Postselection-Free Learning of Measurement-Induced Quantum Dynamics

Max McGinley◉*

T.C.M. Group, Cavendish Laboratory, JJ Thomson Avenue, Cambridge CB3 0HE, United Kingdom

We address how one can empirically infer properties of quantum states generated by dynamics involving measurements. Our focus is on many-body settings where the number of measurements is extensive, making brute-force approaches based on postselection intractable due to their exponential sample complexity. We introduce a general-purpose scheme that can be used to infer any property of the postmeasurement ensemble of states (e.g., the average entanglement entropy, or frame potential) using a scalable number of experimental repetitions. We first identify a general class of "estimable properties" that can be directly extracted from experimental data. Then, based on empirical observations of such quantities, we show how one can indirectly infer information about any particular given nonestimable quantity of interest through classical postprocessing. Our approach is based on an optimization task, where one asks what are the minimum and maximum values that the desired quantity could possibly take, while ensuring consistency with observations. The true value of this quantity must then lie within a feasible range between these extrema, resulting in two-sided bounds. Narrow feasible ranges can be obtained by using a classical simulation of the device to determine which estimable properties one should measure. Even in cases where this simulation is inaccurate, unambiguous information about the true value of a given quantity realized on the quantum device can be learned. As an immediate application, we show that our method can be used to verify the emergence of quantum state designs in experiments. We identify some fundamental obstructions that in some cases prevent sharp knowledge of a given quantity from being inferred, and discuss what can be learned in cases where classical simulation is too computationally demanding to be feasible. In particular, we prove that any observer who cannot perform a classical simulation cannot distinguish the output states from those sampled from a maximally structureless ensemble.

I. INTRODUCTION

In quantum mechanics, measurements serve both as a means to extract information about a system, and as a form of dynamics in of themselves. Not only does the outcome of a measurement provide information about the premeasurement state, but it is also used to update one's knowledge of the postmeasurement state, effectively leading to a stochastic "collapse" of the wave function. This effect is central to a number of longstanding ideas in quantum information science, including active error correction [1] and measurement-based quantum computation [2,3]. In recent years, a great deal of interest has emerged in the study of many-body quantum states that are generated by such dynamics, leading to the discovery of measurement-induced entanglement phase transitions [4–13], emergent quantum state designs [14–18], and protocols for generating long-ranged entangled states via nonunitary dynamics [19–25].

The probabilistic nature of quantum measurements makes probing such phenomena in experiment a considerable challenge. This is because the states of interest cannot be prepared deterministically; rather, in each repetition of the experiment, we obtain a different randomly chosen outcome, and hence a different postmeasurement state. Using conventional learning techniques, any property of a given quantum state can only be inferred through repeated preparation and measurement, which in this context would only be possible if we run the experiment sufficiently many times such that each state is realized on multiple occasions. Such a "postselection"-based approach has a sample complexity that is exponential in the number of measurements [16,26], which is infeasible for many-body systems.

In this paper, we introduce a method by which properties of the postmeasurement ensemble of states can be learned from experimental data, without suffering from the exponential cost of postselection. Although a given ensemble-averaged quantity of interest (e.g., the mean entanglement entropy of the conditional states) may not

be directly accessible in the sense one usually thinks of, we show that information about its value can still be inferred indirectly based on independent observations of certain other quantities, which we call "estimable properties." These latter quantities are constructed such that they can be directly computed using data obtained from a scalable number of experimental repetitions.

The basis of our method is to propose the following optimization task: what are the maximum and minimum values that the quantity of interest could take, based on the empirically observed values of a set of estimable properties? [See Fig. 1(c).] These extrema provide us with two-sided bounds for the desired property, i.e., we can conclude that the true value lies somewhere within this range. Despite the extremely high-dimensional nature of this optimization problem (scaling with the number of possible measurement outcomes), we show using analytical arguments how concrete bounds can be efficiently computed.

The success of our scheme—as quantified by the gap between minimum and maximum—depends on which estimable properties we choose to measure. We propose to use classical simulations of the quantum device to inform this decision. That is, when we run the experiment and obtain a particular measurement outcome, a corresponding simulation of the dynamics is run on a classical device to determine what the best observable to measure is, given that the outcome in question has occurred. As the fidelity of

the simulation improves, the bounds constructed using our method become tighter. We benchmark the performance of the method for particular representative examples, demonstrating that tight bounds can be constructed even when there is appreciable mismatch between the simulated and true states.

This usage of a parallel simulation of the dynamics can be thought of as analogous to constructing "quantum-classical correlators," which have been recently introduced in the context of measurement-induced dynamics [13,27–32] as an alternative to feedback-based approaches [12,33]. In contrast to those works, where the inference one makes is dependent on how closely the quantum and classical devices behave, our approach allows one to extract unambiguous information regarding properties that are intrinsic to the quantum device alone. While bounds of this kind have been proved for specific quantities on a case-by-case basis in Ref. [34], our method can in principle be used to estimate *any* ensemble-averaged quantity, with a guarantee that the sharpness of the bounds is optimal for a given set of estimable parameters. We stress that even though model-based simulations are employed, the inferences we make are not contingent on any assumptions regarding the accuracy of this simulation, thereby allowing one to definitively verify whether or not some phenomenon of interest is actually realized in experiment. As a concrete example, we show how the method introduced here can be used
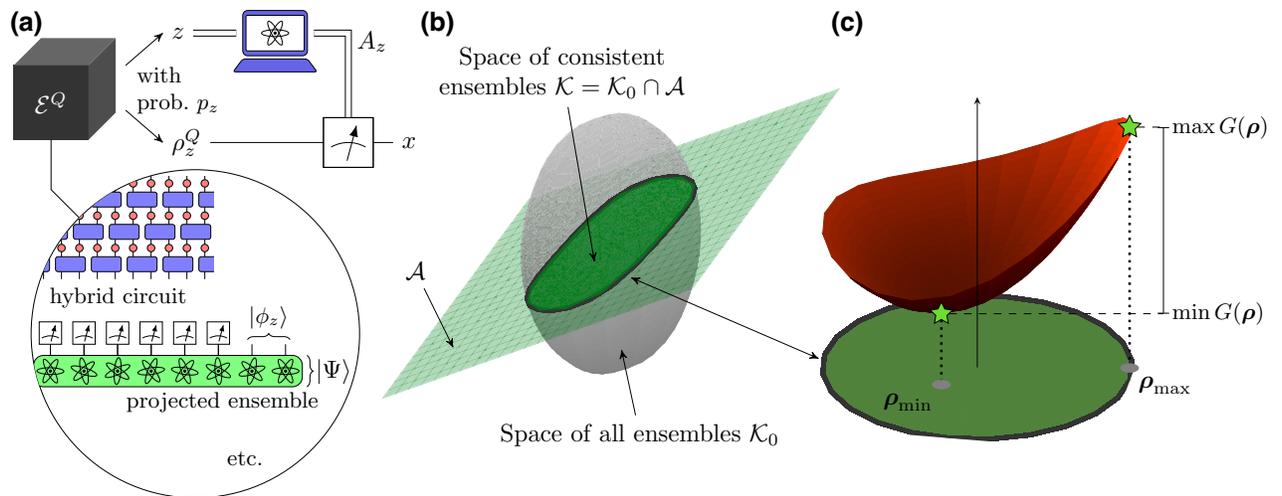
FIG. 1. (a) In this work, we are interested in ensembles of quantum states $\mathcal{E}^Q$ [Eq. (1)] that arise from some quantum dynamics featuring measurement. This could be a hybrid quantum circuit, the projected ensemble, or some other protocol of interest. To be general, we can think of the device as a black box that on a given run of the experiment, with probability p_z, outputs a label z (representing the outcome of all measurements during the dynamics), along with a corresponding quantum state ρ_z^Q (the postmeasurement conditional state). To learn properties of the ensemble, we perform a subsequent measurement on the state, here denoted A_z, which itself can depend on z; after many repetitions this allows us to infer properties of the form (8). A classical simulation of the underlying dynamics can be used to help inform us how to best choose A_z (Sec. IV). (b) Having learned the values of these measurable properties, we can characterize the space of all quantum state ensembles (collections of conditional states $\boldsymbol{\rho} = \{\rho_z^Q\}$) that are consistent with our findings $\boldsymbol{\rho} \in \mathcal{K}$—see Sec. V. (c) For a given ensemble property of interest $\bar{G}$ [Eq. (3)], we can construct minimum and maximum values (green stars) over all ensembles of the empirically deduced set $\mathcal{K}$. We can infer that the true value must lie between these extrema.

to verify the formation of quantum state designs in the projected ensemble [14,17]—a class of states with potential applications in state tomography, benchmarking, and cryptography [35–39].

By analyzing the nature of the information that one gains from performing experiments in general terms, we also identify certain fundamental limitations that in some cases preclude arbitrarily sharp knowledge about a certain property from being known, even in principle. In particular, we prove a result (Theorem 1) that implies that if the conditional postmeasurement states realized by the device are not close to being pure on average, then there will always be some residual uncertainty in the value of the desired quantity, even if the classical simulation is perfect. We also consider the possibility that classical simulations of the device may not be possible due to having too high a computational complexity. Our conclusion is that, without the ability to perform some form of simulation, nothing can be learned about the ensemble of quantum states, besides the averaged state generated by the device (Theorem 2). Put simply, this suggests that we should only think of the postmeasurement conditional quantum states as being physically accessible if we have some means to predict something about their structure in advance. Altogether, our results establish a fundamental separation between what can and cannot be learned about measurement-induced dynamics, and pave the way towards an understanding of how the various phenomena that arise in this context can be leveraged for other purposes, be it quantum communication, cryptography, or computation.

II. SUMMARY OF MAIN RESULTS

A. Setup

Our aim is to understand how one can infer properties of postmeasurement quantum states from experimental data. Specifically, we consider settings where a quantum system Q is subjected to some dynamics featuring measurements, and we are interested in properties of the conditional states of the system at the end of the dynamics. Scenarios that fit into this category include (but are not limited to) the following.

(1) Hybrid quantum circuits: models of dynamics typically defined in discrete time featuring measurements at various points in space and time [4–8,10, 11].
(2) The projected ensemble, where a fixed many-body state is prepared before measurements of a subset of degrees of freedom are made [14–17,25].
(3) Continuously monitored quantum systems [8,40–45], and settings where one is interested in the quantum jump trajectories of open quantum systems [46,47].

(4) Systems undergoing error correction and detection, which involves measuring stabilizer operators [1].
(5) Measurement-based approaches to quantum computation [2,3].

Since measurements are fundamentally stochastic processes, the conditional state of the system after measurements have occurred is itself a random variable. Thus, in contrast to familiar scenarios where the task is to learn about a fixed quantum state ρ that can be prepared deterministically, we are instead concerned with *statistical ensembles* of quantum states

$$\mathcal{E}^Q = \{(p_z, \rho_z^Q)\}_{z \in \mathcal{Z}}. \tag{1}$$

In the above, each $z \in \mathcal{Z}$ is a possible measurement outcome, which occurs with probability p_z, and ρ_z^Q is the (normalized) state of the system conditioned on this outcome—a unit-trace, Hermitian, positive semidefinite operator.

Such ensembles of states can be thought of as an abstracted description of any of the abovementioned examples. We remain indifferent to the exact nature of the underlying dynamics, and instead picture a scenario where in each run of an experiment, some oracle (i.e., a "black box," the implementation of which we disregard) provides us with a label z (the measurement outcomes), sampled from the probability distribution p_z, along with a corresponding conditional state ρ_z^Q. In each run of the actual experiment, the dynamics of interest is executed, which we can think of as a single query to this oracle, after which we can apply some additional measurement whose purpose is to extract information about ρ_z^Q [see Fig. 1(a)].

If the measurement outcome z is ignored then this oracle is equivalent to a device that prepares the averaged state

$$\langle \rho^Q \rangle := \sum_z p_z \rho_z^Q \tag{2}$$

every time. Properties of $\langle \rho^Q \rangle$ can therefore be learned using conventional approaches. In some settings, we may be concerned with properties of this averaged state, while in others—including many of the examples listed above—the physics being investigated may be manifest in the individual conditional states that make up ensemble (1). In the latter case, our wish is to learn properties of this ensemble $\mathcal{E}^Q$ beyond those of the average state, using some fixed number of queries M.

A well-appreciated issue that makes this objective difficult to achieve in the context of measurement-induced dynamics is the *postselection problem*, which we describe in detail later. In brief, the problem stems from the fact that in each run of the experiment, we only get a single copy of the conditional state ρ_z^Q, which is sampled randomly from the distribution p_z. In the regime where the

number of outcomes $|\mathcal{Z}|$ is large (which is to be expected in many-body settings), the state we get will be different every time for any reasonable number of repetitions M. Evidently, we can immediately rule out the possibility of learning properties of any individual conditional state ρ_{z*}^Q, since the probability of this state never occurring is high. One might still hope to be able to estimate ensemble-averaged properties from a finite sample $\{z^{(1)}, \ldots, z^{(M)}\}$, i.e., we look to estimate quantities of the form

$$\mathbb{E}_z[G] := \sum_{z \in \mathcal{Z}} p_z G(\rho_z^Q), \qquad (3)$$

where $G(\sigma)$ is some function of a density matrix σ, and $\mathbb{E}_z$ denotes an expectation value over the distribution $\{p_z\}$. However, in contrast to classical physics, quantum states cannot be copied, and thus having only single-copy access to the conditional states limits the information we can extract about each ρ_z^Q. In particular, as has been discussed previously, there is an apparent obstacle to learning properties of the kind (3) where $G(\sigma)$ is a nonlinear function of σ, since these cannot be expressed as functions of the average state $\langle \rho^Q \rangle$. A central aim of this paper is to critically examine this expectation, which is usually presented in somewhat heuristic terms, and to sharply determine exactly what can and cannot be inferred about a postmeasurement quantum state ensemble from experimental data of a reasonable size M.

B. Results and structure of the paper

Our first step to determine if and how the postselection problem can be circumvented is to establish a general class of ensemble properties that can be directly estimated using a reasonable number of repetitions M. In Sec. III, we demonstrate that expressions of the form (8) constitute such "estimable properties" of the ensemble, in that one can construct a function of the experimental data that equals the property in question on average, without any additional assumptions being made.

Most properties of the ensemble that are of interest do not fall within this class, and hence cannot be directly estimated in the same way. Nevertheless, we demonstrate how information about some nonestimable property can be indirectly inferred using independent measurements of other estimable quantities, using the following logic. Given a collection of estimable quantities and some empirical observations of their values (which we get from running the experiment), we consider the space of all ensembles that are consistent with these observations. We refer to this as the feasible space $\mathcal{K}$ [Fig. 1(b)], and we can guarantee that the true ensemble realized in the experiment lies somewhere within $\mathcal{K}$. We can characterize the best possible state of knowledge about some nonestimable average $\mathbb{E}_z[G]$ by looking at the extremes of this quantity over

the space $\mathcal{K}$. The maximum and minimum possible values of $\mathbb{E}_z[G]$ that are consistent with our observations can be represented as the solutions to an *optimization problem* [Fig. 1(c)]. By solving these optimization problems, we can construct two-sided bounds for the desired quantity, i.e., we infer that $\mathbb{E}_z[G]$ must be between the minimum and maximum, both of which can be computed efficiently using a scalable number of repetitions M. Ideally, the upper and lower bounds constructed using this approach will be close to one another, thereby giving us sharp knowledge about its value. This approach is outlined in detail in Sec. V, and we apply this idea to construct explicit bounds for various quantities of interest in Sec. VI.

We are naturally concerned with how successful this indirect inference scheme can be, as quantified by the width of the feasible range. To address this issue, one must first decide how to choose which estimable parameters to measure. While our analysis works for any such choice, we can make a decision based on an approach introduced in previous works, where one constructs cross-correlations between experimental data and an independent simulation of the underlying dynamics. These "quantum-classical correlators," which we describe in Sec. IV, fall within the set of measurable properties, and hence can be efficiently estimated. (The nature of these simulations need not matter for the purposes of our inference scheme, but we address some specifics in the discussion.)

Using this approach, we argue that the sharpness of the two-sided bounds depend on two key factors. Firstly, the accuracy of the simulation influences the gap between the minimum and maximum: naturally, as the simulation gets closer to the true behavior of the system, the bounds become tighter. Secondly, the nature of the conditional states that are realized on the quantum device ρ_z^Q also plays an important role. We prove an important result—Theorem 1—which states that regardless of which measurable quantities we compute, there will always be a consistent ensemble that is made up of states that are almost all pure. The implication is that, when the actual states realized by the device ρ_z^Q are mixed, we cannot necessarily constrain the range of a desired property $\mathbb{E}_z[G]$ to be within an arbitrarily small window, even if the simulations we use are perfect. This represents a fundamental obstruction to learning mixed state ensembles without postselection, which we discuss in detail in Sec. VIII A.

As an immediate application of our results, we show how the inference scheme developed here can be used to verify the emergence of quantum state designs in chaotic many-body systems [14,17]. We show how one can constrain a quantity known as the frame potential, which can be used to quantify how far an ensemble is from being a k-design (namely, an ensemble whose kth moments coincide with those of the Haar ensemble [48]). In Sec. VII, we present a numerical simulation of an experiment that features both noise and miscalibrations in the Hamiltonian,

and show that despite these imperfections (which are inevitably present in any experiment), the frame potential can be determined to be within a reasonably narrow window, in the sense described above.

In practice, performing a simulation of the quantum device may be a computationally demanding task, and we discuss the feasibility of such simulations for various specific cases in Sec. VIII B. It is therefore natural to consider whether anything can be learned in the case where simulation is not possible. To address this question, we present a result—Theorem 2—the implication of which is that if we do not employ some sort of simulation of the dynamics (broadly defined) then we cannot learn anything about the ensemble of quantum states $\mathcal{E}^Q$ beyond properties of the averaged state (2) using a number of repetitions that scales polynomially with system size. To be specific, in this regime the true ensemble is indistinguishable from one where every conditional state ρ_z^Q is replaced by the averaged state (2). This suggests that an inability to simulate the device in question renders the conditional states inaccessible in experiment. We conclude by discussing this point, along with some of the other broader implications of our results in Sec. IX.

III. POSTSELECTION PROBLEM: WHAT CAN AND CANNOT BE MEASURED

A. The no-coincidence regime

In all the scenarios captured by our generalized setup, a natural task is to infer properties of the ensemble $\mathcal{E}^Q$ using some kind of learning scheme. In particular, for the purposes of this work, we will be interested in estimating averaged properties of the states in the ensemble, i.e., quantities that can be expressed in the form (3).

We start by addressing the postselection problem in detail. When it comes to learning properties of postmeasurement quantum states from experimental data, a fundamental difficulty arises when the number of states in the ensemble $|\mathcal{Z}|$ is large—this is typically the case in the many-body setting, since $|\mathcal{Z}|$ scales exponentially with the number of measurements made, which itself typically scales with system size and/or time. This places us in a regime where, for any reasonable number of experimental repetitions M, the probabilities will satisfy $p_z M \ll 1$, meaning that we typically only get access to at most one instance of each state ρ_z^Q over the whole experiment.

The significance of this "no-coincidence" regime can be appreciated relatively straightforwardly: if we are given a single copy of a given quantum state ρ_z^Q then whatever measurement we subsequently perform on it, the distribution of outcomes will depend linearly on the density matrix ρ_z^Q. Hence, if there are no coincidences (no value of z occurs twice or more) then regardless of how we process the data, the only quantities that we can infer from the observed distribution of measurement outcomes are those

that are themselves linear in ρ_z^Q. If we hastily apply this logic to quantities of the form (3), this forces us to restrict ourselves to functions of the form $G(\rho_z^Q) = \mathrm{Tr}[\rho_z^Q A]$ for some observable A. In this case, we write the property in question as

$$\langle A \rangle^Q := \mathbb{E}_z\{\mathrm{Tr}[\rho_z^Q A]\} = \mathrm{Tr}[\langle \rho^Q \rangle A], \qquad (4)$$

where $\langle \rho^Q \rangle$ is the average state (2) Such quantities evidently give us no information about the nature of individual states in the ensemble, and we only learn about the average state (2).

In contrast, averages of nonlinear functions of the ensemble states, e.g., squared expectation values $G(\rho_z^Q) = \mathrm{Tr}[A\rho_z^Q]^2$, do contain information beyond that contained in the average state, which is why these are the quantities that are of relevance to the various problems described in Sec. II A. If we had access to multiple copies of each state ρ_z^Q then we could in principle learn such nonlinear functions by looking at the full distribution of measurement outcomes for each z separately. However, in the regimes we are interested in this demands a prohibitively large number of experimental repetitions. Our aim is to find a solution to this postselection problem while keeping the query complexity bounded.

