Learnability Transitions in Monitored Quantum Dynamics via Eavesdropper's Classical Shadows

Matteo Ippoliti[1,2,*] and Vedika Khemani[2]

[1] *Department of Physics, University of Texas at Austin, Austin, Texas 78712, USA*
[2] *Department of Physics, Stanford University, Stanford, California 94305, USA*

(Received 13 November 2023; accepted 20 March 2024; published 5 April 2024)

Monitored quantum dynamics—unitary evolution interspersed with measurements—has recently emerged as a rich domain for phase structure in quantum many-body systems away from equilibrium. Here we study monitored dynamics from the point of view of an eavesdropper who has access to the classical measurement outcomes, but not to the quantum many-body system. We show that a measure of information flow from the quantum system to the classical measurement record—the *informational power*—undergoes a phase transition in correspondence with the measurement-induced phase transition (MIPT). This transition determines the eavesdropper's (in)ability to learn properties of an unknown initial quantum state of the system, given a complete classical description of the monitored dynamics and arbitrary classical computational resources. We make this learnability transition concrete by defining classical shadow protocols that the eavesdropper may apply to this problem, and show that the MIPT manifests as a transition in the sample complexity of various shadow-estimation tasks, which become harder in the low-measurement phase. We focus on three applications of interest: Pauli expectation values (where we find the MIPT appears as a point of optimal learnability for typical Pauli operators), many-body fidelity, and global charge in $U(1)$-symmetric dynamics. Our work unifies different manifestations of the MIPT under the umbrella of *learnability* and gives this notion a general operational meaning via classical shadows.

DOI: 10.1103/PRXQuantum.5.020304

I. INTRODUCTION

Recent advances in our ability to address and read out individual degrees of freedom in many-body quantum systems have motivated interest in new types of dynamics where the role of the observer is central. In these *monitored dynamics* [1–5], the observer's measurements shape the evolution of the system and drive it to sharply different possible ensembles of late-time states. These 'measurement-induced phase transitions' (MIPTs) thus define a new paradigm for phase structure in open systems away from equilibrium. At the same time, these technological developments have raised the salience of *quantum state learning*—the general problem of characterizing properties of unknown, potentially complex quantum states with as few measurements as possible [6,7]. In this work we connect these two threads by formulating MIPTs as learnability transitions within the framework of *classical shadows*, a leading practical approach to state learning.

The canonical formulation of the MIPT is in terms of a phase transition in the entanglement properties of ensembles of quantum trajectories [1–3]. In the standard setup, the system evolves through random circuit dynamics composed of local unitary gates interrupted by local projective measurements with probability p. As the measurement rate p is tuned, the ensemble of late-time trajectories undergoes a phase transition from a *disentangling* phase in which trajectories display area-law entanglement (at high p, corresponding to frequent measurements) to an *entangling* phase in which trajectories display volume-law entanglement (at low p, corresponding to infrequent measurements). (Models of dynamics with additional structure, e.g., fermionic Gaussian systems [8–10] or systems that obey a symmetry [11–13], may exhibit different phenomenology and richer phase diagrams. Here we focus on the more generic scenario of non-Gaussian dynamics without symmetries.) While the $p = 0$ and $p = 1$ limits are transparent, the existence of a robust volume-law phase at any finite measurement rate is, *a priori*, surprising. While local unitary gates can generate entanglement only at the boundary of a subsystem, disentangling measurements act everywhere in the bulk: a (naïvely) imbalanced competition,

*Corresponding author: ippoliti@utexas.edu

which should always favor the area-law phase. A key insight for understanding the stability of the volume-law phase was furnished in Refs. [14,15], which posited that the volume-law phase can be understood as a dynamically generated random code in which the correlations between two subsystems are hidden in highly nonlocal degrees of freedom, inaccessible to local measurements. This naturally leads to two complementary information-theoretic perspectives on the MIPT: (a) the *coding* perspective and (b) the *learning* perspective, discussed below. The latter is the focus of this work.

A. Coding

The coding perspective is primarily understood from the point of view of an experimentalist, Alice, controlling a quantum system, see Fig. 1(a). Over the course of the dynamics, measurements are performed on the system (either by Alice herself or by a particular type of "environment" that broadcasts the measurement outcomes). These measurements disturb the initial state of the system. In order to undo this disturbance as much as possible, Alice can perform "recovery" operations on the combined final state of the quantum system and classical measurement apparatus—concretely, she decodes the measurement record $\mathbf{m}$ (which is a binary string of measurement outcomes indexed by spacetime locations) to decide on a unitary operation $U_{\mathbf{m}}$ to apply to the quantum state. As a function of parameters in the monitored dynamics (typically the space-time density of measurements set by p), Alice's ability to recover her initial state undergoes a phase transition: on the entangling side, she can successfully recover an extensive amount of quantum information; on the disentangling side, only a subextensive amount. (This setup can equivalently be formulated in terms of Alice's ability to reconstruct a message sent to her by Bob over a noisy quantum channel.)

We refer to this point of view as the *coding* perspective on the MIPT [14,15], due to its close analogy with quantum error correction (QEC) [16–18]. The measurement record $\mathbf{m}$ serves as the "syndrome" and the conditional unitary $U_{\mathbf{m}}$ as the correction or recovery operation (more accurately, in this setting the measurements play a dual role—both as the "errors" that disturb the encoded information, and as the syndromes that allow for its in-principle recovery); the MIPT arises as a phase transition in the rate of this code, i.e., the ratio of logical qubits to physical qubits, which goes from finite to vanishing. In other words, in the entangling phase an *extensive amount* of quantum information survives in the combined

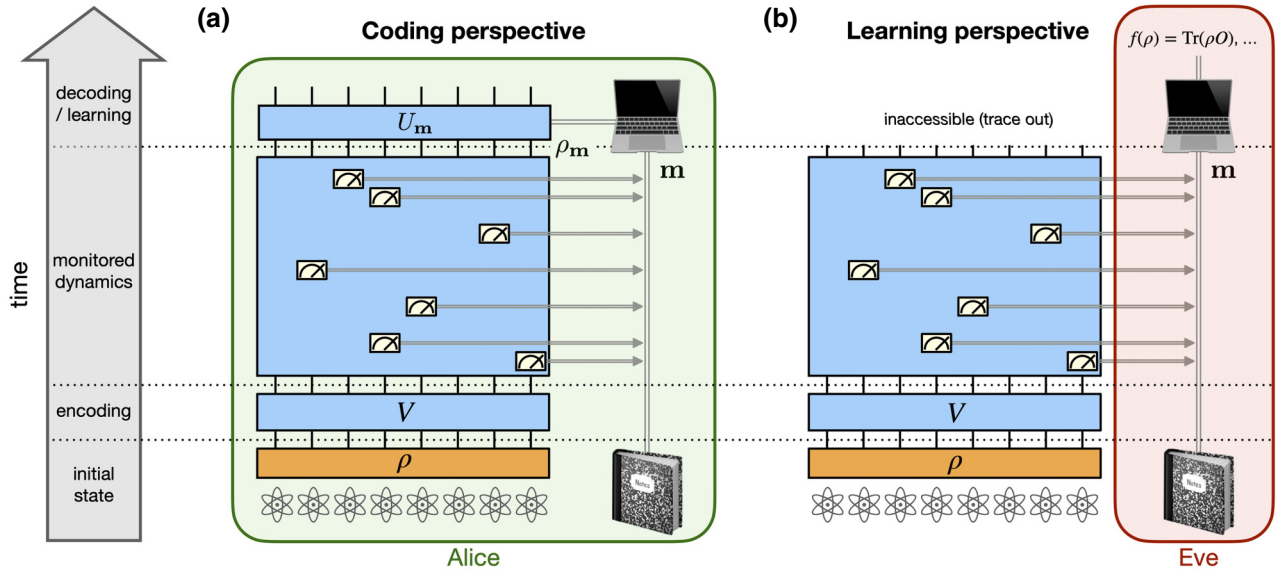

FIG. 1. Different perspectives on measurement-induced phases of quantum information. (a) Coding perspective. A state ρ of a quantum many-body system is subject to monitored dynamics, possibly after an encoding or "prescrambling" step (the global unitary V). The measurement record $\mathbf{m}$ is stored in a classical system, here represented by the lab notebook. The experimentalist, Alice, has access to both the quantum and classical systems (green shaded box). After the dynamics, through classical computation conditioned on the measurement record $\mathbf{m}$, Alice can in principle find a recovery operation $U_{\mathbf{m}}$ and apply it to the quantum system. There is a phase transition in how many qubits from the initial state ρ can be recovered in this way, i.e., in the *coding* properties of monitored dynamics. (b) Learning perspective. An eavesdropper, Eve, has access only to the classical system (red shaded box). She attempts to learn properties of ρ by estimating functions $f(\rho)$ [e.g., expectation values $\mathrm{Tr}(\rho O)$] via classical shadows. There is a phase transition in the sample complexity of these estimation tasks—i.e., in the *learnability* of the state ρ from the measured data $\mathbf{m}$. The two perspectives (a),(b) are dual to each other and the transitions coincide.

quantum-classical state of system and measurement record. While its recovery may be practically hard (in terms of classical computation and quantum circuit complexity), its *in-principle* presence or absence defines sharp phases. This corresponds to a phase transition in the capacity of the channel that maps the initial quantum state to the combined quantum-classical postmeasurement state [14]. This channel capacity also corresponds to the trajectory-averaged entropy of mixed states subject to the monitored evolution, so the coding transition is equally described as a *dynamical purification* transition [15]. Further, the idea of *decoding* implicit in this setup has led to groundbreaking developments in our experimental understanding of the MIPT [19–21].

B. Learning

In this work, we take a complementary perspective of *learning* rather than *coding*, i.e., we focus on the information transmitted to the classical measurement record alone, rather than the combined quantum-classical state. This perspective is centered on an "eavesdropper," Eve, who does not have access to the quantum many-body system, but wants to learn some properties (e.g., observable expectation values) of its unknown initial state. She may try to do so by collecting classical measurement outcomes and performing suitable computation on them, Fig. 1(b).

This perspective has been studied in the literature in two contexts. The first studies the sensitivity of the distribution of measurement outcomes to changes in the initial state. (These are diagnosed either through the Fisher information of the measurement record [22], or through a linear cross-entropy diagnostic that compares the measurement record from a quantum experiment to that of a classical simulation of the same circuit with a different input state [23].) In the second context, the monitored dynamics is enriched with a $U(1)$ symmetry, and an eavesdropper attempts to learn the global charge of a system from measurements of local charge densities [24]. In both cases, the focus is only on the information present in the classical measurement record, as the postmeasurement quantum many-body state is considered inaccessible. This point of view is complementary to the coding perspective. An intuitive expectation is that, if coding is successful, then the "syndrome" measurements $\mathbf{m}$ should reveal no information about the initial state, and thus learning should fail; conversely, if coding fails, information "leaks" into the classical measurement record and learning should become possible.

In this work we sharpen this intuition and make it operationally meaningful. We view the monitored dynamics as a single, complex "randomized measurement" [7] performed on the system. From Eve's point of view, this generalized measurement destroys the quantum state and turns it into classical data. As such, it is impossible for her to recover *quantum* information, as in coding (e.g., any entanglement with an outside reference system is destroyed in the process); but it may still be possible to learn a classical description of the state, as in tomography.

We introduce *classical shadow* protocols [6] that Eve may use to learn various properties of the unknown initial state from the outcomes of many shots of these generalized measurements. We then show that the MIPT manifests as a transition in the sample complexity of these tasks, i.e., how many shots of the experiments Eve needs before she can confidently make predictions. The transition in sample complexity may be from polynomial to exponential, between two exponentials, or between two polynomials, depending on the task at hand. Our framework is very general and furnishes a unified language to describe and study learnability transitions in various different contexts (including previous examples from the literature such as charge learning [24]). Finally, we argue, and prove in some cases, that these learnability transitions reflect a transition in the *informational power* [25] of monitored dynamics—an intrinsic property independent of the chosen learning protocol.

The balance of this paper is structured as follows. In Sec. II we provide a concise, pedagogical review of relevant background topics: generalized measurements, monitored dynamics, and classical shadows. Expert readers may safely skip this section. We then discuss monitored dynamics as a generalized measurement in Sec. III. We review the idea of informational power [25] of generalized measurements and apply it to monitored dynamics, proving under certain assumptions (and conjecturing more generally) that it undergoes a phase transition at the MIPT. In Sec. IV we introduce *"eavesdropper's shadows"*—classical shadows protocols that an eavesdropper may use to learn the state of the system from measurement records. The consequences of the MIPT on these classical shadow protocols are then analyzed in turn: Sec. V on the expectation value of Pauli operators, where we additionally investigate the effect of spatial locality on learnability; Sec. VI on many-body fidelities, which directly connects to recent work on the linear cross-entropy as an order parameter for the transition [23]; and Sec. VII on learning properties of the charge distribution via $U(1)$-symmetric monitored dynamics [12,24], which can also be studied naturally in the formalism of eavesdropper's classical shadows. As we do not impose limits on classical computational resources, we find that the learnability transition coincides with the *charge-sharpening* transition in Ref. [12], as expected [24]. Finally, in Sec. VIII we summarize our results and point out directions for future work.

II. REVIEW

In this section we review essential background topics—generalized measurements (POVMs), monitored

dynamics, and classical shadows—for the sake of a self-contained discussion. This part may be skipped by expert readers.

A. Generalized measurements

Generalized measurements in quantum mechanics are described by positive-operator-valued measures (POVMs), sets of operators $\{E_\alpha\}$ (also known as *effects*) that obey (i) positivity, $E_\alpha \geq 0$ and (ii) normalization $\sum_\alpha E_\alpha = \mathbb{I}$. These predict the probability of observing each outcome α on a given state ρ by the identification $p_\alpha \equiv \mathrm{Tr}(\rho E_\alpha)$, which is guaranteed to be a valid probability distribution by conditions (i–ii). A POVM describes measurement outcomes but not the associated postmeasurement states of the quantum system. That additional data is contained in the *instruments* $\{K_\alpha\}$ (also known as Kraus operators in the context of quantum channels), that obey $E_\alpha = K_\alpha^\dagger K_\alpha$. The state ρ is updated after the measurement as

$$\rho \mapsto \sum_\alpha p_\alpha \rho_\alpha, \quad \rho_\alpha = K_\alpha \rho K_\alpha^\dagger / p_\alpha, \tag{1}$$

where ρ_α is the conditional postmeasurement state of the system given outcome α.

It may further be helpful to view the whole measurement process as a quantum-classical channel

$$\rho \mapsto \sum_\alpha p_\alpha \rho_\alpha \otimes |\alpha\rangle\langle\alpha|_C, \tag{2}$$

where the states $|\alpha\rangle\langle\alpha|_C$ are states of a *classical* register, e.g., Alice's lab notebook. We can always embed such states as orthonormal basis elements [they need to be orthogonal as they are perfectly distinguishable (classical) states] of a sufficiently large Hilbert space.

B. Monitored dynamics

Monitored dynamics is a type of open-system evolution whose quantum trajectories are labeled by a classical "measurement record" $\mathbf{m}$. We take $\mathbf{m} = (m_{t_1,x_1}, \ldots m_{t_M,x_M}) \in \{0,1\}^M$ to be a collection of binary measurement outcomes gathered over different positions and times (t_i, x_i) in the evolution. The quantum many-body state ρ evolves into a quantum-classical state

$$\rho \mapsto \sum_{\mathbf{m}} K_{\mathbf{m}} \rho K_{\mathbf{m}}^\dagger \otimes |\mathbf{m}\rangle\langle\mathbf{m}|_C, \tag{3}$$

which notably is of the same form as Eq. (2): we can in fact view the whole monitored evolution as a single POVM with 2^M possible outcomes $\mathbf{m} \in \{0,1\}^M$ occurring with probabilities $p_{\mathbf{m}} = \mathrm{Tr}(\rho E_{\mathbf{m}})$, $E_{\mathbf{m}} = K_{\mathbf{m}}^\dagger K_{\mathbf{m}}$.

