

# Quantum Adaptive Agents with Efficient Long-Term Memories

Thomas J. Elliott<sup>1,2,3,\*</sup> Mile Gu<sup>3,2,4,†</sup> Andrew J. P. Garner<sup>5,3</sup> and Jayne Thompson<sup>4,6,‡</sup>

<sup>1</sup>*Department of Mathematics, Imperial College London, London SW7 2AZ, United Kingdom*


<sup>2</sup>*Complexity Institute, Nanyang Technological University, Singapore 637335*

<sup>3</sup>*Nanyang Quantum Hub, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore 637371*

<sup>4</sup>*Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, Singapore 117543*

<sup>5</sup>*Institute for Quantum Optics and Quantum Information, Austrian Academy of Sciences, Boltzmannngasse 3, Vienna 1090, Austria*

<sup>6</sup>*Horizon Quantum Computing, 05-22 Alice@Mediapolis, 29 Media Circle, Singapore 138565*

 (Received 21 August 2020; revised 5 October 2021; accepted 2 November 2021; published 11 January 2022)

Central to the success of adaptive systems is their ability to interpret signals from their environment and respond accordingly—they act as agents interacting with their surroundings. Such agents typically perform better when able to execute increasingly complex strategies. This comes with a cost: the more information the agent must recall from its past experiences, the more memory it will need. Here we investigate the power of agents capable of quantum information processing. We uncover the most general form a quantum agent need adopt to maximize memory compression advantages and provide a systematic means of encoding their memory states. We show these encodings can exhibit extremely favorable scaling advantages relative to memory-minimal classical agents, particularly when information must be retained about events increasingly far into the past.

DOI: [10.1103/PhysRevX.12.011007](https://doi.org/10.1103/PhysRevX.12.011007)

Subject Areas: Complex Systems,  
Quantum Information

## I. INTRODUCTION

The world is awash with complex, interacting systems. Predators chasing prey, investors trading stocks, grandmasters playing chess: all share in common that they process information from their environment and act in response, with an eye to achieving some desired outcome. They can be described as adaptive agents [1–5], systems that receive input stimuli and respond with output actions. This framework can be applied to a plethora of problems, including financial markets [6,7], biofilm formation [8], and HIV spread [9].

To be effective, an agent must typically adapt its future behavior based on past experiences. A rudimentary chatbot, for example, would base its response purely on the last phrase it heard—often resulting in wildly out-of-context

output. Meanwhile, a more sophisticated design would extract context from conversational history—in both what they have heard and what they have said. Tracking this contextual data requires a memory and a policy for deciding on what action to take based on the current stimulus and this memory. For agents performing elaborate tasks, effective strategies often require copious information about past data [10]; tools that ameliorate the amount of information agents must retain can thus provide a valuable competitive advantage.

To what extent can agents benefit from quantum technologies? Proof-of-principle quantum agents have demonstrated memory compression beyond classical bounds [11], yet do not make use of the full gamut of possible quantum effects. Here we identify the features of vastly improved quantum adaptive agents that use less memory—and provide a systematic procedure for their design—using insights from quantum stochastic modeling [12–20]. The resulting agents can display extreme scaling advantages over provably minimal classical counterparts [21]. We derive sufficient conditions under which such scaling advantages can occur and illustrate this with a family of scenarios where the agent’s decisions rely on events in the distant past. Complementing techniques for quantum agents to speed up the learning of effective strategies

\*[physics@tjelliott.net](mailto:physics@tjelliott.net)

†[mgu@quantumcomplexity.org](mailto:mgu@quantumcomplexity.org)

‡[thompson.jayne2@gmail.com](mailto:thompson.jayne2@gmail.com)

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

[22–24], our work illustrates that they will also be able to execute them with lower memory overhead. Together, they represent key components of quantum-enhanced artificial intelligences.

## II. FRAMEWORK

### A. Agents and strategies

We describe adaptive agents as automatons that interact with their environment at discrete time steps  $t \in \mathbb{Z}$ . At each time step the agent receives an input stimulus  $x_t \in \mathcal{X}$  and responds with output action  $y_t \in \mathcal{Y}$ , manifest by random variables  $X_t$  and  $Y_t$ , respectively (throughout, upper case indicates random variables and lower case the corresponding variates). Taking  $t = 0$  as the present, we denote the past sequences of stimuli and actions as  $\bar{x} := \dots x_{-2}x_{-1}$  and  $\bar{y} := \dots y_{-2}y_{-1}$ , respectively. For shorthand we denote the pair  $z := (x, y)$ , and similarly  $\bar{z} := (\bar{x}, \bar{y})$  for the entire history. The agent’s choice of action is governed by a strategy, describing the probability that the agent should select action  $y$  in response to stimuli  $x$  given preceding stimuli and actions  $\bar{z}$  [25]. Each strategy  $\mathcal{P}$  is thus defined by the distribution  $P(Y|\bar{Z}, X)$ ; we assume strategies to be time invariant [11,21].

To execute a desired strategy  $\mathcal{P}$ , an agent must be able to execute actions in a manner statistically faithful to the distribution for any sequence of received stimuli. This necessitates that the agent possesses a memory system  $M$  that stores relevant information from the past. A brute-force approach would be to record all past stimuli and actions, allowing a direct sampling from  $P(Y|\bar{z}, x)$ . However, storing the entire history fast becomes prohibitively expensive.

A more refined approach is to use an encoding function  $f$  that maps possible histories  $\{\bar{z}\}$  to a corresponding memory state from the set  $\{\sigma_m\}$ , labeled by  $m \in \mathcal{M}$ . Given a history  $\bar{z}$ , upon receiving any of the possible stimuli  $x$  the agent must be able to use its memory to

- (1) produce output  $y$  with probability  $P(Y|\bar{z}, x)$ ;
- (2) update the state of  $M$  to one consistent with the new history  $\bar{z}z$  [i.e.,  $f(\bar{z}z)$ ].

This process is illustrated schematically in Fig. 1(a). This requires the agent to have a policy  $\Lambda$ —a systematic procedure that governs the internal dynamics of the agent. Repeated application of  $\Lambda$  then allows the agent to execute the strategy over multiple time steps. Provided such a  $\Lambda$  exists for an encoding function  $f$ , this can be used to specify an adaptive agent. That is, the tuple  $(\mathcal{X}, \mathcal{Y}, \{\sigma_m\}, f, \Lambda)$  formally defines an adaptive agent; see the Appendix A for further details.

Since the encoding function is a deterministic mapping from histories to memory states, we are able to succinctly describe the update of the memory according to an update rule  $m' = \lambda(z, m)$ , where  $\sigma_m$  is the memory state corresponding to any given history  $\bar{z}$ ,

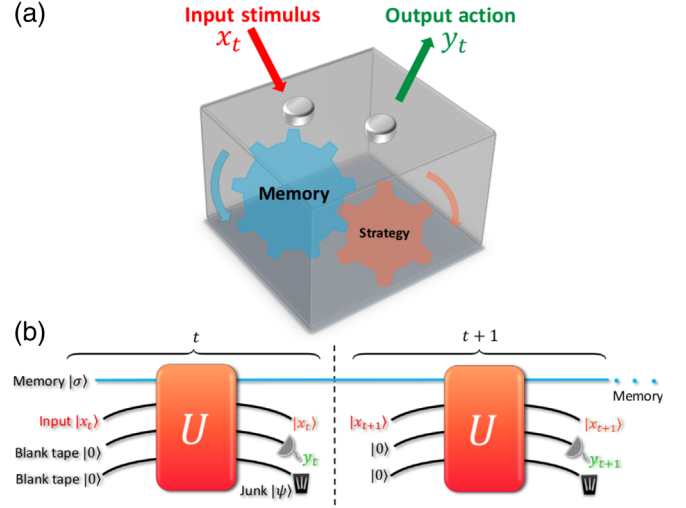


FIG. 1. Agents and their quantum realizations. (a) We consider agents that alternately receive input stimuli and perform output actions. To execute complex behavior, an agent requires a memory to keep track of relevant information about past events (both stimuli and actions) and a strategy for deciding on future actions based on this information together with the current stimulus. (b) A quantum circuit implementing a quantum agent that encompasses all memory-minimal agents (see Theorem 1). At each time step it interacts with an input stimulus encoded in  $|x_t\rangle$  and some blank tape. After the interaction, measurement of the output tape delivers the appropriate action  $y_t$ . In general, an agent must also dispose of additional redundant information, requiring junk tape that is discarded into the environment. This process can be repeated to execute the desired strategic behavior *ad infinitum*.

and  $\sigma_{m'}$  that of  $\bar{z}z$ . We can also replace the distribution  $P(Y|\bar{Z}, X)$  by  $P(Y|M, X)$ , where the substitution of histories by memory state labels is done in accordance with the encoding function [i.e.,  $f(\bar{z}) = \sigma_m$  implies the substitution  $\bar{z} \rightarrow m$ ].

### B. Memory costs

Different choices of  $f$  lead to different memory states, and consequently, agents with different memory requirements. Here we are concerned with memory-minimal agents—those that are able to extract and store the minimal amount of historical information possible whilst still being able to execute a given strategy for any future stimuli. Correspondingly, we take the amount of information stored in the agent’s memory system  $M$  as our metric of performance:

$$C_{f,\mathcal{R}} := S_{\text{vN}}[\sigma_M], \quad (1)$$

where  $S_{\text{vN}}$  is the von Neumann entropy [26] (reducing to the Shannon entropy for classical memory states) of the memory state distribution, here assumed to be their steady-state distribution [11,21]. The second subscript  $\mathcal{R}$

recognizes that this distribution typically depends on how the stimuli the agent receives are selected—be they drawn from a stochastic process or, more generally, by another agent responding to the actions of the agent. The procedure for how the input stimuli are selected is referred to as the input strategy  $\mathcal{R}$ , as is formally defined in Appendix A. It is often useful to also consider a “worst-case” information cost—the necessary amount of memory an agent must have available to be able to respond appropriately to any input strategy.

### C. Memory-minimal classical agents

Using tools from complexity science [27–29], the provably memory-minimal classical adaptive agents can systematically be determined [21]. Consider that if the strategy dictates that two histories  $\vec{z}$  and  $\vec{z}'$  should have statistically identical action responses for all possible future stimuli sequences, there should be no need to distinguish between them in the memory. Similarly, consider that if they do have different action responses, then they have to be represented by different memory states. This rationale, while seemingly simple, directly motivates an encoding function that can be shown to be memory minimal in the design of classical adaptive agents.

This encoding function  $f_\varepsilon$  is thus defined by

$$f_\varepsilon(\vec{z}) = f_\varepsilon(\vec{z}') \Leftrightarrow P(\vec{Y}|\vec{z}, \vec{x}) = P(\vec{Y}|\vec{z}', \vec{x}) \quad \forall \vec{x}. \quad (2)$$

The corresponding memory states  $\{\sigma_s\}$ , labeled by  $s \in \mathcal{S}$ —referred to as the causal states of the strategy [21,27]—are a partitioning of histories into equivalence classes based on their responses to future stimuli. The respective information cost Eq. (1) of this encoding function is given by  $C_{\mu, \mathcal{R}} = -\sum_{s \in \mathcal{S}} P(s) \log_2[P(s)]$ , where  $P(s) = \sum_{\vec{z} \in \mathcal{S}} P(\vec{z})$  for the given input strategy  $\mathcal{R}$ .

The agent as a whole is called the  $\varepsilon$ -transducer of the strategy [21] and, crucially, is classically memory minimal for any nonpathological input strategy. In recognition of this, its memory requirements are seen as fundamental properties of the strategy; in particular, the worst-case information cost is designated as the structural complexity of the strategy [21]. These ideas have seen application in contexts such as agent-based learning [30,31] and energy harvesting [32], and understanding quantum contextuality [33].

## III. QUANTUM ADAPTIVE AGENTS

A quantum adaptive agent is able to store and process quantum information in its memory system  $\mathbb{M}$ , such that the encoding function  $f$  maps histories into quantum states  $\{\rho_m\}$ , and the policy  $\Lambda$  is a quantum channel. As per Eq. (1), the information cost of a quantum encoding function  $q$  is given by  $C_{q, \mathcal{R}} = -\text{Tr}(\rho \log_2[\rho])$ , where  $\rho = \sum_m P(m) \rho_m$ . A specific design for a quantum agent

has already demonstrated the potential for a quantum memory advantage over memory-minimal classical agents [11]. Yet, there is great flexibility in how a quantum agent can be designed beyond these prior proof-of-principle constructions; we now proceed to explore how quantum agents can maximize their advantage.

