Spatiotemporal Quantum Advantages for Rare-Event Sampling

Akshat Gupta,[1,2,*] Jianjun Chen,[3,4,†] Chengran Yang,[3,4,‡] Mile Gu,[3,4,§] Caterina Doglioni,[1] Jayne Thompson,[5,4] and Thomas J. Elliott[1,6,2,‖]

[1]*Department of Physics and Astronomy, University of Manchester, Manchester M13 9PL, United Kingdom*
[2]*Centre for Quantum Science and Engineering, University of Manchester, Manchester M13 9PL, United Kingdom*
[3]*Nanyang Quantum Hub, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore 637371, Singapore*
[4]*Centre for Quantum Technologies, Nanyang Technological University, Singapore 639798, Singapore*
[5]*College of Computing and Data Science, Nanyang Technological University, Singapore 639798, Singapore*
[6]*Department of Mathematics, University of Manchester, Manchester M13 9PL, United Kingdom*

(Received 16 February 2026; revised 11 August 2026; accepted 21 October 2026; published 19 November 2026)

In a world fraught with uncertainty, it behooves us to consider the most disastrous possible outcomes, in order that we can best mitigate against them. Often these are extreme risk events–of high consequence, but low probability. Yet, the rarity of such events limits our ability to study them; there is a paucity of past occurrences to draw upon, and sampling such rare events from models is typically costly in time. Here, we introduce an approach for rare-event sampling based on quantum stochastic models. This integrates the memory advantages of quantum stochastic modeling—and its associated efficient means of preparing quantum sample states—with the "speedups" of quantum amplitude amplification, to yield simultaneous reductions in spatial and temporal computational complexity compared to classical Monte Carlo–based approaches.

DOI: 10.1103/7zh3-zvnd

Global financial crises, major earthquakes, pandemics, and terror attacks are extreme risk events–very rare, yet potentially catastrophic [1–4]. Their scarcity begets uncertainty, and unfamiliar warning signs often escape notice. Yet, early signals often exist: for example, in 2005, a small number of traders spotted surging subprime mortgage delinquencies—ultimately leading to the 2008 crisis—three years in advance [5]. Identifying and monitoring such heralds of catastrophe is of paramount importance, to ensure we can handle these extreme occurrences. This is the remit of rare-event sampling [6–8].

Given a model of a process, a common approach to rare-event sampling is to use Markov chain Monte Carlo (MCMC) [9–12], where many trajectories of the process are generated, and postselected to keep only those trajectories where the rare event occurs for further study. However, if a target set of trajectories occurs with probability p, in order to reach a target number of samples from this set we must generate $O(1/p)$ trajectories. While more sophisticated approaches exist, such as a metadynamics biasing that reweights the model to favor rare event trajectories, these generically also involve a $O(1/p)$ scaling [13–16]. This leads to a prohibitively harsh scaling in computational cost for sampling truly rare events.

With quantum computers on the horizon, promising computational speedups for a range of tasks, it is natural to ask whether a quantum algorithm can yield a better scaling. One may consider the use of the quantum amplitude amplification (QAA) primitive [17] to boost the probability of rare events, which would seemingly allow for a quadratic reduction in temporal complexity. However, the direct naive and unstructured approach to constructing the quantum sample states that could be used as input into QAA would require an exponentially deep circuit with increasing number of time steps, and thus the complexity of state preparation would render any speed-up from QAA useless.

We overcome this limitation by incorporating the framework of quantum stochastic modelling [18–22] with QAA. Specifically, we leverage that quantum models can generate the requisite quantum sample states for QAA with a circuit depth that grows only linearly with the number of timesteps, removing this barrier to the application of QAA. Moreover, quantum models of stochastic processes exhibit spatial complexity advantage, in terms of a reduced memory dimension. Many examples have demonstrated

*Contact author: akshat.gupta-4@postgrad.manchester.ac.uk
†Contact author: chen1554@e.ntu.edu.sg
‡Contact author: yangchengran92@gmail.com
§Contact author: mgu@quantumcomplexity.org
‖Contact author: physics@tjelliott.net

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.

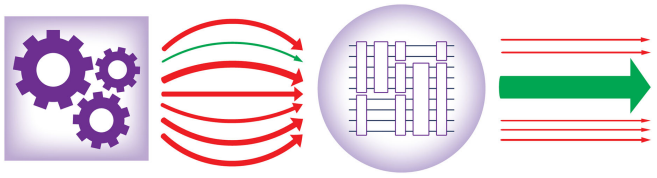


FIG. 1. Conceptual schema of our approach. A quantum model is used to generate a weighted superposition over all trajectories of a stochastic process, including target rare event(s). This is then fed into an algorithm that amplifies the weighting for the target rare event(s) and suppresses the rest.

this quantum advantage can grow without bound; finite-sized quantum models can generate statistics inaccessible with any finite classical model [23–25].

Correspondingly, we show that rare-event trajectories can be sampled with $O(1/\sqrt{p})$ time complexity—a quadratic reduction over the classical scaling—while inheriting the spatial complexity advantages of the quantum models. The conceptual schema is illustrated in Fig. 1. Our approach is highly related, though methodologically distinct, from Ref. [26], which explored the application of quantum stochastic models to rare-event sampling by examining quantum implementations of reweighted classical models, and showing that such "tilted" (i.e., biased) quantum models can exhibit (entropic) quantum memory advantages. By considering memory dimension our Letter introduces a stronger form of spatial complexity advantage, and through incorporating QAA, also yields a temporal complexity advantage.

We begin by introducing the relevant framework, on stochastic processes, quantum models, and QAA. We show how the quantum models can be deployed with QAA to achieve spatiotemporal advantages in rare-event sampling. We then demonstrate our approach with toy model applications. Finally, we conclude with a discussion of extensions, implications, and future work.