Remarkably, it was discovered that the ensemble of quantum trajectories $\rho_{\mathbf{m}} \equiv K_{\mathbf{m}} \rho K_{\mathbf{m}}^\dagger / p_{\mathbf{m}}$ can undergo a sharp phase transition as a function of model parameters, e.g., the density or rate of measurements in the dynamics, from an "entangling" to a "disentangling" phase [1–5,14,15,19–22,26–32]. The phenomenology of these phases is very rich and beyond the scope of this review section. Here we focus only on the aspect most relevant to this work, which is *dynamical purification* [15], closely related to the coding perspective discussed above: due to the nonunitarity of measurements, an initially mixed state at late enough times generically becomes nearly pure, $\mathbb{E}_{\mathbf{m}}[\mathrm{Tr}(\rho_{\mathbf{m}}^2)] \xrightarrow{t\to\infty} 1$. Here the average is taken over the ensemble of trajectories with Born probability $p_{\mathbf{m}} = \mathrm{Tr}(\rho E_{\mathbf{m}})$. The time scale over which this happens varies sharply depending on which phase we are in: it is $O(\log(N))$ in the disentangling or pure phase, and $O(\exp(N))$ in the entangling or mixed phase, where N is the size of the system (e.g., number of qubits). This reflects the emergence of a quantum code in the entangling phase, which protects some information and prevents it from leaking to the environment for very long times. In particular, taking a dynamic limit with $t, N \to \infty$ with $t = \Theta(N)$ ensures that the average purity becomes an order parameter for the two phases (1 in the disentangling phase, <1 in the entangling phase).

A subtle experimental aspect of this physics is that it is revealed only in nonlinear functions of the trajectories $\{\rho_{\mathbf{m}}\}$, such as the aforementioned average purity $\mathbb{E}_{\mathbf{m}}\mathrm{Tr}(\rho_{\mathbf{m}}^2)$. It does not appear in the average of linear functions, which are fully determined by the average final density matrix $\rho' = \mathbb{E}_{\mathbf{m}}[\rho_{\mathbf{m}}] = \sum_{\mathbf{m}} K_{\mathbf{m}} \rho K_{\mathbf{m}}^\dagger$. This is the output of *dissipative* dynamics and thus generically trivial across the phase diagram. A naïve approach based on measuring properties of individual trajectories incurs an exponential sampling overhead, of order 2^M, due to post-selection of the M measurement outcomes (producing the same trajectory twice takes of order 2^M trials).

More sophisticated approaches that ameliorate or avoid this prohibitive sampling cost have been proposed [19,33–35] and implemented experimentally [20,21]. Reference [19] proposed scalable order parameters that can be measured by *decoding* the measurement record and applying feedback control to the quantum system, similar to the setup in Fig. 1(a). Building on this, other approaches have more recently been proposed that bypass the need for quantum feedback by defining "hybrid" quantum-classical order parameters that correlate quantum state readout with the output of classical simulations [21,23,36–38]. These approaches generally trade the exponential sample complexity of order 2^M for the complexity of classical simulation, which may be polynomial or exponential depending on the models. We will show below that our work furnishes a new set of hybrid order parameters for the MIPT, in the form of variances of shadow estimators. We also note that another approach making use of classical shadows to diagnose the MIPT was proposed recently [38], with the

complementary goal of learning the final, *postmeasurement* state. The goal is thus conceptually distinct from our learnability perspective, despite employing similar tools—a testament to the generality and wide applicability of the classical shadow framework, reviewed next.

C. Classical shadows

Classical shadows are a framework for learning properties of quantum states in a relatively sample-efficient manner by making use of randomized measurements [6,7,39–55]. In this work we will make use of classical shadows in order to formulate a general framework that operationalizes the notion of "learnability" in monitored dynamics.

The standard shadow protocol [6,7] attempts to learn properties of an unknown quantum state ρ by making measurements in a complete orthonormal basis, e.g., the computational basis after a suitable random rotation U. From such a measurement one obtains a snapshot $|\mathbf{b}\rangle\langle\mathbf{b}|$ of the rotated state $U\rho U^\dagger$, and thus a snapshot $U^\dagger|\mathbf{b}\rangle\langle\mathbf{b}|U$ of the original state ρ (here $\mathbf{b}\in\{0,1\}^N$ is an output bitstring). Averaging many such snapshots over outcomes of $\mathbf{b}$ and random choices of U yields a "noisy" version of ρ:

$$\mathcal{M}(\rho)=\sum_{\mathbf{b}}\int dU\,\langle\mathbf{b}|U\rho U^\dagger|\mathbf{b}\rangle U^\dagger|\mathbf{b}\rangle\langle\mathbf{b}|U, \tag{4}$$

where $\langle\mathbf{b}|U\rho U^\dagger|\mathbf{b}\rangle$ is the probability of obtaining the measurement outcomes $\mathbf{b}$. Here dU is shorthand for whichever measure on the unitary group we are using (whether continuous or discrete) and $\mathcal{M}$ is known as the *shadow channel*. Therefore, by "denoising" our snapshots we obtain unbiased estimators of the true density matrix:

$$\hat{\rho}=\mathcal{M}^{-1}[U^\dagger|\mathbf{b}\rangle\langle\mathbf{b}|U]. \tag{5}$$

The classical simulation complexity of the protocol is determined by the complexity of obtaining the unrotated snapshot $U^\dagger|\mathbf{b}\rangle\langle\mathbf{b}|U$, and of determining and applying the inverse channel $\mathcal{M}^{-1}$ on a classical computer.

The snapshots in Eq. (5) can then be used to predict many properties of ρ; for example, for linear expectation values $\langle O\rangle$, the estimators $\hat{o}=\mathrm{Tr}(\hat{\rho}O)$ average to the correct answer. The cost of the learning protocol is quantified by the number of samples that are needed to make an accurate prediction. This is dictated by the *shadow norm* $\|O\|_{\mathrm{sh}}^2$ (which controls the variance of $\hat{o}$), whose scaling depends on the chosen ensemble of measurements. Important results are known for the standard protocols based on local random Pauli (i.e., locally randomized) and random Clifford (i.e., globally randomized) measurements. In the former one has $\|P\|_{\mathrm{sh}}^2=3^k$ where P is a Pauli operator of weight k, making the scheme practical for local (few-body) operators. In the latter, $\|O\|_{\mathrm{sh}}^2=\mathrm{Tr}(O^\dagger O)$ for any operator O, making the method suitable for low-rank operators such as pure-state projectors. This has the important application of computing many-body fidelity with pure states. The shadow protocol has recently been generalized to *shallow shadows* [48–51] in which the randomization step is effected by finite-depth circuits, and which yields an exponential gain in sampling complexity for the task of learning large, spatially contiguous k-body operators.

Classical shadows have been recently extended to the case of generalized (i.e., nonprojective) measurements, represented by a POVM $\{E_\alpha\}$ [44,45]. Given a measured outcome α, there are several possibilities for what to use as a "snapshot," giving rise to a family of shadow channels, which generalize Eq. (4) and take the form

$$\mathcal{M}(\rho)=\sum_\alpha\int dU\,\mathrm{Tr}(\rho E_\alpha^U)\eta_\alpha^U. \tag{6}$$

Here we use the superscript U to denote that the effect E_α and the snapshot η_α both in general depend on random unitary rotations that are part of the protocol, again represented by integration over dU, but we do not assume a specific form. We will suppress this dependence on random unitaries from our notation in the following. In Sec. IV we will see three different choices for η_α that are well motivated by conceptual or practical considerations.

III. INFORMATIONAL POWER OF MONITORED DYNAMICS

Taking the *learning* perspective illustrated in Fig. 1(b), monitored dynamics as a whole is effectively a generalized measurement on the system—a process that maps the quantum state ρ to a probability distribution over outcomes $\mathbf{m}$. In particular, if $K_{\mathbf{m}}$ is the evolution operator corresponding to trajectory $\mathbf{m}$, then the process is described by a POVM $\{E_{\mathbf{m}}\equiv K_{\mathbf{m}}^\dagger K_{\mathbf{m}}\}$, which maps the quantum state ρ to the classical probability distribution $\{\mathrm{Tr}(\rho E_{\mathbf{m}})\}$. The question of learning thus boils down to the "strength" of this generalized measurement, or the information content of its outcomes.

To sharpen this notion, let us consider an ensemble of states $\mathcal{E}=\{(p_i,\rho_i)\}$. This is a discrete collection of states ρ_i, each one occurring with probability p_i, e.g., from some classical stochastic process involved in the state preparation. We want to know how well a POVM $\Pi=\{E_{\mathbf{m}}\}$ can distinguish the different states ρ_i in the ensemble $\mathcal{E}$.

If the state ρ_i was drawn, then the outcome $\mathbf{m}$ occurs with probability $p_{\mathbf{m}|i}\equiv\mathrm{Tr}(\rho_i E_{\mathbf{m}})$. Along with p_i (given as part of the definition of $\mathcal{E}$), this defines the joint distribution

$$p_{i,\mathbf{m}}=p_{\mathbf{m}|i}p_i=\mathrm{Tr}(p_i\rho_i E_{\mathbf{m}}) \tag{7}$$

and thus also the marginal $p_{\mathbf{m}}=\mathrm{Tr}(\rho E_{\mathbf{m}})$, with $\rho=\sum_i p_i\rho_i$ the average state of the ensemble. With this data,

we can define the *mutual information* between the POVM Π and the state ensemble $\mathcal{E}$ as the mutual information between variables i and $\mathbf{m}$ in the joint distribution Eq. (7):

$$I(\mathcal{E}:\Pi) = -\sum_i p_i \ln p_i - \sum_{\mathbf{m}} p_{\mathbf{m}} \ln p_{\mathbf{m}} + \sum_{i,\mathbf{m}} p_{i,\mathbf{m}} \ln p_{i,\mathbf{m}} = \sum_{i,\mathbf{m}} p_{i,\mathbf{m}} \ln \frac{p_{i,\mathbf{m}}}{p_i p_{\mathbf{m}}}, \tag{8}$$

where the expression in the second line is in the form of a Kullback-Leibler divergence between the true distribution $p_{i,\mathbf{m}}$ and the product of its marginals $p_i p_{\mathbf{m}}$, in which the two variables are independent. Informally, this mutual information characterizes how much the measurement outcome $\mathbf{m}$ knows about the underlying state i.

Finally, maximizing the mutual information $I(\mathcal{E}:\Pi)$ Eq. (8) over possible choices of the ensemble $\mathcal{E}$ yields an intrinsic property of the POVM Π known as its *informational power* [25,56,57]:

$$W(\Pi) = \max_{\mathcal{E}} I(\mathcal{E}:\Pi). \tag{9}$$

The optimization involved in the definition of informational power, Eq. (9), may be hard in general. However, it becomes trivial if we include a "prescrambling" or "encoding" step—meaning the system is rotated by a random unitary V before being measured, as sketched in Fig. 1. In that case, we show in Appendix A that

$$W(\Pi) = Q(\mathbb{I}/D) - \sum_{\mathbf{m}} \pi_{\mathbf{m}} Q(\sigma_{\mathbf{m}}), \tag{10}$$

where D is the dimension of the many-body Hilbert space [in this work we take $D = q^N$, i.e., a system of N q-state qudits; we also focus on qubits ($q = 2$) when specified], $Q(\rho)$ is a function known as the *subentropy* (see Appendix A 2), and we have introduced a state ensemble $\mathcal{E}_\Pi = \{(\pi_{\mathbf{m}}, \sigma_{\mathbf{m}})\}$ given by

$$\begin{cases} \pi_{\mathbf{m}} &= \mathrm{Tr}(E_{\mathbf{m}})/D, \\ \sigma_{\mathbf{m}} &= E_{\mathbf{m}}/\mathrm{Tr}(E_{\mathbf{m}}). \end{cases} \tag{11}$$

This state ensemble is dual to our POVM $\Pi = \{E_{\mathbf{m}}\}$ [to every state ensemble $\mathcal{E} = \{(p_i, \rho_i)\}$ one can canonically associate the POVM $\Pi_\mathcal{E} = \{p_i \rho^{-1/2} \rho_i \rho^{-1/2}\}$, with $\rho = \sum_i p_i \rho_i$. This is known as the *"pretty good measurement"* [58]. All POVMs Π obey $\Pi_{\mathcal{E}_\Pi} = \Pi$, and all state ensembles $\mathcal{E}$ with $\rho = \mathbb{I}/D$ obey $\mathcal{E}_{\Pi_\mathcal{E}} = \mathcal{E}$]. It is straightforward to check, from the POVM conditions, that this is in fact a state ensemble, i.e., that $\pi_{\mathbf{m}}$ is a valid probability distribution and that the $\sigma_{\mathbf{m}}$ are states. In fact, both objects have intuitive physical interpretations: $\pi_{\mathbf{m}}$ is the probability of obtaining outcome $\mathbf{m}$ when running monitored dynamics on the fully mixed state $\rho = \mathbb{I}/D$; $\sigma_{\mathbf{m}} \propto K_{\mathbf{m}}^\dagger (\mathbb{I}/D) K_{\mathbf{m}}$ is the output of *Heisenberg-picture* monitored evolution $K_{\mathbf{m}}^\dagger$, also acting on the fully mixed state. In the most common models of monitored dynamics, made only of unitary gates and projective measurements [1,2], the Schrödinger and Heisenber pictures are equivalent at the ensemble level (while individual trajectories are generically not self-adjoint, the ensemble—over random unitary operations, locations, and outcomes of measurements—is invariant in those models: $\{K_{\mathbf{m}}\} = \{K_{\mathbf{m}}^\dagger\}$), so that we can interpret $\mathcal{E}_\Pi$ as a valid *ensemble of trajectories* for a monitored mixed-state evolution. Note that this is *not* the ensemble of physical monitored trajectories of the quantum system, $\{(p_{\mathbf{m}}, K_{\mathbf{m}} \rho K_{\mathbf{m}}^\dagger / p_{\mathbf{m}})\}$, with $p_{\mathbf{m}} = \mathrm{Tr}(E_{\mathbf{m}} \rho)$: such states are inaccessible to Eve. The ensemble $\mathcal{E}_\Pi$ instead emerges as a description of the measurement process, built purely from a classical description of the dynamics (the $K_{\mathbf{m}}$ operators) available to Eve. This mapping of the measurement process to an auxiliary ensemble of monitored trajectories plays a key role in this work.

Having established this formalism, we can now interpret the analytical result for the informational power $W(\Pi)$, Eq. (10). An exact analytical expression for the subentropy $Q(\rho)$ in terms of the spectrum of ρ is known [59], but not particularly illuminating. However, for stabilizer states (in fact for all states proportional to projectors), one can analytically obtain a more explicit form for the subentropy (see Appendix A 3),

$$Q(\rho) = 1 - \gamma - \delta H(q^S). \tag{12}$$

Here q is the local Hilbert-space dimension ($D = q^N$), S is the entropy of ρ (in dits), $\gamma = 0.577\ldots$ is Euler-Mascheroni's constant, and $\delta H(x) = H_x - (\ln(x) + \gamma)$ is the deviation of the harmonic sum $H_x = 1 + \frac{1}{2} + \cdots + 1/x$ from its large-x expansion $\ln(x) + \gamma$. δH is nonnegative and bounded above by a constant; it vanishes as approximately $1/2x$ for large x.

