




**Robust inference of memory structure for efficient quantum modeling of stochastic processes**Matthew Ho <sup>1,2,\*</sup> Mile Gu <sup>1,2,3,†</sup> and Thomas J. Elliott <sup>2,1,‡</sup><sup>1</sup>*School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore 637371, Singapore*<sup>2</sup>*Complexity Institute, Nanyang Technological University, Singapore 637335, Singapore*<sup>3</sup>*Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, Singapore 117543, Singapore*

(Received 7 November 2019; revised manuscript received 16 January 2020; accepted 20 February 2020; published 19 March 2020)

A growing body of work has established the modeling of stochastic processes as a promising area of application for quantum technologies; it has been shown that quantum models are able to replicate the future statistics of a stochastic process while retaining less information about the past than any classical model must, even for a purely classical process. Such memory-efficient models open a potential future route to study complex systems in greater detail than ever before and suggest profound consequences for our notions of structure in their dynamics. Yet, to date methods for constructing these quantum models are based on having a prior knowledge of the optimal classical model. Here, we introduce a protocol for blind inference of the memory structure of quantum models—tailored to take advantage of quantum features—direct from time-series data, in the process highlighting the robustness of their structure to noise. This in turn provides a way to construct memory-efficient quantum models of stochastic processes while circumventing certain drawbacks that manifest solely as a result of classical information processing in classical inference protocols.

DOI: [10.1103/PhysRevA.101.032327](https://doi.org/10.1103/PhysRevA.101.032327)**I. INTRODUCTION**

Complex processes are prevalent throughout the world, taking the form of natural processes such as the weather [1,2] and DNA sequences [3], as well as artificial processes like the stock market [4] and traffic [5]. We construct models of these processes in order to better understand their structure and predict their behavior. Within complexity science, the field of computational mechanics [6–8] offers a systematic approach to understanding the intrinsic computation of a process by identifying the causal links between its past and future, and has been used to study a diverse set of dynamics such as deterministic chaos in the logistic map [6,8], cellular automata [9], the dripping faucet experiment [10], stock markets [11], and neural spike trains [12]. A key component of the approach are so-called  $\varepsilon$ -machines, which as a valuable byproduct represent the most parsimonious causal model of a process.

In recent decades, the prospect of using quantum effects in information processing has emerged, promising advantages for a range of applications in terms of algorithmic speedups [13], secure communication [14], and beyond. Stochastic modeling is no exception to this, and a growing body of work has established that when information is encoded into a quantum memory, causal models of a stochastic process can be designed that function while retaining less information about the past than is classically possible [15–22]. This quantum memory advantage can grow

unbounded [19,21,23–26] and has been verified experimentally [27–29]. Like its classical counterpart, the amount of information stored within these quantum models has been suggested as a measure of structural complexity in stochastic dynamics [30–33].

Currently, systematic approaches to constructing such quantum models are predicated on having a prior exact statistical description of the process, or knowledge of its  $\varepsilon$ -machine. As a result, to apply these tools to real-world systems we must first use classical inference protocols to construct an  $\varepsilon$ -machine [6,34,35] and then use this as a basis to construct a corresponding quantum model. It is desirable to instead have a model inference protocol to directly go from data to the quantum model, avoiding any extra computational overhead associated with also determining the classical model. In this vein, here we introduce such a protocol for directly inferring the memory structure of a quantum model of a stochastic process—which we show is robust to statistical noise. The protocol is tailored specifically for quantum models, taking advantage of certain of their features that allow some approximations that must be made in classical information processing to be avoided. Figure 1 provides a schematic of our motivation.

The layout of this article is as follows. In Sec. II we outline the general framework of stochastic processes and computational mechanics as is relevant here, as well as the more efficient quantum models. Section III provides the core of our results, introducing the inference protocol, showing its robustness to statistical fluctuations and justifying its accuracy. The efficacy of our inference protocol is then demonstrated in practice with two toy processes in Sec. IV. Finally, we conclude in Sec. V and discuss some future directions.

\*hosh0021@e.ntu.edu.sg

†mgu@quantumcomplexity.org

‡physics@tjelliott.net

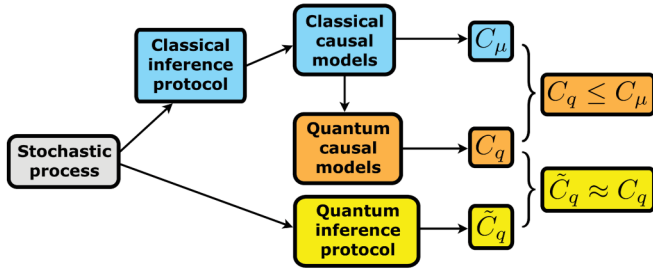


FIG. 1. Schematic context of our work. Quantum information processing has been shown to provide a more memory-efficient route to stochastic modeling than classically possible. However, current approaches to constructing quantum models first require classical models to be inferred; here we introduce a blind inference protocol for going straight from raw data to quantum structure. The quantities  $C_\mu$ ,  $C_q$ , and  $\tilde{C}_q$  represent the information stored by the minimal classical, quantum, and inferred quantum causal models, respectively.

## II. FRAMEWORK

### A. Stochastic processes

We consider discrete-time stochastic processes represented by a bi-infinite probabilistic string of outcomes  $\overleftarrow{X} \equiv X_{-\infty:\infty} = \dots X_{-2}X_{-1}X_0X_1X_2\dots$ , where  $X_t$  are random variables that take on values  $x_t$  drawn from an alphabet  $\mathcal{A}$ , and the subscript  $t$  represents the time step. Consecutive strings  $X_{0:t} := X_0X_1\dots X_{t-1}$  are called *words*, with the left index inclusive and the right exclusive. We consider stationary processes, such that  $P(X_{0:L}) = P(X_{t:t+L}) \forall t, L \in \mathbb{Z}$ . We partition the process into (semi-infinite) *pasts* and *futures*, denoted as  $\overleftarrow{x} \equiv x_{-\infty:0}$  and  $\overrightarrow{x} \equiv x_{0:\infty}$ , respectively, where  $t = 0$  is taken to be the present.

The *Markov order* is an important property of a process that defines an effective history length; a process is said to have *Markov order*  $R$  if  $R$  is the smallest value such that  $P(X_0|\overrightarrow{X}) = P(X_0|X_{-R:0})$  is satisfied [36]. That is, it is the smallest block length of the most recent past that provides a sufficient statistic of the future. When  $R = 1$  the process is said to be Markovian.

