

# Extreme Quantum Memory Advantage for Rare-Event Sampling

Cina Aghamohammadi,\* Samuel P. Loomis,† John R. Mahoney,‡ and James P. Crutchfield§

*Complexity Sciences Center and Physics Department,  
University of California at Davis, One Shields Avenue, Davis, CA 95616*

(Dated: October 30, 2017)

We introduce a quantum algorithm for memory-efficient biased sampling of rare events generated by classical memoryful stochastic processes. Two efficiency metrics are used to compare quantum and classical resources for rare-event sampling. For a fixed stochastic process, the first is the classical-to-quantum ratio of required memory. We show for two example processes that there exists an infinite number of rare-event classes for which the memory ratio for sampling is larger than  $r$ , for any large real number  $r$ . Then, for a sequence of processes each labeled by an integer size  $N$ , we compare how the classical-to-quantum required memory ratio scales with  $N$ . In this setting, since both memories can diverge as  $N \rightarrow \infty$ , the efficiency metric tracks how fast they diverge. An *extreme quantum memory advantage* exists when the classical memory diverges in the limit  $N \rightarrow \infty$ , but the quantum memory has a finite bound. We then show that finite-state Markov processes and spin chains exhibit extreme memory advantage for sampling of almost all of their rare-event classes.

PACS numbers: 05.45.-a 89.75.Kd 89.70.+c 05.45.Tp

Keywords: quantum algorithm, large deviation theory, biased sampling, quantum memory, quantum advantage, stochastic process, hidden Markov model

## I. INTRODUCTION

From earthquakes to financial market crashes, rare events are associated with catastrophe—from decimated social infrastructure and the substantial loss of life to global economic collapse. Though rare, their impact cannot be ignored. Prediction and modeling such rare events is essential to mitigating their effects. However, this is particularly challenging, often requiring huge datasets and massive computational resources, precisely because the events of interest are rare.

Ameliorating much of the challenge, *biased* or *extended sampling* [1, 2] is an effective and now widely-used method for efficient generation and analysis of rare events. The underlying idea is simple to state: transform a given distribution to a new one where previously-rare events are now typical. This concept was originally proposed in 1961 by Miller to probe the rare events generated by discrete-time, discrete-value Markov stochastic processes [3]. It has since been extended to address non-Markovian processes [4]. The approach was also eventually adapted to continuous-time first-order Markov processes [5–7]. Today, the statistical analysis of rare events is a highly developed toolkit with broad applications in sciences and engineering [8]. Given this, it is perhaps not surprising that the idea and its related methods appear under different appellations, depending on the research arena. For

example, large deviation theory refers to the *s-ensemble method* [9, 10], the *exponential tilting algorithm* [11, 12], or as generating *twisted distributions*.

In 1997, building on biased sampling, Torrie and Valleau introduced *umbrella sampling* into Monte Carlo simulation of systems whose energy landscapes have high energy barriers and so suffer particularly from poor sampling [13]. Since then, stimulated by computational problems arising in statistical mechanics, the approach was generalized to *Ferrenberg-Swendsen reweighting*, later still to *weighted histogram analysis* [14], and more recently to *Wang-Landau sampling* [15].

When generating samples for a given stochastic process one can employ alternative types of algorithm. There are two main types—Monte Carlo or finite-state machine algorithms. Here, we consider finite-state machine algorithms based on Markov chains (MC) [16, 17] and hidden Markov models (HMM) [18–20]. For example, if the process is Markovian one uses MC generators and, in more general cases, one uses HMM generators.

When evaluating alternative approaches the key questions that arise concern algorithm speed and memory efficiency. For example, it turns out there are HMMs that are always equally or more memory efficient than MCs. There are many finite-state HMMs for which the analogous MC is infinite-state [21]. And so, when comparing all HMMs that generate the same process, one is often interested in those that are most memory efficient. For a generic stochastic process, the most memory efficient classical HMM known currently is the  $\epsilon$ -*machine* of computational mechanics [22]. The memory it requires is called the process’ *statistical complexity*  $C_\mu$  [23].

---

\* caghamohammadi@ucdavis.edu

† sloomis@ucdavis.edu

‡ jrmahoney@ucdavis.edu

§ chaos@ucdavis.edu

Today, we have come to appreciate that several important mathematical problems can be solved more efficiently using a quantum computer. Examples include quantum algorithms for integer factorization [24], search [25], eigen-decomposition [26], and solving linear systems [27]. Not long ago and for the first time, Ref. [28] provided a quantum algorithm that can perform stochastic process sample-generation using less memory than the best-known classical algorithms. Recently, using a stochastic process’ higher-order correlations, a new quantum algorithm—the *q-machine*—substantially improved this efficiency and extended its applicability [29]. More detailed analysis and a derivation of the closed-form quantum advantage of the *q-machine* is given in a sequel [30]. Notably, the quantum advantage has been verified experimentally for a simple case [31].

The following brings together techniques from large deviation theory, classical algorithms for stochastic process generation, computational complexity theory, and the newly introduced quantum algorithm for stochastic process generation to propose a new, memory efficient quantum algorithm for the biased sampling problem. We show that there can be an extreme advantage in the quantum algorithm’s required memory compared to the best known classical algorithm where the required memory for classical algorithm grows unboundedly with problem size, but is bounded from above for the quantum algorithm. Three examples are analyzed here. The first is the simple, but now well-studied *Perturbed Coin Process*. The second is a more physical example—a stochastic process that arises from the Ising next-nearest-neighbor spin system in contact with thermal reservoir. The third is a sequence of processes generated by a series of Ising  $N$ -nearest-neighbor Hamiltonians.

Today, we know of several different sampling problems for which their best quantum algorithm has an advantage compared to the best classical algorithm. These sampling problems fall into two categories. First are those in which the problem is quantum in nature, such as boson sampling [32]. Second are the ones in which the target problem is classical. Function sampling [33] and mixing [34] are in this category.

On the one hand, the advantage for both boson sampling and mixing appears in shorter run times for the quantum algorithm. For the problem of rare-event sampling we study here the run times for both classical and quantum algorithms are the same. On the other hand, for function sampling advantage appears in smaller required memory; this is similar to our problem. In both boson sampling and function sampling, the advantage appears as an increasing function of problem size. While for mixing it is function of the spectral gap—that is, a property of the problem input—the Markov chain of interest, in that case. The

quantum memory advantage we introduce here is both a function of problem size and a property of the input instance.

In the boson sampling problem, a linear system scatters  $N$  individual bosons into  $M \gg N$  output modes. The goal then is to sample from the output distribution. It is known that for large  $N$  and  $M$ , the run time for the quantum algorithm is much smaller than for the classical algorithm, while both algorithms need memory on the order of the required sample size. In mixing, a Markov chain and an initial state are given and the goal is to sample from the stationary distribution over Markov chain’s states with some acceptable error margin. Denoting the spectral gap for Markov chain’s transition matrix by  $\delta$ , the run time for the best known classical algorithm increases faster than the quantum algorithm when  $\delta \rightarrow 0$  [35]. As a result, the notion of an advantage is captured by a function of  $\delta$ . In function sampling, a function  $f : X \times Y \rightarrow \{0, 1\}$  and a probability distribution  $\mathbb{P}(X, Y)$  over  $X \times Y$  are given. Alice and Bob start with no inputs. The goal then is to sample  $X$ ,  $Y$ , and  $Z$  from the distribution  $(\mathbb{P}, f(\mathbb{P}))$ , where Alice end up with  $X$  and Bob with  $Y$  and  $Z$ . Algorithm efficiency is then defined by how much information Alice and Bob must communicate during the algorithm. It turns out that the best known quantum algorithm has markedly smaller communication costs than the classical. In the function sampling problem, as in many other similar problems, communication cost can be framed as a memory cost, since Alice can always write the message in a memory that Bob reads.

## II. CLASSICAL ALGORITHM

The object for which we wish to generate samples is a discrete-time, discrete-value *stochastic process* [18, 36]: a probability space  $\mathcal{P} = \{\mathcal{A}^\infty, \Sigma, \mathbb{P}(\cdot)\}$ , where  $\mathbb{P}(\cdot)$  is a probability measure over the bi-infinite chain  $\dots X_{-2}X_{-1}X_0X_1X_2\dots$ , each random variable  $X_i$  takes values in a finite, discrete alphabet  $\mathcal{A}$ , and  $\Sigma$  is the  $\sigma$ -algebra generated by the cylinder sets in  $\mathcal{A}^\infty$ . For simplicity we consider only ergodic stationary processes: that is,  $\mathbb{P}(\cdot)$  is invariant under time translation— $\mathbb{P}(X_{i_1}X_{i_2}\dots X_{i_m}) = \mathbb{P}(X_{i_1+n}X_{i_2+n}\dots X_{i_m+n})$  for all  $n$ ,  $m$ —and over successive realizations.

