

Sampling Complexity of Open Quantum Systems


I.A. Aloisio^{1,*}, G.A.L. White^{1,2,†}, C.D. Hill^{3,‡} and K. Modi^{1,4,§}

¹*School of Physics and Astronomy, Monash University, Clayton, Victoria 3800, Commonwealth of Australia*

²*School of Physics, University of Melbourne, Parkville, Victoria 3010, Commonwealth of Australia*

³*Silicon Quantum Computing, University of New South Wales, Sydney, New South Wales 2052, Commonwealth of Australia*

⁴*Centre for Quantum Technology, Transport for New South Wales, Sydney, New South Wales 2000, Commonwealth of Australia*

 (Received 13 October 2022; revised 12 February 2023; accepted 14 March 2023; published 19 April 2023)

Open quantum systems are ubiquitous in the physical sciences, with widespread applications in the areas of chemistry, condensed-matter physics, material science, optics, and many more. Not surprisingly, there is significant interest in their efficient simulation. However, direct classical simulation quickly becomes intractable with coupling to an environment with an effective dimension that grows exponentially. This raises the question: Can quantum computers help model these complex dynamics? A first step in answering this question requires an understanding of the computational complexity of this task. Here, we map the temporal complexity of a process to the spatial complexity of a many-body state using a computational model known as the process-tensor framework. With this, we are able to explore the simulation complexity of an open quantum system as a dynamic sampling problem: a system coupled to an environment can be probed at successive points in time—accessing multitime correlations. The complexity of multitime sampling, which is an important and interesting problem in its own right, contains the complexity of master equations and stochastic maps as a special case. Our results show how the complexity of the underlying quantum stochastic process corresponds to the complexity of the associated family of master equations for the dynamics. We present both analytical and numerical examples where the multitime sampling of an open quantum system is as complex as sampling from a many-body state that is classically hard. This also implies that the corresponding family of master equations is classically hard. Our results pave the way for studying open quantum systems from a complexity-theoretic perspective, highlighting the role quantum computers will play in our understanding of quantum dynamics.

DOI: [10.1103/PRXQuantum.4.020310](https://doi.org/10.1103/PRXQuantum.4.020310)

I. INTRODUCTION

With the rapid improvement in quantum computing technologies, there is increasing interest in finding practical problems that will demonstrate a quantum advantage over classical methods. A leading candidate is the simulation of strongly correlated quantum systems [1]. A compelling subclass of these dynamics is that of open quantum systems, the correlations of which lie between

an accessible system and an inaccessible environment. These constitute some of the most challenging yet fascinating physical phenomena and underlie the fields of quantum chemistry, materials science, and condensed-matter physics [2–6]. Contexts range from the study of strongly correlated materials such as glass-forming liquids [7] to conformational dynamics of proteins [8] and to understanding the efficiency of photosynthesis [9].

Understanding the dynamics of open quantum systems—in terms of master equations, stochastic maps, or their multitime correlations—requires additional considerations over closed ones. Naively, classical simulation of a small system may seem tractable but accounting for the relevant effects of environmental coupling on system evolution increases the problem manifold. From the opposite angle, one might ask whether the simulation of an open quantum system is as difficult as solving the entire dilated dynamics. But this approach is asking too much: the most compressed description of a system can

*isobel.aloisio@monash.edu

†greg.white@monash.edu

‡charles.hill1@unsw.edu.au

§kavan.modi@monash.edu

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

often be ignorant to the specifics of the environment. The true difficulty of the problem lies somewhere in between. Coherent system-environment (*SE*) interactions can lead to an exchange of information between the two and can generate highly complex temporal correlations across the system dynamics. In other words, it is the complexity of the memory—or the non-Markovianity—that governs the difficulty of the problem. At one end of the spectrum, memoryless dynamics are almost as easy to solve as the closed dynamics of the system. At the other extreme, the evolution of the system can be every bit as complex as the system and its environment. Precisely identifying, defining, and exploring this transition is the purpose of this work.

So, then, exactly how complex are non-Markovian processes from a complexity-theoretic perspective? How hard is it to classically simulate a master equation, a stochastic map, or multitime correlations? Answering these questions will inform us about which physical problems may be good candidates for simulating on quantum computers and which are solvable with classical methods. Moreover, this will inform as to how quantum simulations should be designed on real quantum devices, to extract meaningful information from complex quantum processes [10–12]. And, more fundamentally, it is the key to determining if multitime correlations of a single qubit can be complex, if a noisy process can be complex, and, consequently, if we will see any computational advantage using noisy intermediate-scale quantum (NISQ) devices.

The efforts to model the non-Markovian quantum dynamics on classical computers are vast and the *simulation complexity* of many models is well studied. However, typically, the complexity of classically simulating open quantum systems is attributed to the difficulty of solving the closed dynamics of the system and its environment, which we refer to as *dilated complexity*. In this setting, the computational cost of direct simulation scales exponentially with the size of the environment. Numerous classical techniques exist to tackle the simulation of these complex environments [13,14], including automated environment compression [15], time-evolving matrix product operators [16], path-integral-based approximations [17–19], the ensemble of Lindblad’s trajectories [20], and quantum Monte Carlo [21–23].

Specific dimension-based metrics have been proposed to quantify the simulation complexity. This includes the transfer-tensor method [24–27], the matrix bond dimension of the minimum effective environment [28], and the rank of the time-delay matrix [29]. Yet, their relation to specific physics models remains tenuous. It is important to note that the dilated complexity will always be greater than or equal to that of the actual simulation complexity. In other words, the latter should not be burdened by the cumulative task of incorporating all environmental degrees of freedom and then subsequently reducing the problem

back to the open dynamics by means of a partial trace. On the other hand, the classical algorithms listed above are not always accurate with long-time dynamics or strong system-environment coupling [30]. This suggests that the simulation complexity may indeed be on par with the dilated complexity in these settings.

Our approach is to think about simulating open quantum dynamics, which includes master equations, stochastic maps, and multitime correlations, as a quantum sampling problem. Indeed, from a complexity perspective, sampling complexity is used as a definition for simulation complexity, since it replicates what is obtainable in an experiment where a quantum state is prepared, evolved for some fixed time, and then measured [31]. This perspective has led to several key proposals for quantum advantage, including BosonSampling [32], instantaneous quantum polynomial-time (IQP) [33], and random circuit sampling (RCS) [34], as well as more physical models such as driven many-body systems [35], Ising spin models [36], and bosonic lattice models [37]. However, attempts to incorporate the effects of interaction with an uncontrollable environment on the sampling complexity of a quantum system have been restricted to Markovian, or memoryless, processes [38,39]. Once again, the conventional sampling problem considers the dilated simulation complexity of a closed system. Thus, our key proposition for open quantum dynamics is to sample a single system at many consecutive times. The task therefore sheds the environmental baggage and distills the question of open quantum system complexity solely to a single controllable system. The presence of complex temporal correlations directly translates to the complexity of this task. We formalize this approach using a computational model known as the process-tensor framework, a complete mathematical description for quantum stochastic processes [40].

Our work is structured as follows. We begin in Sec. II by describing the process-tensor framework in detail and arguing why it is an ideal computational model for defining the task of sampling an open quantum system over time. Using our model, we show that dynamic multitime sampling of an open quantum system has an equivalent representation as a sampling problem on a many-body state, formalizing a relationship between the complexity of states and the complexity of processes. We then argue that the complexity of the underlying quantum stochastic process is comparable to the complexity of the associated family of master equations. In Sec. III, we define the complexity class *open dynamic quantum polynomial time* (OpenDQP), which is the class of problems solvable within our model, and use its relationship to other complexity classes to classify how complex an open quantum system can be. This complexity class is then examined under several sampling models, including the Heisenberg interaction in a spin chain, IQP, and random processes, where we identify instances of OpenDQP that exist in

$BQP \setminus BPP$ and are thus hard to classically sample from. The backslash denotes the relative complement of BPP with respect to OpenDQP . Ultimately, our results clarify possible avenues to quantum advantage in simulating open quantum systems.

II. COMPUTATIONAL MODEL FOR OPEN QUANTUM DYNAMICS

We begin by outlining the type of problems one usually faces in the area of open quantum dynamics. Most open dynamics problems either involve quantum master equations or stochastic maps (also known as quantum channels or completely positive maps). Master equations are ubiquitous in chemistry, condensed-matter physics, materials science, quantum optics, and statistical physics [41–43]. It turns out that, from a complexity-theory perspective, the complexity of master equations is the same as the complexity of stochastic maps (we elaborate on this below). However, in general, it is difficult to grapple with the computational complexity of these objects. Thus, our approach is to focus on the computational complexity of the underlying process as described by high-order maps known as process tensors [40,44,45]. The latter constitute the most general object in the theory of quantum stochastic processes and thus contain both the master equations and stochastic maps as limiting cases. The process tensor, in its own right, is also of significant interest for understanding and controlling noise in modern quantum devices [46–49].

We start here by stating the most general master equation first, the *Nakajima-Zwanzig master equation*:

$$\partial_t \rho_t = \underbrace{-i[H_t, \rho_t]}_{\text{local driving}} + \underbrace{\int_0^t ds \mathcal{K}_s[\rho_s]}_{\text{memory kernel}} + \underbrace{\mathcal{J}_0}_{\text{i.c.}}. \quad (1)$$

Above, the complexity of the process is entirely in the memory kernel, which accounts for how the information from the past affects the state of the system on aggregate. The local driving term, however, can be responsible for heightening the complexity. In fact, the master equation for different drivings can be drastically different. This master equation also accounts for stochastic interventions from the past; these effects are included in the initial correlation (i.c.) term, as well as the memory kernel.

The dynamics of any open quantum system can be expressed in the form of Eq. (1) via the Nakajima-Zwanzig projection superoperator technique [17]. One can then recover any simplified master equation, including the most general Markov master equation—the Lindblad equation—through an appropriate choice of projection superoperators or by taking approximations of the various terms [13]. For example, expanding the memory kernel

or initial correlation term to second order in the system-environment coupling results in a master equation valid in the weak-coupling regime, which is well suited for the study of many chemical systems [50].

While there is some progress in understanding the computational complexity of the Nakajima-Zwanzig master equation (NZME), it is largely in terms of the dilated complexity. However, there are recent works that bound the dimension of the environment from below by borrowing on similar methods from classical probability theory [29]. We note here in passing that all of our results focus on the dynamics of small systems; it is well known that the open dynamics of many-body quantum systems can encapsulate all of quantum computing [51].

Here, we consider the computational complexity of the underlying quantum stochastic process to determine whether a master equation or a stochastic map is classically tractable. We first introduce the process-tensor framework, which describes quantum stochastic processes in full generality, and then we argue that the sampling complexity of the process is indicative of the complexity of the corresponding family of master equations.

The task of approximately sampling the output probability distribution of a quantum state has been considered a good representation for what it means for a classical computer to simulate a quantum system—capturing the exact obtainable behavior of a quantum experiment. This task is the foundation of the conventional *quantum sampling problem* and has been shown to be classically hard for a number of quantum systems [32,33,52]. To reach our goal, we need to account for the interaction of a system with an environment that is often out of the control of the experimenter, as well as capture interesting dynamical properties of the process, such as temporal entanglement and causal relations. Just as sampling an entire many-body state allows us to infer all k -body spatial relations, we must sample our system across many times to capture all multi-time correlations arising from strong interaction of a system with a complex environment [40,44–46,53,54]. Our first task is, therefore, to generalize state-based sampling complexity to a temporal version capable of capturing the rich multitime physics of open quantum systems. This is precisely what the process tensor enables—namely, it has an equivalent representation as a quantum state. From this, we show that quantum sampling problems are not only a motivating tool for our newly defined temporal sampling problem but can be formally placed on an equal footing.

A. The process-tensor formalism

Consider the setup depicted in Fig. 1(a). An initial system coupled to an environment evolves freely as described by unitary dynamics, $U_{j:j-1}$ between times t_{j-1} and t_j , where $U_{j:j-1}(\cdot) = u_{j:j-1}(\cdot) u_{j:j-1}^\dagger$. Here, we consider the

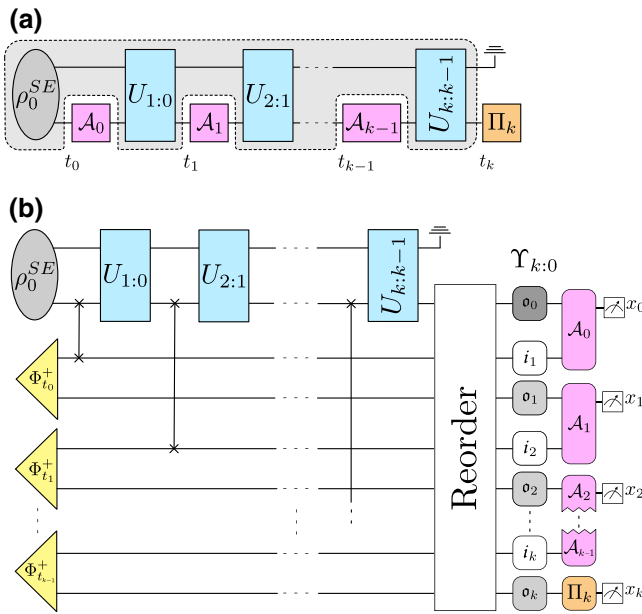


FIG. 1. The process-tensor description of open quantum systems. (a) A circuit-form depiction of the process tensor. The unitaries $U_{j:j-1}$ describe arbitrary SE dynamics. Control operations \mathcal{A}_j are implemented at times t_j . Measurement outcomes at each time generate a probability distribution conditioned on the set of controlled operations. (b) A k -step process-tensor Choi state. One half of a fresh Bell pair $\Phi_{t_j}^+$ is swapped in at each time t_j . Temporal correlations are then mapped to spatial correlations in the Choi state. Projecting the Choi state onto the Choi state of the controlled operations corresponds to sampling from the many-body state.

general case of a d_s -dimensional system interacting with a d_e -dimensional environment. At successive times, $t_0 < t_1 < \dots < t_k$, the experimenter can probe the system with instruments to gain information about its properties. Each manipulation of the system at time t_j corresponds to a trace nonincreasing completely positive (CP) map, \mathcal{A}_j , representing any manner in which the experimenter chooses to probe the system. For example, the experimenter could choose to measure the system according to the positive operator-valued measure (POVM) M_{x_j} and then, on observing outcome x_j , feed forward the quantum state ρ_{x_j} . Here, the map describing the experimental manipulation of the system is given by $\mathcal{A}_j^{(x_j)}[\rho] = \text{Tr}(M_{x_j}\rho)\rho_{x_j}$.