Thus we arrive at the following result for monitored Clifford circuits (where all the trajectories $\sigma_{\mathbf{m}}$ are stabilizer states):

$$W(\Pi) = \mathbb{E}_{\mathbf{m}}[\delta H(q^{S_{\mathbf{m}}})] - \delta H(q^N), \tag{13}$$

with the average over trajectories taken according to the measure $\pi_{\mathbf{m}}$. This explicitly depends on the entropy of monitored trajectories $S_{\mathbf{m}} = -\log_q[\mathrm{Tr}(\sigma_{\mathbf{m}}^2)]$, meaning that the dynamical purification transition manifests as a transition in the informational power. In the entangling phase, the trajectories $\sigma_{\mathbf{m}}$ remain highly mixed, with $S_{\mathbf{m}} \propto N$; for large N, employing the expansion $\delta H(x) \sim 1/(2x)$, we obtain

$$W(\Pi) \simeq \frac{1}{2}\mathbb{E}_{\mathbf{m}}[\mathrm{Tr}(\sigma_{\mathbf{m}}^2)] - \frac{1}{2D} \xrightarrow{D\to\infty} 0, \tag{14}$$

so that the informational power, which quantifies learnability from the measurement record, goes to zero in the entangling phase. In the disentangling phase, the trajectories purify and the entropies $S_{\mathbf{m}}$ quickly decay towards zero [reaching $O(1)$ values at logarithmic depth] and thus $\delta H(q^{S_{\mathbf{m}}})$ remains finite.

To summarize, we have shown that the informational power of prescrambled Clifford monitored dynamics of depth $t = \mathrm{poly}(N)$ for large systems $N \gg 1$ obeys

$$W(\Pi) = \begin{cases} \sim q^{-sN} & \text{(entangling phase)}, \\ \text{const.} > 0 & \text{(disentangling phase)}, \end{cases} \tag{15}$$

with $s \in [0, 1]$ the order parameter of the purification phase transition (entropy density). We conjecture that the same transition holds for generic (non-Clifford) monitored dynamics. A suggestive result to this effect is that a Renyi-2 version of the mutual information $I(\mathcal{E} : \Pi)$ can be computed exactly and depends only on the average purity of the system, thus manifestly displaying the purification transition (see Appendix A 4). While this is not a valid mutual information, in randomized settings it is often a good proxy for the qualitative behavior of the true mutual information.

We note that the informational power is related to the Fisher-information diagnostic in Ref. [22], which measures the susceptibility of the measurement outcome distribution to small changes in the initial state. Like the informational power, the Fisher information aims to quantify how much information about the quantum state flows into the measurement record, thus the two approaches are closely related. A technical difference is that the Fisher information depends on the initial state and the choice of perturbation, while the informational power is an intrinsic property of the monitored dynamics. More importantly, the Fisher-information diagnostic in Ref. [22] generically requires the collection of exponentially many samples (so that probability distributions can be estimated with some accuracy), while in our work we will show that the transition in informational power is reflected in the complexity of classical shadows, and can be determined from few samples—the complexity bottleneck in our scheme lies instead in the classical simulation of the quantum system (needed to carry out classical shadow estimation).

Finally, we note that the informational power is equal to the channel capacity of the quantum-to-classical channel mapping the state ρ to a classical probability distribution over measurement records $\mathbf{m}$ [25]: $\rho \mapsto \sum_{\mathbf{m}} \mathrm{Tr}(E_{\mathbf{m}}\rho)|\mathbf{m}\rangle\langle\mathbf{m}|_C$, where $\{|\mathbf{m}\rangle_C\}$ is an orthonormal basis of classical states of the measurement device, as in Eq. (3). Therefore, $W(\Pi)$ directly quantifies the flow of information from Alice's unknown state to Eve's classical data, sketched in Fig. 1(b). The MIPT arises as a sharp transition in this flow of information. In the rest of this work we examine how this transition affects Eve's ability to learn properties of the quantum state ρ.

IV. EAVESDROPPER'S SHADOWS

The informational power transition discussed above suggests a general characterization of the MIPT as a learnability phase transition. To assign an operational meaning to the transition, one needs to consider concrete protocols that Eve might employ to learn features of Alice's unknown initial state ρ from the eavesdropped classical data $\mathbf{m}$. Classical shadows, reviewed in Sec. II C, have emerged as a general and powerful framework for addressing this type of problem. Here we apply them to the generalized measurement associated with monitored dynamics. We will refer to these protocols as *"eavesdropper's shadows."*

A. Setup

We consider an experimentalist, Alice, who controls a quantum many-body system, and an eavesdropper, Eve, who wants to learn properties of Alice's system without having access to it. Alice prepares an initial state ρ and runs some model of monitored dynamics on it (e.g., a brickwork circuit with single-qubit projective measurements [1,2]). She iterates this process many times, with the same state ρ but a different realization of the dynamics each time. Eve only has access to the following, *purely classical* data, for each run of the experiment:

(i) complete classical description of the monitored dynamics (e.g., circuit architecture, gates, locations and basis of measurements);
(ii) midcircuit measurement record $\mathbf{m}$.

Eve aims to learn as much as she can about the initial state ρ from as few runs of the experiment as possible. This setup in sketched in Fig. 1(b).

Eve's task can be readily cast in the framework of classical shadows with generalized measurements [44,45]. From Eve's point of view, this setup is equivalent to a generalized measurement of ρ, namely the POVM $\Pi = \{E_{\mathbf{m}} = K_{\mathbf{m}}^\dagger K_{\mathbf{m}}\}$, with $K_{\mathbf{m}}$ being the Kraus operator for quantum trajectory $\mathbf{m}$ of the monitored dynamics. While we have left it implicit to lighten our notation, $E_{\mathbf{m}}$ also depends on random unitary gates that vary with each realization; therefore, it is a type of randomized measurement [7]. The problem of learning about ρ from outcomes $\mathbf{m}$ of the randomized measurements is thus formally analogous to the standard classical shadow protocol [6,7], Sec. II C.

B. Protocol

In the standard protocol, based on a projective POVM $\{U^\dagger|\mathbf{b}\rangle\langle\mathbf{b}|U\}$, the choice of a postmeasurement "snapshot" state is automatic—upon getting outcome $\mathbf{b}$, the best guess

for the pre-measurement state is just the POVM element itself, $\sigma_{U,\mathbf{b}} \equiv U^\dagger|\mathbf{b}\rangle\langle\mathbf{b}|U$. In the general case, with a non-projective POVM $\{E_\mathbf{m}\}$, the choice is less obvious. In fact, there are several valid choices motivated on practical or conceptual grounds, as we will see below; all of them recover the "canonical" choice in the limit of the POVM becoming projective. Denoting a choice of snapshot state by $\eta_\mathbf{m}$, we have a measure-and-prepare channel

$$\mathcal{M}(\rho) = \sum_\mathbf{m} \mathrm{Tr}(E_\mathbf{m}\rho)\eta_\mathbf{m}, \tag{16}$$

which can in principle be used for classical shadow estimation, along the same steps outlined in Sec. II C.

A natural choice for $\eta_\mathbf{m}$ is based on a general mapping between POVMs and state ensembles, Eq. (11): we propose setting $\eta_\mathbf{m} = \sigma_\mathbf{m}$. This choice closely mirrors the standard protocol, Sec. II C, upon replacing $U \mapsto K_\mathbf{m}$ (random monitored dynamics instead of random unitary rotation) and $|\mathbf{b}\rangle\langle\mathbf{b}| \mapsto \mathbb{I}/D$. The latter is a uniform average over all possible outcomes $|\mathbf{b}\rangle$, representing Eve's complete ignorance about the final state of the system. Beyond this heuristic reasoning, we note that this choice can be formally motivated as the *Petz recovery map* [60–63] from the space of measurement records to the space of quantum states, see Appendix B.

With this prescription, the shadow channel reads

$$\begin{aligned}\mathcal{M}(\rho) &= \sum_\mathbf{m} \mathrm{Tr}(\rho E_\mathbf{m})\sigma_\mathbf{m} \\ &= D\sum_\mathbf{m} \pi_\mathbf{m}\mathrm{Tr}(\rho\sigma_\mathbf{m})\sigma_\mathbf{m} \\ &= D\mathrm{Tr}_2[(\mathbb{I}\otimes\rho)\sigma^{(2)}],\end{aligned} \tag{17}$$

where Tr_2 is the partial trace over the second replica, and $\sigma^{(2)}$ is the *second moment operator* of the state ensemble $\mathcal{E}_\Pi = \{(\pi_\mathbf{m}, \sigma_\mathbf{m})\}$ dual to our POVM $\Pi = \{E_\mathbf{m}\}$ [see Eq. (11)]:

$$\sigma^{(2)} = \sum_\mathbf{m} \pi_\mathbf{m}\sigma_\mathbf{m}^{\otimes 2}. \tag{18}$$

For most typical models of monitored dynamics (featuring only unitary evolution and projective measurements), the states $\sigma_\mathbf{m}$ are quantum trajectories of a valid monitored evolution $K_\mathbf{m}^\dagger$ acting on the fully mixed state; thus the second-moment operator $\sigma^{(2)}$ is directly sensitive to the MIPT. For instance, the order parameter of *dynamical purification phases* [15,19] (Sec. II B), the trajectory-averaged purity $\mathcal{P}$, can be obtained as an expectation value on the two-replica state $\sigma^{(2)}$:

$$\mathcal{P} = \sum_\mathbf{m} \pi_\mathbf{m}\mathrm{Tr}(\sigma_\mathbf{m}^2) = \mathrm{Tr}(\sigma^{(2)}\hat{\tau}), \tag{19}$$

with $\hat{\tau}$ the replica SWAP operator. Thus $\sigma^{(2)}$ undergoes a sharp change at the MIPT, and by extension so does the shadow channel $\mathcal{M}$, Eq. (17). We will investigate the consequences of this sharp change in terms of *learnability transitions* in the rest of the paper.

C. Alternative prescriptions

Before proceeding, we note that other choices for the "snapshot" $\eta_\mathbf{m}$ are possible and well motivated. In particular, two choices have been considered in the literature in the context of classical shadows with generalized measurements [44,45]. We discuss them in detail in Appendix B, and briefly summarize the results below:

(i) *Least squares* [45]: set $\eta_\mathbf{m} = E_\mathbf{m}$. This is *not* a state (due to trace normalization), and the resulting "shadow channel" $\mathcal{M}(\rho) = \sum_\mathbf{m} \mathrm{Tr}(\rho E_\mathbf{m})E_\mathbf{m}$ is thus not a channel. Classical shadows work regardless. This choice minimizes the two-norm distance between the observed $(\mathrm{Tr}(\rho E_\mathbf{m}))$ and predicted $[\mathrm{Tr}(\hat{\rho}E_\mathbf{m})$, with $\hat{\rho}$ the classical shadow of $\rho]$ measurement outcome distributions. The shadow channel takes a form analogous to Eq. (17) [see Eq. (B7)], but with a modified second-moment operator

$$\tilde{\sigma}^{(2)} = \sum_\mathbf{m} \tilde{\pi}_\mathbf{m}\sigma_\mathbf{m}^{\otimes 2}, \tag{20}$$

where the probabilities are $\tilde{\pi}_\mathbf{m} = \pi_\mathbf{m}^2/\sum_{\mathbf{m}'}\pi_{\mathbf{m}'}^2$. This also features a MIPT, albeit with a different universality class due to a different reweighting of the trajectories [11,22,23,64,65].

(ii) *Maximum fidelity* [44]: set $\eta_\mathbf{m} = |\psi_\mathbf{m}\rangle\langle\psi_\mathbf{m}|$ where $|\psi_\mathbf{m}\rangle$ is the leading eigenvector of $E_\mathbf{m}$ (we neglect degeneracies). This choice maximizes the fidelity $\langle\phi|\mathcal{M}(|\phi\rangle\langle\phi|)|\phi\rangle$ between the input and output of $\mathcal{M}$, on average over Haar-random input states $|\phi\rangle$. Again the shadow channel takes a form analogous to Eq. (17) [see Eq. (B10)], but with $\sigma^{(2)}$ replaced by the state

$$\sigma^{(\infty,1)} = \sum_\mathbf{m} \pi_\mathbf{m}|\psi_\mathbf{m}\rangle\langle\psi_\mathbf{m}| \otimes \sigma_\mathbf{m}. \tag{21}$$

The expectation of the replica SWAP $\hat{\tau}$ on this state yields the average of q^{-S_∞}, meaning this is also sensitive to the MIPT, with the same universality

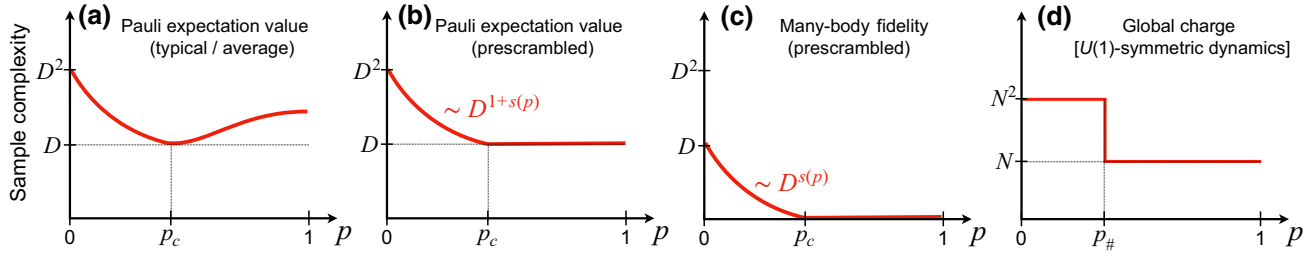

FIG. 2. Schematic summary of main results for the sample complexity of various estimation tasks via "eavesdropper's shadows," as a function of the measurement rate p. N is the system size, $D = q^N$ is the Hilbert-space dimension, and $s(p)$ is the entropy density (order parameter of the entangling phase). (a) Estimation of Pauli expectation values, Sec. V B. Different operators have different complexity; the curve shows the qualitative behavior of average or typical elements of the N-qudit Pauli group. The MIPT ($p = p_c$) emerges as a point of minimum complexity. (b) Estimation of Pauli expectation values with a prescrambling step, Sec. V C. All nonidentity Paulis have the same complexity, of order $D^{1+s(p)}$, which changes nonanalytically at the MIPT. (c) Estimation of many-body fidelity (with a prescrambling step), Sec. VI. The complexity transitions from constant to exponential in N at the MIPT. (d) Estimation of global charge from dynamics with $U(1)$ symmetry, Sec. VII. The complexity transitions from order N to order N^2/t at the charge-sharpening transition, $p = p_\#$.

and critical point as Eq. (17) in this case (all Renyi entropies with indices $n > 1$ are bounded within a multiplicative constant of each other, so S_2 and S_∞ have the same critical properties).

We note that all three prescriptions reduce to the one from Ref. [6] for the case of projective measurements, but they differ for generalized measurements. In the following we will use the prescription $\eta_{\mathbf{m}} = \sigma_{\mathbf{m}}$ unless otherwise specified; the qualitative conclusions would be unchanged with either prescription, since as we saw each of them is sensitive to the MIPT.

To summarize, we have examined different strategies that Eve might use to learn properties of Alice's unknown initial state ρ from the measurement record $\mathbf{m}$ via classical shadows. We have found that the shadow channel is determined by a *moment operator* for an emergent ensemble of monitored quantum trajectories of a mixed-state dynamics, Eq. (11). As the second moment of trajectories is sensitive to the MIPT, we expect this to lead to a qualitative change in the performance of classical shadows. In the following sections we work out the consequences of this observation on a range of different manifestations of the measurement-induced phase transition. Our results are schematically summarized in Fig. 2.

V. LEARNING PAULI EXPECTATION VALUES

In this section we begin to unravel the consequences of the connection between "eavesdropper's shadows" and dynamical purification of mixed states introduced in Sec. IV. We start by focusing on Eve's ability to learn Pauli expectation values on the unknown initial state ρ, namely to estimate $\langle P\rangle = \mathrm{Tr}(\rho P)$ (P is a Pauli operator; we use generalized Pauli operators generated by the "clock" and "shift" operators) to constant additive error ϵ.

