Title
Measurement-induced randomness and structure in controlled qubit processes

Permalink
https://escholarship.org/uc/item/1cj6492c

Journal
Physical Review E, 102(4)

ISSN
2470-0045

Authors
Venegas-Li, Ariadna E
Jurgens, Alexandra M
Crutchfield, James P

Publication Date
2020-10-01

DOI
10.1103/physreve.102.040102

Peer reviewed

Measurement-Induced Randomness and Structure
in Controlled Qubit Processes

Ariadna E. Venegas-Li,[*] Alexandra M. Jurgens,[†] and James P. Crutchfield[‡]
Complexity Sciences Center and Physics Department,
University of California at Davis, One Shields Avenue, Davis, CA 95616

(Dated: August 27, 2019)

When an experimentalist measures a time series of qubits, the outcomes generate a classical stochastic process. We show that measurement induces high complexity in these processes in two specific senses: they are inherently unpredictable (positive Shannon entropy rate) and they require an infinite number of features for optimal prediction (divergent statistical complexity). We identify nonunifilarity as the mechanism underlying the resulting complexities and examine the influence that measurement choice has on the randomness and structure of measured qubit processes. We introduce new quantitative measures of this complexity and provide efficient algorithms for their estimation.

Introduction Temporal sequences of controlled quantum states are key to fundamental physics and its engineering deployment. Quantum entanglement [1] between emitted qubits, such as photons [2], is central to Bell probes [3] of inconsistencies between quantum mechanics and local hidden variable theories [4]. Complementing their scientific role, entangled qubits are now recognized as basic resources for quantum technologies—quantum key distribution [5], teleportation [6], metrology [7], and computing [8]. The quest there is for qubit sources that allow *on-demand* generation: at a certain time a source should emit one and only one pair of entangled photons. Qubit sources should also be *efficient*: qubits emitted and collected with a high success rate. And, individual qubits should have *specified properties*. In EPR experiments photons in emitted pairs should be identical from trial to trial. And, in communication systems polarization states should manipulable at the highest possible rates [9].

Much experimental effort has been invested to develop qubit sources that, for example, extract entangled photons from trapped atoms [10, 11], spontaneous parametric down-conversion [12], quantum dots [13], and related CQED systems [14]. To date, though, there is still no single qubit source that exactly meets the performance desiderata. The on-demand criterion has been particularly vexing [15]. Addressing these challenges leads rather directly to a common question, one that touches on both fundamental physics and quantum engineering, How to characterize the statistical and structural properties of qubit time series? The underlying challenge is that a systematic description of quantum processes with memory in terms of experimental measurements has yet to be given [16].

To address these challenges we first introduce qubit information sources observed through quantum measurement. Noting that their outputs—what an experimentalist sees—are classical stochastic processes, we review measures of complexity for the latter. This leads directly to complexity measures appropriate to qubit processes and to our identifying the mechanism in quantum measurement that generates the observed complexity.

We concern ourselves with qubit sources that generate single qubits at a time, putting aside entangled pairs for now, and we ask the source to determine which qubit property, out of a finite set, is generated at each time. We imagine that the on-demand source is used repeatedly, producing an arbitrarily-long time series, which we call a *qubit process*. A simple example arises when monitoring sequential emissions from a blinking quantum dot [17, 18]. Qubit processes are our main object of study; we refer to their generators as *controlled qubit sources* (CQSs). We ask how random and structured they appear to an experimentalist. (SM I highlights the features of the quantum formalism we use.)

To make headway analyzing qubit process complexity we introduce a qubit source that is classically controlled and classically measured—for short, *classically-controlled qubit source* (cCQS). That is, the qubit generator proper is sandwiched between controller and measurement apparatus, both built out of classical observables. The controller determines the kind of qubit generated at each time. Figure 1 (top) illustrates the setup: a black box, representing a quantum system, emits a qubit in quantum state ρ_t at each time t. More concretely, the lower panel shows an example, revealing that the controller inside the black box is a finite-state hidden Markov model (HMM) that emits qubits in various quantum states.

_____

[*] avenegasli@ucdavis.edu
[†] amjurgens@ucdavis.edu
[‡] chaos@ucdavis.edu

We restrict the qubit states emitted by the cCQS to be pure-state density matrices; that is, $\rho^2 = \rho$. This limits the type of correlations that can be present across the qubit time series to classical correlations and leaves time series with temporal entanglement for future exploration. Since each qubit ρ_t is in a pure state, the quantum state of the random variable chain that forms the time series can be regarded as the tensor product of the individual qubits: $\ldots \rho_{t-2} \otimes \rho_{t-1} \otimes \rho_t \ldots$.
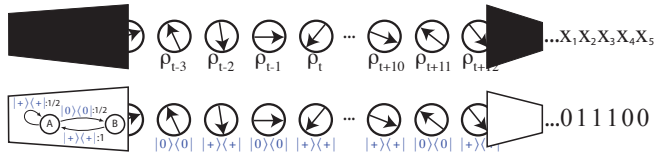

FIG. 1. Two kinds of qubit source: (Top) A general controlled qubit source (CQS) as a discrete-time quantum dynamical system, represented by a black box, generates a time series of qubits $\rho_{t-2}\rho_{t-1}\rho_t \ldots$. (Bottom) A classically-controlled qubit source (cCQS) generates a qubit process $|0\rangle \langle 0| \, |+\rangle \langle +| \ldots$. Measuring each qubit, an observer sees a classical stochastic process: (Top) $\ldots x_1 x_2 x_3 x_4 x_5$, (Bottom) $\ldots 011100$. What can we learn from the classical process about the underlying cCQS?

This simple setup raises several natural questions about characterizing qubit processes generated by CQSs. How random is the qubit process? How much memory does the source use to generate the qubit series? Can we identify the internal control mechanism from the qubit time series alone?

Now, if one replaces the CQS with a classical system that emits symbols X taking values in a discrete alphabet ($x \in \mathcal{A}$), the output is a *stochastic process* $\Pr(X_{-\infty:0}, X_{0:\infty})$ over pasts $X_{-\infty:0}$ and futures $X_{0:\infty}$. Here, X_t denotes the random variable at time t and a block is denoted $X_{t:t+l} = X_t, X_{t+1}, \ldots, X_{t+l-1}$. In this classical setting, all of the questions posed above can be answered constructively. For example, the process' *Shannon entropy rate h_μ*:

$$h_\mu = \lim_{\ell \to \infty} \frac{\mathrm{H}[\Pr(X_{0:\ell})]}{\ell}$$

measures a process' randomness as the rate of increase of information in length-ℓ sequences—in the *Shannon block entropy* $\mathrm{H}[\Pr(X_{0:\ell})]$.

