangle \leq M_\tau(\mathcal{T}_1), \quad (\text{G2})$$

where we called $\mathcal{T}_2 = \mathcal{T}$ and $\mathcal{T}_1 = \min\{\mathcal{T}, t\}$. Note that above $\mathbf{x}_{[0,t]}$ can be replaced by $\mathbf{x}_{[0,\mathcal{T}_1]}$ since $\mathcal{T}_1 \leq t$ by construction and Eq. (G1) is valid for any generic $t \leq \tau$. The average over stopping times of $M_\tau(\mathcal{T}_1)$ then verifies:

$$\begin{aligned} \langle M_\tau(\mathcal{T}_1) \rangle &= \sum_{\mathcal{T}_1=0}^{\tau} \sum_{\mathbf{x}_{[0,\mathcal{T}_1]}} P(\mathbf{x}_{[0,\mathcal{T}_1]}, \mathcal{T}_1) M_\tau(\mathcal{T}_1) \\ &= \sum_{\mathcal{T}_1=0}^{\mathcal{T}_2} \sum_{\mathbf{x}_{[0,\mathcal{T}_1]}} P(\mathbf{x}_{[0,\mathcal{T}_1]}, \mathcal{T}_1) M_\tau(\mathcal{T}_1) \\ &\quad + \sum_{\mathcal{T}_1=\mathcal{T}_2}^{\tau} \sum_{\mathbf{x}_{[0,\mathcal{T}_1]}} P(\mathbf{x}_{[0,\mathcal{T}_1]}, \mathcal{T}_1) M_\tau(\mathcal{T}_1) \\ &\geq \sum_{\mathcal{T}_1=0}^{\mathcal{T}_2} \sum_{\mathbf{x}_{[0,\mathcal{T}_1]}} P(\mathbf{x}_{[0,\mathcal{T}_1]}, \mathcal{T}_1) \langle M_\tau(\mathcal{T}_2) | \mathbf{x}_{[0,\mathcal{T}_1]} \rangle \\ &= \langle M_\tau(\mathcal{T}_2) \rangle, \end{aligned} \quad (\text{G3})$$

where we used that $\mathcal{T}_2 \geq \mathcal{T}_1$ and Eq. (G2) to reach the inequality. The last line follows from the fact that we are summing the conditional average over all possible trajectories for stopping times \mathcal{T}_1 between 0 and \mathcal{T}_2 . Finally, by replacing in Eq. (G3) the explicit expression for $M_\tau(t)$ in Eq. (30), we directly recover Eq. (37). The inequality in Eq. (G3) is saturated either for $\mathcal{T}_2 = \mathcal{T}_1$ or when the

supermartingale $M_\tau(t)$ becomes a martingale, that is, in the absence of absolute irreversibility.

APPENDIX H: SCALING OF THERMODYNAMIC AVERAGES WITH LIMIT TIME IN THE DFA EXAMPLE

As pointed in Sec. VB for the minimal DFA in Fig. 2(a), we observe a tendency for the intrinsic mismatch cost at stopping times $\langle \Sigma(\mathcal{T}) \rangle$ to saturate when increasing the limit time τ . This is in stark contrast with the scaling behavior of the fixed-time average $\langle \Sigma(\tau) \rangle$ with τ .

In this appendix we provide extra numerical evidence for the scaling behavior of mismatch cost at stopping and fixed times, which is shown in Fig. 10. There we observe that indeed the average intrinsic mismatch cost at fixed times scales linearly with τ , that is, $\langle \Sigma(\tau) \rangle \sim \tau$, as illustrated in the inset of Fig. 10. Moreover, we obtain that $\langle \Sigma(\tau) \rangle / \tau$ decreases monotonically with τ up to a saturating positive value, yet its scaling behavior is rather insensitive to the statistics of the input strings (see open symbols in Fig. 10). This point makes us question whether the fixed-time average $\langle \Sigma(\tau) \rangle$, or its rate per iteration $\langle \Sigma(\tau) \rangle / \tau$, is a suitable indicator of the thermodynamic costs of the computation. On the other hand, when considering the stopping-times average $\langle \Sigma(\mathcal{T}) \rangle$, we observe a sub-linear scaling at moderate values of the limit time, i.e., $\langle \Sigma(\mathcal{T}) \rangle \sim \tau^\alpha$, with $|\alpha| < 1$, reaching a plateau at large τ (see filled symbols in Fig. 10).