A central result of this work (proven in Appendix B) is the following set of constraints that a quantum agent can satisfy without penalty to their ability to achieve peak memory compression advantage.

- (i) The agent receives input stimuli  $\{x\}$  encoded in the computational basis states  $\{|x\rangle\}$ .
- (ii) The input stimulus is not consumed by the evolution of the agent;  $\Lambda$  preserves the input tape.
- (iii) The agent delivers output actions  $\{y\}$  via projective measurements in the computational basis states  $\{|y\rangle\}$  of its output tape.
- (iv) The memory states are pure and in one-to-one correspondence with the strategy’s causal states  $\mathcal{S}$ .

That is, generalizing beyond these features cannot provide further memory advantage. With stimuli and actions encoded as classical states, all quantum dynamics occur within the agent’s internal dynamics—the quantum memory advantage is not contingent upon access to a quantum environment. Further, these constraints imply a specific form of memory-minimal quantum agents.

*Theorem 1.*—A provably memory-minimal quantum agent executing any strategy  $\mathcal{P}$ —for any input strategy  $\mathcal{R}$ —can always be realized using the circuit of Fig. 1(b). That is, the policy  $\Lambda$  is realized in two stages. The first stage is a unitary operator  $U$  acting on the joint system of (i) the agent’s memory  $\mathbb{M}$ , (ii) input tape containing stimuli  $x$  encoded as  $|x\rangle$ , (iii) output tape initialized in  $|0\rangle$ , and (iv) “junk” tape also initialized in  $|0\rangle$ . Then, the output action  $y$  is realized by a computational basis measurement of the output tape, and the junk tape is discarded. Moreover, the memory states  $\{|\sigma_s\rangle\}$  are all pure and in one-to-one correspondence with the causal states  $\mathcal{S}$  of the strategy; the encoding function satisfies Eq. (2).

The unitary evolution can be expressed as

$$U|\sigma_s\rangle|x\rangle|0\rangle|0\rangle = \sum_y \sqrt{P(y|x, s)} |\sigma_{\lambda(z, s)}\rangle|x\rangle|y\rangle|\psi(z, s)\rangle, \quad (3)$$

where  $|\psi(z, s)\rangle$  represents the final state of the junk tape before it is discarded.

This implies that the only effective degrees of freedom in designing an agent’s memory encoding lie in the choice of junk states  $\{|\psi(z, s)\rangle\}$ , as  $U$  and the memory states  $\{|\sigma_s\rangle\}$  are then defined implicitly through Eq. (3). However, not every choice of junk states is physically realizable, due to the constraint that  $U$  is unitary. Consider the overlap of two memory states  $s$  and  $s'$ , given by  $c_{ss'} := \langle \sigma_s | \sigma_{s'} \rangle$ . Using the condition  $U^\dagger U = \mathbb{I}$  and defining  $d_{ss'}^z := \langle \psi(z, s) | \psi(z, s') \rangle$ , from Eq. (3) we obtain

$$c_{ss'} = \sum_y \sqrt{P(y|x, s)P(y|x, s')} c_{\lambda(z, s)\lambda(z, s')} d_{ss'}^z. \quad (4)$$

Though expressed for a given stimulus  $x$ , consistency requires that this equation yield identical  $\{c_{ss'}\}$  for all  $x$ . While this constraint can always trivially be satisfied by setting  $|\psi(z, s)\rangle = |s\rangle$  for all  $z$ , this enforces that quantum memory states are mutually orthogonal, recovering the classical  $\varepsilon$ -transducer and removing all quantum memory advantage. The crux of the quantum advantage is thus in finding junk states that admit nonorthogonal memory states, and optimizing their assignment to maximize it.

It is tempting to look for simple junk states that are just complex scalars, removing the need for junk tape altogether (as the corresponding phase can be absorbed by the output tape). However, this is generally impossible.

*Theorem 2.*—Junk states  $\{|\psi(z, s)\rangle\}$  cannot always be assigned as complex scalars. There exist strategies that can only be executed by quantum agents with access to a multidimensional junk tape that is discarded into the environment at each time step.

To prove this theorem, let us first suppose that (contrary to the theorem) the junk states are simply a set of complex scalars  $\{\exp(i\varphi_{zs})\}$ ; i.e.,  $d_{ss'}^z = \exp[i(\varphi_{zs'} - \varphi_{zs})]$ . Substituting into Eq. (4), we obtain

$$c_{ss'} = \sum_y \sqrt{P(y|x, s)P(y|x, s')} e^{i(\varphi_{zs'} - \varphi_{zs})} c_{\lambda(z, s)\lambda(z, s')}. \quad (5)$$

The left-hand side of this equation represents the overlaps of the memory states, and so for consistency we must have that the right-hand side is equal for all possible stimuli  $x$ . To prove the theorem, we need only establish that there is at least one strategy for which no set of phases  $\{\varphi_{zs}\}$  exists that can satisfy this condition.

Consider the strategy illustrated in Fig. 2. For this strategy, Eq. (4) demands  $c_{AB} = 0$ , as there is no overlap in future statistics for stimulus 1. Meanwhile, we must then have that  $d_{AB}^{0,0} = 0$ —clearly this cannot be satisfied if

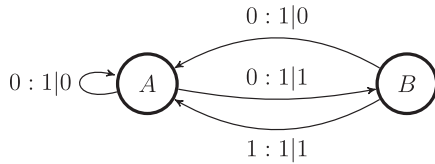


FIG. 2. Agents generally need to discard junk information. Hidden Markov model representation of the example process for Theorem 2. Nodes represent states and edges transitions; the notation  $y:p|x$  denotes that on input stimulus  $x$  the indicated transition occurs with probability  $p$  whilst outputting action  $y$ . The strategy requires two internal states  $\{A, B\}$ , two stimuli  $\{0, 1\}$ , and two actions  $\{0, 1\}$ . On stimulus 0, both states enact 0 with certainty and transition to state A; on stimulus 1, state A acts with 0 and transitions to B, while state B acts with 1 and transitions to A.

$|\psi(0, 0, A)\rangle$  and  $|\psi(0, 0, B)\rangle$  differ only by a phase factor. Thus, Theorem 2 is proven.

This is not an isolated example. In Appendix C we derive a sufficiency condition on the strategy that indicates Eq. (5) cannot be satisfied for any set of phases  $\{\varphi_{zs}\}$ , and hence nonscalar junk is required. Informally, this condition holds when the strategy has two states which must give rise to very similar behavior on one string of possible future stimuli, and very differently on another. The above example represents an extreme case of this. The requirement of nontrivial junk has operational significance, as it mandates that the agent discard information into the environment at each time step, corresponding to a source of thermal dissipation. The next theorem suggests this dissipation manifests from the data processing inequality.

*Theorem 3.*—The magnitude of the overlap between any pair of quantum memory states cannot exceed the overlap of their future statistics for any input strategy  $\mathcal{R}$ :

$$|c_{ss'}| \leq \min_{\vec{y}} \sum_{\vec{X}} \sqrt{P(\vec{y}|s, \vec{X})P(\vec{y}|s', \vec{X})}. \quad (6)$$

Physically, this can be understood as requiring that the future statistics do not provide a means of distinguishing between quantum memory states beyond what is information theoretically possible, imposing a constraint on their maximum fidelity [34]. In Appendix D we show how this bound can be calculated. However, this bound cannot always be saturated; a counterexample is provided in Appendix E.

#### IV. SYSTEMATIC QUANTUM AGENT DESIGN

We now provide a systematic method for assigning junk states such that the corresponding quantum agents achieve superior memory efficiency relative to memory-minimal classical [21] and prior quantum counterparts alike [11]. The design involves an effective representation of each of the memory states as a tensor-product form  $|\sigma_s\rangle = \otimes_x |\sigma_s^x\rangle$ , where the  $\{|\sigma_s^x\rangle\}$  behave as memory states specialized to each input (see Appendix F). These have associated overlaps  $c_{ss'}^x := \langle \sigma_s^x | \sigma_{s'}^x \rangle$ , such that  $c_{ss'} = \prod_x c_{ss'}^x$ . In this representation we identify the junk states as  $|\psi(z, s)\rangle = \otimes_{x' \neq x} |\sigma_s^{x'}\rangle$  and, correspondingly, their overlaps (for pairs with identical  $z$ ) as  $d_{ss'}^z = \prod_{x' \neq x} c_{ss'}^{x'}$ . In Appendix F we prove that for any strategy, a unitary of the form Eq. (3) can always be found that is based on these states. Given a strategy's  $\varepsilon$ -transducer, Algorithm 1 then provides a systematic means of designing quantum agents with this encoding.

In this encoding, any given pair of memory states has nonzero overlap if and only if there is no string of input stimuli for which they are certain to produce distinguishable strings of output actions; provided at least one such pair exists, the quantum agent exhibits a memory advantage

Algorithm 1. Systematic quantum agent encoding.

*Inputs:* Causal states  $\mathcal{S}$ , transition probabilities  $P(Y|X, S)$ , and update rule  $\lambda(z, s)$ .

*Outputs:* Quantum memory states  $\{|\sigma_s\rangle\}$ , evolution operator  $U$ .

1: Construct the set of multivariate polynomial equations

$$c_{ss'}^x = \sum_y \sqrt{P(y|x, s)P(y|x, s')} \prod_{x'} c_{\lambda(z, s)\lambda(z, s')}^{x'} \quad (7)$$

defined  $\forall s, s' \in \mathcal{S}, x \in \mathcal{X}$  and solve to obtain  $\{c_{ss'}^x\}$ .

- 2: Use a reverse Gram-Schmidt procedure [16,35] to construct quantum memory states  $\{|\sigma_s\rangle\}$  from overlaps  $c_{ss'}^x = \prod_x c_{ss'}^x$  and junk states  $\{|\psi(z, s)\rangle\}$  from overlaps  $d_{ss'}^x = \prod_{x' \neq x} c_{ss'}^{x'}$ .
- 3: Construct the columns of  $U$  explicitly defined in Eq. (3).
- 4: Fill the remaining columns of  $U$ , using a Gram-Schmidt procedure to ensure orthogonality with existing columns.

over provably minimal classical counterparts [11,12,26]. Note that despite their factorized representation presenting as an  $|\mathcal{S}||\mathcal{X}|$ -dimensional space, the reverse Gram-Schmidt procedure ensures that memory states can be supported by a memory system of at most  $|\mathcal{S}|$  dimensions.

## V. SCALING ADVANTAGE

The memory advantage of quantum agents can grow without bound. Consider a setting where an agent's optimal strategy depends on tracking some continuous parameter of its environment  $\tau$ . This can occur when naturally continuous parameters are involved, such as spatial position or time. Alternatively, for strategies with a dependence on events long ago in the past, the set of pasts  $\{\tilde{z}\}$  can be mapped to a continuous parameter over the interval  $[0, 1)$ , by taking  $\tilde{z}$  to specify a  $|\mathcal{Z}|$ -ary fraction. In either case, small differences in  $\tau$  often require only slightly different responses to future stimuli. However, if an agent must store  $\tau$  precisely, it requires an unbounded amount of memory.

To circumvent this, the conventional classical method is to adopt coarse graining, in which an approximation of the optimal strategy  $\mathcal{P}$  is executed based on storing  $\tau$  only to some finite precision. That is,  $\tau$  is divided up into a set of discrete bins, and all values of  $\tau$  within a given bin are mapped to the same memory state. An  $n$ -bit precision coarse graining divides  $\tau$  into  $2^n$  such bins, each of width  $\delta\tau^{(n)}$ ; the corresponding coarse graining of the strategy is denoted  $\mathcal{P}^{(n)}$ . For a classical agent, the memory cost then diverges linearly with  $n$  [36], forcing a trade-off between precision and memory cost.

On the other hand, quantum agents may be able to avoid such divergences. Consider a family of quantum agents that implement coarse grainings  $\mathcal{P}^{(n)}$  of a strategy  $\mathcal{P}$  at each level of precision  $n \in \mathbb{N}$ . Consider also the following pair of convergence conditions, defined formally in Appendix G.