More generally, the probability of observing outcomes $\mathbf{x}_{k:0} := \{x_k, x_{k-1}, \dots, x_0\}$ following the application of instruments $\mathbf{J}_{k:0} := \{\mathcal{J}_k, \mathcal{J}_{k-1}, \dots, \mathcal{J}_0\}$ represented by the set of CP maps $\mathbf{A}_{k-1:0} := \{\mathcal{A}_{k-1}, \dots, \mathcal{A}_0\}$ and measurement apparatus $\{\Pi_k^{x_k}\}$ at times $\mathbf{T}_{k:0} := \{t_k, t_{k-1}, \dots, t_0\}$ is given by

$$\begin{aligned} \mathbb{P}(\mathbf{x}_{k:0}|\mathbf{J}_{k:0}) &= \text{Tr}\left(\Pi_k^{(x_k)} U_k \cdots U_1 \mathcal{A}_0^{(x_0)}[\rho_0^{SE}]\right) \\ &:= \text{Tr}\left(\mathcal{T}_k\left[\Pi_k^{(x_k)}, \dots, \mathcal{A}_0^{(x_0)}\right]\right), \end{aligned} \quad (2)$$

where we shorten our notation to $U_{j:j-1} := U_j$ and introduce the multilinear functional \mathcal{T}_k , known as a process tensor. Note that the instruments here only act on the system space, whereas each U_j acts on both the system and the environment.

Equation (2) forms the basis of our definition for simulation complexity. Specifically, these constitute the probability distributions from which we sample. A process tensor, \mathcal{T}_k , represents the uncontrollable underlying evolution of an open quantum system and its environment and is depicted by the light-gray box in Fig. 1(a). The process-tensor formalism is the natural generalization of the theory classical stochastic processes [54]. In other words, it contains all spatiotemporal correlations that the process can exhibit, be they classical or quantum. This has led to a variety of investigations into the nature of temporal correlations, including the following examples: When does a quantum process contain only classical correlations [55,56]? When do the correlations in a quantum process vanish [57–60]? When are the correlations entangling [45,61–63]? In this work, we ask a complementary question: When are the correlations in a quantum process classically hard?

To do so, we consider sampling from a quantum process with operations \mathcal{A}_j . These allow us to infer all statistical information about the underlying process, as well as compute quantities such as local observables of the system at each time, t_j . Sampling from the distribution in Eq. (2), therefore, gives us all the properties of an open quantum system that we could hope to gain with an experiment, making it the ideal candidate for representing the simulation of these systems. We refer to this task as a dynamic sampling problem. Note that characterization of distributions of this type has already been shown to be possible on quantum devices; our work has immediate implications for real-world analysis of open dynamics [47–49].

Given a description of the initial SE state, the set of unitaries $\mathbf{U}_{k:1} := \{U_k, \dots, U_1\}$, and instruments $\mathbf{J}_{k:0}$, we define the complexity of simulating an open quantum system as the classical complexity of sampling from the resulting temporal output distribution, $\mathcal{D}_{\mathbf{J}}^k$, given by Eq. (2). This output distribution is conditioned on both the number of sampling times, $k+1$, and the choice of instruments, \mathbf{J}_k , which form the parameters of our complexity analysis. Consequently, the complexity of simulating an open quantum system may be classified as easy in some regimes and hard in others. In fact, as we analyze different open quantum systems in Secs. IV and V, we identify simulation tasks that depend on the number of sampling times and/or the choice of instruments. The latter is of particular practical importance, since during any simulation of a process on real hardware, it is impossible to separate the operations required for implementing the SE dynamics from the operations required to measure the system and gain useful information about the process.

B. The equivalence of multitime and many-body sampling

Through a generalization of the Choi-Jamiołkowski isomorphism (CJI), the process-tensor mapping \mathcal{T}_k is in one-to-one correspondence with a many-body state $\Upsilon_{k:0}$. Consequently, every dynamic sampling problem has an equivalent representation as a state-based sampling problem. The state representation of the process may be realized via the circuit shown in Fig. 1(b). At each t_j , one half of a fresh maximally entangled pair is passed through the corresponding U_j . Temporal correlations are then mapped onto spatial correlations between each of the subsystems. The resulting Choi state $\Upsilon_{k:0}$ constitutes a $(2k+1)$ -partite state shown by the light-gray box in Fig. 1(b). Similarly, each of the CP maps \mathcal{A}_j correspond via the CJI to a matrix $\hat{\mathcal{A}}_j$. Equation (2) can then be rewritten as [61]

$$\mathbb{P}(x_k, \dots, x_0 | \mathcal{J}_k, \dots, \mathcal{J}_0) = \text{Tr} \left[\Upsilon_{k:0} \left(\Pi_k \otimes \hat{\mathbf{A}}_{k-1:0}^T \right) \right], \quad (3)$$

where $\hat{\mathbf{A}}_{k-1:0} = \bigotimes_{j=0}^{k-1} \hat{\mathcal{A}}_j^{(x_j)}$.

Equation (3) represents the spatiotemporal generalization of Born's rule [53,64] and, consequently, demonstrates that operations on a system at successive points in time correspond to observables on the process-tensor Choi state. The circuit in Fig. 1(b) is, therefore, a sampling problem on a $2k+1$ many-body state that is equivalent to a temporal sampling problem on a k -step process.

C. Master equations, stochastic maps, and process tensor

We are now in position to reexamine the master equations and stochastic maps, with an understanding of the basic structural elements of the process-tensor framework. First, it is clear that given a k -step process-tensor Choi state $\Upsilon_{k:0}$, we can obtain the Choi state of all stochastic maps $\{\hat{\mathcal{E}}_{j:i}\}$ with $0 \leq i \leq j \leq k$ (where the Choi-state representation is indicated by a hat). This is trivially obtained by choosing all \mathcal{J} to be the identity instrument except at times i and j , such that

$$\hat{\mathcal{E}}_{j:i} = \text{Tr}_{\bar{i}_{i+1}, \bar{o}_j} \left[\Upsilon_{k:0} \left(\bigotimes_{i < m < j} \Phi_m^+ \right) \right], \quad (4)$$

where $\hat{\mathcal{E}}$ is the Choi state of the stochastic map, $\Phi^+ = |\Phi^+\rangle\langle\Phi^+|$ where $|\Phi^+\rangle := (|00\rangle + |11\rangle)/\sqrt{2}$ is the Choi state corresponding to the identity instrument, and \bar{i}_{i+1}, \bar{o}_j indicates the trace over all process-tensor legs except i_{i+1} and o_j . Note that there is an implicit identity on each of the traced process-tensor legs. Next, it has recently been shown that this same family of stochastic maps serves as a

discrete version of the Nakajima-Zwanzig master equation by means of the transfer-tensor method [24,65–67]. The precise details of this derivation can be found in Secs. 3 and 4 of Ref. [65]; however, we include a summary of the main ideas here so that the reader can follow the subsequent arguments.

The transfer-tensor method encodes correlations between two discrete times t_j and t_i into a tensor $T_{t_j:t_i}$. Using this, the evolution over the time interval $[0, k]$ can be rewritten as propagation in a multiplicative fashion via the set of transfer tensors over some choice of intermediate times $0 \leq i \leq j \leq k$. Without loss of generality, we can assume that the total time interval is divided into equally spaced discrete times with the time resolution given by $\delta t = t_j - t_{j-1}$. Assuming periodicity of the dynamical maps, the tensor formalism can then be used to propagate the system to later times in an efficient manner. Specifically, the transfer tensors are defined in terms of the operationally determined set of stochastic maps $\{\mathcal{E}_{j:i}\}_{[0,k]}$ as

$$\begin{aligned} T_{t_k:t_{k-1}}^{(1)} &= \mathcal{E}_{t_k:t_{k-1}} \quad \text{and} \\ T_{t_k:t_j}^{(k-j)} &= \mathcal{E}_{t_k:t_j} - \sum_{m=j+1}^k T_{t_k:t_m}^{(k-m)} \mathcal{E}_{t_m:t_j}. \end{aligned} \quad (5)$$

The evolution of a system ρ from t_0 to t_k is then given by

$$\rho_{t_k} = \sum_{j=0}^{k-1} T_{t_k:t_j}^{(k-j)} \rho_{t_j} + \Xi_{t_k:t_0}, \quad (6)$$

where $\Xi_{t_k:t_0}$ is the initial correlation term.

If the stochastic maps are periodic with period $T = c\delta t$ for some $c \in \mathbb{Z}$, such that $\mathcal{E}_{t_k+T:t_j+T} = \mathcal{E}_{t_k:t_j}$, then the transfer tensors will also be periodic, such that

$$T_{t_k:t_j}^{(l)} = T_{t_{k+c}:t_{j+c}}^{(l)}. \quad (7)$$

Then, evolution of the state to any arbitrary time t_m is given by Eq. (6) by associating $T_{t_m:t_{m-1}}^{(l)}$ with $T_{t_{m-1} \bmod T + l\delta t, t_{m-1} \bmod T}^{(l)}$.

Finally, by substituting Eq. (6) into the difference between t_k and t_{k-1} , we recover a discrete version of the NZME. In Ref. [65], this is shown to be identical to the master equation if we take the sampling resolution to be infinitely small. Thus, the process tensor contains both stochastic maps and master equations as limiting cases. One only has to make sure that the time resolution δt is fine enough to approximate the time derivative in Eq. (1) well.

Importantly, Eq. (4) illustrates that if we consider a process tensor as a many-body state, then a master equation is sampling from that state by contracting it with the Choi

states of instruments at all times except i and j . In Eq. (4), the Choi state Φ^+ corresponds to the identity instrument resulting in a nondriven master equation. For local driving, the Choi state of the instrument is a maximally entangled state of the form

$$|\psi^+(H_i)\rangle := (\exp\{-iH_i\} \otimes \mathbb{1}) \sum_i |ii\rangle, \quad (8)$$

where H_i is a time-dependent Hamiltonian, meaning that the instrument will change in time. Thus, construction of any master equation can be represented as sampling from the projected Choi state of the associated process tensor to the spaces of i and j for a set of times $0 \leq i \leq j \leq k$. If this sampling task is hard, then so is simulation of the master equation.

Thus, going forward, we only consider the computational complexity of process tensors. This is because the complexity of a process tensor represents the complexity of the underlying quantum process. If this is hard, then—we posit—so will be the family of master equations that are contained therein. Our strategy is to show that, for a given process, there are sampling regimes that are classically hard. We do this for Shor’s algorithm, the Heisenberg Hamiltonian, and several well-known models of quantum computation. We also then present numerical evidence for the hardness of process tensors generated by random circuits and random Hamiltonians.

III. THE COMPLEXITY CLASS **OpenDQP**

With a clear definition for the task of simulating open quantum systems, we now turn to classifying their complexity. We are particularly interested in studying the boundary between when this task is classically easy versus hard, since we can consider this the point at which a system displays truly complex or quantum features. This corresponds to the boundary between the complexity classes **BPP** and **BQP**, the class of problems solvable by a probabilistic classical and quantum computer, respectively. To formally relate examples of open quantum systems to these classes, we define the complexity class *open dynamic quantum polynomial time* (**OpenDQP**), which encapsulates the set of problems solvable in our model.

The dynamic sampling problem that we describe bears a close resemblance to the complexity class *dynamic quantum polynomial time* (**DQP**), defined by Aaronson in Ref. [68]. **DQP** is the class of problems solvable by a **BQP** machine when given access to an oracle that can return a sample over the probability distribution of the classical histories of a quantum circuit. Using this model, Aaronson shows that **DQP** is slightly more powerful than **BQP**. However, the ability to sample the histories of a circuit without actually implementing a measurement is, of course, unphysical. Motivated by the study of physically

implementing a dynamic sampling problem on a quantum computer, we define the analogous complexity class **OpenDQP** as follows.

Definition 1: OpenDQP. The class of sampling problems $S = \{\mathcal{D}(\mathbf{x})_{\mathbf{J}}^k\}_{x \in \{0,1\}^k}$ for which there exists a quantum polynomial-time algorithm \mathcal{Q} , such that given the initial state ρ_0^{SE} of size n , unitaries $\mathbf{U}_{k:0}$, and instruments $\mathbf{J}_{k:0}$ as input, outputs a sample $\mathbf{x} = \{x_k, \dots, x_0\}$ from distribution $\mathcal{R}(\mathbf{x})$, in time $\text{poly}(n, k)$, such that

$$\|\mathcal{R}(\mathbf{x}) - \mathcal{D}(\mathbf{x})_{\mathbf{J}}^k\|_1 \leq \epsilon \quad \text{for} \quad \epsilon = O\left(\frac{1}{\text{poly}(k)}\right). \quad (9)$$

It is important to note that there must exist an efficient classical description of the initial state, unitaries U_j , and instruments \mathcal{A}_j ; otherwise there could exist an environment that could compute problems outside of **BQP** and return the answer to the system. We therefore require unitaries that decompose into a polynomial number of one- and two-qubit gates. Similarly, we restrict our attention to instruments of the same description. A process is then efficiently classically simulable if there exists a polynomial-time classical algorithm that outputs a sample from distribution $\mathcal{R}(\mathbf{x})$ —i.e., one that can sample from the output distribution of the process up to additive error. This is consistent with the notion of “weak” classical simulation, which achieves polynomial accuracy in polynomial time (for further discussion, see Ref. [69], particularly Sec. 2.3). It follows that our sampling problem can be considered *easy(hard)* if it can(not) be efficiently classically simulated.

It is worth reiterating here that when sampling an open quantum system over time, the resulting probability distribution in Eq. (3) depends on the choice of instruments \mathbf{J}_k and the number of sampling times k . In fact, the instruments themselves may be every bit as complex as a process and have the potential to significantly change the complexity of the resulting sampling task. This is not surprising since, even in the classical case, a measurement apparatus that depends on all previous measurement results would be exponentially hard to simulate. Here, the sequence of control operations may have both classical and quantum correlations, forming *testers* or process-tensor duals, as described in detail in Appendix A. As we show, for a fixed underlying process, these variables have the power to elevate a sampling problem from one in **BPP** to one that is in **BQP** \setminus **BPP**. In what follows, we refer to a process as hard (or complex) if there exist a set of dependencies \mathbf{J}_k or a number of sampling times such that it is classically hard to simulate.

When considering the complexity of master equations, it is useful to define the complexity class **OpenDQP** in terms of the total number of sampling times k . We denote this **OpenDQP** $_k$. Nondriven master equations and those with a constant local field then form the subset **OpenDQP** $_2 \subseteq$

OpenDQP_k. These correspond to a k -step process tensor with identity instruments or identical fixed operations inserted at all times except i and j in order to construct the stochastic map $\mathcal{E}_{j;i}$. We include constant fields in this description, since these time-independent operations could be incorporated into the underlying system-environment interaction. By comparison, a driven master equation will have a set of time-dependent control operations \mathcal{A}_j , and thus exists within **OpenDQP_k**. These sets then form the family of master equations associated with an underlying quantum stochastic process.

IV. COMPLEXITY OF PROCESSES

Since the task of classically simulating any open quantum system can be framed as a problem within **OpenDQP**, the computational power of this complexity class gives an indication of how hard it is to simulate these systems. In this section, we begin by establishing the relationship of **OpenDQP** to existing complexity classes to broadly classify its computational power. Subsequently, we consider a number of specific open quantum systems, in order to identify instances of **OpenDQP** that are classically hard to simulate.

A. **BQP** \subseteq **OpenDQP** = **SampBQP**

It immediately follows from our definition that **BQP** \subseteq **OpenDQP** and thus **OpenDQP** is at least as powerful as standard quantum computation. **BQP** is the class of decision problems solvable by measuring a single qubit with a bounded error rate [31]. Thus, any problem in **BQP** can be represented as a sampling problem in **OpenDQP** restricted to a single time. If we prepare an initial state $|0\rangle^{\otimes n}$, apply a polynomial-size quantum circuit $U_{1;0}$, and then measure in the computational basis on the system at time t_1 , we recover **BQP**.

We then observe that **SampBQP** = **OpenDQP**; that is, we can represent any conventional state-based sampling problem as a dynamic sampling problem on a multitime process and vice versa. Equivalently, this result states that quantum processes and quantum states can be equally complex. We show that **OpenDQP** \subseteq **SampBQP** in Sec. III and thus states are at least as complex as processes. Sampling from any multitime process can be represented as sampling from a many-body state via the CJI, where operations at different times correspond to observables on the Choi state. To show the reverse inclusion—that **SampBQP** \subseteq **OpenDQP**—we show that the output distribution of any state can be represented as the output distribution of a multitime process.

The logic follows the circuit diagram in Fig. 2. Consider an arbitrary quantum sampling problem on a state. This is computed by taking an initial SE n -qubit state $|0\rangle^{\otimes n}$ and evolving it via a circuit that corresponds to unitary U . At this point in time, the **SampBQP** problem involves




FIG. 2. A process with a temporal probability distribution equivalent to the spatial distribution of the initial state ψ_0 . A measurement at time t_j constitutes a measurement of an environment qubit (here indexed as q_{j-1}) and results in outcome x_j . A new state $|0\rangle|0\rangle$ is fed into the subsequent input leg of the process.

sampling from the distribution of the output state given by $p(x) = |\langle x|U|0\rangle^{\otimes n}|^2$, where x is a length- n bit string. We label the resulting state ψ_0 and consider it the initial system-environment state now entering our process. We then measure the system, resulting in outcome x_0 . Next, we insert a SWAP gate acting between the first qubit of the environment, q_0 , and the system, followed by a measurement and result x_1 . This process is repeated for each of the $n - 1$ qubits in the environment, resulting in a n -step process tensor. Following the circuit lines, we can see that the probability of the length- n bit string obtained from sampling the system, $(x_0, x_1, x_2, \dots, x_{n-1})$, is precisely that of the state ψ_0 .

Finally, we note that state complexity is necessary for process complexity. By this, we specifically mean that the system-environment state must be hard to sample from at least a single time during the dynamics. This hardness may be conditioned on a previous operation on the system. If this was not true, then we could directly simulate the SE dynamics and compute the reduced output on the system.

B. Illustrative example: Shor’s algorithm \subseteq **OpenDQP**

Having established that **OpenDQP** is at least as powerful as quantum computation, we are particularly interested in finding problems that exist within **OpenDQP** but that are outside **BPP**. This subset of problems represents quantum stochastic processes that are hard to classically simulate and their existence provides evidence for quantum advantage in simulating open quantum systems. Shor’s algorithm for factoring integers is in the class **BQP** and provides an exponential speed-up over the best-known classical algorithms. As a result, factoring is widely believed to be outside **BPP**, although there is not yet strong complexity-theoretic evidence for this [32].

Here, we show that Shor’s algorithm can, in fact, be cast as a sampling problem on a quantum stochastic process and provides a natural introduction to the conceptual idea of mapping state complexity onto process complexity through the process-tensor framework. This, in turn,

shows that there are quantum stochastic processes that are classically hard, provided that factoring is classically hard. In fact, Kitaev’s one-controlling-qubit trick for factoring an integer already serves as a preexisting example of the conversion from state sampling to a dynamic sampling problem [70–72].

The order-finding circuit for factoring an n -bit integer, N , is shown in Fig. 3(a). Here, the upper register of $2n$ qubits (where $n = \log N$) is a quantum superposition of integers $0, \dots, N^2 - 1$, followed by a modular exponentiation circuit of $2n$ controlled modular multiplications $U_a^{2^j}$, which collectively transform $|x\rangle$ to $|(ax) \bmod N\rangle$. The inverse quantum Fourier transform (QFT) circuit is then applied to the upper register, followed by a measurement

and classical processing that gives the period of the function $f(x) = a^x \bmod N$.

It has been shown that the factoring algorithm can be implemented in a semiclassical manner on a single qubit, plus an ancilla register [71,73]. We summarize the procedure briefly here, as it is important for understanding Shor’s algorithm as a process tensor. The key is to condense the entire $|x\rangle$ register into a single qubit, reset and reused. At each t_j , the system qubit acts as the x_j bit of the register, placed into a $|+\rangle$ state and controlling the unitary $U_a^{2^j}$ on the ancilla space. Collectively, this computes $a^x \bmod N$. The inverse QFT is then performed semiclassically: the control qubit is measured in the Fourier basis and reprepared into a $|+\rangle$ state. Measurement in the Fourier basis involves a series of Z rotations conditioned on outcomes at the previous times, identical to the controlled- Z rotations in the standard inverse QFT circuit. This procedure is summarized in Fig. 3(b). Importantly, the specific choice of instruments and the resulting bit string can be used to efficiently determine the prime factors of a number. This leads us to our first result.

Result 1: Shor’s algorithm to factor an n -bit number can be written as a sampling problem from an open quantum system. This quantum stochastic process is represented by a $2n$ -step single-qubit process tensor.

To see this, we first note that the order-finding circuit of Fig. 3(a) can be represented as a $k = 2n$ step process in its generalized Choi-state form. The n -qubit register is now the environment and the $2n$ qubits of the previous upper register each form one half of a maximally entangled Bell pair $|\Phi_{ij}^+\rangle = 1/\sqrt{2}(|00\rangle + |11\rangle)$, which is fed into the modular exponentiation circuit at time t_j . The resulting $(2k + 1)$ -body state is the process-tensor Choi state $\Upsilon_{k;0}$ of Shor’s algorithm. The inverse QFT coincides with operations on this Choi state or, equivalently, sampling of the many-body state. The mathematical construction of the state $\Upsilon_{k;0}$ is detailed in Appendix C 1.

We can then utilize Kitaev’s proposal to map this state-sampling problem to a dynamic sampling problem on a multitime process. This is represented in Fig. 3(b). The process tensor, shown in light gray, corresponds to unitary evolution of a SE state according to the controlled modular multiplication circuits. After each controlled unitary, we can probe our system with the semiclassical operations corresponding to the inverse QFT on a single qubit. At time t_j , if we observe outcome $x_j = \{0, 1\}$, then the map representing this transformation is given by $\mathcal{A}_j^{x_j}[\rho] = \text{Tr}(M_{x_j}(R\rho R^\dagger))|+\rangle\langle+|$, which accounts for a rotation gate conditioned on all previous classical measurement outcomes, a POVM M_{x_j} corresponding to measuring in the X basis and observing outcome x_j , and feed forward of the state $|+\rangle$ for use at the next time step. Shor’s algorithm is, therefore, a problem that is clearly contained in OpenDQP and provides a tangible example that suggests that within

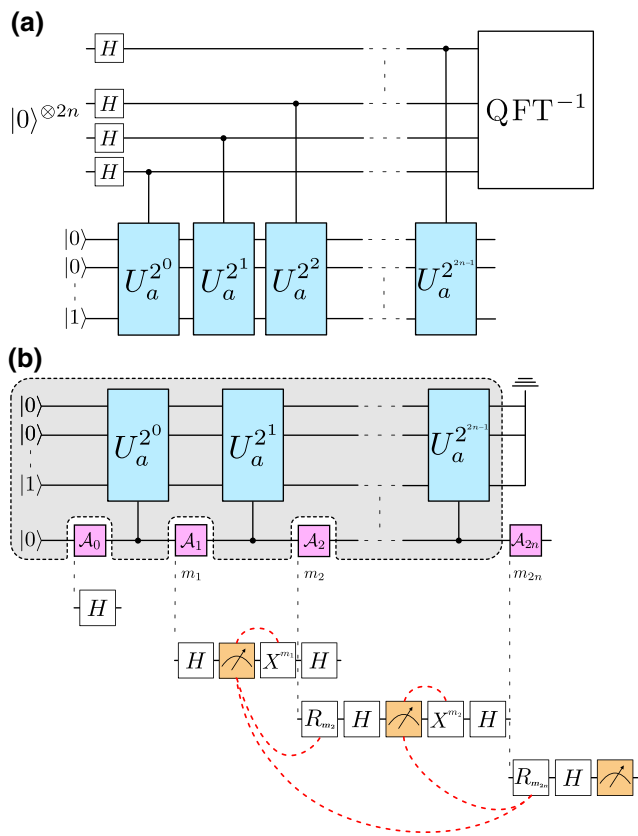


FIG. 3. Shor’s algorithm. (a) The order-finding circuit for Shor’s factoring algorithm of a n -bit integer as described in the text. The unitaries perform modular multiplication and collectively transform $|x\rangle$ to $|(ax) \bmod N\rangle$. The inverse QFT, subsequent measurement, and classical postprocessing give the period of the function $f(x) = a^x \bmod N$. (b) Shor’s algorithm as a process tensor. The $2n$ register is collapsed to a single control or “system” qubit using the one-controlled-qubit trick. The operations \mathcal{A}_j correspond to the gates of the inverse QFT on a single qubit. At each time t_j , a rotation R_{m_j} is performed that depends on all previous measurement results, followed by a measurement and reset X^{m_j} , as described in the text. The classical correlations between these operations are shown by the red dotted line.

this framework there will exist open quantum dynamics that are classically hard to simulate in the worst case.

A nice consequence of framing Shor’s algorithm as a quantum stochastic process is that we identify an instance of **OpenDQP**, where the classical complexity depends on the choice of sampling operations, \mathcal{A}_j . This establishes our second result.

Result 2: The process tensor, from which Shor’s algorithm is derived, is easy to sample from in the computational basis.

This result follows from observing the circuit shown in Fig. 3(b). If instead of performing the series of measurements, rotation gates, and Hadamard, which implements the inverse QFT, we simply measure in the computational basis at each time and feed forward the state $|+\rangle = 1/\sqrt{2}(|0\rangle + |1\rangle)$, then the state of the system following the application of any controlled modular unitary is an equal superposition state. Classical simulation is then given by the outcome of a coin toss at each time. A detailed proof of this is provided in Appendix C 1, where we consider the equivalent problem of sampling in the computational basis on the Choi-state representation of the process-tensor version of Shor’s algorithm.

Ultimately, there exists a process-tensor version of Shor’s algorithm that is easy to sample from in the computational basis but can factor numbers if we sample with the classically correlated control operations described in Fig. 3(b). It is quite remarkable that manipulations on a single qubit alone can elevate the complexity of simulation from a simple classical stochastic process solvable in polynomial time to one solvable only in quantum polynomial time.

C. Physical Hamiltonian: Heisenberg interaction $\subseteq \text{OpenDQP} \setminus \text{BPP}$

Shor’s algorithm provides an introductory example of a sampling problem in **OpenDQP** that is classically hard to simulate given a specific choice of operations. However, we are primarily interested in the simulation of physical systems. Here, we present a physical Hamiltonian for the evolution of a system and environment, which, along with a set of sampling operations \mathcal{A}_j on the system, constitutes a sampling problem in **OpenDQP** that is in **BQP** \setminus **BPP** and is thus classically hard to simulate. Specifically, we show that the Heisenberg interaction in a spin chain coupled with operations on the system results in a temporal probability distribution that is hard to sample from. The general idea is that in this interaction picture, we can use operations on the system to construct a complex system-environment state, and then extract this complexity dynamically.

To see this, we first observe that there exist composite systems that admit algebraic control (AC) and can be framed as problems in the complexity class **OpenDQP**.

The AC approach to quantum control determines the necessary and sufficient conditions such that a composite system can be completely driven by control operations on only a smaller subsystem [74]. In this method, there exists a composite system $V = S \cup E$, with controllable subsystem S , described by the global Hamiltonian $H_V + \sum_j f_j(t) h_S^{(j)} \otimes \mathbb{1}_E$. The time-dependent local controls $h_S^{(j)}$ are achieved through the modulating parameters $f_j(t)$.

It has been shown that Heisenberg-type chains of N spin-1/2 particles with arbitrary coupling strengths admit AC through operations on a single end of the chain [75,76]. This can be represented precisely as a dynamic sampling problem in **OpenDQP**. Here, the combined system and environment form the composite system, with the global Hamiltonian given by the Heisenberg coupling:

$$H_{ij} = J_{ij} (X_i X_j + Y_i Y_j + Z_i Z_j), \quad (10)$$

where the subscripts ij denote neighboring qubits and the J_{ij} are the couplings. The local controls act only on the system and correspond to time-dependent sampling operations. The results of AC mean that in the Heisenberg-interaction picture, operations on the system can generate the entire Lie algebra $SU(2^N)$ on the system-environment state and are thus universal for quantum computing. The net outcome is that we can generate any quantum transformation on the spin chain through operations on the system alone. In particular, we can generate a system-environment state which is classically hard to sample from.