*Sampling* or *generating* a given stochastic process refers to producing a finite realization that comes from the process’ probability distribution. There are two main generation (sampling) problems: *sequential generation* and *simultaneous generation* [37]. In sequential generation or one-shot sampling the goal is to generate one long sample from the given process. However, in simultaneous




FIG. 1. Hidden Markov model generator of a stochastic process with infinite-range statistical dependencies that requires an HMM with only six states. To generate the same process via a Markov chain requires one with an infinite number of states and so infinite memory.

generation the goal is to generate  $M \gg 1$  realizations of a process simultaneously, each of which is statistically independent of the others.

Generally, generating a **memoryful** process via its probability measure  $\mathbb{P}(\cdot)$  is impossible due to the vast number of allowed realizations and, as a result, this prosaic approach requires an unbounded amount of memory. Fortunately, there are more compact ways than specifying in-full the probability measure on the sequence sigma algebra. This recalls the earlier remark that HMMs can be arbitrarily more compact than alternative algorithms for the task of generation.

An HMM is specified by a tuple  $\{\mathcal{S}, \mathcal{A}, \{T^{(x)}, x \in \mathcal{A}\}\}$ . In this,  $\mathcal{S}$  is a finite set of states,  $\mathcal{A}$  is a finite alphabet, and  $\{T^{(x)}, x \in \mathcal{A}\}$  is a set of  $|\mathcal{S}| \times |\mathcal{S}|$  substochastic symbol-labeled transition matrices whose sum  $T = \sum_{x \in \mathcal{A}} T^{(x)}$  is a stochastic matrix.

As an example, consider the HMM state-transition diagram shown in Fig. 1, where  $\mathcal{S} = \{A, B, C, D, E, F\}$ ,  $\mathcal{A} = \{0, 1, 2\}$ , and we have three  $6 \times 6$  substochastic matrices  $T^{(0)}$ ,  $T^{(1)}$ , and  $T^{(2)}$ . Each edge is labeled  $p|x$  denoting the transition probability  $p$  and a symbol  $x \in \mathcal{A}$  which is emitted during the transition. In this HMM, of the two edges exiting state  $C$ , one enters state  $B$  and the other enters state  $A$ . The edges from  $C$  to  $A$  and  $C$  to  $B$  are labeled by  $\frac{1}{2}|1$  and  $\frac{1}{2}|0$ . This simply means that if the HMM is in the state  $C$ , then with probability  $\frac{1}{2}$  it goes to the state  $A$  and emits the symbol  $1$  and with probability  $\frac{1}{2}$  it goes to state  $B$  and emits symbol  $0$ . Following these transition rules in succession generates realizations in the HMM's process.

How does this generation method compare to generating realizations of the same process via a finite Markov chain? (Recall that states in a MC not hidden:  $\mathcal{A} = \mathcal{S}$ ) It turns out that this cannot be implemented, since generating a symbol can depend on the infinite history. That is,

the process has infinite Markov order. As a result, to generate a realization using a Markov chain one needs an infinite number of Markovian states. In other words, implementing the Markov chain algorithm to generate process samples on a conventional computer requires an infinite amount of memory.

To appreciate the reason behind the process' infinite Markov order, refer to Fig. 1's HMM. There are two length-3 state-loops consisting of the edges colored red (right side of state-transition diagram) and those colored maroon (left side). Note that if the HMM generates  $n$  1s in a row, we will not know the HMM's current state, only that it is either  $A$ ,  $D$ , or  $E$ . This state uncertainty (entropy) is bounded away from 0. The observation holds for the other loop and its sequences of symbol  $0$  and the consequent ambiguity among states  $B$ ,  $C$ , and  $F$ . Thus, there exist process realizations from which we cannot determine the future statistics, independent of the number of symbols seen. This means that the process statistics depend on infinite past sequences—the process has infinite Markov order. To emphasize, implementing a MC algorithm for this requires infinite memory. The contrast with the finite HMM method is an important lesson: HMMs are strictly more powerful generators, as a class of algorithms, than Markov chain generators.

For any given process  $\mathcal{P}$ , there are an infinite number of HMMs that generate it. Therefore, one is compelled to ask, which algorithm requires the least memory for implementation? To appreciate the answer, let's first address how much state memory one needs to run an HMM.

Consider sequential generation in which the goal is to produce a very long realization of a process. For this, we use one computer with a code that runs the algorithm (HMM). At each step, the computer must memorize the current HMM state. Assuming the HMM has  $N$  states, this requires  $\log_2(N)$  bits of memory. As a result, if one wishes to implement one-shot sampling using the minimum required memory then, over all the process' HMM generators, one needs to find that with the minimum number of states.

Here, though, we are interested in simultaneous generation for which the goal is to simultaneously generate  $M \gg 1$  process realizations, each of which is statistically independent of the others. The effective implementation uses  $M$  computers each with the above code. Similar to the sequential problem, each computer must memorize the current state of its HMM. If each computer uses its own memory, each needs  $\log_2(N)$  bits of memory as before. The total memory is then  $M \log_2(N)$  bits. However, we can reduce required memory by using one large memory shared among the computers. Figure 2 depicts this




FIG. 2. (Left) Even Process  $\epsilon$ -machine. (Right) Schematic of simultaneous generation problem. Each black box contains an Even Process generator. They all share the same memory for tracking the individual generator states.

schematically. In this way, according to Shannon’s coding theorem [38], we can encode HMM states to reduce the amount of memory down to  $M H(\mathcal{S}) \leq M \log_2(N)$  bits, where  $H(\mathcal{S})$  is the Shannon entropy of the probability distribution over HMM’s states. The memory per sample is then just  $H(\mathcal{S})$ . As a result, if one needs to do simultaneous sampling of a given process using minimum required memory, over all its HMM generators, one needs to find the HMM with the minimum Shannon state entropy.

For both one shot and simultaneous sampling, the best known implementation, and provably the optimal predictor, is known as the  $\epsilon$ -machine  $M$  [22, 39]. Its states are called *causal states*; we denote this set  $\mathcal{S}$ . The average memory required for  $M(\mathcal{P})$  to sequentially sample process  $\mathcal{P}$  is given by the process’ *statistical complexity*  $C_\mu(\mathcal{P})$  [23]. To calculate it:

1. Compute the stationary distribution  $\pi$  over causal states.  $\pi$  is the left eigenvector of the state-transition matrix  $T$  with eigenvalue 1:  $\pi T = \pi$ .
2. Calculate the state’s Shannon entropy  $H[\mathcal{S}] = -\sum_{\sigma \in \mathcal{S}} \pi(\sigma) \log_2 \pi(\sigma)$ .

$C_\mu = H[\mathcal{S}]$  measures the (ensemble average) memory required simultaneous sampling of the process.

Another important, companion measure is  $h_\mu$ , the process’ *metric entropy* (or Shannon entropy rate) [40]:

$$h_\mu(\mathcal{P}) = -\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{w \in \mathcal{A}^n} \mathbb{P}(w) \log_2 \mathbb{P}(w) .$$

Although sometimes confused, it is important to emphasize that  $h_\mu$  describes randomness in realizations, while  $C_\mu$  is the average memory required to generate them.

### III. QUANTUM MEMORY ADVANTAGE

Recently, it was shown that a quantum algorithm for process generation can use less memory than the best known classical algorithm ( $\epsilon$ -machine) [28]. By accounting for a process’ higher-order correlations, a new quantum algorithm—the *q-machine*—was introduced that substantially improved the original and is, to date, the most memory-efficient quantum algorithm known for process generation [29]. We refer to the ratio of required classical memory  $C_\mu$  to quantum memory  $C_q$  as the *quantum memory advantage*. Closed-form expressions for the quantum memory advantage are given in Ref. [30].