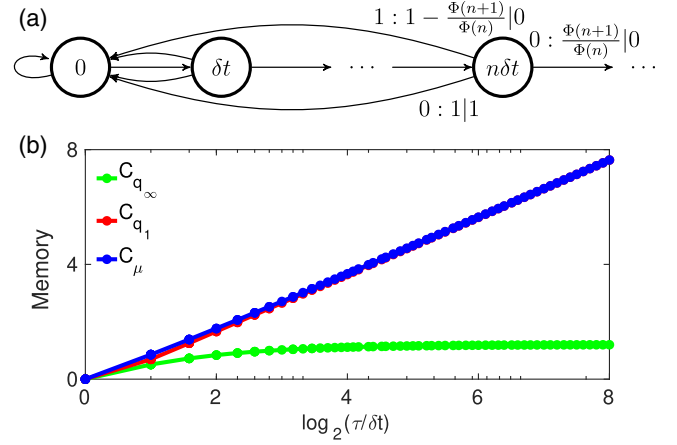


FIG. 3. Unbounded scaling advantage. (a) Hidden Markov model representation of resettable stochastic clocks. State labels represent the number of time steps since the last event;  $\Phi(n) := \int_{n\delta t}^{\infty} \phi(t)dt$  is the survival probability. (b) Associated memory costs for executing such a strategy, showing the advantage of our quantum agent ( $C_{q\infty}$ ) growing unbounded relative to previous quantum ( $C_{q1}$ ) and minimal classical agents ( $C_\mu$ ) with refinement of time steps.

- (i) Distributional convergence: The steady-state probability (densities) of the memory states converge exponentially with increasing precision.
- (ii) Memory-overlap convergence: The overlaps of each pair of memory states converge exponentially with increasing precision.

These convergence conditions encapsulate the intuition that if the strategy varies smoothly with a continuous parameter, then so too may the properties of the memory states of a quantum agent executing the strategy. When these conditions are met, a quantum agent can execute the strategy  $\mathcal{P}$  to arbitrary precision with bounded memory cost, giving rise to a scaling advantage over classical agents. The formal statement of this result is given in Theorem 4, which may be found in Appendix G together with its proof.

We illustrate an example of such scaling advantages occurring for agents tasked with executing certain strategies requiring coordinated stimuli-action responses over an increasingly greater number of time steps. We demonstrate this with an example family of resettable stochastic clocks. In this setting, the agent is tasked with behaving as a clock with stochastic tick events, that may be reset by an external stimulus. This stimulus can take two values:  $x = 0$  for “evolve normally” and  $x = 1$  for “reset,” while possible actions are  $y = 0$  for “no tick” and  $y = 1$  for “tick.” When  $x = 0$ , the agent behaves as a stochastic clock [37,38], modeled by a renewal process [39] where the agent emits a tick at stochastic intervals  $t$  governed by a distribution  $\phi(t)$ . Upon receiving  $x = 1$ , however, the agent must immediately reset its time counter, such that the clock behaves as though it has just ticked. The agent must replicate this behavior to some desired temporal resolution, such that

time is broken into finite time steps  $\delta t$ —as illustrated diagrammatically in Fig. 3(a), with further details in Appendix H. For a given  $\phi(t)$ , this prescribes a family of coarse-grained strategies parametrized by  $\delta t$ .

In Appendix H, we show that our quantum agents satisfy the convergence conditions for a large class of  $\phi(t)$  representing typical resettable stochastic clocks, and thus may execute them to arbitrary precision with a bounded cost. Meanwhile, the memory-minimal classical models must store an ever-increasing amount of information as  $\delta t$  is refined. That is, our quantum agents converge to a finite memory cost, while the classical agents diverge. Figure 3(b) highlights this by comparing the scaling of our quantum agents (labeled  $C_{q_\infty}$ ) with the memory-minimal classical ( $C_\mu$ ) and best prior quantum counterparts ( $C_{q_1}$ ) for the particular case where  $\phi(t)$  is uniformly distributed over the interval  $[0, \tau]$ , and resets are triggered at a constant rate  $1/2\tau$  [40].

## VI. DISCUSSION

We have introduced a general framework for adaptive agents that can capitalize on access to a quantum memory to reduce the information they must track about past stimuli and actions. Key to this, we isolated the features of an agent that are relevant to memory advantages and showed that they are in direct correspondence with the information it discards into its environment. Coupled with this, we provided a systematic algorithm for encoding the memory states of a quantum agent for any strategy, achieving a memory advantage relative to minimal classical and prior state-of-the-art quantum counterparts. Moreover, this advantage can grow without bound. These advantages can be utilized by agents both for executing fixed strategies and in running candidate strategies during their development [41–44], as well as by researchers modeling the behavior of agents. Our systematic quantum agent design may also be used for enhancing mechanical agents, for example, by endowing smart technologies with quantum processors. Our framework is agnostic to the specific engineering details of its implementation, and so can be realized with any quantum architecture that can receive (classical) input, and process and store quantum information according to the required policy evolution of Eq. (3). Proof-of-principle demonstrations are feasible with current setups, by, for example, adapting prior implementations of quantum models of passive stochastic processes in photonic setups [45,46] to undergo different evolutions at each time step conditional on the input.

Our results use entropic benchmarks for the memory, thus naturally assuming an ensemble setting. They describe quantum memory advantages with operational relevance for multiple agents implementing a strategy in parallel with shared memory [15]. This aligns well with scenarios where one wishes to sample over the conditional distributions for various strategies, for example, in Markov chain

Monte Carlo-type methods [47]. A compelling extension is to single-shot settings, where one may instead consider the max entropy—the dimension of the state space inhabited by the memory states. Single-shot advantages have been found for quantum models of passive stochastic processes [18–20,45,48,49] and for specific cases of input-output behavior modeling repeated measurement of a quantum system [50]. Since our general treatment ultimately relates to what can affect memory state overlaps, many of our results will continue to hold in single-shot settings—in particular, our form of the memory-minimal quantum agent—and thus can direct the search for systematic encodings based on other such benchmarks. Based on links established between quantum compression advantages and thermal efficiency in stochastic modeling [51,52], one may expect that our quantum agents are also able to execute their strategies with less thermal dissipation than classical counterparts.

A further enticing extension would be to the case where only near-faithful execution of the strategy is required—that is, some error is tolerated [49,53,54]. Our quantum agents bear a similarity to models of quantum walks with memory [55–57] and other instances of memory compression through quantum processing [58,59] such as quantum autoencoders [60–63]. Moreover, our general form for quantum adaptive agents Eq. (3) produces superpositions of all possible future trajectories for the input [16], potentially allowing for interference experiments that probe the overlap in the distributions of different strategies [46], or different input sequences. One can also consider superpositions of input sequences, akin to algorithms in quantum-enhanced reinforcement learning [23,24,64], where our agents may augment existing quantum speed-ups with extreme memory advantages.

## ACKNOWLEDGMENTS

This work was funded by the Imperial College Borland Fellowship in Mathematics, Grants No. FQXi-RFP-1809 and No. FQXi-RFP-IPW-1903 from the Foundational Questions Institute and Fetzer Franklin Fund (a donor advised fund of Silicon Valley Community Foundation), the Lee Kuan Yew Endowment Fund (Postdoctoral Fellowship), the Ministry of Education Singapore through Tier 1 Grant No. RG190/17, the National Research Foundation (NRF) Singapore under its NRFF Fellow program (Award No. NRF-NRFF2016-02), the NRF-ANR (Agence Nationale de la Recherche) joint program (NRF2017-NRF-ANR004 VanQuTe), and the Quantum Engineering Program QEP-SF3, and the Nanyang Technological University Start-up Grant. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of National Research Foundation or the Ministry of Education Singapore. T.J.E. thanks the Centre for Quantum Technologies for their hospitality.

**APPENDIX A: FRAMEWORK (EXTENDED)**

Here we provide further details of the framework used to describe adaptive agents, containing additional material relevant to the remaining Appendixes. We begin by formally defining an adaptive agent, as introduced in the main text.

*Definition 1.*—(Adaptive agents) An adaptive agent is defined by the tuple  $(\mathcal{X}, \mathcal{Y}, \{\sigma_m\}, f, \Lambda)$ , where

- (i)  $\mathcal{X}$  is the set of stimuli the agent can recognize,
- (ii)  $\mathcal{Y}$  is the set of actions the agent can perform,
- (iii)  $\{\sigma_m\}$  is the set of memory states the agent can store in its memory system  $M$ , labeled by an index  $m \in \mathcal{M}$ ,
- (iv)  $f: \vec{Z} \rightarrow \{\sigma_m\}$  is the encoding function that determines the memory state to which the agent assigns each history  $\vec{z}$ ,
- (v)  $\Lambda: \mathcal{X} \times \{\sigma_m\} \rightarrow \mathcal{Y} \times \{\sigma_m\}$  is the agent’s policy, describing how the agent selects action  $y$  in response to stimulus  $x$  given its current memory state, and how the memory state is updated.

An encoding  $f$  is said to be a valid encoding of a strategy  $\mathcal{P}$  if there exists a policy  $\Lambda$  by which the agent is able to execute actions in a manner statistically faithful to the strategy for every possible history and sequence of future stimuli. That is,  $f$  is valid if and only if  $\exists \Lambda: (P(\vec{Y}|m, \vec{x}) = P(\vec{Y}|\vec{z}, \vec{x})|f(\vec{z}) = \sigma_m) \forall \vec{z}, \vec{x}$ . An agent with such a policy and encoding function is then said to faithfully execute strategy  $\mathcal{P}$ . Hereon, we consider such faithful agents. The physics of the memory states determines the physics of the agent; that is, a classical agent can only store classical states in its memory and use classical dynamics for its policy, while for a quantum agent  $M$  can support quantum states, and  $\Lambda$  takes the form of a quantum channel.

A strategy  $\mathcal{P}$  can be described as a conditional distribution  $P(Y|\vec{Z}, X)$ . Mathematically, this corresponds to a stochastic input-output process [11,21,65,66], where the stimuli are the inputs and the actions the outputs, and the process maps stimuli and past actions to future actions. Consequently, our results encompass as a special case quantum models of passive stochastic processes—stochastic processes that evolve autonomously without environmental input—by taking the input alphabet to consist only of a single symbol (i.e., the strategy does not condition on any observed stimuli).

There are certain conditions implicitly placed on these input-output processes due to the limits of what an agent can predict about the future. That is, an agent cannot leverage information about future events that cannot be deduced from what they have already seen. The two conditions are referred to as the agent being nonanticipatory and causal [21]. The former requires that the strategy for choosing the current action must not depend on future input stimuli whenever these future stimuli are generated

independently of past actions; i.e.,  $P(Y|\vec{Z}, X) = P(Y|\vec{Z}, \vec{X})$  [11,21,67]. The latter requires that the memory of an agent can depend only on the past, and not the future—i.e., that  $f$  is a deterministic map from histories to memory states [18]. We also assume that the strategy is stationary (time invariant), such that the weightings  $P(Y|\vec{Z}, X)$  are independent of the time step  $t$ . Pasts and futures are taken to consist of semi-infinite strings of stimuli and actions. That is, at  $t = 0$  we take  $\vec{x} := \lim_{l \rightarrow \infty} x_{-l:0}$  and  $\vec{x} := \lim_{l \rightarrow \infty} x_{0:l}$ , where  $x_{k:l} := x_k, x_{k+1}, \dots, x_{l-1}$  denotes a contiguous string in the interval  $k \leq t < l$ .

In the main text we note that the input stimuli are in full generality drawn from an input strategy, where the stimuli manifest as actions of the agent’s environment, potentially conditioned on the previous actions of the agent.

*Definition 2.*—(Input strategies) An input strategy  $\mathcal{R}$  is an input-output stochastic process specified by a conditional distribution  $R(X_t|\vec{Z}_t)$  used to generate input stimuli of an adaptive agent. That is, it maps histories  $\{\vec{z}\}$  to the next stimulus received by the agent.

The subscripts indicate that the input strategy can have a temporal dependence (i.e., that it need not be stationary), while the conditioning on the entire history allows the stimuli to have a dependence on the actions of the agent. In the case where the stimuli are generated independent of the agent’s actions,  $\mathcal{R}$  reduces to a passive stochastic process. Note that in previous works the worst-case memory cost was considered only with respect to such input stochastic processes [11,21], rather than the more general input strategies described here.