We now want to show that this composite system under AC permits a dynamic sampling problem in **OpenDQP** that is classically hard, thus demonstrating a physical Hamiltonian of an open quantum system that is hard to simulate. For this problem, it is important that we can efficiently generate complex states using the AC method. That is, if a state can be prepared in polynomial time by a quantum computer through operations on the entirety of the system, then achieving the same control through operations on the subsystem must also scale polynomially. Unfortunately, there is no general solution that determines whether this quantum control can be achieved efficiently. For the general Heisenberg interaction, solving for the controllable operations in itself is inefficient because the Hamiltonian cannot be diagonalized straightforwardly. Thus, commenting on the efficiency of the resulting operations is a highly nontrivial problem. However, progress has been made in presenting scalable implementation with control on two qubits and a tunable magnetic field on the environment [77]. Moreover, local operator control on only a single qubit in a Heisenberg spin chain has been demonstrated to high fidelity on a chain of length 3 [78]. Thus, although implementing any unitary will take an exponentially long time, we instead only require the implementation of a unitary that takes our state to one outside of **BPP**. We anticipate that there will exist complex states that can be

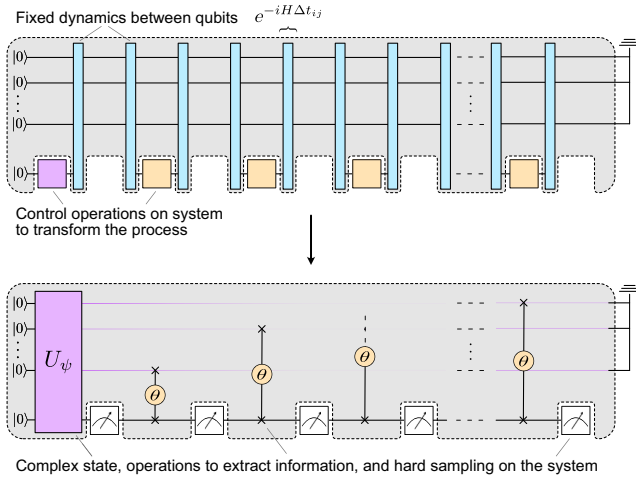


FIG. 4. A complex dynamic sampling problem generated by the Heisenberg interaction on a spin chain. A single-qubit system interacts with a spin chain with fixed Heisenberg coupling, H . The upper figure shows a sequence of control operations on the system in pink and yellow, which can perform arbitrary transformations of the composite system-environment state. In the lower figure, these operations result in evolution of the SE state given by U_ψ , which results in the spatially complex state ψ . This is followed by a series of pSWAP gates between the system and environment. The pSWAP gates transfer information from the complex environment to the system and the corresponding measurements on the system constitute a classically hard sampling problem.

efficiently generated by operations on the system alone and assume this to be true in what follows.

We construct our sampling problem as per the circuit structure shown in Fig. 4. We begin with an initial system-environment state with fixed nearest-neighbor Heisenberg coupling. Using the algebraic controllability of this composite system environment, we can then implement a series of operations \mathcal{A}_j on the system at times t_j until we have transformed the system-environment state into a complex state that is classically hard to sample from, i.e., one that is in $\text{SampBQP} \setminus \text{BPP}$. We label this state ψ .

Using the same Heisenberg interaction, we then leverage the spatial complexity of the state ψ and redistribute the amplitudes of the environment onto the temporally sampled populations. Specifically, we use our control operations to engineer a series of partial-SWAP (pSWAP) gates between each environment qubit and the system, defined as

$$\text{pSWAP}(\theta)_{ij} := \exp \left\{ -i\frac{\theta}{2}(X_i X_j + Y_i Y_j + Z_i Z_j) \right\}, \quad (11)$$

where the subscript ij denotes qubits i and j . These gates pass information directly from the environment to the system with probability depending on θ . We consider a total of k successive pSWAP gates between each environment qubit and the system, with sampling in the computational basis following each gate. This is repeated for all $n - 1$ qubits of the environment, resulting in $(n - 1)k$ sampling times.

The series of pSWAP gates between the environment and system results in a distribution on the system isomorphic to that of the constructed system-environment state ψ and thus sampling from the system over time is as hard as sampling from this state. For simplicity, we first show that this is true for an environment of only one qubit. This can then be repeated iteratively for each qubit that makes up the environment, since we only interact with a single qubit at a time in our procedure. We consider the Choi-state construction of a process with environment in some initial state $|\psi\rangle_E = \alpha|0\rangle + \beta|1\rangle$, where instead of feeding one half of a Bell pair in at each time, we feed in the $|0\rangle$ state. This is equivalent to the construction using Bell pairs and it can be realized with the process-tensor Choi state by projecting all input legs onto $|0\rangle\langle 0|$. At time t_1 , we feed in the first ancilla $a_1 = |0\rangle_1$ and perform a pSWAP_{E,a_1} . Following this, the environment-ancilla state is $|\psi\rangle_{Ea_1} = \alpha|00\rangle - ie^{i\theta}\beta \sin\theta|01\rangle + |10\rangle$ and the probability that the environment is in the zero state is $1 - |\beta|^2 \cos^2\theta$. At time t_2 , the next ancilla $a_2 = |0\rangle_2$ is fed into the process, followed by another pSWAP. The feeding in of ancillas is equivalent to measuring the system and freshly resetting it to the $|0\rangle$ state. This is repeated k times, at which point the probability that the environment is in the zero state is given by $\Pr[|\psi\rangle_E = |0\rangle] = 1 - |\beta|^2 \cos^{2k}\theta$. For a given choice of θ , if $k = O(\log \epsilon)$, where $\epsilon > 0$, then the probability that the environment is in the zero state is

$$\Pr[|\psi\rangle_E = |0\rangle] > 1 - \epsilon. \quad (12)$$

Ultimately, the environment is in the $|0\rangle$ state with probability 1 up to additive error and all operations performed are unitary; thus the resulting distribution on the process-tensor Choi state is isomorphic to the system-environment state ψ .

As a result of this construction, we identify an instance of OpenDQP that is equivalently hard to the sampling problem on the state ψ . Thus, for any spatially complex state, this dynamic sampling problem is classically hard to simulate. This simple circuit is an illustrative example of the types of interactions that generate complex dynamics. Here, information from a complex environment is passed sequentially to the system, which is then measured. This behavior resembles the physical scenario where complex memory effects in the environment can be temporally redistributed, affecting the future dynamics of the system. The temporal complexity here reflects how strongly the environmental complexity is imprinted upon the system.

D. Clifford processes \subseteq OpenDQP \setminus BPP

We can build on this example and decompose the Heisenberg couplings and pSWAP gates into Cliffords and non-Cliffords, which is shown for a single pSWAP in Fig. 5(a). Treating the Cliffords as the background process here, the non-Cliffords can be designated as user-chosen

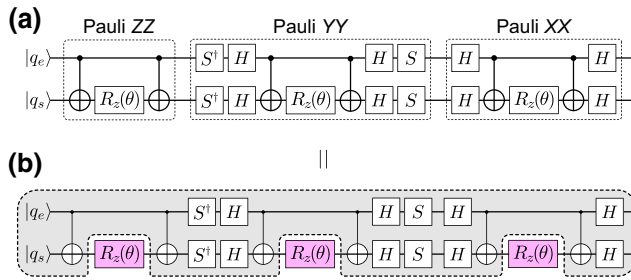


FIG. 5. (a) The decomposition of a pSWAP gate into Clifford and R_z gates. (b) The decomposition results in a process composed of Clifford gates, with sampling operations $R_z(\theta)$ acting on the system. For $\theta = \pi/4$, the sampling operations are T gates.

instruments. As a control problem, this identifies a case in which the process-tensor Choi state can be exactly constructed but from which sampling is hard. This establishes a sampling problem that is in **OpenDQP** but outside of **BPP**. In other words, the degree of complexity of the sampling problem depends on the choice of operations and can be controlled at the system level.

First, we consider the pSWAP gates. A single pSWAP acting between two qubits q_i and q_j can be decomposed as follows. Since XX , YY , and ZZ all commute, we can write the pSWAP as

$$\exp(-i\frac{\theta}{2}X_iX_j) \exp(-i\frac{\theta}{2}Y_iY_j) \exp(-i\frac{\theta}{2}Z_iZ_j). \quad (13)$$

Additionally, the XX and YY terms are locally equivalent to ZZ . A single pSWAP gate can then be represented as shown in Fig. 5(a). We have now decomposed each pSWAP gate into a series of Clifford gates acting on the environment and system, plus three T gates acting on the system only. These can be separated into the underlying SE interaction and sampling operations \mathcal{A}_j , as shown in Fig. 5(b) by the light-gray box and pink R_z gates, respectively. This procedure can be extended to all pSWAP gates and similarly to the Heisenberg interaction of Eq. (10). The overall result is a process tensor composed entirely of Clifford gates and a sequence of sampling operations \mathcal{A}_j made up of AC operations, T gates, and computational-basis measurements.

We identify the underlying process as an instance of **OpenDQP**, where the SE unitaries are Clifford circuits. By the Gottesman-Knill theorem [79], if at each time we choose to measure our system in the computational basis, then sampling from the output of this Clifford process tensor is classically easy. However, as we identify above, sampling with a sequence of AC operations, T gates, and computational-basis measurements is equivalent to sampling from a complex state ψ and thus classically hard. This again identifies an instance of **OpenDQP** where the choice of sampling operations can elevate our problem

from one inside **BPP** to a sampling problem in **BQP** \ **BPP**.

The ability to control complexity at the system level is particularly interesting when considering applications on near-term smaller-scale quantum devices. Currently, we only have access to a small number of qubits, which are too noisy to generate the highly entangled system-environment states required for the simulation of complex dynamics. However, the noise in these devices is partly due to the interaction of each qubit with the surrounding environment, which in itself can generate complex statistics. Being able to control this interaction from a single qubit would allow the experimenter to access this complex noise as a resource, which might aid in the simulation of open quantum systems. In fact, the ability to modify complex temporal correlations has been experimentally verified on IBM devices, where changing control operations on a system qubit interacting with a single-qubit environment allows the experimenter to tune the resulting temporal entanglement [48].

E. **OpenDQP**–(IQP) $\not\subset$ **OpenDQP** \ **BPP**

Having found a member of **OpenDQP** that is outside of **BPP** and having identified how a spatially complex process can be linked to a temporally complex process, we now investigate if all instances of **SampBQP** that demonstrate quantum advantage are progenitors of instances of **OpenDQP** \ **BPP**. By considering the instantaneous quantum polynomial-time (IQP) framework, we show that this is not the case. IQP circuits, introduced and described in detail in Refs. [33,80], are a restricted form of quantum computing comprising two-qubit commuting gates that are diagonal in the X basis. For simplicity, we restrict our attention to an equivalent description of IQP with gates diagonal in the computational basis. The circuit then consists of an initial n -qubit state $|0\rangle^{\otimes n}$, followed by $H^{\otimes n}CH^{\otimes n}$, where C is a quantum circuit comprising gates diagonal in the Z basis. Sampling an IQP circuit amounts to measuring in the computational basis on a designated set of output lines. Bremner, Jozsa, and Shepherd [33] have shown that sampling on all n output lines of an IQP circuit is hard to simulate classically unless the polynomial hierarchy collapses to the third level. However, measuring only $O(\log n)$ qubits is in **P**.

We now turn to our complexity analysis of **OpenDQP**, and consider an instance where the underlying system-environment interaction (the unitaries $U_{j,j-1}$ in Fig. 1) is an IQP circuit on n qubits, described in Fig. 6. Here, we take a single qubit to represent our system and the remaining $n - 1$ qubits constitute the environment. We label this process **OpenDQP**–(IQP). The IQP circuit can be equivalently generated by a Hamiltonian diagonal in the X basis that evolves the input state $|0\rangle^{\otimes n}$ according to [80]

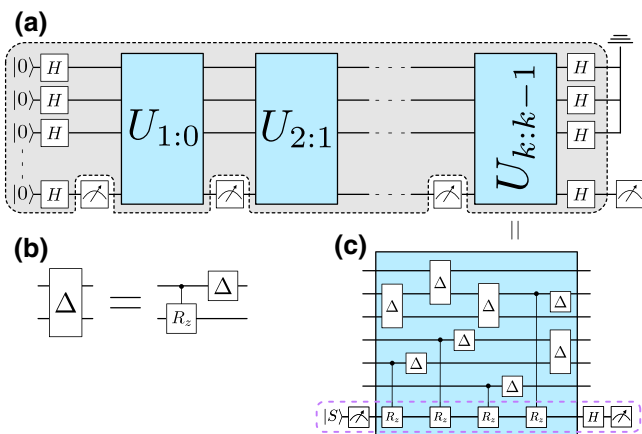


FIG. 6. (a) The circuit construction of OpenDQP-(IQP). (b) The decomposition of a two-qubit Z -diagonal gate (Δ) into a controlled- R_z and single-qubit Z -diagonal gate. (c) Each gate acting between an environment qubit and the system can be decomposed according to the description in (b). The net result is a series of gates acting between environment qubits and a series of R_z gates acting on the system conditioned on environment qubits. The resulting operations acting on the system are outlined with the purple traced line.

$$H = \sum_{\mathbf{p}} \bigotimes_{j:p_j=1} h_j X_j. \quad (14)$$

Here, there are a polynomial number of n bit strings \mathbf{p} that determine which qubits to apply the product of Pauli X operators to in each term of the summation.

In this setup, the results of Bremner, Jozsa, and Shepherd confirm that the dynamics of the entire SE state are hard to simulate classically. That is, the dynamics demonstrate spatial complexity. We now query whether the dynamics demonstrate temporal complexity and, in particular, whether the temporal output distribution generated by the open quantum systems version of IQP is hard to sample from. We know that spatial complexity of the system environment is a necessary condition in order to demonstrate temporal complexity on the system alone. We find that the process version of IQP can be efficiently classically simulated using Monte Carlo methods, resulting in an overall weak classical simulation of the process tensor for any number of times. This leads us to our third result.

Result 3: Temporally complex processes require a spatially complex SE state; spatially complex SE states are not sufficient to guarantee temporally complex processes.

To prove this, we provide an efficient classical algorithm, following the methods in Ref. [33], which outputs a sample from the temporal distribution generated by sampling the process OpenDQP-(IQP). Before proceeding with the algorithm, we provide some intuition for the types of processes that permit a Monte Carlo simulation by considering their representation as a Choi state. First, we

note that Markovian processes with a single-qubit system have a product Choi state and can be efficiently simulated [44,81]. A separable non-Markovian process, then, is a probabilistic mixture of Markovian processes. If the classical distribution governing the mixture can be efficiently sampled from, then it follows that the non-Markovian process can also be efficiently sampled from. This is true for the process OpenDQP-(IQP).

The proof is given here, with the corresponding algorithm in Appendix B. Consider the circuit in Fig. 6(a) on n qubits. The lower qubit represents the system and the upper qubits represent the $(n-1)$ -sized environment. This is a k -step process with $k+1$ possible control operations. The SE unitaries comprise $\text{poly}(n)$ Z diagonal gates, which can be divided into two separate groups of two-qubit diagonal gates: those acting between environment qubits and those acting between the system qubit and an environment qubit. Each of these groups also consists of at most $\text{poly}(n)$ gates. Each gate in the latter group can be efficiently decomposed into a controlled- R_z gate with the phase acting on the system qubit, plus an R_z gate on the environment [82,83], as shown in Fig. 6(b). To compute the output of this process tensor following a measurement in the computational basis at each time, we proceed as follows:

- (1) Compile the Z -diagonal gates acting between the system and an environment qubit into $\text{poly}(n)$ controlled R_z gates and single-qubit diagonal gates, where the R_z unitary acts on the system qubit.
- (2) Consider the total system-environment state at the end of the circuit. Since this is an equal-superposition state, tracing over the environment means that we can stochastically replace it with a bit string of length $n-1$ uniformly at random.
- (3) This length- $(n-1)$ bit string commutes with the controlled R_z system-environment gates at each of the time steps and thus the sample may be brought to the beginning of the circuit. We use this to determine the phase gates applied to the system qubit.
- (4) The evolution of the single-qubit system will then involve $\text{poly}(n)$ single-qubit rotation gates R_z , plus $k+1$ total sampling operations \mathcal{A}_j at each time, which can be classically strongly simulated for each given weak sample of the process.

Due to the ZZ coupling between the system and environment, operations performed on the system do not change the population of the environment. Thus, as long as the sampling operations on the system are strongly simulable—such as a computational-basis measurement and feed-forward—we can always sample over the uniform distribution of the environment and strongly simulate the resulting operations on the system.

We can now write the Choi state of this process as a mixture of separable states. For a k -step process, this is given by

$$\Upsilon_{k:0} = \frac{1}{2^{n-1}} \sum_{j=0}^{2^{n-1}-1} \bigotimes_{i=1}^k \left| \Psi(\theta_j^i) \right\rangle \left\langle \Psi(\theta_j^i) \right| \otimes \rho_0, \quad (15)$$

where ρ_0 is the initial state of the system and the j are the 2^{n-1} different $(n-1)$ -bit binary strings representing the state of the environment. For a fixed bit j at the i th time step, θ_j^i represents the cumulative phase applied to one half of the maximally entangled Bell pair depending on the string j , which determines the control for the controlled- R_z gates. We then sum over all $n-1$ bit binary strings, which corresponds to tracing over the environment.

From a physical perspective, the fact that OpenDQP–(IQP) is easy to sample from for any number of times and any choice of operations is surprising. We are probing a small part of a very complex system–environment interaction across multiple times that is, in itself, hard to sample from. Although there exists an equivalence relation between the complexity of the set of states and processes, there are clearly additional requirements on the system–environment interaction in order to leverage spatial complexity into temporal complexity. From this, we learn, at minimum, that a complex process on a single qubit must either generate temporal entanglement or (if separable) be described as a mixture of trajectories where the probability distribution is hard to sample from.

It is worth emphasizing that these results are only applicable to few-body open quantum systems. Here, we consider only a single-qubit system. In contrast, open many-body systems without temporal entanglement can still be classically hard to simulate, since these many-body states can exhibit complex spatial entanglement. A myriad of sophisticated classical techniques are required to model even Lindblad master equations [84], restricting classical simulation to modest system sizes. In fact, from a computational-complexity perspective, there is evidence that large many-body systems evolving according to an IQP circuit or, equivalently, an Ising model, are classically hard to sample from at a single time in the presence of depolarizing noise; otherwise, the polynomial hierarchy collapses to the third level [85]. We therefore anticipate that multitime sampling in these instances would also remain classically hard.

Additionally, even though sampling from OpenDQP–(IQP) on a single qubit is an easy task, the related problem of learning the output distribution may still be classically hard. The difficulty with which quantum stochastic processes may be learned constitutes an interesting avenue for future work [86–89].

F. The Heisenberg interaction and Ising models

We are now in a position to reexamine the Heisenberg interaction and IQP models in a more general system-plus-bath approach. Here, we explicitly map out the previous results in terms of a Hamiltonian of the form $H_{\text{total}} = H_S + H_{SE} + H_E$, where S and E refer to the system and environment, respectively. In doing so, we clarify which terms may be responsible for the complexity in each of the models.

First, we note that all of our discussions consider the dynamics of small systems; in particular, systems that would otherwise be classically simulable as closed systems. Therefore, the source of sampling complexity comes through interaction with an environment and choice of sampling operations. We begin with the Heisenberg interaction in a spin chain. This sampling task can be represented by the Hamiltonian

$$H_{\text{total}} = H_S(t) + \alpha_{S,q_0} (X_S X_{q_0} + Y_S Y_{q_0} + Z_S Z_{q_0}) + J_{ij} (X_i X_j + Y_i Y_j + Z_i Z_j), \quad (16)$$

where $H_S(t) = \sum_m f_m(t) h_S^{(m)} \otimes \mathbb{1}_E$ is a set of time-dependent operators on the system determined by the modulating parameters $f_m(t)$, α_{S,q_0} is the strength of the coupling between the system and its neighboring qubit in the spin chain, and the final term is the global Heisenberg-like coupling in the environment qubits. Here, the source of complexity can come from local driving on the system alone, which is capable of generating complex system–environment states through the global interaction in the spin chain.

In contrast, we observe that the complexity of the IQP model cannot be determined through operations on the system alone. The IQP circuit can be defined by an Ising Hamiltonian over a graph $G = (V, E)$ such that

$$H_G = \sum_{\{i,j\} \in E} \omega_{ij} X_i X_j + \sum_{k \in V} v_k X_k, \quad (17)$$

where the vertices are qubits, ω_{ij} are the interaction strengths, and the v_k characterize the strength of the external field. In the circuit model, all intermediate Hadamard gates cancel, resulting in an initial layer of Hadamard gates, a series of gates diagonal in the Z basis, and a final layer of Hadamard gates [as depicted in Fig. 6(a)]. Our system–environment unitary interaction then corresponds to a Hamiltonian given by

$$H_{\text{total}} = H_S + Z_S Z_{q_0} + \sum_j Z_j Z_{j+1}, \quad (18)$$

where H_S is a time-dependent control on the system, q_0 is the first qubit of the environment, and the last term is a summation over all qubit interactions in the environment.

As previously noted, due to the ZZ coupling between the system and environment, operations on the system do not change the population of the environment and thus cannot generate complexity. Instead, at a minimum we require terms that change the population of the environment to one that is classically hard to sample from. One such possibility would be the inclusion of local X terms in the IQP circuit corresponding to a transverse field in the environment given by $H_E = \sum_j Z_j Z_{j+1} + X_j$. Globally, this system-environment interaction is then a transverse-field Ising model, where we study the complexity of a small subsystem. Although this does not preclude the existence of other classical algorithms to simulate these dynamics, it does motivate where sources of complexity can come from in these interactions.

V. NUMERICAL EVIDENCE FOR RANDOM PROCESSES \subseteq OpenDQP\BPP

Having determined the general characteristics of system-environment interactions that contribute to the classical difficulty of dynamic sampling, we now return to the study of physical models. In Sec. IV C, we identify a physical Hamiltonian, the Heisenberg interaction, that in principle could result in a dynamic sampling problem that is classically hard. We now consider two models that can be directly investigated for their complexity using numerical benchmarks, thus providing further verification that classical complexity is obtainable when sampling from physical models of open quantum evolution. They are RCS and random matrix sampling. The latter enjoys a wide range of physical applications [90], such as modeling quantum chaos [91], the statistical properties of atomic nuclei [92], and conductance in disordered mesoscopic systems [93], while the former has been proven to demonstrate average-case hardness [52] and is conjectured to be robust to realistic levels of noise [94], which is essential for near-term implementations of quantum advantage on noisy quantum devices. This has led to one of the first claimed demonstrations of quantum computational advantage through RCS on a 53-qubit superconducting quantum computer [95].

Here, we present a series of numerical simulations of an OpenDQP sampling problem where the system-environment dynamics are random processes based on these models. First, we examine SE evolution, where the unitaries are matrices drawn from the circular unitary ensemble (CUE) with Haar measure. Then, we consider random Hamiltonians drawn from the Gaussian unitary ensemble (GUE).

Despite significant progress in the complexity-theoretic arguments for the hardness of RCS [94], it is still unknown whether sampling on a reduced output of a random circuit is hard. Not surprisingly, it is extremely challenging to prove this directly, so instead we rely on a numerical

benchmark to demonstrate classical hardness by showing that the temporal output distribution of our simulations converges to a Porter-Thomas distribution (PTD) consistent with quantum chaos [38].

A. Haar-random unitary evolution

We first describe a process with evolution based on the random circuits utilized in RCS. We consider the process tensor as shown in Fig. 1(a), with a single-qubit system interacting with an environment, of total size n . The initial system-environment state is given by $|0\rangle^{\otimes n}$ and the unitaries $U_{j;j-1}$ at each time step are given by random matrices drawn from the Haar measure on $U(d)$, where $d = 2^n$ is the dimension of the matrix.

Since the Haar measure is invariant under left and right multiplication by an independent unitary matrix, we choose our matrices randomly at each time step so that our cumulative evolution remains Haar random. From a physical perspective, we can consider this process as a system and environment evolving according to a Hamiltonian that changes randomly in time. While evolution according to Haar-random unitaries is in itself not a physical model and can only be achieved by random circuits in exponential time, the products of random unitary matrices appear in chaotic scattering and periodic time-dependent perturbations, which has motivated the study of the properties of composed ensembles of such unitaries [96].

At each of the k times, we make a measurement in the computational basis. The results of these successive measurements form a binary bit string of length k , which is sampled according to the output probability distribution given by Eq. (2).

1. Output statistics

We examine the output probability distribution generated by sampling our system at a total of k times. We consider both the shape of the distribution and the “distance” of the observed distribution to the Porter-Thomas distribution. For this, we use the Kullback-Leibler (K-L) divergence, which for two probability distributions P and Q is given by

$$\mathcal{D}_{\text{K-L}}(P \parallel Q) \equiv \sum_i P_i \ln \frac{P_i}{Q_i}. \quad (19)$$

The Kullback-Leibler divergence provides a measure of how well the Porter-Thomas distribution approximates our output distribution. For our calculations, P is the experimentally determined distribution and Q is the ideal Porter-Thomas distribution.

In Figs. 7(a)–7(c), we plot the rescaled output bit-string probabilities Np of our simulations for varying sampling times k , where $N = 2^k$. For our simulations, we have an environment of five qubits interacting with a single-qubit




FIG. 7. The output probability distribution for dynamic sampling of Haar-random unitary evolution. (a)–(c) The plots represent the rescaled probability distributions obtained by sampling at k times. The simulations comprise a five-qubit environment interacting with a single-qubit system. The sampling times vary from $k = 10$ on the left, to $k = 50$, and to $k = 105$ on the right. These are compared to the orange line, which represents the ideal Porter-Thomas distribution $\text{Pr}(Np) = e^{-Np}$, and the red dashed line, which shows the uniform distribution $\delta(p - 1/N)$. Here, $N = 2^k$ is the dimension of the output bit string of length k . At short sampling times, the distribution is peaked around $Np = 1$ and trending toward incoherent uniform randomness. At $k = 50$ sampling times, the distribution is close to an ideal Porter-Thomas distribution, with an average minimum K-L divergence of 0.04, before pulling away at longer sampling times and demonstrating a characteristic tail suggestive of an insufficiently thermalized system.

system. We begin sampling at $k = 5$ times and increase in increments of 5 until sampling at $k = 105$ times. When sampling at a small number of times [Fig. 7(a)], the distribution is peaked around the uniform distribution and appears to be converging toward incoherent uniform randomness. This is to be expected, since if we prepare an n -qubit random state and sample from only m qubits, then for $m \ll n$ the measured m -qubit state will be close to the maximally mixed state.

As we increase the number of times we sample, the system gradually approaches the Porter-Thomas distribution [Fig. 7(b)]. However, when sampling at a large number of times with $k > O(2^n)$, we observe a larger tail. These results can be confirmed by examining the K-L divergence of the output distribution from a Porter-Thomas distribution as a function of the number of sampling times, as shown in Fig. 8. The K-L divergence reaches an average minimum of approximately 0.04 when the number of sampling times is $k = 50$, before the distributions diverge again at larger sampling times.

The first and most important finding is that we achieve an experimental output distribution that closely resembles a Porter-Thomas distribution when sampling on only a single qubit over time. This result provides numerical support that the RCS version of OpenDQP is classically hard to sample from. However, as we sample beyond the times required for convergence to a Porter-Thomas distribution, our distribution stretches away and demonstrates a large tail. These tails are seen in previous work simulating random circuits [38,97] and are indicative of circuits with insufficient depth to fully thermalize. We provide further analysis of these tails in Sec. VC.

It is worth noting here, however, that our results have significant differences compared to those regarding the efficient classical simulation of recent quantum advantage RCS experiments. Our process involves a single-qubit system evolving according to a random circuit that is interacting with an environment that is purely quantum and random. Our results suggest that a classically simulable




FIG. 8. The K-L divergence from an ideal Porter-Thomas distribution for Haar-random unitary evolution. The K-L divergence for sampling times $k = 5$ to $k = 105$ is shown for simulations of a five-qubit environment interacting with a single-qubit system. Each data point is the average of 30 simulations of the process for the specified sampling time. The shaded region represents the standard deviation. For each of these simulations, we repeat our circuit 10 000 times to compute the probability density function. The average minimum K-L divergence is 0.0378 for $k = 50$.

small closed system can become complex when interacting with a noisy quantum environment. This is in contrast to large complex random circuits executed in the presence of a simple independent identically distributed (i.i.d.) noise model. As a closed system, a many-body system evolving according to a random circuit is classically hard to sample from [52]. Yet in the presence of two-qubit depolarizing noise, this sampling problem becomes classically efficiently simulable [98]. In fact, it appears that all local noise models transform random circuit sampling into an easy problem [99]. Nevertheless, real quantum devices exhibit both classical and complex quantum noise [48] and thus determining the complexity of simulating many-body random systems is not as straightforward as assuming a simple noise model.

B. Random Hamiltonian evolution

We can utilize the same numerical tools to extend our complexity analysis to consider a more physical dynamical model, namely, the evolution of a system and environment according to a random Hamiltonian. Random matrix theory has been frequently applied to the study of nuclear physics, with Hamiltonians drawn from the GUE and Gaussian orthogonal ensemble (GOE) modeling the spectral properties of nucleon excitation, particularly the breaking of time-reversal invariance in nuclei [100].

For our simulations, we consider a simplified model of a single qubit interacting with a fixed environment of total size n , initiated in the state $|0\rangle^{\otimes n}$. The unitary evolution operator is now given by $U_{j:j-1} = \exp(-iH\Delta t_j)$, where Δt_j corresponds to the time step $t_j - t_{j-1}$ and the Hamiltonian, H , is a matrix selected from the GUE so that each entry is a random complex Gaussian number. The output distribution should similarly converge to a Porter-Thomas type distribution in the limit of large $N = 2^n$. To make this model local, we weight the off-diagonal

matrix elements h_{ij} of H with an exponential cutoff, i.e., $h_{ij} \mapsto \exp\{-|i-j|\}h_{ij}$. Each time step Δt_j is then chosen randomly as a fraction of 2π .

1. Output statistics

The output probability distribution for our simulations of random Hamiltonian evolution is plotted for a number of sampling times in Fig. 9. Once again, we see a qualitative transformation in the output distribution from a distribution peaked around incoherent uniform randomness at short sampling times, to convergence to a Porter-Thomas distribution, and ultimately the emergence of tails at long sampling times. Similarly, we see a K-L divergence that reaches an average minimum of 0.14 with sampling time $k = 20$, as shown in Fig. 10.