FIG. 10. Numerical results for the intrinsic mismatch cost with uniform prior as a function of the limit time τ . We show three different values of the probability of symbol 0: $p_0 = 0.75$ (blue squares), $p_0 = 0.5$ (red circles), and $p_0 = 0.3$ (green diamonds). Filled symbols correspond to the stopping-time average $\langle \Sigma(\mathcal{T}) \rangle$. We include for comparison the corresponding fixed-time ensemble average rate $\langle \Sigma(\tau) \rangle / \tau$ without absorbing conditions (open symbols). The open symbols in the inset show the value of $\langle \Sigma(\tau) \rangle$. Parameters of the simulations are as in Fig. 5.

We also find that $\langle \Sigma(T) \rangle$ is more sensitive than $\langle \Sigma(\tau) \rangle / \tau$ to the value of p_0 for all values of τ explored in our simulations. This reveals that the intrinsic mismatch cost at stopping times $\langle \Sigma(T) \rangle$ is a suitable quantity to quantify the average performance of a computation accomplished at a stochastic time. For example, the sensitivity of $\langle \Sigma(T) \rangle$ to string statistics could be fruitfully exploited as a probe of the performance of a DFA in processing different regular languages, an exciting avenue that we leave for future work.

-
- [1] Ken Sekimoto, *Stochastic Energetics*, Vol. 799 (Springer, New York, 2010).
 - [2] Christopher Jarzynski, *Equalities and inequalities: Irreversibility and the second law of thermodynamics at the nanoscale*, *Annu. Rev. Condens. Matter Phys.* **2**, 329 (2011).
 - [3] Udo Seifert, *Stochastic thermodynamics, fluctuation theorems and molecular machines*, *Rep. Prog. Phys.* **75**, 126001 (2012).
 - [4] Naoto Shiraishi, Ken Funo, and Keiji Saito, *Speed limit for classical stochastic processes*, *Phys. Rev. Lett.* **121**, 070601 (2018).
 - [5] Van Tuan Vo, Tan Van Vu, and Yoshihiko Hasegawa, *Unified approach to classical speed limit and thermodynamic uncertainty relation*, *Phys. Rev. E* **102**, 062132 (2020).
 - [6] Tan Van Vu and Keiji Saito, *Thermodynamic unification of optimal transport: Thermodynamic uncertainty relation, minimum dissipation, and thermodynamic speed limits*, *Phys. Rev. X* **13**, 011013 (2023).
 - [7] Andre C. Barato and Udo Seifert, *Thermodynamic uncertainty relation for biomolecular processes*, *Phys. Rev. Lett.* **114**, 158101 (2015).
 - [8] Todd R. Gingrich, Jordan M. Horowitz, Nikolay Perunov, and Jeremy L. England, *Dissipation bounds all steady-state fluctuations*, *Phys. Rev. Lett.* **116**, 120601 (2016).
 - [9] Jordan M. Horowitz and Todd R. Gingrich, *Thermodynamic uncertainty relations constrain non-equilibrium fluctuations*, *Nat. Phys.* **16**, 15 (2020).
 - [10] Yasuhiro Utsumi, Yasuchika Ito, Dimitry Golubev, and Ferdinand Peper, *Computation time and thermodynamic uncertainty relation of Brownian circuits*, [arXiv:2205.10735](https://arxiv.org/abs/2205.10735).
 - [11] Raphaël Chétrite and Hugo Touchette, *Nonequilibrium Markov processes conditioned on large deviations*, in *Annales Henri Poincaré*, Vol. 16 (Springer, 2015), pp. 2005–2057, [10.1007/s00023-014-0375-8](https://doi.org/10.1007/s00023-014-0375-8).
 - [12] Johannes Hoppenau, Daniel Nickelsen, and Andreas Engel, *Level 2 and level 2.5 large deviation functionals for systems with and without detailed balance*, *New J. Phys.* **18**, 083010 (2016).
 - [13] Izaak Neri, Édgar Roldán, and Frank Jülicher, *Statistics of infima and stopping times of entropy production and applications to active molecular processes*, *Phys. Rev. X* **7**, 011019 (2017).
 - [14] Raphaël Chétrite, Shamik Gupta, Izaak Neri, and Édgar Roldán, *Martingale theory for housekeeping heat*, *Europhys. Lett.* **124**, 60006 (2019).
 - [15] Gonzalo Manzano, Rosario Fazio, and Édgar Roldán, *Quantum martingale theory and entropy production*, *Phys. Rev. Lett.* **122**, 220602 (2019).
 - [16] Gonzalo Manzano, Diego Subero, Olivier Maillet, Rosario Fazio, Jukka P. Pekola, and Édgar Roldán, *Thermodynamics of gambling demons*, *Phys. Rev. Lett.* **126**, 080603 (2021).
 - [17] Gonzalo Manzano and Édgar Roldán, *Survival and extreme statistics of work, heat, and entropy production in steady-state heat engines*, *Phys. Rev. E* **105**, 024112 (2022).
 - [18] Izaak Neri and Matteo Poletini, *Extreme value statistics of edge currents in Markov jump processes and their use for entropy production estimation*, *SciPost Phys.* **14**, 131 (2023).
 - [19] Édgar Roldán, Izaak Neri, Raphael Chétrite, Shamik Gupta, Simone Pigolotti, Frank Jülicher, and Ken Sekimoto, *Martingales for physicists*, [arXiv:2210.09983](https://arxiv.org/abs/2210.09983).
 - [20] Marco Baiesi and Gianmaria Falasco, *Inflow rate, a time-symmetric observable obeying fluctuation relations*, *Phys. Rev. E* **92**, 042162 (2015).
 - [21] Ivan Di Terlizzi and Marco Baiesi, *Kinetic uncertainty relation*, *J. Phys. A* **52**, 02LT03 (2018).
 - [22] Christian Maes, *Frenesy: Time-symmetric dynamical activity in nonequilibria*, *Phys. Rep.* **850**, 1 (2020).
 - [23] Harvey Leff and Andrew F. Rex, *Maxwell's Demon 2 Entropy, Classical and Quantum Information, Computing* (CRC Press, Boca Raton, FL, 2002).
 - [24] Ryoichi Kawai, Juan M. R. Parrondo, and Christian Van den Broeck, *Dissipation: The phase-space perspective*, *Phys. Rev. Lett.* **98**, 080602 (2007).
 - [25] Karel Proesmans, Jannik Ehrich, and John Bechhoefer, *Finite-time Landauer principle*, *Phys. Rev. Lett.* **125**, 100602 (2020).
 - [26] Salambô Dago, Jorge Pereda, Nicolas Barros, Sergio Ciliberto, and Ludovic Bellon, *Information and thermodynamics: Fast and precise approach to Landauer's bound in an underdamped micromechanical oscillator*, *Phys. Rev. Lett.* **126**, 170601 (2021).
 - [27] David H. Wolpert and Artemy Kolchinsky, *Thermodynamics of computing with circuits*, *New J. Phys.* **22**, 063047 (2020); see also updated version [arXiv:1806.04103](https://arxiv.org/abs/1806.04103).
 - [28] Thomas E. Ouldridge and David H. Wolpert, *Thermodynamics of deterministic finite automata operating locally and periodically*, *New J. Phys.* **25**, 123013 (2023).
 - [29] Gülce Kardeş and David Wolpert, *Inclusive thermodynamics of computational machines*, [arXiv:2206.01165](https://arxiv.org/abs/2206.01165).
 - [30] Christopher Jarzynski, *Hamiltonian derivation of a detailed fluctuation theorem*, *J. Stat. Phys.* **98**, 77 (2000).
 - [31] Massimiliano Esposito, Katja Lindenberg, and Christian Van den Broeck, *Entropy production as correlation between system and reservoir*, *New J. Phys.* **12**, 013013 (2010).