**APPENDIX B: PROOF OF THEOREM 1**

We begin with the most general form a quantum adaptive agent can take, progressively examining each aspect to ascertain whether it is essential to its function, and whether it offers potential compression advantages—in order to constrain to the most general functional agent.

In full generality, at each time step, we have an evolution (i.e., a quantum channel) that acts on the current memory state  $\rho_m$  and the input stimulus  $x$ , encoded into a state  $\rho_x$ . These are mapped by the policy to an output action  $y$ , extractable from a state  $\rho_Y(x, m)$  with probability  $P(y|x, m)$ , and an updated memory state  $\rho_{m'}$  according to  $m' = \lambda(x, y, m)$ . For a complete accounting, we allow for the inclusion of a “blank” ancilla  $|0\rangle$  tape with the input and a “junk” state  $|\psi(x, y, m)\rangle$  with the output—both may without loss of generality be considered in their purified form [26].

*Lemma 1.*—There is no further quantum advantage from allowing memory states to be nonpure. Moreover, there is no further advantage for the memory states to be anything other than in one-to-one correspondence with the causal states of the  $\epsilon$ -transducer.

These results follow by generalizing the so-called causal state correspondence [34] and mixed state exclusion [18] found for quantum models of passive stochastic processes to the case of strategies. These establish that the memory states of the minimal quantum agents are in one-to-one correspondence with the causal states of the strategy, and can be instantiated as pure states. Our proofs of the generalizations largely follow those of the originals, with the modification to input-conditioned probability distributions.

*Proposition 1.*—(Causal state correspondence) For any strategy  $\mathcal{P}$  with causal encoding function  $f_\epsilon$ , there exists a memory-minimal causal, nonanticipatory quantum agent implementing the strategy with memory encoding function  $f$  that satisfies

$$f_\epsilon(\bar{z}) = f_\epsilon(\bar{z}') \Leftrightarrow f(\bar{z}) = f(\bar{z}') \quad (\text{B1})$$

for all past histories  $\bar{z}$  and  $\bar{z}'$ .

We first prove the reverse direction through its contrapositive. Suppose we had two histories  $\bar{z}$  and  $\bar{z}'$  belonging to different causal states, but mapped to the same memory state by  $f$ . The former condition implies  $P(\bar{Y}|\bar{z}, \bar{X}) \neq P(\bar{Y}|\bar{z}', \bar{X})$ , while the second implies  $f(\bar{z}) = f(\bar{z}')$ . Since the two memory states are identical, there is no quantum operation that could distinguish between them, and hence no operation that could produce different future statistics from them—and thus no quantum agent can generate the correct conditional future statistics for both histories. Therefore, we require  $f(\bar{z}) \neq f(\bar{z}')$  if  $f_\epsilon(\bar{z}) \neq f_\epsilon(\bar{z}')$ .

The forward direction follows from concavity of entropy [26]. Consider the set of histories  $\{\bar{z}\}$  belonging to causal state  $s$ . We define the contribution to the steady state of the memory coming from histories not in this set as  $\rho_{\bar{s}} = \sum_{\bar{z} \notin s} P(\bar{z})f(\bar{z})$ , and hence we can express  $\rho = \sum_{\bar{z} \in s} P'(\bar{z})[P(s)f(\bar{z}) + \rho_{\bar{s}}]$ , where  $P'(\bar{z}) = P(\bar{z})/P(s)$ . From the concavity of entropy, it follows that

$$\begin{aligned} S_{\text{vN}}[\rho] &\geq \sum_{\bar{z} \in s} P'(\bar{z}) S_{\text{vN}}[P(s)f(\bar{z}) + \rho_{\bar{s}}] \\ &\geq \min_{\bar{z} \in s} S_{\text{vN}}[P(s)f(\bar{z}) + \rho_{\bar{s}}]. \end{aligned} \quad (\text{B2})$$

Let  $\bar{z}^*$  be the particular history that minimizes this inequality. We thus have that for any valid quantum agent, an encoding which assigns all histories belonging to  $s$  to  $f(\bar{z}^*)$  will have lower or equal entropy. Moreover, the modified encoding is also a valid encoding: as the future statistics the agent must produce from  $f(\bar{z})$  for any other history  $\bar{z}' \in s$  are the same as those that must be produced from  $f(\bar{z}^*)$ , an encoding with  $f(\bar{z}) = f(\bar{z}^*) \forall \bar{z} \in s$  will produce the correct future statistics. This procedure can be repeated for histories belonging to all other  $s' \neq s$ , and we hence find that for any quantum agent there exists another quantum

agent implementing the same strategy with lower or equal entropy using an encoding function that assigns all histories in the same causal state to the same memory state.

*Proposition 2.*—(Mixed state exclusion) For any quantum agent implementing a strategy  $\mathcal{P}$  using a valid encoding with memory states  $\{\rho_m\}$ , there exists a valid encoding of lower or equal entropy with pure memory states  $\{\sigma_m\}$ .

We start by invoking the causal state correspondence, such that our goal is to show that for any valid quantum encoding for a strategy  $\mathcal{P}$  with memory states  $\{\rho_s\}$ , there exists a valid quantum encoding of lower or equal entropy with pure memory states  $\{\sigma_s\}$ . Suppose a particular memory state  $\rho_s$  is nonpure, such that we can decompose it as  $\rho_s = \sum_j p_j |a_j\rangle\langle a_j|$  for some set of pure states  $\{|a_j\rangle\}$ . Recall that causality demands the memory contain no information about the future that cannot be determined from its past; given the past stimuli and actions, there must be no correlations between the memory states and the futures they produce. This means that each of the states  $\{|a_j\rangle\}$  in our decomposition of  $\rho_s$  must all individually give rise to the same statistical futures as  $\rho_s$ , and thus a valid quantum encoding can be formed by replacing  $\rho_s$  with any of the  $|a_j\rangle$ . We again collect all contributions to the steady state from terms belonging to causal states other than  $s$  as  $\rho_{\bar{s}}$ , such that  $\rho = \sum_j p_j [P(s)|a_j\rangle\langle a_j| + \rho_{\bar{s}}]$ . From concavity of entropy:

$$\begin{aligned} S_{\text{vN}}[\rho] &\geq \sum_j p_j S_{\text{vN}}[P(s)|a_j\rangle\langle a_j| + \rho_{\bar{s}}] \\ &\geq \min_j S_{\text{vN}}[P(s)|a_j\rangle\langle a_j| + \rho_{\bar{s}}]. \end{aligned} \quad (\text{B3})$$

Let  $|a_j\rangle$  be the particular state that minimizes the inequality and designate it as  $|\sigma_s\rangle$ . We can thus obtain a valid quantum encoding of lower or equal entropy after replacing  $\rho_s$  with  $|\sigma_s\rangle$ . We can repeat the procedure for the memory states corresponding to other causal states, thus obtaining a valid encoding of lower or equal entropy where all memory states are pure.

Note that the above lemma is not specific to the von Neumann entropy, and holds for *any* entropy satisfying concavity. Since Lemma 1 allows us to restrict our attention to pure memory states, and the Gram matrix representation [68] of an ensemble of pure quantum states allows us to express the entropy as a function of pairwise overlaps of the states, we can hereon consider features of the agent that can affect the overlap of memory states to be synonymous with those that can (potentially) reduce the memory cost.

*Lemma 2.*—There is no memory advantage to encoding the input stimulus  $x$  as anything other than the computational basis state  $|x\rangle$ . Moreover, the input state need not be consumed by the evolution.

Consider that for each input state  $\rho_x$  there is a computational basis state  $|x\rangle$  appended to it which remains



unchanged by the evolution. Then, since it can be factored out it can be seen that it does not influence the overlaps of the memory states, and hence does not affect the amount of information stored. However, we can perform operations conditioned on the appended state, which, since they are orthogonal, allows us to imprint the  $\rho_x$  directly onto part of the blank ancilla space and proceed as before. Specifically, we can realize this as a unitary operation  $U_{\mathcal{X}}|x\rangle|0\rangle|0\rangle$ , where the third subspace is discarded into the junk and  $\rho_x$  is the resulting state of the second subspace after tracing out the other two. We see that it is sufficient to consider orthogonal input states  $\{|x\rangle\}$ , which can be used to mimic the effect of any set of input states—in effect, accounting for the preprocessing used to create  $\rho_x$  from the input stimulus as part of the evolution. As the appended input space is not affected by the evolution, it can be later used to retrieve the input stimulus.

*Lemma 3.*—There is no memory advantage for the extraction of  $y$  from  $\rho_Y(x, m)$  to be anything other than a projective measurement in the computational basis.

The output action must be extracted from  $\rho_Y(x, m)$  through measurement. Neumark’s dilation theorem allows us to express any quantum measurement as a projective measurement on a purified state in a larger space [69–71]—we can consider any model of the extraction that does not strictly use projective measurements to effectively be relegating this extended space into the junk. This dilation does not change the evolution of the memory state, and hence there is no penalty to working with the projective measurement picture. As Lemma 2 allows us to take the input states  $|x\rangle$  to be orthogonal, we can consider the output subspace to always be conditionally rotated at the end of the evolution such that the appropriate measurement basis is the computational basis, independent of the input stimulus.

With these lemmas, we can express the evolution at each time step by a global unitary operator  $U$  [Eq. (3)]. The amplitudes follow from the requirement that outcome  $y$  must be obtained with probability  $P(y|x, s)$  [16,20], and without loss of generality can be taken to be real by offloading any phase factor into the junk subspace.

### APPENDIX C: SUFFICIENCY CONDITION FOR NECESSITY OF JUNK

Here we provide a sufficient (but not necessary) condition on a strategy upon which no physically realizable quantum agent can implement said strategy without use of discarded junk states.

For a given strategy, consider a pair of states  $s$  and  $s'$  and strings of stimuli  $x_{0:L}$  and actions  $y_{0:L}$  for which  $\lambda(z_{0:L}, s) = \lambda(z_{0:L}, s')$ , where the output of the update function on a string of stimuli and actions is understood to be the sequential application of the update for each time step [i.e.,  $\lambda(z_0 z_1, s) := \lambda(z_1, \lambda(z_0, s))$ ]. Let us for shorthand denote  $p := P(y_{0:L}|x_{0:L}, s)$  and  $p' := P(y_{0:L}|x_{0:L}, s')$ .

Iterating through Eq. (5), we obtain that this provides a contribution of magnitude  $\sqrt{pp'}$  to the overlap of the two states if there is no junk. The magnitude of the remaining terms (corresponding to other action strings) must then be bounded by  $\sqrt{(1-p)(1-p')}$ . If  $p + p' = 1 + \alpha$  for some non-negative  $\alpha$ , we then have that

$$|c_{ss'}| \geq \sqrt{pp'} - \sqrt{(1-p)(1-p')} > \frac{\alpha}{2}. \quad (\text{C1})$$

This can be verified by direct substitution into the condition  $\sqrt{pp'} - \sqrt{(1-p)(1-p')} > \alpha/2$  after rearrangement and squaring, and using that  $p + p' > 1$  implies  $(1-p)(1-p') < p(1-p) < 1/4$ .

Now consider another string of stimuli  $x'_{0:L}$ . Iterating Eq. (5), we have that  $|c_{ss'}| \leq \sum_{y_{0:L}} \sqrt{P(y_{0:L}|x'_{0:L}, s)P(y_{0:L}|x'_{0:L}, s')}$ , placing an upper bound on the permissible overlap of the states. If this is less than  $\alpha/2$ , then a junk-free encoding cannot reduce the overlap as mandated by the  $x_{0:L}$  string of stimuli sufficiently low enough to reach this bound. In such circumstances it is then necessary to have the dissipation of junk states that allow further reductions in the overlaps.

To state this sufficiency condition more directly: If a strategy has a pair of states  $s$  and  $s'$  for which there exists a pair of stimuli-action strings  $x_{0:L}$  and  $y_{0:L}$  satisfying  $\lambda(z_{0:L}, s) = \lambda(z_{0:L}, s')$  and  $P(y_{0:L}|x_{0:L}, s) + P(y_{0:L}|x_{0:L}, s') = 1 + \alpha$  for some non-negative  $\alpha$ , and another stimuli string  $x'_{0:L}$  satisfying  $\sum_{y_{0:L}} \sqrt{P(y_{0:L}|x'_{0:L}, s)P(y_{0:L}|x'_{0:L}, s')} \leq \alpha/2$ , then the strategy cannot be implemented by an agent that does not utilize junk states.