A. Shadow norm and entanglement

In general, the sample complexity of learning the expectation of an operator O on a state ρ is given by its squared *shadow norm* $\|O\|^2_{\mathrm{sh}}$, see Sec. II C. This is given formally by

$$\|O\|^2_{\mathrm{sh}} = D\mathrm{Tr}(\rho \otimes \mathcal{M}^{-1}(O) \otimes \mathcal{M}^{-1}(O)\sigma^{(3)}), \tag{22}$$

where $\sigma^{(3)} = \sum_{\mathbf{m}} \pi_{\mathbf{m}} \sigma_{\mathbf{m}}^{\otimes 3}$ is the third moment operator of the trajectory ensemble $\mathcal{E} = \{(\pi_{\mathbf{m}}, \sigma_{\mathbf{m}})\}$ (note that there is, implicit in this notation, an average over random unitary gates $\{u\}$, which enter the monitored dynamics. An average $\int d\{u\}$ is implicitly present concurrently with each trajectory average $\sum_{\mathbf{m}} \pi_{\mathbf{m}}$; we drop it to lighten the notation). If the ensemble is Pauli invariant [47,66] and O is a Pauli operator, Eq. (22) simplifies to

$$\|O\|^2_{\mathrm{sh}} = \frac{1}{D}\mathrm{Tr}(O\mathcal{M}^{-1}(O)), \tag{23}$$

which is independent of ρ. Furthermore, if the POVM is invariant under multiplication by single-site Pauli operators [66] (as is the case in typical models of monitored dynamics, made with random Haar or Clifford gates), the Pauli operators are eigenmodes of the channel: $\mathcal{M}(P) = \lambda_P P$. Thus, by Eq. (23), their shadow norm is $\|P\|^2_{\mathrm{sh}} = \lambda_P^{-1}$.

We now derive a relationship between the eigenvalues λ_P (controlling the shadow norms of Pauli operators) and the entanglement structure of the ensemble of trajectories $\mathcal{E}$. We have, using operator-to-state notation [$(O|$ for super-bras, $|O)$ for super-kets, with inner product $(A|B) = \mathrm{Tr}(A^\dagger B)$],

$$\lambda_P = \frac{(P|\mathcal{M}|P)}{(P|P)} = \mathrm{Tr}[P^{\otimes 2}\sigma^{(2)}]. \tag{24}$$

At the same time, the averaged purity of a subsystem A in the ensemble of monitored trajectories $\mathcal{E}$ is

$$\mathcal{P}_A = \mathrm{Tr}[\hat{\tau}_A \sigma^{(2)}] \tag{25}$$

with $\hat{\tau}_A$ the replica SWAP operator acting only on subsystem A. We can expand $\hat{\tau}$ in the Pauli basis as

$$\hat{\tau}_A = \frac{1}{D_A} \sum_{P:\,\mathrm{supp}(P)\subseteq A} P^{\otimes 2}, \tag{26}$$

where D_A is the Hilbert-space dimension for subsystem A and $\mathrm{supp}(P)$ denotes the support of P, i.e., the subsystem where P is nonidentity (expanding $\hat{\tau}$ in the two-replica Pauli basis as $\hat{\tau} = \sum_{\alpha\beta} c_{\alpha\beta} P_\alpha \otimes P_\beta$, the coefficients $c_{\alpha\beta}$ are given by $c_{\alpha\beta} = \mathrm{Tr}[\hat{\tau}(P_\alpha \otimes P_\beta)]/D^2 = \mathrm{Tr}(P_\alpha P_\beta)/D^2 = \delta_{\alpha\beta}/D$). A relationship between the entanglement feature $\{\mathcal{P}_A\}$ and the eigenvalues $\{\lambda_A\}$ of channel $\mathcal{M}$ follows:

$$D_A \mathcal{P}_A = \sum_{P:\,\mathrm{supp}(P)\subseteq A} \lambda_P = \sum_{B\subseteq A} (q^2-1)^{|B|} \lambda_B, \tag{27}$$

where the first sum is over Pauli operators P supported inside A, while the second is over subsystems B contained inside A [we use λ_A (A being a subsystem) and λ_P (P being a Pauli operator) interchangeably, with the understanding that $\lambda_P = \lambda_{\mathrm{supp}(P)}$]. The second equality holds because there are $(q^2-1)^{|B|}$ distinct Pauli operators with support B. The inverse of Eq. (27) yields

$$\lambda_A = (1-q^2)^{-|A|} \sum_{B\subseteq A} \mathcal{P}_B (-q)^{|B|}, \tag{28}$$

which is a well-known relationship between entanglement and shadow norm in the theory of classical shadows [47,48,51,52,66]; we see that it straightforwardly extends to our setting of shadows with generalized measurements, when taking $\{\mathcal{P}_A\}$ to be the entanglement feature of monitored trajectories $\sigma_{\mathbf{m}}$.

Equations (27) and (28) connect entanglement properties of the trajectories with shadow norms of Pauli operators. This is interesting as it suggests a sharp change in the performance of classical shadows at the dynamical purification transition, consistent with our prior analysis in Secs. III and IV. The connection between entanglement and shadow norms in Eq. (28) is not straightforward, as it is a sum of exponentially many terms with alternating signs. Nonetheless, a simple exact statement can be made about the *harmonic mean* of $\|P\|_{\mathrm{sh}}^2$ for all Pauli operators supported inside a subsystem A: rewriting Eq. (27), we have

$$\left(\frac{1}{D_A^2} \sum_{P:\,\mathrm{supp}(P)\subseteq A} \|P\|_{\mathrm{sh}}^{-2} \right)^{-1} = D_A/\mathcal{P}_A \sim D_A^{1+s}, \tag{29}$$

where $s \in [0,1]$ is an entropy density defined by the scaling of average purity $\mathcal{P}_A \sim q^{-sN_A} = D_A^{-s}$. This implies a sharp change of the shadow norm distribution at the purification transition, which separates the pure phase ($s = 0$) from the mixed phase ($s > 0$). The harmonic mean of squared shadow norms scales exponentially in subsystem size N_A on both sides of the transition, as $D_A^{1+s} = q^{[1+s(p)]N_A}$, but the coefficient in the exponential changes nonanalytically at the critical point $p = p_c$ [where $s(p) \simeq \Theta(p_c - p)|p - p_c|^\nu$, with Θ the Heaviside theta and $\nu \simeq 1.3$ a critical exponent [15]]. Notably, the harmonic mean Eq. (29) is a lower bound to both the *average* shadow norm (arithmetic mean) and the *typical* shadow norm (geometric mean), implying that both must diverge as $\Omega(2^{(1+s)N_A})$ in the entangling phase and as $\Omega(2^{N_A})$ in the disentangling phase.

B. Optimal learning at the MIPT

To gain more insight on the structure of the shadow norm distribution, beyond the exact harmonic-mean result of Eq. (29) and the bounds it implies, we turn to numerical simulations. We perform exact numerical simulations of mixed-state monitored dynamics (a standard model made of brickwork layers of Haar-random gates and single-qubit Z measurements with probability $p \in [0,1]$) on up to $N = 12$ qubits. Note we are limited to this size by the fact that we simulate the full density matrix dynamics (equivalent to a pure state of $2N$ qubits). We obtain the entanglement feature $\{\mathcal{P}_A\}$ averaged over many realizations of the dynamics; from the entanglement feature we obtain the full set of shadow norms $\{\lambda_A^{-1}\}$ via Eq. (28). Results are shown in Fig. 3. The harmonic mean of shadow norms, as predicted, is a function only of the averaged purity, and is thus monotonically decreasing in the measurement rate p. However the arithmetic mean $D^{-2}\sum_P \|P\|_{\mathrm{sh}}^2$ and geometric mean $\exp(D^{-2}\sum_P \log \|P\|_{\mathrm{sh}}^2)$ alike exhibit nonmonotonic behavior, with a minimum near the purification transition $p \approx 0.16$.

This is explained by the fact that, deep in the pure phase, one recovers *random Pauli shadows* (i.e., shadows with random *local, single-qubit* Pauli measurements [6]). This is exactly true at $p = 1$, and we expect it to be a fairly accurate approximation throughout $p \geq 0.5$, where the circuit is nonpercolating. (For $1/2 < p < 1$, a more precise analogy is with classical shadows with locally entangled measurements [52], as the measurement basis typically breaks up into a tensor product of finite-sized bases, one for each nonpercolating cluster of the circuit.) In this regime, even with complete access to the system, *spatial locality* makes

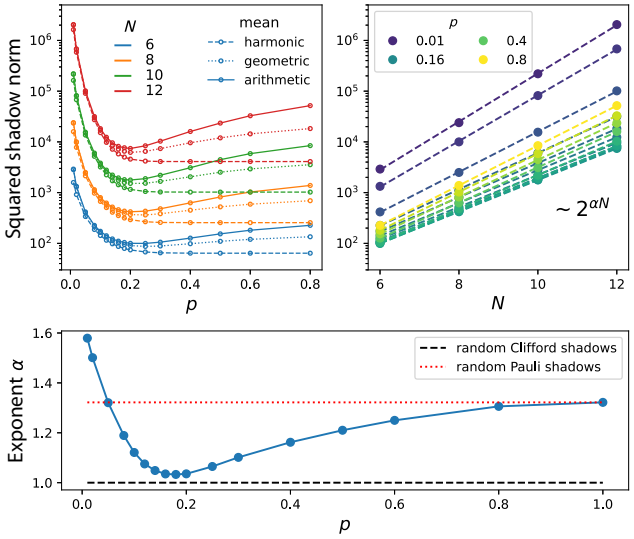

FIG. 3. (a) Squared shadow norms of Pauli operators under eavesdropper's shadows as a function of measurement rate p, for different system sizes N. Data is from exact numerical simulations of quantum trajectory density matrices, averaged over between 2000 and 40,000 random circuit realizations depending on size. The harmonic mean (dotted lines) is monotonically decreasing in p, being completely determined by the system's purity; the arithmetic mean (solid) and geometric mean (dashed) show a minimum at the MIPT. (b) Same data (for the arithmetic mean) plotted as a function of N displays clear exponential scaling proportional to $D^\alpha = 2^{\alpha N}$. Dashed lines are fits. (c) Fit coefficient α versus p shows a minimum at the MIPT.

learning large Pauli operators very inefficient. On the other hand, in the mixed phase, as seen earlier, the informational power of Eve's measurements goes to zero, making any learning inefficient. The MIPT is a "sweet spot" between these two obstructions: the informational power is still finite (at polynomial depth), while the restriction of locality is alleviated.

Approaching the transition from the area-law side, we expect eavesdropper's shadows to perform similarly to *shallow shadows* [48–51] with finite depth: in other words, Eve manages to eventually read out all the information, but this takes an amount of time that grows as the transition is approached, giving information more time to spread, analogous to increasing circuit depth in shallow shadows. At the transition, this "effective depth" should diverge; it is tempting to speculate a relationship with shallow shadows at $\log(N)$ depth, which were shown to give Pauli shadow norms of order q^N for large Pauli operators (consistent with the scaling of both harmonic and arithmetic-mean $\|P\|_{\rm sh}^2$ at the MIPT observed here).

Thus there are two conceptually independent reasons why Eve's task of learning the expectation values of general Pauli operators via classical shadows may be hard: spatial locality at large p (pure, disentangling, or non-coding phase) and a fundamental lack of information at small p (mixed, entangling, or coding phase). Numerical results indicate that an optimum is reached near the purification transition, $p_c \approx 0.16$, where the performance is close to random Clifford shadows ($\|P\|_{\rm sh}^2 \sim D$ for all traceless Paulis). This is a novel characterization of the MIPT, as the optimum rate p for learning Pauli expectation values from the measurement record.

C. Complexity transition

For the rest of this work, we will focus on the presence or absence of *any* information in the measurement record, regardless of its locality structure. For this reason, it is advantageous to "prescramble" the input state ρ with a random global Clifford operation V, as in Sec. III, which eliminates issues with spatial locality. This dramatically simplifies the picture, as the channel $\mathcal{M}$ now depends on a single property of the trajectory ensemble—its average global purity $\mathcal{P}$. We have

$$\mathcal{M}(\rho) = D\int dV \,\mathrm{Tr}_2[(\mathbb{I}\otimes\rho)\, V^{\otimes 2}\sigma^{(2)}(V^\dagger)^{\otimes 2}] = \frac{(D-\mathcal{P})\mathrm{Tr}(\rho)\mathbb{I} + (D\mathcal{P}-1)\rho}{D^2-1}. \tag{30}$$

This reproduces the familiar random-Clifford result $(\mathbb{I}+\rho)/(D+1)$ for $\mathcal{P}=1$. For $\mathcal{P}=1/D$ (no measurements), it correctly gives $\mathbb{I}/D$: there is no information at all about ρ, the measurement channel is a global erasure and is not invertible. Intermediate values of $\mathcal{P}$ interpolate between these two extremes. In particular, the channel is invertible whenever $\mathcal{P} > 1/D$.

It follows that all traceless operators are eigenmodes of $\mathcal{M}$ with eigenvalue

$$\lambda = \frac{D\mathcal{P}-1}{D^2-1} = \frac{D^{1-s}-1}{D^2-1} \sim D^{-(1+s)}, \tag{31}$$

where $\sim$ denotes asymptotic scaling at large D. Thus in particular, all Pauli operators $P \neq I$ have shadow norm

$$\|P\|_{\rm sh}^2 \sim D^{1+s} = q^{(1+s)N}, \tag{32}$$

which indeed changes sharply at the purification transition, as the entropy density goes from $s=0$ (pure phase) to $s>0$ (mixed phase).

D. Information extracted per measurement

The result in Eq. (32) implies that, to learn the expectation of P on the unknown state ρ, Eve needs many more samples in the mixed phase than she does in the pure phase—a factor of approximately q^{sN} more. In other words, the amount of information about ρ leaking into the measurement record $\mathbf{m}$ is suppressed *exponentially* in the mixed phase. This characterization of the mixed or coding

phase [14,15,19] is complementary to the exponentially long lifetime (also scaling as q^{sN}) of information in the system.

In the present framework, however, we can make an even more precise statement about the way in which information on ρ leaks into the measurement record. A key result in dynamical purification [15,29] is that, as a function of circuit depth t, the entropy density s in the mixed phase decreases as

$$s(t) \sim s_0 - \frac{1}{N}\log_q(t) \tag{33}$$

(at times $1 \ll t \ll q^{s_0 N}$), where s_0 is the entropy density "plateau" that characterizes the mixed phase. The ansatz in Eq. (33), plugged into Eq. (32), gives squared shadow norm

$$\|P\|_{\rm sh}^2 \sim \frac{1}{t} D^{1+s_0} \tag{34}$$

for any traceless Pauli P. This quantifies the total number of circuit runs needed to learn $\langle P\rangle$ up to constant error. We can translate this number of circuits into a total number of measurements, $M_{\rm tot}$: with $M \sim pNt$ measurements per circuit, we have

$$M_{\rm tot} \sim M \times \|P\|_{\rm sh}^2 \sim N q^{(1+s_0)N}, \tag{35}$$

which is t independent. [Note this conclusion depends crucially on the coefficient of $\log(t)$ being exactly 1, which was argued, e.g., in Refs. [15,29].]

The fact that $M_{\rm tot}$, the total number of measurements needed, is approximately t independent gives it an invariant meaning: there is, effectively, a fixed amount of information learned by Eve *per measurement*; this amount is approximately $q^{-(1+s_0)N}$ bits. The factor of q^{-N} comes from Haar-random encoding of the initial state (prescrambling), while the factor of $q^{-s_0 N}$ is the additional encoding coming from the dynamics in the mixed phase. This gives a sharp, operational meaning to the idea that measurements fail to read out information about the encoded quantum state in the mixed phase. This sample complexity further saturates the scaling of the informational power in the entangling phase, Sec. III, showing that eavesdropper's shadows are (near) optimal for this task.

VI. LEARNING MANY-BODY FIDELITY

Another observable of interest in many applications, e.g., benchmarking, is the fidelity with a pure many-body state $F = \langle\psi|\rho|\psi\rangle$. The complexity of learning F via classical shadows is given by the shadow norm of the rank-1 projector $|\psi\rangle\langle\psi|$. Notably, this shadow norm is $O(1)$ in random Clifford shadows [39], which makes these observables interesting as practical targets. Reference [23] recently proposed a diagnostic for measurement-induced phases based on a linear cross-entropy function, which intuitively captures the (in)distinguishability between measurement records drawn from monitored dynamics acting on different initial states; here we show that this diagnostic can in fact be readily framed in terms of fidelity estimation via eavesdropper's shadows.