### B. Models

*a. Computational mechanics.* Computational mechanics [6–8] provides a formal statistical framework for identifying and analyzing structure in complex processes. We outline key elements of computational mechanics here, providing further detail in Appendix A. Its modus operandi involves the minimal causal representation<sup>1</sup> of stochastic processes, which may be determined by a systematic clustering of *pasts*. Specifically, the *causal states* of a process are a set of equivalence classes on the *pasts*, defined according to the relation

$$\overleftarrow{x} \sim_\varepsilon \overleftarrow{x}' \Leftrightarrow P(\overrightarrow{X} | \overleftarrow{X} = \overleftarrow{x}) = P(\overrightarrow{X} | \overleftarrow{X} = \overleftarrow{x}'); \quad (1)$$

<sup>1</sup>Here, by causal we mean that the representation stores no information about the future of the process that could not be deduced from past observations.

that is, two *pasts* belong to the same causal state if and only if (iff) they give rise to statistically identical futures. We label the causal states as  $s_j \in \mathcal{S}$ .

Because of the deterministic assignment of *pasts* to causal states, it can be seen that transitions between causal states are also deterministic conditional on the output symbol; that is, given a *past*  $\overleftarrow{x} \in s_j$ , upon emission of the next symbol  $x$  the new *past*  $\overleftarrow{x}x$  must belong to causal state  $s_{\lambda(x,j)}$ , where  $\lambda(x,j)$  is a deterministic update function. This deterministic transition structure is sometimes referred to as *unifilarity* and allows us to represent the process as a deterministic edge-emitting hidden Markov model (HMM), known as the  $\varepsilon$ -machine, where the causal states form the hidden states of the model and the edge emissions the observed symbols [6,7].

The amount of information stored by the  $\varepsilon$ -machine can be quantified by the Shannon entropy of the stationary distribution on causal states:

$$C_\mu := H[P(s_j)] = - \sum_{s_j \in \mathcal{S}} P(s_j) \log_2[P(s_j)], \quad (2)$$

where  $P(s_j) = \sum_{\overleftarrow{x} \in s_j} P(\overleftarrow{x})$ . Across all (classical) causal representations of a process, the  $\varepsilon$ -machine minimizes the information cost of its corresponding memory states, and it is in this sense we refer to it as being minimal (or optimal). Because of this distinguished feature,  $C_\mu$  is called the *statistical complexity* and is considered as a quantifier of structure in the process [8], in some sense representing how much information about the past is needed to produce the future. This quantity is lower bounded by the mutual information between the past and future  $I(\overleftarrow{X}; \overrightarrow{X})$  [7]; in general this bound is not strict, and the difference is referred to as the modeling overhead [37].

When dealing with raw data, one must estimate the probabilities through inference, which will be subject to unavoidable statistical fluctuations due to the finite amount of data. As such, when applying the equivalence relation Eq. (1), a threshold  $\delta$  tolerance must be permitted, where *pasts* are assigned to the same causal state if their conditional future distributions are “close enough” [6].<sup>2</sup> Adjusting the strictness of this tolerance induces different levels of coarse graining: If too narrow then the fluctuations will lead to additional spurious causal states that would have been merged with knowledge of the exact distributions; if too loose then *pasts* with different conditional futures can be merged. Depending on this degree of coarse graining, the obtained value of  $C_\mu$  will vary; the statistical complexity is sensitive to statistical fluctuations and is generally not robust to noise.

*b. Quantum computational mechanics.* When considering quantum methods of information processing, the minimality of  $\varepsilon$ -machines no longer holds; it has been shown that causal quantum models can be found with lower information costs [15] using nonorthogonal memory states to reduce the modeling overhead. The current state-of-the-art quantum models

<sup>2</sup>It should be noted that there is not a fixed definition of how this tolerance should be implemented, but typically it would be appropriate to use some form of statistical distance between the conditional distributions, with some maximal allowed distance for merging parameterized by  $\delta$ .




FIG. 2. Unitary quantum models of stochastic processes. Repeated unitary interactions between a quantum memory and probe ancilla produces a string of stochastic outputs when the probe is measured. The specific form of the interaction depends on the particular process being modeled; the output statistics will then be as specified by this process.

[22] are based on unitary interactions between the memory subsystem and a probe ancilla:

$$U|\sigma_j\rangle|0\rangle = \sum_{x \in \mathcal{A}} \sqrt{P(x|\sigma_j)} e^{i\varphi_{xj}} |\sigma_{\lambda(x,j)}\rangle |x\rangle, \quad (3)$$

where the first subspace contains the memory, the second the probe (measured after the interaction to produce the symbol for that time step),  $\{|\sigma_j\rangle\}$  are the quantum memory states (in one-to-one correspondence with the causal states  $\{s_j\}$ ), and the phase factors  $\{\varphi_{xj}\}$  are tunable parameters. Successive applications of  $U$  on the quantum memory state and probe at each time step will yield a string of outputs from the probe measurement that are statistically distributed according to the modelled process,<sup>3</sup> as depicted in Fig. 2.

The corresponding memory cost is called the *quantum statistical memory*, given by the von Neumann entropy of the quantum memory states:

$$C_q := S[\rho] = -\text{tr}(\rho \log_2[\rho]), \quad (4)$$

where  $\rho = \sum_{s_j \in \mathcal{S}} P(s_j) |\sigma_j\rangle \langle \sigma_j|$ . The title of quantum statistical complexity is reserved for the minimum of this quantity over all causal quantum models; there is as yet no systematic approach to finding this minimal model, however, and optimizing over the phase factors is a cumbersome task. For this reason, we shall here use the best quantum models for which a systematic construction method is known: *phaseless* unitary quantum models [20], given by Eq. (3) with all phase factors set to zero. Despite not generally being minimal, the corresponding quantum statistical memory of these models has still been suggested as quantifier of structure [15,30–33], often emphasizing different features to the classical statistical complexity.

<sup>3</sup>After each time step, the probe ancilla is either reset or a fresh ancilla is introduced.

### III. INFERENCE PROTOCOL

We here introduce an inference protocol for the quantum statistical memory  $\tilde{C}_q$ <sup>4</sup> of the phaseless unitary quantum models that can be used to investigate structure in time-series data. The protocol is tailored specifically to take advantage of features of the specific model and bypassing the need to construct the  $\varepsilon$ -machine as an intermediate step. It is agnostic to the causal architecture of the process and requires only an estimate of the Markov order.

The inference protocol is based on a set of postulated quantum memory states given by clustering pasts in which the last  $L$  symbols are identical; we thus have a memory state  $|\varsigma_{x_{-L:0}}\rangle$  for each of the possible  $L$ -length words  $x_{-L:0}$ . The choice of  $L$  should correspond to the estimated Markov order of the process (or at least the effective Markov order, see below). From the data we then estimate the conditional probabilities  $\tilde{P}(X_0|X_{-L:0})$  and implicitly define the quantum memory states to satisfy the interaction

$$U|\varsigma_{x_{-L:0}}\rangle|0\rangle = \sum_{x_0 \in \mathcal{A}} \sqrt{\tilde{P}(x_0|x_{-L:0})} |\varsigma_{x_{-L+1:0}}\rangle |x_0\rangle. \quad (5)$$

A set of quantum memory states and corresponding interaction can be found through a recursive expression for the overlaps of the states and employing a reverse Gram-Schmidt procedure [20]. The estimated quantum statistical memory is then given by the von Neumann entropy of the corresponding stationary state of the memory:

$$\tilde{C}_q := S[\rho^{(L)}] = -\text{tr}(\rho^{(L)} \log_2[\rho^{(L)}]), \quad (6)$$

where  $\rho^{(L)} = \sum_{x_{-L:0}} \tilde{P}(x_{-L:0}) |\varsigma_{x_{-L:0}}\rangle \langle \varsigma_{x_{-L:0}}|$ .