We note that this condition need only be met for a single pair of stimuli strings on a single pair of states in order for the agent to require junk.

### APPENDIX D: BOUNDING QUANTUM MEMORY STATE OVERLAPS

Equation (6) in the main text places an upper bound on the overlap between any pair of quantum memory states, based on the distinguishability of their future statistics. Here, we provide two methods by which this bound can be calculated: the first method is approximate, with a computational cost that grows quadratically with the number of causal states and linearly with the depth of the approximation; the second is exact, but bears an exponential scaling in cost.

Suppose we are told that the memory has been initialized in one of two memory states  $\{|\sigma_s\rangle, |\sigma_{s'}\rangle\}$ , and we are asked to determine which one with a fixed number of input stimuli  $L$ . Obviously, if  $L = 0$ , we are unable to distinguish between the possible states. With  $L = 1$ , we wish to choose the stimulus  $x$  that minimizes the fidelity of the next output action; i.e.,  $\operatorname{argmin}_x \sum_y \sqrt{P(y|s, x)P(y|s', x)}$ . For  $L = 2$ ,

we are able to choose the second stimulus based on the action output in response to the first, and the first stimulus should be chosen bearing this in mind. Denoting  $F_{ss'}^{(1)} := \min_x \sum_y \sqrt{P(y|s, x)P(y|s', x)}$ , we see that the best strategy for choosing the first stimulus  $x$  is  $\operatorname{argmin}_x \sum_y \sqrt{P(y|s, x)P(y|s', x)} F_{\lambda(z, s)\lambda(z, s')}^{(1)}$ . An iterative strategy can be developed, leading to our first method: define  $F_{ss'}^{(0)} = 1$  and  $F_{ss'}^{(L+1)} = \min_x \sum_y \sqrt{P(y|s, x)P(y|s', x)} F_{\lambda(z, s)\lambda(z, s')}^{(L)}$  for all pairs of causal states  $\{s, s'\}$ ; iterate through to the desired precision. Accounting for the symmetries  $F_{ss'}^{(L)} = F_{s's}^{(L)}$  and  $F_{ss}^{(L)} = 1$ , at each step there are  $|\mathcal{S}|(|\mathcal{S}| - 1)/2$  functions to minimize over  $|\mathcal{X}|$  arguments each, leading to a scaling cost of  $L|\mathcal{X}||\mathcal{S}|(|\mathcal{S}| - 1)/2$ . It is intuitively clear that allowing for longer input strings cannot decrease the ability to distinguish between the states, and indeed,  $F_{ss'}^{(L+1)} \leq F_{ss'}^{(L)}$ . The method thus overestimates the upper bound, and as  $L \rightarrow \infty$ , the estimate converges on the bound.

The second method makes use of the fact that for each pair of memory states there is an optimal choice of next input stimulus, conditional on the number of subsequent input stimuli we are able to make. Observing that in the above iterative procedure we should have  $F_{ss'}^{(\infty)} = \min_x \sum_y \sqrt{P(y|s, x)P(y|s', x)} F_{\lambda(z, s)\lambda(z, s')}^{(\infty)}$ , we can postulate the optimal stimulus for each pair and solve the associated linear equations. Minimizing this over all possible postulates for the optimal stimuli, we obtain the actual bound. However, there are  $|\mathcal{X}|^{|\mathcal{S}|(|\mathcal{S}|-1)/2}$  possible assignments of stimuli, and hence the computational cost of this method scales exponentially with the number of causal states.

We can also consider a hybrid of the two methods, to obtain an improved estimate over the first: begin by carrying out the first method to some desired depth  $L$ , then using the corresponding arguments that minimize the expressions as the postulate, evaluate the recursion relations from the second method.

### APPENDIX E: COUNTEREXAMPLE TO FIDELITY BOUND TIGHTNESS

As noted in the main text, counterexamples to the tightness of the fidelity upper bound on memory state overlap Eq. (6) exist. Here we provide such a counterexample.

Consider an agent with three memory states  $\{s_a, s_b, s_c\}$ , three actions  $\{a, b, c\}$ , and two stimuli  $\{0, 1\}$ . The dynamic is Markovian, such that after action  $y$  the memory

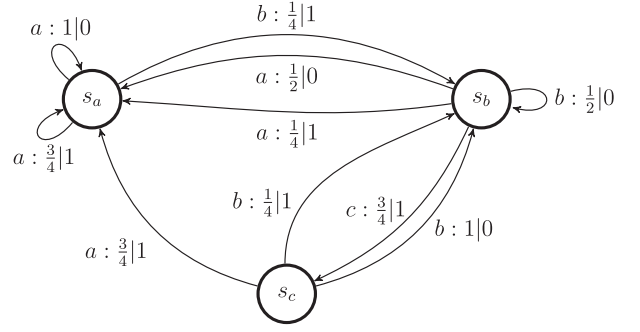


FIG. 4. Nontightness of fidelity bound. Hidden Markov model representation of a counterexample strategy to the tightness of the fidelity bound, as described in Appendix E.

transitions to state  $s_y$ . Let the corresponding strategy be defined by the following probabilities (illustrated in Fig. 4):

$$\begin{aligned}
 P(a|0, s_a) &= 1, & P(a|1, s_a) &= \frac{3}{4}, \\
 P(a|0, s_b) &= \frac{1}{2}, & P(b|1, s_a) &= \frac{1}{4}, \\
 P(b|0, s_b) &= \frac{1}{2}, & P(a|1, s_b) &= \frac{1}{4}, \\
 P(b|0, s_c) &= 1, & P(c|1, s_b) &= \frac{3}{4}, \\
 & & P(a|1, s_c) &= \frac{3}{4}, \\
 & & P(b|1, s_c) &= \frac{1}{4}, \quad (\text{E1})
 \end{aligned}$$

with the remaining unspecified probabilities all zero. Each of the states possess nonequal output responses to the stimuli, and so form the causal states of the strategy.

From stimulus 0 we obtain the following upper bounds on memory state overlaps,

$$|c_{ab}| \leq \frac{1}{\sqrt{2}}, \quad |c_{bc}| \leq \frac{1}{\sqrt{2}}, \quad |c_{ac}| = 0, \quad (\text{E2})$$

while stimulus 1 yields the bounds

$$|c_{ab}| \leq \frac{\sqrt{3}}{4}, \quad |c_{bc}| \leq \frac{\sqrt{3}}{4}, \quad |c_{ac}| \leq 1. \quad (\text{E3})$$

If the fidelity bound is to be saturated, we must have

$$|c_{ab}| = \frac{\sqrt{3}}{4}, \quad |c_{bc}| = \frac{\sqrt{3}}{4}, \quad |c_{ac}| = 0. \quad (\text{E4})$$

For stimulus 1 the evolution must be of the form

$$\begin{aligned}
 U|\sigma_a\rangle|1\rangle|0\rangle|0\rangle &= \frac{\sqrt{3}}{2}|\sigma_a\rangle|1\rangle|a\rangle|\psi(1, a, s_a)\rangle \\
 &\quad + \frac{1}{2}|\sigma_b\rangle|1\rangle|b\rangle|\psi(1, b, s_a)\rangle, \\
 U|\sigma_b\rangle|1\rangle|0\rangle|0\rangle &= \frac{1}{2}|\sigma_a\rangle|a\rangle|1\rangle|\psi(1, a, s_b)\rangle \\
 &\quad + \frac{\sqrt{3}}{2}|\sigma_c\rangle|1\rangle|c\rangle|\psi(1, c, s_b)\rangle, \\
 U|\sigma_c\rangle|1\rangle|0\rangle|0\rangle &= \frac{\sqrt{3}}{2}|\sigma_a\rangle|1\rangle|a\rangle|\psi(1, a, s_c)\rangle \\
 &\quad + \frac{1}{2}|\sigma_b\rangle|1\rangle|b\rangle|\psi(1, b, s_c)\rangle. \quad (\text{E5})
 \end{aligned}$$

To attain the prescribed values of  $|c_{ab}|$  and  $|c_{bc}|$  we must have  $|\psi(1, a, s_a)\rangle = \exp(i\varphi_1)|\psi(1, a, s_b)\rangle = \exp(i\varphi_2)|\psi(1, a, s_c)\rangle$ —i.e., equal up to phase factors. However, the condition on  $|c_{ac}|$  would then require

$$\frac{3}{4}e^{-i\varphi_2} + \frac{1}{4}\langle\psi(1, b, s_c)|\psi(1, b, s_a)\rangle = 0, \quad (\text{E6})$$

which clearly cannot be satisfied. Thus, the fidelity bound cannot be tightly satisfied.

Interestingly, this manifests only for nontrivial strategies; for passive stochastic processes it is always possible to construct a quantum model of the process with trivial (i.e., one-dimensional) junk states that saturates the fidelity bound [16,20].

## APPENDIX F: DETAILS FOR SYSTEMATIC QUANTUM AGENT DESIGN

Recall that our systematic encoding is based on a representation using factorized quantum memory states  $\{|\sigma_s\rangle\} = \{\otimes_x |\sigma_s^x\rangle\}$ , with overlaps  $c_{ss'}^x := \langle\sigma_s^x|\sigma_{s'}^x\rangle$  and  $c_{ss'} = \prod_x c_{ss'}^x$ . Consider input stimuli-specific unitaries  $\{U_x\}$  that act in the following manner [16,20] on the corresponding input-specialized memory substates:

$$U_x|\sigma_s^x\rangle|0\rangle|0\rangle = \sum_y \sqrt{P(y|x, s)}|\sigma_{\lambda(z, s)}\rangle|y\rangle, \quad (\text{F1})$$

where we have combined the first and second subspaces on the left-hand side together on the right; this implicitly defines the memory substates. We also define a selection operation  $U_{\text{select}}$  that ensures that the correct memory state is acted on with the correct  $U_x$ , conditioned on the input state. Specifically, we define this operation to permute the memory substates conditioned on stimulus  $x$  such that the  $x$ th memory substate is in the first position and exchange the remaining memory substates with the junk. We then act with  $U_x$  conditioned on the input state. Defining

$U = (\sum_x U_x \otimes |x\rangle\langle x| \otimes \mathbb{I})U_{\text{select}}$ , we obtain the total evolution consistent with Eq. (3).

In this representation, the junk states are given by  $|\psi(z, s)\rangle = \otimes_{x' \neq x} |\sigma_{s'}^{x'}\rangle$ , i.e., the unused memory substates corresponding to other input stimuli. Using that  $U^\dagger U = \mathbb{I}$ , we obtain

$$c_{ss'} = \sum_y \sqrt{P(y|x, s)P(y|x, s')}c_{\lambda(z, s)\lambda(z, s')} \prod_{x' \neq x} c_{ss'}^{x'}. \quad (\text{F2})$$

This can then be reduced to be purely in terms of the substate overlaps, recovering Eq. (7):

$$c_{ss'}^x = \sum_y \sqrt{P(y|x, s)P(y|x, s')} \prod_{x'} c_{\lambda(z, s)\lambda(z, s')}^{x'}. \quad (\text{F3})$$

As described in Algorithm 1, the overlaps can then be found by solving this set of multivariate polynomial equations. A solution always exists for any process that asymptotically synchronizes [i.e.,  $\lim_{L \rightarrow \infty} H(S_0|Z_{0:L}) = 0$ ]: since a sufficiently long string of past stimuli-action pairs allows the causal state to be determined with certainty, by iterating through the recursion relations we obtain the solution

$$\begin{aligned}
 c_{ss'} &= \lim_{L \rightarrow \infty} \prod_{x_0} \sum_{y_0} \sqrt{P(y_0|x_0, s)P(y_0|x_0, s')} \\
 &\quad \times \prod_{x_1} \sum_{y_1} \sqrt{P(y_1|x_1, \lambda(z_0, s))P(y_1|x_1, \lambda(z_0, s'))} \\
 &\quad \times \dots \\
 &\quad \times \prod_{x_L} \sum_{y_L} \sqrt{P(y_L|x_L, \lambda(z_{0:L}, s))P(y_L|x_L, \lambda(z_{0:L}, s'))}. \quad (\text{F4})
 \end{aligned}$$

The final step is to use forward and reverse Gram-Schmidt procedures [16,35] to construct the memory states, junk states, and evolution operator. Notably, while the factorized memory state representation is specified in terms of an  $|\mathcal{S}||\mathcal{X}|$ -dimensional space, because there are only  $|\mathcal{S}|$  memory states the reverse Gram-Schmidt procedure ensures that the constructed memory states inhabit only an  $|\mathcal{S}|$ -dimensional space. Similarly, because overlaps of junk states corresponding to different  $z$  are irrelevant to the construction, the seemingly  $|\mathcal{S}|(|\mathcal{X}| - 1)$ -dimensional junk states are actually encodable into an  $|\mathcal{S}|$ -dimensional space. The evolution operator  $U$  then acts on this  $|\mathcal{S}|^2|\mathcal{X}||\mathcal{Y}|$ -dimensional joint memory-input-output-junk space.

## APPENDIX G: PROOF OF SCALING ADVANTAGE

To rigorously evaluate the memory costs of agents implementing coarse-grained strategies, we must first

introduce some formal definitions. We provide definitions implicitly in terms of a single continuous parameter; the corresponding definitions for the case of coarse graining multiple continuous parameters straightforwardly follow by nested application of the single parameter definitions. We assume the continuous parameter to be of finite domain, and without loss of generality we can take this domain to be  $[0, 1)$ . We also explicitly consider binary coarse grainings; the definitions and results readily generalize to arbitrary  $d$ -ary coarse grainings.

*Definition 3.*—(Binary coarse graining) An  $n$ -bit precision coarse graining of a continuous parameter  $\tau$  divides  $\tau$  into  $2^n$  bins of equal width  $\delta\tau^{(n)} = 2^{-n}$ . An  $n$ -bit precision coarse graining  $\mathcal{P}^{(n)}$  of a strategy  $\mathcal{P}$  with respect to a continuous parameter  $\tau$  groups together all values of  $\tau$  within each bin into a single memory state.

A continuous parameter over the domain  $[0, 1)$  can be (asymptotically) represented as a binary fraction; i.e.,  $\tau = \sum_{k=1}^{\infty} \tau_k 2^{-k}$ , where  $\tau_k \in \{0, 1\}$ . Correspondingly, an  $n$ -bit precision coarse graining of  $\tau$ , denoted by  $\tau^{(n)}$ , stores only the first  $n$  bits of this expansion; i.e.,  $\tau^{(n)} = \sum_{k=1}^n \tau_k 2^{-k}$ . This also provides a convenient representation for indexing the discretized bin, where the same truncated binary expansion prescribes a unique integer  $\underline{\tau}^{(n)} = \sum_{k=0}^{n-1} \tau_k 2^k$ . Analogous to how the index of a causal state denotes both the label of a memory state and an equivalence class of pasts, we use  $\underline{\tau}^{(n)}$  to denote both the label of the bin and the interval it spans, with the distinction clear in context. Thus, the notation  $\tau \in \underline{\tau}^{(n)}$  indicates  $\tau \in [\tau^{(n)}, \tau^{(n)} + \delta\tau^{(n)})$ . For  $n > n'$ , we also use the notation  $\tau^{(n)} \in \underline{\tau}^{(n')}$  to indicate the set of all possible  $n$ -bit precision coarse grainings of a  $\tau \in \underline{\tau}^{(n')}$ .

In this manner we can construct a family of coarse grainings of a strategy at each level of precision  $\{\mathcal{P}^{(n)}\}$ , where  $n \in \mathbb{N}$ . It is implicitly assumed that such a family should converge upon the behavior of the exact strategy in the infinite precision limit.

When we say that an agent executes a strategy with  $n$ -bit precision, we mean that it has a valid encoding of an  $n$ -bit precision coarse graining of the strategy. Like these coarse grainings, we can similarly define families of agents that implement families of coarse grainings. We denote the  $n$ -bit precision coarse-grained (quantum) memory states—corresponding to the states stored by the agent executing the  $n$ -bit precision coarse-grained strategy—as  $|\sigma_{\tau^{(n)}}^{(n)}\rangle$  for all  $\tau \in \underline{\tau}^{(n)}$ . Correspondingly, we denote the overlaps of these states as  $c_{\tau^{(n)}\tau'^{(n)}}^{(n)} := \langle \sigma_{\tau^{(n)}}^{(n)} | \sigma_{\tau'^{(n)}}^{(n)} \rangle$ .

For notational convenience in the following definitions and proof we will use the notation  $P^{(n')}(\tau^{(n)})$  for  $n > n'$ , which should be interpreted as  $P^{(n')}(\tau^{(n')})$  for all  $\tau^{(n)} \in \underline{\tau}^{(n')}$ . That is, when the argument to a coarse-grained probability is of higher precision than the distribution, then the argument should be further coarse

grained to match the precision of the probability. An analogous interpretation should be made for the coarse-grained memory states and their overlaps, i.e., for  $n > n'$ ,  $c_{\tau^{(n)}\tau'^{(n)}}^{(n)} := c_{\tau^{(n')}\tau'^{(n')}}^{(n')} \quad \forall \tau^{(n)} \in \underline{\tau}^{(n')}, \tau'^{(n)} \in \underline{\tau}'^{(n')}$ .

With this preamble, the memory state convergence conditions can now be formally stated.

*Definition 4.*—(Distributional convergence) A family of coarse-grained strategies  $\{\mathcal{P}^{(n)}\}$  is said to exhibit distributional convergence if for all possible input strategies  $\mathcal{R}$  there exists an  $n_0$  and constant  $K$  such that for all  $n > n_0$  the steady states satisfy  $|P^{(n)}(\tau^{(n)})/\delta\tau^{(n)} - P^{(n-1)}(\tau^{(n)})/\delta\tau^{(n-1)}| < K\delta\tau^{(n)} \quad \forall \tau^{(n)}$ .

A weaker version of this definition can be formulated, where the distributional convergence can be only with respect to a particular input strategy. If only this weaker form is satisfied, then Theorem 4 can be restated in an input strategy-dependent manner. We also note that distributional convergence implies that  $P^{(n)}(\tau^{(n)}) \sim \delta\tau^{(n)}$ .

*Definition 5.*—(Memory-overlap convergence) A family of coarse-grained encoding functions  $\{f^{(n)}\}$  mapping to sets of quantum memory states  $\{|\sigma_{\tau^{(n)}}^{(n)}\rangle\}$  is said to exhibit memory-overlap convergence if there exists an  $n_0$  and constant  $K$  such that for all  $n > n_0$ ,  $|c_{\tau^{(n)}\tau'^{(n)}}^{(n)} - c_{\tau^{(n-1)}\tau'^{(n-1)}}^{(n-1)}| < K\delta\tau^{(n)} \quad \forall \tau^{(n)}, \tau'^{(n)}$ .

Armed with these definitions, we are now in a position to formally state and prove the result given in the main text regarding bounded memory costs for quantum agents executing coarse-grained strategies.

*Theorem 4.*—Consider a strategy  $\mathcal{P}$  that has a valid encoding using memory states labeled by a finite number of continuous parameters of finite domain and a finite set of discrete parameters. A quantum adaptive agent can execute a coarse graining of the strategy to arbitrary precision with bounded memory cost if distributional and memory-overlap convergence are satisfied.

We first prove this for the case where the memory states are labeled by a single continuous parameter, after which we will extend to the general case.

*Lemma 4.*—Consider a strategy  $\mathcal{P}$  that has a valid encoding using memory states labeled by a single continuous parameter of finite domain. A quantum adaptive agent can execute a coarse graining of the strategy to arbitrary precision with bounded memory cost if distributional and memory-overlap convergence are satisfied.

Consider such a quantum encoding at  $n$ -bit precision, where  $n$  is sufficiently large that we are above the  $n_0$  required for the convergence conditions. The steady state of the quantum agent's memory is given by  $\rho^{(n)} = \sum_{\tau^{(n)}} P^{(n)}(\tau^{(n)}) |\sigma_{\tau^{(n)}}^{(n)}\rangle \langle \sigma_{\tau^{(n)}}^{(n)}|$ . Since  $\rho^{(n)}$  is finite dimensional, the associated memory cost is finite.

The Gram matrix [68] of  $\rho^{(n)}$  is given by  $G_{\underline{\tau}^{(n)}\underline{\tau}'^{(n)}}^{(n)} = \sqrt{P^{(n)}(\tau^{(n)})P^{(n)}(\tau'^{(n)})} c_{\tau^{(n)}\tau'^{(n)}}^{(n)}$ , and has the same spectrum (and hence von Neumann entropy) as  $\rho^{(n)}$ . We also define a dilated Gram matrix  $\bar{G}^{(n)}$ :

$$\bar{G}^{(n)} := \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \otimes G^{(n)}, \quad (\text{G1})$$

such that the elements are given by  $\bar{G}_{\underline{\tau}^{(n+1)}\underline{\tau}'^{(n+1)}}^{(n)} = (1/2)\sqrt{P^{(n)}(\underline{\tau}^{(n)})P^{(n)}(\underline{\tau}'^{(n)})}c_{\underline{\tau}^{(n)}\underline{\tau}'^{(n)}}^{(n)}$ . From the properties of the tensor product, it follows that (the nonzero elements of) the spectra of  $G^{(n)}$  and  $\bar{G}^{(n)}$  are identical, and thus they have the same von Neumann entropy.

Consider also the quantum encoding at precision  $n - 1$ . From the distributional and memory-overlap convergences, it follows that  $|G_{\underline{\tau}^{(n)}\underline{\tau}'^{(n)}}^{(n)} - \bar{G}_{\underline{\tau}^{(n-1)}\underline{\tau}'^{(n-1)}}^{(n-1)}| < K\delta\tau^{(n)2} \forall \underline{\tau}^{(n)}, \underline{\tau}'^{(n)}$  for some constant  $K$ . We define this matrix of differences  $\Delta^{(n)} := G^{(n)} - \bar{G}^{(n-1)}$ ; its elements scale as  $O(2^{-2n})$ .

The Schatten  $p$ -norms of a matrix  $A$  are defined  $\|A\|_p := \text{Tr}(|A|^p)^{1/p}$  for  $p \in [1, \infty)$  [72]. They satisfy Hölder's inequality, whereby  $\|AB\|_1 \leq \|A\|_p \|B\|_q$  for  $1/p + 1/q = 1$ . Two special cases of relevance here are  $p = 1$ , also referred to as the trace norm, and  $p = 2$ , which is equivalent to the Frobenius norm  $\|A\|_F := \sqrt{\sum_{jk} |A_{jk}|^2} = \|A\|_2$ . Noting that  $\Delta^{(n)}$  has  $2^n \times 2^n$  elements, we have that  $\|\Delta^{(n)}\|_2 \sim 2^{-n}$ . Then, by applying Hölder's inequality with  $p = 2$ ,  $A = \Delta^{(n)}$ , and  $B$  the identity matrix over the space occupied by  $\Delta^{(n)}$ , we have that  $\|\Delta^{(n)}\|_1 \sim 2^{-n/2}$ .

The Fannes-Audenaert inequality [72] relates the difference in von Neumann entropies of two operators with the trace norm of their difference. For two operators  $\rho_A$  and  $\rho_B$  of dimension  $d$ , it states that

$$|S_{\text{vN}}[\rho_A] - S_{\text{vN}}[\rho_B]| \leq \frac{1}{2} \log_2(d-1) \|\rho_A - \rho_B\|_1 + h\left(\frac{1}{2} \|\rho_A - \rho_B\|_1\right), \quad (\text{G2})$$

where  $h(x) := -x \log_2(x) - (1-x) \log_2(1-x)$ . Setting  $\rho_A = G^{(n)}$  and  $\rho_B = \bar{G}^{(n-1)}$ , together with the above we arrive at

$$|S_{\text{vN}}[\rho^{(n)}] - S_{\text{vN}}[\rho^{(n-1)}]| \leq \sim n 2^{-n/2}. \quad (\text{G3})$$

Thus, beyond a sufficiently high precision, the increase in the quantum agent's memory cost for each extra degree of precision is exponentially decreasing. Correspondingly, the memory cost will eventually converge when the precision is increased an arbitrary number of times, leading to a bounded memory cost at any level of precision.

When we have an additional set of discrete parameters  $m \in \mathcal{M}$  labeling the memory states, such that the pair  $(t, m) \in (\tau, \mathcal{M})$  uniquely specifies the memory state, we effectively have a finite number of sectors for the memory state space, with each sector corresponding to a different  $m$ .