A. Review of linear XEB diagnostic

The linear cross-entropy diagnostic proposed in Ref. [23] reads

$$\mathrm{XEB} = \left\langle \frac{p(\mathbf{m}|\rho_0)}{\sum_{\mathbf{m}'} p(\mathbf{m}'|\rho_0)^2} \right\rangle_{\mathbf{m}\sim p(\mathbf{m}|\rho)}, \tag{36}$$

where $p(\mathbf{m}|\rho)$ is the probability of drawing measurement record $\mathbf{m}$ in an experiment on the initial state ρ. The idea is that ρ_0 is a "simple" (e.g., stabilizer) initial state whose probabilities $p(\mathbf{m}|\rho_0)$ are computed classically, whereas ρ is a generic initial state, and $\mathbf{m} \sim p(\mathbf{m}|\rho)$ denotes samples drawn from an experiment on quantum hardware initialized in state ρ. We again assume the monitored circuit is prefaced by a global random Clifford operation, as in Sec. V C and in Ref. [23].

In the notation of our work, we have $p(\mathbf{m}|\rho) = \mathrm{Tr}(E_{\mathbf{m}}\rho)$. Thus in the limit of a large number of experimental samples, the linear-XEB diagnostic reads

$$\begin{aligned}\mathrm{XEB} &= \frac{\sum_{\mathbf{m}} p(\mathbf{m}|\rho)p(\mathbf{m}|\rho_0)}{\sum_{\mathbf{m}} p(\mathbf{m}|\rho_0)^2} = \frac{\sum_{\mathbf{m}} \pi_{\mathbf{m}}^2 \mathrm{Tr}(\sigma_{\mathbf{m}}^{\otimes 2}\rho\otimes\rho_0)}{\sum_{\mathbf{m}} \pi_{\mathbf{m}}^2 \mathrm{Tr}(\sigma_{\mathbf{m}}^{\otimes 2}\rho_0^{\otimes 2})} \\ &= \mathrm{Tr}[(\rho\otimes\rho_0)\tilde{\sigma}^{(2)}]/\mathrm{Tr}[\rho_0^{\otimes 2}\tilde{\sigma}^{(2)}],\end{aligned} \tag{37}$$

where $\tilde{\sigma}^{(2)}$ is the modified second-moment operator from Eq. (20), defined according to probabilities $\tilde{\pi}_{\mathbf{m}} \propto \pi_{\mathbf{m}}^2$, that also appears in the least-squares shadow prescription. As we discussed in Sec. IV (see also Appendix B), the ensemble of trajectories $\tilde{\mathcal{E}} = \{(\tilde{\pi}_{\mathbf{m}}, \sigma_{\mathbf{m}})\}$ is known to also display an entanglement phase transition, albeit with different universality [22]. Thus the XEB quantity in Eq. (37) is sensitive to a MIPT.

Reference [23] in particular showed that, when including a prescrambling stage (as in Secs. III and V C), in the mixed phase we have XEB $\simeq 1$ independent of ρ, while in the pure phase XEB becomes sensitive to ρ and in particular approaches a finite value < 1 if ρ differs significantly from ρ_0. Aside from the specific details of the protocol, this is suggestive of a phase transition in *learnability of the initial state*: in the pure phase we can successfully tell if ρ and ρ_0 are different, in the mixed phase we fail to do so. In the following, we clarify this connection to a learnability phase transition and frame this result in the language of shadow estimation.

B. Fidelity from a modified linear XEB

To sharpen the connection between the XEB diagnostic of Eq. (36) [23] and learnability, let us start by introducing a slight variation of the quantity, based on how much a new measurement record $\mathbf{m}$ from the experiment updates Eve's belief about the unknown initial state ρ of the system.

We define the quantity

$$\mathrm{XEB}' = \langle p(\rho_0|\mathbf{m})\rangle_{\mathbf{m}\sim p(\mathbf{m}|\rho)}, \tag{38}$$

which differs from Eq. (36) by the order of conditioning [$p(\rho_0|\mathbf{m})$ instead of $p(\mathbf{m}|\rho_0)$ in the numerator of Eq. (36)] and by the absence of a normalization factor, which becomes unnecessary in this case. This quantity has an intuitive interpretation as Eve's updated belief about the initial state, given her new information about a measurement record $\mathbf{m}$ eavesdropped from Alice.

To suitably define the conditional probability $p(\rho_0|\mathbf{m})$, we start with Eve's joint probability distribution over initial (pure) states of the system ρ and measurement records $\mathbf{m}$:

$$p(\rho, \mathbf{m})\, d\rho = \mathrm{Tr}(\rho E_{\mathbf{m}})\, d\rho, \tag{39}$$

where $d\rho$ is a measure over quantum states that reflects Eve's prior beliefs about Alice's initial state ρ (e.g., a uniform measure representing complete ignorance); we require $\int d\rho\, \rho = \mathbb{I}/D$. From this joint distribution we can obtain the marginal

$$p(\mathbf{m}) = \int d\rho\, p(\rho, \mathbf{m}) = \mathrm{Tr}(E_{\mathbf{m}}/D) = \pi_{\mathbf{m}}, \tag{40}$$

and thus the conditional probability via Bayes' rule:

$$p(\rho|\mathbf{m})\, d\rho = \frac{p(\rho, \mathbf{m})\, d\rho}{p(\mathbf{m})} = D\mathrm{Tr}(\rho\sigma_{\mathbf{m}})\, d\rho. \tag{41}$$

In the limit of a large number of experimental shots, we thus obtain

$$\begin{aligned}\mathrm{XEB}' &= \sum_{\mathbf{m}} p(\mathbf{m}|\rho) p(\rho_0|\mathbf{m}) = D\sum_{\mathbf{m}} \mathrm{Tr}(\rho E_{\mathbf{m}})\mathrm{Tr}(\rho_0\sigma_{\mathbf{m}}) \\ &= D\mathrm{Tr}(\rho_0 \mathcal{M}(\rho)),\end{aligned} \tag{42}$$

which is explicitly a function of the shadow channel $\mathcal{M}$ in Eq. (17), and thus sensitive to the standard MIPT ("standard" meaning with correct trajectory weights $\pi_{\mathbf{m}}$).

Finally, with prescrambling, we use Eq. (30) to obtain

$$\mathrm{XEB}' = \frac{D - \mathcal{P} + (D\mathcal{P} - 1)F}{D - 1/D} \simeq 1 + \mathcal{P}F, \tag{43}$$

where $F = \mathrm{Tr}(\rho\rho_0)$ is the fidelity between the true quantum state ρ and the classical guess ρ_0 (taken to be pure). The $\simeq$ denotes the asymptotic scaling at large D, to leading order. Thus in the mixed phase, where $\mathcal{P} \sim D^{-s} \ll 1$, XEB$'$ is exponentially close to 1 regardless of the fidelity F between the unknown state ρ and the guess ρ_0, whereas in the pure phase it approaches a constant value that is informative about F:

$$\overline{\mathrm{XEB}'} \xrightarrow{D\to\infty} \begin{cases} 1 & \text{(mixed phase)} \\ 1 + \text{const.} \times F & \text{(pure phase)} \end{cases} \tag{44}$$

To practically assess the learnability of F, we must also consider the fluctuations of $p(\rho_0|\mathbf{m})$ across experimental shots $\mathbf{m}$. We address this question in Appendix C, where we show that

$$\delta\mathrm{XEB}' \sim D^{-s/2} \tag{45}$$

to leading order in large D. Learning the value of the fidelity F to additive error ϵ requires learning the value of $\mathrm{XEB}' \simeq 1 + \mathcal{P}F$ to additive error $\mathcal{P}\epsilon \sim D^{-s}\epsilon$; this requires a number of repetitions

$$M \sim \frac{(\delta\mathrm{XEB}')^2}{(\mathcal{P}\epsilon)^2} \sim D^s\epsilon^{-2}, \tag{46}$$

which undergoes a sharp change from constant to exponential at the MIPT.

C. Shadow estimation

The same phase transition in sample complexity as Eq. (46) can be straightforwardly obtained from shadow estimation of the fidelity $F = \langle\psi|\rho|\psi\rangle$ with a many-body state $|\psi\rangle$. The sample complexity of learning F is quantified by the squared shadow norm of the operator $O = |\psi\rangle\langle\psi|$:

$$\||\psi\rangle\langle\psi|\|^2_{\mathrm{sh}} = D\mathrm{Tr}(\rho \otimes \mathcal{M}^{-1}(|\psi\rangle\langle\psi|)^{\otimes 2}\sigma^{(3)}), \tag{47}$$

with $\sigma^{(3)}$ the third moment of the ensemble of trajectories:

$$\sigma^{(3)} = \sum_{\mathbf{m}} \pi_{\mathbf{m}}\sigma_{\mathbf{m}}^{\otimes 3}. \tag{48}$$

With prescrambling, Eq. (47) can be computed analytically; see Appendix C for details. In all, we obtain to leading order in large D

$$\begin{aligned}\||\psi\rangle\langle\psi|\|^2_{\mathrm{sh}} &= \mathcal{P}^{-1} + F\mathcal{P}^{-2}\mathcal{P}^{(3)} + O(1/D) \\ &\sim D^s,\end{aligned} \tag{49}$$

where $\mathcal{P}^{(3)} = \mathbb{E}_{\mathbf{m}}\mathrm{Tr}(\sigma_{\mathbf{m}}^3)$ relates to the third Renyi entropy of the ensemble of trajectories, and the second line uses the fact that $\mathcal{P} \geq \mathcal{P}^{(3)}$. This gives the same sample complexity (scaling as D^s) as Eq. (46).

To summarize, we have shown that the cross-entropy diagnostic of Ref. [23] (Sec. VI A) can be reinterpreted,

with a small modification, as a protocol to learn the fidelity of unknown initial state ρ with another state $\rho_0 = |\psi\rangle\langle\psi|$ (Sec. VI B); this task is easy (requires a constant number of experimental samples) in the disentangling phase, and hard (requires an exponential number of samples approximately q^{sN}) in the entangling phase. In turn, the task can be straightforwardly phrased in terms of estimation of the fidelity under eavesdropper's shadows (Sec. VI C); its complexity, given by the squared shadow norm of the projector $\rho_0 = |\psi\rangle\langle\psi|$, exhibits a transition from constant to exponential at the MIPT.

VII. LEARNING CHARGE IN $U(1)$-SYMMETRIC DYNAMICS

A setting where learnability transitions in monitored dynamics are already well established is that of systems with a $U(1)$ symmetry [12,24,67]. There, one may try to learn the global charge Q of the system from measurements of the local charge density q_x. The success of this learning task depends on the decoder (i.e., classical prediction algorithm) used; however, granting Eve arbitrary classical computational resources, the learnability transition [24] was found to coincide with the *charge-sharpening* transition, discussed below [12,67]. Here we recover the same result in the framework of eavesdropper's shadows, where it takes the form of a transition in the shadow norm of the global charge operator $\|\hat{Q}\|_{\rm sh}$.

A. Setup

We consider a system of N qubits with a $U(1)$ symmetry generated by the charge operator

$$\hat{Q} = \frac{1}{2}\sum_i Z_i = \sum_{Q=-N/2}^{N/2} Q\hat{\Pi}_Q, \tag{50}$$

where $\hat{\Pi}_Q$ are orthogonal projectors on the charge sectors, of rank $\binom{N}{Q+N/2}$. The system evolves under a combination of $U(1)$-symmetric unitary gates and measurements of the local charge density Z_i.

Reference [12] identified a *charge-sharpening* transition in this class of models. Charge sharpening is the loss of charge fluctuations over the course of symmetric monitored dynamics; it is analogous to the dynamical purification transition, but restricted to a state's *number entropy*. (Writing a symmetric state ρ as a direct sum of states in each charge block $\rho = \bigoplus_Q p_Q \rho_Q$, the number entropy is $S_n = -\sum_Q p_Q \ln p_Q$.) Purification of the number entropy corresponds to projecting a state into a single charge sector. The time scale for this process (sharpening time, $t_\#$) undergoes a transition from order $\log(N)$ to order N. The transition was located at a critical measurement rate $p = p_\# < p_c$, where p_c is the entanglement (or purification) critical point. In other words, one has a transition between a *fuzzy phase* ($p < p_\#$) and a *sharp phase* ($p > p_\#$) within the entangling phase, while the disentangling phase ($p > p_c$) is always sharp.

The charge-sharpening transition $p = p_\#$ was found to correspond to a transition in learnability of the total charge in the initial state ρ, at least in the limit where Eve has access to complete information about the circuit and unlimited classical computational resources (which is the setting we consider). This result can be recovered straightforwardly in our "eavesdropper's shadows." The formalism of Sec. IV carries over to this case; however, the shadow channel $\mathcal{M}$, Eq. (17), is *not* invertible. To see this, let us note that the states $\sigma_{\mathbf{m}}$ are diagonal in the charge: they are produced by starting from the fully mixed state $\mathbb{I}/D$ and acting with Z measurements and $U(1)$-symmetric gates, neither of which can create coherences between charge sectors. Then, defining for each operator O the *charge-diagonal* component $O_\Delta = \int (d\phi/2\pi) e^{-i\phi\hat{Q}} O e^{i\phi\hat{Q}} = \sum_Q \hat{\Pi}_Q O \hat{\Pi}_Q$, we have

$$\begin{aligned}\mathcal{M}(O_\Delta) &= D\sum_{\mathbf{m}} \pi_{\mathbf{m}} \mathrm{Tr}(\sigma_{\mathbf{m}} O_\Delta)\sigma_{\mathbf{m}} \\ &= D\sum_{\mathbf{m}} \pi_{\mathbf{m}} \mathrm{Tr}(\sigma_{\mathbf{m},\Delta} O)\sigma_{\mathbf{m}} = \mathcal{M}(O)\end{aligned} \tag{51}$$

(we have used cyclicity of the trace and the fact that $\sigma_{\mathbf{m},\Delta} = \sigma_{\mathbf{m}}$ for all $\mathbf{m}$). It follows that all operators that are off diagonal in the charge are in the kernel of $\mathcal{M}$:

$$\mathcal{M}(O - O_\Delta) = 0. \tag{52}$$

This means that such operators are *not learnable*, with any number of samples, under this classical shadow protocol. However, when restricted to the subspace of charge-diagonal operators, $\mathcal{M}$ may become invertible and it may be possible to successfully learn all charge-diagonal operators. Here we focus on the problem of learning total charge Q, which by definition is charge diagonal and so in principle learnable via eavesdropper's shadows. The sample complexity of this task is determined by the squared shadow norm $\|\hat{Q}\|^2_{\rm sh}$, which we study next.

B. Shadow norm of the charge operator

The shadow norm in general is given by

$$\|\hat{Q}\|^2_{\rm sh} = D\mathrm{Tr}(\rho \otimes \mathcal{M}^{-1}(\hat{Q})^{\otimes 2}\sigma^{(3)}), \tag{53}$$

which depends on the initial state ρ. Different choices for the initial state are possible—e.g., Ref. [24] considers initial states that are in either of two charge sectors Q_0, Q_1. Here for simplicity and generality we take an average over all possible initial states (according to any

1-design distribution, such that $\int d\rho\, \rho = \mathbb{I}/D$). This yields the state-averaged shadow norm:

$$\|\hat{Q}\|^2_{\text{sh,avg}} = \text{Tr}(\mathcal{M}^{-1}(\hat{Q})^{\otimes 2}\sigma^{(2)}) = \frac{1}{D}\text{Tr}[\hat{Q}\mathcal{M}^{-1}(\hat{Q})], \tag{54}$$

where the first equality comes from the fact that $\text{Tr}_1(\sigma^{(3)}) = \sigma^{(2)}$ and the second from the definition of $\mathcal{M}$, Eq. (17).