B. Measurable quantities

To get a more precise picture of exactly what quantities are or are not experimentally accessible in the no-coincidence limit, let us consider a general procedure that can be used to extract information about the ensemble $\mathcal{E}^Q$. In a given repetition $r \in \{1, \ldots, M\}$, we obtain a sample from the ensemble $z^{(r)}$, and subsequently apply some (generalized) measurement to the quantum system Q, which can in principle depend on the outcome $z^{(r)}$. This z-dependent measurement scheme can be represented by a positive operator-valued measure (POVM) $\mathcal{F}^Q(z) = \{F_x^Q(z) : x \in \mathcal{X}\}$, where $\mathcal{X}$ is a discrete set of measurement outcomes, and the operators $F_x^Q(z)$ are Hermitian positive semidefinite, satisfying $\sum_{x \in \mathcal{X}} F_x^Q(z) = \mathbb{I}$ for each z. Conditioned on the outcome $z^{(r)}$, the result $x^{(r)} \in \mathcal{X}$ occurs with probability $\mathbb{P}(x^{(r)}|z^{(r)}) = \mathrm{Tr}[F_{x^{(r)}}^Q(z^{(r)}) \cdot \rho_{z^{(r)}}^Q]$, which is linear in $\rho_{z^{(r)}}^Q$, as discussed above, and together with the ensemble probabilities p_z this defines a joint probability distribution for the set of possible outcomes of a single run

$$\mathbb{P}(x, z) = p_z \, \mathrm{Tr}[F_x^Q(z)\rho_z^Q]. \qquad (5)$$

The full set of data that we acquire from the experiment $\{(x^{(r)}, z^{(r)}) : r = 1, \ldots, M\}$ corresponds to a set of M independent, identically distributed samples of pairs $X^{(r)} := (x^{(r)}, z^{(r)})$ drawn according to the probabilities (5). Most obviously, from such a sample we can estimate the average

of an arbitrary function of these pairs $f(X)$:

$$\overline{f(X)} := \mathbb{E}_X f(X) = \sum_{x,z} p_z \operatorname{Tr}[F_x^Q(z)\rho_z^Q] f(x,z) \quad (6)$$

with the overline used as shorthand for expectation values with respect to the samples X. The function $f(X^{(r)})$ is said to be an unbiased estimator of the quantity on the right-hand side of Eq. (6). We focus on the above quantities for the most part, since they are particularly relevant to ensemble averages of the form (3). However, if we wish to be even more general, we could also use the sampled data to estimate functions of $n \leq M$ independently sampled pairs $f_n(X_1, \ldots, X_n)$:

$$\overline{f_n} = \sum_{\{x_i\},\{z_i\}} \left[\prod_{i=1}^{n} p_{z_i} \operatorname{Tr}[F_{x_i}^Q(z_i)\rho_{z_i}^Q] \right] f_n[(x_1,z_1), \ldots, (x_n,z_n)]. \quad (7)$$

Quantities (6) and (7) are referred to as *estimable parameters* of distribution (5), because one can find a function of the sampled data $\{X^{(1)}, \ldots, X^{(M)}\}$ that is equal to these quantities in expectation [49,50]. In fact, any functional over the space of probability distributions that has an unbiased estimator must be expressible in the forms written above [51,52]. Hence, these are the only classes of observables that we can experimentally learn in the no-coincidence limit.

Returning to Eq. (6), we remark that the ensemble states ρ_z^Q only appear through the outcome probabilities (5), which, as discussed in the previous section, are linear in the density matrices. It is therefore helpful to rewrite the average (6) as

$$\overline{f(X)} = \langle A_z \rangle^Q := \sum_z p_z \operatorname{Tr}[A_z \rho_z^Q], \quad (8)$$

where we define the family of operators

$$A_z := \sum_x F_x^Q(z) f(x,z). \quad (9)$$

Equation (8) gives us a succinct characterization of the class of quantities that can be learned experimentally without prohibitive postselection overheads. [The more general estimators (7) can always be decomposed in terms of the above.] We immediately see that Eq. (4) is the special case where we disregard the classical information z when choosing the measurement scheme and postprocessing function f, i.e., $f(x,z) = f(x)$ and $F_x^Q(z) = F_x^Q$, such that A_z becomes z independent. Therefore, quite naturally, we conclude that in order to probe properties of the ensemble that cannot be characterized solely by the average density matrix (2), we must adopt a learning strategy that

itself depends explicitly on z. Notably, to do so necessarily requires us to have some *a priori* knowledge about the relationship between the labels z and the states ρ_z^Q. The essence of our scheme, which we describe in more detail in the following sections, is to use idealized classical computer simulations of the quantum device to inform us as to how A_z should depend on z.

C. Aside: avoiding midcircuit measurement and feedforward using classical shadows

As written, the measurable quantities (8) appear to require a feedforward mechanism, where the outcome z is used to decide what physical measurement to perform. Before describing our postselection-free inference scheme in detail, we briefly pause to explain how this aspect of the measurement procedure, which may be hard to implement in practice, can be avoided using ideas from classical shadow tomography [53]. A similar approach has been outlined in Ref. [34].

Without feedforward, we must fix the POVM $\mathcal{F}^Q(z)$ to be z independent. We make the key assumption that this fixed POVM is *informationally complete* [54], i.e., the collection of operators $\{F_x^Q\}_{x \in \mathcal{X}}$ span the full space of operators over $\mathcal{H}^Q$. In this case, one can find a (not necessarily unique) complementary set of operators $\tilde{F}_x^Q$ that overall have the property

$$\sum_x \operatorname{Tr}[F_x^Q \rho]\tilde{F}_x^Q = \rho \quad \text{for all } \rho \in \mathcal{B}(\mathcal{H}^Q). \quad (10)$$

Using the nomenclature of Ref. [54], the POVM operators F_x^Q constitute an operator frame, while $\tilde{F}_x^Q$ is the corresponding dual frame. With this construction in hand, any observable of the form (8) can be estimated from a set of experimentally measured data $\{(x^{(r)}, z^{(r)}): r = 1, \ldots, M\}$ by using the postprocessing function

$$f(x^{(r)}, z^{(r)}) = \operatorname{Tr}[A_{z^{(r)}} \tilde{F}_{x^{(r)}}^Q]. \quad (11)$$

Combining Eqs. (9) and (10), we see that the average of such a function equals the right-hand side of Eq. (8) in expectation,

$$\overline{f(x,z)} = \langle A_z \rangle^Q. \quad (12)$$

The above prescription (11) gives us an explicit way of computing an unbiased estimator for any quantity of the form (8). Notably, we can do this without deciding on the operators A_z in advance of the physical measurement of the system, i.e., we can "measure first, ask questions later" [55]. This feature of informationally complete POVMs means that we do not need to employ adaptive schemes, where the physical measurement procedure is decided based on the sample z.

A particularly straightforward way to implement an informationally complete POVM is to use classical shadow tomography [53,56], which was proposed as a useful way to study measurement-induced dynamics in Ref. [34]. In each run of the experiment, we apply a randomly chosen unitary U_c from a prechosen ensemble $\mathcal{U} = \{(q_c, U_c)\}_{c \in \mathcal{C}}$, where $\mathcal{C}$ is an arbitrary discrete set, before performing a projective measurement in a fixed basis $\{|m\rangle \langle m|\}$. For certain choices of $\mathcal{U}$, informational completeness is guaranteed, and the dual frame can be computed efficiently.

While there are many such possibilities, as a concrete example, one can take $\mathcal{U}$ to be a uniform distribution over all single-qubit Clifford gates (i.e., operations of the form $U_c = u_{c_1} \otimes \cdots \otimes u_{c_N}$, where $u_{c_i} \in \{I, H_X, H_Y\}$ with H, H_Y Hadamard and Y Hadamard gates, respectively). Then we can use the dual frame

$$\tilde{F}^Q_{(m,c)} = \bigotimes_{i=1}^{N_Q} (3 u^\dagger_{c_i} |m_i\rangle \langle m_i| u_{c_i} - I). \tag{13}$$

Thus, the necessary measurement scheme can be implemented using single-qubit rotations and measurements. In many settings, this means that we can perform all the measurements simultaneously at the end of the experiment (both those that generate the ensemble $\mathcal{E}^Q$ and those we use to extract information). This is illustrated for the case of the projected ensemble of a bipartite state $|\Psi^{AB}\rangle = U(t) |0^{\otimes N}\rangle$ in Fig. 2.

The number of repetitions M required to estimate $\langle A_z \rangle^Q$ to a given accuracy can be expressed in terms of the variance $\text{Var}_{x,z} f(x,z) = \sum_z p_z \text{Tr}[F^Q_x \rho^Q_z] \text{Tr}[\tilde{F}^Q_x A_z]^2$. In the more familiar setting where one wishes to learn the properties of a fixed state ρ, the variance of a shadow tomographic estimator can be bounded using a useful construction, the shadow norm $\|A\|^2_{\text{shadow}}$, which is a function of the observable to be estimated A, as well as the ensemble $\mathcal{U}$; this is defined in Ref. [53]. In the present case, where we are instead dealing with an ensemble of states, it is straightforward to show that $\text{Var}_{x,z} f(x,z)$ can be bounded as

$$\text{Var}_{x,z} f(x,z) \leq \sum_z p_z \|A_z\|^2_{\text{shadow}}. \tag{14}$$

Thus, the property $\langle A_z \rangle^Q$ can be estimated to a good accuracy using a reasonable number of experimental repetitions provided that each of the operators A_z is (with high probability) an operator that itself can be efficiently estimated in ordinary shadow tomography.

The randomized Pauli measurements discussed above constitute one particular example of an informationally complete POVM, but we emphasize that there are many other alternatives, e.g., those based on global Clifford rotations [53], or even generalized methods that use chaotic
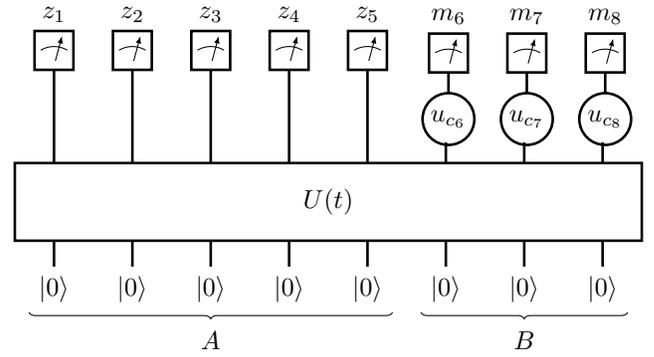
FIG. 2. Protocol for measuring arbitrary estimable properties (8) for the projected ensemble using classical shadows, with all measurements occurring simultaneously at the end of the circuit. The projected ensemble corresponds to the collection of post-measurement states that arise when the bipartite state $|\Psi^{AB}\rangle = U(t) |0^{\otimes N}\rangle$ is prepared, and then qubits in A are measured in the computational basis, with outcomes $\{z_i : i = 1, \ldots, N_A\}$. The postmeasurement states are thus defined on the qubits in B; this figure shows the case $N_A = 5, N_B = 3$. As described in the main text, a shadow-based scheme can be used to probe these states, which corresponds to performing random on-site unitaries u_{c_i} before measuring in the computational basis, with outcomes m_i (to be distinguished from z_i, which label the states we are trying to probe). The dual frame (13) is used when reconstructing estimable properties via Eq. (11).

Hamiltonian evolution and/or ancilla qubits [39,57]. In all the subsequent analysis, we remain agnostic to the exact measurement scheme used, and will simply assume that we have some way to measure the properties (8).

IV. CLASSICAL SIMULATIONS

We have now identified an explicit scheme for measuring quantities of the form $\langle A_z \rangle^Q$ [Eq. (8)] without problematic postselection overheads. However, as mentioned previously, our goal is to infer properties of the ensemble that take the form (3)—specifically, those for which $G(\rho)$ is a nonlinear function. Here we describe the idea of "quantum-classical correlators," where a simulation of the system in question is run in conjunction with the experiment, and describe what they can tell us about such nonlinear averages.

By "classical simulations," we mean the following: each time we perform a run of the experiment, which gives us a label $z \in \mathcal{Z}$ and a measurement outcome $x \in \mathcal{X}$, we also compute and store a representation of some corresponding state ρ^C_z on a classical computer. We use the superscript C to distinguish this "classical" state, which represents the result of some idealized simulation of the experiment, from the "quantum" state that is actually realized on the true device. We emphasize that this simulation can be done "lazily," i.e., we only compute ρ^C_z for those values of z that

happen to arise in the experiment, as opposed to precomputing every state ρ_z^C in advance, the cost of which would scale with $|\mathcal{Z}|$. We also do not need to classically sample from the distribution p_z, which in many cases is itself computationally demanding. The nature, accuracy, and computational cost of the classical simulation may depend on the specific physical setting being considered, and we discuss several particular examples in Sec. VIII B below. However, for the time being, we presume that some form of simulation is possible, and motivate our discussion on the basis that there is some partial correlation between the classical and quantum states, ρ_z^C and ρ_z^Q. The case where a classical simulation is impossible—either due to incomplete knowledge of how the quantum device operates, or prohibitively high computational cost—is discussed in Sec. VIII C below.

Let us take the simplest nontrivial case and suppose that our aim is to learn the average of a particular squared expectation value $\mathrm{Tr}[O\rho_z^Q]^2$ over the ensemble $\mathcal{E}^Q$. We introduce the following notation for such a quantity:

$$\langle O \otimes O \rangle^{QQ} := \sum_z p_z \,\mathrm{Tr}[O^{\otimes 2} \cdot (\rho_z^Q)^{\otimes 2}]$$
$$= \sum_z p_z \,\mathrm{Tr}[O\rho_z^Q]^2. \qquad (15)$$

The superscript QQ is used to emphasize that the quantity in question is a linear functional of the state $\rho_z^Q \otimes \rho_z^Q$. As explained in Sec. III, we cannot directly measure this quantity. However, a natural proxy that has been introduced in several recent works is the "quantum-classical correlator" [13,27–32], which here we define as

$$\langle O \otimes O \rangle^{QC} := \sum_z p_z \,\mathrm{Tr}[O\rho_z^Q]\,\mathrm{Tr}[O\rho_z^C]. \qquad (16)$$

Note that the probabilities p_z appearing in Eq. (16) are the same as those appearing in the fully quantum expression (15). The above quantity can then be cast in the form (8) with

$$A_z = \mathrm{Tr}[O\rho_z^C] \times O. \qquad (17)$$

Thus, once we collect samples $z^{(r)}$ taken on the quantum device, we can compute the corresponding classical observables $\mathrm{Tr}[O\rho_z^C]$ and construct an estimator of the quantum-classical correlator, using the methods prescribed in Sec. III. Finally, we can also consider a "classical-classical" correlator

$$\langle O \otimes O \rangle^{CC} = \sum_z p_z \,\mathrm{Tr}[O\rho_z^C]^2, \qquad (18)$$

defined by analogy to Eq. (15), again with the probabilities p_z set by the quantum device. This is also of the

form (8), with $A_z = \mathrm{Tr}[O\rho_z^C]^2 \times I$ (the quantum states are disregarded here).

Clearly, in the limit where the classical simulation perfectly matches the behavior of the quantum device, $\rho_z^C = \rho_z^Q$, all these quantities are equal. Thus, we hope that if the simulation is good, but not perfect, the proxy quantity (16) will be close in value to the true "quantum-quantum" observable (15), which is the physically relevant quantity. However, at present we cannot make any definitive conclusions about the value of the quantum-quantum correlator without making unsubstantiated assumptions about the accuracy of our classical simulations. Our objective in the following two sections, which form the most technical parts of this paper, is to establish methods that allow one to construct rigorous two-sided bounds for the true value (15) based on experimental observations of the measurable quantities (16) and (18), *without* making any *a priori* assumptions about how accurate the classical simulation is. That is, even though we are using our classical simulation as a form of prior "guess" for the conditional states ρ_z^Q, we allow for the possibility that this guess is incorrect. This skeptical approach to learning means that any conclusions we make about the ensemble will be entirely unambiguous.

V. CONVEX OPTIMIZATION APPROACH TO INFERRING AVERAGES

We have seen in detail how the no-coincidence limit gives rise to a distinction between properties of quantum state ensembles that can be measured straightforwardly—those of the form $\langle A_z \rangle_{\mathcal{E}^Q}$, Eq. (8), which include "quantum-classical" correlators (16)—versus those that cannot be directly measured with a reasonable number of experimental repetitions, e.g., the "quantum-quantum" correlator (15). In the absence of direct estimation schemes for the latter class, we are interested in understanding the best possible state of knowledge that we could in principle have about such unobservable quantities, based on experimentally accessible data. Our intuition, based on the discussion of the previous section, is that by cross-correlating experimental outputs with classical simulations of the quantum system, we can gain some amount of knowledge about these quantities, even though we cannot measure them directly. To make this intuition concrete, in this and the following sections, we aim to address the question: given knowledge of a set of observable quantities of the form (8), what range of values can a particular unobservable quantity take, while ensuring consistency with our observations?

More formally, suppose that from a set of experimental data, we construct estimates of a family of R scalar quantities $\{\langle A_z^{(i)} \rangle : i = 1, \ldots, R\}$, the outcomes of which we denote b_i. For the moment, we presume that any statistical uncertainty in these observations can be neglected (an assumption that we relax later). We wish to determine

the maximum and minimum possible values that a particular average $\mathbb{E}_z G(\rho_z^Q)$ can take over all quantum state ensembles $\mathcal{E}^Q$ that satisfy

$$\langle A_z^{(i)} \rangle_{\mathcal{E}^Q} = b_i \quad \text{for all } i \in \{1, \ldots, R\}. \tag{19}$$

Denoting the minimum and maximum values of $\mathbb{E}_z G(\rho_z^Q)$ as $g_\pm^*$, this will give us a two-sided bound for the desired average, namely,

$$g_-^* \leq \mathbb{E}_z G(\rho_z^Q) \leq g_+^*. \tag{20}$$

The determination of $g_\pm^*$ can be viewed as an *optimization task*, where the object being varied over is the ensemble $\mathcal{E}^Q$ itself, and the function being extremized is the average $\mathbb{E}_z G(\rho_z^Q)$. In this section, we study the structure of such optimization problems at a general level.

A. Set of consistent ensembles

Our first step is to analyze the structure of the space of quantum state ensembles that satisfy Eq. (19). For succinctness of notation, we find it useful to represent the collection of states $\{\rho_z^Q\}_{z \in \mathcal{Z}}$ in terms of a single large block-diagonal matrix $\boldsymbol{\rho} = \bigoplus_{z \in \mathcal{Z}} \rho_z^Q$, where each block contains the density matrix ρ_z^Q for a particular label z. We can then view $\boldsymbol{\rho}$ as an element of the linear space of matrices $\mathcal{M} := \mathcal{B}(\mathcal{H})^{\oplus |\mathcal{Z}|}$. As for the probabilities p_z, while these are not known to us in full in practice, our ability to run the experiment means that we can sample from this distribution; therefore, our approach will be to keep p_z fixed, while allowing the states ρ_z^Q themselves to vary. With all this in mind, we begin by formally specifying the space of valid quantum state ensembles as

$$\mathcal{K}_0 := \left\{ \boldsymbol{\rho} = \bigoplus_{z \in \mathcal{Z}} \rho_z^Q \;\middle|\; \rho_z^Q \in \mathcal{D} \text{ for all } z \in \mathcal{Z} \right\}, \tag{21}$$

where $\mathcal{D} \subset \mathcal{B}(\mathcal{H})$ denotes the space of density matrices for a single copy of the system Hilbert space $\mathcal{H}$.

We are interested specifically in ensembles that are consistent with the observations (19). Again, using the direct sum representation, for each block-diagonal matrix $\mathbf{X} = \bigoplus_z X_z \in \mathcal{M}$, we define the linear function $A^{(i)}(\mathbf{X}) := \sum_z p_z \text{Tr}[A_z^{(i)} X_z]$, and we define the space $\mathcal{A}$ as

$$\mathcal{A} := \{ \mathbf{X} \in \mathcal{M} \mid A^{(i)}(\mathbf{X}) = b_i \text{ for all } i = 1, \ldots, R \}. \tag{22}$$

Since $\mathcal{A}$ is a subspace defined by R linear constraints, $\mathcal{A}$ is a hyperplane of codimension R in $\mathcal{M}$.

Finally, the *feasible space* is given by the intersection $\mathcal{K} = \mathcal{K}_0 \cap \mathcal{A}$. Writing this out in full,

$$\mathcal{K} := \left\{ \boldsymbol{\rho} = \bigoplus_z \rho_z^Q \;\middle|\; \rho_z^Q \in \mathcal{D} \text{ for all } z \in \mathcal{Z}; \right.$$
$$\left. \mathbb{E}_z(\text{Tr}[A_z^{(i)} \rho_z^Q]) = b_i \right\}. \tag{23}$$

Each element of $\mathcal{K}$ corresponds to a particular ensemble that could describe the system, based on the empirical data (19).

Crucially, we note that, when viewed as a subset of the linear space $\mathcal{M}$, the set $\mathcal{K}$ is convex, i.e., if $\boldsymbol{\rho} = \bigoplus_z \rho_z^Q$ and $\boldsymbol{\sigma} = \bigoplus_z \sigma_z^Q$ are two sets of states that both belong to $\mathcal{K}$ then so too does

$$\lambda \boldsymbol{\rho} + (1 - \lambda) \boldsymbol{\sigma} = \bigoplus_{z \in \mathcal{Z}} [\lambda \rho_z^Q + (1 - \lambda) \sigma_z^Q] \in \mathcal{K} \tag{24}$$

for any $\lambda \in [0, 1]$. This is a consequence of the convexity of both the space of density matrices $\mathcal{D}$ and the hyperplane $\mathcal{A}$, along with the fact that the intersection of two convex sets is itself convex.