The state convergence conditions readily generalize to this regime, by imposing the conditions on each sector individually. Then, by applying the above arguments in the proof of Lemma 4 to each sector, we see that the total contribution to the memory cost from the memory states in each sector is bounded. Since there are a finite number of sectors, the total memory cost is thus bounded.

When there are multiple continuous parameters, the conditions on convergence must apply to all such parameters. Beginning from a sufficiently fine discretization of all continuous parameters, we can apply the arguments above to each continuous parameter in turn, to deduce that the memory cost remains bounded at arbitrary precision in all continuous parameters. This completes the proof of Theorem 4.

Finally, we remark that while we have assumed finite, discrete stimulus and action alphabets in the above, the definitions and proofs readily extend to the case where these also are continuous parameters of finite domain.

## APPENDIX H: DETAILS FOR RESETTABLE STOCHASTIC CLOCKS

A renewal process [39] is described by a series of identical events, where the time interval between consecutive events is drawn randomly from a distribution  $\phi(t)$ ; here we focus on the case where this is discretized into time steps of size  $\delta t$ . A resettable renewal process can accept input stimuli that trigger a ‘‘reset’’ of the system to its postevent state, in effect triggering a phantom event and restarting the timer to the next event. We can describe the input stimulus by a two symbol alphabet: 0 (continue) and 1 (reset). Similarly, the output action alphabet can be described by two symbols: 0 for nonevents and 1 for events.  $\Phi(t) := \int_t^\infty \phi(t') dt'$  (and discrete analog thereof) represents the so-called survival probability of the process. Such resettable renewal processes correspond to the strategy of resettable stochastic clocks.

It is clear that since the agent will always behave the same on stimulus 1, the groupings of pasts into causal states depends only on their response to stimulus 0. This recovers the vanilla renewal process case, and we obtain the same causal states as in such settings [15,19,36,73]: outside of specific forms of  $\phi(t)$ —that we shall ignore here, noting that the following analysis can straightforwardly be generalized to encompass them—the causal states  $s_n$  of a renewal process describe the number of time steps  $n$  since the last event (in our case, this also includes the phantom events from resets).

The steady-state distribution of the causal states can be readily calculated for any resettable renewal process where the input stimuli are themselves driven by an input renewal process that resets upon events from either process. We label the event distribution and survival probability of the input process as  $\phi_I(t)$  and  $\Phi_I(t)$ , respectively, and similarly  $\phi_O(t)$  and  $\Phi_O(t)$  for the strategy renewal process. For a

pure renewal process without resettability, the steady-state distribution is given by  $\mu\Phi(t)$ , where the normalization  $\mu := (\int_0^\infty t\phi(t)dt)^{-1} = (\int_0^\infty \Phi(t)dt)^{-1}$  (replace integrals with sums for the discrete-time case) is called the mean firing rate, and represents the average number of events per unit time or time step [15,36,73]. Since both processes are reset by events on the strategy process, we can view the pure output action process without reference to the input as a renewal process in its own right, with an effective event distribution being a function of both stimulus and action event distributions. The effective survival probability is the product of the survival probabilities, as the pure output process will only survive up to a given time provided that neither the underlying renewal process or the input renewal process have fired. Thus, the steady-state probabilities will be proportional to  $\Phi_I(t)\Phi_O(t)$ , and normalized by their sum or integral, which yields the effective mean-firing rate. With these probabilities, the (input-dependent) minimal classical memory cost can be straightforwardly calculated.

To determine the corresponding memory measure for our quantum agent, we must also calculate the memory state overlaps. Using Eq. (7), and noting that all causal states behave identically on input 1, we obtain

$$c_{it'} = \sqrt{\frac{\Phi_O(t+\delta t)\Phi_O(t'+\delta t)}{\Phi_O(t)\Phi_O(t')}} c_{(t+\delta t)(t'+\delta t)} + \sqrt{\frac{\bar{\phi}_O(t)\bar{\phi}_O(t')}{\Phi_O(t)\Phi_O(t')}} \quad (\text{H1})$$

where  $\bar{\phi}_O(t) := \int_t^{t+\delta t} \phi_O(t)dt = \Phi_O(t) - \Phi_O(t+\delta t)$ . From these iterative equations we obtain

$$c_{it'} = \sum_{j=0}^{\infty} \sqrt{\frac{\bar{\phi}_O(t+j\delta t)\bar{\phi}_O(t'+j\delta t)}{\Phi_O(t)\Phi_O(t')}} \quad (\text{H2})$$

These overlaps saturate the fidelity bound Eq. (6). Together with the steady-state probabilities, we can calculate the input-dependent memory cost of our agent.

We also compare with the prior proof-of-principle quantum agent [11]. To determine the overlaps of its memory states  $\{|S_s\rangle := \otimes_x |S_s^x\rangle\}$  we recast this agent in terms of our general form Eq. (3):

$$U_{q_1} |S_s\rangle |x\rangle |0\rangle |0\rangle = \sum_y \sqrt{P(y|x, s)} |S_{\lambda(z, s)}\rangle |x\rangle |y\rangle \otimes_{x' \neq x} |S_s^{x'}\rangle |\lambda(z, s)\rangle \quad (\text{H3})$$

We note that this formulation in terms of unitary evolution differs from the original presentation, though it yields identical memory states. Expressed in this way it is clear where the deficiency of this agent relative to ours lies—in announcing the next causal state in the junk. The corresponding overlaps between the memory states for resettable stochastic clocks are given by

$$c_{it'} = \sqrt{\frac{\Phi_O(t+\delta t)\Phi_O(t'+\delta t)}{\Phi_O(t)\Phi_O(t')}} \delta_{(t+\delta t)(t'+\delta t)} + \sqrt{\frac{\bar{\phi}_O(t)\bar{\phi}_O(t')}{\Phi_O(t)\Phi_O(t')}} \quad (\text{H4})$$

With these overlaps, we can calculate the memory requirement of the agent. It can be seen that the overlaps rely not only on overlap in the output action statistics, but also that the immediately subsequent causal state into which the system transitions is the same. In contrast, our agent relies only on the overlap of output action statistics (over arbitrarily long horizons), which due to asymptotic synchronization to a causal state over sufficiently long pasts requires that the transition into the same causal state is mandated only for arbitrarily far into the future. It is for this reason that we label the quantum agents with subscripts 1 and  $\infty$ , and it becomes clear why our new agent drastically outperforms the prior agent for processes with long historical dependence.

Indeed, the example presented in the main text is not an isolated case of the scaling advantage for a particular resettable stochastic clock. Theorem 4 provides us with a sufficiency condition against which we can verify that typical resettable stochastic clocks with a smooth distribution  $\phi(t)$  require only a bounded amount of memory to execute when driven by a smooth renewal process.

*Corollary 1.*—Consider a resettable stochastic clock with distribution  $\phi_O(t)$  that is either of finite domain or takes the form of a Poisson process at long times. Suppose that it is driven by a renewal process with distribution  $\phi_I(t)$  that resets upon clock ticks. If  $\Phi_O(t)$  and  $\Phi_I(t)$  are infinitely differentiable, then a quantum agent encoded using Algorithm 1 can execute the strategy to arbitrary precision with only a bounded memory cost.

Let us begin by considering the case where  $\phi_O(t)$  is of finite domain. Since  $\phi = -d\Phi/dt$ , we also have that the distributions  $\phi_O(t)$  and  $\phi_I(t)$  are infinitely differentiable, and so can be approximated over small distances by a Taylor expansion [i.e.,  $\phi(t+\delta t) \approx \phi(t) + \partial_t \phi(t)\delta t$ ]. Then, the coarse graining  $P^{(n)}(1|0, t^{(n)}) = \bar{\phi}_O^{(n)}(t^{(n)}) = \int_{t^{(n)}} \phi_O(t)dt$  can readily be verified to exponentially converge toward the exact strategy with increasing precision. We also have that  $\Phi^{(n)}(t^{(n)}) = \Phi(t^{(n)}) \forall n \in \mathbb{N}$ , since

$$\begin{aligned} \Phi^{(n)}(t^{(n)}) &:= \sum_{j=0}^{\infty} \bar{\phi}^{(n)}(t^{(n)} + j\delta t^{(n)}) \\ &= \sum_{j=0}^{\infty} \int_{t^{(n)}+j\delta t^{(n)}} \phi(t') dt' \\ &= \int_{t^{(n)}}^{\infty} \phi(t') dt' = \Phi(t^{(n)}) \quad (\text{H5}) \end{aligned}$$

Consider now the overlaps  $c_{it'}^{(n)}$  as prescribed in Eq. (H2). We can make an expansion  $\bar{\phi}^{(n)}(t^{(n)} + \delta t^{(n)}) \approx \bar{\phi}^{(n)}(t^{(n)}) + \partial_t \bar{\phi}^{(n)}(t^{(n)})\delta t^{(n)}$ , where the correction is

$O(\delta t^{(n)2})$  due to  $\bar{\phi}^{(n)}(t^{(n)})$  scaling as  $O(\delta t^{(n)})$ . Thus, for sufficiently large  $n$ ,  $\bar{\phi}^{(n)}(t^{(n)}) \approx \bar{\phi}^{(n)}(t^{(n)} + \delta t^{(n)}) \approx \bar{\phi}^{(n-1)}(t^{(n)})/2 + O(\delta t^{(n)2})$ . Since there at most  $2^n$  nonzero terms in the sum constituting  $c_{t't'}^{(n)}$ , it then follows that at sufficiently large  $n$ ,  $|c_{t't'}^{(n)} - c_{t't'}^{(n-1)}| \leq \sim \delta t^{(n)} \forall t, t'$ ; memory-overlap convergence is satisfied.

The steady-state probabilities are given by  $P^{(n)}(t^{(n)}) = \mu_{IO}^{(n)} \Phi_I(t^{(n)}) \Phi_O(t^{(n)})$ , where  $\mu_{IO}^{(n-1)} := \sum_{j=0}^{\infty} \Phi_I(j\delta t^{(n)}) \Phi_O(j\delta t^{(n)})$ . For sufficiently large  $n$ , it follows that  $\mu_{IO}^{(n-1)} \approx 2\mu_{IO}^{(n-1)-1} - 1/2$ , and thus  $\mu_{IO}^{(n)} \sim \delta t^{(n)}$ . It can then readily be verified that distributional convergence is satisfied. Hence, the convergence conditions of Theorem 4 are satisfied, and it therefore follows that the memory cost remains bounded, irrespective of the precision.

More generally, when  $\phi_O(t)$  takes the form of a Poisson process at long times (say, for  $t > \tau_0$ ), we can partition the memory states in two, according to whether  $t$  is below or above  $\tau_0$ . This binary classification defines two sectors of memory states. For the former, we can apply the above arguments to show that in this sector the convergence conditions are satisfied. Meanwhile, we can apply known properties of the causal states of renewal processes [36] to see that all memory states in the latter sector belong to the same causal state, and hence Algorithm 1 maps them all to identical states. Thus, the convergence conditions are also satisfied in this sector, and hence Theorem 4 can again be applied.

---

[1] E. Bonabeau, *Agent-Based Modeling: Methods and Techniques for Simulating Human Systems*, *Proc. Natl. Acad. Sci. U.S.A.* **99**, 7280 (2002).

[2] M. W. Macy and R. Willer, *From Factors to Actors: Computational Sociology and Agent-Based Modeling*, *Annu. Rev. Sociol.* **28**, 143 (2002).

[3] C. M. Macal and M. J. North, *Tutorial on Agent-Based Modeling and Simulation*, in *Proceedings of the Winter Simulation Conference, 2005* (IEEE, New York, 2005), p. 14, 10.1109/WSC.2005.1574234.

[4] M. Wooldridge, *An Introduction to Multiagent Systems* (John Wiley & Sons, New York, 2009).

[5] L. An, *Modeling Human Decisions in Coupled Human and Natural Systems: Review of Agent-Based Models*, *Ecol. Model.* **229**, 25 (2012).

[6] J. D. Farmer and D. Foley, *The Economy Needs Agent-Based Modelling*, *Nature (London)* **460**, 685 (2009).

[7] S. Thurner, J. D. Farmer, and J. Geanakoplos, *Leverage Causes Fat Tails and Clustered Volatility*, *Quant. Finance* **12**, 695 (2012).