We consider rare-event sampling in the context of a time series manifest by an underlying stochastic process. Specifically, we consider discrete-time, discrete-event stochastic processes [27], where events $x \in \mathcal{X}$, drawn from a finite alphabet $\mathcal{X}$, occur at timesteps $t \in \mathbb{Z}$. We denote the random variable governing the event at timestep t as X_t, with corresponding variate x_t. The full process $\mathcal{P}$ is defined by a joint distribution over all times, $P(\ldots, X_0, X_1, X_2, \ldots)$. For shorthand, we will denote a contiguous sequence of events between timesteps t_1 and t_2 as $x_{t_1:t_2} = x_{t_1} x_{t_1+1} \ldots x_{t_2-1}$. For ease of notation we will assume stationary (i.e., time-invariant) processes, such that $P(x_{t:t+L}) = P(x_{t':t'+L}) \;\forall\; t, t', L \in \mathbb{Z}$, though our approach readily generalizes by use of timestep-dependent evolution operators for the quantum models.

Such classical stochastic processes can be generated by hidden Markov models (HMMs). A HMM [28,29] $\mathcal{M}$ is a triple $\mathcal{M} = \{\sigma_0, \mathcal{X}, \{T^x\}_{x \in \mathcal{X}}\}$, where $\mathcal{X}$ is the event alphabet, (column) vector σ_0 the initial state, and $\{T^x\}$ a set of transition operators associated with the events [30]. Mechanistically, the system undergoes a series of probabilistic transitions from the initial state according to the transition operators $\{T^x\}$, accompanied by the emission of an output event symbol $x \in \mathcal{X}$. The probability of a particular sequence $w = x_{0:L}$ occurring is $P(w) = \mathbf{1} T^w \sigma_0$, where $\mathbf{1}$ is a (row) vector of ones, and we define $T^w = T^{x_{L-1}} \ldots T^{x_0}$.

Generalizing to the quantum regime, a quantum HMM (QHMM) [22,31–36] is given by the triple $\mathcal{M}_q = \{\rho_0, \mathcal{X}, \{\mathcal{A}_x\}_{x \in \mathcal{X}}\}$, where density operator ρ_0 is the initial state and $\{\mathcal{A}_x(\cdot)\}$ a set of completely positive maps associated with each event [37]. Again, mechanistically the system undergoes transitions starting from its initial state, with the maps $\{\mathcal{A}_x\}$ playing the role of transition matrices; again, each transition is accompanied by emission of an output event $x \in \mathcal{X}$. To preserve probability normalization, the overall channel $\sum_{x \in \mathcal{X}} \mathcal{A}_x(.)$ must be trace preserving. The probability of a particular sequence $w = x_{0:L}$ occurring is then given by $P(w) = \mathrm{Tr}(\mathcal{A}_w(\rho_0))$, where we have defined $\mathcal{A}_w(\cdot) = \mathcal{A}_{x_{L-1}}(\ldots(\mathcal{A}_{x_0}(\cdot))\ldots)$. As a strict generalization, any HMM can be mapped to a corresponding QHMM [22]. Moreover, QHMMs are more expressive than HMMs of the same state space dimension [25,33], leading to (sometimes unbounded) quantum advantages in memory cost for stochastic process generation, in both lossless and lossy simulation [23,24,38–40]. These advantages underpin our claim of consequent spatial advantage in rare-event sampling with quantum stochastic models.

The approach we will introduce for rare-event sampling can be applied with any QHMM—irrespective of whether it is defined through a mapping from a HMM, or learnt directly from data—as QHMMs in general can be constructed as recurrent unitary circuits that assemble (generalized) quantum sample states. To avoid notational clutter we shall focus much of our exposition in terms of QHMMs as derived from a particular class of HMMs known as deterministic HMMs [41]. We will cover the extension to general QHMMs later in the text.

Deterministic (or "unifilar") HMMs [42,43] are HMMs for which the current state can always be uniquely determined given the previous memory state and event. Given a deterministic HMM $\mathcal{M}$ with states $\{S_j\}$, a corresponding QHMM $\mathcal{M}_q$ always can be constructed with states $\{|\sigma_j\rangle\}$ and evolution operator U according to the following implicit relation [19,20]:

$$U|\sigma_j\rangle|0\rangle = \sum_x \sqrt{P(x|j)}\, e^{i\varphi_{xj}} |\sigma_{\lambda(x,j)}\rangle|x\rangle, \qquad (1)$$

where $\lambda(x, j)$ is the update rule for the state, $P(x|j) = T^x_{\lambda(x,j)j}$, and $\varphi_{xj} \in [0, 2\pi)$ are a set of tunable phases that we shall set here to zero for ease of notation. Measurement of the second subsystem yields the event for the timestep, and

by iterating repeatedly with fresh ancillae for the second subsystem each time, events for arbitrarily many consecutive timesteps can be generated (see Supplemental Material [44]).

If the event ancillae are not measured, then this procedure generates a superposition over all possible event trajectories of the corresponding length [40],

$$\mathcal{U}_L|\sigma_j\rangle|0\rangle^{\otimes L} = \sum_{x_{0:L}} \sqrt{P(x_{0:L}|j)}\,|\sigma_{\lambda(x_{0:L},j)}\rangle|x_{0:L}\rangle, \quad (2)$$

where $\mathcal{U}_L$ is the evolution corresponding to L applications of U across L ancillary subsystems. This final state, the superposition over trajectories, is often referred to as a *quantum sample state* [45], or "qsample" for short. Leveraging the recurrent structure of quantum models, these states are comparatively efficient to generate. Let $d \leq |\{S_j\}|$ be the dimension of the QHMM memory system; as noted above, d can often be made substantially smaller than the dimension of the HMM state space. Then, using that a generic n qubit gate decomposition into one and two qubit gates typically results in $O(4^n)$ gates [46], the recurrent structure corresponds to a total depth $O(Ld^2|\mathcal{X}|^2)$ (see Supplemental Material A [44]). Compared to the naïve unstructured approach that does not leverage the recurrent structure, generically with depth $O(|\mathcal{X}|^{2L})$, this is an exponentially shorter depth with respect to the number of timesteps studied, and is core to making the prospect of temporal complexity advantage for rare-event sampling using QAA viable. As can be seen from the form of the qsample [Eq. (2)], the total space complexity involved is $\log_2 d + L\log_2|\mathcal{X}|$, the first term coming from the memory system, and the latter the L timesteps of output events. The spatial quantum advantage is with respect to the first term, while the second term is identical for both classical and quantum approaches.