Importantly, the quantum advantage was recently verified experimentally for the simple *perturbed coins process* [31]. It was also discovered that the q-machine can confer an extreme quantum memory advantage. For example, when generating ground-state configurations (in a Dyson-type spin model with  $N$ -nearest-neighbor interactions at temperature  $T$ ), the quantum advantage scales as  $NT^2/\log_2 T$  [41, 42]. One consequence of quantum advantage arises in model selection [43]. Statistical inference of models for stochastic systems often involves controlling for model size or memory. And so, quantum statistical inference may see large improvements.

The following determines the quantum advantage in biased sampling of a process’ rare events. In particular, we develop tools to determine how the memory requirements of classical and quantum algorithms vary over rare-event classes.

### IV. QUANTUM ALGORITHM

We define a stochastic process  $\mathcal{P}$ ’s *quantum machine* by  $Q(\mathcal{P}) = \{\mathcal{H}, \mathcal{A}, \{K_x, x \in \mathcal{A}\}\}$ , where  $\mathcal{H}$  denotes the Hilbert space with dimension  $|\mathcal{S}|$  in which quantum states reside,  $\mathcal{A}$  is the same alphabet as the given process’, and  $\{K_x, x \in \mathcal{A}\}$  is a set of Kraus operators we use to specify the measurement protocol for states [44].<sup>1</sup> Assume we have the *quantum state* (density matrix)  $\rho_0 \in \mathcal{B}(H)$  in hand. We perform a measurement by applying Kraus operators and, as a result, measure  $X$ . The probability of yielding symbol  $X = x$  is:

$$\mathbb{P}(X = x_0 | \rho_0) = \text{tr}(K_{x_0} \rho_0 K_{x_0}^\dagger) .$$

<sup>1</sup> We adopt a particular form for the Kraus operators. In general, they are not unique.

After measurement with outcome  $X = x_0$ , the new quantum state is:

$$\rho_1 = \frac{K_{x_0} \rho_0 K_{x_0}^\dagger}{\text{tr}(K_{x_0} \rho_0 K_{x_0}^\dagger)} .$$

Repeating these measurements generates a stochastic process. The process potentially could be nonergodic, depending on the initial state  $\rho_0$ . However, starting the  $q$ -machine in the stationary state defined by:

$$\rho_s = \sum_{x \in \mathcal{A}} K_x \rho_s K_x^\dagger ,$$

and repeatedly measuring generates a stationary stochastic process over  $x \in \mathcal{A}$ . For any given process,  $\rho_s$  can be calculated by the method introduced in Ref. [30].

Our immediate goal is to design a quantum generator of a given classical process. (Section VI will then take the given process to represent a rare-event class of another process.) For now, we start with the process'  $\epsilon$ -machine. The construction consists of three steps, as follows.

*First step* Map every causal state  $\sigma_i \in \mathcal{S}$  to a pure quantum state  $|\eta_i\rangle$ . Each signal state  $|\eta_i\rangle$  encodes the set of length- $R$  sequences that may follow  $\sigma_i$ , as well as each corresponding conditional probability:

$$|\eta_i\rangle \equiv \sum_{w \in \mathcal{A}^R} \sqrt{\mathbb{P}(w|\sigma_i)} |w\rangle ,$$

where  $w$  denotes a length- $R$  sequence,  $\mathbb{P}(w|\sigma_i) = \mathbb{P}(X_0 \cdots X_{R-1} = w | \mathcal{S}_0 = \sigma_i)$ , and  $R$  is the process' the Markov order. The resulting Hilbert space is  $\mathcal{H}_w$  with size  $|\mathcal{A}|^R$ , the number of length- $R$  sequences, with basis elements  $|w\rangle = |x_0\rangle \otimes \cdots \otimes |x_{R-1}\rangle$ . From here on out, assume all  $|\eta_i\rangle$ s are linearly independent.<sup>2</sup> As a result, one defines  $|\mathcal{S}|$  new states  $|\xi_i\rangle$  that reside in a Hilbert space of size  $|\mathcal{S}|$ . (This is much smaller than the  $|\eta_i\rangle$ 's Hilbert space, which has size  $|\mathcal{A}|^R$ .) Moreover, the  $|\xi_i\rangle$ s have the same pairwise overlaps as the  $|\eta_i\rangle$ s. That is, for all  $i, j$ :

$$\langle \xi_i | \xi_j \rangle = \langle \eta_i | \eta_j \rangle .$$

Reference [30] developed a method to calculate all the overlaps  $\langle \eta_i | \eta_j \rangle$  for a given process in closed form. Since the  $|\eta_i\rangle$ s are linearly independent one can use the overlaps to construct the  $|\xi_i\rangle$ s.

*Second step* Form a matrix  $\Xi$  by assembling the signal states:

$$\Xi = [ |\xi_0\rangle \quad |\xi_1\rangle \quad \cdots \quad |\xi_{|\mathcal{S}|-1}\rangle ] .$$

Define  $|\mathcal{S}|$  new bra states  $|\tilde{\xi}_i\rangle$ :

$$\begin{bmatrix} \langle \tilde{\xi}_0 | \\ \langle \tilde{\xi}_1 | \\ \vdots \\ \langle \tilde{\xi}_{|\mathcal{S}|-1} | \end{bmatrix} = \Xi^{-1} .$$

That is, we design the new bra states such that we obtain the identity:

$$\begin{bmatrix} \langle \tilde{\xi}_0 | \\ \langle \tilde{\xi}_1 | \\ \vdots \\ \langle \tilde{\xi}_{|\mathcal{S}|-1} | \end{bmatrix} [ |\xi_0\rangle \quad |\xi_1\rangle \quad \cdots \quad |\xi_{|\mathcal{S}|-1}\rangle ] = \mathbf{I} .$$

*Third step* Define  $|\mathcal{A}|$  Kraus operators  $K_x$  for all  $x \in \mathcal{A}$  via:

$$K_x = \sum_{i,j} \sqrt{T_{ij}^x} |\xi_j\rangle \langle \tilde{\xi}_i | .$$

By applying Kraus operators repeatedly one generates the target stochastic process. For example, consider the case in which the  $q$ -machine is in state  $|\xi_i\rangle \langle \xi_i|$  and we apply the Kraus operators. Then, if we do not make a measurement, the next state is  $\sum_{j,x} T_{ij}^x |\xi_j\rangle \langle \xi_j|$ . Let us say, though, that we make a measurement and the result is  $x$ . The next state is  $|\xi_j\rangle \langle \xi_j|$ , where  $j$  is the unique index where  $T_{ij}^x$  is nonzero.

The Hilbert space in which the algorithm operates has dimension  $|\mathcal{S}|$ . Since the operation is not unitary in this space and measurements are not projective, one may argue that  $|\mathcal{S}|$  is not the actual size of the Hilbert space needed to physically implement the algorithm. However, it was recently shown that the algorithm presented here can be implemented by unitary operations and projective measurements in a Hilbert space with dimension  $|\mathcal{S}| * |\mathcal{A}|$  [45]. This new result gives an experimental implementation of our algorithm.

Using the quantum generator  $Q(\mathcal{P})$ , the required average memory for simultaneous generation of process  $\mathcal{P}$  is  $C_q(\mathcal{P}) = S(\rho_s)$ , where  $S(\rho) = -\text{tr}(\rho \log \rho)$  denotes the *von Neumann entropy* [44], and the required average memory for sequential generation of  $\mathcal{P}$  is the Hilbert space dimension  $|\mathcal{S}|$ .

Comparing memory efficiencies of classical and quantum algorithms requires an efficiency metric. Depending on the

<sup>2</sup> The procedure requires only slight modification for linearly dependent causal states and, in any case, does not affect the results.

setting, there are two one can use. In the *single-process* case,  $\mathcal{P}$  is given and memory efficiency is defined as a the ratio of required memory for the classical algorithm to the quantum algorithm. Here, since we only explore finite-size stochastic processes, both memories are finite and the ratio is a good quantitative efficiency measure.

In the *multi-process* case, a series of stochastic processes is given, with each process labeled by an integer  $N$  that measures process size. Then, memory efficiency is defined by how the memory scales in  $N$  for the classical algorithm compared to the quantum. This metric is closer to relative-complexity definitions familiar in computation complexity theory. In the present case, since both memories are allowed to diverge when  $N \rightarrow \infty$ , the quantitative measure of efficiency tracks how fast they diverge. We say we have *extreme quantum memory advantage* when the classical memory diverges as  $N \rightarrow \infty$ , but quantum memory converges to a finite quantity.

## V. TYPICAL REALIZATIONS

At this point, we established classical and quantum representations of processes and characterized their respective memory requirements. Using this, our purpose now shifts to monitor classical and quantum resources required to generate rare-event classes of a process' realizations. Our first task is to review the theory of typical events and their complement—rare events.

The concept of a stochastic process is quite general. Any physical system that exhibits stochastic dynamics in time or space may be thought of as *generating* a stochastic process. In the spatial setting one considers not time evolution, but rather the spatial “dynamic”. For example, consider a one-dimensional noninteracting Ising spin- $\frac{1}{2}$  chain with classical Hamiltonian  $H = -\sum_{i=1}^n h\sigma_i$  in contact with a thermal reservoir at temperature  $T$ . After thermalizing, a spin configuration at one instant of time may be thought of as having been generated by a process that scans the lattice left-to-right (or, equivalently, right-to-left). The probability distribution over these spatial-translation invariant configurations defines a stationary stochastic process—a simple Markov random field.

For  $n \gg 1$ , one can ask for the probability of seeing  $k$  up spins. The Strong Law of Large Numbers [46] guarantees that for large  $n$ , the ratio  $k/n$  almost surely converges to  $p_\uparrow = \frac{1}{2}(1 + \tanh(h/k_B T))$ . That is:

$$\mathbb{P}\left(\lim_{n \rightarrow \infty} \frac{k}{n} = p_\uparrow\right) = 1.$$

Informally, a *typical sequence* is one that has close to  $p_\uparrow n$  up-spins. However, this does not preclude seeing other




FIG. 3. For a given process, the space  $\mathcal{A}^\infty$  of all sequences is partitioned into forbidden sequences, sequences in the typical set, and sequences neither forbidden nor typical—the *atypical* or rare sequences.

kinds of long runs, e.g., all up-spins or all down-spins. It simply means that the latter are *rare events*, compared to the typical ones.

Now, let us formally define the concept of typical realizations and, consequently, rare ones. Consider a given process  $\mathcal{P}$  and let  $\mathcal{A}^n$  denote the set of all possible length- $n$  realizations. Then, for an arbitrary  $0 < \epsilon \ll 1$  the process' *typical set* [38, 47, 48] is defined:

$$\mathcal{A}_\epsilon^n \equiv \{w : 2^{-n(h_\mu + \epsilon)} \leq \mathbb{P}(w) \leq 2^{-n(h_\mu - \epsilon)}, w \in \mathcal{A}^n\}, \quad (1)$$

where  $h_\mu$  is the process' Shannon entropy rate, introduced above.

According to the *Shannon-McMillan-Breiman theorem* [49–51], for a given  $\epsilon \ll 1$ , sufficiently large  $n^*$ , and  $w \in \mathcal{A}^n$ :

$$\mathbb{P}(w \notin \mathcal{A}_\epsilon^n) \leq \epsilon, \quad (2)$$

for all  $n \geq n^*$ . There are two important lessons here. First, from Eq. (1) we see that all sequences in the typical set have approximately the same probability. More precisely, the probability of typical sequences decays at the *same exponential rate*. The following adapts this to use decay rates to identify distinct sets of rare events. Second, coming from Eq. (2), for large  $n$  the probability of sequences falling outside the typical set is close to zero—these are the sets of rare sequences.

Another important consequence of the theorem is that sequences generated by a stationary ergodic process fall into one of three partitions; see Fig. 3. The first contains sequences that are never generated; they fall in the the *forbidden set*. For example, the HMM in Fig. 1 never generates sequences that have consecutive 2s. The second partition consists of those in the typical set—the set with probability close to one, as in Eq. (1). And, the last contains sequences in a family of atypical sets—realizations

that are rare to different degrees. We now refine this classification by dividing the atypical sequences into identifiable subsets, each with their own characteristic rarity. Mirroring the familiar *Boltzmann weight* in statistical physics [52], in the  $n \rightarrow \infty$  limit, we define the subsets  $\Lambda_U^{\mathcal{P}} \subset \mathcal{A}^\infty$  for a process  $\mathcal{P}$  as:

$$\Lambda_{U,n}^{\mathcal{P}} = \left\{ w : -\frac{\log_2 \mathbb{P}(w)}{n} = U, w \in \mathcal{A}^n \right\} \quad \text{and} \quad (3)$$

$$\Lambda_U^{\mathcal{P}} = \lim_{n \rightarrow \infty} \Lambda_{U,n}^{\mathcal{P}}.$$

This partitions  $\mathcal{A}^\infty$  into disjoint subsets  $\Lambda_U^{\mathcal{P}}$  in which all  $w \in \Lambda_U^{\mathcal{P}}$  have the same probability decay rate  $U$ . Physics vernacular would speak of the sequences having the same *energy density*  $U$ .<sup>3</sup> Figure 4 depicts these subsets as “bubbles” of equal energy. Equation (1) says the typical set is that bubble with energy equal to the process’ Shannon entropy rate:  $U = h_\mu$ . All the other bubbles contain rare events, some rarer than others. They exhibit faster or slower probability decay rates.

Employing a process’ HMM to generate realizations produces sequences in the typical set with probability close to one and, rarely, atypical sequences. Imagine that one is interested in a particular class  $\Lambda_U^{\mathcal{P}}$  of rare sequences, say, those with energy  $U$ . (One might be concerned about the class of large-magnitude earthquakes or the emergence of major instabilities in the financial markets, for example.) How can one efficiently generate these rare sequences? We now show that there is a new process  $\mathcal{P}^U$  whose typical set is  $\Lambda_U^{\mathcal{P}}$  and this returns us directly to the challenge of biased sampling.

## VI. BIASED SAMPLING

Consider a finite set of configurations  $\{c_i\}$  with probabilities specified by distribution  $\mathbb{P}(\cdot)$  and an associated set  $\{\omega_i\}$  of *weighting factors*. Consider the procedure of reweighting that introduces a new distribution  $\tilde{\mathbb{P}}(\cdot)$  over configurations where:

$$\tilde{\mathbb{P}}(c_i) = \frac{\mathbb{P}(c_i) \exp(\omega_i)}{\sum_i \mathbb{P}(c_i) \exp(\omega_i)}.$$

Given a process  $\mathcal{P}$  and its  $\epsilon$ -machine  $M(\mathcal{P})$ , How do we construct an  $\epsilon$ -machine  $M(\mathcal{P}^U)$  that generates  $\mathcal{P}$ ’s atypical sequences at some energy density  $U \neq h_\mu$  or, as




FIG. 4. Space of all sequences  $\mathcal{A}^\infty$  partitioned into  $\Lambda_U$ —isoenergy density or equal probability-decay-rate bubbles—in which all sequences in the same  $\Lambda_U$  have the same energy density  $U$ . The typical set is one such bubble with energy equal to Shannon entropy rate:  $U = h_\mu$ . Another important class is the forbidden set, in which all sequences do not occur. The forbidden set can also be interpreted as the subset of sequences with infinite positive energy. By applying the map  $\mathcal{B}_\beta$  to the process and changing  $\beta$  continuously from  $-\infty$  to  $+\infty$  (excluding  $\beta = 0$ ) one can generate any rare class of interest  $\Lambda_U^{\mathcal{P}}$ .  $\beta \rightarrow -\infty$  corresponds to the most probable sequences with the largest energy density  $U_{\max}$ ,  $\beta = 1$  corresponds to the typical set and  $\beta \rightarrow +\infty$  corresponds to the least probable sequences with the smallest energy density  $U_{\min}$ .

we denoted it, the set  $\Lambda_U^{\mathcal{P}}$ ? Here, we answer this question by constructing a map  $\mathcal{B}_\beta : \mathcal{P} \rightarrow \mathcal{P}_\beta$  from the original process  $\mathcal{P}$  to a new one  $\mathcal{P}_\beta$ . The map is parametrized by  $\beta \in \mathbb{R}/\{0\}$  which indexes the rare set of interest. (We use  $\beta$  for convenience here, but it is related to  $U$  by a function introduced shortly.) Both processes  $\mathcal{P} = \{\mathcal{A}^\infty, \Sigma, \mathbb{P}(\cdot)\}$  and  $\mathcal{P}_\beta = \{\mathcal{A}^\infty, \Sigma, \mathbb{P}_\beta(\cdot)\}$  are defined on the same measurable sequence space. The measures differ, but their supports (allowed sequences) are the same. For simplicity we refer to  $\mathcal{B}_\beta$  as the  $\beta$ -map.