B. Convex functions

Concretely, our aim is to determine the range of *feasible values* that a particular average $\mathbb{E}_z G$ can take over the space of feasible ensembles $\mathcal{K}$. Namely, we wish to characterize the set

$$\mathcal{G} = \{ g \mid \exists \boldsymbol{\rho} \in \mathcal{K} \text{ such that } \bar{G}(\boldsymbol{\rho}) = g \}, \tag{25}$$

where we introduce the shorthand

$$\bar{G}(\boldsymbol{\rho}) := \sum_z p_z G(\rho_z^Q). \tag{26}$$

The convexity of $\mathcal{K}$ will prove useful for this purpose, in particular for determining the extremal feasible values $g_+^* = \max_{g \in \mathcal{G}} g$ (similar for the minimum g_-^*). Indeed, we can formulate such a task in terms of the following optimization problems (which we write out in longhand momentarily):

$$g_\pm^* := \max_{\bigoplus_z \rho_z^Q \in \mathcal{K}} / \min \left[\sum_z p_z G(\rho_z^Q) \right] \tag{27}$$

(with g_+^* corresponding to the maximum, and g_-^* to the minimum). One can then immediately use these extrema to bound the average for the true ensemble $\mathcal{E}^Q$ on both sides as

$$g_-^* \leq \bar{G}(\mathcal{E}^Q) \leq g_+^*. \tag{28}$$

This reflects the best possible state of knowledge we could have about the average $\bar{G}$, based on our observations.

Optimization problems over convex sets such as Eq. (27) have been well studied in a wide variety of contexts. Most prominently, much work has gone into the study of *convex optimization*, which is concerned with minimizing convex functions (equivalently, maximizing concave functions) over convex sets. Recall that a function $\bar{G}: \mathcal{K} \to \mathbb{R}$ is convex if, for any two $\boldsymbol{\rho}, \boldsymbol{\rho}' \in \mathcal{K}$ and any $\lambda \in [0, 1]$, we have

$$\bar{G}[\lambda \boldsymbol{\rho} + (1 - \lambda)\boldsymbol{\rho}'] \le \lambda \bar{G}(\boldsymbol{\rho}) + (1 - \lambda)\bar{G}(\boldsymbol{\rho}'). \quad (29)$$

Convex optimization tasks enjoy many useful properties, which can be exploited to gain analytical insight into their solutions, and to design efficient algorithms. For this reason, from hereon we specialize to cases where the function $\bar{G}(\boldsymbol{\rho})$ is convex, unless mentioned otherwise. From our construction of $\bar{G}$ [Eq. (26)], we see that $\bar{G}$ is convex if the function $G(\rho_z^Q)$, the average of which we are interested in, is itself a convex function over density matrices. Some particularly pertinent functions that arise in quantum many-body physics include the following.

(1) *Powers of expectation values.*—While expectation values of observables $\langle O \rangle = \text{Tr}[\rho O]$ are linear in ρ, and hence directly estimable without any additional analysis, often we are interested in integer powers of expectation values, $\text{Tr}[\rho O]^k$. The quantum-quantum correlator (15) is an example of this with $k = 2$. We would need to compute such a quantity if we wanted to know the variance of $\langle O \rangle$ over the ensemble $\mathcal{E}^Q$. We note that $\text{Tr}[\rho O]^k$ is convex for even $k \ge 2$, or for any positive integer when O is semipositive definite.

(2) *Entropies.*—Often we are interested in an entropy associated with a quantum state ρ, or a entropy of a subsystem of Q. The von Neumann entropy $S(\rho) = -\text{Tr}[\rho \log \rho]$ is a concave function of ρ, as are Rényi entropies $S_\alpha(\rho) = (1 - \alpha)^{-1} \log(\text{Tr}[\rho^\alpha])$ for $\alpha \in (0, 1)$ [58]. Since our subsequent analysis refers explicitly to convex functions, one can simply work with $-S(\rho)$, which is convex. One should then bear in mind that the role of minimization and maximization of the objective function are exchanged.

(3) *Purities.*—A closely related object is the purity $\text{Tr}[\rho^2]$ and generalizations to higher powers $\text{Tr}[\rho^k]$. These are equal to exponentials of Rényi entropies $S_{\alpha=k}(\rho)$, and hence give information on how mixed the states are, either globally or within a subsystem. It is straightforward to show that $\text{Tr}[\rho^k]$ is convex in ρ for $k \ge 1$.

While our specialization to convex functions may seem restrictive, we note that any function $G(\rho)$ whose Hessian (matrix of second derivatives) is bounded can be decomposed as a sum of a convex and a concave function [59],

and hence each component can be bounded separately using the methods described in the following.

C. Duality and certificates for convex optimization problems

Working with the understanding that $\bar{G}(\boldsymbol{\rho})$ is a convex function, we now describe our approach to finding the extrema $g_\pm^*$, or approximations thereof. Several standard techniques and results from the field of convex optimization will be used in the following; we refer the interested reader to Ref. [60] for an introduction to the field and proofs of various pertinent results.

The standard approach to optimization problems with equality constraints of the form (19) is to make use of Lagrange multipliers. For each constraint $i = 1, \ldots, R$, we introduce a scalar Lagrange multiplier λ_i, and define the Lagrange dual functions as [60]

$$h_-(\lambda_i) = \inf_{\boldsymbol{\rho} \in \mathcal{K}_0} L(\boldsymbol{\rho}, \lambda_i), \quad (30\text{a})$$

$$h_+(\lambda_i) = \sup_{\boldsymbol{\rho} \in \mathcal{K}_0} L(\boldsymbol{\rho}, \lambda_i), \quad (30\text{b})$$

where the Lagrangian $L(\boldsymbol{\rho}, \lambda_i)$ is given by

$$L(\boldsymbol{\rho}, \lambda_i) := \bar{G}(\boldsymbol{\rho}) - \sum_i \lambda_i [A^{(i)}(\boldsymbol{\rho}) - b_i]. \quad (31)$$

We emphasize that the domain in Eqs. (30) is $\mathcal{K}_0$, without the linear constraints (19). This is significant because $L(\boldsymbol{\rho}, \lambda_i)$ is a sum over z, and thanks to the structure of $\mathcal{K}_0$, we can perform the extremization for each block ρ_z^Q separately.

An important concept in the study of optimization problems is that of strong and weak dualities, which relate the original optimization task (the "primal problem") to a particular secondary task (the "dual problem"). The dual problem is to find the extreme points

$$h_-^* := \max_{\lambda_i} h_-(\lambda_i), \qquad h_+^* := \min_{\lambda_i} h_+(\lambda_i). \quad (32)$$

Since $h_\pm(\lambda_i)$ are, by definition, the pointwise supremum and infimum of a family of affine functions, they are concave and convex functions, respectively. The dual problems are therefore convex optimization problems, even if the primal is not.

Strong duality is the statement that the solutions to the primal and dual problems are identical. In our case, by Slater's condition [60], strong duality holds for the minimization problem, i.e., $h_-^* = g_-^*$, provided that $\bar{G}(\boldsymbol{\rho})$ is convex. As for the maximization problem, this is not a convex optimization problem, since it is equivalent to minimizing $-\bar{G}$, which is concave by assumption. There is still some useful structure exhibited by convex maximization problems, which we discuss in Sec. V D below, but

for the time being, we instead rely on weak duality, which holds irrespective of the nature of $\bar{G}$. Weak duality states that $h_-^* \leq g_-^*$ and $h_+^* \geq g_+^*$. This general property can be proved using the min-max inequality.

While many efficient algorithms exist that allow one to solve convex optimization problems numerically, the high dimensionality of the primal problem (as well as the stochastic nature of the observations we make) mean that we cannot directly employ such algorithms as they stand. One option is to apply such methods to solve the dual problem, which has dimensionality R, rather than $|\mathcal{Z}| \times d^2$, thus making the problem more manageable. While this numerical approach is perfectly feasible, here we choose to study these problems analytically, in order to gain more insight into the structure of these problems. While exact solutions to the optimization problems are not always obtainable using analytic methods, we can still invoke strong or weak duality, which allow us to find *primal-dual certificates*: that is, even if we only have an approximate solution $\{\lambda_i\}$ to the dual problem, we always have the following series of inequalities:

$$h_-(\lambda_i) \leq h_-^* = g_-^*, \qquad h_+(\lambda_i) \geq h_+^* \geq g_+^*. \qquad (33)$$

Thanks to these relations, if we can evaluate $h_-(\lambda_i)$ for some set of Lagrange multipliers, we can certify that the true minimum is no less than $h_-(\lambda_i)$. Thus, even if we cannot determine the exact value of the extrema, we can use suboptimal solutions of the dual problem to obtain conservative estimates of $g_\pm^*$, i.e., a rigorous bound that is guaranteed *not* to be an overestimate (underestimate) of the minimum (maximum). In lieu of an exact solution to either the primal or dual problems, our aim is then to find a way of obtaining $h_\pm(\lambda_i)$ for near optimal λ_i, so as to obtain as tight a bound as possible.

By treating these optimization problems analytically, rather than numerically, we gain valuable insight that informs us how to choose the operators $A_z^{(i)}$ from the very beginning so as to obtain small feasible ranges; we have this in mind in Sec. VI below and beyond.

D. Minimization versus maximization

Before we begin the process of constructing explicit bounds for specific observables, it is important to point out a fundamental difference between minimization versus maximization problems for a given convex average $\bar{G}(\boldsymbol{\rho})$. Namely, the former is a convex optimization problem (minimizing a convex function over a convex set), while the latter is not. This difference between the two will turn out to have important consequences for the tightness of the bounds that one can infer based on experimental observations.

While the maximization problem cannot be cast into a standard convex optimization problem, we can still invoke a useful property that results from its special

structure: the maximum of a convex function over a convex set is always attained at at least one extreme point of the set. Recall that an extreme point of a convex set $\mathcal{C}$ is an element $\tau \in \mathcal{C}$ that cannot be expressed as a nontrivial convex combination of two other elements. Specifically, τ is extreme if and only if $\tau = \lambda \tau' + (1 - \lambda)\tau''$ for some $\lambda \in (0, 1)$ implies that $\tau' = \tau'' = \tau$ [60]. If $G(\tau)$ is a convex function over $\mathcal{C}$ then, for any nonextreme τ, we can find appropriate elements $\tau', \tau'' \neq \tau$ such that $G(\tau) = G[\lambda \tau' + (1 - \lambda)\tau''] \leq \lambda G(\tau') + (1 - \lambda)G(\tau'') \leq \max[G(\tau'), G(\tau'')]$, which proves the claim stated above.

By considering the eigenstate decomposition of a density matrix ρ, one can see that the extreme points of the space $\mathcal{D}$ are pure states $\rho = |\phi\rangle \langle \phi|$. This structure is naturally reflected in the extreme points of $\mathcal{K}$. In Appendix A, we prove the following result.

Theorem 1. Any ensemble $\boldsymbol{\rho}$ that is an extreme point of the set $\mathcal{K}$ [Eq. (23)] has at most R states that are not globally pure, where R is the number of linear constraints in Eq. (19).

Since $\mathcal{K}$ is necessarily nonempty (the true ensemble $\mathcal{E}^Q$ lies within $\mathcal{K}$), and is a closed, linearly bounded subset of $\mathcal{M} = \mathcal{B}(\mathcal{H})^{\oplus|\mathcal{Z}|}$, at least one extreme point must exist, and hence we immediately have the following result.

Corollary 1. Given the values of R scalar measurable properties (19), there exists an ensemble consistent with these observations for which at most R states are globally nonpure.

As a result, in the no-coincidence limit, where $Rp_z \ll 1$, we can infer that the solution to the maximization problem is extremely close to the solution of the same problem with the additional constraint that all states are pure. This fact imposes fundamental limitations on how wide the feasible range (20) can be made in cases where the actual state ensemble $\mathcal{E}^Q$ is significantly mixed, $\sum_z p_z \operatorname{Tr}[(\rho_z^Q)^2] < 1$. We consider this issue in more detail in Sec. VIII A below.

VI. CONSTRUCTING TWO-SIDED BOUNDS

Having outlined the general structure of our optimization-based approach in the previous section, we can now consider a range of different averaged properties that are of particular physical relevance, and construct explicit upper and lower bounds for each. Because the minimization and maximization problems have distinct characters, it is helpful to consider the two separately for each observable. The collection of quantities considered here is by no means exhaustive, and analogous bounds can be constructed for other observables, but the examples we choose encompass a broad range of quantities that are pertinent to measurement-induced dynamics.

TABLE I.   Examples of ensemble-averaged quantities for which we derive upper and lower bounds, and references to the relevant inequalities derived below. Here, Q_1 denotes a subsystem of Q, $\mathcal{N}$ is an arbitrary linear map over the space of operators, and the notation $\mathcal{N} \succeq 0$ indicates that this map is semipositive definite, i.e., $\langle\!\langle X|\mathcal{N}|X\rangle\!\rangle \geq 0$ for all operators X. In some cases, two bounds are quoted, one of which is simple to evaluate, with the other being more versatile and amenable to optimization. The asterisk indicates that the bound is vacuous.

Quantity	Expression	Concave or convex?	Lower bound	Upper bound
Global purity	$\mathbb{E}_z\{\mathrm{Tr}[(\rho_z^Q)^2]\}$	Convex	Eqs. (38), (45)	Eq. (46)*
Subsystem purity or quadratic observable	$\mathbb{E}_z\{\mathrm{Tr}[\rho_z^Q O]^2\}$ or $\mathbb{E}_z\{\mathrm{Tr}[\rho_z^Q \mathcal{N}(\rho_z^Q)]\}$, $\mathcal{N} \succeq 0$	Convex	Eq. (49)	Eq. (52)
von Neumann entropy	$\mathbb{E}_z\{-\mathrm{Tr}[\rho_z^{Q_1} \log \rho_z^{Q_1}]\}$	Concave	Eqs. (63), (64)	Eqs. (56), (58)
Frame potential	$\mathbb{E}_{z,z'} \mathrm{Tr}[\rho_z^Q \rho_{z'}^Q]^k$, $k \in \mathbb{N}$	Convex	Eq. (67a)	Eq. (67b)

The logical arguments used to derive bounds for each quantity follow much the same pattern, and so after deriving the first several cases, we simply quote the remainder of our results, leaving the proofs to Appendix B. Readers who are mainly concerned with the results of our calculations, rather than the detailed derivations, may skip the bulk of this section, and instead consult Table I, where references to specific bounds are listed.

A. Global purity lower bound

One of the simplest possible nonlinear averages that we can consider is the averaged global purity, $G(\rho_z^Q) = \mathrm{Tr}[(\rho_z^Q)^2]$ (to be distinguished from the purity of a subsystem of Q, which we treat in Sec. VI C). As mentioned above, this is a convex function of ρ_z^Q, and hence the minimization problem can be solved using convex optimization techniques. The dual function (30a) for this problem can be formally defined as

$$h_-(\lambda_i) = \inf_{\rho \in \mathcal{K}_0} \sum_z p_z \left(\mathrm{Tr}[(\rho_z^Q)^2] - \sum_i \lambda_i \mathrm{Tr}[\rho_z^Q A_z^{(i)}] \right) + \sum_i \lambda_i b_i. \tag{34}$$

By virtue of the product structure of $\mathcal{K}_0 \equiv \mathcal{D}^{\oplus|\mathcal{Z}|}$, we can minimize with respect to each conditional state ρ_z^Q separately, giving

$$h_-(\lambda_i) = \sum_z p_z F_{2,-}\left(\sum_i \lambda_i \tilde{A}_z^{(i)}\right) + \sum_i \lambda_i(b_i - a_i), \tag{35}$$

where $a_i := \sum_z p_z \mathrm{Tr}[A_z^{(i)}]$, and $\tilde{A}_z^{(i)} = A_z^{(i)} - \mathrm{Tr}[A_z^{(i)}]\mathbb{I}/d$ is the traceless part of $A_z^{(i)}$. Here we have defined a function over the space of traceless Hermitian operators C,

$$F_{2,-}(C) := \inf_{\rho \in \mathcal{D}}(\mathrm{Tr}[\rho^2] - \mathrm{Tr}[\rho C]). \tag{36}$$

Evidently, the above depends only on the eigenvalues of C. In principle, a fully general expression for the above can be

found; however, to make progress in the following, we use a simple lower bound $F_{2,-}(C) \geq 1/d - \mathrm{Tr}[C^2]/4$. (This actually becomes an equality if $|\min \mathrm{eig}C| \leq 2/d$.) Invoking this bound, we are left with a manageable expression for the dual function

$$h_-(\lambda_i) \geq -\frac{1}{4}\sum_z p_z \mathrm{Tr}\left[\left(\sum_i \lambda_i \tilde{A}_z^{(i)}\right)^2\right] + \sum_i \lambda_i(b_i - a_i). \tag{37}$$

The right-hand side is readily maximized over the Lagrange parameters λ_i, which gives us a lower bound for the solution of the dual problem, i.e.,

$$h_-^* \geq \tilde{h}_- := \frac{1}{d} + \sum_{ij}(b_i - a_i)[L^{-1}]_{ij}(b_j - a_j), \tag{38}$$

where we define the matrix

$$L_{ij} = \sum_z p_z \mathrm{Tr}[\tilde{A}_z^{(i)}\tilde{A}_z^{(j)}]. \tag{39}$$

The lower bound (38) corresponds to the set of dual parameters $\lambda_i = 2\sum_j [L^{-1}]_{ij} b_j$, and thus the bound becomes optimal ($\tilde{h}_- = h_-^*$) if the condition $\min \mathrm{eig}(\sum_{ij} [L^{-1}]_{ij} b_j A_z^{(i)}) \geq -1/d$ is met for all z. This condition may or may not be met for any given set of measurements; regardless, Eq. (38) constitutes a viable certificate by way of strong duality (32), in that the averaged global purity of the true ensemble $\mathcal{E}^Q$ can be no less than the right-hand side, i.e., $\mathbb{E}_z \mathrm{Tr}[(\rho_z^Q)^2] \geq \tilde{h}_-$.

1. Incorporating classical simulations

Having reached this point, we can revisit the original problem and ask: how could we have chosen $A_z^{(i)}$ in the first place in order to make our bound (38) as close as possible to the true value of the average $\mathbb{E}_z G(\rho_z^Q)$? We can use the discussion of Sec. IV to guide our intuition. There, we posited that the quantum-classical correlator (16), which is a measurable quantity, would serve as a good proxy for

the quantum-quantum correlator (15), on account of the fact that the two coincide in the limit of perfect classical simulation $\rho_z^Q = \rho_z^C$. Here, we can define an analogous "quantum-classical purity,"

$$P^{QC} := \sum_z p_z \operatorname{Tr}[\rho_z^Q \rho_z^C], \qquad (40)$$

which can evidently be cast in the form of an estimable quantity (8), and indeed is equal to the desired average when $\rho_z^Q = \rho_z^C$. Accordingly, it is instructive to consider an example where this is the only constraint we use, i.e., $R = 1$, with $A_z^{(1)} = \rho_z^C$ and $b_1 = P^{QC}$. Then, altogether, bound (B9) reduces to

$$\mathbb{E}_z\{\operatorname{Tr}[(\rho_z^Q)^2]\} \geq \frac{(P^{QC})^2}{P^{CC}}, \qquad (41)$$

where we define $P^{CC} = \sum_z p_z \operatorname{Tr}[(\rho_z^C)^2]$ by analogy to Eq. (40). The above inequality can also be proved by independent means using the Cauchy-Schwartz inequality, first applied to the Hilbert-Schmidt inner product $\operatorname{Tr}[\rho_z^Q \rho_z^C] \leq \sqrt{\operatorname{Tr}[(\rho_z^Q)^2] \operatorname{Tr}[(\rho_z^C)^2]}$, and then to the average $(\mathbb{E}_z \sqrt{a_z b_z})^2 \leq (\mathbb{E}_z a_z)(\mathbb{E}_z b_z)$. This serves as a useful sanity check for our series of bounds $\tilde{h}_- \leq g_-^* \leq \mathbb{E}_z \operatorname{Tr}[(\rho_z^Q)^2]$. We observe that, as the quality of the classical simulation improves, the corresponding lower bound should increase, resulting in tighter bounds.