Determining a process' memory requires analyzing a particular kind of edge-labeled HMM generator. First, this HMM is *unifilar*: for each state $s_k \in \mathcal{S}$ and each symbol x there is at most one transition from s_k that emits x. Second, its states are *probabilistically distinct*: For every pair of states $s_k, s_j \in \mathcal{S}$ there exists some finite word $w = x_0 x_1 \ldots x_{\ell-1}$ such that $\Pr(w|s_k) \neq \Pr(w|s_j)$. These properties define a process' *causal states*, they capture the

minimal amount of information from the past to optimally predict the process' future. Together with their transition dynamic $\{T^{(x)} : x \in \mathcal{A}\}$ the causal states form the process' minimal optimally-predictive model—its ϵ-*machine* [19]. (Figures 2 (a), (b), and (c) give examples of unifilar and nonunifilar HMMs.) Though a seemingly-innocent structural property, we show that unifilarity plays a decisive role in quantum measurement.

Most immediately, in the classical setting unifilarity allows one to calculate the entropy rate directly from the ϵ-machine as the state-averaged symbol uncertainty:

$$h_\mu = -\sum_{s \in \mathcal{S}} \Pr(s) \sum_{x \in \mathcal{A}} \sum_{s' \in \mathcal{S}} T_{ss'}^{(x)} \log T_{ss'}^{(x)} . \qquad (1)$$

Given that a process generates information at rate h_μ bits per measurement, one is next interested in the resources required to predict measurements. This is given by the *statistical complexity C_μ*—the Shannon information or average memory in the causal states:

$$C_\mu = -\sum_{s \in \mathcal{S}} \Pr(s) \log_2 \Pr(s) . \qquad (2)$$

(See SM II A and II B.)

In light of these complexity metrics, we can now state our main result: Even with a finite-state control, *generically a CQS produces a measured qubit process whose minimal predictor requires an infinite number of causal states.* Prediction resources (C_μ) diverge; though at a quantifiable rate. We establish the result constructively, by determining h_μ and exploring C_μ's divergence for qubit processes and by identifying the driving mechanism as measurement-induced nonunifilarity. These steps require introducing novel concepts from ergodic theory and dynamical systems and new efficient algorithms, whose development appears in SM II C and II E.

Qubit Processes Generating the qubits in a time series is governed by a cCQS that, without loss of generality, we take to be an ϵ-machine for which the symbols emitted during state-to-state transitions consist of qubit-states. This choice ensures that the source's internal complexity used in generating the qubit process can be quantified. Both the entropy rate h_μ^g and the statistical complexity C_μ^g of the generating, internal-state process can be exactly computed.

The states of the qubits output by a cCQS form a stochastic process; two examples of the latter are shown in Figs. 2 (a) and (b). The caption explains how these qubit-generating state machines operate. The cCQS in (a) generates a qubit time series of orthogonal pure states $|0\rangle \langle 0|$ and $|1\rangle \langle 1|$. The cCQS in (b) generates a qubit time series of nonorthogonal pure states $|0\rangle \langle 0|$ and $|+\rangle \langle +|$, where $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$.
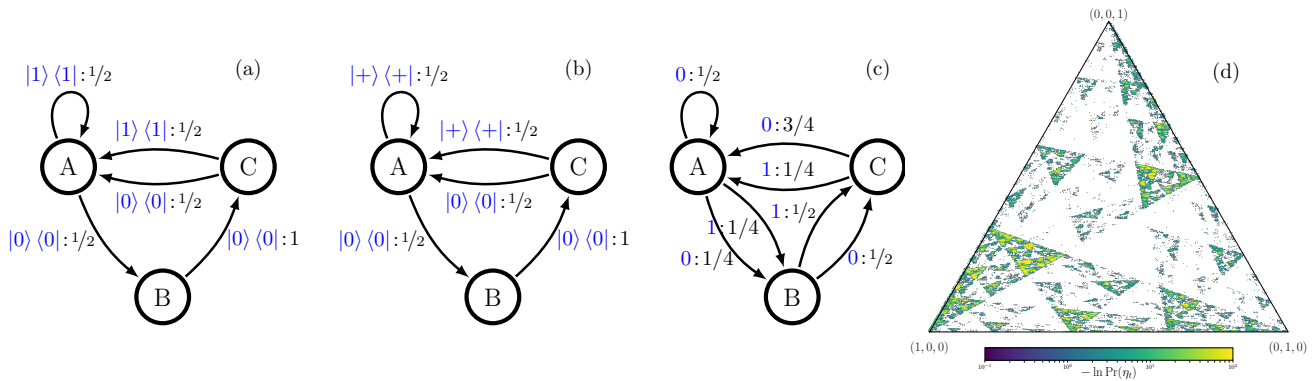
FIG. 2. (a) Three-state classically-controlled qubit source (cCQS) that generates a process consisting of orthogonal qubits: (i) controller states $\mathcal{S} = \{A, B, C\}$; (ii) orthogonal qubit alphabet $\mathcal{A}_Q = \{\rho_0 = |0\rangle\langle0|, \rho_1 = |1\rangle\langle1|\}$; (iii) labeled transition matrices $\mathbb{T}^{\rho_0}$ and $\mathbb{T}^{\rho_1}$, whose state-transition probability components $(\mathbb{T}^{\rho_k})_{ij}$ can be read from the diagram; and (iv) stationary state distribution $\pi = \begin{pmatrix} 1/2 & 1/4 & 1/4 \end{pmatrix}$. If the controller is in state A it has equal probabilities of staying in that state or transitioning to state B. If it transitions to state B, then the system emits a qubit in pure state $|0\rangle\langle0|$. In the next time step the machine must transition to state C and emits another $|0\rangle\langle0|$ qubit. Then, in state C it transitions back to state A, emitting either pure state $|0\rangle\langle0|$ or $|1\rangle\langle1|$ with equal probability. As the cCQS runs, a time series of orthogonal qubits is generated. (b) Nonorthogonal-qubit cCQS: Three-state HMM that outputs nonorthogonal qubits in pure states $|0\rangle\langle0|$ and $|+\rangle\langle+|$, where $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$. (c) HMM presentation for the classical stochastic process resulting from measurement ($\theta = \pi/2$) of the quantum process generated by (b). (d) Mixed states for the stochastic process generated by (c) in that HMM's state distribution simplex. Each mixed state is a point of the form (p_A, p_B, p_C) with probabilities of being in state A, B, or C of (c).

Measured Qubit Processes The observer interacts with such processes by applying to each qubit a projective measurement, consisting of the set of orthonormal measurement operators $\{E_0, E_1\}$ with measurement basis $E_0 = |\psi_0\rangle\langle\psi_0|$ and $E_1 = |\psi_1\rangle\langle\psi_1|$ parametrized by the Bloch angles θ and ϕ via:

$$|\psi_0\rangle = \cos\frac{\theta}{2}|0\rangle + e^{i\phi}\sin\frac{\theta}{2}|1\rangle \text{ and} \tag{3a}$$

$$|\psi_1\rangle = \sin\frac{\theta}{2}|0\rangle - e^{i\phi}\cos\frac{\theta}{2}|1\rangle \ . \tag{3b}$$

The outcome of each measurement can then be labeled 0 or 1, respectively, resulting in a binary classical stochastic process.