QAA [17] generalizes Grover's quantum search algorithm [47], inheriting the provable quadratic speedup over optimal classical approaches. In QAA, the state space is partitioned into "good" and "bad" subspaces, the former containing all target states (as defined according to the particular problem at hand), and the latter the remainder of the state space. An initial state $|\psi\rangle$ can then without loss of generality be written as $|\psi\rangle = \sin(\theta)|\text{good}\rangle + \cos(\theta)|\text{bad}\rangle$. Following repeated application of the operator $\mathcal{G} = -R_\psi R_\tau$ to $|\psi\rangle$, where $R_\psi = I - 2|\psi\rangle\langle\psi|$, $R_\tau = I - 2\Pi_\tau$, I is the identity operator, and Π_τ the projection operator onto the good subspace, the resulting state is

$$\mathcal{G}^n|\psi\rangle = \sin((2n+1)\theta)|\text{good}\rangle + \cos((2n+1)\theta)|\text{bad}\rangle. \quad (3)$$

Given an appropriate n, this rotates the initial state toward its component in the good subspace. In the case that one has access to the operator V that generates $|\psi\rangle$ from some fiducial state (i.e., $V|0\rangle = |\psi\rangle$), and its inverse, then we can replace $R_\psi = VR_0V^\dagger$, where $R_0 = I - |0\rangle\langle 0|$.

We combine QAA with quantum stochastic models by using the quantum model evolution operator and its inverse as the state preparation operator for QAA. Specifically, let U_{init} represent the unitary operator that prepares the initial memory state from the fiducial $|0\rangle$ state, e.g., $|\sigma_j\rangle = U_{\text{init}}|0\rangle$ [48]. Then, we can perform quantum model QAA by setting $V = \mathcal{U}_L U_{\text{init}}$, with L the desired length of the sequence we wish to sample over.

The target projection operator Π_τ is set according to the desired rare event, which we assume to have been determined in advance. For example, if this rare event is a particular sequence $x_{0:L}$, then $\Pi_\tau = I_{\text{mem}} \otimes |x_{0:L}\rangle\langle x_{0:L}|$, where I_{mem} is the identity operator over the memory subspace; R_τ can then be implemented as a single multicontrolled gate for each rare event; this is called once per Grover iteration. Likewise for R_0. These multicontrolled gates can be decomposed into at most $O((\log_2 d + L\log_2|\mathcal{X}|)^2)$ elementary one and two qubit gates [49]. The target rare event sequence can be shorter than L by setting the remaining components of the projector on the output subspace to identity; similarly, the rare events can consist of particular non-contiguous outputs, interspersed by identity operators between target positions. Further, we can sample over multiple sequences in our rare event by taking a combined target projection operator that sums over individual projectors.

If the total probability of the rare event(s) happening is p, then the amplitude of the "good" subspace is $\sqrt{p}$; note that we may neglect phase, as this can be rolled into $|\text{good}\rangle$ without loss of generality. For $p \ll 1$, this corresponds to $\theta \approx \sqrt{p}$. Thus, $n \sim O(1/\sqrt{p})$ Grover iterations are required to bring the target "good" subspace to approximately unit amplitude. Direct measurement can then be made to verify that we are in the target subspace. This corresponds to a quadratic reduction in query complexity compared to classical Monte Carlo–based methods, inherited from QAA. In the particular context of this problem and approach, this can also be reframed in terms of temporal complexity. In classical Monte Carlo, trajectories are generated and each checked to see if they contain rare events; the temporal complexity is linearly related to query complexity. Similarly, the gate depth (and temporal complexity) of QAA implemented in this manner is also linearly related to query complexity, with depth $O(Ld^2|\mathcal{X}|^2 + (\log_2 d + L\log_2|\mathcal{X}|)^2)/\sqrt{p}$ [50]. We can thus alternatively consider the advantage as a quadratic quantum reduction in temporal complexity. Coupled with the reduction in spatial complexity in the size of the memory inherited from the quantum models, this results in a quantum spatiotemporal advantage in rare-event sampling. We note again that use of the structure of quantum stochastic models results in a potentially exponential reduction in circuit depth, for each call of the preparation unitary, compared to a naïve unstructured assembly of the quantum sample state for the future trajectories.

This approach also directly offers a quantum advantage in estimating the probability of rare events [51,52]. Estimating the rare event probability corresponds to estimating the success probability of a binomial distribution, as each trajectory corresponds to either a rare event or not. Since the variance of a binomial distribution with success probability p is $p(1-p)/n$ for n trials [53], the confidence interval Δp for an estimate of p scales as $\sqrt{p(1-p)/n}$, with the prefactor dependent on the particular confidence interval. To constrain the estimate to a fixed proportional accuracy, i.e., the relative error, $\delta = \Delta p/p \propto \sqrt{(1-p)/np}$ thus requires $n \sim O(1/p\delta^2)$ trials, for $p \ll 1$. With QAA we can amplify the good subspace until its modified probability $p' \approx 1/2$. Then, the number of trials for fixed accuracy $\delta = \Delta p'/p'$ scales as $O(1/\delta^2)$. Considering that QAA requires $O(1/\sqrt{p})$ calls to the quantum sample state preparation unitary, this leads to a total query complexity $O(1/\sqrt{p}\delta^2)$, a quadratic advantage in p compared to the classical approach. The spatial complexity reduction of the quantum models also applies.

As an example, we apply our approach to a well-studied toy model, the Nemo process [54], modified to include an "error" symbol ("2") with small probability. The associated HMM is shown in Fig. 2(a). For certain values of p and q, this process admits a quantum model with reduced memory dimension $d = 2$ (compared to the classical $d = 3$). For particular ratios of ϵ_B and ϵ_C, the error terms can be incorporated within the $d = 2$ quantum model, with a free parameter controlling their absolute size. Details are given in the End Matter.

We take $p = 0.5$, $q \approx 0.432$, and $\epsilon_B, \epsilon_C \sim O(10^{-6})$, and consider trajectories of length 5. Starting from state σ_A, we seek trajectories for which the final symbol is an error. Given these occur with probability $O(10^{-6})$, we require $O(10^3)$ Grover iterations to maximize the rare event trajectory probabilities. For 10^4 shots, we show the progression of amplification in Figs. 2(b)–2(e) with increasing numbers of Grover iterations, observing the expected reweighting in favor of the rare events.