To show that this protocol provides a faithful estimate of the quantum statistical memory, we first prove two properties of the above construction:

A: *Self-merging of quantum memory states*

We show that when  $L$  is at least as large as the Markov order the overlap of quantum memory states assigned to different pasts in the same causal state is unity when exact probabilities are used.

B: *Robustness of quantum statistical memory.*

We show that the quantum statistical memory of the phaseless unitary quantum model is insensitive to small perturbations in the probabilities.

With these properties we can then discuss the accuracy of the inference protocol.

C: *Inference protocol.*

We indicate how the accuracy of our estimate of quantum statistical memory scales with the amount of data and how it converges for sufficiently large data streams.

#### A. Self-merging of quantum memory states

We first show that the blind construction Eq. (5) will automatically adopt the causal architecture of the process (i.e., that pasts belonging to the same causal state are assigned to the same memory state) without explicit need to apply the causal equivalence relation [21], provided that exact probabilities are

<sup>4</sup>We use tildes to represent estimated quantities.

used and the chosen  $L$  is at least as large as the Markov order of the process. That is, the quantum memory states we construct will correspond to the same states as would be obtained from the phaseless form of the model Eq. (3) but without prior knowledge of how the pasts are clustered into causal states. In turn, this means the blind construction will faithfully replicate the process, with the same quantum statistical memory.

To see this, let  $R$  denote the Markov order of the process and recall that this means  $P(\vec{X} | X_{-R:0}) = P(\vec{X} | \vec{X})$ . Since the Markov order can alternatively be expressed as the longest history length needed to determine the causal state (i.e.,  $R = \min\{r : H(S_0 | x_{-r:0}) = 0\}$ ), all pasts where the latest  $R$  symbols are identical belong to the same causal state [16]. We can see that if  $L \geq R$ , the construction already correctly merges all pasts where the latest  $L$  symbols are identical.

By analogy with the corresponding methods for phaseless unitary models [20], we can express

$$\begin{aligned} \langle \mathcal{S}_{x_{-L:0}} | \mathcal{S}'_{x'_{-L:0}} \rangle &= \langle \mathcal{S}_{x_{-L:0}} | U^\dagger U | \mathcal{S}'_{x'_{-L:0}} \rangle \\ &= \sum_{x_0 \in \mathcal{A}} \sqrt{P(x_0 | x_{-L:0}) P(x_0 | x'_{-L:0})} \langle \mathcal{S}_{x_{-L+1:1}} | \mathcal{S}'_{x'_{-L+1:1}} \rangle. \end{aligned} \quad (7)$$

Iteratively applying this relation, we obtain that

$$\langle \mathcal{S}_{x_{-L:0}} | \mathcal{S}'_{x'_{-L:0}} \rangle = \sum_{\vec{x}} \sqrt{P(\vec{x} | x_{-L:0}) P(\vec{x} | x'_{-L:0})}, \quad (8)$$

and if  $L \geq R$ , we then have

$$\langle \mathcal{S}_{x_{-L:0}} | \mathcal{S}'_{x'_{-L:0}} \rangle = \sum_{\vec{x}} \sqrt{P(\vec{x} | \vec{x}) P(\vec{x} | \vec{x}')}, \quad (9)$$

where the full pasts  $\vec{x}$  and  $\vec{x}'$  can be taken as any pasts with the correct corresponding last  $L$  symbols. We are thus able to conclude that

$$\langle \mathcal{S}_{x_{-L:0}} | \mathcal{S}'_{x'_{-L:0}} \rangle = 1 \Leftrightarrow P(\vec{X} | \vec{x}) = P(\vec{X} | \vec{x}'), \quad (10)$$

which can be seen as an instantiation of the causal equivalence relation, i.e., two pasts are mapped to the same memory state iff they have the same conditional future statistics.

### B. Robustness of quantum statistical memory

We next show that the quantum statistical memory of our construction is robust to small perturbations of the probabilities. Consider mapping

$$P(\vec{X}) \rightarrow P^\epsilon(\vec{X}) = P(\vec{X}) + \epsilon \Delta P(\vec{X}), \quad (11)$$

where  $\Delta P$  governs the relative changes in the distribution for each string and  $\epsilon$  the strength of the perturbation. We here outline a proof that the perturbation to  $C_q$  scales smoothly with  $\epsilon$ ; full details may be found in Appendix B.

The Gram matrix  $G$  of a quantum state  $\rho = \sum_j P_j |\sigma_j\rangle \langle \sigma_j|$  is defined as  $G_{jk} = \sqrt{P_j P_k} \langle \sigma_j | \sigma_k \rangle$  and can be shown to have the same spectrum as  $\rho$  [17, 38, 39]. As such, it is possible to define the Gram matrix of our construction as

$$G_{x_{-L:0} x'_{-L:0}} = \sqrt{P(x_{-L:0}) P(x'_{-L:0})} \langle \mathcal{S}_{x_{-L:0}} | \mathcal{S}'_{x'_{-L:0}} \rangle, \quad (12)$$

and correspondingly, from its spectrum calculate  $C_q$ .

Consider that we have  $L \geq R$  such that it is possible to express the overlaps of the quantum memory states as

$$\langle \mathcal{S}_{x_{-L:0}} | \mathcal{S}'_{x'_{-L:0}} \rangle = \sum_{x_{0:L}} \sqrt{P(x_{0:L} | x_{-L:0}) P(x_{0:L} | x'_{-L:0})}. \quad (13)$$

Note that we only need consider  $L$  steps into the future, as this uniquely determines the subsequent memory state, independent of the past. Now replace each of the probability distributions in this expression by their corresponding perturbed forms, which may be obtained from the marginals of Eq. (11). We can then calculate the perturbed form of the corresponding Gram matrix using this expression for the overlaps and show that its spectrum varies smoothly with  $\epsilon$ . Since the von Neumann entropy is a continuous function of the spectrum of a state [40], we thus find that it, too, smoothly deforms with  $\epsilon$ .