We emphasize, however, that Eq. (38) is far more versatile as a bound than the simple inequality (41). In particular, we can incorporate multiple constraints (19) that allows us to use more information than just the averaged overlap between classical and quantum states. In particular, since the purity is a quadratic function of the density matrix, this suggests using more general quantum-classical correlators that are bilinear in ρ_z^Q and ρ_z^C. For this purpose, we introduce the superoperators (linear maps over the space of operators)

$$\eta^{QC} = \sum_z p_z |\rho_z^Q \rangle\!\rangle \langle\!\langle \rho_z^C|, \qquad (42a)$$

$$\eta^{CC} = \sum_z p_z |\rho_z^C \rangle\!\rangle \langle\!\langle \rho_z^C|. \qquad (42b)$$

Thinking of these as $(d^2 \times d^2)$-dimensional matrices, one can see that one can extract all possible quantum-classical and classical-classical correlators from the above objects: for any operators A, B, one has $\langle\!\langle A|\eta^{QC}|B\rangle\!\rangle = \langle A^\dagger \otimes B \rangle^{QC}$, and similar for η^{CC}. Therefore, if we were to construct a complete basis of operators $\{\sigma^\mu\}$ and measure all correlators of the form $\langle \sigma^\mu \otimes \sigma^\nu \rangle^{QC}$ for $\mu, \nu = 1, \ldots, d^2$, we could fully reconstruct η^{QC}. For each one of these measured correlators, we have a Lagrange multiplier λ_i, and these can also be organized into a superoperator form that

we call $\tilde{\zeta}^{CQ}$. When arranged in this way, the dual function (37) then becomes

$$h_-(\tilde{\zeta}^{CQ}) = 2\operatorname{STr}[\eta^{QC}\tilde{\zeta}^{CQ}] - \operatorname{STr}[(\tilde{\zeta}^{CQ})^\dagger \tilde{\zeta}^{CQ}]. \qquad (43)$$

Here, $\operatorname{STr}[\eta]$ denotes the trace of a superoperator η, which we could write in terms of a complete orthonormal basis of operators σ_μ as $\operatorname{STr}[\eta] = \sum_\mu \langle\!\langle \sigma_\mu|\eta|\sigma_\mu\rangle\!\rangle$. While the above holds for any choice of $\tilde{\zeta}^{CQ}$, by analogy to Eq. (38) we can find the optimum choice ζ^{CQ} (written without a tilde), which is the solution to the superoperator equation (which is guaranteed to exist)

$$\eta^{CC}\zeta^{CQ} = (\eta^{QC})^\dagger. \qquad (44)$$

[If η^{CC} has an inverse, we could write $\zeta^{CQ} = (\eta^{CC})^{-1}(\eta^{QC})^\dagger$.] At this dual-optimum point, we obtain the bound

$$\mathbb{E}_z \operatorname{Tr}[(\rho_z^Q)^2] \geq g_-^* \geq \operatorname{STr}[\eta^{QC}\zeta^{CQ}]. \qquad (45)$$

Equation (45) is our first concrete inequality that can be straightforwardly calculated using classical-quantum correlators. Evidently, in the case of perfect classical simulation $\rho_z^C = \rho_z^Q$, we have $\zeta^{CQ} = \mathrm{id}$, and the inequalities in Eq. (45) are both saturated. Hence, although we have lost some tightness in our calculations, we expect the bound to be near optimal when the classical simulation is not quite perfect. We benchmark how tight these inequalities are in various scenarios in Sec. VII below.

B. Global purity upper bound

In searching for an upper bound for the global purity, we could in principle set up the maximization problem (27) in an entirely analogous way to the above, finding an upper bound for h_+^* and exploiting weak duality to obtain a corresponding bound for g_+^*. However, the global purity has a special significance in this context, which means that an immediate answer can be obtained by invoking the corollary of Theorem 1: we know that there exists at least one ensemble consistent with measurements (19) for which no more than R of the states ρ_z^Q are mixed. The existence of such an ensemble immediately implies that g_+^* must be at least

$$g_+^* \geq 1 - \left(1 - \frac{1}{d}\right) R \max_z p_z. \qquad (46)$$

In the no-coincidence limit $p_z \ll 1$, the above is very close to 1, which is itself a universally applicable upper bound for the averaged purity. We conclude that regardless of which observables $A_z^{(i)}$ we measure, we cannot obtain a nonvacuous upper bound for the averaged global purity. This fundamental obstruction is related to the special significance that global purity has for the structure of $\mathcal{K}$, and in particular its extreme points. We discuss this issue in depth in Sec. VIII A below.

C. Lower bound for subsystem purities and quadratic observables

Moving beyond the global purity, we can consider more general functions $G(\rho_z^Q)$ that are quadratic in the conditional states. Most generally, one can write these as

$$G(\rho_z^Q) = \langle\!\langle \rho_z^Q | \mathcal{N} | \rho_z^Q \rangle\!\rangle, \qquad (47)$$

where $\mathcal{N}$ is a Hermitian superoperator, i.e., a linear map over the space of operators satisfying $\langle\!\langle C | \mathcal{N} | D \rangle\!\rangle = \langle\!\langle D | \mathcal{N} | C \rangle\!\rangle^*$ for operators C, D. To ensure the convexity of G, we insist that $\mathcal{N}$ is positive semidefinite: $\langle\!\langle C | \mathcal{N} | C \rangle\!\rangle \geq 0$ for any operator C; we write this condition as $\mathcal{N} \succeq 0$. Examples of this include the quantum-quantum correlator $\langle O \otimes O \rangle^{QQ}$ [Eq. (15)], which corresponds to $\mathcal{N} = |O\rangle\!\rangle\langle\!\langle O|$. Moreover, the purity of some subsystem of Q, which we denote Q_1 with complement $Q_2 = Q \backslash Q_1$, can be written in this form: one can express purity of Q_1 as

$$\text{Tr}_{Q_1}\{(\text{Tr}_{Q_2}[\rho_z^Q])^2\} = \sum_{\nu \in Q_1} \langle\!\langle \rho_z^Q | \sigma_\nu \rangle\!\rangle \langle\!\langle \sigma_\nu | \rho_z^Q \rangle\!\rangle, \qquad (48)$$

where $\{\sigma_\nu\}$ is a basis of operators that respects the tensor product structure of the Hilbert space $\mathcal{H}^Q = \mathcal{H}^{Q_1} \otimes \mathcal{H}^{Q_2}$, and the notation $\nu \in Q_1$ denotes that the sum is restricted to those ν for which σ_ν acts as the identity on Q_2. The above implicitly defines the superoperator $\mathcal{N}$ corresponding to the subsystem purity.

To solve the minimization problem for this class of observables, we use similar logic to that described in Sec. VI A, with some modifications. Detailed arguments are presented in Appendix B 1, which lead to the bound

$$\sum_z p_z \langle\!\langle \rho_z^Q | \mathcal{N} | \rho_z^Q \rangle\!\rangle \geq \text{STr}[\mathcal{N} \eta^{QC} \zeta^{CQ}], \qquad (49)$$

where ζ^{CQ} is the superoperator that was defined in Eq. (44). The global purity bound (45) corresponds to the special case where $\mathcal{N} = \text{id}$, the identity superoperator. Again, the above inequality is saturated in the limit of perfect classical simulation.

D. Upper bound for quadratic observables

The corresponding maximization problem for averages of the form (47) is not a convex optimization problem, and so cannot be solved in a fully analogous way. Rather than directly solving the dual problem, we instead choose to re-express the problem by first trivially rewriting

$$\mathcal{N} = \|\mathcal{N}\|_\infty \text{id} + \overline{\mathcal{N}}, \qquad (50)$$

where we define

$$\overline{\mathcal{N}} = (\mathcal{N} - \|\mathcal{N}\|_\infty \text{id}), \qquad (51)$$

and $\|\mathcal{N}\|_\infty = \max_{\langle\!\langle C | C \rangle\!\rangle = 1} \langle\!\langle C | \mathcal{N} | C \rangle\!\rangle$ is the spectral norm of $\mathcal{N}$ when viewed as a matrix. The significance of the

above is that $\overline{\mathcal{N}} \preceq 0$ by construction, and hence, when we substitute this into Eq. (47), the second term constitutes a concave function of ρ_z^Q. The first term simply gives us a term proportional to the purity $\text{Tr}[(\rho_z^Q)^2]$, which by Theorem 1 we know to be very close to unity at the point where the maximum is achieved. Hence, we lose little tightness by replacing the first term with the constant $\|\mathcal{N}\|_\infty$.

Being concave, the average of $\langle\!\langle \rho_z^Q | \overline{\mathcal{N}} | \rho_z^Q \rangle\!\rangle$ can be upper bounded, using the same method as the lower bound for the convex function (49). Altogether, we obtain

$$\mathbb{E}_z \langle\!\langle \rho_z^Q | \mathcal{N} | \rho_z^Q \rangle\!\rangle \leq \|\mathcal{N}\|_\infty + \text{STr}[\overline{\mathcal{N}} \eta^{QC} \zeta^{CQ}]$$
$$= \|\mathcal{N}\|_\infty (1 - \text{STr}[\eta^{QC} \zeta^{CQ}])$$
$$+ \text{STr}[\mathcal{N} \eta^{QC} \zeta^{CQ}]. \qquad (52)$$

From Eqs. (49) and (52), it becomes clear that we can constrain the average of Eq. (47) to within a window whose width is determined by the quantity $(1 - \text{STr}[\eta^{QC} \zeta^{CQ}])$. Indeed, from Secs. VI A and VI B, this quantity is itself equal to the range of values that the global purity can take. This highlights the significance of Theorem 1: if the true states ρ_z^Q realized by the device are themselves mixed then the purity cannot be constrained to be within an arbitrarily narrow window, and in turn, quadratic observables of this kind cannot be tightly constrained either. Indeed, even before we performed any of the manipulations in this section, one could straightforwardly show that

$$\text{ran}(\mathcal{N}) + \text{ran}(\overline{\mathcal{N}}) \geq \|\mathcal{N}\|_\infty \text{ran}(\text{id})$$
$$= \|\mathcal{N}\|_\infty (1 - \text{STr}[\eta^{QC} \zeta^{CQ}]), \qquad (53)$$

where we used the shorthand $\text{ran}(\mathcal{N}) = g_+^* - g_-^*$ for the observable $\mathbb{E}_z[\langle\!\langle \rho_z^Q | \mathcal{N} | \rho_z^Q \rangle\!\rangle]$. Hence, if there is large uncertainty in the purity then there must also be large uncertainty in certain quadratic observables (see also Sec. VIII A below).

E. von Neumann entropy upper bound

The von Neumann entropy $S(\rho_z^Q) = -\text{Tr}[\rho_z^Q \log \rho_z^Q]$ is another quantity that can be used to characterized mixedness and/or entanglement of quantum states, which has particular information-theoretic significance. This function is concave in ρ_z^Q, and so in this instance the maximization problem is most easily addressed. The dual function for the maximization problem is

$$h_+(\lambda_i) = \sum_z p_z F_{S,+}\left(\sum_i \lambda_i A_z^{(i)}\right) + \sum_i \lambda_i b_i, \qquad (54)$$

where by analogy to $F_{2,-}(C)$ we define $F_{S,+}(C) = \sup_{\rho^Q \in \mathcal{D}} \{S(\rho^Q) - \mathrm{Tr}[\rho^Q C]\}$. A fairly straightforward calculation reveals

$$F_{S,+}(C) = \log \mathrm{Tr}[e^{-C}], \qquad (55)$$

which is attained for $\rho = e^{-C}/\mathrm{Tr}[e^{-C}]$. Hence, we have

$$\mathbb{E}_z S(\rho_z^Q) \le h_+^* = \min_{\lambda_i} \left\{ \sum_i \lambda_i b_i \right.$$
$$\left. + \sum_z p_z \log \mathrm{Tr}_Q \left[\exp\left(-\sum_i \lambda_i A_z^{(i)} \right) \right] \right\}. \qquad (56)$$

This upper bound can be applied quite generally, and can even be used to rederive a result that was obtained in Ref. [34] using alternative arguments based on the quantum relative entropy: if we have a single constraint with $A_z = -\log \rho_z^C$ then

$$h_+(\lambda) = \sum_z p_z \{\log \mathrm{Tr}[(\rho_z^C)^\lambda] - \lambda \mathrm{Tr}[\rho_z^Q \log \rho_z^C]\}. \qquad (57)$$

Setting $\lambda = 1$ in the above, we obtain a certificate

$$\mathbb{E}_z S(\rho_z^Q) \le h_+(\lambda = 1) = -\sum_z p_z \mathrm{Tr}[\rho_z^Q \log \rho_z^C], \qquad (58)$$

the right-hand side of which was introduced in Ref. [34] as the "quantum-classical entanglement entropy." Evidently, we could use exactly the same inequality for the average von Neumann entropy of a subsystem of Q, provided one replaced all density matrices with the corresponding reduced density matrices. The significance of the choice $\lambda = 1$ can be understood by recognizing that $\lambda = 1$ is dual optimal in the limit of perfect simulation $\rho_z^C = \rho_z^Q$, since inequality (58) is then saturated. If there are small discrepancies in the quantum and classical states, one expects that the optimal choice of λ will be shifted slightly, and hence this bound can in principle be tightened by optimizing over λ.

However, a more serious problem to address is the instability of the quantum-classical entanglement entropy for singular or near-singular classical states ρ_z^C, since $\log \rho_z^C$ diverges when the eigenvalues of ρ_z^C are small. This is a particularly important problem when the ρ_z^C are close to being pure, as noted in Ref. [34]. Here, our more general approach allows one to get around this issue in a systematic way, since the choice $A_z = -\log \rho_z^C$ can be easily altered in a way that guarantees numerical stability. A natural choice of regularization is to work directly with the eigenvalue decomposition $\rho_z^C = \sum_n q_{z,n} |\chi_{z,n}\rangle \langle \chi_{z,n}|$, and to separate out the problematic eigenvalues—namely, those that

are below some threshold ϵ. We are free to separately measure the two quantities

$$A_z^{(1)} = -\Pi_z^> \log \rho_z^C \Pi_z^>, \qquad A_z^{(2)} = \Pi_z^\le, \qquad (59)$$

where we define the projector $\Pi_z^\le = \sum_{q_n \le \epsilon} |\chi_{z,n}\rangle \langle \chi_{z,n}|$ and its complement $\Pi_z^> = \mathbb{I} - \Pi_z^\le$. The first of these is similar to the original choice $A_z = -\log \rho_z^C$, but with near-singular eigenvalues removed, while the second measures the average weight of ρ_z^Q lying within the near-singular subspace $\Pi_z^\le$. We can then proceed as before and find an optimal bound based on empirical values of the above two observables. In Appendix B 2, we work out such an upper bound explicitly. As a particularly simple special case, if the classical states are all pure, $\rho_z^C = |\phi_z^C\rangle \langle \phi_z^C|$, then this improved bound can be expressed in terms of the quantity $\overline{\delta^Q} = 1 - \sum_z p_z \langle \phi_z^C | \rho_z^Q | \phi_z^C \rangle$ as

$$\mathbb{E}_z S_{\mathrm{vN}}(\rho_z^Q) \le H_2(\overline{\delta^Q}) + \overline{\delta^Q} \log(d-1), \qquad (60)$$

where $H_2(p) = -p \log p - (1-p)\log(1-p)$ is the Shannon entropy for a binary random variable.

F. von Neumann entropy lower bound

If we focus on the global von Neumann entropy then arguments similar to those in Sec. VI B can be used to show that a nonvacuous lower bound cannot be obtained: Theorem 1 implies the existence of a feasible ensemble whose average entropy is no greater than $R \max_z p_z \log d$, which is small in the no-coincidence limit $p_z \ll 1$.

As for the von Neumann entropy of a subsystem Q_1, this can in principle be lower bounded, but a direct analysis of the minimization problem is not straightforward, due to the concavity of $S(\rho_z^{Q_1})$. Instead, we use a similar approach to Sec. VI D, where we exploited the fact that the maximum of a convex function is attained at an extreme point, which we know to be a (mostly) pure ensemble. There, we rewrote the desired function in terms of the difference between the purity and a convex function [Eq. (50)], thereby allowing the maximization problem to be transformed into a convex optimization task. Here, we use a similar line of reasoning, by trivially rewriting

$$S(\rho_z^{Q_1}) = S(\rho_z^Q) - S(\rho_z^Q | Q_1), \qquad (61)$$

where $S(\rho^{AB}|B) := S(\rho^{AB}) - S(\rho^B)$ is the conditional quantum entropy for a bipartite state ρ^{AB}. From Theorem 1 and the concavity of $S(\rho_z^{Q_1})$, the desired minimum is attained for an ensemble that is mostly globally pure, and so the average of $S(\rho_z^Q)$ will be close to 0 at this point. This suggests that we do not lose much tightness in employing

the bound

$$
\begin{aligned}
g_-^* &= \min_{\rho \in \mathcal{K}} \mathbb{E}_z S(\rho_z^{Q_1}) \\
&\leq \min_{\rho \in \mathcal{K}} [-\mathbb{E}_z S(\rho_z^Q | Q_1)] \\
&= -\max_{\rho \in \mathcal{K}} [\mathbb{E}_z S(\rho_z^Q | Q_1)]. \quad (62)
\end{aligned}
$$

Crucially, $S(\rho_z^Q | Q_1)$ is a concave function of the global state ρ_z^Q (a statement that is equivalent to strong subadditivity of the von Neumann entropy [61]). Thus, we are left with a convex optimization problem, as desired. In Appendix B 3, we derive the concrete bound

$$
\mathbb{E}_z[S(\rho_z^{Q_1})] \geq \sum_i \lambda_i b_i - \mathbb{E}_z \log \| \mathrm{Tr}_{Q_2} e^{\sum_i \lambda_i A_z^{(i)}} \|_\infty \quad (63)
$$

for any choice of the Lagrange multipliers λ_i. Again, we can specialize to a particular choice of operators $A_z^{(i)}$ to rederive a result of Ref. [34]: taking $A_z^{(1)} = -\log \rho_z^C$ and $A^{(2)} = +\log \rho_z^{C_1}$, one can set $\lambda_1 = \lambda_2 = 1$. With this particular choice, the argument of the logarithm in the above becomes $\| \mathrm{Tr}_{Q_2} e^{\log \rho_z^C - \log \rho_z^{C_1} \otimes \mathbb{I}_{Q_2}} \|_\infty$, which can be shown to be no greater than unity using Theorem 11.29 of Ref. [62]; we then have

$$
\mathbb{E}_z[S(\rho_z^{Q_1})] \geq \mathbb{E}_z \{ \mathrm{Tr}[\rho_z^Q \log \rho_z^C] - \mathrm{Tr}[\rho^{Q_1} \log \rho_z^{C_1}] \}. \quad (64)
$$

As with the other optimization problems considered in this section, we can choose to use bounds that are simple and easy to evaluate, such as Eq. (64), or to use the more versatile expression (63), which is more cumbersome, but can in principle be numerically optimized to obtain a tighter inequality.

G. Frame potential

The frame potential is a property of a quantum state ensemble that characterizes the variation between different quantum states in the ensemble. For any integer k, it is defined as

$$
F^{(k)}(\mathcal{E}^Q) := \sum_{z,z'} p_z p_{z'} \mathrm{Tr}[\rho_z^Q \rho_{z'}^Q]^k. \quad (65)
$$

For pure state ensembles, the frame potential can be used to characterize how far away an ensemble is from being a quantum state k-design [63,64], namely, an ensemble for which the kth moments $\sum_z p_z (\rho_z^Q)^{\otimes k}$ coincide with those of the Haar ensemble [48,65] (see Sec. VII B below).

Although the quantity (65) is not manifestly in the form of the averages (3) considered so far, we can use a simple trick to bring it into the appropriate form: focusing on $k = 2$ for now, we can write $F^{(2)} = \sum_z p_z \langle\langle \rho_z^Q | \eta^{QQ} | \rho_z^Q \rangle\rangle$,

where η^{QQ} is defined by analogy to Eq. (42). We do not have access to η^{QQ} in advance, but from bounds (49) and (52) we can infer that

$$
\eta^{QQ} \succeq \frac{\eta^{QC} \zeta^{CQ} + \mathrm{H.c.}}{2}, \quad (66a)
$$

$$
\eta^{QQ} \preceq \frac{\eta^{QC} \zeta^{CQ} + \mathrm{H.c.}}{2} + (1 - \mathrm{STr}[\eta^{QC} \zeta^{CQ}]) \times \mathrm{id}, \quad (66b)
$$

where the notation $\mathcal{N} \preceq \mathcal{N}'$ for two superoperators $\mathcal{N}$, $\mathcal{N}'$ means that $\langle\langle C | \mathcal{N} | C \rangle\rangle \leq \langle\langle C | \mathcal{N}' | C \rangle\rangle$ for all operators C. Moreover, by its definition we have $\|\eta^{QQ}\|_\infty = \max_{\|C\|_2^2 = 1} \sum_z p_z \mathrm{Tr}[\rho_z^Q C]^2 \leq 1$. Hence, successive application of bounds (49) and (52) yields

$$
F^{(2)}(\mathcal{E}^Q) \geq \mathrm{STr}[(\eta^{QC} \zeta^{CQ})^2], \quad (67a)
$$

$$
F^{(2)}(\mathcal{E}^Q) \leq \mathrm{STr}[(\eta^{QC} \zeta^{CQ})^2] + (1 - \mathrm{STr}[\eta^{QC} \zeta^{CQ}]^2). \quad (67b)
$$

More generally, for higher k, the frame potential can be expressed in terms of a new "doubled" ensemble $\mathcal{E}^{QQ}$ with the structure

$$
\mathcal{E}^{QQ} = \{ (p_z p_{z'}, \rho_z^Q \otimes \rho_{z'}^Q) \}_{(z,z') \in \mathcal{Z} \times \mathcal{Z}}. \quad (68)
$$

In words, samples from this ensemble correspond to independently sampled pairs (z, z') of the original ensemble $\mathcal{E}^Q$, and the corresponding states are tensor products $\rho_z^Q \otimes \rho_{z'}^Q$. Then, the frame potential (65) of $\mathcal{E}^Q$ can be interpreted as an average of the form (3) for $\mathcal{E}^{QQ}$, with the function G given by

$$
G(\sigma) = \mathrm{Tr}[\sigma \pi_S]^k, \quad (69)
$$

where σ is a density operator on the doubled space, and $\pi_S = \sum_{a,b} |a \otimes b\rangle \langle b \otimes a|$ is the swap operator. Bounds for averages of kth powers of expectations, such as the above, can in principle be derived using similar approaches to those described in this section.