We then use these amplified trajectories in order to sample the distribution of the preceding symbols when such a rare event occurs. We compare this with the exact values, calculated analytically [Fig. 2(f)]. We also compare this with a classical Monte Carlo–based approach, which, for fair comparison, is given a number of shots equal to the number of times the quantum model evolution unitary is called, i.e., 2× number of Grover iterations (1024)× number of shots of quantum circuit (10^4), corresponding to 2.048×10^7 shots in total. For the classical approach, only 13 rare event trajectories were generated from these shots; correspondingly, reconstruction of the distribution of symbols leading up to rare events is much poorer than that achieved by our quantum approach, for which 8796 trajectories were sampled from after amplification.

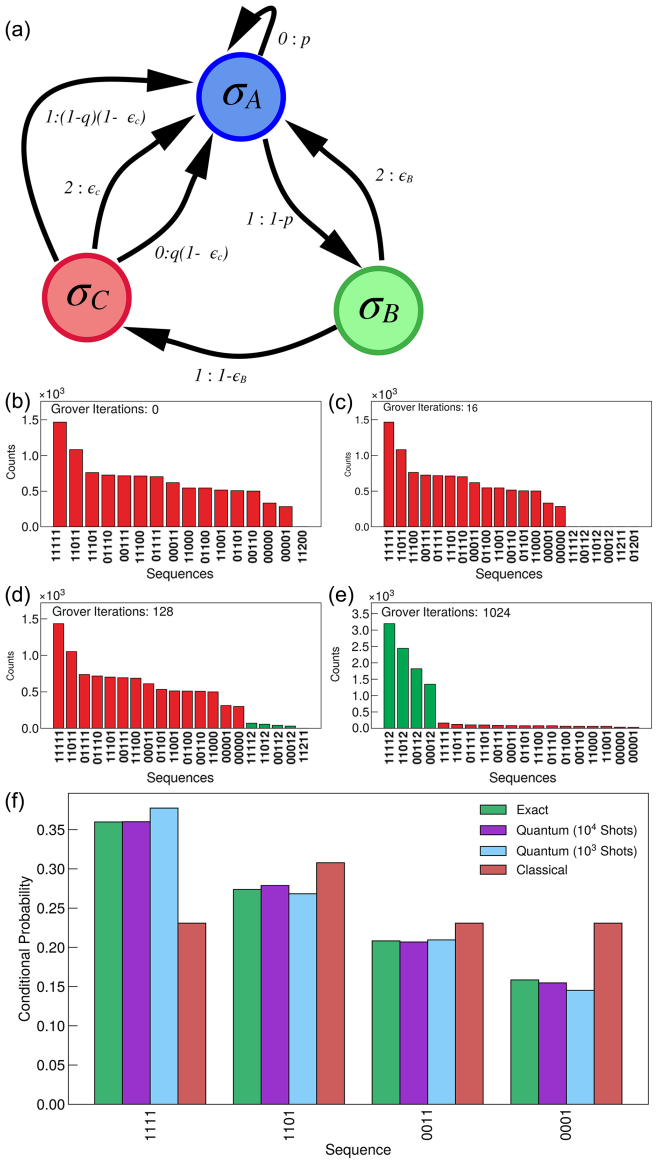


FIG. 2. (a) Hidden Markov model representation of the Nemo process with error symbols; notation $x\!:p$ denotes the probability p of the indicated transition occurring, accompanied by output symbol x. Progression of reweighting of rare events for (b) 0, (c) 16, (d) 128, and (e) 1024 Grover iterations; green denotes target rare events, red the remaining trajectories. (f) Comparison of reconstructed conditional probabilities for preceding symbols to a rare event for classical and quantum approaches; sequences with multiple errors omitted due to vanishingly small probabilities of occurrence.

Moreover, we see that even with an order of magnitude fewer shots of the quantum model a reasonably accurate reconstruction is achieved; in comparison, the analogous number of classical shots yielded only a single rare event trajectory.

In the End Matter, we also demonstrate application of our approach to a further toy model based on a renewal process—representative of example application areas such

as queueing theory [55] and reliability analysis [56]—for which the corresponding QHMM has been experimentally implemented [24]. This latter example demonstrates a starker spatial complexity advantage ($d = 2$, while $|\{S_j\}| = 25$), a much rarer event ($p \sim O(10^{-15})$) corresponding to a greater temporal complexity advantage, and a larger number of timesteps.

The approach is readily extended to general HMMs, even when the property of determinism is not met. In such circumstances, to convert a HMM with transition matrices $\{T^x\}$ into a QHMM we must modify Eq. (1) with an additional "junk" ancillary system. A valid approach for any HMM is as follows [22]:

$$U|\sigma_j\rangle|0\rangle|0\rangle = \sum_{xk} \sqrt{T^x_{kj}}\,|\sigma_k\rangle|x\rangle|k\rangle, \tag{4}$$

where the third subsystem is the junk, discarded at each timestep. Keeping the junk, along with the outputs, results in a modified quantum sample state after L steps:

$$\mathcal{U}_L|\sigma_j\rangle|0\rangle^{\otimes 2L} = \sum_{x_{0:L}\boldsymbol{k}} \sqrt{T^{x_{L-1}}_{k_L k_{L-1}} \ldots T^{x_0}_{k_1 j}}\,|\sigma_{k_L}\rangle|x_{0:L}\rangle|\boldsymbol{k}\rangle, \tag{5}$$

where $\boldsymbol{k} = (k_1, k_2, \ldots k_L)$ is the trajectory through the HMM's state space. The quantum model QAA approach then proceeds as above, but with the new quantum sample state preparation unitary as prescribed in Eq. (5) above, and an extension of the projection operator Π_τ to include an identity operator for the junk subsystems carrying the state space trajectory. We remark that the specific assignment of the junk subsystems as prescribed above will generally negatively impact the quantum advantage in spatial complexity, while the temporal/query complexity advantage remains unaffected. There often exists alternative junk encoding strategies that *do* yield a spatial complexity advantage [22,57]; Ref. [57] provides stricter bounds on the dimension of the junk system based on the "maximal branching" of the underlying HMM. Moreover, we emphasize that nondeterminism is a property of a HMM, *not* the underlying process, which permits multiple alternative, distinct HMM representations yielding identical statistics.