Hence, we can conclude that the quantum statistical memory is robust to small perturbations in the probability distributions. Due to the self-merging of our quantum memory states, we can also see that the quantum statistical memory of phaseless unitary quantum models is similarly robust in general. This is in contrast to the classical statistical complexity, which can vary discontinuously with the probabilities— notably, whenever the perturbation triggers a new merging of pasts into a causal state, or conversely, the splitting of a causal state.

### C. Inference protocol

With these two results in hand, we are now in a position to argue the accuracy of our inference protocol. By parsing the data and adopting a frequentist approach to estimate the marginal distribution of words of length  $L+1$  (for some chosen  $L$ ), we are able to estimate the conditional probabilities needed for Eq. (5). Moreover, we can use these marginals to construct estimates for the probabilities needed to evaluate Eq. (12) [the conditional probabilities  $\tilde{P}(X_{0:L} | X_{-L:0})$  can be obtained from multiplying  $\tilde{P}(X_0 | X_{-L:0})$  with the assumption that  $L$  is at least as large as the Markov order] and in turn, estimate  $C_q$ . From the results of the previous two sections, we can be assured that this should be a faithful estimate, provided that the estimated marginals are close to the exact distributions, and  $L \geq R$ .

Let us first consider the error in our estimates of the marginals. With any inference from finite data there will of course be statistical fluctuations; this is not problematic for our inference protocol provided that these fluctuations are sufficiently small. These fluctuations can be treated in the same manner as the perturbations of the previous section but with the perturbation terms  $\Delta P$  now being stochastic variables. In Appendix C we show that the size of the error in the inferred Gram matrix and the estimated quantum statistical memory  $\hat{C}_q$  approximately scale as  $O(|\mathcal{A}|^L / \sqrt{N})$ , where  $|\mathcal{A}|$  is the size of the alphabet and  $N$  the size of the data stream.

The next question is how we determine the choice of  $L$ . From the above, we see that taking too large an  $L$  will lead to untenably large fluctuations. We must therefore cap it at some value  $L_{\max}$  where  $|\mathcal{A}|_{\max}^L / \sqrt{N} \ll 1$ ; as a rough guideline we suggest  $L_{\max} \lesssim \log_{|\mathcal{A}|}(N/1000)$ . On the other hand, we require  $L$  to be large enough that it matches or



exceeds the Markov order of the process in order to effect the (approximate) self-merging of quantum memory states.

When a process has a large, or even infinite Markov order, it may not be possible to have an  $L$  that satisfies both of these requirements. Nevertheless, while the Markov order tells us how far back into the past memory effects can persist, it does not inform us how strong they are. It is often the case that these long-range historical dependencies only have minor influence on the future and that the recent past is much more relevant. In such instances, the influence of the distant past can be thought of as a small perturbation to the statistics with respect to the recent past alone, and so from the previous section we can expect that they have minimal impact on the requisite quantum statistical memory. We therefore introduce the concept of an *effective Markov order*  $R_{\text{eff}}$  that encapsulates the idea that a sufficiently long string of past observations that is less than the Markov order may nevertheless still be “good enough” to capture most of the predictive information contained in the past.

We define the effective Markov order as the smallest length of a string of past observations for which the influence of considering an additional symbol one step further into the past does not exceed some threshold. Specifically, we define  $R_{\text{eff}}$  as the smallest integer  $r$  that satisfies

$$\max_{xx'} \langle D(P(X_0|X_{-r:0}), P(X_0|x'X_{-r:0})) \rangle < \delta, \quad (14)$$

where  $\delta$  is the parameter defining the threshold,<sup>5</sup> the expectation value is taken over the distribution of strings of length  $r$ , and  $D$  is some distance measure between probability distributions; for the purposes of this work we will use the trace distance  $D(P, Q) := \sum_x |P(x) - Q(x)|/2$ . We can estimate the effective Markov order from this data and use this as a guide for choosing a value for  $L$  in the inference protocol.

#### IV. EXAMPLES

We now demonstrate the efficacy of our model with two toy example processes: The so-called  $R$ - $k$  golden mean and nemo processes. We use an exact HMM representation of the processes to generate a representative string of outputs and infer  $\tilde{C}_q$  directly from this time series. We choose the initial state according to the stationary distribution of the HMM such that the output statistics are representative of the stationary state of the process. We will look at how the estimate for the quantum statistical memory varies both with the amount of data  $N$  and the history length  $L$ , highlighting the range of  $L$  values that would be considered appropriate given the above considerations regarding the (effective) Markov order and expected size of fluctuations.

##### A. Golden mean process

We first look at the  $R$ - $k$  golden mean process family. Here,  $R$  and  $k$  are tunable parameters that correspond to

<sup>5</sup>Strictly, we have a family of effective Markov orders for the process, parameterized by  $\delta$ .

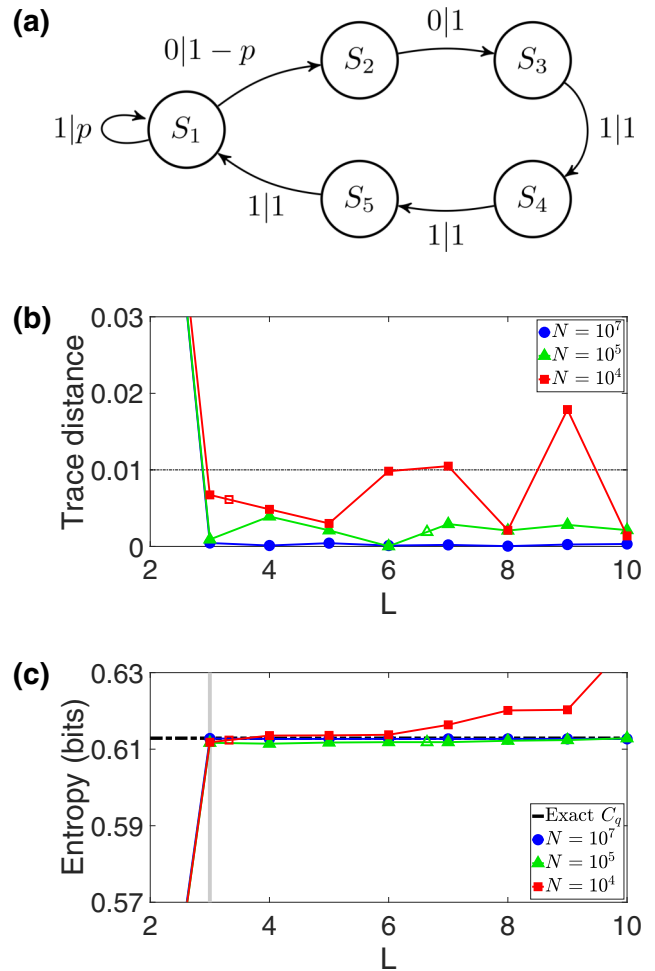


FIG. 3. Golden mean process. (a)  $\varepsilon$ -machine for the 3-2 golden mean process; the notation  $x|P$  denotes that the indicated transition between states involves output of symbol  $x$  and occurs with probability  $P$ . (b) Average trace distance for variation  $L + 1$  steps in the past; hollow circles denote the points where  $L = \log_2(N/1000)$ , and the dashed line  $\delta = 0.01$ . (c) Comparison of exact and estimated quantum statistical memory for different lengths of data stream; the variation with  $L$  is shown for the estimated quantities, and the vertical line indicates the effective Markov order for  $\delta = 0.01$ . For plots (b) and (c) we take  $p = 0.9$ .

the Markov order and cryptic order<sup>6</sup> of the process, respectively [16,37,41,42]. Here, we will consider the 3-2 golden mean process specifically, as represented by its  $\varepsilon$ -machine in Fig. 3(a).

In Fig. 3(b) we plot the expectation of the trace distance between differing symbols increasingly far into the past [i.e., between  $\tilde{P}(X_0|0X_{-R:0})$  and  $\tilde{P}(X_0|1X_{-R:0})$ , from which we can infer an effective Markov order for the process as defined in Eq. (14)]. The limits of finite data are already visible in this plot, with the instability clear when  $L$  is too large relative

<sup>6</sup>The cryptic order is a counterpart to the Markov order, describing the minimum length of past observations required to be certain of the present causal state, given that one knows the entire future, i.e., the smallest  $k$  satisfying  $H[S|X_{-k:\infty}] = 0$ .

to  $N$ , due to undersampling of the process statistics. The hollow circle on each plot represents the point at which  $L = \log_2(N/1,000)$ —beyond this point we consider the statistics to be undersampled. Setting a threshold  $\delta = 0.01$ , we would assign an effective Markov order of  $R_{\text{eff}} = 3$ , which aligns with the true Markov order of the process. Figure 3(c) displays the estimated  $\tilde{C}_q$ ; we see that at the Markov order the estimate is very close to the exact value  $C_q$ , with statistical noise gradually degrading the quality of the estimate at larger  $L$  when we have insufficient data. This highlights both the efficacy of protocol and importance of selecting an appropriate value for  $L$ . The corresponding statistical complexity of the process  $C_\mu \approx 1.435$  is omitted as it is much larger than both  $C_q$  and  $\tilde{C}_q$ .

### B. Nemo process

As a second example, we consider the nemo process, which can be represented by its  $\varepsilon$ -machine as in Fig. 4(a). A key feature of this process is that it has infinite Markov order: A contiguous string of zeros of any length cannot be exactly synchronized to a causal state. As such, with this example it is not possible to choose an  $L$  that matches the Markov order of the process. Nevertheless, we will show that the effective Markov order can provide a suitable proxy.

Figure 4(b) shows the expected trace distance for variation in the symbol  $L + 1$  steps into the past; setting a tolerance  $\delta = 0.01$ , we assign an effective Markov order  $R_{\text{eff}} = 6$ . Examining the estimated quantum statistical memory  $\tilde{C}_q$  [Fig. 4(c)], we see that setting  $L = 6$  does indeed appear to provide an accurate estimate of  $C_q$ , striking a balance between allowing sufficiently long histories to capture most of the past dependency while not going as far as to underfit. Note that in this case we should consider  $N = 10000$  to be insufficient data to provide a good estimate, as the prescribed  $L_{\text{max}}$  is significantly smaller than the effective Markov order—practically, this could be deduced from the trace distance, which never drops below the threshold value. The corresponding statistical complexity of the process  $C_\mu \approx 1.583$  is again omitted from the figure.

## V. DISCUSSION

We have introduced a protocol for estimating the information cost of quantum simulation of stochastic processes. We have shown that both this quantity and our protocol are robust to small statistical perturbations. This provides a means to characterize structure in a process according to the amount of (quantum) resources needed to capture its behavior, analogous to corresponding classical quantities [6,34,35]. Moreover, this provides a key step towards blind construction of quantum models that efficiently replicate the behavior of such processes.

An essential consideration to be made in this latter vein is the capabilities of current and near-term quantum technologies. Our inference protocol accurately captures the information that must be stored by a quantum model of a process—appropriately indicating the amount of structure in the process—at the expense of indicating a multitude of memory states, typically  $\sim |\mathcal{A}|^L$ , in excess of the number

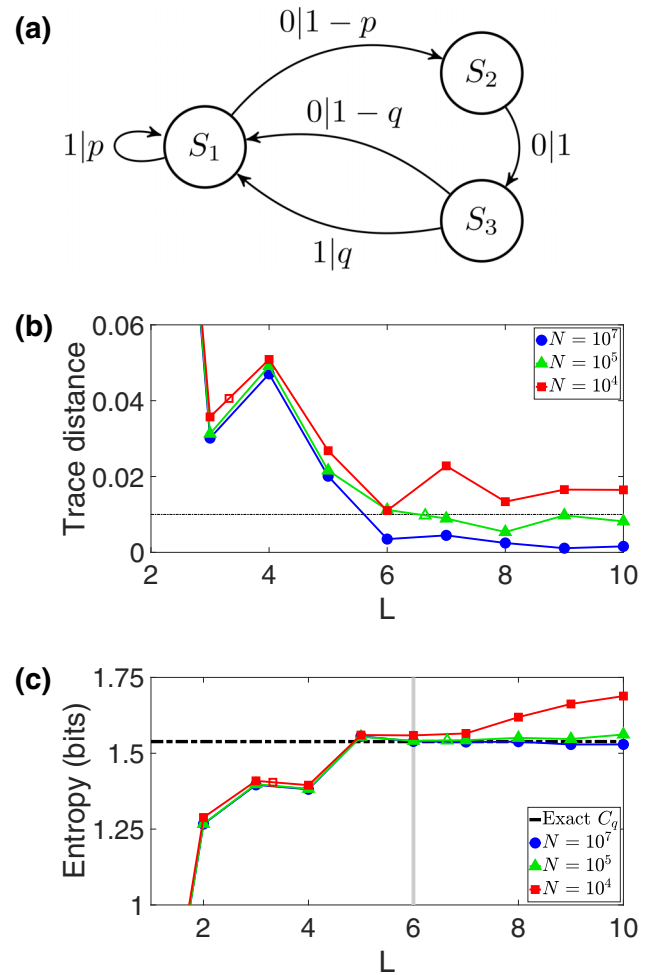


FIG. 4. Nemo process. (a)  $\varepsilon$ -machine for the nemo process. (b) Average trace distance for variation  $L + 1$  steps in the past; hollow circles denote the points where  $L = \log_2(N/1000)$ , and the dashed line  $\delta = 0.01$ . (c) Comparison of exact and estimated quantum statistical memory for different lengths of data stream; the variation with  $L$  is shown for the estimated quantities, and the vertical line indicates the effective Markov order for  $\delta = 0.01$ . For plots (b) and (c) we take  $p = 0.1$  and  $q = 0.9$ .

of causal states. The number of memory states is parameterized by a companion metric, the *topological memory*  $D_q = \log_2[\dim(\rho)]$ ; quantum advantages can also be found in terms of this measure [22,25,26,29,43]. Future work will investigate methods of compression in this parameter via truncation in terms of the quantum state space, opening up the possibility to implement the inferred constructions experimentally. The accuracy of these inferred models can then be explored using recently developed quantifiers of process distinguishability [44].