While ensembles of large Gaussian random matrices can be used to model the statistical properties of complex closed quantum systems, we show here that the reduced dynamics of a single qubit of these larger systems remains chaotic when sampling over time. If we assume convergence of a probability distribution to the PTD as a measure of complexity, then this is a further step toward showing that more realistic physical models of open quantum systems may be classically hard to sample from and thus simulate.

C. Understanding the tails with random MPS

We now provide a heuristic analysis of the emergence of tails seen at long sampling times in our simulations. We specifically consider the Haar-random evolution, although the explanation equally applies to random Hamiltonian evolution. Since at each time step we implement a Haar-random unitary, which corresponds to a circuit of sufficient depth to demonstrate Porter-Thomas statistics on the entire system-environment state, we cannot immediately explain the tails as resulting from a non-fully thermalized circuit,




FIG. 9. The output probability distribution for dynamic sampling of random Hamiltonian evolution. (a)–(c) The plots represent the rescaled probability distributions obtained by sampling at k times. The simulations comprise a five-qubit environment interacting with a single-qubit system. The sampling times vary from $k = 5$ on the left, to $k = 15$, and to $k = 35$ on the right. The orange line represents the ideal Porter-Thomas distribution $\Pr(Np) = e^{-Np}$ and the red dashed line is the uniform distribution $\delta(p - 1/N)$.




FIG. 10. The K-L divergence from a Porter-Thomas distribution for random Hamiltonian evolution. The K-L divergence is plotted as a function of the total number of sampling times, k . A minimum of 0.142 is reached at $k = 20$. The shaded region represents the standard deviation computed over 30 runs of the simulation for each value of k .

although we expect that a similar phenomenon is occurring and that the circuit construction of the Choi state for this process results in a circuit with insufficient depth.

To see this, we first note that the dimension of the process-tensor Choi state grows with the number of sampling times k . However, the bond dimension of the state will be bounded by the effective dimension of the environment. We hypothesize that as we increase our sampling times, the bond dimension of our process saturates according to the random SE interaction of our unitaries and yet the dimension of our Choi state grows. Ultimately, as we sample at further times, we cannot explore any more of the environment and sampling from the resulting Choi state

is akin to sampling from a random matrix product state (MPS) with nonmaximal bond dimension.

We can confirm this intuition by comparing our results in Fig. 7 with the sampling of a Haar-random MPS with varying bond dimension. We demonstrate this in Figs. 11(a)–11(c), where we reproduce the qualitative transition in output distribution seen in our dynamic RCS simulations. Of particular note in Fig. 11(c) is the fact that we are able to reproduce the tail by sampling all qubits of a 12-qubit MPS with nonmaximal bond dimension $\chi = 20$.

VI. DISCUSSION AND CONCLUSIONS

In this work, we initiate the study of open quantum systems from a complexity-theoretic perspective. Specifically, we examine the classical complexity of sampling an open quantum system at successive points in time. We show the existence of quantum stochastic processes that, when formulated as multitime sampling problems of an open quantum system, are hard to classically simulate. This includes both circuit models and Hamiltonian dynamics.

The typical consideration in open quantum dynamics is to solve the corresponding master equation. We initially postulate that if the underlying quantum stochastic process is classically complex, then a member of the associated family of master equations will also be hard—which includes both driven and nondriven master equations. Our results hence provide a strong foundation for the existence of master equations that are classically hard to simulate. In particular, given that complexity in many of our models can be controlled by system operations, then it is highly likely that many driven master equations will also be hard. In fact, our results already enable us to comment on the




FIG. 11. The output probability distribution for sampling of a random MPS. We plot the rescaled probability distributions obtained from sampling Haar-random MPSS, varying the bond dimension and number of qubits sampled. Once again, the orange line represents the ideal Porter-Thomas distribution $\text{Pr}(Np) = e^{-Np}$ and the red dashed line is the uniform distribution $\delta(p - 1/N)$. Qualitatively, these distributions follow the same transition seen in the simulation of Haar-random unitary evolution and random Hamiltonian evolution for varying sampling times. (a) Sampling on seven qubits of a 12-qubit Haar-random MPS with maximum bond dimension $\chi = 64$. (b) Sampling on all 12 qubits of a 12-qubit Haar-random MPS with maximum bond dimension $\chi = 64$. (c) Sampling on all 12 qubits of a 12-qubit Haar-random MPS with nonmaximal bond dimension $\chi = 20$.

types of processes that will admit hard driven master equations. If a process is Markovian, then adding a local field via unitary operations should not change the complexity of this task and the family of master equations may be efficiently simulated [81]. However, for a non-Markovian process, unitary control on the system has the potential to drastically change the complexity of simulation.

Finding complex nondriven or constant local field master equations is perhaps more challenging, since reducing the number of sampling times requires that the complexity obtained by multitime sampling is propagated when the remaining process-tensor legs are projected onto identity operations. In itself, this task is already anticipated to be a hard problem in the worst case, as shown by the complexity of the DQC1 model [101–103], where a qubit is prepared, interacts with an infinite temperature bath, and is then sampled at some later time. In this setting, the set of maps is limited to those from time $t = 0$ to any later time $t = t_j$ and the resulting master equation is only valid for uncorrelated initial states with time-homogeneous underlying dynamics and stationary initial-environment states. Moreover, the specific circuit used to prove the hardness of the DQC1 model is highly engineered [102] and it is known that without entanglement, the DQC1 model is easy to simulate [104]. Nonetheless, even if a single example of this restricted model is likely to be hard, then it is promising that more complicated nondriven master equations with initial correlations or inhomogeneous dynamics will also be classically hard to simulate.

To support this, consider the quantum stochastic process with initial system-environment state $\rho_0 = |+\rangle\langle +|_S \otimes |0\rangle\langle 0|_E$ and Hamiltonian $H_{SE} = |0\rangle\langle 0|_S \otimes H_E$, where H_E is drawn from the Gaussian unitary ensemble as per Sec. V B. The complexity of the associated nondriven master equation corresponds to the set of two-time samplings within OpenDQP_2 . Thus, we evolve our state for some time t according to $U_{SE}^t = \exp\{-iH_{SE}t\} = |0\rangle\langle 0|_S \otimes w_t + |1\rangle\langle 1|_S \otimes \mathbb{1}_E$ and sample the system of the resulting system-environment state, $|\Psi_{SE}\rangle = 1/\sqrt{2}(|0\rangle_S |w_t\rangle_E + |1\rangle_S |0\rangle_E)$, where $|w_t\rangle_E = w_t|0\rangle_E$. The corresponding density matrix is $\Psi_{SE} = |\Psi_{SE}\rangle\langle \Psi_{SE}|$, such that the reduced density matrix of the system is given by

$$\Psi_S = \frac{1}{2} \begin{bmatrix} 1 & {}_E\langle 0|w_t\rangle_E \\ {}_E\langle w_t|0\rangle_E & 1 \end{bmatrix}. \quad (20)$$

Repeating this for varying values of t amounts to sampling from the set of stochastic maps $\{\mathcal{E}_{t,0}\}$. When sampling, we can choose to measure in the X and Y bases with outcomes x and y , respectively; then, $\langle X \rangle^2 + \langle Y \rangle^2 = |\langle 0|w_t\rangle|^2$. But this is precisely the probability of obtaining the all-zero bit string when sampling the environment state that has evolved according to $U_E^t = \exp\{-iH_E t\}$. And for randomly chosen t , the outcomes will form a probability density function with an exponential shape like the

Porter-Thomas distribution. In this instance, the complexity of the master equation corresponds to sampling from the ensemble $\{|w_t\rangle\}$, which is equivalent to the ensemble of unitary matrices $\{U_E^t\}$. In fact, for any complex U^t , this master equation will also be complex. The foundational arguments made here pave the way for determining more general master equations that are classically hard and, furthermore, to delineate the subsets OpenDQP_2 and OpenDQP_k .

This work bridges the gap from spatial to temporal sampling problems, which establishes the pathway to systematically examining the complexity of physically significant models of open quantum systems—e.g., understanding the paradigmatic spin-boson model. We anticipate that our results will serve as a firm foundational basis for the exploration of open quantum system simulation as an avenue to quantum advantage. The ideal application will be in understanding the complexity of important problems in the field, such as the effect of coherent driving on dissipative environments. As we show, quantum computers may find utility in determining features of interesting open systems: a straightforward task such as finding the idle dynamics of a state or as exotic as learning complex multitime features in a quantum stochastic process.

ACKNOWLEDGMENTS

I.A.A. is supported by an Australian Government Research Training Program Scholarship, a Monash Graduate Excellence Scholarship, and the Alan P. Roberts Doctoral Scholarship. G.A.L.W. is supported by an Australian Government Research Training Program Scholarship. K.M. acknowledges support from the Australian Research Council Future Fellowship under Grant No. FT160100073. K.M. and C.D.H. acknowledge the support of the Australian Research Council’s Discovery Project under Grant No. DP210100597. K.M. and C.D.H. were recipients of the International Quantum U Tech Accelerator award by the U.S. Air Force Research Laboratory.

APPENDIX A: THE PROCESS-TENSOR FRAMEWORK

Process tensors are a mathematical framework for describing arbitrary quantum stochastic processes. The general setup, as described in the main text in Fig. 1(a), is that of an accessible system coupled to an inaccessible environment, which evolve together according to a series of unitaries $U_{j,j-1}$. The experimenter can drive the system with a sequence of control operations $\mathbf{A}_{k-1,0}$, which may take the form of measurements, unitaries, or more generally correlated instruments or testers, which we discuss below. A process tensor represents all of the uncontrollable dynamics of the process.

Formally, a process tensor is a multilinear map that takes in a sequence of operations at different times and outputs

a final state conditioned on this choice of interventions. In doing so, process tensors account for the fact that a choice of operation on the system may affect the future evolution of the system. However, it is important to note that the process exists independent of the control operations and, instead, the control operations, which serve as input variables into the process, are a choice made by the experimenter. Mathematically, the controlled dynamics are given by

$$\begin{aligned} \rho_k(\mathbf{A}_{k-1:0}) &= \text{tr}_E(U_{k:k-1}\mathcal{A}_{k-1} \dots U_{1:0}\mathcal{A}_0(\rho_0^{SE})) \\ &:= \mathcal{T}_{k:0}[\mathbf{A}_{k-1:0}], \end{aligned} \quad (\text{A1})$$

where $\mathcal{T}_{k:0}$ is the process-tensor representation of the quantum stochastic process.

In Fig. 1(a), the sequence of operations \mathcal{A}_j is depicted as independent; however, the choice of instruments can be more general. Specifically, the operations can be testers with classical or quantum correlations. This can be achieved through the use of an ancilla interacting with the system, which may act as a quantum memory. The resulting tester is then a higher-order quantum map, given by

$$\mathbf{A}_{k-1:0} = (\Pi_k^{SA} \mathcal{A}_{k-1}^{SA} \dots \mathcal{A}_0^{SA}(\rho_0^A)), \quad (\text{A2})$$

where the superscript A denotes action on the ancilla. Equation (A1) can then be rewritten as

$$\begin{aligned} \rho_k(\mathbf{A}_{k-1:0}) &= \text{tr}_{EA}(\Pi_k^{SA} U_{k:k-1}^{SE} \mathcal{A}_{k-1}^{SA} \dots U_{1:0}^{SE} \mathcal{A}_0^{SA}(\rho_0^{SE} \otimes \rho_0^A)). \end{aligned} \quad (\text{A3})$$

Similarly, the resulting probability distribution is given by

$$\mathbb{P}(\mathbf{x}_{k:0} | \mathbf{J}_{k:0}) = \text{Tr}(\Pi_k^{SA} U_k^{SE} \dots U_1^{SE} \mathcal{A}_0^{SA}(\rho_0^{SE} \otimes \rho_0^A)). \quad (\text{A4})$$

Once again, we note that the operations constitute a set of dependencies for the task of sampling from this probability distribution. In particular, in non-Markovian evolution, the choice of operations may influence future statistics.

1. Choi-state representation of multitime processes

Figure 12(a) provides a more detailed circuit diagram for construction of the process-tensor Choi state using the generalized Choi-Jamiolkowski isomorphism. At each time, t_j , one half of a new Bell pair is swapped into the circuit and interacts with the environment according to the SE dynamics $U_{j:j-1}$. The corresponding output leg from the process tensor and the input leg to the next time step are labeled o_j and i_{j+1} , respectively. The resulting state on the system and k Bell pairs constitutes the k -step process tensor, $\Upsilon_{k:0}$, with the tensor legs labeled $\{o_k, i_k, o_{k-1}, \dots, i_1, o_0\}$.

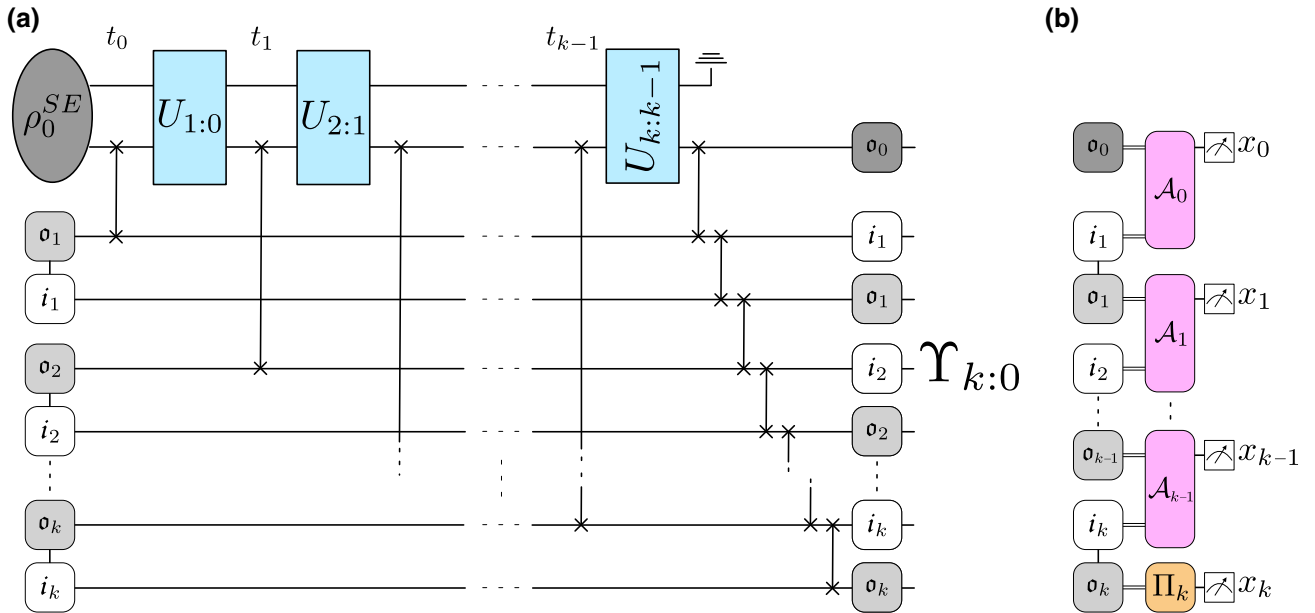


FIG. 12. Sampling the process-tensor Choi state. (a) The circuit construction of the Choi state of a k -step process using a generalization of the Choi-Jamiolkowski isomorphism. One half of k Bell pairs is swapped into the process at each time. At the end of the circuit, a series of SWAP gates reorder the indices to correspond to causal ordering. (b) The resulting Choi state can be sampled with operations \mathcal{A}_j and measurement apparatus Π_k , resulting in measurement outcomes x_j .