 - [32] Sebastian Deffner and Christopher Jarzynski, *Information processing and the second law of thermodynamics: An inclusive, Hamiltonian approach*, *Phys. Rev. X* **3**, 041003 (2013).

- [33] Krzysztof Ptaszyński and Massimiliano Esposito, *Entropy production in open systems: The predominant role of intraenvironment correlations*, *Phys. Rev. Lett.* **123**, 200603 (2019).
- [34] Saar Rahav and Upendra Harbola, *An integral fluctuation theorem for systems with unidirectional transitions*, *J. Stat. Mech.* (2014) P10044.
- [35] Arnab Pal and Saar Rahav, *Integral fluctuation theorems for stochastic resetting systems*, *Phys. Rev. E* **96**, 062135 (2017).
- [36] Daniel M. Busiello, Deepak Gupta, and Amos Maritan, *Entropy production in systems with unidirectional transitions*, *Phys. Rev. Res.* **2**, 023011 (2020).
- [37] Ken Hiura and Shin-ichi Sasa, *Kinetic uncertainty relation on first-passage time for accumulated current*, *Phys. Rev. E* **103**, L050103 (2021).
- [38] Charles Moslonka and Ken Sekimoto, *Martingale-induced local invariance in progressive quenching*, *Phys. Rev. E* **105**, 044146 (2022).
- [39] Kiyoshi Kanazawa, Tomohiko G. Sano, Takahiro Sagawa, and Hisao Hayakawa, *Minimal model of stochastic athermal systems: Origin of non-Gaussian noise*, *Phys. Rev. Lett.* **114**, 090601 (2015).
- [40] Massimiliano Esposito and Juan M. R. Parrondo, *Stochastic thermodynamics of hidden pumps*, *Phys. Rev. E* **91**, 052114 (2015).
- [41] Gonzalo Manzano, Jordan M. Horowitz, and Juan M. R. Parrondo, *Quantum fluctuation theorems for arbitrary environments: Adiabatic and nonadiabatic entropy production*, *Phys. Rev. X* **8**, 031037 (2018).
- [42] Lennart Dabelow, Stefano Bo, and Ralf Eichhorn, *Irreversibility in active matter systems: Fluctuation theorem and mutual information*, *Phys. Rev. X* **9**, 021009 (2019).
- [43] Laura Tociu, Étienne Fodor, Takahiro Nemoto, and Suriyanarayanan Vaikuntanathan, *How dissipation constrains fluctuations in nonequilibrium liquids: Diffusion, structure, and biased interactions*, *Phys. Rev. X* **9**, 041026 (2019).
- [44] However, we emphasize that our theory is also applicable to the more standard scenario, such as when LDB is verified.
- [45] See Chap. 4 in Ref. [19] for rigorous mathematical definitions and mathematical properties of stopping times, and examples of stopping times in physics (e.g., first-passage times).
- [46] A version of this result even holds for quantum thermodynamic processes. For further discussion of the concept of mismatch cost in stochastic thermodynamics of computation, see Refs. [47–50].
- [47] Artemy Kolchinsky and David H. Wolpert, *Dependence of dissipation on the initial distribution over states*, *J. Stat. Mech.* (2017) 083202.
- [48] D. H. Wolpert, A. Kolchinsky, and J. A. Owen, *A space-time tradeoff for implementing a function with master equation dynamics*, *Nat. Commun.* **10**, 1727 (2019).
- [49] Artemy Kolchinsky and David H. Wolpert, *Dependence of integrated, instantaneous, and fluctuating entropy production on the initial state in quantum and classical processes*, *Phys. Rev. E* **104**, 054107 (2021).
- [50] Paul M. Riechers and Mile Gu, *Initial-state dependence of thermodynamic dissipation for any quantum process*, *Phys. Rev. E* **103**, 042145 (2021).
- [51] An analogous notion of stochastic distinguishability was introduced in Ref. [16], in the context of forward and time-reversed nonstationary dynamics.