VII. NUMERICAL BENCHMARKING

Having derived various bounds for certain ensemble-averaged quantities of interest, we now present results of some numerical experiments where we simulate the full procedure described in this work, including estimation of the measurable parameters (19) to constructing the bounds.

A. Projected ensemble

As a testbench for our method, we consider the projected ensemble of a many-body state generated by (possibly noisy) finite time evolution from a product state.

For the noiseless case, we take as the premeasurement state $|\Psi(t)\rangle = U(t)|0^{\otimes N}\rangle$, where t is an integer-value discrete time, and $U(t) = U_F^t$ is a unitary generated by Floquet evolution with Floquet unitary $U_F = e^{-it_2 H_2} e^{-it_1 H_1}$. We choose the Hamiltonians $H_{1,2}$ to be given by the one-dimensional tilted-field Ising model with open boundary conditions $H_\alpha = \sum_{j=1}^{N-1} J_{j,\alpha} Z_j Z_{j+1} + h_{j,\alpha}^x X_j + h_{j,\alpha}^z Z_j$ for $\alpha = 1, 2$, where (X_j, Y_j, Z_j) are Pauli operators on site j. For the purposes of this subsection, we fix $J_{j,\alpha} = 1$, $h_{j,\alpha}^z = (1 + \sqrt{5})/2$, $h_{j,1}^x = 0.4$, $h_{j,2}^x = -0.6$. The projected ensemble considered here is defined as the set of postmeasurement states arising when $N - 1$ qubits are measured in the computational (Pauli-Z) basis, with the qubit on site $j = 1$ acting as the "unmeasured" qubit. In cases where noise is present, we apply an amplitude damping channel of strength p_{dec} on every qubit after each Floquet period. Concretely, the system density matrix evolves as $\rho(t + 1) = (\mathcal{N}_{\text{damp}} \circ \mathcal{N}_{\text{unit}})[\rho(t)]$, where $\mathcal{N}_{\text{unit}}[\rho] = U_F \rho U_F^\dagger$ and $\mathcal{N}_{\text{damp}}[\rho] = K_0 \rho k_0^\dagger + K_1 \rho K_1^\dagger$ with

$$K_0 = \text{diag}(1, \sqrt{1 - p_{\text{dec}}}) \quad \text{and} \quad K_1 = \begin{pmatrix} 0 & \sqrt{p_{\text{dec}}} \\ 0 & 0 \end{pmatrix}.$$

In the noisy case, where we must compute the full evolution of the density matrix, our exact diagonalization results are limited to relatively small system sizes. To ensure cross-comparability between all cases, we fix $N = 10$ throughout. Note that the specific form of evolution we choose, along with the parameters selected, do not have a great deal of bearing on the performance on our method—indeed, we have considered several different settings and found quantitatively similar performance. For the purposes of this section, we focus on the particular average $\bar{G} = \mathbb{E}_z \text{Tr}[\rho_z^Q Z_1]^2$, for which we can use bounds (49) and (52); again, this choice is not particularly important.

We simulate the full protocol using the shadow tomography–based method described in Sec. III C (Fig. 2) to extract the values of the estimable parameters (8). Specifically, the procedure involves (1) preparation of the premeasurement state, (2) applying random on-site Clifford unitaries to the unmeasured qubits, (3) projectively measuring all qubits, (4) processing the measurement outcomes via the dual frame (13), in terms of which estimators for the quantum-classical correlators can be obtained, and, finally, (5) using the bounds derived in the previous section to constrain the value of the desired quantity. This is repeated a finite number of times M.

The finite number of repetitions M means that there will be some residual uncertainty in the values of the chosen estimable parameters. Taking the example of correlators (42), we obtain an estimate $\hat{\eta}^{QC} = \bar{\eta}^{QC} + \hat{\epsilon}^{QC}$, where $\bar{\eta}^{QC}$ is the true value, and $\hat{\epsilon}^{QC}$ is a zero-mean error, whose variance decays as $1/\sqrt{M}$ (hats are used to denote random variables here). Despite this uncertainty, we can still

employ the bounds in a rigorous manner as follows. We use the estimated value $\hat{\eta}^{QC}$ as parameters for finding the optimal Lagrange multipliers [in this case we substitute $\hat{\eta}^{QC}$ into Eq. (44) to find $\tilde{\zeta}^{CQ}$]. This will give us a set of Lagrange multipliers that are approximately dual-optimal point for the true problem, which features $\bar{\eta}^{QC}$ rather than $\hat{\eta}^{QC}$. These Lagrange multipliers can then be substituted into the dual function [in this case Eq. (43)], which by virtue of being linear in b_i (equivalently, η^{QC}) can be evaluated with some error bars based on estimates of the standard deviation of b_i, along with some prechosen confidence level—we pick 99% confidence throughout. One can then guarantee that the true value is no smaller or greater (as appropriate) than the constructed value with probability at least 99%.

As a first check, we can consider the best-case scenario, where dynamics is noiseless and the classical simulation is perfect, $\rho_z^C = \rho_z^Q$. We show upper and lower bounds for the target function $\bar{G} = \mathbb{E}_z (\text{Tr}[\rho_z^Q Z_1]^2)$ as calculated via our method, along with the true value, in Fig. 3. Because of the fact that the ensemble states are pure and simulation is perfect, the only source of uncertainty in the inferred value of $\bar{G}$ is from the finite number of samples M, which we deal with as described above. We show bounds empirically constructed for two different sample sizes M, and both lower and upper bounds will tend towards the true value as M is increased indefinitely, with corrections scaling as $1/\sqrt{M}$.

We can now introduce some inaccuracy in the classically simulated states ρ_z^C. For this purpose, we consider the projected ensemble at a fixed time $t = 8$, using the same "quantum" states ρ_z^Q as before. To construct the imperfect classical states, we employ the same Floquet
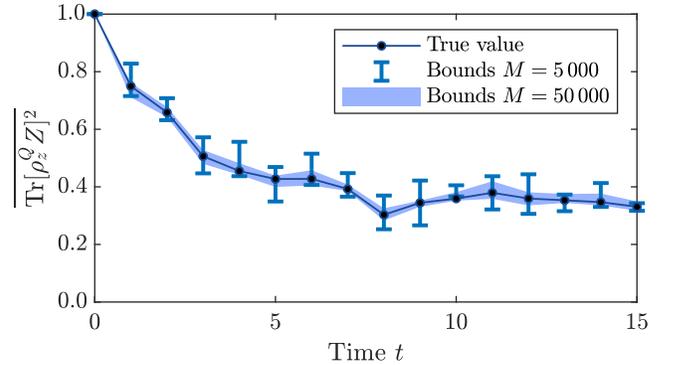
FIG. 3. Estimates of the averaged quantity $\bar{G} = \mathbb{E}_z (\text{Tr}[\rho_z^Q Z_1]^2)$ for the projected ensemble generated from noiseless dynamics as described in the main text. The true value is shown, along with bounds constructed using our method for two different choices of the number of experimental repetitions $M = 5000$ and $M = 50000$. In this case, classical simulation is perfect and the states in the ensemble are pure, so as M is increased, both upper and lower bounds converge asymptotically towards the true value.

evolution as above, but now with some additional spatially dependent randomness in each of the Hamiltonian parameters $J_{j,\alpha}, h_{j,\alpha}^{x,z}$. For each parameter, we pick a perturbed value $\tilde{J}_{j,\alpha} = J_{j,\alpha}(1 + fn_{j,\alpha})$ (similar for $\tilde{h}_{j,\alpha}^{x,z}$), where $n_{j,\alpha} = \pm 1/2$ are chosen independently at random for each parameter, but kept fixed in all of the data shown here. The free parameter f represents a fractional uncertainty in the Hamiltonian parameters; thus, as it is increased, the quantum and classical states ρ_z^Q, ρ_z^C become less correlated. We can keep track of accuracy of the classical states in our numerical experiment by evaluating the quantity $\Delta^{QC} := \sum_z p_z \|\rho_z^Q - \rho_z^C\|_1$, where $\|\rho - \sigma\|_1$ is the trace distance.

Figure 4 shows data obtained using these inaccurate classical states. We show both the bounds constructed based on data from a finite number of experimental samples $M = 50000$, along with "asymptotic" (i.e., $M \to \infty$) bounds, which we obtain by evaluating the right-hand sides of Eqs. (49) and (52), thus eliminating any statistical uncertainty. We see that, as the fractional uncertainty f is increased, the distance between the quantum and classical states increases, and the bounds become less tight. For reference, the value of Δ^{QC} that would be obtained if ρ_z^C were chosen as independent random states would be 1; thus, the data towards the right of this plot represent

fairly poor simulation. In the absence of noise and for the specific quantity considered here, crude estimates for the deviation between the bounds and the true value scale as $\sum_z p_z \|\rho_z^Q - \rho_z^C\|^2$ (at least in the regime of small Δ^{QC}), which is upper bounded by $(\Delta^{QC})^2$. This quadratic dependence on the distance between quantum and classical states explains why we see good performance, even when there is appreciable difference between the two states.

Finally, we consider the case where both noise and inaccuracies in the classical states are present, which reflects the nature of realistic experiments. Again, we fix $t = 8$, and now consider amplitude damping of strength $p_{\text{dec}} = 0.002$ for each qubit and each timestep. Although this value of p_{dec} appears to be small, the expected number of errors in the full circuit can be estimated to be $N \times t \times p_{\text{dec}} = 0.24$, which has an appreciable effect on the conditional states, as can be seen in their mean purity, which here is 0.95. The classical states are generated in the same way as the ensemble states, including the noise, but with additional randomness in the Hamiltonian parameters of strength f, introduced in the same manner as above.

The results are shown in Fig. 5. The data follow most of the same trends as before, with higher fractional uncertainty leading to higher deviation Δ^{QC}, and in turn less tight bounds. However, because the dynamics is noisy, the states in the ensemble ρ_z^Q are no longer pure—in this case their average purity is $\sum_z p_z \text{Tr}[(\rho_z^Q)^2] = 0.95$. Unlike the lower bound, the upper bound saturates at a value that is strictly above the true value of $\bar{G}$ even in the limit of perfect simulation ($f = 0$). This can be seen as a consequence of Theorem 1: we saw already in Sec. VI B that the global purity cannot be tightly bounded, and the same goes

FIG. 4. Effect of the fidelity of the classically simulated states ρ_z^C on the tightness of the bounds for the noiseless projected ensemble at time $t = 8$. The classical states ρ_z^C are perturbed away from the ensemble states ρ_z^Q by introducing some fractional uncertainty $f \sim \delta J/J$ in the Hamiltonian parameters, as described in the main text. In the top panel, we include both bounds constructed based on shadow tomographic data from $M = 50000$ repetitions, along with asymptotic bounds, where we evaluate the right-hand sides of inequalities (49) and (52) without any statistical uncertainty—this is the value that one would obtain as $M \to \infty$. The bottom panel shows the average trace distance between quantum and classical states $\Delta^{QC} = \sum_z p_z \|\rho_z^Q - \rho_z^C\|_1$, which acts as a measure of how faithful the simulated states are.

FIG. 5. As in Fig. 4, but with noisy dynamics ($p_{\text{dec}} = 0.002$), resulting in a projected ensemble with mixed states, with average purity $\sum_z p_z \text{Tr}[(\rho_z^Q)^2] = 0.95$. The lower bound approaches the true value in the limit of perfect simulation, but the upper bound does not, as a consequence of Theorem 1.

for other quadratic observables. This issue is discussed in detail in Sec. VIII A below.

B. Application: verifying emergent quantum state designs

One feature of the projected ensemble that has attracted much interest recently is that, under rather generic conditions, the ensemble of postmeasurement states turns out to form an (approximate) quantum state k-design [14,15,17]. That is, the kth moments of the ensemble

$$\rho^{(k)} := \sum_z p_z (\rho_z^Q)^{\otimes k} \tag{70}$$

agree (closely) with the corresponding moments of the Haar ensemble

$$\rho_{\text{Haar}}^{(k)} := \int d\mu_{\text{Haar}}(\psi)(|\psi\rangle \langle\psi|)^{\otimes k} = \frac{\Pi_{\text{sym}}^{(k)}}{\binom{k+d-1}{k}}, \tag{71}$$

where $\Pi_{\text{sym}}^{(k)}$ is the projector onto the permutation-symmetric subspace of $(\mathcal{H}^Q)^{\otimes k}$. Such ensembles of states are in a certain sense "maximally random," and this makes them useful for certain tasks including quantum state and channel tomography [35,36,38,39,54] and cryptography [37]. Here we demonstrate how our method can be employed to verifiably conclude whether or not the projected ensemble realized in an experiment is an approximate k-design.

For ensembles of pure states, the frame potential constitutes a measure of how far the ensemble is from being a k-design. Specifically, in Ref. [64] it was shown that the frame potential (65) is related to the (normalized) Frobenius distance between the moments (70), (71) via

$$\|\rho^{(k)} - \rho_{\text{Haar}}^{(k)}\|_2^2 \stackrel{\text{(pure)}}{=} F^{(k)}(\mathcal{E}^Q) - F_{\text{Haar}}^{(k)} \tag{72}$$

with $F_{\text{Haar}}^{(k)} = \binom{k+d-1}{k}^{-1}$ the frame potential for the Haar ensemble. Thus, when the frame potential is minimized, a k-design can be formed. However, in experiment the ensemble states can be mixed, and this leads to a suppression of the frame potential that can mimic the effect of forming a quantum state design—indeed, the right-hand side of Eq. (72) can even be negative when the states are mixed. For this case, we need a more generally applicable bound. By expanding the left-hand side of Eq. (72), we find that the distance can more generally be re-expressed as

$$\|\rho^{(k)} - \rho_{\text{Haar}}^{(k)}\|_2^2 = F^{(k)} + F_{\text{Haar}}^{(k)}\left(1 - \frac{2}{k!}\sum_{\tau \in \Sigma_k} \mathbb{E}_z[P_z^{(\tau)}]\right), \tag{73}$$

where we use the shorthand

$$P_z^{(\tau)} = \prod_{c \in \tau} \text{Tr}[(\rho_z^Q)^{|c|}]. \tag{74}$$

Here, each $c \in \tau$ is a cycle of the permutation group element τ, whose length is $|c|$. In the case $k = 2$, which we focus on here, this gives

$$\|\rho^{(2)} - \rho_{\text{Haar}}^{(2)}\|_2^2 = F^{(k)}(\mathcal{E}^Q) - \frac{2}{d(d+1)}\mathbb{E}_z\{\text{Tr}[(\rho_z^Q)^2]\}. \tag{75}$$

The right-hand side of the above contains the frame potential and the average global purity. Interestingly, these appear with opposite signs, and thus we have a difference between two convex functions, which itself is not convex. Nevertheless, we can employ the upper and lower bounds derived in Sec. VI for each term separately.

We can test this method using the same dynamics that we considered in the previous subsection, including both noise and inaccuracies in the classical states. The tilted-field Ising Hamiltonian that generates the Floquet unitary is understood to be chaotic [66], and hence we expect to see emergent state designs in the projected ensemble [17]. The results of our numerical simulations are shown in Fig. 6. We plot both the true value of distance (75) along with bounds constructed from simulated experimental data using $M = 50000$ repetitions. In this case, the classically simulated states are generated from noiseless evolution, but with a fractional uncertainty in the parameters of $f = 0.5\%$. Evidently, the distance does indeed decay towards zero as time increases. For small values of noise, the bounds we obtain are relatively tight, meaning that in a real experiment (where one would not have access to the true value), a definitive conclusion regarding the formation of approximate state designs could be made. As the noise rate and/or time t increases, the total number of errors accrued during the circuit increases, and the states in the ensemble become less pure, as evidenced in the bottom panels of Fig. 6. This leads to bounds that become less tight, in particular the upper bound.

These results demonstrate that even in the presence of noise and miscalibrations between the quantum and classically simulated states, one can make concrete inferences regarding the closeness of the projected ensemble realized in experiment to being a quantum state design.

VIII. FUNDAMENTAL LIMITATIONS

So far, we have introduced an optimization-based approach to inferring properties of postmeasurement quantum state ensembles, constructed explicit two-sided bounds for various quantities, and demonstrated an immediate application of our method for the verification of emergent quantum state designs. Evidently, the method
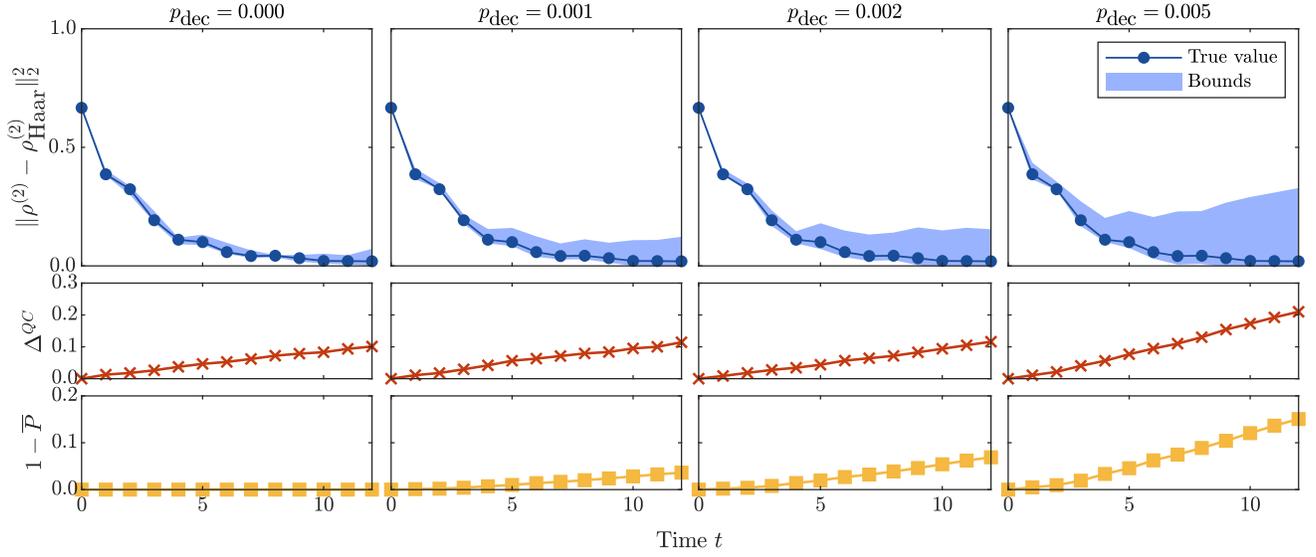
FIG. 6. Distance (75) between the $k = 2$ moments of the projected ensemble from those of the Haar ensemble, as calculated using our method with $M = 50000$ experimental repetitions (classical shadows) for various values of p_{dec} (top panel). Classical states are generated using a fractional parameter uncertainty $f = 0.5\%$, and without noise. For reference, we also show the average trace distance between quantum and classical states Δ^{QC} (middle panel), along with $1 - \sum_z p_z \, \mathrm{Tr}[(\rho_z^Q)^2]$, which measures how far the states are from being mixed on average (bottom panel).

described in this paper gives us some degree of ability to learn quantities of interest in the context of measurement-induced dynamics, but we have already seen that in some cases there are unavoidable limitations in terms of what can be unambiguously inferred from experimental data. In this section, we discuss some of these limitations in detail, with the aim to more precisely characterize the boundary between properties of the postmeasurement conditional quantum states that can or cannot be inferred from experiment.

A. Interpretation and consequences of Theorem 1

Theorem 1 and its corollary immediately tell us something regarding what can be inferred about certain quantities—most obviously those that measure how mixed the conditional states ρ_z^Q are. The average global purity $\mathbb{E}_z \, \mathrm{Tr}[(\rho_z^Q)^2]$ and the average global von Neumann entropy $\mathbb{E}_z S(\rho_z^Q)$ are both examples of this; let us consider the former for concreteness. Suppose that the true ensemble $\mathcal{E}^Q$ features states that are appreciably mixed on average, $\mathbb{E}_z \, \mathrm{Tr}[(\rho_z^Q)^2] = 1 - \delta$, with $\delta > 0$. Even in the best-case scenario, where we have perfect classical simulations, we cannot make the range of feasible values [Eqs. (20) and (23)] significantly narrower than δ, since this range must include both the true value $(1 - \delta)$ and the value $1 - O(e^{-H_{min}(p)})$ implied by Theorem 1. [Here $H_{min}(p) = -\log \max_z p_z$ is the min-entropy of the distribution p_z, which typically scales linearly with the number of measurements.] So, if the ensemble being probed is not close to being pure, $\delta > 0$, then although we might hope to obtain

a good lower bound for the averaged purity, we cannot obtain a good upper bound, and hence there will always be some uncertainty in our conclusions. We can hope to learn the purity to a good accuracy if the ensemble being measured is itself close to pure (something we would not know in advance, but could verify using our bounds). Indeed, for perfect simulation, bound (45) tends to the true purity $1 - \delta$, and hence the range of feasible values (23) has an optimal width δ.