It is now helpful to introduce super-ket and -bra notation and write Eq. (54) as $\|\hat{Q}\|^2_{\text{sh,avg}} = (\hat{Q}|\mathcal{M}^{-1}|\hat{Q})/D$. For all invertible Hermitian forms H and all nonzero complex vectors v we have, by convexity, $v^\dagger H^{-1} v/(v^\dagger v) \geq [v^\dagger H v/(v^\dagger v)]^{-1}$. Therefore, the following bound holds:

$$\|\hat{Q}\|^2_{\text{sh,avg}} \geq \frac{(\hat{Q}|\hat{Q})^2}{D(\hat{Q}|\mathcal{M}|\hat{Q})}. \tag{55}$$

The numerator is computed straightforwardly:

$$(\hat{Q}|\hat{Q}) = \text{Tr}(\hat{Q}^2) = \frac{1}{4}\sum_{i,j}\text{Tr}(Z_i Z_j) = \frac{N}{4}D. \tag{56}$$

For the denominator, we note

$$\frac{(\hat{Q}|\mathcal{M}|\hat{Q})}{D} = \sum_{\mathbf{m}} \pi_{\mathbf{m}} \text{Tr}(\hat{Q}\sigma_{\mathbf{m}})^2 = \mathbb{E}_{\mathbf{m}}[\langle\hat{Q}\rangle^2_{\mathbf{m}}]. \tag{57}$$

This is the *variance across trajectories* of the charge expectation value. (Since $\mathbb{E}_{\mathbf{m}}\sigma_{\mathbf{m}} = \mathbb{I}/D$, we have $\mathbb{E}_{\mathbf{m}}[\langle\hat{Q}\rangle_{\mathbf{m}}] = \text{Tr}(\hat{Q})/D = 0$ and thus $\mathbb{E}_{\mathbf{m}}[\langle\hat{Q}\rangle^2_{\mathbf{m}}] = \text{var}_{\mathbf{m}}\langle\hat{Q}\rangle_{\mathbf{m}}$.) This is directly related to the order parameter of the charge-sharpening transition used in Ref. [12], the *trajectory-averaged charge fluctuation* $\delta Q = \mathbb{E}_{\mathbf{m}}[\langle\hat{Q}^2\rangle_{\mathbf{m}} - \langle\hat{Q}\rangle^2_{\mathbf{m}}]$: in particular, we have

$$\text{var}_{\mathbf{m}}[\langle\hat{Q}\rangle_{\mathbf{m}}] = \frac{N}{4} - \delta Q. \tag{58}$$

This follows from the fact that the first term in δQ is taken in the fully mixed state and gives $\text{Tr}(\hat{Q}^2)/D = N/4$.

We can thus recast the bound Eq. (57) in terms of the charge-sharpening order parameter:

$$\|\hat{Q}\|^2_{\text{sh,avg}} \geq \frac{\delta Q_0}{1 - \delta Q(t)/\delta Q_0}, \tag{59}$$

where $\delta Q_0 = N/4 = \text{Tr}(\hat{Q}^2)$ is the quantum fluctuation of charge in a completely mixed state. Below we work out the consequences of this bound in each phase.

Sharp phase. We have $\delta Q(t) \sim \delta Q_0 e^{-ct}$ with $c > 0$ a constant. At sufficiently large constant depth $t \sim \frac{1}{c}\log(1/\epsilon)$ the charge sharpens to within tolerance ϵ, and we have $\|\hat{Q}\|^2_{\text{sh}} \geq (1-\epsilon)N/4$: it takes $\Omega(N)$ experiments, or a total of $\Omega(N^2)$ measurements, to learn the charge of the initial state within error ϵ. This bound is expected from a simple central limit theorem argument [24] and would apply even upon measuring all qubits immediately, so the bound in this phase is trivial [Note that the random initial state we consider need not have definite charge. It has mean charge 0 and fluctuations $O(\sqrt{N})$, so lowering the uncertainty to $O(1)$ takes of order N experiments. This is unlike Ref. [24], where the initial state is promised to be of definite charge and so a single shot may suffice]. It is easy to see that, if each outcome $\mathbf{m}$ uniquely specifies a value Q of the charge (as is the case after sharpening), then $O(N)$ experiments are also sufficient.

Fuzzy phase. We have [12] $\delta Q(t) \sim \delta Q_0 e^{-ct/N}$—i.e., the sharpening time $t_\#$ diverges linearly in system size N. This gives

$$\|\hat{Q}\|^2_{\text{sh,avg}} \geq \frac{N/4}{1 - e^{-ct/N}} \simeq \frac{N^2}{4ct}, \tag{60}$$

where the approximation holds at times $t \ll N$. At finite t, this proves that of order N^2/t samples are needed, for a total of $M_{\text{tot}} \sim (N^2/t) \times (Nt) \sim N^3$ local measurements. This result is parametrically larger than in the sharp phase, scaling as N^3 rather than N^2. Furthermore, it is analogous to our previous result on learning Pauli operators, Sec. V D, in that an invariant unit of information extracted per measurement emerges. [This holds as long as t is not much larger than the sharpening time $t_\# \sim N$ so that the scaling ansatz for $\delta Q(t)$ applies.] In the fuzzy phase, this unit is suppressed by a factor of N relative to the sharp phase.

To summarize, we have analyzed the state-averaged shadow norm of the charge operator $\|\hat{Q}\|^2_{\text{sh,avg}}$, which quantifies the number of experimental repetitions needed to learn the charge $\langle\hat{Q}\rangle$ of an unknown initial state from the measurement record and knowledge of the circuit. We have shown that this behaves differently in the two measurement-induced phases: it is $O(N)$ throughout the sharp phase, whereas it is bounded below by approximately N^2/t in the fuzzy phase. The total number of measurements (there are of order Nt measurements per experimental repetition) thus transitions from $O(N^2)$ (sharp phase) to $\Omega(N^3)$ (fuzzy phase). Thus, in the same way in which the entanglement phase transition was shown to be a transition in learnability of generic properties such as Pauli expectation values (Sec. V) and many-body fidelities (Sec. VI), the charge-sharpening transition emerges as a transition in learnability of the charge $\hat{Q}$ in symmetric monitored dynamics.

VIII. DISCUSSION

A. Summary

We have presented a *learnability* perspective on measurement-induced phases of quantum information, based on the ability of an eavesdropper to learn properties of an unknown quantum state of the system from classical midcircuit measurement data. The learnability perspective is complementary to more well-established perspectives on the MIPT based on the entanglement properties of postmeasurement states of the quantum many-body system, or to the *coding* perspective, which focuses on recovering quantum information from the combined quantum-classical state of the many-body system and measurement device. In our setting, the eavesdropper, Eve, collects measurement records $\mathbf{m}$ from multiple repetitions of the experiment. In combination with a complete classical description of the dynamics and unconstrained classical computing resources, Eve can try to learn properties of the system's unknown initial state ρ—*without* access to the final, postmeasurement state.

We have shown that the MIPT generically coincides with a complexity phase transition for these tasks, i.e., a transition in the number of measurement outcomes needed to predict properties of ρ to a fixed accuracy. This is underpinned by a transition in the *informational power* of the POVM associated to monitored dynamics. This is an invariant measure of the information flow from the quantum state to the classical measurement record (technically a channel capacity), which sharply changes from finite to exponentially small at the MIPT.

We give an operational meaning to this transition through the framework of classical shadows, which furnishes a unified language to describe learnability transitions in various different contexts. To this end, we introduced a family of classical shadows protocols that the eavesdropper may use to concretely predict properties of ρ, and analytically showed that they carry signatures of the MIPT; namely the shadow channel used in the estimation process depends on the second moment of an associated ensemble of monitored quantum trajectories, which can undergo a MIPT. We have then unpacked the consequences of this result on several estimation tasks of interest.

Pauli expectation values. We have found the critical point $p = p_c$ to be (on average) optimal for the estimation of Pauli operators—the critical point balances the negative effects of spatial locality in the disentangling phase (which makes learning large Pauli operators more difficult) against the overall lack of information in the entangling phase. Washing out the effects of locality with a prescrambling step (i.e., a sufficiently deep random unitary circuit preceding the monitored dynamics), we analytically derived the sample complexity of Pauli estimation for any traceless Pauli to be proportional to $q^{(1+s)N}$, with q the local Hilbert-space dimension, N the size of the system, and $s \in [0, 1]$ the order parameter of the entangling phase (entropy density). The MIPT thus manifests as a nonanalyticity in the coefficient of the exponential in this case.

Many-body fidelity. We analytically derived the sample complexity of fidelity estimation (with prescrambling) to be proportional to q^{sN}, transitioning from constant to exponential at the MIPT. Furthermore, we found a close connection between shadow estimation of the fidelity and a previously proposed order parameter for the MIPT based on a linear cross-entropy diagnostic [23].

Charge. We considered models of monitored dynamics with a $U(1)$ symmetry and derived the sample complexity of learning the global charge expectation $\langle \hat{Q} \rangle$ on the initial state. We have found a transition, this time between distinct power laws in N, at the charge-sharpening transition [12]. Thus we have recast previous results on charge learnability transitions in the unified language of shadow estimation, on the same footing as the other learnability transitions identified above. A striking result, derived both for the entanglement and charge-sharpening transitions, is the emergence of an invariant amount of information extracted *per measurement* by the eavesdropper.

B. Experimental implications

In this work we have focused uniquely on the *sample complexity* of the various learning tasks: how many repetitions of the quantum experiment are necessary for learning. We have intentionally neglected the issue of *classical* computational complexity. This is key to any practical considerations.

Several recent works have, in various ways, introduced "hybrid" quantum-classical order parameters for the MIPT that trade experimental sample complexity (associated to a "postselection overhead" of obtaining postmeasurement properties) for classical computational complexity [19,21,23,36–38]. This is generally advantageous as (i) classical resources are cheaper and more available than quantum ones, (ii) the required classical simulation may in fact be efficient (e.g., in Clifford or matchgate circuits, etc.), and (iii) even when that is not the case, the exponential barrier for classical simulation [typically of order $\exp(N)$] may be more favorable than that of quantum sampling [typically of order $\exp(NT)$, T being the duration of the dynamics].

Our perspective in this work automatically leads to a family of these quantum-classical order parameters, namely the *variances* of shadow estimators for various properties of ρ. Such variances can be estimated from a small number of experimental datapoints in both phases. For example, learning a Pauli expectation value $\langle P \rangle$ in the mixed phase is hard due to the large variance $\mathrm{var}(\hat{p}) \sim D^{1+s}$ of the shadow estimator $\hat{p}$; our estimate of $\langle P \rangle$ after M iterations of the experiment, $(1/M)\sum_{i=1}^{M} \hat{p}_i$, carries an uncertainty approximately $\sqrt{\mathrm{var}(\hat{p})/M}$, so that $M \sim D^{1+s}$ experiments are needed in order to make the error small.

However, the variance $\mathrm{var}(\hat{p})$ itself is easily estimated from $M = O(1)$ iterations of the experiment, and its scaling with system size N (proportional to $q^{[1+s]N}$) serves as an order parameter of the MIPT. We note that all samples generated by the quantum experiment are valid and can be used for shadow estimation. The practical complexity reduces entirely to *classical computation*, as various steps of shadow estimation (computation of the snapshots, the shadow channel, the inverted snapshots, and the observable estimators) may require exponential classical resources.

C. Outlook

Our work opens some interesting directions for future work. First of all, we have chosen to focus only on the *sample complexity* of learning, neglecting the classical computational complexity. It would be interesting, especially with an eye to practical applications, to revisit our results with restrictions on classical computation resources: how much can Eve learn, e.g., with only polynomial-time classical algorithms? The methods used in this work generically require exponential-time computation; are there other more efficient methods that can still capture the learnability transition?

For the problem of learning the global charge from $U(1)$-symmetric dynamics [24], it was found that polynomial-time decoders still give rise to learnability transitions, albeit at a larger measurement rate p; with increasing computational resources, one eventually recovers the "intrinsic" transition at $p = p_\#$ (the charge-sharpening transition [12]). Does a similar picture hold for the entanglement transition? We have found that, with unlimited classical computation, the MIPT $p = p_c$ yields a learnability transition; would the transition move to some larger measurement rate $p > p_c$ upon restricting classical computational resources? The percolation threshold (e.g., $p = 1/2$ in one-dimensional brickwork circuits [1]), above which the monitored dynamics breaks down into finite-sized space-time regions, may be an upper bound for such a transition in polynomial-time learnability. We leave this question as an interesting direction for follow-up research.

In this work we have focused on *non-Gaussian* systems. However, Gaussian systems can exhibit potentially richer types of measurement-induced criticality and different entanglement phase diagrams (e.g., without a stable volume-law phase) [8,10,68]. It would be interesting to understand how such behaviors map onto the performance of eavesdropper's shadows, relative to standard *"matchgate shadows"* for fermionic systems [46].

Another open question is the precise nature of eavesdropper's shadows across the phase diagram, and especially at the critical point. We have found that, without prescrambling, the MIPT appears as an optimum in the learnability of typical or average Pauli operators on the system, see Sec. V B and Fig. 3. This is due to the combination of two effects: the lack of informational power in the entangling phase, and the role of locality in the disentangling phase (the latter makes large Pauli operators particularly hard to learn). We have conjectured that in the entangling phase, eavesdropper's shadows function similarly to random Clifford shadows [6], up to an overall inflation of sample complexity by the exponential factor q^{sN} (due to the suppressed informational power); whereas in the disentangling phase, they function similarly to *shallow shadows* [48–51] of variable depth (recovering the zero-depth limit, i.e., random Pauli measurements, at $p = 1$). This leaves the question of eavesdropper's shadows at the MIPT. The logarithmic scaling of entanglement at the critical point is expected to yield a distinctive fingerprint on the shadow norm distribution [via Eq. (28)], perhaps similar to classical shadows based on tree tensor networks. Testing and quantifying these conjectures, and finding potential applications, are exciting directions for future research.

Note added.—Recently, we became aware of a related work on classical shadows from monitored dynamics [69].

ACKNOWLEDGMENTS

We thank Ehud Altman, Bryan Clark, Xiaozhou Feng, Sam Garratt, Sarang Gopalakrishnan, and Tibor Rakovszky for helpful discussions. We are especially grateful to Yaodong Li for insightful discussions in the early stages of this project. M.I. was partly supported by the Gordon and Betty Moore Foundation's EPiQS Initiative through Grant GBMF8686. V.K. acknowledges support from the US Department of Energy, Office of Science, Basic Energy Sciences, under Early Career Award No. DE-SC0021111, from the Alfred P. Sloan Foundation through a Sloan Research Fellowship and from the Packard Foundation through a Packard Fellowship in Science and Engineering. Numerical simulations were carried out on Stanford Research Computing Center's Sherlock cluster. This project originated at the KITP program *"Quantum Many-Body Dynamics and Noisy Intermediate-Scale Quantum Systems"*; KITP is supported by the National Science Foundation under Grant No. NSF PHY-1748958.

APPENDIX A: INFORMATIONAL POWER

In this Appendix we collect various technical results related to the computation of the informational power, Sec. III.

1. Derivation of Eq. (10)

Here we derive the analytical expression for the prescrambled informational power in terms of the subentropy, Eq. (10).