The controlled operations on the system, which the experimenter performs at each time, are represented by a set of CP maps $\{\mathcal{A}_0, \mathcal{A}_1, \dots, \mathcal{A}_{k-1}\}$ plus a final measurement apparatus $\{\Pi_k\}$. Each operation has an associated measurement outcome x_j . When considering the process tensor in its Choi-state representation, these operations constitute observables of the many-body state and the resulting temporal probability distribution conditioned on the choice of instruments is found by projection onto the Choi state of the operators:

$$\begin{aligned} & \mathbb{P}(x_k, \dots, x_0 | \mathcal{J}_k, \dots, \mathcal{J}_0) \\ &= \text{Tr} \left[\Upsilon_{k:0} \left(\Pi_k \otimes \left(\hat{\mathcal{A}}_{k-1}^{(x_{k-1})} \otimes \dots \otimes \hat{\mathcal{A}}_0^{(x_0)} \right)^T \right) \right]. \end{aligned} \quad (\text{A5})$$

Ultimately, the act of sampling a system sequentially over a set of times $\{t_0, t_1, \dots, t_k\}$ is transformed into sampling of the output of a many-body state, as shown in Fig. 12(b).

APPENDIX B: SAMPLING FROM THE OPENDQP-(IQP) PROCESS

There exists an efficient classical algorithm to sample the output of the process-tensor version of IQP. Full details of this proof are given in the text. Here, we provide the pseudocode for this algorithm.

APPENDIX C: SHOR'S ALGORITHM

1. Choi state

Here, we provide a detailed construction of the process-tensor Choi state for the order-finding circuit of Shor's algorithm as given by the circuit in Fig. 3(b). First, we consider a Choi state with indices ordered as $\{\mathbf{o}_k, \mathbf{o}_{k-1}, \dots, \mathbf{o}_0, \mathbf{i}_k, \mathbf{i}_{k-1}, \dots, \mathbf{i}_0\}$. This ordering is consistent with the state that results from the circuit in Fig. 3(a), where the $2n$ upper qubit register corresponds to the output

Input: Unitaries $\mathbf{U}_{k:1}$, number of sampling times k , error ϵ

Output: sample \mathbf{x} from distribution $\mathcal{R}(\mathbf{x})$ such that $\|\mathcal{R}(\mathbf{x}) - \mathcal{D}(\mathbf{x})\|_1 \leq \epsilon$

- 1: **for** each U_j **do**
- 2: decompose $poly(n)$ 2-qubit diagonal gates into controlled- R_z acting on the system, and single qubit diagonal gates on the environment
- 3: **end for**
- 4: Pick an $n - 1$ length bit-string uniformly at random
- 5: Use the bit-string to fix the controlled- R_z gates acting on the system qubit
- 6: Apply $poly(n)$ single qubit R_z gates and k measurement operators to the system qubit
- 7: **return** k length bit-string \mathbf{x} as sample

Algorithm 1. Sampling algorithm

legs of each Bell pair and there are an additional $2n$ input legs corresponding to the other half of each Bell pair. We introduce the construction of the Choi state in this manner first, to familiarize the reader. Following the application of $2n$ controlled unitaries, the entire system-environment Bell-pair state is given by

$$|\psi\rangle = \frac{1}{\sqrt{2^{2n}}} \sum_{x=0}^{2^{2n}-1} |xx\rangle |a^x \bmod N\rangle, \quad (\text{C1})$$

where $n = \log(N)$ and the modular exponential function is defined as $f(x) = a^x \bmod N$. Tracing over the environmental degrees of freedom gives the k -step process-tensor Choi state:

$$\begin{aligned} \Upsilon_{k:0} &= \text{tr}_E (|\psi\rangle\langle\psi|) \\ &= \frac{1}{2^{2n}} \sum_{z \in M} \langle z | \sum_{x=0}^{2^{2n}-1} \sum_{x'=0}^{2^{2n}-1} |xx\rangle\langle x'x'| \otimes |f(x)\rangle\langle f(x')| |z\rangle, \end{aligned} \quad (\text{C2})$$

where M is the codomain of $f(x)$; that is, all y such that $f(x) = y$. If we define $|\Psi_i\rangle := \sum_x |xx\rangle$, summed over all x for which $f(x) = y_i$, then we arrive at the final representation:

$$\Upsilon_{k:0} = \frac{1}{2^{2n}} \sum_i |\Psi_i\rangle\langle\Psi_i|. \quad (\text{C3})$$

This is the Choi state of Fig. 3(b), up to a reordering of indices. If we consider indices ordered according to time evolution, we can construct the associated Choi state that we project onto our controlled sequence of operations. In this setting, the final system-environment Bell-pair state is given by

$$\begin{aligned} |\psi\rangle &= \frac{1}{\sqrt{2^{2n}}} \sum_{x_0 x_1 x_2 \dots x_{2n}=0}^1 |x_0 x_0\rangle |x_1 x_1\rangle \dots |x_{2n} x_{2n}\rangle \\ &\otimes |a^{2^{2n} x_{2n}} a^{2^{2n-1} x_{2n-1}} \dots a^{2^0 x_0}\rangle. \end{aligned} \quad (\text{C4})$$

Once again, we take the trace over the environment to retrieve the process-tensor Choi state and define $|\phi_i\rangle := \sum_{x_i} \bigotimes_i |x_i x_i\rangle$, where the summation is over all x_i such that $f(x) = y_i$. The resulting Choi state can then be expressed as

$$\Upsilon_{k:0} = \frac{1}{2^{2n}} \sum_i |\phi_i\rangle\langle\phi_i|. \quad (\text{C5})$$

2. Sampling from the Choi state

Sampling from the process-tensor Choi state corresponds to projection of the state onto the controlled

sequence of operations $\mathbf{A}_{k:0}$. For the inverse QFT, the general operation at time t_j is defined as

$$\mathcal{A}_j^{x_j} := \text{Tr} \left(M_{x_j} R_{\mathbf{m}_j} \rho_j R_{\mathbf{m}_j}^\dagger \right) |+\rangle\langle +|_{j+1}. \quad (\text{C6})$$

Here, $\mathbf{m}_j := \{m_j, m_{j-1}, \dots, m_1\}$, to indicate that the rotation operation depends on all previous measurement results. We exclude the index m_0 , since the first operation is simply a Hadamard gate to prepare the $|+\rangle$ state before the first controlled modular multiplication. With respect to the Choi state, we can see that the operation in Eq. (C6) acts on legs $\{\sigma_j, i_{j+1}\}$. Specifically, at time t_j , the operation \mathcal{A}_j takes in the output leg from the process σ_j , given by state ρ_j , rotates it according to prior measurements, and performs a measurement in the X basis, before preparing the $|+\rangle_{j+1}$ state as input into the process leg i_{j+1} . The rotation operation $R_{\mathbf{m}_j}$ is given by

$$R_{\mathbf{m}_j} = \begin{bmatrix} 1 & 0 \\ 0 & \phi_j \end{bmatrix}, \quad (\text{C7})$$

with $\phi_j = e^{-2\pi i \sum_{k=1}^{j-1} m_{j-k}/2^k}$.

Sampling in the computational basis then corresponds to operations \mathcal{A}_j , with $x_j = 0$, and thus the rotation operation simply corresponds to the identity. Using the construction of the Choi state for Shor’s algorithm, we can provide an alternative proof demonstrating that it is classically easy to sample in the computational basis at each time of the process. Consider the state in Eq. (C3). Sampling in the computational basis on the resulting process $\Upsilon_{k:0}$ then corresponds to measuring the output legs $\{\sigma_k, \sigma_{k-1}, \dots, \sigma_0\}$ of this state and projecting the input legs $\{i_k, i_{k-1}, \dots, i_0\}$ onto the postmeasurement state, which, without loss of generality, we can choose to be $|0\rangle\langle 0|$. This is, therefore, equivalent to tracing the lower register of n -qubits in Fig. 3(a) and measuring the upper register of $2n$ qubits before the inverse QFT is performed. For simplicity, we consider the complexity of this equivalent sampling problem. To simulate the outcome of this sampling problem, we proceed as follows. The function f is periodic with period r and the number of periods is therefore given by $A = 2^{2n}/r$. Taking the partial trace over the environment amounts to measuring the environment and discarding the result. Since each possible value of the function f occurs the same number of times, precisely A times, we can select one outcome, y_0 , uniformly at random. This projects the register of $2n$ qubits into an equal superposition of the total A values of x such that $f(x) = y_0$. The resulting state is therefore

$$|\psi_{2n}\rangle = \frac{1}{\sqrt{A}} \sum_{j=0}^{A-1} |x_0 + jr\rangle. \quad (\text{C8})$$

Measuring this register then gives a bit string $x_0 + j_0 r$, where once again each value of j_0 is equally likely and

so is chosen uniformly at random. Overall, this returns a number between 0 and 2^{2n} uniformly at random.

-
- [1] J. R. McClean, N. C. Rubin, J. Lee, M. P. Harrigan, T. E. O’Brien, R. Babbush, W. J. Huggins, and H.-Y. Huang, What the foundations of quantum computer science teach us about chemistry, *J. Chem. Phys.* **155**, 150901 (2021).
 - [2] A. Smith, M. Kim, F. Pollmann, and J. Knolle, Simulating quantum many-body dynamics on a current digital quantum computer, *npj Quantum Inf.* **5**, 106 (2019).
 - [3] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, and N. P. Sawaya, *et al.*, Quantum chemistry in the age of quantum computing, *Chem. Rev.* **119**, 10856 (2019).
 - [4] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, *Rev. Mod. Phys.* **92**, 015003 (2020).
 - [5] G. Cohen and E. Rabani, Memory effects in nonequilibrium quantum impurity models, *Phys. Rev. B* **84**, 075150 (2011).
 - [6] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature* **549**, 242 (2017).
 - [7] K. Kim and S. Saito, Multi-time density correlation functions in glass-forming liquids: Probing dynamical heterogeneity and its lifetime, *J. Chem. Phys.* **133**, 044511 (2010).
 - [8] J. Ono, S. Takada, and S. Saito, Couplings between hierarchical conformational dynamics from multi-time correlation functions and two-dimensional lifetime spectra: Application to adenylate kinase, *J. Chem. Phys.* **142**, 06B601_1 (2015).
 - [9] D. A. Mazziotti, Effect of strong electron correlation on the efficiency of photosynthetic light harvesting, *J. Chem. Phys.* **137**, 074117 (2012).
 - [10] X. Li and C. Wang, Succinct description and efficient simulation of non-Markovian open quantum systems (2021), [ArXiv:2111.03240](https://arxiv.org/abs/2111.03240).
 - [11] R. Cleve and C. Wang, in *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017* (2017).
 - [12] R. Trivedi, Description and complexity of non-Markovian open quantum dynamics (2022), [ArXiv:2204.06936](https://arxiv.org/abs/2204.06936).
 - [13] I. De Vega and D. Alonso, Dynamics of non-Markovian open quantum systems, *Rev. Mod. Phys.* **89**, 015001 (2017).
 - [14] R. Rosenbach, J. Cerrillo, S. F. Huelga, J. Cao, and M. B. Plenio, Efficient simulation of non-Markovian system-environment interaction, *New J. Phys.* **18**, 023035 (2016).
 - [15] M. Cygorek, M. Cosacchi, A. Vagov, V. M. Axt, B. W. Lovett, J. Keeling, and E. M. Gauger, Simulation of open quantum systems by automated compression of arbitrary environments, *Nat. Phys.* **18**, 662 (2022).
 - [16] A. Strathearn, P. Kirton, D. Kilda, J. Keeling, and B. W. Lovett, Efficient non-Markovian quantum dynamics using

- time-evolving matrix product operators, *Nat. Commun.* **9**, 3322 (2018).
- [17] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, Oxford, 2002).
- [18] A. Strathearn, B. W. Lovett, and P. Kirton, Efficient real-time path integrals for non-Markovian spin-boson models, *New J. Phys.* **19**, 093009 (2017).
- [19] M. R. Jørgensen and F. A. Pollock, Exploiting the Causal Tensor Network Structure of Quantum Processes to Efficiently Simulate Non-Markovian Path Integrals, *Phys. Rev. Lett.* **123**, 240602 (2019).
- [20] K. Head-Marsden, S. Krastanov, D. A. Mazziotti, and P. Narang, Capturing non-Markovian dynamics on near-term quantum computers, *Phys. Rev. Res.* **3**, 013182 (2021).
- [21] A. Nagy and V. Savona, Variational Quantum Monte Carlo Method with a Neural-Network Ansatz for Open Quantum Systems, *Phys. Rev. Lett.* **122**, 250501 (2019).
- [22] H.-T. Chen, G. Cohen, and D. R. Reichman, Inchworm Monte Carlo for exact non-adiabatic dynamics. I. Theory and algorithms, *J. Chem. Phys.* **146**, 054105 (2017).
- [23] G. Cohen, E. Gull, D. R. Reichman, A. J. Millis, and E. Rabani, Numerically exact long-time magnetization dynamics at the nonequilibrium Kondo crossover of the Anderson impurity model, *Phys. Rev. B* **87**, 195108 (2013).
- [24] J. Cerrillo and J. Cao, Non-Markovian Dynamical Maps: Numerical Processing of Open Quantum Trajectories, *Phys. Rev. Lett.* **112**, 110401 (2014).
- [25] S. Gherardini, A. Smirne, S. F. Huelga, and F. Caruso, Transfer-tensor description of memory effects in open-system dynamics and multi-time statistics, *Quantum Sci. Technol.* **7**, 025005 (2022).
- [26] M. Buser, J. Cerrillo, G. Schaller, and J. Cao, Initial system-environment correlations via the transfer-tensor method, *Phys. Rev. A* **96**, 062122 (2017).
- [27] A. A. Kananenka, C.-Y. Hsieh, J. Cao, and E. Geva, Accurate long-time mixed quantum-classical Liouville dynamics via the transfer tensor method, *J. Phys. Chem. Lett.* **7**, 4809 (2016).
- [28] I. Luchnikov, S. Vintskevich, H. Ouerdane, and S. Filipov, Simulation Complexity of Open Quantum Dynamics: Connection with Tensor Networks, *Phys. Rev. Lett.* **122**, 160401 (2019).
- [29] A. Khan, D. Quigley, M. Marcus, E. Thyraug, and A. Datta, Model-Independent Simulation Complexity of Complex Quantum Dynamics, *Phys. Rev. Lett.* **126**, 150402 (2021).
- [30] K. A. Jung, P. E. Videla, and V. S. Batista, Ring-polymer, centroid, and mean-field approximations to multi-time Matsubara dynamics, *J. Chem. Phys.* **153**, 124112 (2020).
- [31] A. P. Lund, M. J. Bremner, and T. C. Ralph, Quantum sampling problems, bosonsampling and quantum supremacy, *npj Quantum Inf.* **3**, 1 (2017).
- [32] S. Aaronson and A. Arkhipov, in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (2011), p. 333.
- [33] M. J. Bremner, R. Jozsa, and D. J. Shepherd, Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy, *Proc. R. Soc. A: Math. Phys. Eng. Sci.* **467**, 459 (2011).
- [34] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, *Nat. Phys.* **14**, 595 (2018).
- [35] J. Tangpanitanon, S. Thanasilp, M.-A. Lemonde, N. Dangiam, and D. G. Angelakis, Quantum supremacy in driven quantum many-body systems (2020), [ArXiv:2002.11946](https://arxiv.org/abs/2002.11946).
- [36] B. Fefferman, M. Foss-Feig, and A. V. Gorshkov, Exact sampling hardness of Ising spin models, *Phys. Rev. A* **96**, 032324 (2017).
- [37] A. Deshpande, B. Fefferman, M. C. Tran, M. Foss-Feig, and A. V. Gorshkov, Dynamical Phase Transitions in Sampling Complexity, *Phys. Rev. Lett.* **121**, 030501 (2018).
- [38] E. Kapit, P. Roushan, C. Neill, S. Boixo, and V. Smelyanskiy, Entanglement and complexity of interacting qubits subject to asymmetric noise, *Phys. Rev. Res.* **2**, 043042 (2020).
- [39] O. Shtanko, A. Deshpande, P. S. Julienne, and A. V. Gorshkov, Complexity of Fermionic Dissipative Interactions and Applications to Quantum Computing, *PRX Quantum* **2**, 030350 (2021).
- [40] F. A. Pollock, C. Rodríguez-Rosario, T. Frauenheim, M. Paternostro, and K. Modi, Non-Markovian quantum processes: Complete framework and efficient characterization, *Phys. Rev. A* **97**, 012127 (2018).
- [41] Y. Tanimura, Stochastic Liouville, Langevin, Fokker-Planck, and master equation approaches to quantum dissipative systems, *J. Phys. Soc. Jpn.* **75**, 082001 (2006).
- [42] W.-M. Zhang, Exact master equation and general non-Markovian dynamics in open quantum systems, *Eur. Phys. J. Spec. Top.* **227**, 1849 (2019).
- [43] D. Brian and X. Sun, Generalized quantum master equation: A tutorial review and recent advances, *Chin. J. Chem. Phys.* **34**, 497 (2021).
- [44] F. A. Pollock, C. Rodríguez-Rosario, T. Frauenheim, M. Paternostro, and K. Modi, Operational Markov Condition for Quantum Processes, *Phys. Rev. Lett.* **120**, 040405 (2018).
- [45] F. Costa and S. Shrapnel, Quantum causal modelling, *New J. Phys.* **18**, 063032 (2016).
- [46] S. Milz and K. Modi, Quantum Stochastic Processes and Quantum Non-Markovian Phenomena, *PRX Quantum* **2**, 030201 (2021).
- [47] G. A. L. White, C. D. Hill, F. A. Pollock, L. C. L. Hollenberg, and K. Modi, Demonstration of non-Markovian process characterisation and control on a quantum processor, *Nat. Commun.* **11**, 6301 (2020).
- [48] G. A. L. White, F. A. Pollock, L. C. L. Hollenberg, C. D. Hill, and K. Modi, From many-body to many-time physics (2021), [ArXiv:2107.13934](https://arxiv.org/abs/2107.13934).
- [49] G. A. L. White, F. A. Pollock, L. C. L. Hollenberg, K. Modi, and C. D. Hill, Non-Markovian Quantum Process Tomography, *PRX Quantum* **3**, 020344 (2022).
- [50] C. Meier and D. J. Tannor, Non-Markovian evolution of the density operator in the presence of strong laser fields, *J. Chem. Phys.* **111**, 3365 (1999).

- [51] F. Verstraete, M. M. Wolf, and J. Ignacio Cirac, Quantum computation and quantum-state engineering driven by dissipation, *Nat. Phys.* **5**, 633 (2009).
- [52] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, On the complexity and verification of quantum random circuit sampling, *Nat. Phys.* **15**, 159 (2019).
- [53] S. Shrapnel, F. Costa, and G. Milburn, Updating the Born rule, *New J. Phys.* **20**, 053010 (2018).
- [54] S. Milz, F. Sakuldee, F. A. Pollock, and K. Modi, Kolmogorov extension theorem for (quantum) causal modelling and general probabilistic theories, *Quantum* **4**, 255 (2020).
- [55] P. Strasberg and M. G. Díaz, Classical quantum stochastic processes, *Phys. Rev. A* **100**, 022120 (2019).
- [56] S. Milz, D. Egloff, P. Taranto, T. Theurer, M. B. Plenio, A. Smirne, and S. F. Huelga, When Is a Non-Markovian Quantum Process Classical?, *Phys. Rev. X* **10**, 041049 (2020).
- [57] P. Figueroa-Romero, K. Modi, and F. A. Pollock, Almost Markovian processes from closed dynamics, *Quantum* **3**, 136 (2019).
- [58] P. Figueroa-Romero, F. A. Pollock, and K. Modi, Markovianization with approximate unitary designs, *Commun. Phys.* **4**, 127 (2021).
- [59] N. Dowling, P. Figueroa-Romero, F. A. Pollock, P. Strasberg, and K. Modi, Relaxation of multitime statistics in quantum systems (2021), [ArXiv:2108.07420](https://arxiv.org/abs/2108.07420).
- [60] N. Dowling, P. Figueroa-Romero, F. A. Pollock, P. Strasberg, and K. Modi, Equilibration of non-Markovian quantum processes in finite time intervals (2021), [ArXiv:2112.01099](https://arxiv.org/abs/2112.01099).
- [61] S. Milz, C. Spee, Z.-P. Xu, F. Pollock, K. Modi, and O. Gühne, Genuine multipartite entanglement in time, *SciPost Phys.* **10**, 141 (2021).
- [62] M. Ringbauer, F. Costa, M. E. Goggin, A. G. White, and A. Fedrizzi, Multi-time quantum correlations with no spatial analog, *npj Quantum Inf.* **4**, 1 (2018).
- [63] C. Giarmatzi and F. Costa, Witnessing quantum memory in non-Markovian processes, *Quantum* **5**, 440 (2021).
- [64] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Informational derivation of quantum theory, *Phys. Rev. A* **84**, 012311 (2011).
- [65] F. A. Pollock and K. Modi, Tomographically reconstructed master equations for any open quantum dynamics, *Quantum* **2**, 76 (2018).
- [66] M. R. Jørgensen and F. A. Pollock, Discrete memory kernel for multitime correlations in non-Markovian quantum processes, *Phys. Rev. A* **102**, 052206 (2020).
- [67] S. Milz, F. A. Pollock, and K. Modi, Reconstructing non-Markovian quantum dynamics with limited control, *Phys. Rev. A* **98**, 012108 (2018).
- [68] S. Aaronson, Quantum computing and hidden variables, *Phys. Rev. A* **71**, 032325 (2005).
- [69] X. Ni and M. Van den Nest, Commuting quantum circuits: Efficiently classical simulations versus hardness results, *Quantum Inf. Comput.* **13**, 54 (2013).
- [70] A. Y. Kitaev, in *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, 2022), p. 1308.
- [71] S. Parker and M. B. Plenio, Efficient Factorization with a Single Pure Qubit and $\log n$ Mixed Qubits, *Phys. Rev. Lett.* **85**, 3049 (2000).
- [72] S. Beauregard, Circuit for Shor's algorithm using $2n + 3$ qubits (2002), [ArXiv:quant-ph/0205095](https://arxiv.org/abs/quant-ph/0205095).
- [73] M. Mosca and A. Ekert, in *NASA International Conference on Quantum Computing and Quantum Communications* (Springer, 1998), p. 174.
- [74] R. Romano and D. D'Alessandro, Incoherent control and entanglement for two-dimensional coupled systems, *Phys. Rev. A* **73**, 022323 (2006).
- [75] D. Burgarth, S. Bose, C. Bruder, and V. Giovannetti, Local controllability of quantum networks, *Phys. Rev. A* **79**, 060305 (2009).
- [76] L. Banchi, D. Burgarth, and M. J. Kastoryano, Driven Quantum Dynamics: Will It Blend?, *Phys. Rev. X* **7**, 041015 (2017).
- [77] D. Burgarth, K. Maruyama, M. Murphy, S. Montangero, T. Calarco, F. Nori, and M. B. Plenio, Scalable quantum computation via local control of only two qubits, *Phys. Rev. A* **81**, 040303 (2010).
- [78] R. Heule, C. Bruder, D. Burgarth, and V. M. Stojanović, Local quantum control of Heisenberg spin chains, *Phys. Rev. A* **82**, 052333 (2010).
- [79] D. Gottesman, The Heisenberg representation of quantum computers (1998), [ArXiv:quant-ph/9807006](https://arxiv.org/abs/quant-ph/9807006).
- [80] D. Shepherd and M. J. Bremner, Temporally unstructured quantum computation, *Proc. R. Soc. A: Math. Phys. Eng. Sci.* **465**, 1413 (2009).
- [81] R. Jozsa and N. Linden, On the role of entanglement in quantum-computational speed-up, *Proc. R. Soc. London Ser. A: Math. Phys. Eng. Sci.* **459**, 2011 (2003).
- [82] V. V. Shende, S. S. Bullock, and I. L. Markov, in *Proceedings of the 2005 Asia and South Pacific Design Automation Conference* (2005), p. 272.
- [83] A. M. Krol, A. Sarkar, I. Ashraf, Z. Al-Ars, and K. Bertels, Efficient decomposition of unitary matrices in quantum circuit compilers, *Appl. Sci.* **12**, 759 (2022).
- [84] H. Weimer, A. Kshetrimayum, and R. Orús, Simulation methods for open quantum many-body systems, *Rev. Mod. Phys.* **93**, 015008 (2021).
- [85] M. J. Bremner, A. Montanaro, and D. J. Shepherd, Achieving quantum supremacy with sparse and noisy commuting quantum computations, *Quantum* **1**, 8 (2017).
- [86] S. Shrapnel, F. Costa, and G. Milburn, Quantum Markovianity as a supervised learning task, *Int. J. Quantum Inf.* **16**, 1840010 (2018).
- [87] C. Guo, K. Modi, and D. Poletti, Tensor-network-based machine learning of non-Markovian quantum processes, *Phys. Rev. A* **102**, 062414 (2020).
- [88] I. Luchnikov, S. Vintskevich, D. Grigoriev, and S. Filippov, Machine learning of Markovian embedding for non-Markovian quantum dynamics (2019), [ArXiv:1902.07019](https://arxiv.org/abs/1902.07019).
- [89] I. Luchnikov, S. Vintskevich, D. Grigoriev, and S. Filippov, Machine Learning Non-Markovian Quantum Dynamics, *Phys. Rev. Lett.* **124**, 140502 (2020).
- [90] T. Guhr, A. Müller-Groeling, and H. A. Weidenmüller, Random-matrix theories in quantum physics: Common concepts, *Phys. Rep.* **299**, 189 (1998).

- [91] R. Blümel and U. Smilansky, Random-Matrix Description of Chaotic Scattering: Semiclassical Approach, *Phys. Rev. Lett.* **64**, 241 (1990).
- [92] J. French, V. Kota, A. Pandey, and S. Tomsovic, Statistical properties of many-particle spectra V. Fluctuations and symmetries, *Ann. Phys. (NY)* **181**, 198 (1988).
- [93] C. W. Beenakker, Random-matrix theory of quantum transport, *Rev. Mod. Phys.* **69**, 731 (1997).
- [94] A. Bouland, B. Fefferman, Z. Landau, and Y. Liu, in *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, 2022), p. 1308.
- [95] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, and D. A. Buell, *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [96] M. Pozniak, K. Zyczkowski, and M. Kus, Composed ensembles of random unitary matrices, *J. Phys. A: Math. Gen.* **31**, 1059 (1998).
- [97] B. Villalonga, S. Boixo, B. Nelson, C. Henze, E. Rieffel, R. Biswas, and S. Mandrà, A flexible high-performance simulator for verifying and benchmarking quantum circuits implemented on real hardware, *npj Quantum Inf.* **5**, 86 (2019).
- [98] K. Noh, L. Jiang, and B. Fefferman, Efficient classical simulation of noisy random quantum circuits in one dimension, *Quantum* **4**, 318 (2020).
- [99] A. M. Dalzell, N. Hunter-Jones, and F. G. Brandão, Random quantum circuits transform local noise into global white noise (2021), arXiv preprint [ArXiv:2111.14907](https://arxiv.org/abs/2111.14907).
- [100] Y. V. Fyodorov and H.-J. Sommers, Statistics of resonance poles, phase shifts and time delays in quantum chaotic scattering: Random matrix approach for systems with broken time-reversal invariance, *J. Math. Phys.* **38**, 1918 (1997).
- [101] T. Morimae, K. Fujii, and J. F. Fitzsimons, Hardness of Classically Simulating the One-Clean-Qubit Model, *Phys. Rev. Lett.* **112**, 130502 (2014).
- [102] K. Fujii, H. Kobayashi, T. Morimae, H. Nishimura, S. Tamate, and S. Tani, Impossibility of Classically Simulating One-Clean-Qubit Model with Multiplicative Error, *Phys. Rev. Lett.* **120**, 200502 (2018).
- [103] J. Xuereb, S. Campbell, J. Goold, and A. Xuereb, DQC1 as an open quantum system (2022), [ArXiv:2209.03947](https://arxiv.org/abs/2209.03947).
- [104] M. Yoganathan and C. Cade, The one clean qubit model without entanglement is classically simulable (2019), [ArXiv:1907.08224](https://arxiv.org/abs/1907.08224).