It may seem counterintuitive that there remains an ambiguity in the average purity even when the simulation being used is perfect. This stems from the difference between knowing that the simulation is perfect versus having a perfect simulation but not knowing (or assuming) that it is so. To learn something definitive about the ensemble, we cannot make such an assumption, and this means that we may not be able to make sharp conclusions even for perfect classical simulation. Note also that this limitation is not intrinsic to the inference method we are proposing in this paper: it is a fundamental obstruction, in that whatever strategy we employ, we cannot rule out the possibility of a mostly pure state ensemble.

As a simple example that allows one to appreciate this idea intuitively, consider an ensemble defined for a single qubit ($d = 2$), where every conditional state is the same, $\rho_z^Q = \rho_0^Q$, and this fixed state ρ_0^Q is mixed. This is represented by the blue arrow on the Bloch sphere in Fig. 7(a). If our classical simulation were perfect, $\rho_z^C = \rho_0^Q$, then any quantum-classical correlator we measure corresponds to a z-independent choice of A_z in Eq. (8), since the classical states are themselves z independent. Once we measure
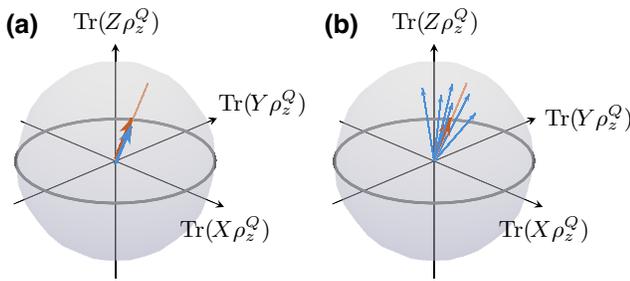
FIG. 7. An illustrative example that demonstrates that Theorem 2 implies a limitation to our ability to sharply determine the purity of an ensemble. We consider a case where the quantum ensemble is given by the same mixed state for each z, i.e., $\rho_z^Q = \rho_0^Q$, represented by a light blue arrow on the Bloch ball, with subunit length. (a) Perfect classical simulation also implies that $\rho_z^C = \rho_0^Q$ (dark orange arrow). (b) When we measure quantum-classical correlations using $\rho_z^C = \rho_0^Q$, there is also another candidate ensemble made up of pure states that vary between each z (multiple light blue arrows), distributed such that the same correlations are observed. Since we cannot *a priori* know that our simulations are perfect (even if they are), we cannot use our experimental data to rule out the possibility that the true scenario is (b).

some collection of quantum-classical correlators, we can consider which ensembles are consistent with these observations. The true ensemble is of course one possibility, but there are also alternative ensembles, made up of pure states that vary between different z, such that they average out to the same mixed state ρ_0^Q. This is shown by the many blue arrows in Fig. 7(b). Put simply, Theorem 1 reflects the fact that based on our experimental observations, we cannot determine whether the true states are mixed, or whether the states are pure but our classical simulations are inaccurate: indeed, both of these have the effect of reducing the values of the quantum-classical correlators compared to the case of pure states with perfect simulation. To rule out the possibility of inaccurate classical simulation, one would have to estimate a number of quantities R that itself scales with the cardinality of the ensemble $|\mathcal{Z}|$—effectively, one for each possible value of z. This is of course not possible unless one uses a number of experimental repetitions M that also scales with $|\mathcal{Z}|$.

Looking at Fig. 7, we can roughly characterize the difference between panels (a) and (b) by saying that the mixed state in (a) is realized as a convex combination of the many different pure states in (b). When we consider more general averages of the form (3) with convex functions G (beyond those like the average global purity), the ensembles (b) will have a higher value for this average compared to the true scenario (a), since taking convex combinations of states will reduce the value of G. Therefore, the existence of consistent pure-state ensembles implied by Theorem 1 means that, when the true ensemble is mixed, the upper bounds

of convex functions cannot be made arbitrarily tight. As for lower bounds of convex averages, these can indeed be made tight as the fidelity of classical simulation is improved. Indeed, this is borne out in the various bounds we derived in Sec. VI: compare, e.g., Eqs. (49) and (52) in the limit of $\rho_z^Q = \rho_z^C$. The lower bound tends to the true value, while the upper bound deviates from the true value by an amount $\|\mathcal{N}\|_\infty (1 - \mathbb{E}_z \operatorname{Tr}[(\rho_z^Q)^2])$. One can see this asymmetry between the upper and lower bounds explicitly in Fig. 5.

To conclude, the arguments presented above illustrate the conditions that must be satisfied if we wish to constrain the value of a particular nonlinear convex average (3) to within a small window: not only should the classical simulation be accurate, but also the conditional states of the true ensemble should be close to pure—otherwise, it will be impossible to make the upper bound tight.

Purification transition.—One case where bounds on the global purity of the system are required is when probing purification dynamics [9], and so this case deserves special attention. Here, an initially mixed state is subjected to a hybrid unitary-projective circuit, and one asks how the global purity increases over time. Since the states of interest are necessarily mixed even in the idealized case where the hybrid circuit is noiseless, one may conclude on the basis of Theorem 1 that purification dynamics cannot be probed experimentally. This is indeed the case if one actually prepares maximally mixed states as the inputs to the circuit; however, if one instead uses a maximally entangled state between the system and a set of ancillas as the initial state then the global state will be pure. The purity of the system can then be thought of as the purity of a subsystem of the global state, which can be bounded using the approach outlined in Sec. VI B.

Conveniently, using methods based on classical shadow tomography one can avoid having to use any ancillas in an actual experiment: by preparing randomized initial states and measuring the output states in random bases, it is possible to construct a classical shadow of the Choi state describing the hybrid dynamics. This Choi state is precisely the global state one would obtain by performing conventional shadow tomography on a purified system-plus-ancilla state. See, for instance, Refs. [67,68].

B. Computational cost of classical simulation

So far, we have considered the classical states ρ_z^C in quite general terms, without specifying how (and if) the states ρ_z^C are constructed. Recall that the specific computational task required by our protocol is that upon obtaining an outcome z in the experiment, we must compute and store a state ρ_z^C, which we construct on the basis of some model of how we expect the quantum device to behave. Whether or not this task is (even approximately) achievable is an intricate question, and has been discussed to

some extent in other recent works where quantum-classical correlators have been introduced [13,27–32]. Here we explore some examples where we can answer this question in the affirmative or negative.

Most obviously, some clear examples where such a classical simulation can be readily performed are in few-body systems, or many-body systems where the full underlying dynamics is efficiently simulable, e.g., for Clifford circuits or free fermionic systems. Going beyond these simple cases, the first issue to address is whether a representation of the simulated state can be stored using a scalable amount of classical memory (irrespective of whether it can be computed or not). One case where this is certainly possible is when the conditional states are defined on a $\mathcal{O}(1)$ number of qubits, even if the full quantum device is a many-body system. For example, in the projected ensemble [14–17], or teleportation transition [25], we can choose to measure many qubits, leaving only a few unmeasured. If, on the other hand, the postmeasurement states consist of extensively many qubits then the states themselves must have some appropriate structure that allows them to be represented using some variational ansatz, e.g., a tensor network.

In terms of the simulation strategy, the approach chosen will naturally depend on the specifics of the system, and so to make any further statements, we must specialize to particular types of measurement-induced dynamics. One example where full simulation can be done efficiently is in the projected ensemble for a wave function generated by evolving a product state under some finite-time dynamics in one spatial dimension (see, e.g., Refs. [14,64]). There, the state of the unmeasured qubits conditioned on a particular measurement outcome can be constructed using a transfer matrix technique, the cost of which is linear in system size for any fixed time of evolution. Since the convergence towards emergent state designs is exponential in time for one-dimensional (1D) systems [64], this means that, for the scheme described in Sec. VII B, the computational cost can remain small. The same goes for many-body teleportation protocols in one dimension [25]. If the time of evolution were to scale with system size, or we move to higher dimensions, then efficient simulation may not be possible, although depending on the circuit being simulated there may still be viable options; see, e.g., Ref. [69].

For the case of one-dimensional hybrid quantum circuits, results on the measurement-induced phase transition indicate that the dynamics is efficiently simulable in the area-law phase using matrix product state–based techniques, but not in the volume-law phase [5,70,71]. This suggests that experimentally learning properties of the postmeasurement ensemble of conditional states without employing brute-force postselection is only possible in one phase. In the volume-law phase, where simulation is presumably not possible, one cannot use the scheme described in this paper in practice, and we must instead ask whether anything nontrivial can be learned in the absence of a simulation; see Sec. VIII C below. Put simply, in these cases there exists a set of measurable properties (8) that could in principle be used to determine some ensemble property of interest, but to determine the correct operators A_z is a computationally intractable task. See also Ref. [34], where the question of whether entanglement can be probed in area-versus volume-law phases is considered.

Beyond these examples, it is interesting to consider whether there may be scenarios in which the computational task required here (calculating conditional states for specific values of z) can be done efficiently, even if full simulation, which would also involve sampling from the probability distribution—a potentially hard task classically [72]—is not. One might also consider the possibility of using a second quantum device to perform the simulation itself: this certainly would be possible for the projected ensemble in 1D dual-unitary circuits by leveraging space-time rotation [73], but whether this can be done efficiently in other cases is unclear. Finally, we highlight that in settings where the dynamics features some particular structure, such as a continuous global symmetry, then it may be possible to make concrete inferences about the states in question by employing simulations that only use partial knowledge about the dynamics (e.g., the distribution of gates employed rather than the exact gates chosen), which are computationally scalable. Indeed, this was applied in Refs. [28,74] to demonstrate the feasibility of probing charge-sharpening transitions in hybrid U(1)-symmetric circuits. We leave these ideas to future work.

C. Can simulations be avoided?

Having highlighted the fact that constructing the simulated states ρ_z^C may not always be possible, this raises the question of what can be learned without any advanced knowledge of the structure of the conditional states. To address this question in a sharp way, we consider a thought experiment, the implication of which is that if we cannot perform such a simulation, no information can be gained about the ensemble beyond the properties of the averaged state (2) using a reasonable number of experimental repetitions M.

Let us introduce the following hypothetical scenario: we are given access to a device (oracle) that, when queried, outputs a label z and state ρ_z^Q sampled from an ensemble. The ensemble realized by the oracle is one of the following, chosen with equal *a priori* probability:

$$\mathcal{E}_1^Q = \{(p_z, \rho_z^Q)\}_{z \in \mathcal{Z}}, \tag{76a}$$

$$\mathcal{E}_2^Q = \{(p_z, \langle \rho^Q \rangle)\}_{z \in \mathcal{Z}}, \tag{76b}$$

with $\langle \rho^Q \rangle$ the average state (2). After querying the oracle M times, we are asked to determine whether the samples

came from $\mathcal{E}_1^Q$ or $\mathcal{E}_2^Q$. We want to consider this problem because if we cannot reliably distinguish these two ensembles then we cannot hope to learn any properties of $\mathcal{E}_1^Q$ besides those that depend only on $\langle \rho^Q \rangle$.

For the present purpose, we broadly define a simulation as any method that allows us to obtain some (partial) prior knowledge of how the labels z map onto properties of the individual states ρ_z^Q. Thus, in the scenario we are thinking about where we do not have access to a simulation, we can only employ strategies that treat each z on an equal footing, i.e., we cannot exploit of any contextual information about what the labels z represent. To make this idea concrete, we define a simulation-free strategy as one that works the same if all labels z output by the oracle are first permuted by some arbitrary $\tau \in \Sigma_{|\mathcal{Z}|}$. In Appendix C, we prove the following result.

Theorem 2. Given coherent access to M independent samples of quantum states from one of the two ensembles $\mathcal{E}_1^Q$ or $\mathcal{E}_2^Q$ [Eqs. (76)], chosen with equal a priori probability, then any simulation-free strategy to distinguish $\mathcal{E}_1^Q$ from $\mathcal{E}_2^Q$ succeeds with probability no greater than

$$p_{\text{succ}} \leq \frac{1}{2} + M^2 \sum_z p_z^2. \tag{77}$$

Thus, in the regime where $M \lesssim 2^{H^{(2)}[p_z]}$, with $H^{(2)}[p_z] := -\log_2(\sum_z p_z^2)$ the collision entropy of distribution p_z, which typically scales linearly with the number of measurements, we cannot reliably distinguish a given ensemble $\mathcal{E}_1^Q$ from one where the average state $\langle \rho^Q \rangle$ is supplied independently of z. As a consequence, we cannot hope to learn any property of an ensemble that is not contained within the average state (2), unless we obtain a number of samples that scales exponentially with the collision entropy of p_z. This is even the case when we can access all M copies of the state simultaneously, which encompasses situations where we employ adaptive measurement strategies, i.e., protocols where the gates and measurements that we apply can be chosen in a way that depends on measurement outcomes that occurred earlier. Hence, Theorem 2 establishes that the postselection problem can only be avoided if we incorporate some prior knowledge about the states ρ_z^Q into our learning strategy.

IX. CONCLUSION AND OUTLOOK

In this work, we have introduced a scheme that allows one to infer properties of an ensemble of quantum states generated by dynamics that involve measurements. We avoid postselection by insisting that the quantities we directly measure from the experiment are "estimable properties" [Eq. (8)], which can be computed using a scalable number of experimental repetitions. Then, information about an ensemble averaged property (3) can be indirectly inferred by solving an optimization problem, as defined in Sec. V. Our method gives one a lower and upper bound for the desired quantity, which can be made narrow by employing simulations of the quantum device on a classical computer. Crucially, the conclusions we make regarding the ensemble of states generated by the quantum device are not contingent on any *a priori* assumptions about the accuracy of the simulation: any bounds we construct are entirely rigorous.

Our results clearly have immediate implications for near-term experiments that probe measurement-induced dynamics. In Sec. VII B, we saw how the method used here can be used to infer the emergence of quantum state designs in the projected ensemble. For other types of measurement-induced phenomena beyond this example, once the relevant order parameters and/or figures of merit are identified, and corresponding inequalities of the kind presented in Sec. VI are derived, one can immediately start using these bounds in the spirit of this work to indirectly infer the value of this quantity realized in some experiment. In particular, by virtue of having postmeasurement states defined on a small number of qubits, we anticipate that experiments for witnessing many-body teleportation transitions [13,25] should be ideal settings for the application of our method.

Thinking more generally, the arguments presented in Sec. VIII point to a sharp distinction between scenarios where properties of the ensemble of postmeasurement conditional states can or cannot be inferred from experimental data. One such condition is on the states being probed: if these are not close to being pure then there will inevitably be some uncertainty in the value of the desired quantity. The other pertains to the feasibility of classical simulation: if simulation is not possible then the states generated by the dynamics are indistinguishable from the case where the average state (2) is realized every time. The implication is that if we hope to use the conditional states ρ_z^Q as resources for some useful task then we must necessarily have some prior knowledge about their structure, i.e., some model to guide our expectation for how ρ_z^Q depends on z. This basic expectation should be borne in mind when considering possible extrinsic applications of measurement-induced physics in future studies.

While the approach we introduce here is designed with the aim to characterize properties of the postmeasurement states, it is helpful to make comparisons to other scenarios where one wishes to make other kinds of inferences about processes that generate both quantum and classical data. One example that fits particularly well into this category is the problem of parameter estimation for continuously monitored few-body systems [75–79], which has applications in quantum metrology. There, one again has to deal with the fact that every repetition of the experiment results in a different nondeterministic outcome, which can

be handled by using a parallel classical simulation of the dynamics conditioned on those quantum trajectories that arise in the experiment. A key difference is that in parameter estimation, a particular starting assumption is made that the generator of dynamics is given by some model with a fixed number of unknown parameters, which we wish to infer based on the outcomes of the measurements we make. Because of this strong assumption, inferences can be made purely from the observed distribution of outcomes, without measuring the state of the system at the end of the experiment. In cases where such model-based assumptions can be reliably applied, it may well be possible to incorporate these into our scheme through modification of the space $\mathcal{K}$, and this might perhaps allow one to overcome some of the limitations of Theorem 1. In fact, in the even more structured setting where one wishes to learn the best candidate out of a set of hypotheses for a quantum-classical process, general bounds on the required sample complexity are known [80,81], and so insights from these works could suggest strategies for overcoming these limitations.

One type of scenario that has not been explicitly considered in this work is where the outcomes of measurements are used to determine some subsequent dynamics, i.e., feedback is employed. In such scenarios, which include error correction [1] and measurement-based computation [2,3] as special cases, nontrivial behavior can arise even in the averaged state [19,82–87], and this can be probed in experiment using conventional learning techniques (though one should note that the physics of the averaged state is still distinct from those of the individual trajectories ρ_z^Q [86–88]). Still, if one is interested in using such "interactive dynamics" to generate states with some desired property then the choice of feedback applied could in principle be made based on some knowledge of the postmeasurement states prior to the conditional unitary operation. The scheme introduced here can in principle be used to characterize these "prefeedback" states, which can then be used to make a constructive choice of feedback algorithm. Indeed, this type of experimental flow arises in certain proposals of measurement-based computation [89].

ACKNOWLEDGMENTS

APPENDIX A: PROOF OF THEOREM 1

To establish Theorem 1, we need the following result, which is a restatement of claim (5.8) in Ref. [90].

Lemma 1 (Dubins [90]). Let $\mathcal{K}_0$ be a linearly bounded convex subset of a finite-dimensional vector space $\mathcal{M}$. Let $\mathcal{A}$ be an intersection of R hyperplanes in $\mathcal{M}$, i.e., a linear subspace of codimension R. Then the extreme points of the intersection $\mathcal{K} = \mathcal{K}_0 \cap \mathcal{A}$ are elements of the R-skeleton of $\mathcal{K}_0$, namely, the union of all faces of $\mathcal{K}$ of dimension less than or equal to R.

When applied to the feasible space $\mathcal{K}$ defined in Eq. (23), we can reduce our problem to a study of the R-skeleton of $\mathcal{K}_0$. Recall that a face of a convex set $\mathcal{K}_0$ is a subset $\mathcal{F} \subset \mathcal{K}_0$ with the property that if $x \in \mathcal{F}$ is a convex combination of two other elements $x = \lambda x' + (1 - \lambda)x''$, where $x', x'' \in \mathcal{K}_0$, then this implies that $x', x'' \in \mathcal{F}$. All faces have a dimension, which is the dimension of the smallest affine set containing it—for example, the faces of dimension zero are singletons, each containing an extreme point of $\mathcal{K}_0$. Here we are interested in the faces of dimension at most R.

The set $\mathcal{K}_0$ in question is a Cartesian product of $|\mathcal{Z}|$ copies of the space of density matrices $\mathcal{D}$ over a Hilbert space of dimension d. Then, our first step is to show that the faces of $\mathcal{K}_0$ are products of faces of $\mathcal{D}$. To demonstrate this, consider a face $\mathcal{F} \subset \mathcal{C}$, where $\mathcal{C} = \mathcal{C}_1 \times \mathcal{C}_2$, with $\mathcal{C}_{1,2}$ arbitrary convex sets. For a given $x_1 \in \mathcal{C}_1$, denote by $\mathcal{F}_2(x_1) \subset \mathcal{C}_2$ the set of points x_2 such that the pair (x_1, x_2) is an element of $\mathcal{F}$. Note that $\mathcal{F}_2(x_1)$ is a face of $\mathcal{C}_2$, since if $x_2 = \alpha y_2 + (1 - \alpha)z_2$ then we have $(x_1, y_2) \in \mathcal{F}$ and $(x_1, z_2) \in \mathcal{F}$, which, by definition, implies that $y_2, z_2 \in \mathcal{F}_2(x_1)$. Now take two points x_1, x_1' for which $\mathcal{F}_2(x_1)$ is nonempty. By convexity of $\mathcal{F}$, if $x_2 \in \mathcal{F}_2(x_1)$ and $x_2' \in \mathcal{F}_2(x_1')$ then, for any $\alpha \in [0, 1]$, we have $[\alpha x_1 + (1 - \alpha)x_1', \alpha x_2 + (1 - \alpha)x_2'] \in \mathcal{F}$. Then, by the definition of a face we have $x_2, x_2' \in \mathcal{F}_2[\alpha x_1 + (1 - \alpha)x_1']$. This implies that $\mathcal{F}_2(x_1) = \mathcal{F}_2(x_1')$, and hence, for any x_1, $\mathcal{F}_2(x_1)$ is either empty or equal to a fixed set $\mathcal{F}_2$, and we conclude that $\mathcal{F} = \mathcal{F}_1 \times \mathcal{F}_2$, where $\mathcal{F}_2$ is a face of $\mathcal{C}_2$. The same logic applied to the analogously defined set $\mathcal{F}_1(x_2)$ gives us that $\mathcal{F}_1$ must be a face of $\mathcal{C}_1$. If we apply this to the multiple product $\mathcal{K}_0 = \times_{z \in \mathcal{Z}} \mathcal{D}$, we conclude that faces of $\mathcal{K}_0$ are products of $|\mathcal{Z}|$ faces of $\mathcal{D}$, as claimed above.

With this understood, note that the n-dimensional faces of $\mathcal{K}_0$ are products of faces of $\mathcal{D}$ whose dimensions sum to n. Thus, at least $(|\mathcal{Z}| - n)$ factors of any n-dimensional face of $\mathcal{K}_0$ will be zero dimensional. Since a zero-dimensional face is an extreme point, and the extreme points of $\mathcal{D}$ are pure states [91], then any element of an n-dimensional face of $\mathcal{K}_0$ is an ensemble with no more than n nonpure states. We conclude that extreme points of $\mathcal{K}$, which according to the lemma quoted above belong to the R-skeleton of $\mathcal{K}_0$, must have no more than R nonpure states, thus completing the proof of Theorem 2.