Knowledge of the controller and the measurement basis allows us to construct an HMM describing the measured qubit process itself. This *measured* cCQS has the same states and stationary distribution π as the original cCQS. It has labeled transition matrices $\{T^{(x)}\}$ with $x \in \mathcal{A}$:

$$T^{(x)} = \sum_{\rho_j} \mathbb{T}^{\rho_j} \Pr(x|\rho_j) \ , \tag{4}$$

where $\Pr(i|\rho_j) = tr(E_i \rho_j E_i^\dagger)$ and the cCQS labeled transition matrices $\mathbb{T}^{\rho_j}$ are defined in Fig. 2. A key step is that one can determine the HMM of the measured process by composing the measurement operator with the the qubit controller HMM that governs the cCQS. See the HMM in Fig. 2 (c). It generates the classical process resulting from measuring the qubit process generated by Fig. 2 (b)

with angles $\phi = 0$ and $\theta = \pi/2$.

Uncountable Predictive Features One would hope that, since here we know the measured cCQS—e.g., shown in Fig.2(c) for the example there—we can apply Eqs. (1)-(2) to calculate our measures directly from that model. Unfortunately, a problem arises. The measured cCQS is not an ϵ-machine since the generated measurement sequences are not in one-to-one correspondence with the internal state sequences. This is the problem of cCQS *nonunifilarity* and it stymies any attempt to directly calculate cCQS randomness and structure. In fact, and this is our first result, nonunifilarity is generic to cCQSs and even to more general qubit sources. Thus, measurement induces nonunifilarity and engenders, therefore, all of the resulting complications, whose consequences we now explore in detail.

The puzzle is that we have a model in hand and it—the measured cCQS of Fig. 2(c)—generates the measured qubit process, but we cannot use it to directly determine even the most basic process properties. What we would like and could use is an ϵ-machine that represents the measured process itself. Fortunately, the measured cCQS can be converted to an ϵ-machine by calculating the cCQS's *mixed states*; see SM II C. Here, we give a synopsis.

As first formalized by Blackwell [20], an N-state HMM's mixed states are conditional probability distributions $\eta(x_{-\ell:0}) = \Pr(\mathcal{R}_0|X_{-\ell:0} = x_{-\ell:0})$ over the measured HMM's internal states $\mathcal{R}$ given all sequences $x_{-\ell:0} \in \mathcal{A}^\ell$.

The collection over all of a process' allowed sequences induces a (Blackwell) measure μ on the state distribution $\Pr(\mathcal{R})$ N-dimensional simplex $\mathcal{R}$. The mixed states together with the mixed-state transition dynamic (see SM II C) give an HMM's *mixed-state presentation* (MSP).

A mixed state answers the question, given that one knows the HMM structure and has seen a particular sequence, what is the best guess of the internal state probabilities? Transient mixed states are those state distributions after having seen finite-ℓ sequences, while recurrent mixed states are those remaining with positive probability in the limit that $\ell \to \infty$. Recurrent mixed states exactly correspond to causal states $\mathcal{S}$ [21]. When C_μ diverges, recurrent mixed states lay in an uncountable (Cantor-like) set $S \subseteq \mathcal{R}$; see Fig. 2(d).

To emphasize, and this is our second result, measurement-induced nonunifilarity results in the number of causal states diverging. That is, despite the internal controller having only a finite number of states and the controlling cCQS being unifilar, predicting the observed process requires an uncountable number of features in the typical case. In this way, measurement induces infinite complexity that confronts an observer of a quantum process. It also introduces a new and fundamental challenge: How to define and quantify the resulting randomness and complexity?

Measurement-Induced Randomness and Statistical Complexity The uncountable number of mixed states also renders the complexity measure expressions in Eqs. (1)-(2) unusable. Fortunately, Blackwell provided a formal expression for the entropy rate [20] by showing that an HMM's mixed-state presentation is unifilar. The entropy rate is then an integral over the invariant Blackwell measure $\mu(\eta)$ in the mixed-state simplex $\mathcal{R}$:

$$h_\mu^B = -\int_{\mathcal{R}} d\mu(\eta) \sum_{x \in \mathcal{A}} \Pr(x|\eta) \log_2 \Pr(x|\eta) . \quad (5)$$

Recently, Ref. [22] introduced a constructive approach to evaluate this integral by establishing contractivity of the simplex maps—the substochastic transition matrices of Eq. (4)—and showing that the mixed-state process is ergodic. Rather than integrate over the Blackwell measure $\mu(\eta)$, such as in Fig. 2(d), these properties say that we can average over a time-series of mixed states μ_t to get the measured CQSs entropy rate:

$$\widehat{h_\mu^B} = -\lim_{\ell \to \infty} \frac{1}{\ell} \sum_{i=0}^{\ell} \sum_{x \in \mathcal{A}} \Pr(x|\eta_i) \log_2 \Pr(x|\eta_i) , \quad (6)$$

where $\Pr(x|\eta_i) = \eta(x_{0:i}) \cdot T^{(x)} \cdot \mathbf{1}$, $x_{0:i}$ is the first i symbols of an arbitrarily long sequence $x_{0:\infty}$ generated by the process, and $\mathbf{1}$ is a column-vector of all 1s. And so, we can now quantify the measured qubit process' randomness— the entropy rate of the cCQS considered as a Shannon information source.

As a CQS generates its output process, how much memory (statistical complexity C_μ) does it use? Since the number of causal states diverges, the answer requires care. Generically, in this case the causal state set $\mathcal{S}$ is uncountable and, visually, a rather complicated-looking self-similar set; again see Fig. 2(d). The consequence is that the amount memory, as monitored by C_μ, diverges. And so, instead, we track its rate of divergence—the *statistical complexity dimension* d_μ of the Blackwell measure μ on $\mathcal{R}$ [23]:

$$d_\mu = \lim_{\epsilon \to 0} -\frac{\mathrm{H}_\epsilon[\mathcal{R}]}{\log_2 \epsilon} , \quad (7)$$

where $\mathrm{H}_\epsilon[Q]$ is the Shannon entropy (in bits) of the continuous-valued random variable Q coarse-grained at size ϵ and $\mathcal{R}$ is the random variable associated with the mixed states $\eta \in \mathcal{R}$. SM II C and II E develop an upper bound on d_μ that can be accurately determined from the measured process' entropy rate $\widehat{h_\mu^B}$ above and the mixed-state process' Lyapunov characteristic exponent spectrum Λ. As discussed in SM II E, this upper bound may be close approximation to d_μ, but may also be a strict inequality. Which is the case can be easily determined from the HMM's form, as well as through comparison to direct numerical estimation of d_μ. And so, one can determine that there is a nonzero, but finite C_μ divergence rate.

Measurement Dependence Equations (3) and (4) indicate that the choice of measurement basis alters the observed process. This, in turn, implies that the process entropy rate and statistical complexity dimension also depend on measurement. To explore this with an example, we calculate the dependence of the above complexity measures as a function of measurement angle θ, with fixed ϕ for the cCQS of Fig. 2(b), determining the measured cCQS at each measurement setting.

Figure 3 (top) shows the results. The cCQS is measured in 500 different bases, holding $\phi = 0$ fixed and varying $\theta \in [0, \pi]$ uniformly. For each measured cCQS the MSP is computed and the resulting series of mixed state sets is plotted. Figure 3 (bottom) plots $\widehat{h_\mu^B}(\theta)$ and highlights three particular measurement angles $\{\theta_a, \theta_b, \theta_c\}$, showing the recurrent states found in latter's mixed-state simplices. MSP entropy rate and the statistical complexity dimension are estimated using Eqs. (6) and (7), respectively.

Common characteristics are apparent, such as a smooth behavior of $h_\mu(\theta)$ with well-defined maxima and minima and the systematic change in the MSP structure a function of θ which is consistent with the quoted dimensions d_μ. Angles $\theta = 0$ and $\theta = \pi$ give particularly simple behaviors with finite statistical complexity and $d_\mu = 0$, in accord with the countable MSPs there. The measured machines at these two values of θ are identical, aside from a symbol
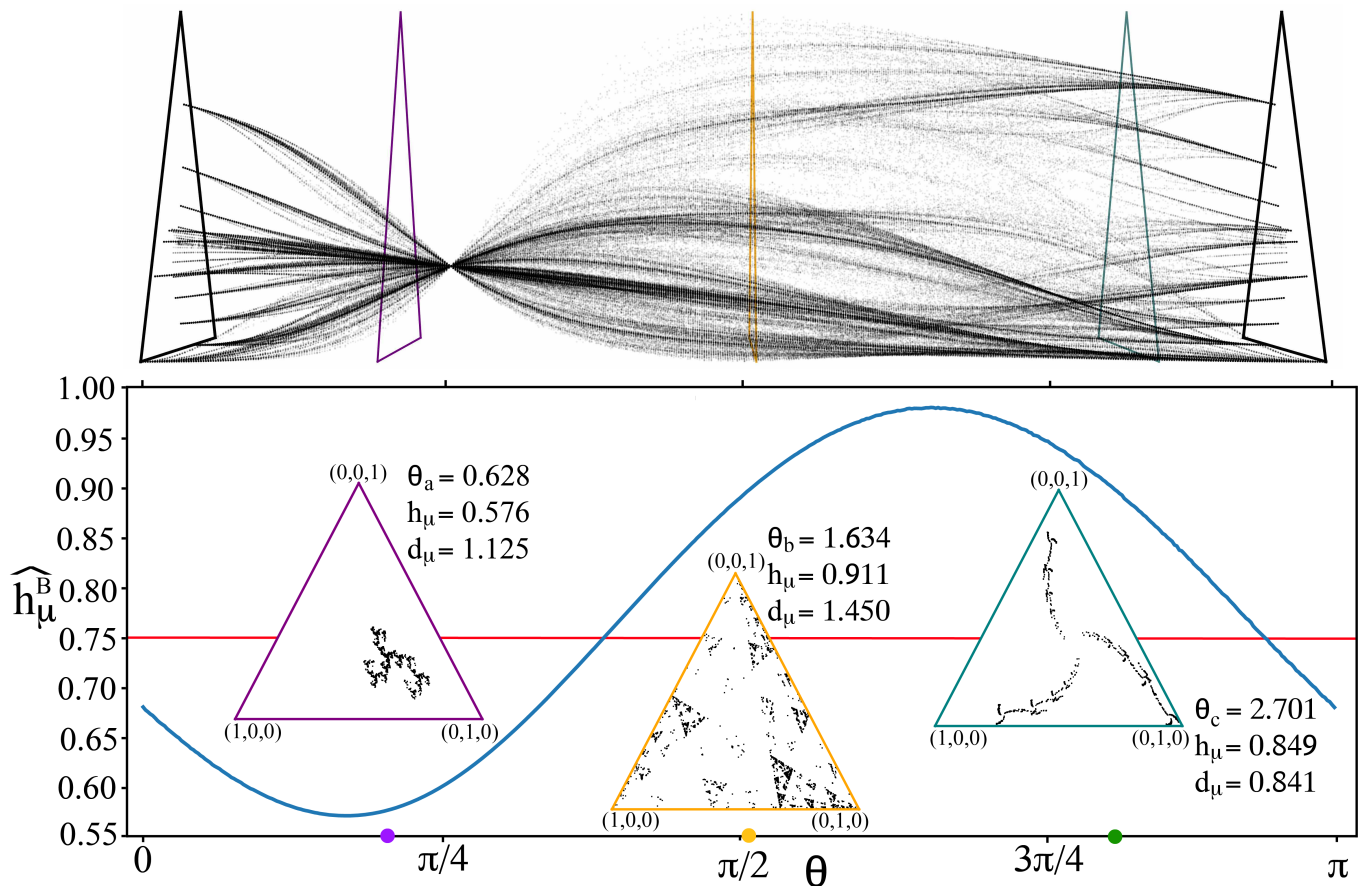
FIG. 3. Measurement-induced randomness and structure in a qubit process: (Top) Mixed state sets when measuring the qubit process of Fig. 2(b) as a function of measurement angle $\theta \in [0, \pi]$. (Bottom) Entropy rate $\widehat{h_\mu^B}$ (blue curve) as a function of angle. Horizontal line (red) is the entropy rate of the (unmeasured) qubit sequences: $h_\mu^g = 3/4$ bit per output qubit. (Insets) Mixed states of the qubit process observed at three measurement angles: (a) $\theta_a = 0.628$ (purple) (green), (b) $\theta_b = 1.634$ (orange), and (c) $\theta_c = 2.701$. The measured process entropy rates h_μ and statistical complexity dimensions d_μ given there. Both mixed states and complexity measures computed with $\ell = 10^6$ iterates.

swap—all 0's become 1's and vice versa. As such, they both have $C_\mu = 0.6813$ bits.

Figure 3 (top) exhibits a case of interest at $\theta = \pi/4$. The mixed states converge to a single point: a single-state machine that represents a biased coin. This occurs since the underlying cCQS has a binary quantum alphabet $\mathcal{A}_Q = \{\rho_0, \rho_+\}$ and the measurement basis corresponding to $\phi = 0$ and $\theta = \frac{\pi}{4}$ with basis vectors $|\psi_0\rangle$ and $|\psi_1\rangle$ is such that $\Pr(0|\rho_0) = \Pr(0|\rho_+)$ and $\Pr(1|\rho_0) = Pr(1|\rho_+)$. This basis is equidistant from both quantum states in $\mathcal{A}_Q$. Therefore, applying the measurement to one state or the other yields the same probability distribution over outcomes. One loses all information about the underlying structure and the measured cCQS generates an *independent identically distributed* process.

To compare the randomness and organization of the underlying generator process, the horizontal line in the $\widehat{h_\mu^B}(\theta)$ plot gives the entropy rate of the (unmeasured) qubit process: $h_\mu^g = 3/4$ bit per output qubit. Its statistical complexity is $C_\mu^g = \mathrm{H}[\pi] = 1.5$ bits. The differences between these constant values and those of the measured cCQS values makes it clear that quantum measurement can both add or remove randomness and structure.

Conclusion That randomness and complexity arise when observing qubit processes can be too facilely appreciated. Indeed, quantum measurement often comes steeped in mystery. We dispelled some of that mystery by showing that (i) an infinite number of predictive features are required to describe measured qubit time series and (ii) measurement both introduces and subtracts information and correlation. These characters of measurement greatly complicate learning about the informational and dynamical organization of quantum systems. However, at least now, we can appreciate more fully what the task is, what mechanism drives it (nonunifilarity), and why it is challenging.

Precisely stating how this occurs and quantitatively identifying what is added and removed was nontrivial, however, requiring several innovations. We needed to introduce new methods to estimate qubit process randomness (Shannon-Kolmogorov-Sinai entropy rate). We then had to introduce a wholly new quantity—the statistical complexity dimension—to track stored information and memory resources. And, to constructively work with these quantities required extending results from ergodic theory, abstract dynamical systems, and information theory to cCQSs. Thus, analyzing the quantum physics led us to introduce novel theory and efficient algorithms for quantifying the randomness and complexity of ergodic, stationary processes generated by nonunifilar hidden Markov models. Mathematically, these gave a constructive answer to the longstanding information-theoretic problem of identifying functions of Markov chains—a problem that until now had only been formally, not constructively, solved [20].

Supplementary Materials Materials and Methods: Reviews classical and quantum stochastic processes, information measures, and mixed states and provides details of the numerical calculations.

[1] E. Schrödinger. *Proc. Am. Philos. Soc.*, 124:323, 1935.

[2] C. S. Wu and I. Shaknov. *Phys. Rev.*, 77:136, 1950.

[3] J. S. Bell. On the Einstein Podolsky Rosen paradox. *Physics*, 1:195, 1964.

[4] A. Einstein, B. Podolsky, and N. Rosen. *Phys. Rev.*, 47:777, 1935.

[5] A. K. Ekert. *Phys. Rev. Lett.*, 67:661, 1991.

[6] C. H. Bennett, G. Brassard, C. Crepeau, R. Jozsa, A. Peres, and W. K. Wootters. *Phys. Rev. Lett.*, 70:1895, 1993.

[7] V. Giovannetti, S. Lloyd, and L. Maccone. *Phys. Rev. Lett.*, 96:010401, 2006.

[8] R. Raussendorf and H. J. Briegel. *Phys. Rev. Lett.*, 86:5188, 2001.

[9] M. Lindemann, G. Xu, T. Pusch, R. Michalzik, M. R. Hofmann, I. Zutic, and N. C. Gerhardt. *Nature*, 5678:212, 2019.

[10] S. J. Freedman and J. F. Clauser. *Phys. Rev. Lett.*, 28:938, 1972.

[11] R. Miller, T. E. Northup, K. M. Birnbaum, A. Boca, A. D. Boozer, and H. J. Kimble. *J. Phys. B: At. Mol. Opt. Phys.*, 38(9):S551–S565, 2005.

[12] P. G. Kwiat, K. Mattle, H. Weinfurter, A. Zeilinger, A. V. Sergienko, and Y. Shih. *Phys. Rev. Lett.*, 75:4337, 1995.

[13] O. Benson, C. Santori, M. Pelton, and Y. Yamamoto. *Phys. Rev. Lett.*, 84:2513, 2000.

[14] H. Walther, B. T. H. Varcoe, B.-G. Englert, and T. Becker. *Rep. Prog. Phys.*, 69(5):1325–1382, 2006.

[15] M. D. Eisaman, J. Fan, A. Migdall, and S. V. Polyakov. *Rev. Sci. Instr.*, 82(071101), 2011.

[16] F. A. Pollock, C. Rodriguez-Rosario, T. Frauenheim, M. Paternostro, and K. Modi. *Phys. Rev. A*, 97:012127, 2018.

[17] C. Galland, Y. Ghosh, A. Steinbru, M. Sykora, J. A. Hollingsworth, V. I. Klimov, and H. Htoon. *Nature*, 479:203–207, 2011.

[18] A. L. Efros and D. J. Nesbitt. *Nature Nanotech.*, 11:661–671, 2016.

[19] J. P. Crutchfield. *Nature Physics*, 8(January):17–24, 2012.

[20] D. Blackwell. In *Transactions of the first Prague conference on information theory, Statistical decision functions, Random processes*, volume 28, pages 13–20, Prague, Czechoslovakia, 1957. Publishing House of the Czechoslovak Academy of Sciences.

[21] C. J. Ellison, J. R. Mahoney, and J. P. Crutchfield. *J. Stat. Phys.*, 136(6):1005–1034, 2009.

[22] A. Jurgens and J. P. Crutchfield. *in preparation*, 2019.

[23] S. E. Marzen and J. P. Crutchfield. *Phys. Rev. E*, 95(5):051301(R), 2017.

[24] M. Gu, K. Wiesner, E. Rieper, and V. Vedral. *Nature Comm.*, 3(762):1–5, 2012.

[25] J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield. *Scientific Reports*, 6:20495, 2016.

[26] P. M. Riechers, J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield. *Phys. Rev. A*, 93(5):052317, 2016.

[27] C. Aghamohammadi, J. R. Mahoney, and J. P. Crutchfield. *Physics Letters A*, 381(14):1223–1227, 2017.

[28] S. Loomis and J. P. Crutchfield. Strong and weak optimizations in classical and quantum models of stochastic processes. *J. Stat. Phys.*, pages 1–26, 2019. arXiv:1808.08639.

[29] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, New York, 1991.

[30] C. E. Shannon. *Bell Sys. Tech. J.*, 27:379–423, 623–656, 1948.

[31] J. P. Crutchfield and D. P. Feldman. *CHAOS*, 13(1):25–54, 2003.

[32] R. G. James, C. J. Ellison, and J. P. Crutchfield. *CHAOS*, 21(3):037109, 2011.

[33] P. Frederickson, J. Kaplan, E Yorke, and J. Yorke. *J. Diff. Eqs*, 49(2):185–207, 1983.

[34] J. Kaplan and J. Yorke. In *Functional Differential Equations and Approximation of Fixed Points*, volume 730 of

Lecture Notes in Mathematics, pages 204–227. Springer, 1979.

[35] D. Feng and H. Hu. *Comm. Pure App. Math.*, 62(11):1435–1500, 2009.

[36] K. Falconer. *Fractal geometry: mathematical foundations and applications.* John Wiley, Chichester, 1990.

Supplementary Materials

Measurement-Induced Randomness and Structure in Controlled Qubit Processes

Ariadna Venegas-Li, Alexandra Jurgens, and James P. Crutchfield

The Supplementary Materials review the quantum and classical formalisms with which we work and present details, derivations, and explanations for the theoretical claims and simulation results.

I. QUANTUM FORMALISM

The quantum evolution of a qubit is typically considered to occur in the Hilbert space H_2 which contains the one-parameter (time) family of states $\rho(t)$. While this representation is appropriate for many problems, it does not accurately describe the time series of qubits that concern us. In our time series, a different qubit is emitted at each time step t. The time parameter t is discrete and labels the qubit state ρ_t emitted at time t. A different Hilbert space H_2^t contains the state of each distinct qubit—the state of the qubit emitted at time t belongs to the Hilbert space: $\rho_t \in H_2^t$. The state of the entire bi-infinite time series lies in the Hilbert space is $\mathcal{H} = \lim_{\ell \to \infty} \bigotimes_{t=-\ell}^{+\ell} H_2^t$. And, when considering a finite part of the time series of length ℓ, the Hilbert space of interest is a truncation of $\mathcal{H}$ denoted $\mathcal{H}_{t:t+l} = \bigotimes_{k=t}^{t+l-1} H_2^k$.

In our setting, a *classically controlled qubit source* (cCQS) emits a qubit at each time step. The cCQS also determines state ρ_t of each output qubit. The latter is taken to be a pure state and it remains constant until measured. The qubit chain's state then is $\ldots \rho_{t-1} \otimes \rho_t \otimes \rho_{t+1} \otimes \ldots$. These restrictions guarantee that there is no temporal entanglement and that one can apply a single-qubit projective measurement E to each output qubit without affecting the states of the other qubits in the time series.

We define measurement operators $E_0 = |\psi_0\rangle \langle\psi_0|$ and $E_1 = |\psi_1\rangle \langle\psi_1|$ with their basis vectors parametrized as in Eq. (3). As depicted in Fig. S1, applying measurement E to each qubit ρ_t yields a classical time series realization $\ldots x_{t-1} x_t x_{t+1} \ldots$ which depends on both the internal quantum process and the projective measurement E. We concentrate on quantifying randomness and structure in these observed classical stochastic processes, as they capture the classically-observed statistical properties of the internal quantum process.

The developments here complement recent progress on representing classical processes via quantum channels, which showed that quantum representations can be markedly smaller than classical [24–26] and that classical and quantum physics are at odds when it comes to measures of organization [27, 28]. In a similar, complementary way, we hope that our results aid in developing a systematic description of memoryful quantum processes built on experimentally accessible quantities [16].
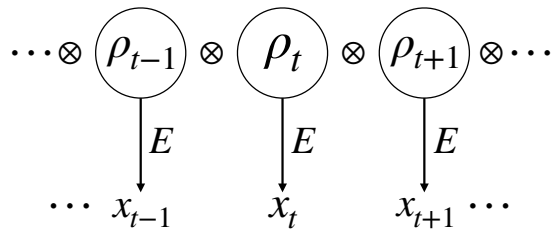

FIG. S1. A controlled qubit source emits a discrete-time stochastic process of qubits in pure states ρ_t. Applying measurement E to the qubit emitted at time t in state ρ_t yields outcome x_t. Measuring each qubit at each time step yields a classical stochastic process.

II. CLASSICAL FORMALISM AND METHODS

First, we recall notation for classical stochastic processes and their HMM generators, explicitly calling out the key subclass of unifilar HMMs. The sections then turn to information theory and computational mechanics for stationary,

ergodic processes. They review mixed-state presentations and measures of randomness and structure which are needed to characterize processes emitted by nonunifilar generators. And, they end outlining further explorations that take into account the several disparate fields involved.

A. Stochastic Processes

As many of the tools used here come from the theory of classical stochastic processes, we introduce several definitions and notation for the reader less familiar with it. A classical stochastic process $\mathbf{X}$ is a series of random variables and a specification of the probabilities of their realizations. The random variables corresponding to the behaviors are denoted by capital letters $\ldots X_{t-2}, X_{t-1}, X_t, X_{t+1}, X_{t+2} \ldots$. Their realizations are denoted by lowercase letters $\ldots x_{t-2}, x_{t-1}, x_t, x_{t+1}, x_{t+2} \ldots$, with the x_t values are drawn from a discrete alphabet $\mathcal{A}$. Blocks are denoted as: $X_{t:t+l} = X_t, X_{t+1}, \ldots X_{t+l-1}$, the left index is inclusive and the right one exclusive.

For our purposes, we consider *stationary* stochastic processes, in which the probability of observing behaviors is time-translation invariant:

$$\Pr(X_{t:t+\ell} = x_{t:t+\ell}) = \Pr(X_{0:\ell} = x_{0:\ell}) \ ,$$

for all t and ℓ. We concentrate, in particular, on processes that can be generated by a hidden Markov model.

1. Hidden Markov Models and Unifilarity

A *hidden Markov model* (HMM) is a quadruple $(\boldsymbol{\mathcal{S}}, \mathcal{A}, \{T^x\}, \pi)$ consisting of:

- $\boldsymbol{\mathcal{S}}$ is the set of hidden states.

- $\mathcal{A}$ the alphabet of symbols that the HMM emits on state-to-state transitions at each time step.

- $\{T^x : x \in \mathcal{A}\}$ is the set of labeled transition matrices such that $T_{ij}^x = \Pr(x, s_j|s_i)$ with $s_i, s_j \in \boldsymbol{\mathcal{S}}$. That is, T_{ij}^x denotes the probability of the HMM transitioning from state s_i to state s_j while emitting symbol x.

- π is the stationary state distribution determined from the left eigenvector of $T = \sum_{x \in \mathcal{A}} T^x$ normalized in probability.

Figures 2 (a), (b), and (c) are three-state HMMs over qubit alphabets $\mathcal{A}_Q = \{\rho_0 = |0\rangle \langle 0|, \rho_1 = |1\rangle \langle 1|\}$ and $\mathcal{A}_Q = \{\rho_0 = |0\rangle \langle 0|, \rho_1 = |+\rangle \langle +|\}$, and over binary alphabet $\mathcal{A} = \{0, 1\}$, respectively.

An HMM property that proves to be essential is unifilarity. An HMM is *unifilar* if, given a hidden state, the emitted symbol $x \in \mathcal{A}$ uniquely identifies the next state. Equivalently, for each labeled transition matrix T^x, there is at most one nonzero entry in each row. Unifilarity ensures that a sequence of emitted symbols has a one-to-finite correspondence with sequences of hidden-state paths. The HMMs in Figs. 2 (a) and (b) are unifilar.

In contrast, if an HMM is nonunifilar the set of allowed hidden-state state paths corresponding to a sequence of emitted symbols grows exponentially with sequence length. The HMM in Fig. 2 (c) is nonunifilar. This can be easily checked by noting that state sequences AAA, AAB, and ABC all emit the single output 000. Most basically, nonunifilarity makes inferring the underlying states and transitions from the generated output process a difficult task.

B. Measures of Complexity

We consider two complexity measures that have clear operational meanings: a process' intrinsic randomness and the minimal memory resources required to predict its behavior.

The intrinsic randomness of a classical stochastic process $\mathbf{X}$ is measured by its *entropy rate* [29]:

$$h_\mu = \lim_{\ell \to \infty} \frac{H(\ell)}{\ell} \ , \tag{S1}$$

where $H(\ell) = \mathrm{H}[\mathrm{Pr}(X_{0:\ell-1})]$ is the Shannon entropy for length-ℓ blocks. That is, a process' intrinsic randomness is the asymptotic average Shannon entropy per emitted symbol—a process' entropy growth rate.

Shannon showed that this is the same as the asymptotic value of the entropy of the next symbol conditioned on the past [30]:

$$h_\mu = \lim_{\ell \to \infty} \mathrm{H}[X_0|X_{-\ell:0}] \ . \tag{S2}$$

This can be interpreted as how much information is gained per measurement once all the possible structure in the sequence has been captured.

Determining h_μ is possible only for a small subset of stochastic processes. Shannon [30] gave closed-form expressions for processes generated by Markov Chains (MC), which are "unhidden" HMMs—they emit their states as symbols. Making use of Eq. (S2), he proved that for MC-generated processes the entropy rate is simply the average uncertainty in the next state:

$$h_\mu = -\sum_{s \in \boldsymbol{\mathcal{S}}} \mathrm{Pr}(s) \sum_{s' \in \boldsymbol{\mathcal{S}}} T_{ss'} \log T_{ss'} \ , \tag{S3}$$

where T is the MC's transition matrix and $\boldsymbol{\mathcal{S}}$ its set of states.

Another special case for which the entropy rate can be exactly computed is for processes generated by unifilar HMMs (uHMMs) [29]. This class generates an exponentially larger set of processes than possible from MCs. Since each infinite sequence of emitted symbols corresponds to a unique sequence of internal states, or at most a finite number, the process entropy rate is that of the internal MC. And so, one (slightly) adapts Eq. (S3) to calculate h_μ for these processes. The expression is presented in Eq. (1).

A process' structure is most directly analyzed by determining its minimal predictive presentation, its ϵ-machine. A simple measure of structure is then given by the number of causal states $|\boldsymbol{\mathcal{S}}|$ or by the *statistical complexity* C_μ defined in Eq. (2), which is the Shannon information $\mathrm{H}[\boldsymbol{\mathcal{S}}]$ stored in the causal states. Since the set of causal states is minimal, C_μ measures of how much memory about the past a process remembers. Said differently, C_μ quantifies the minimum amount of memory necessary to optimally predict the process' future.

However, for processes generated by nonunifilar HMMs—such as generic measured cCQSs—both h_μ and C_μ given by Eqs. (1) and (2) are incorrect. The former overestimates the generated process' h_μ, since uncertainty in the next symbol is not in direct correspondence with the uncertainty in the next internal state. In fact, there is no exact general method to compute the entropy rate of a process generated by a generic nonunifilar HMM. One has only the formal expression of Eq. (5) which refers to a abstract measure that, until now, was not constructively determined. For related reasons, the statistical complexity C_μ given by Eq. (2) applied to that abstract measure is useless—it simply diverges.

For processes generated by nonunifilar HMMs one can take a very pragmatic approach to estimate randomness and structure from process realizations (measured or simulated time series) using information measures for sequences of finite-length (ℓ), such as reviewed in Refs. [31, 32]. This approximates the sequence statistics as an order-ℓ Markov process. The associated conditional distributions capture only finite-range correlations, becoming: $\mathrm{Pr}(X_{t:\infty}|x_{-\infty:t}) = \prod_{i=t}^{\infty} \mathrm{Pr}(X_i|X_{i-\ell} \ldots X_{i-1})$. This approach is data-intensive and the complexity estimators have poor convergence.

Addressing the shortcomings for processes generated by nonunifilar HMMs requires introducing the fundamental concepts of predictive features and a process' mixed-state presentation.

C. Calculating Mixed States

The finite Markov-order approach seems to make sense empirically. However, one would hope that, if we know the nonunfilar HMM and therefore have a model (states and transitions) that generates the process at hand, we can calculate randomness and structure directly from that model or, at least, do better than by using slowly-converging order-ℓ Markov approximations. The approach is to construct a unifilar HMM—the process' ϵ-machine—from the the nonunfilar HMM. This is done by calculating the latter's *mixed states*.

Each mixed state tracks the probability distribution over the the nonunifilar HMM's internal states, conditioned on the possible sequences of observed symbols. In other words, the mixed states represent *states of knowledge* of the nonunifilar HMM's internal states. This also allows one to compute the transition dynamic between mixed states, forming a unifilar model for the same process as generated by the original nonunifilar HMM.

Explicitly, assume that an observer has an HMM presentation M for a process $\mathcal{P}$, and before making any observations has a probabilistic knowledge of the current state—the state distribution $\eta_0 = \Pr(\mathcal{S})$. Typically, prior to observing any system output the best guess is $\eta_0 = \pi$.

Once M generates a length-ℓ word $w = x_0 x_1 \ldots x_{\ell-1}$ the observer's *state of knowledge* of M's current state can be updated to $\eta(w)$, that is:

$$\eta_s(w) \equiv \Pr(\mathcal{S}_\ell = s | X_{0:\ell} = w, \mathcal{S}_0 \sim \pi) . \tag{S4}$$

The collection of possible *states of knowledge* $\eta(w)$ form the set $\boldsymbol{\mathcal{R}}$ of M's *mixed states*:

$$\boldsymbol{\mathcal{R}} = \{\eta(w) : w \in \mathcal{A}^+, \Pr(w) > 0\} .$$

And, we have the mixed-state measure $\mu(\eta)$—the probability of being in a mixed state:

$$\Pr(\eta(w)) = \Pr(\mathcal{S}_\ell | X_{0:\ell} = w, \mathcal{S}_0 \sim \pi) \Pr(w) .$$

From this follows the probability of transitioning from $\eta(w)$ to $\eta(wx)$ on observing symbol x:

$$\Pr(\eta(wx)|\eta(w)) = \Pr(x|\mathcal{S}_\ell \sim \eta(w)) .$$

This defines the mixed-state dynamic in terms of the original process, not in terms of an HMM presentation of the latter. Together the mixed states and their dynamic give the HMM's *mixed state presentation* (MSP) $\mathcal{U} = \{\boldsymbol{\mathcal{R}}, \mathcal{W}\}$ [20].

Given an HMM presentation, though, we can explicitly calculate its MSP. The probability of generating symbol x when in mixed state η is:

$$\Pr(x|\eta) = \eta \cdot T^{(x)} \cdot \mathbf{1} , \tag{S5}$$

with $\mathbf{1}$ a column vector of 1s. Upon seeing symbol x, the current mixed state η_t is updated:

$$\eta_{t+1}(x) = \frac{\eta_t \cdot T^{(x)}}{\eta \cdot T^{(x)} \cdot \mathbf{1}} , \tag{S6}$$

with $\eta_0 = \eta(\lambda) = \pi$ and λ the null sequence.

Thus, given an HMM presentation we can calculate the mixed state of Eq. (S4) via:

$$\eta(w) = \frac{\pi \cdot T^{(w)}}{\pi \cdot T^{(w)} \cdot \mathbf{1}} .$$

The mixed-state transition dynamic is then:

$$\Pr(\eta_{t+1}, x|\eta_t) = \Pr(x|\eta_t)$$
$$= \eta_t \cdot T^{(x)} \cdot \mathbf{1} ,$$

since Eq. (S6) tells us that, by construction, the MSP is unifilar. That is, the next mixed state is a function of the previous and the emitted (observed) symbol.

Now, with a unifilar presentation one is tempted to directly apply Eqs. (1) and (2) to compute measures of randomness and structure, but another challenge prevents this. With a small number of exceptions, the MSP of a process generated by a nonunifilar HMM has an uncountable infinity of states η [23]. Practically, this means that one cannot construct the full MSP, that direct application of Eq. (1) to compute the entropy rate is not feasible, and that $|\boldsymbol{\mathcal{S}}|$ diverges and, typically, so does C_μ.

D. Entropy Rate of Nonunifilar Processes

Fortunately, when working with ergodic processes, such as those addressed here, one can accurately estimate the MSP by generating a word w_ℓ of sufficiently long length [22]. The main text addresses in some detail how to use this to circumvent the complications of uncountable mixed states when computing the entropy rate. Specifically, with the mixed states in hand computationally, accurate numerical estimation of the entropy rate of a process generated by a nonunifilar HMM is given by using the temporal average specified in Eq. (6). The development of that expression and the proof that it is correct is given in Ref. [22].

This handily addresses accurately estimating the entropy rate of nonunifilar processes. And so, we are left to tackle the issue of these process' structure with the *statistical complexity dimension.* This requires a deeper discussion.

E. Statistical Complexity Dimension

C_μ diverges for processes generated by generic HMMs, as they are nonunifilar with probability one and that, in turn, leads to an uncountable infinity of mixed states. To quantify these processes' memory resources one tracks the rate of divergence—the *statistical complexity dimension* d_μ of the Blackwell measure μ on $\mathcal{R}$:

$$d_\mu = \lim_{\epsilon \to 0} -\frac{\mathrm{H}_\epsilon[\mathcal{R}]}{\log_2 \epsilon} \ , \tag{S7}$$

where $\mathrm{H}_\epsilon[Q]$ is the Shannon entropy (in bits) of the continuous-valued random variable Q coarse-grained at size ϵ and $\mathcal{R}$ is the random variable associated with the mixed states $\eta \in \mathcal{R}$.

d_μ is determined by the measured process' entropy rate $\widehat{h_\mu^B}$, as given by Eq. (6), and the mixed-state process' *spectrum of Lyapunov characteristic exponents* (LCEs). The latter is calculated from Eq. (4)'s labeled transition matrices which map the mixed states $\eta_t \in \mathcal{R}$ according to Eq. (S6). The LCE spectrum $\Lambda = \{\lambda_1, \lambda_2, \ldots, \lambda_N : \lambda_i \geq \lambda_{i+1}\}$ is determined by time-averaging the contraction rates along the N eigendirections of this map's Jacobian. The statistical complexity dimension is then bounded by a modified form of the LCE dimension [33]:

$$d_\mu \leq d_{\mathrm{LCE}} \ , \tag{S8}$$

where:

$$d_{\mathrm{LCE}} = k - 1 + \frac{\widehat{h_\mu^B} + \sum_{i=1}^{k} \lambda_i}{|\lambda_{k+1}|} \tag{S9}$$

and k is the greatest index for which $\widehat{h_\mu^B} + \sum_{i=1}^{k} \lambda_k > 0$. Reference [22] introduces this bound for an HMM's statistical complexity dimension, interprets the conditions required for its proper use, and explains in fuller detail how to calculate an HMM's LCE spectrum.

In short, the set of mixed states generated by a generic HMM is equivalent to the fractal set defining the attractor of a nonlinear, *place-dependent iterated function system (IFS)*. Exactly calculating dimensions—say, d_μ—of such sets is known to be difficult. This is why here we adapt d_{LCE} to iterated function systems. The estimation is conjectured to be accurate in "typical systems" [34, 35]. Even so, in certain cases where the IFS does not meet the *open set condition* [35]—for example, the inset in Fig. 3 at $\theta_b = 1.634$—the relationship becomes an inequality: $d_\mu < d_{\mathrm{LCE}}$. This case, which is easily detected from an HMM's form, is discussed in more detail in Ref. [22].

That these subtleties require care is borne out by numerical checks on the estimated dimensions. We estimated the box-counting dimension d_{BC} [36], a finite-ϵ approximation to d_μ, for the inset MPSs in Fig. 3: for $\theta_a = 0.628$ $d_{\mathrm{BC}} = 1.12$, for $\theta_b = 1.634$ $d_{\mathrm{BC}} = 1.18$, and for $\theta_c = 2.701$ $d_{\mathrm{BC}} = 0.85$. Thus, for measurement angles θ_a and θ_c, there is good agreement between d_{BC} and the d_μ bound. For angle θ_b, the estimated d_μ bound is greater than the box-counting dimension; but this is as expected. To be clear, the qualitative conclusions that we draw about randomness and complexity in quantum measurement do not depend on these differences in estimation. However, we do see a time when experimentally exploring these results demands refined accuracy.

Let's close by noting that the LCE spectrum's advantage in estimating d_μ is threefold. First, LCEs are computationally faster and cheaper to numerically calculate than the box-counting dimension. Second, we conjecture that d_μ of the attractors in Fig. 3 varies smoothly with angle in general. We verified that their d_{LCE} does. While the box-counting algorithm also finds this behavior, d_{BC}'s dependence on box shape and position in these algorithms induces substantial estimation errors when sweeping through a Cantor-set family. Third, the difference between d_{LCE} and the true d_μ is meaningful, as it captures the "compressibility" of a process with uncountably infinite states. This makes the d_{LCE} bound on d_μ useful beyond merely its role as a dimension estimator.