Assume we are given  $M(\mathcal{P}) = \{\mathcal{S}, \mathcal{A}, \{T^{(x)}, x \in \mathcal{A}\}\}$ . We showed that for every probability decay rate or energy density  $U$ , there exists a particular  $\beta$  such that  $M(\mathcal{P}_\beta)$  typically generates the words in  $\Lambda_{U,n}^{\mathcal{P}}$  for large  $n$  [37]. The  $\beta$ -map which establishes this is calculated by a construction that relates  $M(\mathcal{P})$  to  $M(\mathcal{P}_\beta) = \{\mathcal{S}, \mathcal{A}, \{\mathbf{S}_\beta^{(x)}, x \in \mathcal{A}\}\}$ —the HMM that generates  $\mathcal{P}_\beta$ :

1. For each  $x \in \mathcal{A}$ , construct a new matrix  $\mathbf{T}_\beta^{(x)}$  for which  $(\mathbf{T}_\beta^{(x)})_{ij} = (\mathbf{T}^{(x)})_{ij}^\beta$ .

<sup>3</sup>  $U$ , considered as a random variable, is sometimes called a *self process* [53].

2. Form the matrix  $\mathbf{T}_\beta = \sum_{x \in \mathcal{A}} T_\beta^{(x)}$ .
3. Calculate  $\mathbf{T}_\beta$ 's maximum eigenvalue  $\hat{\lambda}_\beta$  and corresponding right eigenvector  $\hat{\mathbf{r}}_\beta$ .
4. For each  $x \in \mathcal{A}$ , construct new matrices  $\mathbf{S}_\beta^{(x)}$  for which:

$$(\mathbf{S}_\beta^{(x)})_{ij} = \frac{(\mathbf{T}_\beta^{(x)})_{ij} (\hat{\mathbf{r}}_\beta)_j}{\hat{\lambda}_\beta (\hat{\mathbf{r}}_\beta)_i}.$$

Having constructed the new process  $\mathcal{P}_\beta$  by introducing its generator, we use the latter to produce some rare set of interest  $\Lambda_{U,n}^\mathcal{P}$ .

**Theorem 1.** *In the limit  $n \rightarrow \infty$ , within the new process  $\mathcal{P}_\beta$  the probability of generating realizations from the set  $\Lambda_{U,n}^\mathcal{P}$  converges to one:*

$$\lim_{n \rightarrow \infty} \mathbb{P}_\beta(\Lambda_{U,n}^\mathcal{P}) = 1,$$

where:

$$U = \beta^{-1} (h_\mu(\mathcal{P}_\beta) - \log_2 \hat{\lambda}_\beta). \quad (4)$$

In addition, in the same limit the process  $\mathcal{P}_\beta$  assigns equal energy densities over all the members of the set  $\Lambda_{U,n}^\mathcal{P}$ .

**Proof.** See Ref. [37].

As a result, for large  $n$  the process  $\mathcal{P}_\beta$  typically generates the set  $\Lambda_{U,n}^\mathcal{P}$  with the specified energy  $U$ . The process  $\mathcal{P}_\beta$  is sometimes called the *auxiliary*, *driven*, or *effective* process [54–56]. Examining the form of the energy in Eq. (4), one sees that there is a one-to-one relationship between  $\beta$  and  $U$ . And so, we can equivalently denote the process  $\mathcal{P}_\beta$  by  $\mathcal{P}^U$ . More formally, every word in  $\Lambda_U^\mathcal{P}$  with probability measure one is in the typical set of process  $\mathcal{P}_\beta$ .

The  $\beta$ -map construction guarantees that the HMMs  $M(\mathcal{P})$  and  $M(\mathcal{P}_\beta)$  have the same states and transition topology:  $(\mathbf{T}_\beta^{(x)})_{ij} \neq 0 \iff (\mathbf{S}_\beta^{(x)})_{ij} \neq 0$ . The only difference is in their transition probabilities.  $M(\mathcal{P}_\beta)$  is not necessarily an  $\epsilon$ -machine—the most memory-efficient classical algorithm that generates the process. Typically, though,  $M(\mathcal{P}_\beta)$  is an  $\epsilon$ -machine. Moreover, there are only finitely many  $\beta$ s for which it is not. (More detailed development along these lines will appear in a sequel.)

## VII. QUANTUM AND CLASSICAL MEMORY FOR BIASED SAMPLING

Having introduced the necessary background to compare classical versus quantum models and to appreciate typical versus rare realizations, we are ready to investigate the




FIG. 5.  $\epsilon$ -Machine generator of the Perturbed Coins Process. Edges are labeled with conditional transition probabilities and emitted symbols. For example, for the self-loop on state  $A$ ,  $p|0$  indicates the transition is taken with probability  $\mathbb{P}(0|A) = p$  and the symbol  $0$  is emitted.

quantum memory advantage when generating a given process' rare events.

The last section concluded that the memory required by the classical algorithm to generate  $\mathcal{P}$ 's rare sequences with energy density  $U$  is:

$$C_\mu(\mathcal{P}_\beta) = C_\mu(\mathcal{B}_\beta(\mathcal{P})),$$

where  $U$  and  $\beta$  are related via  $U = \beta^{-1} (h_\mu(\mathcal{P}_\beta) - \log_2 \hat{\lambda}_\beta)$ . Similarly, the memory required by the quantum algorithm to generate the rare class with energy density  $U$  is:

$$C_q(\mathcal{B}_\beta(\mathcal{P})).$$

For simplicity, we denote these two quantities by  $C_\mu(\beta) \equiv C_\mu(\mathcal{P}_\beta)$  and  $C_q(\beta) \equiv C_q(\mathcal{P}_\beta)$ .

The following analyzes the advantage for three example processes—two in the *single-process* setting and one in the *multi-process* setting. For the first two, we consider particular given stochastic processes and study the advantage (memory ratio) as the metric of memory efficiency. In the third example, we consider a series of stochastic processes labeled by their size  $N$  and compare how both classical and quantum memories scale with  $N$ . The ratio of scaling then is the metric for memory efficiency.

### A. Quantum Memory Advantage for a Simple Markov Process

Let's start in the single-process setting in which an individual stochastic process is given. Consider the case where we have two biased coins, call them  $A$  and  $B$ , and each has a different bias  $p$  and  $1 - q$  both for Heads (symbol  $0$ ), respectively. When we flip a coin, if the result is Heads, then on the next flip we choose coin  $A$ . If the result is Tails, we choose coin  $B$ . Flipping the coins over and over again results in a process  $\mathcal{P}^{\text{pc}}$  called the *Perturbed Coins Process* [28]. Figure 5 shows the process'  $\epsilon$ -machine generator  $M(\mathcal{P}^{\text{pc}})$ , where  $\mathcal{S} = \{A, B\}$  and  $\mathcal{A} = \{0, 1\}$ .



One can also generate this process with a quantum machine  $Q(\mathcal{P}^{\text{pc}})$ . Using the construction introduced in

Sec. IV, it has two Kraus operators corresponding to symbols 0 and 1:

$$K_0 = \frac{1}{d} \begin{bmatrix} \sqrt{q(1-q)p} - p\sqrt{1-p} & p\sqrt{p} - (1-q)\sqrt{p} \\ \sqrt{q(1-q)(1-p)} - (1-p)\sqrt{p} & p\sqrt{1-p} - (1-q)\sqrt{1-p} \end{bmatrix}$$

and:

$$K_1 = \frac{1}{d} \begin{bmatrix} q\sqrt{(1-q)} - (1-p)\sqrt{1-q} & \sqrt{p(1-p)(1-q)} - (1-q)\sqrt{q} \\ q\sqrt{q} - (1-p)\sqrt{q} & \sqrt{p(1-p)q} - q\sqrt{1-q} \end{bmatrix},$$




FIG. 6. Classical memory  $C_\mu(\beta)$  and quantum memory  $C_q(\beta)$  versus  $\beta$  for biased sampling of Perturbed Coins Process' rare sequence classes: See Fig. 5, with  $p = 0.6$  and  $q = 0.8$ . As the inset shows, for large  $\beta$  both classical and quantum memories decay exponentially with  $\beta$ , but the quantum memory decays faster.

where  $d = \sqrt{pq} + \sqrt{(1-p)(1-q)}$ . For its stationary state distribution we have:

$$\rho_s = \frac{1}{2-p-q} \begin{bmatrix} 1-p & \alpha \\ \alpha & 1-q \end{bmatrix},$$

where  $\alpha = (1-q)\sqrt{p(1-p)} + (1-p)\sqrt{q(1-q)}$ .

Figure 6 shows the classical and quantum memory costs to generate rare realizations:  $C_\mu(\beta)$  and  $C_q(\beta)$  versus  $\beta$  for different  $\beta$ -classes. Surprisingly, the two costs exhibit completely different behaviors. For example,  $\lim_{\beta \rightarrow 0} C_q = 0$ , while  $\lim_{\beta \rightarrow 0} C_\mu = 1$ . More interestingly, as the inset demonstrates, even though both  $C_\mu(\beta)$  and  $C_q(\beta)$  vanish exponentially fast, in the limit of  $\beta \rightarrow \infty$   $C_q(\beta)$  goes to zero noticeably faster.

We define the *memory efficiency*  $\eta(\beta)$  of biased sampling

in the single-process setting as the ratio of classical to quantum memory:

$$\eta(\beta) \equiv \frac{C_\mu(\beta)}{C_q(\beta)}.$$

Figure 7 graphs  $\eta(\beta)$ , revealing how it divides into three distinct scaling regimes.

First, for small  $|\beta|$  the memory ratio scales as  $\mathcal{O}(\beta^{-2})$ . Second, for large positive  $\beta$  the memory ratio scales exponentially  $\mathcal{O}(\exp(c\beta))$  as one increases  $\beta$ , where  $c$  is a function of  $p$  and  $q$ . Third, for large negative  $\beta$ , there is no quantum advantage. Since we are analyzing finite-state processes, this regime appears and is the analog of population inversion. And so, formally there are  $\beta$ -class events with negative  $\beta$ .

Such is the quantum advantage for the Perturbed Coins Process at  $p = 0.6$  and  $q = 0.8$ . The features exhibited—the different scaling regimes—are generic for any  $p > 1-q$ , though. Moreover, for Perturbed Coins Processes with  $p < 1-q$ , the positive and negative low temperature behaviors switch.

## B. Quantum Memory Advantage for Next-Nearest-Neighbor Spin Systems

Again, consider the *single-process* setting in which an individual stochastic process is given. Let us, however, analyze the quantum advantage in a more familiar physical system. Consider a general one-dimensional ferromagnetic next-nearest-neighbor Ising spin- $\frac{1}{2}$  chain [57, 58] defined by the Hamiltonian:

$$\mathcal{H} = - \sum_i (s_i s_{i+1} + \frac{1}{4} s_i s_{i+2}), \quad (5)$$

in contact with thermal bath at temperature  $k_B T = 1$ . The spin  $s_i$  at site  $i$  takes on values  $\{+1, -1\}$ .




FIG. 7. The classical to quantum memory ratio for generating rare realizations of the Perturbed Coins Process with  $p = 0.6$  and  $q = 0.8$  when employing its  $q$ -machine instead of its (classical)  $\epsilon$ -machine. Three different advantages occur: (i) near  $\beta = 0$  the ratio scales as  $\mathcal{O}(\beta^{-2})$ , (ii) for large positive  $\beta$ , it scales exponentially with  $\beta$ ,  $\mathcal{O}(\exp(f(q, p)\beta))$ , and (iii) no advantage at large negative  $\beta$ .

After thermalizing, a spin configuration at one instant of time may be thought of as having been generated left-to-right (or, equivalently, right-to-left) along the lattice. The probability distribution over these spatial-translation invariant configurations defines a stationary stochastic process. Reference [59, Eqs. (84) – (91)] showed that for any finite and nonzero temperature  $T$ , this process has Markov order  $R = 2$ . More to the point, the  $\epsilon$ -machine that generates this process has four causal states and those states are in one-to-one correspondence with the set of length-2 spin configurations.

Figure 8 displays the parametrized  $\epsilon$ -machine that generates this family of spin-configuration processes. To simulate the process, the generator need only remember the last two spins generated. This means the  $\epsilon$ -machine has four states,  $\downarrow\downarrow$ ,  $\downarrow\uparrow$ ,  $\uparrow\downarrow$ , and  $\uparrow\uparrow$ . If the last two observed spins are  $\uparrow\uparrow$  for example, then the current state is  $\uparrow\uparrow$ . We denote the probability of generating a  $\downarrow$  spin given that the previous two spins were  $\uparrow\uparrow$  by  $p_{\uparrow\uparrow}^{\downarrow}$ . If the generator is in the  $\uparrow\uparrow$  state and generates a  $\downarrow$  spin, then the generator state changes to  $\uparrow\downarrow$ .

To determine the  $\epsilon$ -machine transition probabilities  $\{T^{(x)}\}_{x \in \mathcal{A}}$ , we first compute the transfer matrix  $V$  for the Hamiltonian of Eq. (5) at temperature  $T$  and then extract conditional probabilities, following Ref. [59] and Ref. [41]’s appendix.

What are the classical and quantum memory costs for bias sampling of the rare spin-configuration class with decay rate  $U$ , as defined in Eq. (3)? First, note that  $U$  is not a configuration’s actual energy density. If we assume




FIG. 8.  $\epsilon$ -Machine that generates the spin configurations occurring in the one-dimensional ferromagnetic next-nearest-neighbor Ising spin chain with the Hamiltonian in Eq. (5).

the system is in thermal equilibrium and thus exhibits a Boltzmann distribution over configurations, then  $U$  and  $E$  are related via:

$$U = \frac{\log_2(e)}{k_B T} (E - \mathcal{F}(T)),$$

where:

$$\mathcal{F}(T) = -k_B T \lim_{n \rightarrow \infty} \frac{1}{n} \ln \left( \sum_{\{w \in \mathcal{A}^n\}} e^{-\frac{nE(w)}{k_B T}} \right).$$

This simply tells us that if a stochastic process describes thermalized configurations of a physical system with some given Hamiltonian, then every rare-event bubble in Fig. 4 can be labeled either with  $\beta$ ,  $U$ , or  $E$ . Moreover, there is a one-to-one mapping between every such variable pair.

Figure 9 plots  $\eta(U)$  versus  $U$ —the memory ratio for generating rare configurations with decay rate  $U$ . To calculate  $\eta(U)$  for a given process  $\mathcal{P}$ , first we determine the process’ classical generator  $M(\mathcal{P})$  using the method introduced in Ref. [39]. Second, for every  $\beta \in \mathbb{R}/\{0\}$ , using the map introduced in Sec. VI, we find the new classical generator  $M(\mathcal{P}_\beta)$ . Third, using the construction introduced in Sec. III, we find  $Q(\mathcal{P}_\beta)$ . Fourth, using Thm. 1 we find the corresponding  $U$  for the chosen  $\beta$ . Using these results gives  $\eta(U(\beta)) = C_\mu(\beta)/C_q(\beta)$ . By varying  $\beta$  in the range  $\mathbb{R}/\{0\}$  we cover all the energy density  $Us$ . Practically, to calculate  $\eta(U)$  in Fig. (9), we chose 2000  $\beta \in [-10, 7.5]$ .

As pointed out earlier,  $\beta = 1$  always corresponds to the process itself. And, one obtains its typical sequences. As one sees in Fig. 9, the memory ratio  $\eta(1) < 2$ . This simply means that, though there is a quantum advantage generating typical sequences, it is not that notable. However, the figure highlights four other interesting regimes.

First, there is the  $\beta \rightarrow \infty$  limit corresponding to the rare class with minimum decay rate equal to  $U_{\min} = -\log_2(p_{\uparrow\uparrow}^{\downarrow}) = -\log_2(p_{\uparrow\uparrow}^{\uparrow})$ . From Eq. (5) it is easy to




FIG. 9. Classical to quantum memory ratio for biased sampling of Ising spin configurations:  $\eta(U)$  versus decay rate  $U$  for bias sampling of equal-energy density spin configurations. Vertical lines locate  $\beta$ s corresponding to particular  $U$ s. Note the memory ratio  $\eta(U)$  diverges at  $U = u_0 \approx 1.878$  corresponding to  $\beta = 0$ . Note that the quantum memory advantage is not limited to  $\beta = 0$ , but occurs in a large neighborhood around it. Quantitatively, for any arbitrary large number  $r$  there exist a number  $\epsilon$  for which the rare class  $\beta_0 = \epsilon$  has the memory ratio  $\eta(\beta_0) > r$ .

see that this rare bubble only has two configurations as members: all up-spins or all down-spins. Let us consider finite but large  $\beta \gg 1$  that corresponds to the rare class with a low energy density close to  $U_{\min}$ . Figure 10(top-left) shows a general  $\epsilon$ -machine for this process. Low color intensity for both edges and states means that the process rarely visits them during generation. This means, in turn, that a typical realization consists of large blocks of all up-spins and all down-spins. These large blocks are joined by small segments.

Second, there is the  $\beta \rightarrow -\infty$  limit that corresponds to the rare class with maximum decay rate equal to  $U_{\max} = -\frac{1}{2} \log_2(p_{\downarrow\uparrow} p_{\uparrow\downarrow})$ . From Eq. (5) it is easy to see that this rare bubble only has one configuration as a member: a periodic repetition of spin down and spin up. Consider finite  $\beta \ll 1$  corresponding to a rare class with a high energy density close to  $U_{\max}$ . Figure 10(top-right) shows the general  $\epsilon$ -machine for the associated process. The typical configuration consists of large blocks tiled with spin-up and spin-down pairs that are connected by other short segments.

Third, there is the  $\beta \rightarrow 0^+$  limit. In this limit we expect to see completely random spin-up/spin-down configurations. Figure 10(bottom-right) shows the  $\epsilon$ -machine for this class labeled with nonzero small  $\beta$ . The transition probability for the edges labeled  $+$  is  $1/2 + \epsilon$  and for the edges labeled  $-$  is  $1/2 - \epsilon$ , where  $\epsilon$  is a small positive number. As one




FIG. 10. Classical generators of four important rare classes: (Top-left) Negative zero-temperature limit. (Top-right) positive zero temperature limit. (Bottom-left) Negative infinite temperature limit. (Bottom-right) positive temperature limit. Gray edges and states denotes them being rarely visited.

can see, even though each transition probability is close to one-half, the self-loops are slightly favored.

Fourth and finally, there is the  $\beta \rightarrow 0^-$  limit. The generator here, Figure 10(bottom-left), is similar to that at positive infinite temperature, except that the edge-sign labels are reversed. This means that the self-loops are slightly less favored.

Remarkably, the memory ratio  $\eta(U)$  diverges at  $U = u_0 \approx 1.878$ , where  $u_0 = \lim_{\beta \rightarrow 0} U$ —that is, at both the positive and negative high-temperature limit. Moreover, the memory ratio  $\eta(U)$  diverges as  $(U - u_0)^{-2}$  in both limits.

### C. Extreme Quantum Memory Advantage for $N$ -Nearest-Neighbor Spin Systems

Now, consider the *multi-process* setting in which we specify a series of stochastic processes labeled by an integer  $N$  that determines the size of each. In this, efficiency is defined by how the memory scales in  $N$  for the quantum algorithm compared to the classical.

We explore a general one-dimensional ferromagnetic  $N$ -nearest-neighbors Ising spin- $1/2$  chain defined by the Hamiltonian:

$$\mathcal{H}_N = - \sum_i \sum_{k=1}^N \frac{1}{k^\delta} S_i S_{i+k},$$

in contact with a thermal bath at temperature  $k_B T = 1$  and for which there is monopole-dipole coupling ( $\delta = 2$ ).




FIG. 11. Classical memory  $C_\mu(N, \beta)$  (solid lines) and quantum memory  $C_q(N, \beta)$  (dashed lines) required for generating process  $\mathcal{P}(N, \beta)$  for interaction ranges  $N = 1, \dots, 4$ , a range of  $\beta \in [-5, 3]$ , and  $\delta = 2$ . At both limits  $\beta \rightarrow 0^+$  and  $\beta \rightarrow 0^-$ ,  $C_\mu(N, \beta)$  scales linearly with  $N$  while  $C_q(N, \beta)$  vanishes. For any finite  $\beta$ , for sufficiently large  $N$ ,  $C_\mu(N, \beta)$  is an increasing function of  $N$ , while  $C_q(N, \beta)$  is bounded above by 1.

As in the nearest-neighbor spin system, after thermalizing the probability over configurations at one instant of time defines a spatially-stationary stochastic process. We denote the process generated by this Hamiltonian by  $\mathcal{P}(N)$ .  $\mathcal{P}(N)$  has Markov order  $R = N$  and also cryptic order  $\chi = N$ . More to the point, the  $\epsilon$ -machine that generates this process has  $2^N$  causal states and those states are in one-to-one correspondence with the set of length- $N$  spin configurations. To determine the  $\epsilon$ -machine transition probabilities  $\{T^{(x)}\}_{x \in \mathcal{A}}$ , one can use Ref. [59] and Ref. [41]’s appendix.

Let  $\mathcal{P}(N, \beta)$  denote the process that typically generates the rare  $\beta$ -class of process  $\mathcal{P}(N)$ . Now, for an arbitrary fixed  $\beta$ , one can ask how the required classical memory  $C_\mu(N, \beta)$  and quantum memory  $C_q(N, \beta)$  for generating  $\mathcal{P}(N, \beta)$  scales with  $N$ .

Figure 11 shows  $C_\mu(N, \beta)$  and  $C_q(N, \beta)$  versus  $\beta$  for different  $N$ s.  $C_\mu(\cdot)$ s are plotted by solid-line and  $C_q$ s by dashed-line. To make them distinguishable, curves at different  $N$ s are displayed with different colors. Interestingly, in both  $\beta \rightarrow 0^+$  and  $\beta \rightarrow 0^-$  limits,  $C_\mu(N, \beta)$

scales linearly with  $N$ , while  $C_q(N, \beta)$  goes to zero. More importantly, one can also check that for any finite nonzero  $\beta$  and sufficiently large  $N$ ,  $C_\mu(N, \beta)$  is an increasing function of  $N$ . Surprisingly, it can be shown that for any nonzero  $\beta$  and any  $N$ ,  $C_q(N, \beta)$  is bounded above by 1. The result is extreme quantum memory advantage for rare-event sampling of this series of stochastic processes.

## VIII. CONCLUSIONS

We introduced a new quantum algorithm for simultaneous sampling rare events in classical stochastic processes. We showed that it confers a significant memory advantage when compared to the best known classical algorithm.

We explored two settings: single-process and multi-process sampling. For single processes an individual stochastic process is given and memory efficiency is defined as the ratio of memory required by the classical algorithm compared to that by the quantum one. For two example systems, we showed that for any large real number  $r$  there exist infinite classes of rare events for which the classical-quantum memory ratio for sampling is larger than  $r$ . In the multi-process setting, a series of stochastic processes each labeled by an integer size  $N$  is given. There, the memory efficiency is defined by how required memory scales in  $N$  for the classical algorithm compared to the quantum algorithm. In this setting we demonstrated an extreme quantum memory advantage in which the classical memory grows with  $N$  unboundedly, but the quantum memory is bounded.

## ACKNOWLEDGMENTS

The authors thank Leonardo Duenas-Osorio for stimulating discussions on risk estimation in networked infrastructure and the referees for valuable suggestions. JPC thanks the Santa Fe Institute for its hospitality during visits as an External Faculty member. This material is based upon work supported by, or in part by, the John Templeton Foundation grant 52095, the Foundational Questions Institute grant FQXi-RFP-1609, and the U. S. Army Research Laboratory and the U. S. Army Research Office under contracts W911NF-13-1-0390 and W911NF-13-1-0340.

[1] A. R. Leach. *Molecular modelling: Principles and applications*. Pearson Education, Boston, Massachusetts, 2001.

[2] D. Frenkel and B. Smit. *Understanding Molecular Simulation: From Algorithms to Applications*. Academic Press, New York, second edition, 2007. 1

- [3] H. D. Miller. A convexity property in the theory of random variables defined on a finite Markov chain. *An. Math. Stat.*, 32(4):1260–1270, 1961. 1
- [4] K. Young and J. P. Crutchfield. Fluctuation spectroscopy. *Chaos, Solitons, and Fractals*, 4:5 – 39, 1994. 1
- [5] V. Lecomte, C. Appert-Rolland, and F. van Wijland. Chaotic properties of systems with Markov dynamics. *Phys. Rev. Lett.*, 95(1):010601, 2005. 1
- [6] V. Lecomte, C. Appert-Rolland, and F. Van Wijland. Thermodynamic formalism for systems with markov dynamics. *J. Stat. Physics*, 127(1):51–106, 2007.
- [7] R. Chetrite and H. Touchette. Nonequilibrium micro-canonical and canonical ensembles and their equivalence. *Phys. Rev. Lett.*, 111(12):120601, 2013. 1
- [8] S. R. S. Varadhan. *Large deviations and applications*. SIAM, Philadelphia, Pennsylvania, 1984. 1
- [9] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland. First-order dynamical phase transition in models of glasses: an approach based on ensembles of histories. *J. Phys. A: Math. Theo.*, 42(7):075007, 2009. 1
- [10] L. O. Hedges, R. L. Jack, J. P. Garrahan, and D. Chandler. Dynamic order-disorder in atomistic models of structural glass formers. *Science*, 323(5919):1309–1313, 2009. 1
- [11] J. Van Campenhout and T. M. Cover. Maximum entropy and conditional probability. *IEEE Trans. Info. Th.*, 27(4):483–489, 1981. 1
- [12] I. Csiszár. Sanov property, generalized I-projection and a conditional limit theorem. *Ann. Prob.*, 12(3):768–793, 1984. 1
- [13] G. M. Torrie and J. P. Valleau. Nonphysical sampling distributions in monte carlo free-energy estimation: Umbrella sampling. *J. Comp. Physics*, 23(2):187–199 1
- [14] S. Kumar, J. M. Rosenberg, D. Bouzida, R. H. Swendsen, and P. A. Kollman. The weighted histogram analysis method for free energy calculations on biomolecules. I. The method. *J. Comp. Chemistry*, 13(8):1011–1021 1
- [15] F. Wang and D. P. Landau. Efficient, multiple-range random walk algorithm to calculate the density of states. *Phys. Rev. Let.*, 86(10):2050, 2001. 1
- [16] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, Providence, Rhode Island, 2009. 1
- [17] J. R. Norris. *Markov Chains*, volume 2. Cambridge University Press, Cambridge, United Kingdom, 1998. 1
- [18] D. R. Upper. *Theory and Algorithms for Hidden Markov Models and Generalized Hidden Markov Models*. PhD thesis, University of California, Berkeley, 1997. Published by University Microfilms Intl, Ann Arbor, Michigan. 1, 2
- [19] L. R. Rabiner and B. H. Juang. An introduction to hidden Markov models. *IEEE ASSP Magazine*, January, 1986.
- [20] L. R. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE*, 77(2):257–286 1
- [21] J. P. Crutchfield and D. P. Feldman. Regularities unseen, randomness observed: Levels of entropy convergence. *CHAOS*, 13(1):25–54, 2003. 1
- [22] J. P. Crutchfield. Between order and chaos. *Nature Physics*, 8(1):17–24, 2012. 1, 4
- [23] J. P. Crutchfield and K. Young. Inferring statistical complexity. *Phys. Rev. Let.*, 63:105–108, 1989. 1, 4
- [24] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Review*, 41(2):303–332 2
- [25] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, pages 212–219 2
- [26] D. S. Abrams and S. Lloyd. Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors. *Phys. Rev. Let.*, 83(24):5162, 1999. 2
- [27] A. W. Harrow, A. Hassidim, and S. Lloyd. Quantum algorithm for linear systems of equations. *Phys. Rev. Let.*, 103(15):150502, 2009. 2
- [28] M. Gu, K. Wiesner, E. Rieper, and V. Vedral. Quantum mechanics can reduce the complexity of classical models. *Nature Comm.*, 3:762, 2012. 2, 4, 8
- [29] J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield. Occam’s quantum strop: Synchronizing and compressing classical cryptic processes via a quantum channel. *Sci. Reports*, 6, 2016. 2, 4
- [30] P. M. Riechers, J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield. Minimized state complexity of quantum-encoded cryptic processes. *Phys. Rev. A*, 93(5):052317, 2016. 2, 4, 5
- [31] M. S. Palsson, M. Gu, J. Ho, H. M. Wiseman, and G. J. Pryde. Experimental quantum processing enhancement in modelling stochastic processes. *Science Advances*, 3(2):1601302, 2017. 2, 4
- [32] S. Aaronson and A. Arkhipov. The computational complexity of linear optics. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 333–342. ACM, 2011. 2
- [33] A. Ambainis, L. J. Schulman, A. Ta-Shma, U. Vazirani, and A. Wigderson. The quantum communication complexity of sampling. *SIAM Journal on Computing*, 32(6):1570–15857, 2003. 2
- [34] M. Szegedy. Quantum speed-up of Markovchain based algorithms. In *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on*, pages 32–41. IEEE, 2004. 2
- [35] P. Wocjan and A. Abeyesinghe. Speedup via quantum sampling. *Phys. Rev. A*, 78(4):042336, 2008. 2
- [36] N. F. Travers. Exponential bounds for convergence of entropy rate approximations in hidden Markov models satisfying a path-mergeability condition. *Stochastic Proc. Appln.*, 124(12):4149–4170, 2014. 2
- [37] C. Aghamohammadi and J. P. Crutchfield. Minimum memory for generating rare events. *Phys. Rev. E*, 95(3):032101, 2017. 2, 7, 8
- [38] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, New York, second edition, 2006. 4, 6
- [39] C. R. Shalizi and J. P. Crutchfield. Computational mechanics: Pattern and prediction, structure and simplicity.

- J. Stat. Phys.*, 104:817–879, 2001. 4, 10
- [40] G. Han and B. Marcus. Analyticity of entropy rate of hidden Markov chains. *IEEE Trans. Info. Th.*, 52(12):5251–5266, 2006. 4
- [41] C. Aghamohamdi, J. R. Mahoney, and J. P. Crutchfield. Extreme quantum advantage when simulating strongly coupled classical systems. *Sci. Reports*, 7(6735):1–11, 2017. 4, 10, 12
- [42] A. J. P. Garner, Q. Liu, J. Thompson, V. Vedral, and M. Gu. Provably unbounded memory advantage in stochastic simulation using quantum mechanics. *New J. Physics*, 2017. 4
- [43] C. Aghamohammadi, J. R. Mahoney, and J. P. Crutchfield. The ambiguity of simplicity in quantum and classical simulation. *Phys. Lett. A*, 381(14):1223–1227, 2017. 4
- [44] J. Preskill. *Lecture notes for physics 229: Quantum information and computation*, volume 16. California Institute of Technology, Pasadena, California, 1998. 4, 5
- [45] F. C. Binder, J. Thompson, and M. Gu. A practical, unitary simulator for non-Markovian complex processes. *arXiv:1709.02375*, 2017. 5
- [46] R. Durrett. *Probability: theory and examples*. Cambridge University Press, Cambridge, United Kingdom, 2010. 6
- [47] S. Kullback. *Information Theory and Statistics*. Dover, New York, 1968. 6
- [48] R. W. Yeung. *Information Theory and Network Coding*. Springer, New York, 2008. 6
- [49] C. E. Shannon. A mathematical theory of communication. *Bell Sys. Tech. J.*, 27:379–423, 623–656, 1948. 6
- [50] B. McMillan. The basic theorems of information theory. *Ann. Math. Stat.*, 24:196–219, 1953.
- [51] L. Breiman. The individual ergodic theorem of information theory. *Ann. Math. Statistics*, 28(3):809–811, 1957. 6
- [52] L. Boltzmann. *Lectures on gas theory*. University of California Press, Berkeley, California, 1964. 7
- [53] H. Touchette. The large deviation approach to statistical mechanics. *Physics Reports*, 478:1–69, 2009. 7
- [54] R. L. Jack and P. Sollich. Large deviations and ensembles of trajectories in stochastic models. *Prog. Theo. Physics Suppl.*, 184:304–317, 2010. 8
- [55] J. P. Garrahan and I. Lesanovsky. Thermodynamics of quantum jump trajectories. *Phys. Rev. Lett.*, 104(16):160601, 2010.
- [56] R. Chetrite and H. Touchette. Nonequilibrium Markov processes conditioned on large deviations. *Annales Henri Poincaré*, 16(9):2005–2057, 2015. 8
- [57] R. J. Baxter. *Exactly solved models in statistical mechanics*. Academic Press, New York, New York, 2007. 9
- [58] A. Aghamohammadi, C. Aghamohammadi, and M. Khorrami. Externally driven one-dimensional Ising model. *J. Stat. Mech.*, 2012(02):P02004 9
- [59] D. P. Feldman and J. P. Crutchfield. Discovering non-critical organization: Statistical mechanical, information theoretic, and computational views of patterns in simple one-dimensional spin systems. Technical Report 98-04-026, Santa Fe Institute, 1998. 10, 12