- [52] Massimiliano Esposito and Christian Van den Broeck, *Three detailed fluctuation theorems*, *Phys. Rev. Lett.* **104**, 090601 (2010).
- [53] Takahiro Hatano and Shin-ichi Sasa, *Steady-state thermodynamics of Langevin systems*, *Phys. Rev. Lett.* **86**, 3463 (2001).
- [54] In general, in this paper, $\langle A(t)|x_{[0,s]} \rangle = \mathbb{E}[A(t)|x_{[0,s]}]$ is an average of functional A over all computational trajectories with the condition of $x_{[0,s]} = x_0, \dots, x_s$ to be fixed to a given sequence.
- [55] Here $D(\rho_0||\bar{\rho}_\tau) = \sum_x \rho_0(x) \ln[\rho_0(x)/\bar{\rho}_\tau(x)] \geq 0$ is the Kullback-Leibler (KL) quantifying the “distance” between the initial distribution of the computer’s state ρ_0 and the distribution $\bar{\rho}_\tau$ of the auxiliary computational process at the limit time τ .
- [56] I. Di Terlizzi, M. Gironella, D. Herráez-Aguilar, T. Betz, F. Monroy, M. Baiesi, and F. Ritort, *Variance sum rule for entropy production*, *Science* **383**, 971 (2024).
- [57] T. Nishiyama and Y. Hasegawa, *Upper bound for entropy production in Markov processes*, *Phys. Rev. E* **108**, 044139 (2023).
- [58] Domingos S. P. Salazar, *Upper bound for quantum entropy production from entropy flux*, *Phys. Rev. E* **105**, L042101 (2022).
- [59] Juan M. R. Parrondo, Jordan M. Horowitz, and Takahiro Sagawa, *Thermodynamics of information*, *Nat. Phys.* **11**, 131 (2015).
- [60] Takahiro Sagawa, *Thermodynamic and logical reversibilities revisited*, *J. Stat. Mech.* (2014) P03025.
- [61] Philipp Strasberg, Javier Cerrillo, Gernot Schaller, and Tobias Brandes, *Thermodynamics of stochastic Turing machines*, *Phys. Rev. E* **92**, 042104 (2015).
- [62] Artemy Kolchinsky and David H. Wolpert, *Thermodynamic costs of Turing machines*, *Phys. Rev. Res.* **2**, 033312 (2020).
- [63] Udo Seifert, *Stochastic thermodynamics, fluctuation theorems and molecular machines*, *Rep. Prog. Phys.* **75**, 126001 (2012).
- [64] R. Kawai, J. M. R. Parrondo, and C. Van den Broeck, *Dissipation: The phase-space perspective*, *Phys. Rev. Lett.* **98**, 080602 (2007).
- [65] J. A. Owen, A. Kolchinsky, and D. H. Wolpert, *Number of hidden states needed to physically implement a given conditional distribution*, *New J. Phys.* **21**, 013022 (2019).
- [66] Note that formula for F is indeed linear in q_i .
- [67] David H. Wolpert, *The free energy requirements of biological organisms; implications for evolution*, *Entropy* **18**, 219 (2016).
- [68] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. (John Wiley and Sons, New York, 2012).
- [69] Here we are interested in the marginal prior distribution over computational states $x \in \mathcal{X}$ only, not over the set \mathcal{Y} containing hidden states.

- [70] Farita Tasnim and David H. Wolpert, *Stochastic thermodynamics of multiple co-evolving systems—Beyond multipartite processes*, *Entropy* **25**, 1078 (2023).
- [71] Yoshitsugu Oono and Marco Paniconi, *Steady state thermodynamics*, *Prog. Theor. Phys. Suppl.* **130**, 29 (1998).
- [72] Massimiliano Esposito and Christian Van den Broeck, *Three faces of the second law. I. Master equation formulation*, *Phys. Rev. E* **82**, 011143 (2010).
- [73] Here we will not consider the eventual erasure of the input tape after the computation ends, but we focus on the thermodynamic costs and dissipation due to the computing process itself, using the information contained in the tape. Such eventual erasure costs are independent of the mismatch costs considered here and might be calculated on a separate basis.