For a general QHMM $\mathcal{M}_q$, given the decomposition of the maps $\{\mathcal{A}_x\}$ in terms of their Kraus decomposition $\{\mathcal{A}_x(\cdot) = \sum_y K_{xy} \cdot K^\dagger_{xy}\}$, the corresponding unitary implementation is given by

$$U|\cdot\rangle|0\rangle|0\rangle = \sum_{xy} (K_{xy}|\cdot\rangle)|x\rangle|y\rangle. \tag{6}$$

Iteratively applying this with fresh ancillae at each timestep, one obtains the requisite $\mathcal{U}_L$ needed for input into QAA. Thus, our framework can be applied with *any* QHMM.

In summary, we have introduced an approach that combines quantum stochastic modelling with QAA to realize a spatiotemporal reduction in complexity for rare-event sampling, incorporating memory advantages from the former with speedups from the latter. Such memory-efficient quantum models moreover enable a comparatively low-circuit depth means to assemble quantum sample states for QAA compared to generic unstructured approaches, making practical realization of the approach more amenable. While our exposition is, for ease of presentation, in terms of HMMs with known transition matrices, the approach can be deployed with any QHMM. Correspondingly, lossy compression of quantum models [39,40] may be incorporated, thus further enhancing the spatial complexity advantage, as well as quantum models inferred directly from data. As remarked above, our approach is fully complementary to that of Ref. [26]; using such biased quantum models could be employed as the input models for our approach, further integrating and enhancing quantum spatial and temporal complexity advantages.

QAA and QAE are central algorithms in proposals for the utility of quantum computing in enhancing computational power in a broad range of fields, such as finance [58–62]. By overcoming a significant obstacle in such algorithms—that of efficiently preparing the input quantum sample state—we offer a significant step-change toward the practical deployment of such approaches. One may readily expect that this approach readily generalizes to other related quantum algorithms, such as quantum-enhanced rare event discovery and Monte Carlo integration [63–65]. For very rare events, our approach may still require deep circuits. However, it may be possible to use circuit recompilation [66] to overcome this, leveraging the repeating block nature of the circuits involved. We leave detailed exploration of this to future work.

An enticing further generalization is to input-output processes [67], where the observed behaviour is conditioned upon input stimuli presented to the system. Quantum models of input-output processes with memory advantages have been discovered [68,69]; using such a construction in place of the quantum stochastic models would allow quantum-enhanced rare event sampling to be performed for such processes. Notably, this offers the ability to input superpositions over inputs, which may be then sampled over with respect to a specific output string; this may find application in reinforcement learning [70,71], by allowing an agent to deduce the input sequence(s) that have the highest probability to yield a target output.

In our motivation we have focussed on extreme risk events, such as natural disasters and global pandemics, i.e., catastrophic events that can lead to significant loss of life, substantial economic damage, or large-scale social upheaval. Yet, the importance of rare-event sampling stretches far beyond this, with applications across the sciences. For example, in high-energy physics, experiments at the Large Hadron Collider analyze proton-proton collisions where the probabilities of physical processes span

over 10 orders of magnitude, requiring simulation of enormous background processes to contextualize rare signals [72]. Monte Carlo generators [73] employ sampling methods to simulate these interactions, generating collision events used to compare experimental data with theoretical predictions and search for evidence of new physics phenomena [74–76]. Yet, these generators face a severe computational challenge: generating simulated datasets statistically comparable to data is prohibitive, especially if samples predominantly contain high-probability, known processes. This leads to a need for effective and efficient rare-event sampling algorithms that can bias samples toward generation toward low-probability, high-impact regions of phase space. Indeed, with many big data fields beset by bottlenecks due to limits of computational power, attention is turning to the prospect of quantum technologies to overcome such limits. Our approach may find further utility in such scenarios.

Acknowledgments—This work was funded by EPSRC Studentship EP/W524347/1 (Project No. 2932638); the University of Manchester Dame Kathleen Ollerenshaw Fellowship; the National Research Foundation, Singapore through the National Quantum Office, hosted in A*STAR, under its Centre for Quantum Technologies Funding Initiative (S24Q2d0009); the RIE 2025 AQAS projects S25Q9DA001 and S25Q9DA002; the National Research Foundation of Singapore through the NRF Investigatorship Program (Award No. NRF-NRFI09-0010); and the Singapore Ministry of Education Tier 1 Grant RT4/23 and RG91/25.