### ACKNOWLEDGMENTS

We thank Chew L. Y., Suen W. Y., and Suryadi for discussions. This work was funded by the Lee Kuan Yew Endowment Fund (Postdoctoral Fellowship), the Singapore Ministry of Education through Tier 1 Grant No. RG190/17, the Foundational Questions Institute and the Fetzer Franklin




FIG. 5. Pasts are clustered into groups—causal states—based on the causal equivalence relation Eq. (A1); this illustration is sometimes termed “Occam’s Pool” [7].

Fund (a donor-advised fund of the Silicon Valley Community Foundation) through Grant No. FQXi-RFP-1809, and the Singapore National Research Foundation through Fellowship No. NRF-NRFF2016-02, and NRF-ANR Grant No. NRF2017-NRF-ANR004 VanQuTe. T.J.E. thanks the Centre for Quantum Technologies for their hospitality.

#### APPENDIX A: COMPUTATIONAL MECHANICS PRIMER

Computational mechanics [6–8] is a branch of complexity theory originating from studies of structure and intrinsic computation in dynamical systems and stochastic processes. The core of computational mechanics is in isolating information contained in the past of the process that is relevant to its future. This is achieved through the *causal equivalence relation*:

$$\overleftarrow{x} \sim_{\varepsilon} \overleftarrow{x}' \Leftrightarrow P(\overrightarrow{X} | \overleftarrow{X} = \overleftarrow{x}) = P(\overrightarrow{X} | \overleftarrow{X} = \overleftarrow{x}'). \quad (\text{A1})$$

This relation groups together two pasts  $\overleftarrow{x}$  and  $\overleftarrow{x}'$  iff their future outputs are statistically indistinguishable. The resulting equivalence classes are called *causal states*  $s_j \in \mathcal{S}$ ; this grouping is illustrated in Fig. 5. Put simply, if two pasts give rise to identical future statistics, retaining any further information that would distinguish between them yields no predictive power. It transpires that this is the optimal grouping of pasts: Causal states are the minimal sufficient statistic of the past with respect to the future [7].

When considering bi-infinite, discrete-time stationary stochastic processes (as done in this work), the causal states can be represented as the latent states of an edge-emitting hidden Markov model (HMM), where the emissions on edges correspond to the output symbols of the process, and the probabilistic transition structure is defined by the process. Specifically, if a past  $\overleftarrow{x} \in s_j$  has a probability  $P(x|j)$  of emitting  $x$ , and  $\overleftarrow{x}x \in s_k$ , then an edge  $s_j \rightarrow s_k$  exists with emission  $x$  and with transition probability  $P(x|j)$ . Note that when appending a particular symbol  $x$  to any of the pasts  $\{\overleftarrow{x}\}$  in a given causal state  $s_j$  their concatenated pasts  $\{\overleftarrow{x}x\}$  will all belong to the same causal state  $s_k$ . That is, given the initial causal state and output symbol, the subsequent causal state is uniquely determined. This follows directly from the causal equivalence relation and within computational mechanics is referred to as unifilarity.

The corresponding HMM is referred to as the  $\varepsilon$ -machine of the process and is the (classically) provably minimal—in terms of number of states and their information content—representation of the process. Due to its privileged position, this information content—which can be seen as the amount of information that must be tracked about the past in order to




FIG. 6. General form of a two-symbol Markov chain. Aside from special parameter sets the states of the chain represent the causal states of the process. The notation  $x|p$  on an edge indicates the transition occurs with probability  $p$ , accompanied by an emission  $x$ .

produce a statistically accurate future set of outputs—is taken in computational mechanics as a measure of structure called the *statistical complexity*  $C_{\mu}$ :

$$C_{\mu} := H[P(s_j)] = - \sum_{s_j \in \mathcal{S}} P(s_j) \log_2[P(s_j)], \quad (\text{A2})$$

where  $P(s_j) = \sum_{\overleftarrow{x} \in s_j} P(\overleftarrow{x})$  is the probability that an observed past belongs to  $s_j$ . This quantity can also more abstractly be thought of as the information communicated from the past of the process to the future.

There exist a number of methods for reconstructing  $\varepsilon$ -machines from data [6,34,35]; we do not detail these here. Further background and details on computational mechanics can be found within, e.g., Refs. [6–8,45].

As an example of how computational mechanics can be deployed, consider a two-symbol Markov chain (sometimes referred to as the perturbed coin [15]), as illustrated in Fig. 6. Given that the process is Markovian, we can immediately conclude that every past for which the most recent symbol is the same can be grouped into the same causal state; in other words, the causal states are either the states of the Markov chain or a coarse graining thereof. Indeed, we can see that with the exception of certain special cases, the two states in the chain provide different future statistics and so form the two causal states of the process.

The aforementioned special cases are when either of the following occur:

- (i)  $q_0 = 0$  or  $q_1 = 0$ , wherein the process consists of either an infinite string of 1s, or an infinite string of 0s, respectively;
- (ii)  $q_0 = q_1 = 1$ , wherein the process will be an infinite string of alternating 0s and 1s (i.e.,  $\dots 0101010101 \dots$ );
- (iii)  $q_0 = 1 - q_1$  (and  $q_1 = 1 - q_0$ ), wherein the outputs are stochastic but both states produce identical future statistics.

In all these cases, there is only a single causal state, and trivially  $C_{\mu} = 0$ .

For all other parameter sets, the steady-state distribution of the causal states can be found from the (normalized) eigenvector of the transition matrix  $T$  of the  $\varepsilon$ -machine with unit eigenvalue. Specifically, we have

$$T = \begin{pmatrix} 1 - q_0 & q_1 \\ q_0 & 1 - q_1 \end{pmatrix}, \quad (\text{A3})$$

for which the corresponding eigenvector is  $(q_1, q_0)^T / (q_0 + q_1)$ . The Shannon entropy of the elements of this vector

then yield the statistical complexity

$$C_\mu = \log_2(q_0 + q_1) - \frac{q_0}{q_0 + q_1} \log_2(q_0) - \frac{q_1}{q_0 + q_1} \log_2(q_1). \quad (\text{A4})$$

Notably, whenever  $q_0 = q_1 \neq \{0, 0.5, 1\}$  we have that  $C_\mu = 1$ .