- [74] Andre C. Barato and Udo Seifert, *Stochastic thermodynamics with information reservoirs*, *Phys. Rev. E* **90**, 042150 (2014).
- [75] Notice that the DTMC transition probabilities $P(x_{t+1}|x_t)$ among the (coarse-grained) computational states in \mathcal{X} are given by marginalizing a fine-grained conditional distribution $P(y_{t+1}|y_t)$, where $y = (x, s, \dots)$ over s , namely the current symbol being read, and where the update function f defining the DFA specifies $P[x_{t+1}|(x_s, s_t)]$. As a consequence, the transition probabilities $P(x_{t+1}|x_t)$ are time independent and fixed along the entire computational process.
- [76] This means that one cannot model a given DFA as some specific (and therefore fixed width) circuit, in general. The DFA will have properties that are not captured by that circuit. In this sense, individual DFAs are computationally more powerful than individual circuits.
- [77] Note that by “ $P(j|i)$ ” we do *not* mean the Bayesian inverse of the forward transition matrix $P(i|j)$; rather, we mean that forward transition matrix evaluated for a transposed choice of the initial and final states.
- [78] This can be easily checked from $\sum_j \bar{P}(i|j)r'(j) = \sum_j P(j|i)r(i) = r(i)$, which immediately implies $[\bar{\mathbf{W}}]r' = r$.
- [79] E. T. Jaynes, *Probability Theory. The Logic of Science* (Cambridge University Press, New York, USA, 2003).
- [80] Francesco Buscemi and Valerio Scarani, *Fluctuation theorems from Bayesian retrodiction*, *Phys. Rev. E* **103**, 052111 (2021).
- [81] Jacopo Surace and Matteo Scandi, *State retrieval beyond Bayes’ retrodiction*, *Quantum* **7**, 990 (2023).
- [82] Gonzalo Manzano, Jordan M. Horowitz, and Juan M. R. Parrondo, *Nonequilibrium potential and fluctuation theorems for quantum maps*, *Phys. Rev. E* **92**, 032129 (2015).
- [83] Herbert Spohn, *Entropy production for quantum dynamical semigroups*, *J. Math. Phys.* **19**, 1227 (1978).
- [84] Yūto Murashita, Ken Funo, and Masahito Ueda, *Nonequilibrium equalities in absolutely irreversible processes*, *Phys. Rev. E* **90**, 042110 (2014).
- [85] Ken Funo, Yūto Murashita, and Masahito Ueda, *Quantum nonequilibrium equalities with absolute irreversibility*, *New J. Phys.* **17**, 075005 (2015).
- [86] Y. Masuyama, K. Funo, Y. Murashita, A. Noguchi, S. Kono, Y. Tabuchi, R. Yamazaki, M. Ueda, and Y. Nakamura, *Information-to-work conversion by Maxwell’s demon in a superconducting circuit quantum electrodynamical system*, *Nat. Commun.* **9**, 1291 (2018).
- [87] David Williams, *Probability with Martingales* (Cambridge University Press, Cambridge, England, 1991).
- [88] J. L. Doob, *Stochastic Processes* (Wiley, New York, 1990).
- [89] John A. C. Albay, Yonggun Jun, and Pik-Yin Lai, *Winning strategies of a gambling demon in a Brownian particle under a squeezing potential*, *Phys. Rev. Res.* **5**, 023115 (2023).
- [90] M. Ribezzi-Crivellari and F. Ritort, *Large work extraction and the Landauer limit in a continuous Maxwell demon*, *Nat. Phys.* **15**, 660 (2019).
- [91] M. Rico-Pasto, R. K. Schmitt, M. Ribezzi-Crivellari, J. M. R. Parrondo, H. Linke, J. Johansson, and F. Ritort, *Dissipation reduction and information-to-measurement conversion in DNA pulling experiments with feedback protocols*, *Phys. Rev. X* **11**, 031052 (2021).
- [92] Ilya Prigogine, *Etude thermodynamique des phénomènes irréversibles*, *Bull. Acad. R. Belg. Cl. Sci.* **31**, 600 (1945).
- [93] L. Onsager and S. Machlup, *Fluctuations and irreversible processes*, *Phys. Rev.* **91**, 1505 (1953).
- [94] L. Bertini, A. De Sole, D. Gabrielli, G. Jona-Lasinio, and C. Landim, *Minimum dissipation principle in stationary non-equilibrium states*, *J. Stat. Phys.* **116**, 831 (2004).
- [95] J. W. Carlyle and A. Paz, *Realizations by stochastic finite automata*, *J. Comput. Syst. Sci.* **5**, 26 (1971).
- [96] Jorge Castro and Ricard Gavaldà, *Learning probability distributions generated by finite-state machines*, in *Topics in Grammatical Inference* (Springer, Berlin, 2016), pp. 113–142.
- [97] Izaak Neri, Édgar Roldán, Simone Pigolotti, and Frank Jülicher, *Integral fluctuation relations for entropy production at stopping times*, *J. Stat. Mech.* (2019) 104006.
- [98] Martin R. Evans, Satya N. Majumdar, and Grégory Schehr, *Stochastic resetting and applications*, *J. Phys. A* **53**, 193001 (2020).
- [99] Shamik Gupta and Arun M. Jayannavar, *Stochastic resetting: A (very) brief review*, *Front. Phys.* **10**, 789097 (2022).
- [100] Ofir Tal-Friedman, Arnab Pal, Amandeep Sekhon, Shlomi Reuveni, and Yael Roichman, *Experimental realization of diffusion with stochastic resetting*, *J. Phys. Chem. Lett.* **11**, 7350 (2020).
- [101] Édgar Roldán, Ana Lisica, Daniel Sánchez-Taltavull, and Stephan W. Grill, *Stochastic resetting in backtrack recovery by RNA polymerases*, *Phys. Rev. E* **93**, 062411 (2016).
- [102] Paul C. Bressloff, *Modeling active cellular transport as a directed search process with stochastic resetting and delays*, *J. Phys. A* **53**, 355001 (2020).
- [103] Antoine Bérut, Artak Arakelyan, Artyom Petrosyan, Sergio Ciliberto, Raoul Dillenschneider, and Eric Lutz, *Experimental verification of Landauer’s principle linking information and thermodynamics*, *Nature (London)* **483**, 187 (2012).
- [104] Yonggun Jun, Momčilo Gavrilov, and John Bechhoefer, *High-precision test of Landauer’s principle in a feedback trap*, *Phys. Rev. Lett.* **113**, 190601 (2014).
- [105] S. Ciliberto, *Experiments in stochastic thermodynamics: Short history and perspectives*, *Phys. Rev. X* **7**, 021051 (2017).

- [106] Ignacio A. Martínez, Édgar Roldán, Luis Dinis, and Raúl A. Rica, *Colloidal heat engines: A review*, *Soft Matter* **13**, 22 (2017).
- [107] J. P. Pekola, D. S. Golubev, and D. V. Averin, *Maxwell's demon based on a single qubit*, *Phys. Rev. B* **93**, 024501 (2016).
- [108] Salambô Dago and Ludovic Bellon, *Dynamics of information erasure and extension of Landauer's bound to fast processes*, *Phys. Rev. Lett.* **128**, 070604 (2022).
- [109] Sidney Redner, *A Guide to First-Passage Processes* (Cambridge University Press, Cambridge, England, 2001).
- [110] Vladimir Y. Chernyak, Michael Chertkov, and Christopher Jarzynski, *Path-integral analysis of fluctuation theorems for general Langevin processes*, *J. Stat. Mech.* (2006) P08001.
- [111] Gatién Verley, Raphaël Chétrite, and David Lacoste, *Inequalities generalizing the second law of thermodynamics for transitions between nonstationary states*, *Phys. Rev. Lett.* **108**, 120601 (2012).
- [112] Arya Datta, Patrick Pietzonka, and Andre C. Barato, *Second law for active heat engines*, *Phys. Rev. X* **12**, 031034 (2022).