Data availability—There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

[1] Y. Hong, Y. Liu, and S. Wang, Granger causality in risk and detection of extreme risk spillover between financial markets, J. Econ. **150**, 271 (2009).
[2] T. Aven, On the meaning of a black swan in a risk context, Saf. Sci. **57**, 44 (2013).
[3] L. E. McPhillips, H. Chang, M. V. Chester, Y. Depietri, E. Friedman, N. B. Grimm, J. S. Kominoski, T. McPhearson, P. Méndez-Lázaro, E. J. Rosi *et al.*, Defining extreme events: A cross-disciplinary review, Earth's Future **6**, 441 (2018).
[4] C. Raymond, R. M. Horton, J. Zscheischler, O. Martius, A. AghaKouchak, J. Balch, S. G. Bowen, S. J. Camargo, J. Hess, K. Kornhuber *et al.*, Understanding and managing connected extreme events, Nat. Clim. Change **10**, 611 (2020).
[5] M. Lewis, *The Big Short: Inside the Doomsday Machine* (Penguin, United Kingdom, 2011).
[6] J. A. Bucklew and J. Bucklew, *Introduction to Rare Event Simulation* (Springer, New York, 2004), Vol. 5.
[7] R. J. Allen, C. Valeriani, and P. R. Ten Wolde, Forward flux sampling for rare event simulations, J. Phys. Condens. Matter **21**, 463102 (2009).
[8] F. Bouchet, J. Rolland, and J. Wouters, Rare event sampling methods, Chaos **29** (2019).
[9] C. J. Geyer, Practical Markov chain Monte Carlo, Stat. Sci. **7**, 473 (1992).
[10] S. Brooks, Markov chain Monte Carlo method and its application, J. R. Stat. Soc. Ser. D **47**, 69 (1998).
[11] J. C. Chan and D. P. Kroese, Rare-event probability estimation with conditional monte carlo, Ann. Oper. Res. **189**, 43 (2011).
[12] F. Cérou, P. Del Moral, T. Furon, and A. Guyader, Sequential Monte Carlo for rare event estimation, Stat. Comput. **22**, 795 (2012).
[13] S. Chib and E. Greenberg, Understanding the metropolis-Hastings algorithm, Am. Stat. **49**, 327 (1995).
[14] C. P. Robert, G. Casella, and G. Casella, *Monte Carlo Statistical Methods* (Springer, New York, 1999), Vol. 2.
[15] A. Laio and F. L. Gervasio, Metadynamics: A method to simulate rare events and reconstruct the free energy in biophysics, chemistry and material science, Rep. Prog. Phys. **71**, 126601 (2008).
[16] O. Valsson, P. Tiwary, and M. Parrinello, Enhancing important fluctuations: Rare events and metadynamics from a conceptual viewpoint, Annu. Rev. Phys. Chem. **67**, 159 (2016).
[17] G. Brassard, P. Høyer, M. Mosca, and A. Tapp, Quantum amplitude amplification and estimation, arXiv:quant-ph/0005055.
[18] M. Gu, K. Wiesner, E. Rieper, and V. Vedral, Quantum mechanics can reduce the complexity of classical models, Nat. Commun. **3**, 762 (2012).
[19] F. C. Binder, J. Thompson, and M. Gu, Practical unitary simulator for non-Markovian complex processes, Phys. Rev. Lett. **120**, 240502 (2018).
[20] Q. Liu, T. J. Elliott, F. C. Binder, C. Di Franco, and M. Gu, Optimal stochastic modeling with unitary quantum dynamics, Phys. Rev. A **99**, 062110 (2019).
[21] 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).
[22] T. J. Elliott, Memory compression and thermal efficiency of quantum implementations of nondeterministic hidden Markov models, Phys. Rev. A **103**, 052615 (2021).
[23] 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).
[24] K.-D. Wu, C. Yang, R.-D. He, M. Gu, G.-Y. Xiang, C.-F. Li, G.-C. Guo, and T. J. Elliott, Implementing quantum dimensionality reduction for non-Markovian stochastic simulation, Nat. Commun. **14**, 2624 (2023).
[25] M. Fanizza, J. Lumbreras, and A. Winter, Quantum theory in finite dimension cannot explain every general process with finite memory, Commun. Math. Phys. **405**, 50 (2024).
[26] 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).

[27] A. Khintchine, Korrelationstheorie der stationären stochastischen Prozesse, Math. Ann. **109**, 604 (1934).

[28] L. Rabiner and B. Juang, An introduction to hidden Markov models, IEEE ASSP Mag. **3**, 4 (1986).

[29] D. R. Upper, Theory and algorithms for hidden Markov models and generalized hidden Markov models, Ph.D. thesis, University of California, Berkeley, 1997, Published by University Microfilms International, Ann Arbor, Michigan.

[30] An implicit state space for the HMM is defined according to an orthonormal basis that spans the reachable vectors $T^w\sigma_0$, for $w \in \mathcal{X}^*$.

[31] A. Monras, A. Beige, and K. Wiesner, Hidden quantum Markov models and non-adaptive read-out of many-body states, Appl. Math. Comp. Sci. **3**, 93 (2011).

[32] L. A. Clark, W. Huang, T. M. Barlow, and A. Beige, Hidden quantum Markov models and open quantum systems with instantaneous feedback, in *ISCS 2014: Interdisciplinary Symposium on Complex Systems* (Springer, New York, 2015), pp. 143–151.

[33] A. Monras and A. Winter, Quantum learning of classical stochastic processes: The completely positive realization problem, J. Math. Phys. (N.Y.) **57**, 015219 (2016).

[34] S. Adhikary, S. Srinivasan, G. Gordon, and B. Boots, Expressiveness and learning of hidden quantum Markov models, in *International Conference on Artificial Intelligence and Statistics* (PMLR, 2020), pp. 4151–4161, https://proceedings.mlr.press/v108/adhikary20a.html.

[35] M. Zonnios, F. C. Binder, and A. Boyd, Quantum generation of stochastic processes: Spectral invariants and memory bounds, New J. Phys. **27**, 064507 (2025).

[36] P. M. Riechers and T. J. Elliott, Identifiability and minimality bounds of quantum and post-quantum models of classical stochastic processes, arXiv:2509.03004.

[37] As with HMMs, an implicit state space for a QHMM is defined according to the reachable states $\mathcal{A}_w(\rho_0)$, for $w \in \mathcal{X}^*$.

[38] 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).

[39] T. J. Elliott, Quantum coarse graining for extreme dimension reduction in modeling stochastic temporal dynamics, PRX Quantum **2**, 020342 (2021).

[40] C. Yang, M. Florido-Llinàs, M. Gu, and T. J. Elliott, Dimension reduction in quantum sampling of stochastic processes, npj Quantum Inf. **11**, 34 (2025).

[41] Note that *deterministic* here refers to the update rule for the model's state given the current state and output; the model's output is still stochastic.

[42] J. P. Crutchfield and K. Young, Inferring statistical complexity, Phys. Rev. Lett. **63**, 105 (1989).

[43] C. R. Shalizi and J. P. Crutchfield, Computational mechanics: Pattern and prediction, structure and simplicity, J. Stat. Phys. **104**, 817 (2001).

[44] See Supplemental Material at http://link.aps.org/supplemental/10.1103/7zh3-zvnd for additional technical details supporting the claims of the manuscript.

[45] M. Schuld, F. Petruccione, M. Schuld, and F. Petruccione, Quantum advantages, in *Supervised Learning with Quantum Computers* (Springer, Cham, 2018), pp. 127–137.

[46] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).

[47] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (1996), pp. 212–219, 10.1145/237814.237866.

[48] In the case where we wish to initialize in a distribution over memory states, e.g., $\rho = \sum_j P(j)|\sigma_j\rangle\langle\sigma_j|$, then we instead use U_{init} to prepare a purification of ρ, such as $\sum_j \sqrt{P(j)}|j\rangle|\sigma_j\rangle$, and use the second subsystem for the memory in the remainder of the algorithm.

[49] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, Phys. Rev. A **52**, 3457 (1995).

[50] The first term in the numerator corresponds to the depth in preparing (and inverting) the sample state, the second the cost of a multicontrol gate, and the denominator coming from the requisite number of iterations.

[51] Y. Suzuki, S. Uno, R. Raymond, T. Tanaka, T. Onodera, and N. Yamamoto, Amplitude estimation without phase estimation, Quantum Inf. Process. **19**, 75 (2020).

[52] S. Simon, M. Degroote, N. Moll, R. Santagati, M. Streif, and N. Wiebe, Amplified amplitude estimation: Exploiting prior knowledge to improve estimates of expectation values, arXiv:2402.14791.

[53] M. L. Boas, *Mathematical Methods in the Physical Sciences* (Wiley, New York, 2006).

[54] 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).

[55] J. F. Shortle, J. M. Thompson, D. Gross, and C. M. Harris, *Fundamentals of Queueing Theory* (John Wiley & Sons, New York, 2018).

[56] R. Arnold, S. Chukova, Y. Hayakawa, and S. Marshall, Geometric-like processes: An overview and some reliability applications, Reliab. Eng. Syst. Saf. **201**, 106990 (2020).

[57] R. Sundar and T. J. Elliott, Quantum dimension reduction of hidden Markov models, Phys. Rev. A **113**, 062461 (2026).

[58] R. Orús, S. Mugel, and E. Lizaso, Quantum computing for finance: Overview and prospects, Rev. Phys. **4**, 100028 (2019).

[59] S. Woerner and D. J. Egger, Quantum risk analysis, npj Quantum Inf. **5**, 15 (2019).

[60] D. J. Egger, C. Gambella, J. Marecek, S. McFaddin, M. Mevissen, R. Raymond, A. Simonetto, S. Woerner, and E. Yndurain, Quantum computing for finance: State-of-the-art and future prospects, IEEE Trans. Quantum Eng. **1**, 3101724 (2020).

[61] S. Wilkens and J. Moorhouse, Quantum computing for financial risk measurement, Quantum Inf. Process. **22**, 51 (2023).

[62] K. Ghosh, K. Yogaraj, G. Agliardi, P. Sabino, M. Fernández-Campoamor, J. Bernabé-Moreno, G. Cortiana, O. Shehab, and C. O'Meara, Energy risk analysis with dynamic amplitude estimation and piecewise approximate quantum compiling, IEEE Trans. Quantum Eng. **5**, 3102717 (2024).

[63] A. Montanaro, Quantum speedup of Monte Carlo methods, Proc. R. Soc. A **471**, 20150301 (2015).
[64] C. Blank, D. K. Park, and F. Petruccione, Quantum-enhanced analysis of discrete stochastic processes, npj Quantum Inf. **7**, 126 (2021).
[65] S. Herbert, Quantum Monte Carlo integration: The full advantage in minimal circuit depth, Quantum **6**, 823 (2022).
[66] T. Jones and S. C. Benjamin, Robust quantum compilation and circuit optimisation via energy minimisation, Quantum **6**, 628 (2022).
[67] N. Barnett and J. P. Crutchfield, Computational mechanics of input–output processes: Structured transformations and the ϵ-transducer, J. Stat. Phys. **161**, 404 (2015).
[68] 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).
[69] T. J. Elliott, M. Gu, A. J. Garner, and J. Thompson, Quantum adaptive agents with efficient long-term memories, Phys. Rev. X **12**, 011007 (2022).
[70] L. P. Kaelbling, M. L. Littman, and A. W. Moore, Reinforcement learning: A survey, J. Artif. Intell. Res. **4**, 237 (1996).
[71] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction* (MIT press, Boston, Massachusetts, 2018).
[72] L. Clissa, M. Lassnig, and L. Rinaldi, How big is big data? A comprehensive survey of data production, storage, and streaming in science and industry, Front. Big Data **6**, 1271639 (2023).
[73] A. Buckley, J. Butterworth, S. Gieseke, D. Grellscheid, S. Höche, H. Hoeth, F. Krauss, L. Lönnblad, E. Nurse, P. Richardson, S. Schumann, M. H. Seymour, T. Sjöstrand, P. Skands, and B. Webber, General-purpose event generators for LHC physics, Phys. Rep **504**, 145 (2011).
[74] Particle Data Group, Review of particle physics, Phys. Rev. D **110**, 030001 (2024).
[75] CMS Collaboration, Dark sector searches with the CMS experiment, Phys. Rep. **1115**, 448 (2025).
[76] ATLAS Collaboration, The quest to discover supersymmetry at the ATLAS experiment, Phys. Rep. **1116**, 261 (2025).
[77] Specifically, we obtain $c_{AB} = (\alpha\beta^* + \sqrt{2}\beta)/(2 - |\alpha|^2)$.
[78] S. E. Marzen and J. P. Crutchfield, Informational and causal architecture of discrete-time renewal processes, Entropy **17**, 4891 (2015).

End Matter

Here we provide further details on the example in the main text. In the absence of the error symbols, the quantum model consists of three quantum memory states $\{|\sigma_A\rangle, |\sigma_B\rangle, |\sigma_C\rangle\}$, and an evolution operator U, defined implicitly through

$$\begin{aligned} U|\sigma_A\rangle|0\rangle &= \sqrt{p}|\sigma_A\rangle|0\rangle + \sqrt{1-p}|\sigma_B\rangle|1\rangle \\ U|\sigma_B\rangle|0\rangle &= |\sigma_C\rangle|1\rangle \\ U|\sigma_C\rangle|0\rangle &= \sqrt{q}e^{i\phi}|\sigma_A\rangle|0\rangle + \sqrt{1-q}e^{i\varphi}|\sigma_A\rangle|1\rangle, \end{aligned} \quad \text{(A1)}$$

where ϕ and φ are free, tunable parameters. Quantum dimension reduction is possible if the quantum memory states are linearly dependent, i.e., we can express $|\sigma_C\rangle = \alpha|\sigma_A\rangle + \beta|\sigma_B\rangle$, for some α and β. By direct substitution into the right-hand side of Eq. (A1), it can be shown that this is satisfied when $\alpha = \sqrt{q/p}$ and $\beta = i(q(1-p)/p)^{1/4}$, under the condition that $(1-q)^2 = q^3(1-p)/p^3$. For our particular choice of $p = 0.5$, this requires that $q = 1/12 + \sqrt[3]{181/1728 + \sqrt{13/864}} + \sqrt[3]{181/1728 - \sqrt{13/864}} \approx 0.432$.