The corollary quoted in the main text follows since $\mathcal{K}$ is a compact, convex subset of $\mathcal{M} = \bigoplus_{z \in \mathcal{Z}} \mathcal{B}(\mathcal{H})$, which is a finite-dimensional vector space, and hence the Krein-Milman theorem applies. Because $\mathcal{K}$ is necessarily nonempty, at least one extreme point of $\mathcal{K}$ must exist.

APPENDIX B: DERIVATIONS OF LOWER AND UPPER BOUNDS

In Sec. VI, we outlined how approximate solutions to the relevant optimization problems could be found for various specific quantities. Here we include some of the details of these derivations.

1. Quadratic observables

Here we prove the lower bound (49), which is obtained using a similar line of reasoning to that discussed in Sec. VI A. We start with the case where the superoperator $\mathcal{N} = |O\rangle\rangle\langle\langle O|$, i.e., we are considering the average of $G(\rho_z^Q) = \text{Tr}[\rho_z^Q O]^2$. The case of general positive semidefinite $\mathcal{N} \succeq 0$ will follow quickly from this case.

We can assume that $\text{Tr}[O] = 0$ without loss of generality, since any component of O proportional to $\mathbb{I}$ can be subtracted off, yielding an additional term $\propto \text{Tr}[\langle\rho^Q\rangle O]$, which is a property of the average state and hence can be measured directly. Then, the dual function $h_-(\lambda_i)$ corresponding to this particular primal problem takes the form

$$h_-(\lambda_i) = \sum_z p_z F_{2,-}\left(O, \sum_i \lambda_i \tilde{A}_z^{(i)}\right) + \sum_i \lambda_i(b_i - a_i),$$
(B1)

where a_i, $\tilde{A}_z^{(i)}$ are as in Eq. (35), and we have defined the function

$$F_{2,-}(O, C) := \inf_{\rho \in \mathcal{D}}(\text{Tr}[\rho O]^2 - \text{Tr}[\rho C])$$
(B2)

by analogy to Eq. (36). While we cannot exactly solve this problem analytically in full generality, it is possible in the case where $\mathcal{H}$ is a single qubit (Hilbert space dimension $d = 2$). Our solution to this case will prove instructive when it comes to treating the more general scenario $d \geq 2$ later; we thus briefly specialize in the following.

Single qubit.—When we fix $d = 2$, since both C and O are traceless, one can rescale and perform rotations in Hilbert space appropriately to map the problem of evaluating $F_{2,-}(O, C)$ to the case $O = Z$, and $C = \alpha_1 X + \alpha_3 Z$, where (X, Y, Z) are Pauli matrices, and $\alpha_1, \alpha_3 \geq 0$ are non-negative scalars. Evidently, α_1 and α_3 are the components of C orthogonal and parallel to O, respectively. This effectively reduces $F_{2,-}$ to a function of two variables (α_1, α_3).

Another useful aspect of the qubit problem that helps us here is that the space of density matrices $\mathcal{D}$ has a simple geometry that can be straightforwardly characterized: the conditions that ρ is Hermitian and satisfies $\text{Tr}[\rho] = 1$, $\text{Tr}[\rho^2] \leq 1$ are necessary and sufficient for ρ to be a valid density matrix. This leads to the notion of the Bloch ball, which is a helpful visualization of the space $\mathcal{D}$: any density matrix can be written as $|\rho\rangle\rangle = (|\mathbb{I}\rangle\rangle + n_1|X\rangle\rangle + n_2|Y\rangle\rangle + n_3|Z\rangle\rangle)/2$, with $\vec{n} = (n_1, n_2, n_3)$ a three-dimensional vector that specifies the state, which belongs to the unit ball $|\vec{n}| \leq 1$. Since the function to be extremized is independent of n_2, we can reduce this to

$$F_{2,-}^{(d=2)}(\alpha_1, \alpha_3) = \inf_{n_1^2 + n_3^2 \leq 1}[n_3^2 - \alpha_1 n_1 - \alpha_3 n_3].$$
(B3)

A solution to this problem can be formally written down in terms of the roots of a certain polynomial equation, but the resulting expression is rather cumbersome and not particularly informative. Rather, it is useful to consider the behavior of the above function in the vicinity of $\alpha_1 = 0$. By considering small perturbations around the $\alpha_1 = 0$ solution, we obtain the expansion

$$F_{2,-}^{(d=2)}(\alpha_1, \alpha_3) = \begin{cases} -\dfrac{\alpha_3^2}{4} - |\alpha_1|\sqrt{1 - \dfrac{\alpha_3^2}{4}} & \text{for } |\alpha_3| \leq 2, \\ 1 - |\alpha_3| & \text{for } |\alpha_3| > 2 \end{cases}$$
$$+ \mathcal{O}(\alpha_1^2),$$
(B4)

which is achieved at

$$n_3 = \text{sgn}(\alpha_3) \times \min(1, |\alpha_3|/2) + \mathcal{O}(\alpha_1).$$
(B5)

Looking at Eq. (B4), one notes that the behavior of $F_{2,-}^{(d=2)}$ in the directions α_1 and α_3 is markedly different near the point $\alpha_{1,3} = 0$: the function decreases linearly along α_1, and quadratically along α_3. Bearing in mind that $\alpha_{1,3}$ are linear combinations of the Lagrange multipliers λ_i, we consider a stability analysis of the function $h_-(\lambda_i)$, which is to be maximized, around the candidate point $\lambda_i = 0$. Qualitatively speaking, we find that it is favorable to increase λ_i away from zero for those i where the operators $A_z^{(i)}$ are predominately along the Z direction in operator space (α_3 dominates), but not those along orthogonal directions (α_1 dominates). This is because, for the former subset of λ_i, the second term in Eq. (B1) will increase faster than the first term decreases as λ_i is varied away from zero.

This rough intuition can be made concrete most straightforwardly if we assume that, for each i, $A_z^{(i)}$ is either orthogonal to or parallel to Z. [Such an assumption is not particularly restrictive—if the $A_z^{(i)}$ are measurable then we can always decompose $A_z^{(i)} = \sum_\mu a_{i\mu z}\sigma^\mu$ in terms of components along some basis of operators σ^μ, which we are free to choose such that one of the σ^μ is proportional to O. Then, using our shadow-based scheme, we can always choose to measure the enlarged set of operators $A_z^{(i,\mu)} := a_{i\mu z}\sigma^\mu$ for each pair (i, μ) separately without losing any information, simply by altering the classical postprocessing.] Let us denote the set of i for which $\langle\langle A_z^{(i)}|Z\rangle\rangle = 0$ as $\mathcal{I}_\perp$, and those for which $A_z^{(i)} = c_{zi}Z$ as $\mathcal{I}_\parallel$. We temporarily fix $\lambda_{i \in \mathcal{I}_\perp} = 0$ by hand and then optimize over $\lambda_{i \in \mathcal{I}_\parallel}$.

Using expansion (B4) with the appropriate replacement of parameters, we obtain

$$
h_-(\lambda_i)|_{\lambda_i \in \mathcal{I}_\perp = 0} = \sum_z p_z F_{2,-}^{(d=2)}\left(0, \sum_{i \in \mathcal{I}_\parallel} \lambda_i c_{zi}\right) + \sum_{i \in \mathcal{I}_\parallel} \lambda_i b_i. \tag{B6}
$$

(Note that we have $a_i = 0$ for $i \in \mathcal{I}_\parallel$.) From this point, our arguments follow a similar structure to those of Sec. VI A. We can use the bound $F_{2,-}^{(d=2)}(0, \alpha_3) \geq -\alpha_3^2/4$ to obtain an analogous expression to Eq. (37), i.e.,

$$
h_-(\lambda_i)|_{\lambda_i \in \mathcal{I}_\perp = 0} \geq -\frac{1}{4} \sum_{i,j \in \mathcal{I}_\parallel} \lambda_i J_{ij} \lambda_j + \sum_{i \in \mathcal{I}_\parallel} b_i \lambda_i, \tag{B7}
$$

where we have introduced J_{ij}, an analogue of L_{ij} in Eq. (38),

$$
J_{ij} = \sum_z p_z c_{zi} c_{zj}. \tag{B8}
$$

[Equation (B7) actually becomes an equality if we have $|\sum_{i \in \mathcal{I}_\parallel} \lambda_i c_{zi}| \leq 2$ for all z.] As before, this can be maximized to obtain a certificate

$$
\max_{\lambda_i : i \in \mathcal{I}_\parallel} h_-(\lambda_i) \geq \sum_{i,j \in \mathcal{I}_\parallel} b_i [L^{-1}]_{ij} b_j =: \tilde{h}_-. \tag{B9}
$$

Again, optimality of this candidate solution to the dual problem is not guaranteed, i.e., $\tilde{h}_-$ might be strictly smaller than the true optimal value g_-^*. Even so, the above can be evaluated straightforwardly using experimental data, and can be used as a lower bound for the quantum-quantum correlator $\mathbb{E}_z \operatorname{Tr}[\rho_z^Q O]^2$.

Beyond a single qubit.—When we generalize beyond the case of a single qubit $d = 2$, the optimization over density matrices in higher-dimensional Hilbert spaces is made more complicated by the geometric structure of the space $\mathcal{D}$, which does not have simple interpretations such as the Bloch ball. For instance, expressions analogous to Eq. (B3) are not so straightforward. Nevertheless, we can carry the intuition gained above forward to this case to anticipate that reasonable bounds can be found by optimizing only λ_i for $i \in \mathcal{I}_\parallel$, where now $\mathcal{I}_\parallel$ denotes the set of is for which $A_z^{(i)} = c_{zi} O$ for scalars c_{iz}, whose complement $\mathcal{I}_\perp$ is made up of operators orthogonal to O in operator space, namely, $\operatorname{Tr}[A_z^{(i)} O] = 0$ for $i \in \mathcal{I}_\perp$.

Following the same logic as before, we set $\lambda_i = 0$ for $i \in \mathcal{I}_\perp$, after which the second argument of the function $F_{2,-}$ in Eq. (B1) becomes proportional to O. The resulting minimization problem (B2) can then be directly evaluated

for arbitrary d,

$$
F_{2,-}(O, \alpha O) = -\|O\|_\infty^2 \times \min(\alpha^2/4, |\alpha|), \tag{B10}
$$

where $\|O\|_\infty$ denotes the spectral norm, equal to the largest singular value of O. This solution has the same structure as the qubit case (B4), and hence the same logic can be used as before to reproduce bound (B9). That is, the same expression (B9) for $\tilde{h}_-$ as a lower bound for g_-^* for arbitrary d.

Now, following the same approach as in Sec. VI A, we can use quantum-classical correlators to choose the operators $A_z^{(i)}$. The simplest case, where we use the standard quantum-classical correlator (16) as the only constraint, yields

$$
\langle O \otimes O \rangle^{QQ} \geq \frac{[\langle O \otimes O \rangle^{QC}]^2}{\langle O \otimes O \rangle^{CC}}, \tag{B11}
$$

which can be interpreted as a result of the Cauchy-Schwartz inequality, or, equivalently, the standard inequality $\operatorname{Var}(X)\operatorname{Var}(Y) \geq \operatorname{Cov}(X, Y)^2$ for classical random variables X, Y. If we make use of quantum-classical correlators beyond just $\langle O \otimes O \rangle^{QC}$ then we can in principle use all the information contained in the superoperators η^{QC}, η^{CC} [Eq. (42)]. This allows us to improve the simple bound above to

$$
\langle O \otimes O \rangle^{QQ} \geq \langle\!\langle O | \eta^{QC} \zeta^{CQ} | O \rangle\!\rangle, \tag{B12}
$$

where ζ^{CQ} is the superoperator implicitly defined in Eq. (44).

Now, for more general convex quadratic observables ($\mathcal{N} \succeq 0$), we can always perform a decomposition of the superoperator $\mathcal{N} = \sum_a \mu_a |C_a\rangle\!\rangle\langle\!\langle C_a|$, where the $\mu_a \geq 0$ are eigenvalues and the C_a are eigenoperators. Because of the non-negativity of μ_a, we can apply bound (B12) with $O = C_a$ for each a separately, which gives Eq. (49).

2. Numerically stable upper bound for the von Neumann entropy

In the main text, we derived a general upper bound for the average von Neumann entropy (56), which, for the special case $A_z = -\log \rho_z^C$, reduces to the quantum-classical entropy introduced in Ref. [34]. Unfortunately, this leads to an expression that is numerically unstable when the classical states ρ_z^C are near singular, due to the need to take the logarithm of the operators ρ_z^C. However, thanks to the flexibility of our optimization-based approach, we can devise a simple solution, where we modify the chosen operators A_z to obtain a better bound. Here we consider the case where we measure the two observables defined in Eq. (59), denoted $A_z^{(1,2)}$. For each of these two observables,

we have a Lagrange multiplier $\lambda_{1,2}$. In terms of the expectation values $\langle A_z^{(1,2)} \rangle$, the dual function can be obtained by minimizing with respect to ρ_z^Q, giving

$$h_+(\lambda_{1,2}) = \mathbb{E}_z(\log\{\text{Tr}[(\Pi_z^> \rho_z^C \Pi_z^>)^{\lambda_1}] + e^{-\lambda_2} \text{rank}(\Pi_z^{\leq})\} + \lambda_1 \langle A_z^{(1)} \rangle + \lambda_2 \langle A_z^{(2)} \rangle)). \quad \text{(B13)}$$

Optimizing for $\lambda_{1,2}$ cannot be done exactly using analytical methods; however, based on the same considerations as for Eq. (57), it is reasonable to set $\lambda_1 = 1$ in the above, after which the minimization over the remaining Lagrange parameter λ_2 can be done approximately. In the limit where $\overline{\delta^Q} := \langle A_z^{(2)} \rangle = \mathbb{E}_z \text{Tr}[\rho_z^Q \Pi_z^{\leq}]$ is small (i.e., the overlap of the quantum states with the near-singular eigenstates of ρ_z^C is small on average), we find that the value $\lambda_2 = \log[\overline{r}(1 - \overline{\delta^Q})/\overline{\delta^Q}(1 - \overline{\delta^C})]$ is approximately optimal, where we use the shorthand $\overline{r} = \mathbb{E}_z r_z = \mathbb{E}_z \text{rank}[\Pi_z^{\leq}]$ and $\overline{\delta^C} = \mathbb{E}_z \delta_z^C = \mathbb{E}_z \text{Tr}[\rho_z^C \Pi_z^{\leq}]$. This gives us a regularized bound

$$\mathbb{E}_z S_{\text{vN}}(\rho_z^Q)$$
$$\leq \sum_z p_z \bigg[-\langle A_z^{(1)} \rangle + \log\bigg(1 - \delta_z^C + \frac{\overline{\delta^Q}(1 - \overline{\delta^C})}{1 - \overline{\delta^Q}} \frac{r_z}{\overline{r}}\bigg)$$
$$- \overline{\delta^Q} \log\bigg(\frac{\overline{\delta^Q}(1 - \overline{\delta^C})}{1 - \overline{\delta^Q}}\bigg)\bigg]. \quad \text{(B14)}$$

The above simplifies considerably when the classical states are pure, where we have $\log' \rho_z^C = 0$, $r_z = (d-1)$, and $\delta_z^C = 0$, whence the result given in Eq. (60).

3. Lower bound for the subsystem von Neumann entropy

In the main text, we showed that the minimization problem for the average von Neumann entropy of a subsystem Q_1 could be turned into a convex optimization problem by way of the conditional entropy (61), which is concave, and hence can be maximized. To make progress, it will be useful to introduce a variational characterization of the conditional quantum entropy [92]

$$S(\rho^Q|Q_1) = \sup_{\sigma^{Q_1} \in \mathcal{D}_1} [-S(\rho^Q \| \sigma^{Q_1} \otimes \mathbb{I}_{Q_2})], \quad \text{(B15)}$$

where $S(\rho\|\sigma) = \text{Tr}[\rho \log \rho] - \text{Tr}[\rho \log \sigma]$ is the quantum relative entropy. After constructing the dual function for maximizing the conditional entropy, we can use the above representation, after which result (55) can be used to

perform the optimization over ρ_z^Q, thus giving

$$\max_{\rho \in \mathcal{K}}[\mathbb{E}_z S(Q_2|Q_1)_{\rho_z^Q}] = \min_{\lambda_i} \bigg\{\sum_i \lambda_i b_i$$
$$+ \max_{\sigma_z^{Q_1} \in \mathcal{D}_1} \mathbb{E}_z \log \text{Tr}\bigg[\exp\bigg(\log \sigma_z^{Q_1} \otimes \mathbb{I}^{Q_2} - \sum_i \lambda_i A_z^{(i)}\bigg)\bigg]\bigg\}. \quad \text{(B16)}$$

Because of the form of Eq. (B15), we have an extra maximization step to perform over the states $\sigma_z^{Q_1}$. This cannot be solved exactly in full generality, but we can use the Golden-Thompson inequality $\text{Tr}[e^{A+B}] \leq \text{Tr}[e^A e^B]$ for Hermitian matrices A, B to obtain bound (63).

APPENDIX C: PROOF OF THEOREM 2

Here we provide a proof of Theorem 2, which is a statement regarding our ability to distinguish between the two ensembles (76). First, we note that having coherent access to M samples from the ensembles $\mathcal{E}_1^Q$, $\mathcal{E}_2^Q$ is equivalent to owning a single copy of the following respective quantum-classical states:

$$\rho_{1,M} = \bigg(\sum_z p_z |z\rangle\langle z| \otimes \rho_z^Q\bigg)^{\otimes M}, \quad \text{(C1a)}$$

$$\rho_{2,M} = \bigg(\sum_z p_z |z\rangle\langle z| \otimes \langle\rho^Q\rangle\bigg)^{\otimes M}. \quad \text{(C1b)}$$

In this representation, the labels are stored in M classical registers, each prepared in a state $|z\rangle\langle z|$, and the conditional states are simultaneously stored in M separate quantum registers. The problem of deciding which of the two scenarios (C1) is realized by the oracle is a form of hypothesis testing, which is a well-studied problem. Without any constraints on our hypothesis testing strategy, a standard result due to Helstrom [93] says that the best strategy is equivalent to performing a generalized measurement described by a two-component POVM $E_1 + E_2 = \mathbb{I}$, where E_1 is the projector onto the space of positive eigenvalues of $\rho_{1,M} - \rho_{2,M}$. When one employs this strategy, one can successfully determine which of the ensembles is realized with probability $p_{\text{succ}} = 1/2 + \|\rho_{1,M} - \rho_{2,M}\|_1/4$, where $\|\rho - \sigma\|_1$ is the trace distance between states ρ, σ. As the number of copies M increases, the success probability tends towards unity exponentially quickly [94,95], and so there always exists a strategy that allows us to distinguish between the ensembles (76) reliably.

However, even though such an optimal hypothesis testing strategy may exist, an observer with limited computational power may not be able to determine what the correct measurement procedure should be. That is, although the ensembles may be distinguishable in an

information-theoretic sense, they may not be computationally distinguishable. Since the question of the computational complexity of finding the optimal strategy will depend on the specifics of the ensemble $\mathcal{E}^Q$ in question, our starting point will be to assume that simulations of the device of any kind are not available. This motivates our definition of a simulation-free strategy as defined in the main text.

In this context, any simulation-free strategy can still be represented as a two-component POVM $E_{1,2}$ acting jointly on the M classical and M quantum registers that store states (C1), in the same way as above—this encompasses protocols that feature measurements in entangled bases, adaptive strategies, and stochastic operations. Regardless, we want to restrict ourselves to operations that are symmetric among the different labels z, i.e., we treat all labels equivalently. This implies that we should make a restriction

$$(\pi_\tau^{\otimes M} \otimes \mathbb{I}_Q) E_1 (\pi_\tau^{\otimes M} \otimes \mathbb{I}_Q) = E_1 \quad \text{for all } \tau \in \Sigma_{|\mathcal{Z}|}, \tag{C2}$$

where $\pi_\tau |z\rangle = |\tau(z)\rangle$ is a permutation operator on a single classical register, and $\mathbb{I}_Q$ is the identity operator on all M quantum registers. Making this restriction forbids us from using some prior knowledge of how the ensemble states ρ_z^Q depend on the states z, i.e., a means to simulate the dynamics.