With prescrambling by a random many-body unitary V, our POVM is given by $\Pi = \{V^\dagger E_{\mathbf{m}} V\}$, with effects indexed by the pair $(\mathbf{m}, V)$, which plays the role of a generalized "outcome." The pair $(\mathbf{m}, V)$ occurs with probability $\mathrm{Tr}(V\rho V^\dagger E_{\mathbf{m}})$ (note this is a probability density over the continuous variable V with the Haar measure dV. The POVM normalization condition reads $\sum_{\mathbf{m}} \int dV\, (V^\dagger E_{\mathbf{m}} V) = \sum_{\mathbf{m}} \mathrm{Tr}(E_{\mathbf{m}}) \mathbb{I}/D = \mathbb{I}$.). Due to prescrambling, *all* pure-state ensembles $\mathcal{E} = \{(p_i, |\psi_i\rangle\langle\psi_i|)\}$ that unravel the same density matrix ρ have the same mutual information with Π [59]: indeed, we have

$$I(\mathcal{E} : \Pi) = \sum_i p_i \sum_{\mathbf{m}} \int dV\, \langle\psi_i|V^\dagger E_{\mathbf{m}} V|\psi_i\rangle \times \ln \frac{\langle\psi_i|V^\dagger E_{\mathbf{m}} V|\psi_i\rangle}{\mathrm{Tr}(E_{\mathbf{m}}\rho)}; \tag{A1}$$

then, introducing a Haar-random state $|\phi\rangle \equiv V|\psi_i\rangle$ and replacing integration over the unitary group (Haar measure dV) with integration over the Hilbert space (Haar measure $d\phi$), we obtain

$$I(\mathcal{E} : \Pi) = \sum_{\mathbf{m}} \pi_{\mathbf{m}} \int d\phi\, f\left(\langle\phi|D\sigma_{\mathbf{m}}|\phi\rangle\right) - \sum_{\mathbf{m}} \pi_{\mathbf{m}} \int dV f\left[\mathrm{Tr}(D\sigma_{\mathbf{m}} V\rho V^\dagger)\right], \tag{A2}$$

where $\sigma_{\mathbf{m}}$ is as in Eq. (11) and we have introduced the shorthand $f(x) = x\ln(x)$. This shows that all pure-state ensembles $\mathcal{E}$ that unravel the same ρ have the same mutual information $I(\mathcal{E} : \Pi)$, as claimed.

We can now proceed to maximize the mutual information. It can be shown that the second line of Eq. (A2) is ≤ 0 [by convexity of $f(x)$], thus it can be maximized by setting $\rho = \mathbb{I}/D$. Next, we note that the mutual information is nondecreasing under replacement of a mixed state ρ_i in the ensemble with a pure decomposition, $(p_i, \rho_i) \mapsto \{(p_{ij}, |\psi_{ij}\rangle\langle\psi_{ij}|)\}$ such that $\sum_j p_{ij} |\psi_{ij}\rangle\langle\psi_{ij}| = p_i\rho_i$ [This follows from convexity of $f(x) = x\ln(x)$]. Thus the optimal ensemble can always be written in terms of pure states only. It follows that the optimization is trivial—any 1-design ensemble of pure states (e.g., the computational basis) maximizes the mutual information. The informational power, following Eq. (A2), is thus given by

$$W(\Pi) = \sum_{\mathbf{m}} \pi_{\mathbf{m}} \mathcal{G}(\sigma_{\mathbf{m}}), \tag{A3}$$

$$\mathcal{G}(\sigma) = \int d\phi\, \langle\phi|D\sigma|\phi\rangle \ln \langle\phi|D\sigma|\phi\rangle. \tag{A4}$$

The Haar integral $\mathcal{G}$ has been worked out for general σ in Ref. [59], and gives

$$\mathcal{G}(\sigma) = Q(\mathbb{I}/D) - Q(\sigma), \tag{A5}$$

where $Q(\sigma)$ is the subentropy, see below. This yields Eq. (10).

2. Subentropy

Here we provide some more details about the subentropy $Q(\rho)$ [59].

Like the von Neumann entropy $S(\rho)$, the subentropy $Q(\rho)$ is solely a function of the spectrum of ρ, $\{\lambda_j\}_{j=1}^D$:

$$Q(\rho) = -\sum_{j=1}^{D} \frac{\lambda_j \ln \lambda_j}{\prod_{k\neq j}(1-\lambda_k/\lambda_j)}. \tag{A6}$$

If the spectrum is degenerate with $\lambda_i = \lambda_j$, the formula can be regularized by taking a limit $\lambda_i \to \lambda_j$; $Q(\rho)$ is finite and well defined.

Another connection between entropy and subentropy is that they bound (above and below, respectively) the "accessible information" of a state ensemble $\mathcal{E}$, defined as

$$A(\mathcal{E}) \equiv \max_{\Pi} I(\mathcal{E} : \Pi). \tag{A7}$$

Note the duality with the informational power of a POVM, cf. Eq. (9). As shown in Ref. [59], one has $Q(\rho) \le A(\mathcal{E}) \le S(\rho)$, where $\rho = \sum_i p_i\rho_i$ is the average density matrix of the state ensemble. The ensemble $\mathcal{E}$ attaining the lower bound $A(\mathcal{E}) = Q(\rho)$ for a given ρ is known as the *Scrooge ensemble* and has recently emerged as a candidate universal distribution for postmeasurement states of subsystems in chaotic dynamics [70–72].

An important difference with the entropy is that the subentropy cannot be extensive. In fact it is bounded above by a constant: we have $Q(\rho) \le Q(\mathbb{I}/D) = 1 - \gamma - \delta H(D) < 1 - \gamma < 0.424$. [Here $\gamma = 0.577\ldots$ is Euler-Mascheroni's constant and $\delta H(x) = \sum_{j=1}^{x} 1/j - \ln(x) - \gamma$, as in the main text.]

3. Informational power of Clifford monitored dynamics

Here we derive Eq. (12), which underlies the result for the informational power of Clifford monitored circuits, Eq. (13).

We consider the Haar integral $\mathcal{G}(\rho)$ from Eq. (A4) for the case in which ρ is proportional to a projector, $\rho = \Pi/r$ where $\Pi^2 = \Pi$ and r is the rank of Π. This includes the case of stabilizer states. First we use a standard replica trick to write

$$\mathcal{G}(\rho) = \partial_n \int d\phi\, \langle\phi|D\rho|\phi\rangle^{1+n}\Big|_{n=0}. \tag{A8}$$

We then evaluate the integral for integer values of n by using the form of the $(n+1)$th moment of the Haar

measure on a D-dimensional Hilbert space [70,72]:

$$\rho^{(n+1)}_{\mathrm{Haar},D} \equiv \int d\phi |\phi\rangle\langle\phi|^{\otimes n+1} = \frac{(D-1)!}{(D+n)!} \sum_{\sigma\in S_{n+1}} \hat{\sigma}, \quad \text{(A9)}$$

where σ is a permutation of $n+1$ elements and $\hat{\sigma}$ is the associated replica permutation operator. We obtain

$$\mathcal{G}(\rho) = \partial_n \left[\frac{(D-1)!}{(D+n)!} \sum_{\sigma\in S_{n+1}} \mathrm{Tr}((D\rho)^{\otimes n+1}\hat{\sigma}) \right]_{n=0}. \quad \text{(A10)}$$

Using the fact that $\rho = \Pi/r$ (r being the rank of the projector Π) we have $\mathrm{Tr}(\Pi^{\otimes n+1}\hat{\sigma}) = \prod_{i=1}^{|\sigma|} \mathrm{Tr}(\Pi^{n_i}) = r^{|\sigma|}$, with $|\sigma|$ the number of cycles in the permutation σ, and n_i the length of each cycle. Thus

$$\mathcal{G}(\rho) = \partial_n \left[\frac{(D-1)!}{(D+n)!} \left(\frac{D}{r}\right)^{n+1} \sum_{\sigma\in S_{n+1}} r^{|\sigma|} \right]_{n=0}. \quad \text{(A11)}$$

The summation over σ can be done exactly by noting that, upon taking the trace of Eq. (A9) and replacing $D \mapsto r$, one has

$$\frac{(r-1)!}{(r+n)!} \sum_{\sigma\in S_{n+1}} r^{|\sigma|} = \mathrm{Tr}(\rho^{(n+1)}_{\mathrm{Haar},r}) = 1. \quad \text{(A12)}$$

It follows that

$$\mathcal{G}(\rho) = \partial_n \left[\left(\frac{D}{r}\right)^n \frac{D!}{(D+n)!} \frac{(r+n)!}{r!} \right]_{n=0}, \quad \text{(A13)}$$

where the Hilbert-space replicas are gone and we can now take the derivative (i.e., replica limit).

Analytically continuing the factorial to the Γ function, we have $(x+n)! = x! + x!(H_x - \gamma)n + O(n^2)$, with $H_x = \sum_{j=1}^{x} 1/j$ the harmonic sum. We conclude

$$\mathcal{G}(\rho) = \ln(D/r) + H_r - H_D = \delta H(r) - \delta H(D), \quad \text{(A14)}$$

with $\delta H(x) = H_x - (\ln(x) + \gamma)$ as in the main text. Finally, writing the rank r of the state in terms of the entropy as q^S we obtain Eq. (12).

4. Renyi-2 informational power of general monitored dynamics

Here we introduce a Renyi-2 version of the informational power, which is computable for general (nonstabilizer) states. Note that the Renyi-2 version of the mutual information on which this construction is based is *not* a valid mutual information (e.g., does not obey positivity); nonetheless in randomized settings it often behaves in a qualitatively similar way to the true mutual information, so our result here is suggestive of the presence of an informational power transition in general (nonstabilizer) monitored dynamics.

We define the Renyi-2 informational power W_2 by maximizing (over state ensembles $\mathcal{E}$) the Renyi-2 "mutual information"

$$I_2(\mathcal{E},\Pi) = S_2(p_i) + S_2(p_\alpha) - S_2(p_{i,\alpha}) = \ln \frac{\sum_{i,\alpha} p_{i,\alpha}^2}{\sum_{i,\alpha} p_i^2 p_\alpha^2}. \quad \text{(A15)}$$

Going through the same manipulations as in Eq. (A2), we have

$$I_2(\mathcal{E},\Pi) = \ln \frac{\sum_i p_i^2 \sum_{\mathbf{m}} \int d\phi\, \langle\phi|E_{\mathbf{m}}|\phi\rangle^2}{\sum_i p_i^2 \sum_{\mathbf{m}} \mathrm{Tr}(E_{\mathbf{m}})^2/D^2}. \quad \text{(A16)}$$

Defining modified probabilities $\tilde{p}_i = p_i^2/\sum_j p_j^2$ and $\tilde{\pi}_{\mathbf{m}} = \pi_{\mathbf{m}}^2/\sum_{\mathbf{m}'} \pi_{\mathbf{m}'}^2$ we arrive at

$$I_2(\mathcal{E},\Pi) = \ln \sum_i \tilde{p}_i \sum_{\mathbf{m}} D^2 \tilde{\pi}_{\mathbf{m}} \int d\phi\, \langle\phi|\sigma_{\mathbf{m}}|\phi\rangle^2. \quad \text{(A17)}$$

The integrand is independent of i, which again shows that the mutual information is independent of $\mathcal{E}$ owing to prescrambling. We get

$$W_2(\Pi) = \ln \sum_{\mathbf{m}} \tilde{\pi}_{\mathbf{m}} \frac{D}{D+1} \left[1 + \mathrm{Tr}(\sigma_{\mathbf{m}}^2)\right] = \ln \frac{1+\mathcal{P}}{1+1/D}, \quad \text{(A18)}$$

where $\mathcal{P}$ is the average purity of the trajectories $\sigma_{\mathbf{m}}$, averaged over the modified distribution $\tilde{\pi}_{\mathbf{m}}$.

It follows that the above-defined "Renyi-2 informational power" exhibits a MIPT: in the mixed phase, with $\mathcal{P} = D^{-s}$ and $s > 0$, we have $W_2(\Pi) \to 0$ in the large-system limit; in the pure phase, with $\mathcal{P} = q^{-S}$ and S finite, we have $W_2(\Pi) \to \ln(1+q^{-S}) > 0$. Note that, due to the modified measure over trajectories, the transition is in a different universality class from the standard one [22,23].

APPENDIX B: ALTERNATIVE CONSTRUCTIONS FOR EAVESDROPPER'S SHADOWS

Here we complete the discussion in Sec. IV by providing details on multiple options for classical shadows protocols based on generalized measurements, and how they relate to the MIPT. We start from the prescription discussed in the main text, and how it can be formally interpreted as *Petz recovery* of the quantum state ρ from the measurement record $\mathbf{m}$. We then review two existing approaches for

shadows based on generalized measurements—one based on least squares, one on maximum fidelity. We obtain the associated shadow channels and discuss how they are sensitive to the MIPT.

1. Petz recovery

We begin with the prescription followed in Sec. IV, i.e., $\eta_{\mathbf{m}} = \sigma_{\mathbf{m}}$. To complement the heuristic justification based on time-reversed monitored dynamics, we show that the same prescription arises as the *Petz recovery map* [60–63] relative to the channel $\mathcal{N}(\rho) = \sum_{\mathbf{m}} \mathrm{Tr}(E_{\mathbf{m}}\rho)|\mathbf{m}\rangle\langle\mathbf{m}|$ (mapping quantum states to classical measurement records) and the reference state $\rho_0 = \mathbb{I}/D$. The Petz map for a noise channel $\mathcal{N}$ and reference state ρ_0 is defined in general as

$$\mathcal{R}^{\rho_0,\mathcal{N}}_{\mathrm{Petz}}(\bullet) = \rho_0^{1/2}\mathcal{N}^\dagger\left[\mathcal{N}(\rho_0)^{-1/2}\bullet\,\mathcal{N}(\rho_0)^{-1/2}\right]\rho_0^{1/2}. \tag{B1}$$

The Petz map tries to undo the action of the channel $\mathcal{N}$, which plays the role of noise; it manifestly succeeds for the reference state ρ_0, as seen by plugging in $\bullet = \mathcal{N}(\rho_0)$ in Eq. (B1). It also succeeds for any other states whose relative entropy with ρ_0 is nondecreasing under the action of $\mathcal{N}$ [60,62]. These properties have made it a useful tool from formal quantum information to applications in error correction [73] and gravity [63,74]. For the task at hand, $\mathcal{R}_{\mathrm{Petz}}$ maps classical states (i.e., probability distributions over measurement records) to quantum states, which is indeed our goal: given some eavesdropped measurement record $\mathbf{m}$, predict the quantum state ρ it came from.

Using the fact that $\mathcal{N}(\rho_0) = \sum_{\mathbf{m}} \pi_{\mathbf{m}}|\mathbf{m}\rangle\langle\mathbf{m}|$ and $\mathcal{N}^\dagger[|\mathbf{m}\rangle\langle\mathbf{m}|] = E_{\mathbf{m}}$, explicit calculation of the Petz recovery on a classical state $p^{\mathrm{exp}} = \sum_{\mathbf{m}} p^{\mathrm{exp}}_{\mathbf{m}}|\mathbf{m}\rangle\langle\mathbf{m}|$ yields

$$\mathcal{R}^{\rho_0,\mathcal{N}}_{\mathrm{Petz}}(p^{\mathrm{exp}}) = \sum_{\mathbf{m}} p^{\mathrm{exp}}_{\mathbf{m}}\frac{E_{\mathbf{m}}}{D\pi_{\mathbf{m}}} = \sum_{\mathbf{m}} p^{\mathrm{exp}}_{\mathbf{m}}\sigma_{\mathbf{m}}. \tag{B2}$$

In conclusion, the Petz recovery prescription says that, given an experimental outcome $\mathbf{m}$ (representable as a δ-function distribution $p^{\mathrm{exp}}_{\mathbf{m}'} = \delta_{\mathbf{m},\mathbf{m}'}$), Eve should prepare the state $\sigma_{\mathbf{m}}$.