### APPENDIX B: DETAILED PROOF OF ROBUSTNESS OF QUANTUM STATISTICAL MEMORY

Recall from the main text that the Gram matrix  $G$  of a quantum state  $\rho = \sum_j P_j |\sigma_j\rangle\langle\sigma_j|$  is defined as  $G_{jk} = \sqrt{P_j P_k} \langle\sigma_j|\sigma_k\rangle$  [17,38,39]. If one considers a purification of the original state  $|\Psi\rangle = \sum_j \sqrt{P_j} |\sigma_j\rangle |j\rangle$ , then  $\rho$  can be recovered by taking a partial trace over the second subsystem and  $G$  by tracing out the first. The Gram matrix thus has the same spectrum as the original state and so can be used to calculate functions of this spectrum, such as the entropy. As stated in Eq. (12), for our construction the Gram matrix is given by

$$G_{x_{-L:0}x'_{-L:0}} = \sqrt{P(x_{-L:0})P(x'_{-L:0})} \langle\zeta_{x_{-L:0}}|\zeta_{x'_{-L:0}}\rangle. \quad (\text{B1})$$

Using Eq. (13), this can be expanded as

$$\begin{aligned} G_{x_{-L:0}x'_{-L:0}} &= \sqrt{P(x_{-L:0})P(x'_{-L:0})} \\ &\quad \times \sum_{x_{0:L}} \sqrt{P(x_{0:L}|x_{-L:0})P(x_{0:L}|x'_{-L:0})} \\ &= \sum_{x_{0:L}} \sqrt{P(x_{-L:0}x_{0:L})P(x'_{-L:0}x_{0:L})}. \end{aligned} \quad (\text{B2})$$

We now examine how this changes when the probabilities  $P$  are replaced by their perturbed versions  $P^\epsilon := P + \epsilon\Delta P$ . Consider expanding out the square root of the product of two such perturbations:

$$\begin{aligned} \sqrt{P^\epsilon Q^\epsilon} &= \sqrt{(P + \epsilon\Delta P)(Q + \epsilon\Delta Q)} \\ &= \sqrt{PQ} \sqrt{1 + \epsilon\left(\frac{\Delta P}{P} + \frac{\Delta Q}{Q}\right) + \epsilon^2 \frac{\Delta P \Delta Q}{PQ}} \\ &\approx \sqrt{PQ} + \epsilon \frac{\sqrt{PQ}}{2} \left(\frac{\Delta P}{P} + \frac{\Delta Q}{Q}\right) + O(\epsilon^2). \end{aligned} \quad (\text{B3})$$

Substituting this into Eq. (B2), we obtain

$$\begin{aligned} G_{x_{-L:0}x'_{-L:0}}^\epsilon &= \sum_{x_{0:L}} \sqrt{P^\epsilon(x_{-L:0}x_{0:L})P^\epsilon(x'_{-L:0}x_{0:L})} \\ &\approx \sum_{x_{0:L}} \sqrt{P(x_{-L:0}x_{0:L})P(x'_{-L:0}x_{0:L})} \\ &\quad + \epsilon \sum_{x_{0:L}} \frac{\sqrt{P(x_{-L:0}x_{0:L})P(x'_{-L:0}x_{0:L})}}{2} \\ &\quad \times \left(\frac{\Delta P(x_{-L:0}x_{0:L})}{P(x_{-L:0}x_{0:L})} + \frac{\Delta P(x'_{-L:0}x_{0:L})}{P(x'_{-L:0}x_{0:L})}\right) + O(\epsilon^2). \end{aligned} \quad (\text{B4})$$

Thus, we can write

$$G_{x_{-L:0}x'_{-L:0}}^\epsilon \approx G_{x_{-L:0}x'_{-L:0}} + \epsilon \Delta G_{x_{-L:0}x'_{-L:0}} + O(\epsilon^2), \quad (\text{B5})$$

where

$$\begin{aligned} \Delta G_{x_{-L:0}x'_{-L:0}} &= \sum_{x_{0:L}} \frac{\sqrt{P(x_{-L:0}x_{0:L})P(x'_{-L:0}x_{0:L})}}{2} \\ &\quad \times \left(\frac{\Delta P(x_{-L:0}x_{0:L})}{P(x_{-L:0}x_{0:L})} + \frac{\Delta P(x'_{-L:0}x_{0:L})}{P(x'_{-L:0}x_{0:L})}\right). \end{aligned} \quad (\text{B6})$$

From Weyl's inequality, it then follows that the perturbation to the eigenvalues of  $G$  are bounded by the spectral norm of  $\epsilon\Delta G$  [39]. Clearly, this norm scales with  $\epsilon$ , and so the perturbation to the spectrum of  $G$  varies continuously with  $\epsilon$ . Finally, since the von Neumann entropy of a quantum state is given by the Shannon entropy of its spectrum—and is a continuous function of it [40]—the quantum statistical memory is smoothly deformed by the perturbation and so is robust.

### APPENDIX C: SCALING OF STATISTICAL NOISE

We now examine the effects of statistical noise in our estimates of word probabilities on  $\tilde{C}_q$ . These fluctuations can be considered as a (stochastic) perturbation, i.e.,  $\tilde{P} = P + \epsilon\Delta P$ , allowing us to employ the results above. We will set  $\epsilon$  to 1, folding the full scaling of the perturbation with  $L$  and  $N$  into  $\Delta P$ .

Recall that the corrections to the eigenvalues arising from the perturbation are bounded by the spectral norm of  $\Delta G$ , i.e., its largest eigenvalue, which in turn is bounded by the product of the dimension of the matrix with its largest element. The elements of this matrix are given in Eq. (B6); to assess how they scale we will replace the statistical variables by their standard errors. Since the word probabilities  $P(X_{-L:L})$  are described by binomial distributions (a randomly selected string can be assigned as either being the given word  $x_{-L:L}$  or not), the standard error is given by

$$\sigma_{\tilde{P}(x_{-L:L})} = \sqrt{\frac{P(x_{-L:L})[1 - P(x_{-L:L})]}{N}}. \quad (\text{C1})$$

Inserting this into Eq. (B6) we obtain

$$\begin{aligned} \Delta G_{x_{-L:0}x'_{-L:0}} &= \frac{1}{2\sqrt{N}} \sum_{x_{0:L}} \sqrt{P(x'_{-L:0}x_{0:L})[1 - P(x_{-L:0}x_{0:L})]} \\ &\quad + \sqrt{P(x_{-L:0}x_{0:L})[1 - P(x'_{-L:0}x_{0:L})]}. \end{aligned} \quad (\text{C2})$$

The probability of obtaining a given word of length  $L$  falls off approximately exponentially with  $L$ . Let us assume all long words are roughly evenly distributed and take  $P(X_{-L:L}) \sim |\mathcal{A}|^{-2L}$ . Considering that there are roughly  $|\mathcal{A}|^L$  terms in the sum, we have  $\Delta G_{x_{-L:0}x'_{-L:0}} \sim 1/\sqrt{N}$ . Finally, since the dimension of the matrix scales as  $|\mathcal{A}|^L$ , the spectral norm (and thus the bound on the size of perturbations to the spectrum of the Gram matrix) scales  $\sim |\mathcal{A}|^L/\sqrt{N}$ .



- [1] P. Lynch, *J. Comput. Phys.* **227**, 3431 (2008).
- [2] P. Bauer, A. Thorpe, and G. Brunet, *Nature (London)* **525**, 47 (2015).
- [3] G. P. Pavlos, L. P. Karakatsanis, A. C. Iliopoulos, E. G. Pavlos, M. N. Xenakis, P. Clark, J. Duke, and D. S. Monos, *Physica A* **438**, 188 (2015).
- [4] T. Preis, D. Y. Kenett, H. E. Stanley, D. Helbing, and E. Ben-Jacob, *Sci. Rep.* **2**, 752 (2012).
- [5] B. S. Kerner and H. Rehborn, *Phys. Rev. E* **53**, R4275(R) (1996).
- [6] J. P. Crutchfield and K. Young, *Phys. Rev. Lett.* **63**, 105 (1989).
- [7] C. R. Shalizi and J. P. Crutchfield, *J. Stat. Phys.* **104**, 817 (2001).
- [8] J. P. Crutchfield, *Nat. Phys.* **8**, 17 (2011).
- [9] J. E. Hanson and J. P. Crutchfield, *Physica D* **103**, 169 (1997).
- [10] W. M. Gonçalves, R. D. Pinto, J. C. Sartorelli, and M. J. De Oliveira, *Physica A* **257**, 385 (1998).
- [11] J. B. Park, J. Won Lee, J. S. Yang, H. H. Jo, and H. T. Moon, *Physica A* **379**, 179 (2007).
- [12] R. Haslinger, K. L. Klinkner, and C. R. Shalizi, *Neural Comput.* **22**, 121 (2010).
- [13] A. Montanaro, *npj Quantum Inf.* **2**, 15023 (2016).
- [14] C. H. Bennett and G. Brassard, *Theor. Comput. Sci.* **560**, 7 (2014).
- [15] M. Gu, K. Wiesner, E. Rieper, and V. Vedral, *Nat. Commun.* **3**, 762 (2012).
- [16] J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield, *Sci. Rep.* **6**, 20495 (2016).
- [17] P. M. Riechers, J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield, *Phys. Rev. A* **93**, 052317 (2016).
- [18] J. Thompson, A. J. P. Garner, V. Vedral, and M. Gu, *npj Quantum Inf.* **3**, 6 (2017).
- [19] T. J. Elliott and M. Gu, *npj Quantum Inf.* **4**, 18 (2018).
- [20] F. C. Binder, J. Thompson, and M. Gu, *Phys. Rev. Lett.* **120**, 240502 (2018).
- [21] T. J. Elliott, A. J. P. Garner, and M. Gu, *New J. Phys.* **21**, 013021 (2019).
- [22] Q. Liu, T. J. Elliott, F. C. Binder, C. Di Franco, and M. Gu, *Phys. Rev. A* **99**, 062110 (2019).
- [23] A. J. P. Garner, Q. Liu, J. Thompson, and V. Vedral, *New J. Phys.* **19**, 103009 (2017).
- [24] C. Aghamohammadi, J. R. Mahoney, and J. P. Crutchfield, *Sci. Rep.* **7**, 6735 (2017).
- [25] J. Thompson, A. J. P. Garner, J. R. Mahoney, J. P. Crutchfield, V. Vedral, and M. Gu, *Phys. Rev. X* **8**, 031013 (2018).
- [26] T. J. Elliott, C. Yang, F. C. Binder, A. J. P. Garner, J. Thompson, and M. Gu, [arXiv:1909.02817](https://arxiv.org/abs/1909.02817).
- [27] M. S. Palsson, M. Gu, J. Ho, H. M. Wiseman, and G. J. Pryde, *Sci. Adv.* **3**, e1601302 (2017).
- [28] F. G. Jouneghani, M. Gu, J. Ho, J. Thompson, W. Y. Suen, H. M. Wiseman, and G. J. Pryde, [arXiv:1711.03661](https://arxiv.org/abs/1711.03661).
- [29] F. Ghafari, N. Tischler, J. Thompson, M. Gu, L. K. Shalm, V. B. Verma, S. W. Nam, R. B. Patel, H. M. Wiseman, and G. J. Pryde, *Phys. Rev. X* **9**, 041013 (2019).
- [30] R. Tan, D. R. Terno, J. Thompson, V. Vedral, and M. Gu, *Eur. Phys. J. Plus* **129**, 191 (2014).
- [31] W. Y. Suen, J. Thompson, A. J. P. Garner, V. Vedral, and M. Gu, *Quantum* **1**, 25 (2017).
- [32] C. Aghamohammadi, J. R. Mahoney, and J. P. Crutchfield, *Phys. Lett. A* **381**, 1223 (2017).
- [33] W. Y. Suen, T. J. Elliott, J. Thompson, A. J. P. Garner, J. R. Mahoney, V. Vedral, and M. Gu, [arXiv:1812.09738](https://arxiv.org/abs/1812.09738).
- [34] C. R. Shalizi and K. L. Klinkner, in *Uncertainty in Artificial Intelligence: Proceedings of the Twentieth Conference (UAI 2004)*, edited by M. Chickering and J. Y. Halpern (AUAI Press, Arlington, VA, 2004), pp. 504–511.
- [35] C. C. Strelloff and J. P. Crutchfield, *Phys. Rev. E* **89**, 042119 (2014).
- [36] E. Racca, F. Laio, D. Poggi, and L. Ridolfi, *Phys. Rev. E* **75**, 011126 (2007).
- [37] J. R. Mahoney, C. J. Ellison, R. G. James, and J. P. Crutchfield, *Chaos* **21**, 037112 (2011).
- [38] R. Jozsa, *Phys. Rev. A* **62**, 012301 (2000).
- [39] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed. (Cambridge University Press, Cambridge, England, 2012).
- [40] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, UK, 2010).
- [41] J. P. Crutchfield, C. J. Ellison, and J. R. Mahoney, *Phys. Rev. Lett.* **103**, 094101 (2009).
- [42] J. R. Mahoney, C. J. Ellison, and J. P. Crutchfield, *J. Phys. A: Math. Theor.* **42**, 362002 (2009).
- [43] S. P. Loomis and J. P. Crutchfield, *J. Stat. Phys.* **176**, 1317 (2019).
- [44] C. Yang, F. C. Binder, M. Gu, and T. J. Elliott, [arXiv:1909.08366](https://arxiv.org/abs/1909.08366).
- [45] C. R. Shalizi, Ph.D. Thesis, University of Wisconsin, 2001.