Whilst ensuring that Eq. (C2) holds, we want to maximize the success probability $p_{\text{succ}} = 1/2 + \text{Tr}[E_1(\rho_{1,M} - \rho_{2,M})]/2$, as in standard hypothesis testing. Evidently, for any such operator E_1, we have $\text{Tr}[E_1 \sigma] = \text{Tr}\{E_1(\mathcal{S} \otimes \text{id}_Q)[\sigma]\}$, where id_Q is the identity superoperator on the quantum register, and

$$\mathcal{S}[\sigma^C] := \frac{1}{|\mathcal{Z}|!} \sum_{\tau \in \Sigma_{|\mathcal{Z}|}} \pi_\tau^{\otimes M} \sigma^C \pi_\tau^{\otimes M} \tag{C3}$$

projects the operator σ^C corresponding to the classical registers onto the subspace that is invariant under permutations of the labels z. Using this construction, it is straightforward to show that

$$p_{\text{succ}} = \frac{1}{2} + \frac{1}{2} \|(\mathcal{S} \otimes \text{id}_Q)[\rho_{1,M} - \rho_{2,M}]\|_1. \tag{C4}$$

Now, we can write the argument of the trace norm as

$$(\mathcal{S} \otimes \text{id}_Q)[\rho_{1,M} - \rho_{2,M}] = \sum_{\vec{z}} p(\vec{z}) \mathcal{S}[|\tau(\vec{z})\rangle \langle \tau(\vec{z})|]$$

$$\otimes \left(\bigotimes_{i=1}^M \rho_{z_i}^Q - \langle \rho^Q \rangle^{\otimes M} \right), \tag{C5}$$

where we used the shorthand $p(\vec{z}) = \prod_{i=1}^M p_{z_i}$ and $|\tau(\vec{z})\rangle = |\tau(z_1)\rangle \otimes \cdots \otimes |\tau(z_M)\rangle$. Now, for a string of labels $\vec{z}$,

define the matrix $R_{ij}(\vec{z}) = \delta_{z_i, z_j}$, which specifies which pairs (z_i, z_j) are equal for the given string $\vec{z}$. It holds that $\mathcal{S}[|\vec{z}\rangle \langle \vec{z}|]$ is equal to $\mathcal{S}[|\vec{z}\,'\rangle \langle \vec{z}\,'|]$ if $R_{ij}(\vec{z}) = R_{ij}(\vec{z}\,')$ for all $i,j = 1, \ldots, M$, and they are orthogonal otherwise. The trace norm in Eq. (C4) then becomes a sum over contributions from each possible value of R_{ij},

$$\|(\mathcal{S} \otimes \text{id}_Q)[\rho_{1,M} - \rho_{2,M}]\|_1$$

$$= \sum_{R_{ij} \in \{0,1\}^{\times M^2}} \left\| \sum_{\vec{z} \in R_{ij}} p(\vec{z}) \left(\bigotimes_{i=1}^M \rho_{z_i}^Q - \langle \rho^Q \rangle^{\otimes M} \right) \right\|_1, \tag{C6}$$

where the notation $\vec{z} \in R_{ij}$ denotes a sum over all strings $\vec{z}$ for which $R_{ij}(\vec{z}) = R_{ij}$. At this point, we can start to apply some upper bounds. We separate out the term in the sum for which $R_{ij} = 0$, and by the triangle inequality we have

$$\|(\mathcal{S} \otimes \text{id}_Q)[\rho_{1,M} - \rho_{2,M}]\|_1$$

$$\leq \sum_{z \notin 0} p(\vec{z}) + \left\| \sum_{\vec{z} \in 0} p(\vec{z}) \left(\bigotimes_{i=1}^M \rho_{z_i}^Q - \langle \rho^Q \rangle^{\otimes M} \right) \right\|_1. \tag{C7}$$

Now we use the definition of the average state, which implies that $\sum_{\vec{z}} p(\vec{z})(\bigotimes_{i=1}^M \rho_{z_i}^Q) = \langle \rho^Q \rangle^{\otimes M}$. Hence, we can infer that

$$\|(\mathcal{S} \otimes \text{id}_Q)[\rho_{1,M} - \rho_{2,M}]\|_1$$

$$\leq \sum_{z \notin 0} p(\vec{z}) + \left\| \sum_{\vec{z} \notin 0} p(\vec{z}) \bigotimes_{i=1}^M \rho_{z_i}^Q \right\|_1, \tag{C8}$$

from which we obtain

$$p_{\text{succ}} \leq \frac{1}{2} + \sum_{\vec{z} \notin 0} p(\vec{z}). \tag{C9}$$

The second term in the above corresponds to the probability of at least one pair of labels (z_i, z_j) being equal. A simple upper bound for this probability is $\binom{M}{2} \sum_z p_z^2$ [96], and hence we obtain the bound quoted in Theorem 2.

---

[1] D. A. Lidar and T. A. Brun, *Quantum Error Correction* (Cambridge university press, Cambridge, 2013).

[2] R. Raussendorf and H. J. Briegel, A one-way quantum computer, Phys. Rev. Lett. **86**, 5188 (2001).

[3] R. Raussendorf, D. E. Browne, and H. J. Briegel, Measurement-based quantum computation on cluster states, Phys. Rev. A **68**, 022312 (2003).

[4] Y. Li, X. Chen, and M. P. A. Fisher, Quantum Zeno effect and the many-body entanglement transition, Phys. Rev. B **98**, 205136 (2018).

[5] B. Skinner, J. Ruhman, and A. Nahum, Measurement-induced phase transitions in the dynamics of entanglement, Phys. Rev. X **9**, 031009 (2019).

[6] A. Chan, R. M. Nandkishore, M. Pretko, and G. Smith, Unitary-projective entanglement dynamics, Phys. Rev. B **99**, 224307 (2019).

[7] Y. Li, X. Chen, and M. P. A. Fisher, Measurement-driven entanglement transition in hybrid quantum circuits, Phys. Rev. B **100**, 134306 (2019).

[8] M. Szyniszewski, A. Romito, and H. Schomerus, Entanglement transition from variable-strength weak measurements, Phys. Rev. B **100**, 064204 (2019).

[9] M. J. Gullans and D. A. Huse, Dynamical purification phase transition induced by quantum measurements, Phys. Rev. X **10**, 041020 (2020).

[10] Y. Bao, S. Choi, and E. Altman, Theory of the phase transition in random unitary circuits with measurements, Phys. Rev. B **101**, 104301 (2020).

[11] R. Fan, S. Vijay, A. Vishwanath, and Y.-Z. You, Self-organized error correction in random unitary circuits with measurement, Phys. Rev. B **103**, 174309 (2021).

[12] C. Noel, P. Niroula, D. Zhu, A. Risinger, L. Egan, D. Biswas, M. Cetina, A. V. Gorshkov, M. J. Gullans, D. A. Huse, and C. Monroe, Measurement-induced quantum phases realized in a trapped-ion quantum computer, Nat. Phys. **18**, 760 (2022).

[13] J. C. Hoke, *et al.*, Measurement-induced entanglement and teleportation on a noisy quantum processor, Nature **622**, 481 (2023).

[14] W. W. Ho and S. Choi, Exact emergent quantum state designs from quantum chaotic dynamics, Phys. Rev. Lett. **128**, 060601 (2022).

[15] P. W. Claeys and A. Lamacraft, Emergent quantum state designs and biunitarity in dual-unitary circuit dynamics, Quantum **6**, 738 (2022).

[16] J. Choi, A. L. Shaw, I. S. Madjarov, X. Xie, R. Finkelstein, J. P. Covey, J. S. Cotler, D. K. Mark, H.-Y. Huang, A. Kale, *et al.*, Preparing random states and benchmarking with many-body quantum chaos, Nature **613**, 468 (2023).

[17] J. S. Cotler, D. K. Mark, H.-Y. Huang, F. Hernández, J. Choi, A. L. Shaw, M. Endres, and S. Choi, Emergent quantum state designs from individual many-body wave functions, PRX Quantum **4**, 010311 (2023).

[18] M. Lucas, L. Piroli, J. De Nardis, and A. De Luca, Generalized deep thermalization for free fermions, Phys. Rev. A **107**, 032215 (2023).

[19] L. Piroli, G. Styliaris, and J. I. Cirac, Quantum circuits assisted by local operations and classical communication: Transformations and phases of matter, Phys. Rev. Lett. **127**, 220503 (2021).

[20] R. Verresen, N. Tantivasadakarn, and A. Vishwanath, Efficiently preparing Schrödinger's cat, fractons and non-Abelian topological order in quantum devices, arXiv:2112.03061.

[21] N. Tantivasadakarn, R. Thorngren, A. Vishwanath, and R. Verresen, Long-range entanglement from measuring symmetry-protected topological phases, arXiv:2112.01519.

[22] T.-C. Lu, L. A. Lessa, I. H. Kim, and T. H. Hsieh, Measurement as a shortcut to long-range entangled quantum matter, PRX Quantum **3**, 040337 (2022).

[23] M. Iqbal, N. Tantivasadakarn, T. M. Gatterman, J. A. Gerber, K. Gilmore, D. Gresh, A. Hankin, N. Hewitt, C. V. Horst, M. Matheny, T. Mengle, B. Neyenhuis, A. Vishwanath, M. Foss-Feig, R. Verresen, and H. Dreyer, Non-Abelian topological order and anyons on a trapped-ion processor, Nature **626**, 505 (2024).

[24] M. Foss-Feig, A. Tikku, T.-C. Lu, K. Mayer, M. Iqbal, T. M. Gatterman, J. A. Gerber, K. Gilmore, D. Gresh, A. Hankin, N. Hewitt, C. V. Horst, M. Matheny, T. Mengle, B. Neyenhuis, H. Dreyer, D. Hayes, T. H. Hsieh, and I. H. Kim, Experimental demonstration of the advantage of adaptive quantum circuits, arXiv:2302.03029.

[25] Y. Bao, M. Block, and E. Altman, Finite-time teleportation phase transition in random quantum circuits, Phys. Rev. Lett. **132**, 030401 (2024).

[26] J. M. Koh, S.-N. Sun, M. Motta, and A. J. Minnich, Measurement-induced entanglement phase transition on a superconducting quantum processor with mid-circuit readout, Nat. Phys. **19**, 1314 (2023).

[27] Y. Li and M. P. A. Fisher, Decodable hybrid dynamics of open quantum systems with $\mathbb{Z}_2$ symmetry, Phys. Rev. B **108**, 214302 (2023).

[28] F. Barratt, U. Agrawal, A. C. Potter, S. Gopalakrishnan, and R. Vasseur, Transitions in the learnability of global charges from local measurements, Phys. Rev. Lett. **129**, 200602 (2022).

[29] J. Y. Lee, W. Ji, Z. Bi, and M. P. A. Fisher, Decoding measurement-prepared quantum phases and transitions: from Ising model to gauge theory, and beyond, arXiv:2208.11699.

[30] Y. Li, Y. Zou, P. Glorioso, E. Altman, and M. P. A. Fisher, Cross entropy benchmark for measurement-induced phase transitions, Phys. Rev. Lett. **130**, 220404 (2023).

[31] X. Feng, B. Skinner, and A. Nahum, Measurement-induced phase transitions on dynamical quantum trees, PRX Quantum **4**, 030333 (2023).

[32] S. J. Garratt, Z. Weinstein, and E. Altman, Measurements conspire nonlocally to restructure critical quantum states, Phys. Rev. X **13**, 021026 (2023).

[33] M. J. Gullans and D. A. Huse, Scalable probes of measurement-induced criticality, Phys. Rev. Lett. **125**, 070606 (2020).

[34] S. J. Garratt and E. Altman, Probing post-measurement entanglement without post-selection, arXiv:2305.20092.

[35] P. Hayden, D. Leung, P. W. Shor, and A. Winter, Randomizing quantum states: Constructions and applications, Commun. Math. Phys. **250**, 371 (2004).

[36] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Randomized benchmarking of quantum gates, Phys. Rev. A **77**, 012307 (2008).

[37] J. Radhakrishnan, M. Rötteler, and P. Sen, Random measurement bases, quantum state distinction and applications to the hidden subgroup problem, Algorithmica **55**, 490 (2009).

[38] F. G. S. L. Brandão, A. W. Harrow, and M. Horodecki, Efficient quantum pseudorandomness, Phys. Rev. Lett. **116**, 170502 (2016).

[39] M. McGinley and M. Fava, Shadow tomography from emergent state designs in analog quantum simulators, Phys. Rev. Lett. **131**, 160601 (2023).

[40] X. Cao, A. Tilloy, and A. D. Luca, Entanglement in a fermion chain under continuous monitoring, SciPost Phys. **7**, 024 (2019).

[41] X. Chen, Y. Li, M. P. A. Fisher, and A. Lucas, Emergent conformal symmetry in nonunitary random dynamics of free fermions, Phys. Rev. Res. **2**, 033017 (2020).

[42] O. Alberton, M. Buchhold, and S. Diehl, Entanglement transition in a monitored free-fermion chain: from extended criticality to area law, Phys. Rev. Lett. **126**, 170602 (2021).

[43] Y. Li, X. Chen, A. W. W. Ludwig, and M. P. A. Fisher, Conformal invariance and quantum nonlocality in critical hybrid circuits, Phys. Rev. B **104**, 104305 (2021).

[44] H. Lóio, A. De Luca, J. De Nardis, and X. Turkeshi, Purification timescales in monitored fermions, Phys. Rev. B **108**, L020306 (2023).

[45] S. Leontica and M. McGinley, Purification dynamics in a continuous-time hybrid quantum circuit model, Phys. Rev. B **108**, 174308 (2023).

[46] J. Dalibard, Y. Castin, and K. Mølmer, Wave-function approach to dissipative processes in quantum optics, Phys. Rev. Lett. **68**, 580 (1992).

[47] H. M. Wiseman and G. J. Milburn, Interpretation of quantum jump and diffusion processes illustrated on the Bloch sphere, Phys. Rev. A **47**, 1652 (1993).

[48] A. Ambainis and J. Emerson, Quantum t-designs: t-wise independence in the quantum world, *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)* (IEEE, San Diego, CA, USA, 2007), pp. 129–140.

[49] A. Bose and S. Chatterjee, *U-Statistics, Mm-Estimators and Resampling* (Springer, Singapore, 2018).

[50] T. S. Ferguson, *U*-statistics, lecture notes for statistics 200B, UCLA (2003), Accessed August 2023. https://stt.msu.edu/Academics/ClassPages/uploads/SS20/953-1/Ustat.pdf.

[51] P. R. Halmos, The theory of unbiased estimation, Ann. Math. Stat. **17**, 34 (1946).

[52] P. J. Bickel and E. L. Lehmann, Unbiased estimation in convex families, Ann. Math. Stat. **40**, 1523 (1969).

[53] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, Nat. Phys. **16**, 1050 (2020).

[54] A. J. Scott, Tight informationally complete quantum measurements, J. Phys. A **39**, 13507 (2006).

[55] A. Elben, S. T. Flammia, H.-Y. Huang, R. Kueng, J. Preskill, B. Vermersch, and P. Zoller, The randomized measurement toolbox, Nat. Rev. Phys. **5**, 9 (2023).

[56] M. Paini and A. Kalev, An approximate description of quantum states, arXiv:1910.10543.

[57] M. C. Tran, D. K. Mark, W. W. Ho, and S. Choi, Measuring arbitrary physical properties in analog quantum simulation, Phys. Rev. X **13**, 011049 (2023).

[58] A. E. Rastegin, Some general properties of unified entropies, J. Stat. Phys. **143**, 1120 (2011).

[59] A. L. Yuille and A. Rangarajan, The Concave-Convex Procedure, in *Advances in Neural Information Processing Systems*, Vol. 14, edited by T. Dietterich, S. Becker, and Z. Ghahramani (MIT Press, Cambridge, MA, 2001). http://papers.nips.cc/paper/2125-the-concave-convex-procedure-cccp.pdf.

[60] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, Cambridge, 2004).

[61] M. A. Nielsen and I. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2010).

[62] D. Petz, *Quantum Information Theory and Quantum Statistics* (Springer Berlin, 2008).

[63] D. A. Roberts and B. Yoshida, Chaos and complexity by design, J. High Energy Phys. **2017**, 1 (2017).

[64] M. Ippoliti and W. W. Ho, Dynamical purification and the emergence of quantum state designs from the projected ensemble, PRX Quantum **4**, 030322 (2023).

[65] J. M. Renes, R. Blume-Kohout, A. J. Scott, and C. M. Caves, Symmetric informationally complete quantum measurements, J. Math. Phys. **45**, 2171 (2004).

[66] H. Kim and D. A. Huse, Ballistic spreading of entanglement in a diffusive nonintegrable system, Phys. Rev. Lett. **111**, 127205 (2013).

[67] S. Chen, W. Yu, P. Zeng, and S. T. Flammia, Robust shadow estimation, PRX Quantum **2**, 030348 (2021).

[68] M. McGinley, S. Leontica, S. J. Garratt, J. Jovanovic, and S. H. Simon, Quantifying information scrambling via classical shadow tomography on programmable quantum simulators, Phys. Rev. A **106**, 012441 (2022).

[69] J. C. Napp, R. L. La Placa, A. M. Dalzell, F. G. S. L. Brandão, and A. W. Harrow, Efficient classical simulation of random shallow 2D quantum circuits, Phys. Rev. X **12**, 021021 (2022).

[70] H. Dehghani, A. Lavasani, M. Hafezi, and M. J. Gullans, Neural-network decoders for measurement induced phase transitions, Nat. Commun. **14**, 2918 (2023).

[71] R. Suzuki, J. Haferkamp, J. Eisert, and P. Faist, Quantum complexity phase transitions in monitored random circuits, arXiv:2305.15475.

[72] B. M. Terhal and D. P. DiVincenzo, Adptive quantum computation, constant depth quantum circuits and Arthur-Merlin games, Quantum Inf. Comput. **4**, 134 (2004).

[73] M. Ippoliti and V. Khemani, Postselection-free entanglement dynamics via spacetime duality, Phys. Rev. Lett. **126**, 060501 (2021).

[74] U. Agrawal, J. Lopez-Piqueres, R. Vasseur, S. Gopalakrishnan, and A. C. Potter, Observing quantum measurement collapse as a learnability phase transition, arXiv:2311.00058.

[75] H. Mabuchi, Dynamical identification of open quantum systems, Quantum Semiclass. Opt.: J. Eur. Opt. Soc. Part B **8**, 1103 (1996).

[76] J. Gambetta and H. M. Wiseman, State and dynamical parameter estimation for open quantum systems, Phys. Rev. A **64**, 042105 (2001).

[77] M. Tsang, Continuous quantum hypothesis testing, Phys. Rev. Lett. **108**, 170502 (2012).

[78] S. Gammelmark and K. Mølmer, Bayesian parameter inference from continuously monitored quantum systems, Phys. Rev. A **87**, 032115 (2013).

[79] S. Gammelmark and K. Mølmer, Fisher information and the quantum Cramér-Rao sensitivity limit of continuous measurements, Phys. Rev. Lett. **112**, 170401 (2014).

[80] M. C. Caro, Binary classification with classical instances and quantum labels, Quantum Mach. Intell. **3**, 18 (2021).

[81] M. Fanizza, Y. Quek, and M. Rosati, Learning quantum processes without input control, arXiv:2211.05005.

[82] M. McGinley, S. Roy, and S. A. Parameswaran, Absolutely stable spatiotemporal order in noisy quantum systems, Phys. Rev. Lett. **129**, 090404 (2022).

[83] T. Iadecola, S. Ganeshan, J. H. Pixley, and J. H. Wilson, Dynamical entanglement transition in the probabilistic control of chaos, arXiv:2207.12415.

[84] A. J. Friedman, O. Hart, and R. Nandkishore, Measurement-induced phases of matter require feedback, arXiv:2210.07256.

[85] M. Buchhold, T. Müller, and S. Diehl, Revealing measurement-induced phase transitions by pre-selection, arXiv:2208.10506.

[86] N. O'Dea, A. Morningstar, S. Gopalakrishnan, and V. Khemani, Entanglement and absorbing-state transitions in interactive quantum dynamics, arXiv:2211.12526.

[87] P. Sierant and X. Turkeshi, Controlling entanglement at absorbing state phase transitions in random circuits, Phys. Rev. Lett. **130**, 120402 (2023).

[88] C. LeMaire, A. A. Allocca, J. H. Pixley, T. Iadecola, and J. H. Wilson, Separate measurement- and feedback-driven entanglement transitions in the stochastic control of chaos, arXiv:2309.04520.

[89] D. T. Stephen, D.-S. Wang, A. Prakash, T.-C. Wei, and R. Raussendorf, Computational power of symmetry-protected topological phases, Phys. Rev. Lett. **119**, 010504 (2017).

[90] L. Dubins, On extreme points of convex sets, J. Math. Anal. Appl. **5**, 237 (1962).

[91] I. Bengtsson, S. Weis, and K. Życzkowski, in *Geometric Methods in Physics*, edited by P. Kielanowski, S. T. Ali, A. Odzijewicz, M. Schlichenmaier, and T. Voronov (Springer Basel, 2013), p. 175.

[92] M. Tomamichel, M. Berta, and M. Hayashi, Relating different quantum generalizations of the conditional Rényi entropy, J. Math. Phys. **55**, 082206 (2014).

[93] C. W. Helstrom, Quantum detection and estimation theory, J. Stat. Phys. **1**, 231 (1969).

[94] T. Ogawa and M. Hayashi, On error exponents in quantum hypothesis testing, IEEE Trans. Inf. Theory **50**, 1368 (2004).

[95] K. M. R. Audenaert, J. Calsamiglia, R. Muñoz Tapia, E. Bagan, L. Masanes, A. Acin, and F. Verstraete, Discriminating states: The quantum Chernoff bound, Phys. Rev. Lett. **98**, 160501 (2007).

[96] S. Mase, Approximations to the birthday problem with unequal occurrence probabilities and their application to the surname problem in Japan, Ann. Inst. Stat. Math. **44**, 479 (1992).