Without loss of generality, we can set $|\sigma_A\rangle = |0\rangle$, and $|\sigma_B\rangle = c_{AB}|0\rangle + \sqrt{1-|c_{AB}|^2}|1\rangle$, where $c_{AB} = \langle\sigma_A|\sigma_B\rangle$, as can be calculated using Eq. (A1) [77]. Following the procedure of Refs. [19,20], we can determine the action of U on the basis states. We then introduce the error symbol by modifying U into a new unitary operator U_ϵ as follows:

$$\begin{aligned} U_\epsilon|0\rangle|0\rangle &= U|0\rangle|0\rangle \\ U_\epsilon|1\rangle|0\rangle &= \sqrt{1-\delta}U|1\rangle|0\rangle + \sqrt{\delta}|0\rangle|2\rangle, \end{aligned} \quad \text{(A2)}$$

where the remaining columns can be set arbitrarily, provided they maintain orthogonality with each other and the prescribed columns. Here, δ is a parameter that controls the strength of the error; for our simulation, we set $\delta = 10^{-6}$. We can relate this to the error probabilities as follows: $\epsilon_B = (1 - |c_{AB}|^2)\delta$ and $\epsilon_C = |\beta|^2\epsilon_B$. For our parameters, this corresponds to $\epsilon_B \approx 0.88 \times 10^{-6}$ and $\epsilon_C \approx 0.58 \times 10^{-6}$.

As a second example, we consider a periodically modulated decay process [24]. This is a renewal process [78], which can be represented as a deterministic HMM with alphabet $\mathcal{X} = \{0, 1\}$, N states $\{S_n\}$ $n = 0, 1, \ldots, N-1$, with transitions satisfying $P(0|S_n) = \Phi(n+1)/\Phi(n)$, where emissions of 0 are accompanied by a transition $S_n \to S_{(n+1) \bmod N}$ and emissions of 1 $S_n \to S_0$. The "survival probability" $\Phi(n) = \Gamma^n(1 - V\sin^2(n\pi/N))$. This process permits a significant quantum advantage in spatial complexity; while the minimal classical model requires a $\log_2 N$ bit memory, a quantum model can be constructed that requires only a single qubit of memory. We consider the case where $N = 25$, $V = 1/3.05$, and $\Gamma \approx 0.9511$. Specifically, we follow the procedure detailed in Supplemental Material of Ref. [24],

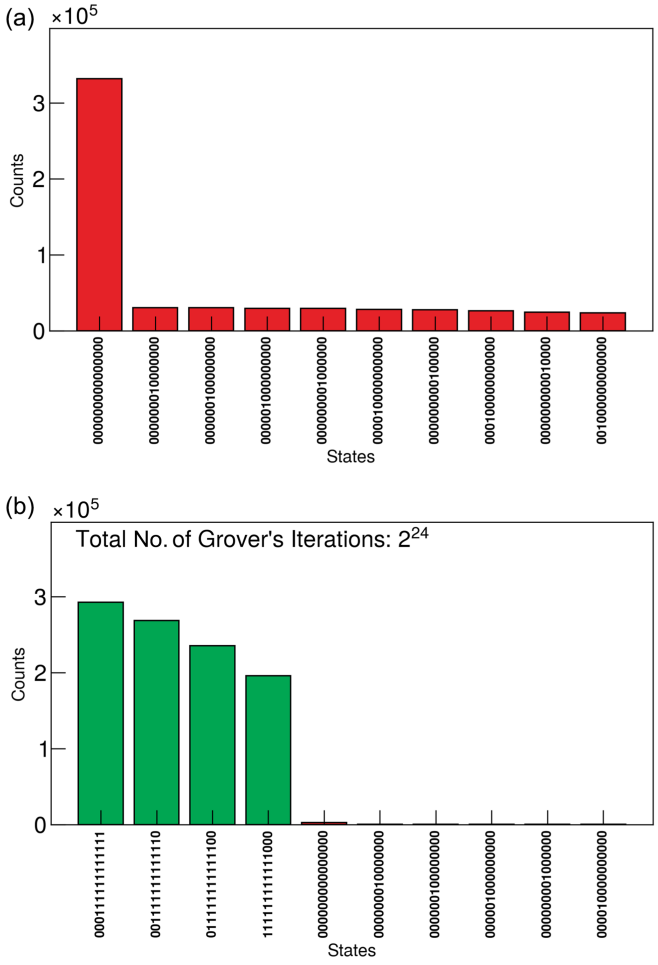


FIG. 3. Sampled trajectories for the PMD process (a) without amplification, and (b) with 2^{24} Grover iterations; green denotes target rare events, red the remaining trajectories. Circuits evaluated with 10^6 shots. Plots limited to ten highest-probability sequences. Plots of the progression of amplification are given in Supplemental Material [44].

resulting in Kraus operators

$$K_0 = \begin{pmatrix} 0.97525 & -0.0061278 + 0.048507i \\ 00.94461 + 0.24253i & \end{pmatrix}$$
$$K_1 = \begin{pmatrix} -0.16987 & -0.020763 + 0.16436i \\ -0.14156 & -0.017302 + 0.13696i \end{pmatrix} \quad \text{(A3)}$$

(up to 5 significant figures).

We consider 15 timesteps of evolution, starting from the S_0 state. As our rare event, we consider sequences containing exactly 12 consecutive 1 s, with the remaining events all 0. There are four such sequences (000111111111111, 001111111111110, 011111111111100, 111111111111000), with total probability $p \approx 2.445 \times 10^{-15}$. Thus, classical Monte Carlo sampling has a below 25% probability of yielding a single rare event sequence in 10^{14} trajectories. In contrast, we see that after $2^{24} \approx 1.68 \times 10^7$ rounds of amplitude amplification, sampling from the quantum state with high probability ($\approx$0.992) yields a rare event sequence [Fig. 3]. Thus, there is a substantial quantum advantage also in query and temporal complexity.