2. Least squares

Reference [45] proposes using $\eta_{\mathbf{m}} = E_{\mathbf{m}}$ based on a least-squares criterion. Namely, given an experimentally observed measurement record distribution $p^{\mathrm{exp}}_{\mathbf{m}}$, we can try to reconstruct the unknown state ρ by minimizing the cost function

$$\mathcal{L}(\rho) = \sum_{\mathbf{m}}\left[p^{\mathrm{exp}}_{\mathbf{m}} - \mathrm{Tr}(E_{\mathbf{m}}\rho)\right]^2, \tag{B3}$$

i.e., the two-norm of the distance between observed distribution $p^{\mathrm{exp}}_{\mathbf{m}}$ and predicted distribution $\mathrm{Tr}(E_{\mathbf{m}}\rho)$. At this stage it is convenient to introduce some extra notation: we use $|\mathbf{m}\rangle\!\rangle = |\mathbf{m}\rangle\langle\mathbf{m}|$ to denote classical states of the measurement record, and $|A)$ to denote quantum operators as states in a doubled Hilbert space. Defining again the quantum-to-classical channel $\mathcal{N} = \sum_{\mathbf{m}}|\mathbf{m}\rangle\!\rangle(E_{\mathbf{m}}|$ already encountered in the discussion of the Petz recovery above, we can write the predicted distribution for a given ρ as $\mathrm{Tr}(\rho E_{\mathbf{m}}) = \langle\!\langle\mathbf{m}|\mathcal{N}|\rho)$. The cost function thus reads

$$\mathcal{L}(\rho) = \left\||p^{\mathrm{exp}}\rangle\!\rangle - \mathcal{N}|\rho)\right\|^2, \tag{B4}$$

with $|p^{\mathrm{exp}}\rangle\!\rangle = \sum_{\mathbf{m}} p^{\mathrm{exp}}_{\mathbf{m}}|\mathbf{m}\rangle\!\rangle$ for short, and its optimization reduces to usual least squares, with the well-known result

$$|\hat{\rho}) = (\mathcal{N}^\dagger\circ\mathcal{N})^{-1}\mathcal{N}^\dagger|p^{\mathrm{exp}}\rangle\!\rangle. \tag{B5}$$

In a nutshell, this prescription says that for every run of the experiment, giving some outcome $\mathbf{m}$, we should construct a "snapshot" $\mathcal{N}^\dagger|\mathbf{m}\rangle\!\rangle = |E_{\mathbf{m}})$ and then an "inverted snapshot" $\mathcal{M}^{-1}(E_{\mathbf{m}})$, where the "measurement channel" $\mathcal{M}$ is given by

$$\mathcal{M} = \mathcal{N}^\dagger\circ\mathcal{N} = \sum_{\mathbf{m}}|E_{\mathbf{m}})(E_{\mathbf{m}}|. \tag{B6}$$

The operator $E_{\mathbf{m}}$ is evidently not a state, due to its trace normalization: we have $E_{\mathbf{m}} = D\pi_{\mathbf{m}}\sigma_{\mathbf{m}}$ in the notation introduced in Sec. II A. It follows that $\mathcal{M}$ is not a channel (it is not trace preserving). The shadow protocol works regardless, as the application of the inverse shadow channel $\mathcal{M}^{-1}$ takes care of the incorrect normalization.

Finally, it is helpful to rewrite the shadow channel in terms of a state ensemble dual to our POVM $\{E_{\mathbf{m}}\}$, in analogy with Eq. (17). We have

$$\begin{aligned}\mathcal{M}(\rho) &= \sum_{\mathbf{m}}\mathrm{Tr}(\rho E_{\mathbf{m}})E_{\mathbf{m}} = D^2\sum_{\mathbf{m}}\pi_{\mathbf{m}}^2\mathrm{Tr}(\rho\sigma_{\mathbf{m}})\sigma_{\mathbf{m}}\\ &= D^2\left(\sum_{\mathbf{m}}\pi_{\mathbf{m}}^2\right)\mathrm{Tr}[(\mathbb{I}\otimes\rho)\tilde{\sigma}^{(2)}],\end{aligned} \tag{B7}$$

where $\tilde{\sigma}^{(2)}$ is the second moment operator of the ensemble $\tilde{\mathcal{E}} = \{(\tilde{\pi}_{\mathbf{m}}, \sigma_{\mathbf{m}})\}$ defined by the usual states $\sigma_{\mathbf{m}} = E_{\mathbf{m}}/\mathrm{Tr}(E_{\mathbf{m}})$ [cf. Equation (11)] but with a modified probability distribution $\tilde{\pi}_{\mathbf{m}} = \pi_{\mathbf{m}}^2/\sum_{\mathbf{m}'}\pi_{\mathbf{m}'}^2$:

$$\tilde{\sigma}^{(2)} = \sum_{\mathbf{m}}\tilde{\pi}_{\mathbf{m}}\sigma_{\mathbf{m}}^{\otimes 2}. \tag{B8}$$

This trajectory ensemble also features a MIPT, but due to the modified weights it has a different universality [22] and generally occurs at a different (though empirically very close [23]) measurement rate.

3. Maximum fidelity

Finally, Ref. [44] proposes using the pure state corresponding to the leading eigenvalue (we neglect degeneracies at this stage) in $E_{\mathbf{m}}$: $|\psi_{\mathbf{m}}\rangle\langle\psi_{\mathbf{m}}| = \lim_{n\to\infty} E_{\mathbf{m}}^n/\mathrm{Tr}(E_{\mathbf{m}}^n)$. This choice maximizes the the Haar-averaged fidelity between input and output states of the shadow channel, $F = \int d\phi\, \langle\phi|\mathcal{M}(|\phi\rangle\langle\phi|)|\phi\rangle$: we have

$$F = D\sum_{\mathbf{m}} \pi_{\mathbf{m}} \int d\phi \langle\phi|\sigma_{\mathbf{m}}|\phi\rangle\langle\phi|\eta_{\mathbf{m}}|\phi\rangle = \sum_{\mathbf{m}} \pi_{\mathbf{m}} \frac{1+\mathrm{Tr}(\sigma_{\mathbf{m}}\eta_{\mathbf{m}})}{D+1}. \tag{B9}$$

This is maximized by taking $\eta_{\mathbf{m}} = |\psi_{\mathbf{m}}\rangle\langle\psi_{\mathbf{m}}|$, the projector on the leading eigenvector of $\sigma_{\mathbf{m}}$, as claimed. We note that this construction is closely related to a previous proposal for a notion of "quantumness" of a Hilbert space [75,76], based on the ability of a classical eavesdropper to read and resend the information without being detected.

The resulting shadow channel is given by

$$\mathcal{M}(\rho) = D\mathrm{Tr}[(\mathbb{I}\otimes\rho)\sigma^{(\infty,1)}], \tag{B10}$$

where we defined a generalized "moment operator"

$$\sigma^{(\infty,1)} = \sum_{\mathbf{m}} \pi_{\mathbf{m}} |\psi_{\mathbf{m}}\rangle\langle\psi_{\mathbf{m}}| \otimes \sigma_{\mathbf{m}} \tag{B11}$$

for the ensemble of trajectories. This operator is sensitive to the MIPT; for example, the expectation value of the replica SWAP operator $\hat{\tau}$ yields

$$\mathrm{Tr}(\sigma^{(\infty,1)}\hat{\tau}) = \sum_{\mathbf{m}} \pi_{\mathbf{m}} \langle\psi_{\mathbf{m}}|\sigma_{\mathbf{m}}|\psi_{\mathbf{m}}\rangle = \mathbb{E}_{\mathbf{m}}[e^{-S_{\infty,\mathbf{m}}}], \tag{B12}$$

the "annealed average" over trajectories of the Renyi-∞ entropy. The measure over trajectories in this case is the conventional one, $\pi_{\mathbf{m}}$.

APPENDIX C: XEB CALCULATIONS

Here we derive two results relating to the sample complexity of fidelity estimation from Sec. VI—Eq. (45) and (49)—both of which involve a third-moment expectation value.

1. Computation of third-moment quantity

We start by obtaining an auxiliary result: evaluating the third-moment quantity

$$\Gamma \equiv \mathrm{Tr}(\rho\otimes\rho_0^{\otimes 2}\sigma^{(3)}), \tag{C1}$$

with $\rho_0 = |\psi\rangle\langle\psi|$ and $\sigma^{(3)}$ the third-moment operator given in Eq. (48). Due to prescrambling, the latter is expressed as a sum of three-replica permutations,

$$\sigma^{(3)} = \sum_{\nu\in S_3} c_\nu \hat{\nu} = c_e\hat{e} + c_\tau(\hat{\tau}_{1,2}+\hat{\tau}_{2,3}+\hat{\tau}_{3,1}) + c_\chi(\hat{\chi}_+ + \hat{\chi}_-), \tag{C2}$$

where e is the identity permutation, $\tau_{i,j}$ is the transposition of elements i, j, and $\chi_\pm$ are the cyclical permutations. By plugging Eq. (C2) into the definition of Γ, we get

$$\Gamma = c_e + c_\tau + 2(c_\tau + c_\chi)F, \tag{C3}$$

where $F = \mathrm{Tr}(\rho\rho_0)$ is the fidelity between the true state ρ and the guess $\rho_0 = |\psi\rangle\langle\psi|$.

By making use of Weingarten functions [77]

$$\begin{cases} \mathrm{Wg}(e) = (D^2-2)/g(D) \\ \mathrm{Wg}(\tau_{ij}) = -D/g(D) \\ \mathrm{Wg}(\chi_\pm) = 2/g(D) \end{cases} \tag{C4}$$

with $g(D) = D(D^2-1)(D^2-4)$, we get

$$\begin{pmatrix} c_e \\ c_\tau \\ c_\chi \end{pmatrix} = \frac{1}{g(D)} \begin{pmatrix} D^2-2 & -3D & 4 \\ -D & D^2+2 & -2D \\ 2 & -3D & D^2 \end{pmatrix} \begin{pmatrix} 1 \\ \mathcal{P} \\ \mathcal{P}^{(3)} \end{pmatrix}, \tag{C5}$$

where $\mathcal{P}^{(3)} = \mathrm{Tr}(\chi_\pm\sigma^{(3)}) = \mathbb{E}_{\mathbf{m}}\mathrm{Tr}(\sigma_{\mathbf{m}}^3)$ is related to the third Renyi entropy of the trajectories. Therefore we have

$$c_e + c_\tau = \frac{D^2-D-2+\mathcal{P}(D^2-3D+2)+\mathcal{P}^{(3)}(4-2D)}{D(D^2-1)(D^2-4)} \simeq D^{-3}(1+\mathcal{P}) \tag{C6}$$

and

$$c_\tau + c_\chi = \frac{2-D+\mathcal{P}(D^2-3D+2)+\mathcal{P}^{(3)}(D^2-2D)}{D(D^2-1)(D^2-4)} \simeq D^{-3}(\mathcal{P}+\mathcal{P}^{(3)}). \tag{C7}$$

Here $\simeq$ denotes leading order in D. With this we conclude

$$\Gamma \simeq D^{-3}[1+\mathcal{P}+2F(\mathcal{P}+\mathcal{P}^{(3)})]. \tag{C8}$$

2. Statistical fluctuations of modified linear XEB

Here we derive Eq. (45). The standard deviation $\delta\mathrm{XEB}'$ is given by

$$(\delta\mathrm{XEB}')^2 = \langle p(\rho_0|\mathbf{m})^2\rangle_{\mathbf{m}\sim p(\mathbf{m}|\rho)} - \langle p(\rho_0|\mathbf{m})\rangle^2_{\mathbf{m}\sim p(\mathbf{m}|\rho)}. \tag{C9}$$

The second term is simply $(\mathrm{XEB}')^2$, already computed in Eq. (43). Focusing on the first, we have

$$\langle p(\rho_0|\mathbf{m})^2\rangle_{\mathbf{m}\sim p(\mathbf{m}|\rho)} = D^3\mathrm{Tr}(\rho\otimes\rho_0\otimes\rho_0\,\sigma^{(3)}), \tag{C10}$$

which is proportional to the Γ quantity in Eq. (C8). It follows that

$$(\delta\mathrm{XEB}')^2 = \mathcal{P} + 2F\mathcal{P}^{(3)} - (F\mathcal{P})^2. \tag{C11}$$

In the pure phase, both $\mathcal{P}$ and $\mathcal{P}^{(3)}$ are constant. In the mixed phase they both vanish asymptotically, with $\mathcal{P}\sim D^{-s}$ and $\mathcal{P}^{(3)}\leq\mathcal{P}$. Equation (45) follows.

3. Shadow norm computation

Here we derive Eq. (49). The inverse channel $\mathcal{M}^{-1}$, derivable from Eq. (30), reads $\mathcal{M}^{-1}(\rho) = \lambda^{-1}\rho - c\mathbb{I}$ with $\lambda = (D\mathcal{P}-1)/(D^2-1)$ and $c = (D-\mathcal{P})/(D\mathcal{P}-1)$. This yields

$$\|\rho_0\|^2_{\mathrm{sh}} = D\left\{c^2\mathrm{Tr}(\rho\sigma^{(1)}) - 2c\lambda^{-1}\mathrm{Tr}(\rho\otimes\rho_0\sigma^{(2)}) + \lambda^{-2}\mathrm{Tr}(\rho\otimes\rho_0^{\otimes 2}\,\sigma^{(3)})\right\}, \tag{C12}$$

where we have used the fact that moment operators obey $\mathrm{Tr}_1(\sigma^{(k)}) = \sigma^{(k-1)}$ (tracing over one of the k replicas yields the moment operator on $k-1$ replicas). The first two terms are straightforwardly evaluated by noting that, due to prescrambling, $\sigma^{(1)} = \mathbb{I}/D$ and $\sigma^{(2)} = [(1-\mathcal{P}/D)e + (\mathcal{P}-1/D)\chi]/(D^2-1)$ (e, χ are the two replica permutations, identity and swap). The last term, involving the third moment, is again given by the Γ quantity evaluated in Eq. (C8). Explicit evaluation yields Eq. (49).

[1] Brian Skinner, Jonathan Ruhman, and Adam Nahum, Measurement-induced phase transitions in the dynamics of entanglement, Phys. Rev. X **9**, 031009 (2019).

[2] Yaodong Li, Xiao Chen, and Matthew P. A. Fisher, Quantum Zeno effect and the many-body entanglement transition, Phys. Rev. B **98**, 205136 (2018).

[3] Yaodong Li, Xiao Chen, and Matthew P. A. Fisher, Measurement-driven entanglement transition in hybrid quantum circuits, Phys. Rev. B **100**, 134306 (2019).

[4] Andrew C. Potter and Romain Vasseur, in *Entanglement in Spin Chains: From Theory to Quantum Technology Applications*, edited by A. Bayat, S. Bose, and H. Johannesson (Quantum Science and Technology, Cham, 2022) pp. 211–249.

[5] Matthew P. A. Fisher, Vedika Khemani, Adam Nahum, and Sagar Vijay, Random quantum circuits, arXiv:2207.14280.

[6] Hsin-Yuan Huang, Richard Kueng, and John Preskill, Predicting many properties of a quantum system from very few measurements, Nat. Phys. **16**, 1050 (2020).

[7] Andreas Elben, Steven T. Flammia, Hsin-Yuan Huang, Richard Kueng, John Preskill, Benoit Vermersch, and Peter Zoller, The randomized measurement toolbox, Nat. Rev. Phys. **5**, 9 (2023).

[8] Xiangyu Cao, Antoine Tilloy, and Andrea De Luca, Entanglement in a fermion chain under continuous monitoring, SciPost Phys. **7**, 024 (2019).

[9] Lukasz Fidkowski, Jeongwan Haah, and Matthew B. Hastings, How dynamical quantum memories forget, Quantum **5**, 382 (2021).

[10] Michele Fava, Lorenzo Piroli, Tobias Swann, Denis Bernard, and Adam Nahum, Nonlinear sigma models for monitored dynamics of free fermions, Phys. Rev. X **13**, 041045 (2023).

[11] Yimu Bao, Soonwon Choi, and Ehud Altman, Symmetry enriched phases of quantum circuits, Ann. Phys. **435**, 168618 (2021).Special issue on Philip W. Anderson.

[12] Utkarsh Agrawal, Aidan Zabalo, Kun Chen, Justin H. Wilson, Andrew C. Potter, J. H. Pixley, Sarang Gopalakrishnan, and Romain Vasseur, Entanglement and charge-sharpening transitions in $U(1)$ symmetric monitored quantum circuits, Phys. Rev. X **12**, 041002 (2022).

[13] Shayan Majidy, Utkarsh Agrawal, Sarang Gopalakrishnan, Andrew C. Potter, Romain Vasseur, and Nicole Yunger Halpern, Critical phase and spin sharpening in SU(2)-symmetric monitored quantum circuits, Phys. Rev. B **108**, 054307 (2023).

[14] Soonwon Choi, Yimu Bao, