

Causal Asymmetry in a Quantum World

Jayne Thompson,^{1,*} Andrew J. P. Garner,^{1,7} John R. Mahoney,² James P. Crutchfield,² Vlatko Vedral,^{3,1,4} and Mile Gu^{5,6,1,†}¹Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, 117543, Singapore²Complexity Sciences Center and Physics Department, University of California at Davis, One Shields Avenue, Davis, California 95616, USA³Atomic and Laser Physics, University of Oxford, Clarendon Laboratory, Oxford OX1 3PU, United Kingdom⁴Department of Physics, National University of Singapore, 3 Science Drive 2, Singapore 117543, Singapore⁵School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore 639673, Singapore⁶Complexity Institute, Nanyang Technological University, Singapore 639673, Singapore⁷Institute for Quantum Optics and Quantum Information, Austrian Academy of Sciences, Boltzmanngasse 3, A-1090 Vienna, Austria (Received 30 November 2017; revised manuscript received 8 May 2018; published 18 July 2018)

Causal asymmetry is one of the great surprises in predictive modeling: The memory required to predict the future differs from the memory required to retrodict the past. There is a privileged temporal direction for modeling a stochastic process where memory costs are minimal. Models operating in the other direction incur an unavoidable memory overhead. Here, we show that this overhead can vanish when quantum models are allowed. Quantum models forced to run in the less-natural temporal direction not only surpass their optimal classical counterparts but also any classical model running in reverse time. This holds even when the memory overhead is unbounded, resulting in quantum models with unbounded memory advantage.

DOI: [10.1103/PhysRevX.8.031013](https://doi.org/10.1103/PhysRevX.8.031013)Subject Areas: Quantum Physics,
Quantum Information

I. INTRODUCTION

How can we observe an asymmetry in the temporal order of events when physics at the quantum level is time symmetric? The source of time's barbed arrow is a long-standing puzzle in foundational science [1–4]. Causal asymmetry offers a provocative perspective [5]. It asks how Occam's razor—the principle of assuming no more causes of natural things than are both true and sufficient to explain their appearances—can privilege one particular temporal direction over another. In other words, if we want to model a process causally—such that the model makes statistically correct future predictions based only on information from the past—what is the minimum past information we must store? Are we forced to store more data if we model events in one particular temporal order over the other (see Fig. 1)?

Consider a cannonball in free fall. To model its future trajectory, we need only its current position and velocity. This remains true even when we view the process in reverse time.

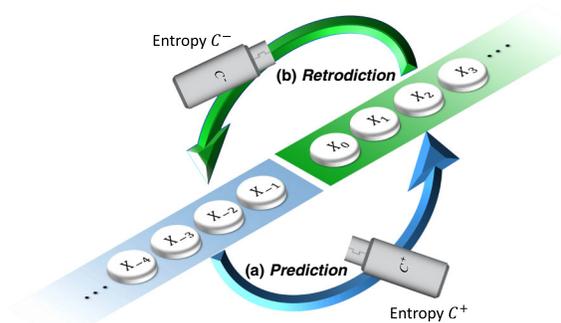


FIG. 1. A stochastic process can be modeled in either temporal order. (a) A causal model takes information available in the past \vec{x} and uses it to make statistically accurate predictions about the process's conditional future behavior $P(\vec{X}|\vec{X} = \vec{x})$. (b) A retrocausal model replicates the system's behavior, as seen by an observer who scans the outputs from right to left, encountering X_{t+1} before X_t . Thus, it stores relevant future information \vec{x} , in order to generate a statistically accurate retrodiction of the past $P(\vec{X}|\vec{X} = \vec{x})$. Causal asymmetry implies a nonzero gap between the minimum memory required by any causal model C^+ and its retrocausal counterpart C^- .

* thompson.jayne2@gmail.com

† gumile@ntu.edu.sg

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

This exemplifies causal symmetry. There is no difference in the amount of information we must track for prediction versus retrodiction. However, this is not as obvious for more complex processes. Take a glass shattering upon impact with the floor. In one temporal direction, the future distribution of shards depends only on the glass's current position, velocity, and orientation. In the opposite direction, we may need to track relevant information regarding each glass shard to infer the glass's prior trajectory. Does this require more or less information? This potential divergence is quantified in the theory of computational mechanics [6]. It is not only generally nonzero, but it can also be unbounded. This phenomenon implies that a simulator operating in the "less-natural" temporal direction is penalized with potentially unbounded memory overhead and is cited as a candidate source of time's barbed arrow [5].

These studies assumed that all models are implemented using classical physics. Could the observed causal asymmetry be a consequence of this classicality constraint? Here, we first consider a particular stochastic process that is causally asymmetric. We determine the minimal information needed to model the same process in forward versus reverse time using quantum physics, and we prove that these quantities exactly coincide. More generally, we present systematic methods to model any causally asymmetric stochastic process quantum mechanically. Critically, the resulting quantum models not only use less information than any classical counterpart but also than any classical model of the time-reversed process. Thus, quantum models can field a memory advantage that always exceeds the memory overhead incurred by causal asymmetry. Our work indicates that this overhead can emerge when imposing classical causal explanations. These results remain true even in cases where causal asymmetry becomes unbounded.

II. BACKGROUND

A. Framework

Consider a system that emits an output x_t governed by some random variable X_t at each discrete point in time t . This behavior can be described by a stochastic process \mathcal{P} —a joint probability distribution $P(\vec{X}, \vec{X})$ that correlates past behavior, $\vec{X} = \dots X_{-2} X_{-1}$, with future expectations, $\vec{X} = X_0 X_1 \dots$. Each instance of the past $\vec{x} = \dots x_{-2} x_{-1}$ exhibits a conditional future $\vec{x} = x_0 x_1 \dots$ with probability $P(\vec{X} = \vec{x} | \vec{X} = \vec{x})$.

Suppose that a model for this system can replicate this future statistical behavior using only H bits of past information. Then, this model can be executed by encoding the past \vec{x} into a state $s(\vec{x}) \in \mathcal{S}$ of a physical system Ξ of entropy H , such that repeated application of a systematic action \mathcal{M} on Ξ sequentially generates $x_0, x_1 \dots$ governed by the conditional future $P(\vec{X} | \vec{X} = \vec{x})$. The model is *causal* if, at each instance of time, all the information Ξ contains about the future can be

obtained from the past [7]. Implementing it on a computer then gives us a statistically faithful simulation of the process's realizations. The simplest causal model for a process $P(\vec{X}, \vec{X})$ is the model that minimizes H .

The statistical complexity C^+ is defined as the entropy H of this simplest model—it is the minimal amount of past information needed to make statistically correct future predictions [13,14]. This measure is used to quantify structure in diverse settings [15–17], including hidden variable models emulating quantum contextuality [18]. Here, C^+ also fields thermodynamic significance, having been linked to the minimal heat dissipation in stochastic simulation and the minimal structure a device needs to fully extract free energy from nonequilibrium environments [19–22].

Causal asymmetry captures the discrepancy in statistical complexity when a process is viewed in forward versus reverse time [23]. Consider an observer that encounters X_{t+1} before X_t . Their observations are characterized by the time-reversed stochastic process $\mathcal{P}^- = P^-(\vec{Y}, \vec{Y})$ where past and future are interchanged, such that $\vec{Y} = \dots X_1 X_0$, while $\vec{Y} = X_{-1} X_{-2} \dots$ and $Y_t = X_{-(t+1)}$. A causal model for the time-reversed process then corresponds to a retrocausal model for the forward process $P(\vec{X}, \vec{X})$. It generates a statistically accurate retrodiction of the conditional past $P(\vec{X} | \vec{X} = \vec{x})$, using only information contained in the future \vec{x} . The statistical complexity of this time-reversed process C^- (referred to as the retrodictive statistical complexity for \mathcal{P}) quantifies the minimal amount of causal information we must assign to model $P(\vec{X}, \vec{X})$ in order of decreasing t . Causal asymmetry captures the divergence $\Delta C = |C^- - C^+|$. When $\Delta C > 0$, a particular temporal direction is privileged, such that modeling the process in the other temporal direction incurs a memory overhead of ΔC .

Note that the definitions above are entropic measures and thus take operational meaning at the i.i.d. limit—i.e., modeling N instances of a stochastic process with statistical complexity C^+ requires NC^+ bits of past information, in the limit of large N . While this is the most commonly adopted measure in computational mechanics, single-shot variants do exist. The topological state complexity D^+ is particularly noteworthy [13]. It captures the minimum number of dimensions (max entropy) Ξ must have to generate future statistics. A single-shot variant of causal asymmetry can thus be defined by the difference $\Delta D = |D^- - D^+|$ between the topological state complexities of \mathcal{P}^+ and \mathcal{P}^- . Here, we focus on statistical complexity for clarity. However, many of our results also hold in this single-shot regime. We return to this when relevant.

B. Classical models

Prior studies of causal asymmetry assumed all models were classical. In this context, causal asymmetry can be explicitly demonstrated using ε -machines, the provably

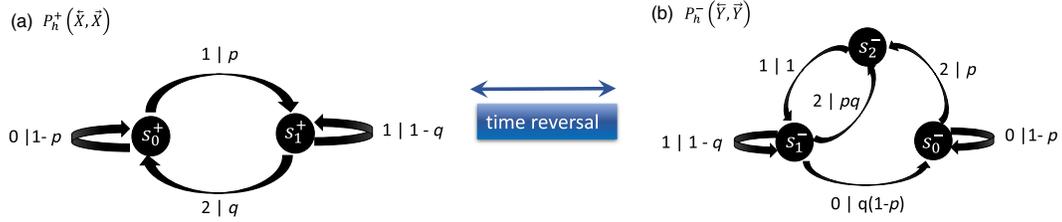


FIG. 2. (a) The ε -machine for the process $P_h^+(\bar{X}, \bar{X})$, created by flipping a biased coin and emitting outcome 2 when $H \rightarrow T$, 0 when $T \rightarrow T$, and 1 when $T/H \rightarrow H$. This process has two causal states s_1^+ and s_0^+ , where the latter includes all pasts ending in either 0 or 2. (b) The time-reversed process $P_h^-(\bar{Y}, \bar{Y})$. Here, pasts ending in 0, 1, and 2 now all lead to qualitatively different future behavior and must be stored in distinct causal states s_0^- , s_1^- , and s_2^- , respectively, which occur with respective probabilities $\pi_0^- = (q - pq)/(p + q)$, $\pi_1^- = \pi_1^+ = p/(p + q)$, and $\pi_2^- = pq/(p + q)$.

optimal classical causal models [13,14]. This involves dividing the set of pasts into equivalence classes, such that two pasts, \bar{x} and \bar{x}' , lie in the same class if and only if they have coinciding future behavior, i.e., $P(\bar{X}|\bar{X} = \bar{x}) = P(\bar{X}|\bar{X} = \bar{x}')$. Instead of recording the entire past, an ε -machine records only which equivalence class \bar{x} lies within—inducing an encoding function $\varepsilon: \bar{X} \rightarrow \mathcal{S}$ from the space of pasts \bar{X} onto the space of equivalence classes $\mathcal{S} = \{s_i\}$, known as causal states. At each time step, the machine operates according to a collection of transition probabilities T_{ij}^x : the probability that an ε -machine initially in s_i will transition to s_j while emitting output x . The classical statistical complexity thus coincides with the amount of information needed to store the current causal state

$$C_\mu^+ = -\sum_i \pi_i \log \pi_i, \quad (1)$$

where π_i is the probability the past lies within s_i . Note that ε -machines are also optimal with respect to the max entropy [24], such that the topological state complexity D_μ of a process is the logarithm of the number of causal states [13]. Despite their provable optimality, ε -machines still appear to waste memory. The amount of past information they demand typically exceeds the amount the past contains about the future—the mutual information $E = I(\bar{X}, \bar{X})$. Observing an ε -machine's entire future is insufficient for deducing its initial state. Some of the information it stores in the present is never reflected in future statistics and is thus effectively erased during operation. In general, this waste differs between prediction and retrodiction, inducing nonzero causal asymmetry.

C. Examples

We illustrate this by examples, starting with the perturbed coin. Consider a box containing a single biased coin. At each time step, the box is perturbed, causing the coin to flip with probability p if it is in heads (0) and q if it is in tails (1). The coin's state is then emitted as output. This describes a stochastic process \mathcal{P}_h^+ . As only the last output is necessary for generating correct future statistics, \mathcal{P}_h^+ has

two causal states, corresponding to the states of the coin. The statistical complexity $h(\pi_1^+)$ thus represents the entropy of the biased coin, where $\pi_1^+ = [p/(p + q)]$ is the probability the coin is in heads and $h(x) = -x \log x - (1 - x) \log(1 - x)$ is the binary entropy. Furthermore, \mathcal{P}_h^+ is clearly symmetric under time reversal (i.e., $\mathcal{P}_h^+ = \mathcal{P}_h^-$) and thus trivially causally symmetric.

Suppose we postprocess the output of the perturbed coin, replacing the first 0 of each consecutive substring of 0s with a 2 (for example, ...1000110100... becomes ...1200112120...). This results in a new stochastic process, $\mathcal{P}_h^+(\bar{X}, \bar{X})$, called the heralding coin \mathcal{P}_h^+ , which also has two causal states, $s_1^+ = \{\bar{x}|x_{-1} = 1\}$ and $s_0^+ = \{\bar{x}|x_{-1} \neq 1\}$. In fact, one can model $\mathcal{P}_h^+(\bar{X}, \bar{X})$ by perturbing the same biased coin in a box and modifying it to output 2—instead of 0—when it transitions from heads to tails (see Fig. 2). Thus, the heralding coin also has classical statistical complexity $C_\mu^+ = h(\pi_1^+)$.

Its retrodictive statistical complexity, however, is higher. The time-reversed process $\mathcal{P}_h^-(\bar{Y}, \bar{Y})$ represents an alternative postprocessing of the perturbed coin—replacing the *last* 0 in each consecutive substring of 0s with a 2. Now, 0 can be followed by 0 or 2, while 1 can be followed by anything, and 2 can only be followed by 1, inducing three causal states $s_j^- = \{\bar{y}|y_{-1} = j\}$ (see Fig. 2). This immediately establishes a difference in the number of distinct configurations needed for causal versus retrocausal modeling. Indeed, \mathcal{P}_h^+ fields causal asymmetry,

$$\Delta C_\mu = C_\mu^- - C_\mu^+ = (1 - \pi_1^-)h(\gamma), \quad (2)$$

where $\gamma = \pi_2^-/(1 - \pi_1^-)$ and $\pi_j^- = P_h^-(\bar{y} \in s_j^-)$. To understand this asymmetry, note that when modeling \mathcal{P}_h^+ , we need only know if the previous output was 1 (i.e., current state of the coin) to decide whether a 0 should be replaced by a 2. To model \mathcal{P}_h^- , however, one cannot simply look into the “future” to see if the system will output 1 next. Causal asymmetry thus captures the overhead required to accommodate this restriction.

In general, causal asymmetry can be unbounded. In Appendix D, we describe the class of n - m flower processes,

where C_μ^+ scales as $O(\log n)$ while C_μ^- scales as $O(\log m)$. Note that n and m can be adjusted independently, allowing construction of processes where $\Delta C_\mu > K$ for any given constant K . Setting $m = 2$, for example, can yield a process where C_μ^+ can be made arbitrarily high, while $C_\mu^- \leq \log 3$. When this occurs, the memory overhead incurred for modeling the process in the less-natural direction scales towards infinity.

D. Quantum models

A quantum causal model is described formally by an ordered tuple $\mathcal{Q} = (f, \Omega, \mathcal{M})$, where Ω is a set of quantum states; $f: \tilde{\mathcal{X}} \rightarrow \Omega$ defines how each past \tilde{x} is encoded into a state $f(\tilde{x}) = |s_{\tilde{x}}\rangle$ of a physical system Ξ ; and \mathcal{M} is a quantum measurement process. To model $P(\tilde{\mathcal{X}}, \tilde{\mathcal{X}})$, repeated applications of \mathcal{M} on Ξ must generate correct conditional future behavior. In other words, application of \mathcal{M} on a system Ξ in state $|s_{\tilde{x}}\rangle$ must (i) generate an output x with probability $P(X_0 = x | \tilde{\mathcal{X}} = \tilde{x})$ and (ii) transition Ξ into a new state $f(\tilde{x}') = |s_{\tilde{x}'}\rangle$, where $\tilde{x}' = \tilde{x}x$, such that L -repeated applications of \mathcal{M} will generate x_0, \dots, x_{L-1} with correct probability $P(X_{0:L} | \tilde{\mathcal{X}} = \tilde{x})$ for any desired $L \in \mathbb{Z}^+$ [25]. The entropy of a model \mathcal{Q} is given by the von Neumann entropy $S(\rho) = -\text{Tr}(\rho \log \rho)$, where $\rho = \sum P(\tilde{\mathcal{X}} = \tilde{x}) |s_{\tilde{x}}\rangle \langle s_{\tilde{x}}|$. Thus, the quantum statistical complexity C_q^+ of a process can be computed by minimizing $S(\rho)$ over all valid models [26].

This optimization is highly nontrivial. There exists no systematic techniques for constructing optimal quantum models or for proving the optimality of a given candidate model. To date, C_q^+ has only been evaluated for the Ising chain [25]. However, this process is symmetric under time reversal, implying that ΔC_μ is trivially zero. Nevertheless, recent advances show multiple settings where quantum models outperform optimal classical counterparts [27–31]. In fact, for every stochastic process where the optimal classical models are wasteful (i.e., $C_\mu^+ > E$), it is always possible to design a simpler quantum model [27]. Indeed, sometimes the quantum memory advantage $C_\mu^+ - C_q^+$ can be unbounded [32]. Could quantum models mitigate the memory overhead induced by causal asymmetry?

III. RESULTS

We study this question via two complementary approaches. The first is a case study of the heralding coin—the aforementioned process that exhibits causal asymmetry. We pioneer methods to establish its provably optimal quantum causal and retrocausal models and thus produce a precise picture of how quantum mechanics mitigates all present causal asymmetry. The second case studies quantum modeling of arbitrary processes with causal asymmetry. Here, C_q^+ and C_q^- cannot be directly evaluated but can nevertheless be bounded. In doing so, we

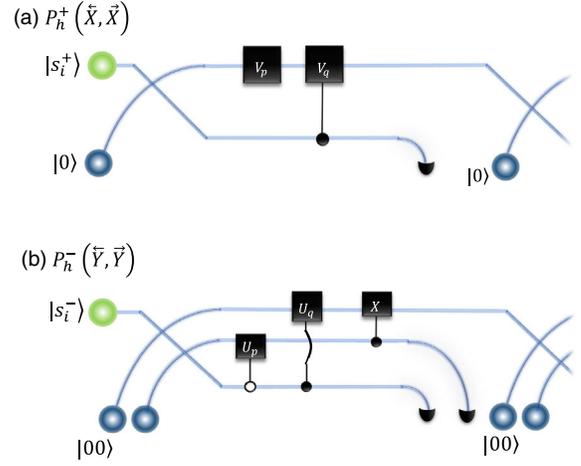


FIG. 3. Quantum circuits for generating (a) $P_h^+(\tilde{\mathcal{X}}, \tilde{\mathcal{X}})$ and (b) $P_h^-(\tilde{\mathcal{Y}}, \tilde{\mathcal{Y}})$. Here, C_U (black circle and line) is the standard control gate $C_U: |w\rangle|\psi\rangle \rightarrow |w\rangle U^{(w \bmod 2)}|\psi\rangle$. Meanwhile, \bar{C}_U (white circle, black line) is defined as $\bar{C}_U|w\rangle|\psi\rangle = |0\rangle U^{(w+1 \bmod 2)}|\psi\rangle$. (a) To simulate $P_h^+(\tilde{\mathcal{X}}, \tilde{\mathcal{X}})$, we initialize a qubit in state $|s_i^+\rangle$ and an ancilla in state $|0\rangle$. Executing the local unitary $V_p|0\rangle \rightarrow |s_0^+\rangle$, followed by the two-qubit gate C_{U_q} , where $V_q V_p|0\rangle = |s_1^+\rangle$, creates a suitable entangled state—such that a computation basis measurement of the top qubit yields x_t and simultaneously collapses the bottom qubit into the causal state for the next time step. (b) To simulate $P_h^-(\tilde{\mathcal{Y}}, \tilde{\mathcal{Y}})$, we prepare state $|s_i^-\rangle|0\rangle|0\rangle$ as input. Execution of \bar{C}_{U_p} , where $U_p|0\rangle = \sqrt{1-p}|0\rangle + \sqrt{p}|1\rangle$, followed by C_{U_q} , where U_q satisfies $U_q|0\rangle = \sqrt{q}|0\rangle + \sqrt{1-q}|1\rangle$, and finally, C_X , where X is the Pauli X operator, generates a suitable entangled state—such that measuring the first two qubits yields y_t (provided we identify measurement outcome $00 \rightarrow y_t = 0$, $10 \rightarrow y_t = 1$, and $01 \rightarrow y_t = 2$) and collapses the remaining qubit into the quantum causal state for the next time step. In either circuit, retaining only the state of Ξ (green circle) at each time step is sufficient for generating statistically correct predictions or retrodictions.

show that when forced to model such a process in the less-natural direction, the quantum advantage always exceeds the memory overhead ΔC_μ .

A. The heralding coin

Let \mathcal{P}_h^+ denote the heralding coin process. Here, we first state the optimal quantum models of \mathcal{P}_h^+ and \mathcal{P}_h^- . We then outline how their optimality is established, leaving details of the formal proof to Appendix B. The optimal causal model \mathcal{Q}^+ has two internal states:

$$\begin{aligned} |s_0^+\rangle &= \sqrt{1-p}|0\rangle + \sqrt{p}|1\rangle, \\ |s_1^+\rangle &= \sqrt{q}|2\rangle + \sqrt{1-q}|1\rangle, \end{aligned} \quad (3)$$

with associated encoding function $e_q^+(\tilde{x}) = |s_i^+\rangle$ if and only if $\tilde{x} \in s_i^+$. Given a qubit in state $e_q^+(\tilde{x})$, Fig. 3 establishes the sequential procedure that replicates expected future behavior, i.e., samples $P_h^+(\tilde{\mathcal{X}} | \tilde{\mathcal{X}} = \tilde{x})$.

Meanwhile, the optimal quantum retrocausal model \mathcal{Q}^- has encoding function $\epsilon_q^-(\vec{y}) = |s_i^-\rangle$ if and only if $\vec{y} \in s_i^-$, where

$$\begin{aligned} |s_0^-\rangle &= |0\rangle, \\ |s_1^-\rangle &= \sqrt{q}|0\rangle + \sqrt{1-q}|1\rangle, \\ |s_2^-\rangle &= |1\rangle. \end{aligned} \quad (4)$$

The associated procedure for sequential generation of \vec{y} as governed by $P_h^-(\vec{Y}|\vec{Y} = \vec{y})$ is outlined in Fig. 3.

To establish optimality, we first invoke the causal-state correspondence: For any stochastic process with causal states $\{s_i\}$ that occur with probability π_i , there exists an optimal model $\mathcal{Q} = (\epsilon_q, \Omega, \mathcal{M})$, where the elements of Ω are in one-to-one correspondence with $\{s_i\}$ (see Lemma 1 of Appendix A). Since the heralding coin process has two forward causal states, we can restrict our computation of C_q^+ to quantum models where $\Omega = \{|\psi_0^+\rangle, |\psi_1^+\rangle\}$. Moreover, we can show that the data-processing inequality implies $|\langle\psi_0^+|\psi_1^+\rangle| \leq \sqrt{p(1-q)} \equiv F$ (see Lemma 2 of Appendix A). The monotonicity between $|\langle\psi_0^+|\psi_1^+\rangle|$ and the entropy of the resulting model, together with the observation that $|\langle s_0^+|s_1^+\rangle| = F$, then implies optimality of \mathcal{Q}^+ (see Theorem 1 in Appendix). This result establishes $C_q^+ = S(\rho^+)$ for $\rho^+ = \sum_i \pi_i^+ |s_i^+\rangle\langle s_i^+|$.

Proving the optimality of \mathcal{Q}^- is more involved. First, note that the causal-state correspondence allows us to consider only candidate models $\mathcal{Q} = (f, \Omega, \mathcal{M})$, where $\Omega = \{|\psi_k^-\rangle\}_{k=0\dots 2}$ has three elements. The data processing inequality can then be used to establish the fidelity

constraints $|\langle\psi_j^-|\psi_k^-\rangle| \leq |\langle s_j^-|s_k^-\rangle|$ (see Lemma 2 of Appendix A). Let $\sigma = \sum \pi_k^- |\psi_k^-\rangle\langle\psi_k^-|$, with eigenvalues λ_k , and $\rho^- = \sum \pi_k^- |s_k^-\rangle\langle s_k^-|$, with eigenvalues λ_k^- . In Lemma 4 of Appendix B, we prove that for all choices of $|\psi_k^-\rangle$ satisfying the fidelity constraint, λ_k^- majorizes λ_k . Thus, ρ^- has minimal entropy among all valid retrocausal quantum models.

Note that \mathcal{Q}^+ and \mathcal{Q}^- exhibit different encoding functions (one maps onto two code words, the other onto three) and invoke seemingly unrelated quantum circuits for generating future statistics (see Fig. 3). Nevertheless, direct computation yields

$$C_q^+ = C_q^- = h\left(\frac{1+\sqrt{c}}{2}\right), \quad (5)$$

where $c = (p^2(1+4(1-q)q) - 2pq + q^2)/(p+q)^2$ and $h(\cdot)$ is the binary entropy. Thus, $\Delta C_q = 0$ for all values of p and q . This establishes our first result:

Result 1. There exists stochastic processes that are causally asymmetric ($C_\mu^+ \neq C_\mu^-$) but exhibit no such asymmetry when modeled quantum mechanically ($C_q^+ = C_q^-$).

This vanishing of causal asymmetry at the quantum level is not simply the result of saturating the bound given by E . Figure 4 shows that $E < C_q^+ = C_q^- < C_\mu^+ < C_\mu^-$ for almost all values of p and q . While both quantum causal and retrocausal models reduce memory resources beyond classical limits [i.e., $C_q^+ < C_\mu^+$ and $C_q^- < C_\mu^-$ —see Figs. 4(f) and 4(g)], they each still store some unnecessary information [$C_q^+, C_q^- > E$ —see Fig. 4(i)].

Our results persist when considering minimal dimensions, rather than minimal entropy required for causal modeling.

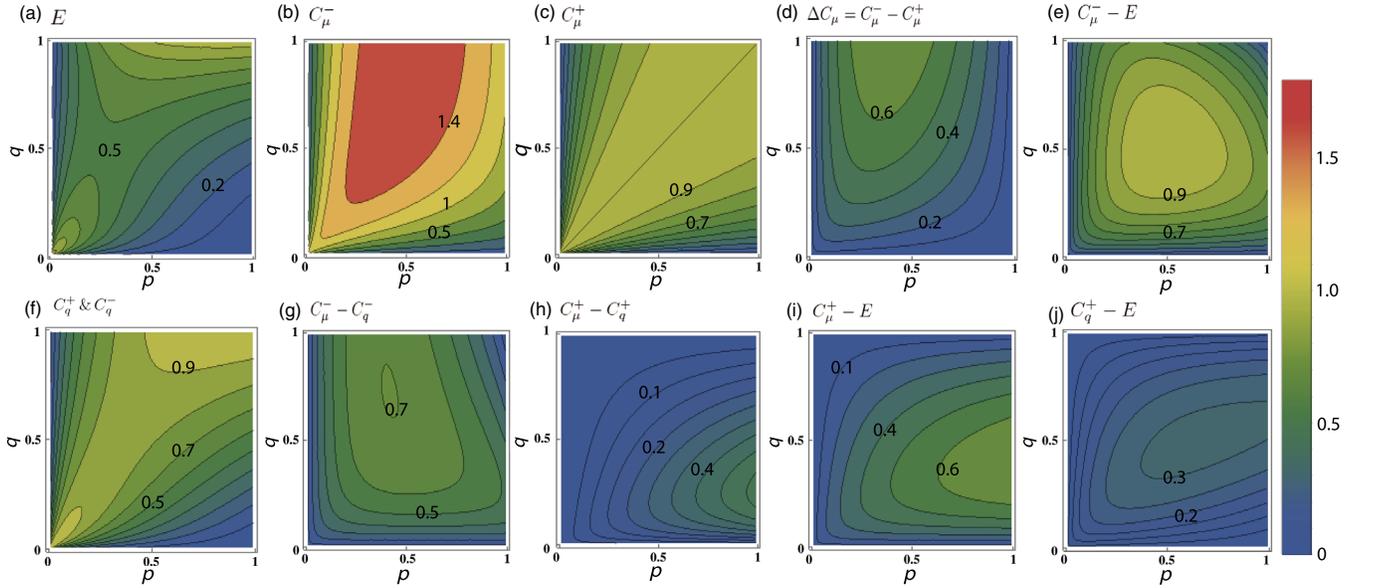


FIG. 4. Complexity of the heralding coin plotted against p and q . The figure illustrates $E \leq C_q^+ = C_q^- \leq C_\mu^+ \leq C_\mu^-$ across all values of the parameter space ($0 \leq p, q \leq 1$). Panel (d) depicts the classical causal asymmetry ΔC_μ , and panel (f) effectively demonstrates $C_q^+ = C_q^-$ and thus $\Delta C_q = 0$.

Here, \mathcal{P}_h^+ requires only two causal states and thus can be modeled using a two-level system ($D_\mu^+ = \log 2$). However, \mathcal{P}_h^- has three causal states. Modeling it thus requires a three-level system ($D_\mu^- = \log 3$). In contrast, the three quantum causal states of \mathcal{P}_h^- can be embedded within a single qubit; thus, the dynamics of the heralding coin can be modeled using a single qubit in either temporal direction. Therefore, this vanishing of causal asymmetry also applies in single-shot settings.

B. General processes

We now study quantum mitigation of causal asymmetry for general stochastic processes by bounding C_q^+ and C_q^- from above. Let $C_\mu^{\min} = \min(C_\mu^+, C_\mu^-)$ represent the minimum amount of information we need to classically model $P(\vec{X}, \vec{X})$ when allowed to optimize over the temporal direction. Meanwhile, let $C_q^{\max} = \max(C_q^+, C_q^-)$ be the minimal memory a quantum system needs when forced to model the process in the least favorable temporal direction. In Appendix C, we establish the following:

Result 2. For any stochastic process \mathcal{P} ,

$$\max(C_q^+, C_q^-) \leq \min(C_\mu^+, C_\mu^-). \quad (6)$$

Equality occurs only if $C_\mu^+ = C_\mu^- = E$, such that \mathcal{P} is causally symmetric.

Consider any causally asymmetric process \mathcal{P} , such that modeling it in the less-favorable temporal direction incurs memory overhead ΔC_μ . Result 2 implies that this overhead can be entirely mitigated by quantum models. There exists a quantum model that is not only provably simpler than its optimal classical counterpart but is also simpler than any classical model of the time-reversed process \mathcal{P}^- . In Lemma 7 (see Appendix C), we show that such models can be systematically constructed and aligned with the simplest, currently known, quantum models—q-machines [33,34]. As a corollary, causal asymmetry guarantees both $C_q^+ < C_\mu^+$ and $C_q^- < C_\mu^-$: i.e., a nonzero quantum advantage exists when modeling in either causal direction.

A variant of these results also applies to topological state complexity. Suppose the number of causal states for \mathcal{P} and its time reversal \mathcal{P}^- differ, such that $D_\mu^+ \neq D_\mu^-$. Let D_q^+ and D_q^- , respectively, be the logarithm of the minimal dimensions needed to model \mathcal{P} and \mathcal{P}^- quantum mechanically. Appendix C also establishes the following result.

Result 3. For any stochastic process \mathcal{P} ,

$$\max(D_q^+, D_q^-) \leq \min(D_\mu^+, D_\mu^-). \quad (7)$$

Given that there exists stochastic processes where predictive and retrodictive topological complexities differ (e.g., the heralding coin), this immediately implies the following corollary.

Result 4. The quantum topological complexity D_q can be strictly less than the classical topological complexity D_μ .

This solves an open question in quantum modeling—whether quantum mechanics allows for models that simulate stochastic processes using not only reduced memory but also reduced dimensions.

These results have a particular impact when ΔC_μ is exceedingly large. Recall that in the case of the $n - 2$ flower process, $C_\mu^{\min} \leq \log 3$, while C_μ^+ scales as $O(\log n)$. Our theorem then implies that $C_q^\pm \leq C_\mu^{\min} \leq \log 3$. Thus, we immediately identify a class of processes whose optimal classical models require a memory that scales as $O(\log n)$ and yet can be modeled quantum mechanically using a single qubit.

IV. FUTURE DIRECTIONS

There are a number of potential relations between causal asymmetry and innovations on the arrow of time and retrodictive quantum theory. In this section, we survey some of these connections and highlight promising future research directions.

A. Retrodictive quantum mechanics

Consider the evolution of an open quantum system that is monitored continuously in time. Standard quantum trajectory theory describes how the system’s internal state $\rho(t)$ evolves, encapsulating how our expectations of future measurement outcomes are updated based on past observations. Retrodictive quantum mechanics introduces the effect matrix $E(t)$ —a time-reversed analogue of the density matrix $\rho(t)$ [35–37]. Note that $E(t)$ propagates backwards through time, representing how our expectations of the past change as we scan future measurement outcomes in time-reversed order. The original motivation was that $\rho(t)$ and $E(t)$ combined yield a more accurate estimate of the measurement statistics at time t than $\rho(t)$ alone, allowing improved smoothing procedures [38–41].

While this framework and causal asymmetry differ in motivation and details (e.g., monitoring is done in continuous time, whereas we have so far only considered discrete time), there are also notable coinciding concepts. The standard propagation equation for $\rho(t)$ parallels a causal model for observed measurement statistics, while its time-reversed counterpart governing $E(t)$ parallels a corresponding retrocausal model. It would certainly be interesting to see if such systems exhibit either classical or quantum causal asymmetry. For example, does the resource cost of tracking $E(t)$ differ from that of $\rho(t)$ under some appropriate measure [42]?

Answering these questions will likely involve significant extensions of current results. Our framework presently assumes that the process evolves autonomously and that time is divided into discrete steps. These restrictions will need to be lifted by combining present results with recent

generalizations of classical and quantum computational mechanics to continuum time [43,44] and input-dependent regimes [21,45,46]. More generally, such developments will enable a formal study of causal asymmetry in the quantum-trajectory formulation of open quantum systems.

B. Arrow of time in quantum measurement

Related to such open systems are recent proposals for inferring an arrow of time from continuous measurement [47]. These proposals consider continuously monitoring a quantum system initialized in state ρ_i , resulting in a measurement record $r(t)$ with some probability $P[r(t)|\rho_i]$. Concurrently, the state of the system evolves through a quantum trajectory $\rho(t)$ into some final configuration $\rho(T) = \rho_f$. The goal is to identify an alternative sequence of measurements, such that for at least one possible outcome record $r'(t)$ occurring with nonzero probability $P[r'(t)|\rho_f]$, the trajectory rewinds. In other words, a system initially in state ρ_f will evolve into ρ_i , passing through all intermediary states in time-reversed order. An arrow of time emerges as $P[r(t)|\rho_i]$ and $P[r'(t)|\rho_f]$ generally differ, such that one of the two directions occurs with greater probability. An argument via Bayes' theorem then assigns different probabilistic likelihoods towards whether $\rho(t)$ occurred in forward or reverse time.

This framework provides a complementary perspective to our results. It aims to reverse the trajectory of the system's internal state $\rho(t)$, placing no constraints on the relation between the measurement statistics governing $r(t)$ and $r'(t)$. In contrast, causal asymmetry deals with reversing the observed measurement statistics (as described by some stochastic process \mathcal{P}) while placing no restrictions on the internal dynamics of the causal and retrocausal models (the two models may even field different Hilbert space dimensions, such as in the heralding coin example).

We also observe some striking parallels. Both works start out with some sequential data but no knowledge about whether the sequence occurred in forward or reverse time. Both ask the following question: Is there some sort of asymmetry singling out one temporal direction over the other? In the emerging arrow of time from quantum measurement, we are given a trajectory $\rho(t)$, and asymmetry arises from the difficulty (in terms of success probability) of realizing this trajectory in forward versus reverse time. Meanwhile, in causal asymmetry, we are given the observed measurement statistics, and an arrow of time arises from the difference in resource costs needed to realize these statistics causally in forward versus reverse time. It would then be interesting to see if a similar argument via Bayes' theorem can be adapted to causal asymmetry. If we suppose that more complex machines are less likely to exist in nature (e.g., due to dimensional or entropic constraints), could we then argue whether a given stochastic process is more likely to occur in one causal direction versus the other?

V. DISCUSSION

Causal asymmetry captures the memory overhead incurred when modeling a stochastic process in one temporal order versus the other. This induces a privileged temporal direction when one seeks the simplest causal explanation. Here, we demonstrate a process where this overhead is nonzero when using classical models and yet vanishes when quantum models are allowed. For arbitrary processes exhibiting causal asymmetry, we prove that quantum models forced to operate in a given temporal order always require less memory than classical counterparts, even when the latter are permitted to operate in either temporal direction. The former result represents a concrete case where causal asymmetry vanishes in the quantum regime. The latter implies that the more causally asymmetric a process, the greater the resource advantage of modeling it quantum mechanically.

Our results also hold when memory is quantified by max entropy. They thus establish that quantum mechanics can reduce the dimensionality needed to simulate a process beyond classical limits. Indeed, our results isolate families of processes whose statistical complexity grows without bound but can nevertheless be modeled exactly by a quantum system of bounded dimension. These features make such processes ideal for demonstrating the practical benefits of quantum models—allowing us to verify an arbitrarily large quantum advantage in single-shot regimes [24,48] and avoiding the need to measure von Neumann entropy as in current state-of-the-art experiments [29].

One compelling open question is the potential thermodynamic consequences of causal asymmetry. In computational mechanics, C_μ^+ has thermodynamical relevance in the contexts of prediction and pattern manipulation [19–22,49]. For instance, the minimum heat one must dissipate to generate future predictions based on only past observations is given by $W_{\text{diss}}^+ = k_B T (C_\mu^+ - E)$, where k_B is Boltzmann's constant, T is the environmental temperature, and the excess entropy E is symmetric with respect to time reversal. Therefore, nonzero causal asymmetry implies that flipping the temporal order in which we ascribe predictions incurs an energetic overhead of $\Delta W_{\text{diss}} = k_B T \Delta C_\mu$. In processes where ΔC_μ scales without bound, this cost may become prohibitive. Could our observation that $\Delta C_q \leq C_\mu^{\text{min}}$ imply that such energetic penalties become strongly mitigated when quantum simulators are taken into account?

A second direction is to isolate what properties of quantum processing enable it to mitigate causal asymmetry. In Appendix C, we establish that all deterministic processes are causally symmetric, such that $C_\mu^\pm = C_q^\pm = E$ (see Lemma 6 of Appendix C). Randomness is therefore essential for causal asymmetry. Observe also that the provably optimal quantum causal and retrocausal models for the heralding coin both operated unitarily—such that their dynamics are entirely deterministic (modulo

measurement of outputs). Indeed, such unitary quantum models can always be constructed [34], and we conjecture that this unitarity implies causal symmetry. However, it remains an open question as to whether the optimal quantum model is always unitary.

Insights here will ultimately help answer the big outstanding question of whether the quantum statistical complexity ever displays asymmetry under time reversal. Identifying any process for which such asymmetry persists implies that Occam’s preference for minimal cause can privilege a temporal direction in a fully quantum world. Proof that no such process exists would be equally exciting, indicating that causal asymmetry is a consequence of enforcing all causal explanations to be classical in a fundamentally quantum world.

ACKNOWLEDGMENTS

The authors appreciated the feedback and input received from: Yang Chengran, Suen Whei Yeap, Liu Qing, Alec Boyd, Varun Narasimhachar, Felix Binder, Thomas Elliott, Howard Wiseman, Geoff Pryde, Nora Tischler, Farzad Ghafari, Chiara Marletto. This work was supported by the National Research Foundation of Singapore and, in particular, NRF Awards No. NRF-NRFF2016-02, No. NRF-CRP14-2014-02, and No. RF2017-NRF-ANR004 VanQuTe, the John Templeton Foundation Grants No. 52095 and No. 54914, Foundational Questions Institute Grant No. FQXi-RFP-1609, and Physics of the Observer Grant (Observer-Dependent Complexity: The Quantum-Classical Divergence over ‘What Is Complex?’) No. FQXi-RFP-1614, the Oxford Martin School, the Singapore Ministry of Education Tier 1 RG190/17, and the U.S. Army Research Laboratory and the U.S. Army Research Office under Contracts No. W911NF-13-1-0390, No. W911NF-13-1-0340, and No. W911NF-18-1-0028. Much of the collaborative was also made possible by the “Interdisciplinary Frontiers of Quantum and Complexity Science” workshop held in Singapore, funded by the John Templeton Foundation, the Centre for Quantum Technologies, and the Lee Foundation of Singapore.

APPENDIX A: TECHNICAL DEFINITIONS

We first introduce further technical notation and background that will be used for subsequent proofs.

Definition 1 (Quantum causal model) Consider an ordered tuple $\mathcal{Q} = (f, \Omega, \mathcal{M})$, where Ω is a set of quantum states; $f: \vec{x} \rightarrow \Omega$ is an encoding function that maps each \vec{x} onto a state $f(\vec{x}) = |s_{\vec{x}}\rangle$ of a physical system Ξ ; and \mathcal{M} is a quantum process. \mathcal{Q} is a quantum model for $P(\vec{X}, \vec{X})$ if and only if, for any $\vec{x} \in \vec{\mathcal{X}}$, whenever Ξ is prepared in $f(\vec{x})$, subsequent application of \mathcal{M} (i) generates an output x with probability $P(X_0 = x | \vec{X} = \vec{x})$ and (ii) transitions Ξ into a new state $f(\vec{x}') = |s_{\vec{x}'}\rangle$, where $\vec{x}' = \vec{x}x$ [25].

Condition (i) guarantees that if a quantum model is initialized in state $f(\vec{x})$, then the model’s future output $X_0 = x$ will be statistically indistinguishable from the output of the process itself. Condition (ii) ensures that the internal memory of the quantum model is updated to record the event $X_0 = x$, allowing the model to stay synchronized with the sequence of outputs it has generated so far. Thus, a series of L repeated applications of \mathcal{M} acting on Ξ generates output $x_{0:L} = x_0 \dots x_{L-1}$ with probability $P(X_{0:L} = x_{0:L} | \vec{X} = \vec{x})$ and simultaneously transitions Ξ into the state $f(\vec{x}x_{0:L})$. In the limit $L \rightarrow \infty$, the model produces a sequence of outputs $\vec{x} = x_0 x_1 \dots$ with probability $P(\vec{X} | \vec{X} = \vec{x})$.

The entropy of a quantum model \mathcal{Q} is given by

$$C_q(\mathcal{Q}) = S(\rho) = -\text{Tr}(\rho \log \rho), \quad (\text{A1})$$

where $S(\cdot)$ is the von Neumann entropy, $\rho = \sum_{\vec{x}} \pi_{\vec{x}} \rho_{\vec{x}}$ for $\rho_{\vec{x}} = |s_{\vec{x}}\rangle\langle s_{\vec{x}}|$, and $\pi_{\vec{x}} = P(\vec{X} = \vec{x})$.

Definition 2 \mathcal{Q} is an optimal quantum model for a process $P(\vec{X}, \vec{X})$ if, given any other model \mathcal{Q}' , we have $C_q(\mathcal{Q}') \geq C_q(\mathcal{Q})$.

Consider a stationary stochastic process $P(\vec{X}, \vec{X})$, such that $P(X_{0:L}) = P(X_{t:t+L})$ for any $L \in \mathbb{Z}^+$, $t \in \mathbb{Z}$. Let $P(\vec{X}, \vec{X})$ have causal states $\mathcal{S} = \{s_i\}$ each occurring with stationary probability π_i . Define the conditional distribution $P_i(\vec{X}) = P(\vec{X} | \vec{X} = \vec{x} \in s_i)$ as the future morph of causal state s_i . We make use of the following two results derived in Ref. [25].

Lemma 1 (Causal-state correspondence). Let $P(\vec{X}, \vec{X})$ be a stochastic process with causal states $\{s_i\}$. There exists an optimal model $\mathcal{Q} = (e_q, \Omega, \mathcal{M})$, where $\Omega = \{|s_i\rangle\}$ and $e_q(\vec{x}) = |s_i\rangle$ if and only if $\vec{x} \in s_i$.

This implies that we can limit our search for optimal models $\mathcal{Q} = (f, \Omega, \mathcal{M})$ to those whose internal states $\Omega = \{|\psi_i\rangle\}$ are in one-to-one correspondence with the classical causal states. In addition, it can be shown that Ω must satisfy the following constraint:

Lemma 2 (Maximum fidelity constraint). Let $P(\vec{X}, \vec{X})$ be a stochastic process with causal states $\{s_i\}$, and let $\mathcal{Q} = (f, \Omega, \mathcal{M})$ be a valid quantum model satisfying $f(\vec{x}) = |\psi_i\rangle$ iff $\vec{x} \in s_i$. Then, $|\langle \psi_i | \psi_j \rangle| \leq F_{ij}$, where $F_{ij} = \sum_{\vec{x}} [P_i(\vec{x}) P_j(\vec{x})]^{\frac{1}{2}}$ is the fidelity between the future morphs of s_i and s_j .

These definitions assume that all elements of Ω are pure. This is because computational mechanics considers only causal models—models whose internal states do not store more information about the future than what is available from the past. Specifically, let R be a random variable governing the state of a model at $t = 0$. Then, $I(R, \vec{X} | \vec{X})$ is known as the oracular information, and it represents the amount of extra information R contains about the future \vec{X} that is not contained in the past \vec{X} . For causal models,

$I(R, \vec{X}|\vec{X}) = 0$ [8]. In Appendix E, we show that this allows us to assume that all elements of Ω are pure without loss of generality.

APPENDIX B: PROOFS OF OPTIMALITY

Here, we formally prove that the quantum models for the heralding coin given in Eqs. (3) and (4) are optimal.

1. Optimality of the causal model

Let \mathcal{P}_h^+ denote the heralding coin process, with the corresponding ε -machine depicted in Fig. 2(a).

Theorem 1. Consider $\mathcal{Q}^+ = (\varepsilon_q^+, \Omega^+, \mathcal{M}^+)$, where $\varepsilon_q^+(\vec{x}) = |s_i^+\rangle$ if and only if $\vec{x} \in s_i^+$, with

$$\begin{aligned} |s_0^+\rangle &= \sqrt{1-p}|0\rangle + \sqrt{p}|1\rangle, \\ |s_1^+\rangle &= \sqrt{q}|2\rangle + \sqrt{1-q}|1\rangle, \end{aligned} \quad (\text{B1})$$

$\Omega^+ = \{|s_0^+\rangle, |s_1^+\rangle\}$, and \mathcal{M}^+ described by the quantum circuit in Fig. 3(a). Here, \mathcal{Q}^+ is an optimal quantum model for \mathcal{P}_h^+ .

Proof.—We prove this by contradiction. Assume there exists some $\mathcal{Q} = (f, \Omega, \mathcal{M})$ such that $C_q(\mathcal{Q}) < C_q(\mathcal{Q}^+)$. Lemma 1 implies that we can assume $\Omega = \{|\psi_0\rangle, |\psi_1\rangle\}$ for some $|\psi_0\rangle$ and $|\psi_1\rangle$ and encoding function $f(\vec{x}) = |\psi_i\rangle$ if and only if $\vec{x} \in s_i^+$, without loss of generality. Here, $C_q(\mathcal{Q})$, the von Neumann entropy of the ensemble $\{|\psi_i\rangle, \pi_i^+\}$, is a monotonically decreasing function of $|\langle\psi_0|\psi_1\rangle|$ [50]. Thus, $C_q(\mathcal{Q}) < C_q(\mathcal{Q}^+)$ implies that $|\langle\psi_0|\psi_1\rangle| > |\langle s_0^+|s_1^+\rangle| = \sqrt{p(1-q)}$. Meanwhile, Lemma 2 implies

$$|\langle\psi_0|\psi_1\rangle| \leq \sum_{\vec{x}} [P_0^+(\vec{x})P_1^+(\vec{x})]^{\frac{1}{2}} = \sqrt{p(1-q)}. \quad (\text{B2})$$

This is a contradiction. Thus, no such \mathcal{Q} exists. \square

2. Optimality of the retrocausal model

Let \mathcal{P}_h^- denote the time reversal of the heralding coin process, with the corresponding ε -machine in Fig. 2(b).

Theorem 2. Define $\mathcal{Q}^- = (\varepsilon_q^-, \Omega^-, \mathcal{M}^-)$, where $\varepsilon_q^-(\vec{y}) = |s_i^-\rangle$ if and only if $\vec{y} \in s_i^-$, with

$$\begin{aligned} |s_0^-\rangle &= |0\rangle, \\ |s_1^-\rangle &= \sqrt{q}|0\rangle + \sqrt{1-q}|1\rangle, \\ |s_2^-\rangle &= |1\rangle, \end{aligned} \quad (\text{B3})$$

$\Omega^- = \{|s_i^-\rangle\}$, and the measurement process \mathcal{M}^- given in Fig. 3(b). Here, \mathcal{Q}^- is an optimal quantum model for \mathcal{P}_h^- .

Below, we break down the proof of this theorem into a series of small steps. Each step is phrased as a lemma.

Lemma 3. Let $\mathcal{Q} = (f, \Omega, \mathcal{M})$ be a quantum model for \mathcal{P}_h^- satisfying $f(\vec{y}) = |\psi_i\rangle$ iff $\vec{y} \in s_i^-$. Then, up to a unitary rotation,

$$\begin{aligned} |\psi_0\rangle &= |0\rangle, \\ |\psi_1\rangle &= r \sin \theta e^{i\omega}|0\rangle + \sqrt{1-r^2}e^{i\alpha}|1\rangle + r \cos \theta|2\rangle, \\ |\psi_2\rangle &= |1\rangle, \end{aligned} \quad (\text{B4})$$

for some $\theta \in [0, \pi/2]$, $0 \leq r \leq 1$, $\alpha, \omega \in [0, 2\pi]$, such that $r \sin \theta \leq \sqrt{q}$ and $\sqrt{1-r^2} \leq \sqrt{1-q}$.

Proof.—Set $F_{ij}^- = \sum_{\vec{y}} [P_i^-(\vec{y})P_j^-(\vec{y})]^{\frac{1}{2}}$. Explicit evaluation yields $F_{01}^- = \sqrt{q}$, $F_{02}^- = 0$, $F_{12}^- = \sqrt{1-q}$. By the maximum fidelity constraint, $|\langle\psi_i|\psi_j\rangle| \leq F_{ij}^-$. Thus, $\langle\psi_0|\psi_2\rangle = 0$. Therefore, $|\psi_0\rangle = |0\rangle$ and $|\psi_2\rangle = |1\rangle$ up to a unitary rotation. We can then write $|\psi_1\rangle$ in the form above without loss of generality, as the coefficient of $|2\rangle$ can be made real and positive by choosing a suitable definition of basis element $|2\rangle$. Meanwhile, constraints on $|\langle\psi_0|\psi_1\rangle| \leq \sqrt{1-q}$ and $|\langle\psi_0|\psi_1\rangle| \leq \sqrt{q}$ imply $r \sin \theta \leq \sqrt{q}$ and $\sqrt{1-r^2} \leq \sqrt{1-q}$. \square

Our models described in Eq. (4) can be obtained by setting $r \sin \theta e^{i\omega} = \sqrt{q}$ and $\sqrt{1-r^2}e^{i\alpha} = \sqrt{1-q}$ in Eq. (B4) (i.e., this corresponds to choosing $|\psi_0\rangle = |s_0^-\rangle$, $|\psi_1\rangle = |s_1^-\rangle$, and $|\psi_2\rangle = |s_2^-\rangle$). The subsequent lemma then establishes that this is the optimal choice.

Lemma 4. For any quantum model $\mathcal{Q} = (f, \Omega, \mathcal{M})$ of \mathcal{P}_h^- satisfying $f(\vec{y}) = |\psi_i\rangle$ if and only if $\vec{y} \in s_i^-$,

$$C_q(\mathcal{Q}) \geq C_q(\mathcal{Q}^-). \quad (\text{B5})$$

In other words, \mathcal{Q}^- , as described by Eq. (4), is the lowest entropy (optimal) model that satisfies the causal-state correspondence.

Proof.—By definition, $C_q(\mathcal{Q}^-) = S(\rho^-)$ for $\rho^- = \sum_i \pi_i^- |s_i^-\rangle\langle s_i^-|$, where $\pi_i^- = P_h^-(\vec{y} \in s_i^-)$ and the states $|s_i^-\rangle$ are given in Eq. (4). We label the eigenvalues of this state from largest to smallest by $\lambda_0^-, \lambda_1^-, \lambda_2^-$. Meanwhile, by the above lemma, $C_q(\mathcal{Q}) = S(\rho^{\psi_1})$, where

$$\rho^{\psi_1} = \pi_0^- |0\rangle\langle 0| + \pi_2^- |1\rangle\langle 1| + \pi_1^- |\psi_1\rangle\langle \psi_1|,$$

and $|\psi_1\rangle$ is described by Eq. (B4). We label the eigenvalues of ρ^{ψ_1} from largest to smallest by $\lambda_0^{\psi_1}, \lambda_1^{\psi_1}, \lambda_2^{\psi_1}$. To establish that $C_q(\mathcal{Q}) \geq C_q(\mathcal{Q}^-)$, it is sufficient to show $\lambda^- \succ \lambda^{\psi_1}$, where \succ denotes majorization [51]. This is established by proving that (1) $\lambda_0^- \geq \lambda_0^{\psi_1}$ and (2) $\lambda_0^- + \lambda_1^- \geq \lambda_0^{\psi_1} + \lambda_1^{\psi_1}$.

We begin by establishing $\lambda_0^- \geq \lambda_0^{\psi_1}$. By the minimax principle [52], the largest eigenvalue for ρ^{ψ_1} is

$$\lambda_0^{\psi_1} = \max_{|x\rangle} \langle x | \rho^{\psi_1} | x \rangle. \quad (\text{B6})$$

Suppose that this maximum is attained for some $|x\rangle = |x(t, \phi, \kappa, \eta)\rangle$ such that

$$|x\rangle = t \sin \phi e^{i\eta}|0\rangle + e^{i\kappa} \sqrt{1-t^2}|1\rangle + t \cos \phi|2\rangle, \quad (\text{B7})$$

where $\phi \in [0, \pi/2]$, $0 \leq t \leq 1$, and $\eta, \kappa \in [0, 2\pi]$. We can assume the coefficient of $|2\rangle$ is real and positive because Eq. (B6) remains unchanged when $|x\rangle \rightarrow e^{i\psi}|x\rangle$. Substituting Eq. (B7) into Eq. (B6) yields

$$\lambda_0^{\psi_1} = \pi_0^- |\langle x|0\rangle|^2 + \pi_2^- |\langle x|1\rangle|^2 + \pi_1^- |rt \sin \theta \sin \phi e^{i(\omega-\eta)} + e^{i(\alpha-\kappa)} \sqrt{1-r^2} \sqrt{1-t^2} + rt \cos \theta \cos \phi|^2. \quad (\text{B8})$$

We defined $|x(t, \phi, \kappa, \eta)\rangle$ to be the vector that maximizes Eq. (B8); thus, we have implicitly optimized over κ and η in Eq. (B8). This optimization will automatically set $e^{i(\alpha-\kappa)} = e^{i(\omega-\eta)} = 1$ [since any two complex numbers $c_1, c_2 \in \mathbb{C}$ satisfy $|c_1 + c_2|^2 \leq (|c_1| + |c_2|)^2$]. Using this and trigonometry identities to simplify Eq. (B8), we obtain

$$\lambda_0^{\psi_1} = \pi_0^- t^2 \sin^2 \phi + \pi_2^- (1 - t^2) + \pi_1^- |rt \cos(\phi - \theta) + \sqrt{1-r^2} \sqrt{1-t^2}|^2. \quad (\text{B9})$$

We now show that there always exists some $\lambda_0^{\psi_1}$ such that $\lambda_0^- \geq \lambda_0^{\psi_1} \geq \lambda_0^+$. The maximum fidelity constraint implies $r \sin(\theta) \leq \sqrt{q}$. Thus, there exists some $d\theta$ such that $r \sin(\theta + d\theta) = \sqrt{q}$ (in particular, we choose the solution of this equation where $0 < \theta + d\theta \leq \pi/2$). Consider

$$\lambda_0^{\psi_1} = \max_{|\langle x|x\rangle|^2=1} \langle x|\rho^{\psi_1}|x\rangle, \\ \rho^{\psi_1} = \pi_0^- |0\rangle\langle 0| + \pi_2^- |1\rangle\langle 1| + \pi_1^- |\psi_1'\rangle\langle \psi_1'|,$$

where

$$|\psi_1'\rangle = r \sin(\theta + d\theta)|0\rangle + \sqrt{1-r^2}|1\rangle + r \cos(\theta + d\theta)|2\rangle \\ = \sqrt{q}|0\rangle + \sin \chi \sqrt{1-q}|1\rangle + \cos \chi \sqrt{1-q}|2\rangle, \quad (\text{B10})$$

for $\sin \chi = \sqrt{1-r^2}/\sqrt{1-q}$. Furthermore, let $|x'\rangle = |x(t, \beta, 0, 0)\rangle$ be defined as

$$|x'\rangle = t \sin \beta |0\rangle + \sqrt{1-t^2} |1\rangle + t \cos \beta |2\rangle \quad (\text{B11})$$

for $\beta = \min(\pi/2, \phi + d\theta)$. Then, we have

$$\lambda_0^{\psi_1} \geq \langle x'|\rho^{\psi_1}|x'\rangle \\ = \pi_0^- t^2 \sin^2 \beta + \pi_2^- (1 - t^2) \\ + \pi_1^- |rt \cos(\beta - \theta - d\theta) + \sqrt{1-r^2} \sqrt{1-t^2}|^2 \\ \geq \pi_0^- t^2 \sin^2 \phi + \pi_2^- (1 - t^2) \\ + \pi_1^- |rt \cos(\phi - \theta) + \sqrt{1-r^2} \sqrt{1-t^2}|^2 \\ = \lambda_0^{\psi_1}, \quad (\text{B12})$$

where we have used the fact that $0 \leq \phi \leq \beta \leq \pi/2$ and $|\beta - \theta - d\theta| \leq |\phi - \theta| \leq \pi/2$. Specifically, these two conditions imply $\sin \beta \geq \sin \phi$ and $\cos(\beta - \theta - d\theta) \geq \cos(\phi - \theta) \geq 0$. Thus, we have $\lambda_0^{\psi_1} \geq \lambda_0^+$.

To show $\lambda_0^- \geq \lambda_0^{\psi_1}$, we define $|y\rangle$ to be the state satisfying $\lambda_0^{\psi_1} = \langle y|\rho^{\psi_1}|y\rangle$. In general, we can parametrize

$$|y\rangle = w e^{ia} |0\rangle + \sqrt{1-w^2} \sin \xi e^{ib} |1\rangle + \sqrt{1-w^2} \cos \xi |2\rangle$$

for $0 \leq w \leq 1$, $\xi \in [0, \pi/2]$, and $a, b \in [0, 2\pi]$. Using the same argument as in Eq. (B9), we can show $a = b = 0$ and thus

$$\lambda_0^{\psi_1} = \pi_0^- w^2 + \pi_2^- (1 - w^2) \sin^2 \xi \\ + \pi_1^- |\sqrt{1-q} \sqrt{1-w^2} \cos(\chi - \xi) + \sqrt{q} w|^2. \quad (\text{B13})$$

Define $|y'\rangle = w|0\rangle + \sqrt{1-w^2}|1\rangle$. By mirroring the analysis in Eq. (B12), we find

$$\lambda_0^- \geq \langle y'|\rho^-|y'\rangle \\ = \pi_0^- w^2 + \pi_2^- (1 - w^2) \\ + \pi_1^- |\sqrt{1-q} \sqrt{1-w^2} + \sqrt{q} w|^2 \\ \geq \lambda_0^{\psi_1}. \quad (\text{B14})$$

Together, the above results imply $\lambda_0^- \geq \lambda_0^{\psi_1} \geq \lambda_0^+$, establishing step (1), $\lambda_0^- \geq \lambda_0^{\psi_1}$.

For step (2), we must show $\lambda_0^- + \lambda_1^- \geq \lambda_0^{\psi_1} + \lambda_1^{\psi_1}$. However, by construction, ρ^- only spans a two-dimensional Hilbert space; thus, we have $\lambda_0^- + \lambda_1^- = 1$. It follows that $\lambda_0^- + \lambda_1^- \geq \lambda_0^{\psi_1} + \lambda_1^{\psi_1}$. Together, these results imply $\lambda^- \succ \lambda^{\psi_1}$ and therefore $C_q(\mathcal{Q}^-) \leq C_q(\mathcal{Q})$. \square

By Lemma 1, \mathcal{P}_h^- has an optimal quantum model that satisfies the causal-state correspondence. Meanwhile, by Lemma 4, any \mathcal{Q} satisfying the causal-state correspondence must have $C_q(\mathcal{Q}) \geq C_q(\mathcal{Q}^-)$. It follows that \mathcal{Q}^- is an optimal quantum model for \mathcal{P}_h^- .

APPENDIX C: PROOF OF RESULT 2

Here, we prove Result 2. To do this, we require some preliminary lemmas. The first connects the capacity for quantum models to improve upon their optimal classical counterparts with causal asymmetry.

Lemma 5. If the classical and quantum statistical complexities of a process \mathcal{P} coincide, such that $C_q^+ = C_\mu^+$, then \mathcal{P} is causally symmetric and $C_\mu^+ = C_\mu^- = E$.

Proof.—We first make use of the prior results, showing that whenever classical models waste information, more efficient quantum models exist [27]. Specifically, $C_\mu^+ > E$ if and only if $C_q^+ < C_\mu^+$. Thus, $C_q^+ = C_\mu^+$ implies that

$C_\mu^+ = E$. It is therefore sufficient to show that $C_\mu^+ = E$ implies $C_\mu^- = E$.

We prove this by contradiction. Assume $C_\mu^+ = E$ but $C_\mu^- > E$. Now, $C_\mu^+ = E$ implies $H(S_{-1}|\vec{X}) = 0$, where S_{-1} is the random variable governing the causal state at $t = -1$ [14]. Thus, given \vec{x} , we can find a unique s_i such that $P(\vec{X} = \vec{x}|\vec{X} = \vec{x})$ is only nonzero when $\vec{x} \in s_i$. It follows that the sets $\tau_i = \{\vec{x}|P_i(\vec{X} = \vec{x}) \neq 0\}$ form a partitioning on the space of all futures (i.e., $\tau_i \cap \tau_j = \emptyset$ for $i \neq j$).

Furthermore, any two $\vec{x}, \vec{x}' \in s_i$ satisfy $P(\vec{X}|\vec{X} = \vec{x}) = P(\vec{X}|\vec{X} = \vec{x}')$, by definition of s_i . Thus, Bayes' theorem implies that the τ_i partition the future into equivalence classes $\vec{x} \sim \vec{x}'$ if and only if $P(\vec{X}|\vec{X} = \vec{x}) = P(\vec{X}|\vec{X} = \vec{x}')$ [53]. Hence, $\{\tau_i\}$ constitute the retrocausal states. Bayes' theorem also yields $P(\vec{X} = \vec{x}|\vec{X} = \vec{x} \in \tau_i) \neq 0$ only when $\vec{x} \in s_i$. This implies $H(S_{-1}|\vec{X} = \vec{x}) = 0$, where S_{-1} governs the retrocausal state at time $t = -1$. Hence, $C_\mu^- = E$, which is a contradiction. \square

It follows as a direct corollary of this result that causal asymmetry vanishes for deterministic processes [i.e., processes where $H(\vec{X}|\vec{X}) = 0$].

Lemma 6. Any deterministic process $P(\vec{X}, \vec{X})$ has $\Delta C_\mu = 0$.

Proof.—Any deterministic process has $E = C_\mu^+$ [14,54]. Since $E \leq C_q^+ \leq C_\mu^+$, it follows that $E = C_\mu^+ = C_q^+$; thus, according to the above lemma, $\Delta C_\mu = 0$. \square

Our next lemma makes use of q-machines [33], the simplest currently known quantum models. Consider a process $\mathcal{P} = P(\vec{X}, \vec{X})$ whose classical ε -machine has a collection of causal states $\mathcal{S} = \{s_i\}$ and transition probabilities T_{ij}^x . Let k denote the cryptic order of $P(\vec{X}, \vec{X})$, defined as the smallest l such that $H(S_l|X_{0:\infty}) = 0$ [28,33,54]. The q-machine of \mathcal{P} has internal states $|S_i\rangle$ defined by a recursive relation

$$|S_i\rangle = |S_i(l=k)\rangle, \quad \text{where} \quad (\text{C1})$$

$$|S_i(l)\rangle = \sum_{x_j} \sqrt{T_{ij}^x} |x\rangle |S_j(l-1)\rangle, \quad (\text{C2})$$

and $|S_i(0)\rangle = |i\rangle$. The associated encoding function satisfies $f(\vec{x}) = |S_i\rangle$ whenever $\vec{x} \in s_i$ [28,33]. Let $\bar{C}_q = S(\rho)$ be the q-machine complexity—the amount of information a q-machine stores about the past, where $\rho = \sum_i \pi_i |S_i\rangle\langle S_i|$. Meanwhile, let the max entropy $\bar{D}_q = \log \text{tr}[\rho]^0$ be the q-machine state complexity—the minimum dimensionality of any quantum system Ξ capable of storing these internal states. Note that since q-machines are valid quantum models, $C_q^+ \leq \bar{C}_q^+$ and $C_q^- \leq \bar{C}_q^-$. Likewise, $D_q^+ \leq \bar{D}_q^+$ and $D_q^- \leq \bar{D}_q^-$. We now establish that the q-machine for \mathcal{P} and its time reversal \mathcal{P}^- have coinciding von Neumann entropies and coinciding max entropies.

Lemma 7. Let $P(\vec{X}, \vec{X})$ be a stationary stochastic process and $P^-(\vec{Y}, \vec{Y})$ its time reversal, with q-machine complexity \bar{C}_q^+ and \bar{C}_q^- , and q-machine state complexity \bar{D}_q^+ and \bar{D}_q^- , respectively. Then, $\bar{C}_q^+ = \bar{C}_q^-$ and $\bar{D}_q^+ = \bar{D}_q^-$.

Proof.—We first introduce some compact notation. Let $P(\vec{X} = \vec{x}, \vec{X} = \vec{x}) = P_{\vec{x}}$ and, similarly, $P(\vec{X} = \vec{x}|\vec{X} = \vec{x}) = P_{\vec{x}|\vec{x}}$, $P(\vec{X} = \vec{x}|S_{-1} = s_i) = P_i(\vec{x})$, as well as $P(\vec{X} = \vec{x}) = P_{\vec{x}}$ and $P(\vec{x} \in s_i) = \pi_i$.

Now let $|S_i^+\rangle$ denote the internal states of the q-machine for $P(\vec{X}, \vec{X})$, such that $\rho^+ = \sum_i \pi_i |S_i^+\rangle\langle S_i^+|$ and $\bar{C}_q^+ = S(\rho^+)$. From existing work [28,33], we know that $\lim_{l \rightarrow \infty} \langle S_i(l)|S_j(l)\rangle = \langle S_i(k)|S_j(k)\rangle$. Thus, let $\omega^+(l) = \sum_i \pi_i |S_i^+(l)\rangle\langle S_i^+(l)|$ and $\omega^+ = \lim_{l \rightarrow \infty} \omega^+(l)$ such that $\bar{C}_q^+ = S(\omega^+)$. Then,

$$\begin{aligned} \omega^+ &= \lim_{l \rightarrow \infty} \sum_i \pi_i |S_i(l)\rangle\langle S_i(l)|, \\ &= \sum_i \pi_i \sum_{\vec{x}, \vec{x}'} \sqrt{P_i(\vec{x})P_i(\vec{x}')} |\vec{x}\rangle\langle \vec{x}'|, \\ &= \sum_i \sum_{\vec{x} \in s_i} P_{\vec{x}} \sum_{\vec{x}, \vec{x}'} \sqrt{P_{\vec{x}|\vec{x}}P_{\vec{x}'|\vec{x}}} |\vec{x}\rangle\langle \vec{x}'|, \\ &= \sum_{\vec{x}, \vec{x}', \vec{x}, \vec{x}'} \sqrt{P_{\vec{x}|\vec{x}}P_{\vec{x}'|\vec{x}}P_{\vec{x}'|\vec{x}}\delta_{\vec{x}, \vec{x}'}} |\vec{x}\rangle\langle \vec{x}'|. \end{aligned} \quad (\text{C3})$$

Furthermore, the forward q-machine complexity is given by $\bar{C}_q^+ = S(\omega^+)$. A similar argument shows that \bar{C}_q^- is given by $\bar{C}_q^- = S(\omega^-)$, where

$$\omega^- = \sum_{\vec{x}, \vec{x}', \vec{x}, \vec{x}'} \sqrt{P_{\vec{x}|\vec{x}}P_{\vec{x}'|\vec{x}}P_{\vec{x}'|\vec{x}}\delta_{\vec{x}, \vec{x}'}} |\vec{x}\rangle\langle \vec{x}'|. \quad (\text{C4})$$

Consider now the pure state

$$|\psi\rangle_{\vec{X}, \vec{X}} = \sum_{\vec{x}, \vec{x}'} \sqrt{P(\vec{x}, \vec{x})} |\vec{x}\rangle |\vec{x}'\rangle, \quad (\text{C5})$$

which represents the quantum superposition, or the q sample [55], over all possible output strings of the stochastic process $P(\vec{X}, \vec{X})$, with associated density operator

$$\rho_{\vec{X}, \vec{X}} = \sum_{i,l} \sum_{\vec{x} \in s_i} \sum_{\vec{x}' \in s'_i} \sqrt{P_{\vec{x}} P_{\vec{x}'}} |\vec{x}\rangle |\vec{x}'\rangle\langle \vec{x}'| \langle \vec{x}|. \quad (\text{C6})$$

We can verify that $\omega^+ = \text{Tr}_{\vec{X}}[\rho_{\vec{X}, \vec{X}}]$ and $\omega^- = \text{Tr}_{\vec{X}}[\rho_{\vec{X}, \vec{X}}]$. Thus, $S(\omega^-) = S(\omega^+)$, and therefore, $\bar{C}_q^+ = \bar{C}_q^-$. The q-machine complexities of the forward and backward processes thus coincide.

Note that the ranks of ω^+ and ω^- must also coincide. Thus, an analogous argument establishes that

$\log \text{tr}[\rho^+]^0 = \log \text{tr}[\rho^-]^0$, indicating the two models also have the same dimensionality. Therefore, $\bar{D}_q^+ = \bar{D}_q^-$. \square

We now prove Result 2. Consider any stochastic process \mathcal{P} . First, assume \mathcal{P} is causally asymmetric, such that $\Delta C_\mu \neq 0$. Note that this implies $C_\mu^+, C_\mu^- > E$ (by Lemma 5). Meanwhile, Lemma 7 implies that $\bar{C}_q^+ = \bar{C}_q^-$. Thus, it is sufficient to show that $\bar{C}_q^+ < C_\mu^+$ and $\bar{C}_q^- < C_\mu^-$, whenever $C_\mu^+, C_\mu^- > E$.

Note that for a general process, $\bar{C}_q < C_\mu$ if and only if the q-machine has two internal states with nonzero overlap $\langle S_i | S_j \rangle > 0$ [56]. It was also previously established that whenever $C_\mu > E$, we can find some $\langle S_i(1) | S_j(1) \rangle > 0$ [27], as defined by Eq. (C1). It follows from the iterative construction that $\langle S_i | S_j \rangle > 0$, and thus $\bar{C}_q < C_\mu$. Therefore, $C_\mu^+ > E$ implies $\bar{C}_q^+ < C_\mu^+$ and $C_\mu^- > E$ implies $\bar{C}_q^- < C_\mu^-$. Hence, for any causally asymmetric \mathcal{P} ,

$$\max(C_q^+, C_q^-) < \min(C_\mu^+, C_\mu^-). \quad (\text{C7})$$

Conversely, suppose $\max(C_q^+, C_q^-) = \min(C_\mu^+, C_\mu^-)$. Without loss of generality, we can assume $C_\mu^+ \leq C_\mu^-$. This implies either (i) $C_q^+ = C_\mu^+$ or (ii) $C_q^- = C_\mu^+$. In the case of (i), direct application of Lemma 5 implies $C_\mu^+ = C_\mu^- = E$. In the case of (ii), we have $C_\mu^+ \geq \bar{C}_q^+ = \bar{C}_q^- \geq C_q^- = C_\mu^+$, which implies $C_\mu^+ = \bar{C}_q^+$. In other words, q-machines are not more efficient than ε -machines in modeling \mathcal{P} . This is true if and only if $C_q^+ = C_\mu^+$ [27,28]. Thus, Lemma 5 again implies $C_\mu^+ = C_\mu^- = E$. This completes the proof.

APPENDIX D: n - m FLOWER PROCESS

The family of n - m flower processes demonstrates how causal asymmetry can be potentially unbounded (see Fig. 5). The process has statistical complexity $C_\mu^+ = 1 + \frac{1}{2} \log[n]$. In contrast, the time-reversed process will have at most $m+1$ causal states, and thus $C_\mu^- \leq \log[m+1]$. Meanwhile, the predictive and retrodictive topological state complexities satisfy $D_\mu^+ = \log[n+1]$ and $D_\mu^- \leq \log[m+1]$. Note that n and m can be adjusted independently. Setting $m = 2$, and allowing $n \rightarrow \infty$, yields diverging C_μ^+ but finite C_μ^- . Thus, ΔC_μ also diverges to infinity. A similar divergence is witnessed for topological state complexity.

Applying Result 2, we see that C_q^+ and C_q^- are both bounded above by $\log 3$. The same is also true for D_q^+ and D_q^- . Thus, quantum models of this process can fit within a single qutrit, whether modeling in forward or reverse time. In the specific case of the former, C_μ^+ and D_μ^+ diverge to infinity. Thus, we obtain a family of processes whose quantum models field an unbounded memory advantage— in both the entropic and single-shot sense.

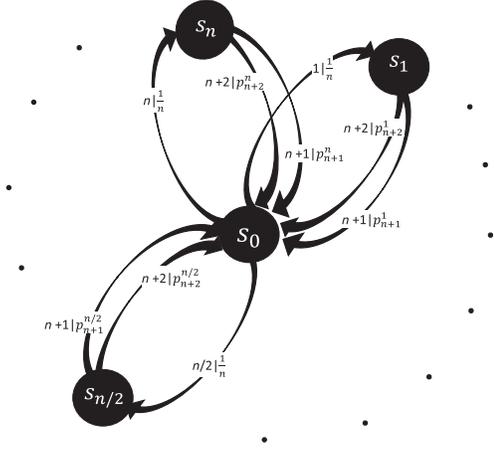


FIG. 5. The n - m flower process, illustrated for the case $m = 2$, and n even. Physically, this process can be generated by a set $\{d_1, \dots, d_n\}$ of m -sided dice, where each die d_i is biased so that it lands on side $j \in \{1, \dots, m\}$ with probability p_j^i (and, in general, the bias on each die is different, such that $p_j^i \neq p_j^k$ for $i \neq k$). We randomly select a die d_i , recording the choice $x_t = i$. Afterwards, we role the die, transcribing the outcome j as $x_{t+1} = j + n$.

APPENDIX E: EXCLUDING MIXED-STATE MODELS

In this appendix, we consider more general causal models $\mathcal{Q} = (f, \Omega, \mathcal{M})$, which have the freedom to encode pasts

$$f(\bar{x}) = \omega_{\bar{x}} = \sum_k q_k(\bar{x}) |\psi_k^{\bar{x}}\rangle \langle \psi_k^{\bar{x}}| \quad (\text{E1})$$

into mixed quantum states. We show that this does not allow for models that are more optimal than those that only encode pasts into pure quantum states.

Theorem 3. Consider a stochastic process $P(\vec{X}, \vec{X})$, with a causal model $\mathcal{Q} = (f, \Omega, \mathcal{M})$. If the internal states of \mathcal{Q} are mixed, such that $f(\bar{x}) = \sum_i q_i(\bar{x}) |\psi_i^{\bar{x}}\rangle \langle \psi_i^{\bar{x}}|$, then we can always find a causal model $\mathcal{Q}' = (f', \Omega', \mathcal{M}')$ such that $f'(\bar{x}) = |s_{\bar{x}}\rangle \langle s_{\bar{x}}|$ and $C_q(\mathcal{Q}') \leq C_q(\mathcal{Q})$.

Proof.—Let $P(\vec{X}, \vec{X})$ have causal states $\mathcal{S} = \{s_i\}$. Suppose $\mathcal{Q} = (f, \Omega, \mathcal{M})$ is an optimal causal model for $P(\vec{X}, \vec{X})$, with mixed internal states.

It is trivial to generalize the causal-state correspondence to mixed-state models. Thus, we can assume that \mathcal{Q} has an encoding function, where

$$f(\bar{x}) = \omega_i = \sum_k q_k(s_i) |\psi_k^{(i)}\rangle \langle \psi_k^{(i)}| \quad (\text{E2})$$

if and only if $\bar{x} \in s_i$. Therefore, the internal states $\Omega = \{\omega_i\}$ are in one-to-one correspondence with the classical causal states.

Our proof makes use of the requirement that causal models store no oracular information, i.e., $I(R, \vec{X} | \vec{X}) = \sum_{\bar{x}} P(\bar{x}) I(R, \vec{X} | \vec{X} = \bar{x}) = 0$, where R is the random variable

governing the memory. Regrouping the pasts into causal-state equivalence classes yields $\sum_{s_i \in \mathcal{S}} \pi_i I(R, \vec{X} | \vec{x} \in s_i) = 0$, where π_i is the probability that the past belongs to s_i . Thus, $I(R, \vec{X} | \vec{x} \in s_i) = 0$ for every $s_i \in \mathcal{S}$.

We have assumed some elements of Ω are mixed. In particular, suppose we have a specific $\omega_{\gamma} = \omega \in \Omega$ with $S(\omega) > 0$, which occurs with probability $\pi_{\gamma} = \pi$. Let \vec{x} be a particular past such that $f(\vec{x}) = \omega$, and let $\Psi = \{|\psi_k\rangle\}$ be a set of pure states that form an unraveling of ω . There must exist some $q_k \in [0, 1]$ such that $\omega = \sum_k q_k |\psi_k\rangle\langle\psi_k|$. Now, let $\mathcal{O}_{\mathcal{M}}$ be a quantum process that maps ω to a classical random variable \vec{X} governed by probability distribution $P(\vec{X} | \vec{x} = \vec{x})$. By definition of a quantum model, this process can always be constructed by concatenations of \mathcal{M} acting on a physical system Ξ .

Let A represent the state of Ξ , and let B be the random variable that governs the resulting output of $\mathcal{O}_{\mathcal{M}}$ acting on Ξ . Zero oracular information implies that A and B must be uncorrelated when conditioned on observing past \vec{x} . Therefore, $\mathcal{O}_{\mathcal{M}}(|\psi_k\rangle\langle\psi_k|) = \mathcal{O}_{\mathcal{M}}(|\psi_j\rangle\langle\psi_j|) = \mathcal{O}_{\mathcal{M}}(\omega)$ for all $|\psi_k\rangle, |\psi_j\rangle \in \Psi$.

Now, consider the entropy of \mathcal{Q} . By concavity of entropy,

$$\begin{aligned} S\left(\sum_j \pi_j \omega_j\right) &= S\left(\sum_k q_k \left(\pi |\psi_k\rangle\langle\psi_k| + \sum_{\omega_j \neq \omega} \pi_j \omega_j\right)\right) \\ &\geq \sum_k q_k S\left(\pi |\psi_k\rangle\langle\psi_k| + \sum_{\omega_j \neq \omega} \pi_j \omega_j\right) \\ &\geq \min_k S\left(\pi |\psi_k\rangle\langle\psi_k| + \sum_{\omega_j \neq \omega} \pi_j \omega_j\right). \end{aligned} \quad (\text{E3})$$

Without loss of generality, we can assume that this minimum is obtained for $k = 0$. Let $\Omega'' = (\Omega \setminus \omega) \cup \{|\psi_0\rangle\langle\psi_0|\}$ be a set of internal states, where ω is replaced with $|\psi_0\rangle\langle\psi_0|$, and define the encoding function f'' such that $f''(\vec{x}) = f(\vec{x})$, except when $f(\vec{x}) = \omega$, whereby $f''(\vec{x}) = |\psi_0\rangle\langle\psi_0|$. Define a new quantum model $\mathcal{Q}'' = (f'', \Omega'', \mathcal{M})$. Clearly, $C_q(\mathcal{Q}'') \leq C_q(\mathcal{Q})$.

If any of the states in Ω'' are still mixed, then by repeating the above procedure, we can replace them with pure states, thereby constructing a model $\mathcal{Q}' = (f', \Omega', \mathcal{M})$ with pure internal states such that $C_q(\mathcal{Q}') \leq C_q(\mathcal{Q})$. \square

[1] N. Linden, S. Popescu, A. J. Short, and A. Winter, *Quantum Mechanical Evolution Towards Thermal Equilibrium*, *Phys. Rev. E* **79**, 061103 (2009).

[2] S. Popescu, A. J. Short, and A. Winter, *Entanglement and the Foundations of Statistical Mechanics*, *Nat. Phys.* **2**, 754 (2006).

- [3] D. N. Page and W. K. Wootters, *Evolution without Evolution: Dynamics Described by Stationary Observables*, *Phys. Rev. D* **27**, 2885 (1983).
- [4] H. Price, *Time's Arrow & Archimedes' Point: New Directions for the Physics of Time* (Oxford University Press, New York, 1997).
- [5] J. P. Crutchfield, C. J. Ellison, and J. R. Mahoney, *Time's Barbed Arrow: Irreversibility, Crypticity, and Stored Information*, *Phys. Rev. Lett.* **103**, 094101 (2009).
- [6] J. P. Crutchfield, *Between Order, and Chaos*, *Nat. Phys.* **8**, 17 (2012).
- [7] Formally, this implies that the conditional mutual information $I(R, \vec{X} | \vec{X})$ is zero, where R is the random variable governing the state of Ξ . In the literature, $I(R, \vec{X} | \vec{X})$ is called oracular information, as it captures the amount of information a state of the memory s stores about the future that is not contained in the past [8,9]. Note that this notion of causal model differs from the recent developments in quantum causal models [10,11] and from foundation work on causal inference [12].
- [8] J. P. Crutchfield, C. J. Ellison, R. G. James, and J. R. Mahoney, *Synchronization and Control in Intrinsic and Designed Computation: An Information-Theoretic Analysis of Competing Models of Stochastic Computation*, *Chaos* **20**, 037105 (2010).
- [9] C. J. Ellison, J. R. Mahoney, R. G. James, J. P. Crutchfield, and J. Reichardt, *Information Symmetries in Irreversible Processes*, *Chaos* **21**, 037107 (2011).
- [10] J.-M. A. Allen, J. Barrett, D. C. Horsman, C. M. Lee, and R. W. Spekkens, *Quantum Common Causes and Quantum Causal Models*, *Phys. Rev. X* **7**, 031021 (2017).
- [11] F. Costa and S. Shrapnel, *Quantum Causal Modelling*, *New J. Phys.* **18**, 063032 (2016).
- [12] J. Pearl, *Causal Inference in Statistics: An Overview*, *Statistics Surveys* **3**, 96 (2009).
- [13] J. P. Crutchfield and K. Young, *Inferring Statistical Complexity*, *Phys. Rev. Lett.* **63**, 105 (1989).
- [14] C. R. Shalizi and J. P. Crutchfield, *Computational Mechanics: Pattern and Prediction, Structure and Simplicity*, *J. Stat. Phys.* **104**, 817 (2001).
- [15] W. M. Gonalves, R. D. Pinto, J. C. Sartorelli, and M. J. de Oliveira, *Inferring Statistical Complexity in the Dripping Faucet Experiment*, *Physica (Amsterdam)* **257A**, 385 (1998).
- [16] J. B. Park, J. W. Lee, J.-S. Yang, H.-H. Jo, and H.-T. Moon, *Complexity Analysis of the Stock Market*, *Physica (Amsterdam)* **379A**, 179 (2007).
- [17] P. Tino and M. Koteles, *Extracting Finite-State Representations from Recurrent Neural Networks Trained on Chaotic Symbolic Sequences*, *IEEE Trans. Neural Networks* **10**, 284 (1999).
- [18] A. Cabello, M. Gu, O. Gühne, and Z.-P. Xu, *Optimal Classical Simulation of State-Independent Quantum Contextuality*, *Phys. Rev. Lett.* **120**, 130401 (2018).
- [19] A. J. P. Garner, J. Thompson, V. Vedral, and M. Gu, *Thermodynamics of Complexity and Pattern Manipulation*, *Phys. Rev. E* **95**, 042140 (2017).
- [20] K. Wiesner, M. Gu, E. Rieper, and V. Vedral, *Information-theoretic lower bound on energy cost of stochastic computation*, in *Proceedings of the Royal Society of London*

- A: Mathematical, Physical and Engineering Sciences*, Vol. 468 (The Royal Society, 2012) pp. 4058–4066.
- [21] A. Cabello, M. Gu, O. Gühne, J.-Å. Larsson, and K. Wiesner, *Thermodynamical Cost of Some Interpretations of Quantum Theory*, *Phys. Rev. A* **94**, 052127 (2016).
- [22] A. B. Boyd, D. Mandal, P. M. Riechers, and J. P. Crutchfield, *Transient Dissipation and Structural Costs of Physical Information Transduction*, *Phys. Rev. Lett.* **118**, 220602 (2017).
- [23] Note that in computational mechanics literature, causal asymmetry is often referred to as causal irreversibility [5]. Here, we choose the term causal asymmetry to avoid confusion with standard notions of irreversibility used in the physics community.
- [24] R. König, R. Renner, and C. Schaffner, *The Operational Meaning of Min- and Max-Entropy*, *IEEE Trans. Inf. Theory* **55**, 4337 (2009).
- [25] W. Y. Suen, J. Thompson, A. J. P. Garner, V. Vedral, and M. Gu, *The Classical-Quantum Divergence of Complexity in the Ising Spin Chain*, *Quantum* **1**, 25 (2017).
- [26] Note that these definitions implicitly assume we can only encode onto pure states Ω . We show in Appendix E that even when we allow the encoding function to map directly onto mixed quantum states, it does not help—there is always a pure state quantum causal model with entropy $S(\rho) = C_q^+$ (see Theorem 3).
- [27] M. Gu, K. Wiesner, E. Rieper, and V. Vedral, *Quantum Mechanics Can Reduce the Complexity of Classical Models*, *Nat. Commun.* **3**, 762 (2012).
- [28] P. M. Riechers, J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield, *Minimized State Complexity of Quantum-Encoded Cryptic Processes*, *Phys. Rev. A* **93**, 052317 (2016).
- [29] M. S. Palsson, M. Gu, J. Ho, H. M. Wiseman, and G. J. Pryde, *Experimentally Modeling Stochastic Processes with Less Memory by the Use of a Quantum Processor*, *Sci. Adv.* **3**, e1601302 (2017).
- [30] R. Tan, D. R. Terno, J. Thompson, V. Vedral, and M. Gu, *Towards Quantifying Complexity with Quantum Mechanics*, *Eur. Phys. J. Plus* **129**, 191 (2014).
- [31] A. Monras and A. Winter, *Quantum Learning of Classical Stochastic Processes: The Completely Positive Realization Problem*, *J. Math. Phys. (N.Y.)* **57**, 015219 (2016).
- [32] A. J. P. Garner, Q. Liu, J. Thompson, V. Vedral, and M. Gu, *Provably Unbounded Memory Advantage in Stochastic Simulation Using Quantum Mechanics*, *New J. Phys.* **19**, 103009 (2017).
- [33] J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield, *Occam’s Quantum Strop: Synchronizing and Compressing Classical Cryptic Processes via a Quantum Channel*, *Sci. Rep.* **6**, 20495 (2016).
- [34] F. C. Binder, J. Thompson, and M. Gu, *A Practical, Unitary Simulator for Non-Markovian Complex Processes*, arXiv:1709.02375.
- [35] J. Dressel, *Weak Values as Interference Phenomena*, *Phys. Rev. A* **91**, 032116 (2015).
- [36] S. Gammelmark, B. Julsgaard, and K. Mølmer, *Past Quantum States of a Monitored System*, *Phys. Rev. Lett.* **111**, 160401 (2013).
- [37] H. M. Wiseman, *Weak Values, Quantum Trajectories, and the Cavity-QED Experiment on Wave-Particle Correlation*, *Phys. Rev. A* **65**, 032111 (2002).
- [38] D. Tan, S. J. Weber, I. Siddiqi, K. Mølmer, and K. W. Murch, *Prediction and Retrodiction for a Continuously Monitored Superconducting Qubit*, *Phys. Rev. Lett.* **114**, 090403 (2015).
- [39] T. Rybarczyk, B. Peaudecerf, M. Penasa, S. Gerlich, B. Julsgaard, K. Mølmer, S. Gleyzes, M. Brune, J. M. Raimond, S. Haroche *et al.*, *Forward-Backward Analysis of the Photon-Number Evolution in a Cavity*, *Phys. Rev. A* **91**, 062116 (2015).
- [40] S. J. Weber, A. Chantasri, J. Dressel, A. N. Jordan, K. W. Murch, and I. Siddiqi, *Mapping the Optimal Route between Two Quantum States*, *Nature (London)* **511**, 570 (2014).
- [41] P. Campagne-Ibarcq, L. Bretheau, E. Flurin, A. Auffèves, F. Mallet, and B. Huard, *Observing Interferences between Past and Future Quantum States in Resonance Fluorescence*, *Phys. Rev. Lett.* **112**, 180402 (2014).
- [42] Note that since $\rho(t)$ and $E(t)$ vary over a continuum, the memory cost of tracking either is likely to be unbounded. Thus, one may need to modify present approaches. Existing approaches include the use of differential entropies [43] and studying scaling as we track the process to greater precision [32].
- [43] S. Marzen and J. P. Crutchfield, *Informational and Causal Architecture of Continuous-Time Renewal Processes*, *J. Stat. Phys.* **168**, 109 (2017).
- [44] T. J. Elliott and M. Gu, *Superior Memory Efficiency of Quantum Devices for the Simulation of Continuous-Time Stochastic Processes*, *npj Quantum Inf.* **4**, 18 (2018).
- [45] N. Barnett and J. P. Crutchfield, *Computational Mechanics of Input–Output Processes: Structured Transformations and the ϵ -Transducer*, *J. Stat. Phys.* **161**, 404 (2015).
- [46] J. Thompson, A. J. P. Garner, V. Vedral, and M. Gu, *Using Quantum Theory to Simplify Input–Output Processes*, *npj Quantum Inf.* **3**, 6 (2017).
- [47] J. Dressel, A. Chantasri, A. N. Jordan, and A. N. Korotkov, *Arrow of Time for Continuous Quantum Measurement*, *Phys. Rev. Lett.* **119**, 220507 (2017).
- [48] O. C. O. Dahlsten, *Non-equilibrium Statistical Mechanics Inspired by Modern Information Theory*, *Entropy* **15**, 5346 (2013).
- [49] S. Still, D. A. Sivak, A. J. Bell, and G. E. Crooks, *Thermodynamics of Prediction*, *Phys. Rev. Lett.* **109**, 120604 (2012).
- [50] R. Jozsa and J. Schlienz, *Distinguishability of States and von Neumann Entropy*, *Phys. Rev. A* **62**, 012301 (2000).
- [51] M. A. Nielsen and G. Vidal, *Majorization and the Interconversion of Bipartite States*, *Quantum Inf. Comput.* **1**, 76 (2001).
- [52] R. Courant and D. Hilbert, *Methods of Mathematical Physics* (Interscience, New York, 1953), Vol. 1.
- [53] To derive this, simply write $P(\vec{X} = \vec{x} | \vec{X} = \vec{x} \in \tau_i) = \{ [P(\vec{X} = \vec{x} | \vec{X} = \vec{x})P(\vec{x})] / [\sum_{\vec{x} \in \tau_i} P(\vec{X} = \vec{x} | \vec{X} = \vec{x})P(\vec{x})] \} = \{ [P(\vec{X} = \vec{x} | \vec{X} = \vec{x})P(\vec{x})] / [\sum_{\vec{x} \in \tau_i} P(\vec{X} = \vec{x} | \vec{X} = \vec{x})P(\vec{x})] \}$.

- [54] J. R. Mahoney, C. J. Ellison, R. G. James, and J. P. Crutchfield, *How Hidden Are Hidden Processes? A Primer on Crypticity and Entropy Convergence*, *Chaos* **21**, 037112 (2011).
- [55] D. Aharonov and A. Ta-Shma, *Adiabatic Quantum State Generation and Statistical Zero Knowledge*, in *Proceedings of the 35th Annual ACM Symposium on Theory of Computing* (ACM, New York, 2003), pp. 20–29.
- [56] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2010).