[8] L. A. Lardon, B. V. Merkey, S. Martins, A. Dötsch, C. Picioreanu, J.-U. Kreft, and B. F. Smets, *iDynoMiCS: Next-Generation Individual-Based Modelling of Biofilms*, *Environ. Microbiol.* **13**, 2416 (2011).

[9] S. Mei, P. M. A. Sloot, R. Quax, Y. Zhu, and W. Wang, *Complex Agent Networks Explaining the HIV Epidemic*

*among Homosexual Men in Amsterdam*, *Math. Comput. Simul.* **80**, 1018 (2010).

[10] A. R. Cassandra, L. P. Kaelbling, and M. L. Littman, *Acting Optimally in Partially Observable Stochastic Domains*, in *Proceedings of the Twelfth AAAI National Conference on Artificial Intelligence, Seattle WA, 1994* (AAAI Press, 1994), pp. 1023–1028.

[11] 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).

[12] M. Gu, K. Wiesner, E. Rieper, and V. Vedral, *Quantum Mechanics Can Reduce the Complexity of Classical Models*, *Nat. Commun.* **3**, 762 (2012).

[13] 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).

[14] C. Aghamohammadi, S. P. Loomis, J. R. Mahoney, and J. P. Crutchfield, *Extreme Quantum Memory Advantage for Rare-Event Sampling*, *Phys. Rev. X* **8**, 011025 (2018).

[15] 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).

[16] F. C. Binder, J. Thompson, and M. Gu, *Practical Unitary Simulator for non-Markovian Complex Processes*, *Phys. Rev. Lett.* **120**, 240502 (2018).

[17] T. J. Elliott, A. J. P. Garner, and M. Gu, *Memory-Efficient Tracking of Complex Temporal and Symbolic Dynamics with Quantum Simulators*, *New J. Phys.* **21**, 013021 (2019).

[18] J. Thompson, A. J. P. Garner, J. R. Mahoney, J. P. Crutchfield, V. Vedral, and M. Gu, *Causal Asymmetry in a Quantum World*, *Phys. Rev. X* **8**, 031013 (2018).

[19] T. J. Elliott, C. Yang, F. C. Binder, A. J. P. Garner, J. Thompson, and M. Gu, *Extreme Dimensionality Reduction with Quantum Modeling*, *Phys. Rev. Lett.* **125**, 260501 (2020).

[20] Q. Liu, T. J. Elliott, F. C. Binder, C. D. Franco, and M. Gu, *Optimal Stochastic Modeling with Unitary Quantum Dynamics*, *Phys. Rev. A* **99**, 062110 (2019).

[21] N. Barnett and J. P. Crutchfield, *Computational Mechanics of Input-Output Processes: Structured Transformations and the  $\epsilon$ -Transducer*, *J. Stat. Phys.* **161**, 404 (2015).

[22] D. Dong, C. Chen, H. Li, and T.-J. Tarn, *Quantum Reinforcement Learning*, *IEEE Trans. Syst. Man Cybern. B* **38**, 1207 (2008).

[23] G. D. Paparo, V. Dunjko, A. Makmal, M. A. Martin-Delgado, and H. J. Briegel, *Quantum Speedup for Active Learning Agents*, *Phys. Rev. X* **4**, 031002 (2014).

[24] V. Dunjko and H. J. Briegel, *Machine Learning & Artificial Intelligence in the Quantum Domain: A Review of Recent Progress*, *Rep. Prog. Phys.* **81**, 074001 (2018).

[25] Here, we assume that agents must be causal and non-anticipatory, such that their decisions and internal memory can only be based on past information, and not as-yet-unseen future information; see Appendix A for details.

[26] M. A. Nielsen and I. Chuang, *Quantum Computation and Quantum Information* (2000), 10.1017/CBO9780511976667.

- [27] J. P. Crutchfield and K. Young, *Inferring Statistical Complexity*, *Phys. Rev. Lett.* **63**, 105 (1989).
- [28] C. R. Shalizi and J. P. Crutchfield, *Computational Mechanics: Pattern and Prediction, Structure and Simplicity*, *J. Stat. Phys.* **104**, 817 (2001).
- [29] J. P. Crutchfield, *Between Order and Chaos*, *Nat. Phys.* **8**, 17 (2012).
- [30] S. E. Marzen and J. P. Crutchfield, *Optimized Bacteria Are Environmental Prediction Engines*, *Phys. Rev. E* **98**, 012408 (2018).
- [31] A. Zhang, Z. C. Lipton, L. Pineda, K. Azizzadenesheli, A. Anandkumar, L. Itti, J. Pineau, and T. Furlanello, *Learning Causal State Representations of Partially Observable Environments*, arXiv:1906.10437.
- [32] A. B. Boyd, D. Mandal, and J. P. Crutchfield, *Leveraging Environmental Correlations: The Thermodynamics of Requisite Variety*, *J. Stat. Phys.* **167**, 1555 (2017).
- [33] A. Cabello, M. Gu, O. Gühne, and Z.-P. Xu, *Optimal Classical Simulation of State-Independent Quantum Contextuality*, *Phys. Rev. Lett.* **120**, 130401 (2018).
- [34] W. Y. Suen, J. Thompson, A. J. P. Garner, V. Vedral, and M. Gu, *The Classical-Quantum Divergence of Complexity in Modelling Spin Chains*, *Quantum* **1**, 25 (2017).
- [35] P. Denny and A. Krzywicki, *Mathematics for Physicists* (Dover Publications, New York, 1996).
- [36] S. E. Marzen and J. P. Crutchfield, *Informational and Causal Architecture of Discrete-Time Renewal Processes*, *Entropy* **17**, 4891 (2015).
- [37] M. P. Woods, R. Silva, G. Pütz, S. Stupar, and R. Renner, *Quantum Clocks Are More Accurate Than Classical Ones*, arXiv:1806.00491.
- [38] Y. Yang and R. Renner, *Ultimate Limit on Time Signal Generation*, arXiv:2004.07857.
- [39] W. L. Smith, *Renewal Theory and Its Ramifications*, *J. R. Stat. Soc. Ser. B* **20**, 243 (1958).
- [40] Note that though we are considering the memory costs with respect to a particular input strategy, we drop the  $\mathcal{R}$  from the subscript for notational brevity.
- [41] L. P. Kaelbling, M. L. Littman, and A. W. Moore, *Reinforcement Learning: A Survey*, *J. Artif. Intell. Res.* **4**, 237 (1996).
- [42] L. Bu, R. Babu, and B. De Schutter, *A Comprehensive Survey of Multiagent Reinforcement Learning*, *IEEE Trans. Syst. Man Cybern. C* **38**, 156 (2008).
- [43] J. Kober, J. A. Bagnell, and J. Peters, *Reinforcement Learning in Robotics: A Survey*, *International Journal of Robotics Research* **32**, 1238 (2013).
- [44] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, MA, 2018).
- [45] F. Ghafari, N. Tischler, J. Thompson, M. Gu, L. K. Shalm, V. B. Verma, S. W. Nam, R. B. Patel, H. M. Wiseman, and G. J. Pryde, *Dimensional Quantum Memory Advantage in the Simulation of Stochastic Processes*, *Phys. Rev. X* **9**, 041013 (2019).
- [46] F. Ghafari, N. Tischler, C. Di Franco, J. Thompson, M. Gu, and G. J. Pryde, *Interfering Trajectories in Experimental Quantum-Enhanced Stochastic Simulation*, *Nat. Commun.* **10**, 1630 (2019).
- [47] W. R. Gilks, S. Richardson, and D. Spiegelhalter, *Markov Chain Monte Carlo in Practice* (Taylor & Francis, London, 1995).
- [48] S. P. Loomis and J. P. Crutchfield, *Strong and Weak Optimizations in Classical and Quantum Models of Stochastic Processes*, *J. Stat. Phys.* **176**, 1317 (2019).
- [49] T. J. Elliott, *Quantum Coarse Graining for Extreme Dimension Reduction in Modeling Stochastic Temporal Dynamics*, *PRX Quantum* **2**, 020342 (2021).
- [50] 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).
- [51] S. P. Loomis and J. P. Crutchfield, *Thermal Efficiency of Quantum Memory Compression*, *Phys. Rev. Lett.* **125**, 020601 (2020).
- [52] T. J. Elliott, *Memory Compression and Thermal Efficiency of Quantum Implementations of Nondeterministic Hidden Markov Models*, *Phys. Rev. A* **103**, 052615 (2021).
- [53] M. Ho, M. Gu, and T. J. Elliott, *Robust Inference of Memory Structure for Efficient Quantum Modeling of Stochastic Processes*, *Phys. Rev. A* **101**, 032327 (2020).
- [54] C. Yang, F. C. Binder, M. Gu, and T. J. Elliott, *Measures of Distinguishability between Stochastic Processes*, *Phys. Rev. E* **101**, 062137 (2020).
- [55] A. P. Flitney, D. Abbott, and N. F. Johnson, *Quantum Walks with History Dependence*, *J. Phys. A* **37**, 7581 (2004).
- [56] M. McGettrick, *One Dimensional Quantum Walks with Memory*, *Quantum Inf. Comput.* **10**, 0509 (2010).
- [57] P. P. Rohde, G. K. Brennen, and A. Gilchrist, *Quantum Walks with Memory Provided by Recycled Coins and a Memory of the Coin-Flip History*, *Phys. Rev. A* **87**, 052302 (2013).
- [58] L. A. Rozema, D. H. Mahler, A. Hayat, P. S. Turner, and A. M. Steinberg, *Quantum Data Compression of a Qubit Ensemble*, *Phys. Rev. Lett.* **113**, 160504 (2014).
- [59] Y. Yang, G. Chiribella, and M. Hayashi, *Quantum Stopwatch: How to Store Time in a Quantum Memory*, *Proc. R. Soc. A* **474**, 20170773 (2018).
- [60] J. Romero, J. P. Olson, and A. Aspuru-Guzik, *Quantum Autoencoders for Efficient Compression of Quantum Data*, *Quantum Sci. Technol.* **2**, 045001 (2017).
- [61] K. H. Wan, O. Dahlsten, H. Kristjánsson, R. Gardner, and M. S. Kim, *Quantum Generalisation of Feedforward Neural Networks*, *npj Quantum Inf.* **3**, 36 (2017).
- [62] A. Pepper, N. Tischler, and G. J. Pryde, *Experimental Realization of a Quantum Autoencoder: The Compression of Qutrits via Machine Learning*, *Phys. Rev. Lett.* **122**, 060501 (2019).
- [63] C.-J. Huang, H. Ma, Q. Yin, J.-F. Tang, D. Dong, C. Chen, G.-Y. Xiang, C.-F. Li, and G.-C. Guo, *Realization of a Quantum Autoencoder for Lossless Compression of Quantum Data*, *Phys. Rev. A* **102**, 032412 (2020).
- [64] V. Dunjko, J. M. Taylor, and H. J. Briegel, *Quantum-Enhanced Machine Learning*, *Phys. Rev. Lett.* **117**, 130501 (2016).
- [65] A. Khintchine, *Korrelationstheorie der Stationären Stochastischen Prozesse*, *Math. Ann.* **109**, 604 (1934).
- [66] J. P. Crutchfield, *The Calculi of Emergence: Computation, Dynamics and Induction*, *Physica (Amsterdam)* **75D**, 11 (1994).
- [67] F. Caruso, V. Giovannetti, C. Lupo, and S. Mancini, *Quantum Channels and Memory Effects*, *Rev. Mod. Phys.* **86**, 1203 (2014).



- [68] R. A. Horn and C. R. Johnson, *Matrix Analysis* (Cambridge University Press, Cambridge, England, 1990).
- [69] M. Neumark, *On Spectral Functions of a Symmetric Operator*, *Izv. Ross. Akad. Nauk Ser. Mat.* **7**, 285 (1943).
- [70] W. F. Stinespring, *Positive Functions on C\*-Algebras*, *Proc. Am. Math. Soc.* **6**, 211 (1955).
- [71] N. I. Akhiezer and I. M. Glazman, *Theory of Linear Operators in Hilbert Space* (Courier Corporation, New York, 2013).
- [72] J. Watrous, *The Theory of Quantum Information* (Cambridge University Press, Cambridge, England, 2018).
- [73] S. Marzen and J. P. Crutchfield, *Informational and Causal Architecture of Continuous-Time Renewal Processes*, *J. Stat. Phys.* **168**, 109 (2017).