- [113] Faezeh Khodabandehlou, Christian Maes, and Karel Netočný, *A Nernst heat theorem for nonequilibrium jump processes*, *J. Chem. Phys.* **158**, 204112 (2023).
- [114] Pritha Dolai, Christian Maes, and Karel Netočný, *Calorimetry for active systems*, *SciPost Phys.* **14**, 126 (2023).
- [115] Yaakov Benenson, *Biomolecular computing systems: Principles, progress and potential*, *Nat. Rev. Genet.* **13**, 455 (2012).
- [116] Athel Cornish-Bowden, *Fundamentals of Enzyme Kinetics* (John Wiley and Sons, New York, 2013).
- [117] Julien O. Dubuis, Gašper Tkačik, Eric F. Wieschaus, Thomas Gregor, and William Bialek, *Positional information, in bits*, *Proc. Natl. Acad. Sci. U.S.A.* **110**, 16301 (2013).
- [118] Chunhe Li and Jin Wang, *Quantifying cell fate decisions for differentiation and reprogramming of a human stem cell network: Landscape and biological paths*, *PLoS Comput. Biol.* **9**, e1003165 (2013).
- [119] Jin Wang, *Landscape and flux theory of non-equilibrium dynamical systems with application to biology*, *Adv. Phys.* **64**, 1 (2015).
- [120] Ivana Ban, Lucija Tomašić, Marianna Trakala, Iva M. Tolić, and Nenad Pavin, *Proliferative advantage of specific aneuploid cells drives evolution of tumor karyotypes*, *Biophys. J.* **122**, 632 (2023).
- [121] Christopher Battle, Chase P. Broedersz, Nikta Fakhri, Veikko F. Geyer, Jonathon Howard, Christoph F. Schmidt, and Fred C. MacKintosh, *Broken detailed balance at mesoscopic scales in active biological systems*, *Science* **352**, 604 (2016).
- [122] Étienne Fodor, Cesare Nardini, Michael E. Cates, Julien Tailleur, Paolo Visco, and Frédéric van Wijland, *How far from equilibrium is active matter?*, *Phys. Rev. Lett.* **117**, 038103 (2016).
- [123] Édgar Roldán, Jérémie Barral, Pascal Martin, Juan M. R. Parrondo, and Frank Jülicher, *Quantifying entropy production in active fluctuations of the hair-cell bundle from time irreversibility and uncertainty relations*, *New J. Phys.* **23**, 083013 (2021).
- [124] José A. Morin, Francisco J. Cao, José M. Lázaro, J. Ricardo Arias-Gonzalez, José M. Valpuesta, José L. Carrascosa, Margarita Salas, and Borja Ibarra, *Mechanochemical kinetics of DNA replication: Identification of the translocation step of a replicative DNA polymerase*, *Nucleic Acids Res.* **43**, 3643 (2015).
- [125] Annwasha Dutta, Gunter M. Schütz, and Debashish Chowdhury, *Stochastic thermodynamics and modes of operation of a ribosome: A network theoretic perspective*, *Phys. Rev. E* **101**, 032402 (2020).
- [126] David H. Wolpert, *The stochastic thermodynamics of computation*, *J. Phys. A* **52**, 193001 (2019).
- [127] Philipp Strasberg and Massimiliano Esposito, *Stochastic thermodynamics in the strong coupling regime: An unambiguous approach based on coarse graining*, *Phys. Rev. E* **95**, 062101 (2017).
- [128] Robert S. Whitney, *Non-Markovian quantum thermodynamics: Laws and fluctuation theorems*, *Phys. Rev. B* **98**, 085415 (2018).
- [129] Cillian Cockrell and Ian J. Ford, *Stochastic thermodynamics in a non-Markovian dynamical system*, *Phys. Rev. E* **105**, 064124 (2022).
- [130] A. Gomez-Marin, J. M. R. Parrondo, and C. Van den Broeck, *Lower bounds on dissipation upon coarse graining*, *Phys. Rev. E* **78**, 011107 (2008).
- [131] Édgar Roldán and Juan M. R. Parrondo, *Estimating dissipation from single stationary trajectories*, *Phys. Rev. Lett.* **105**, 150607 (2010).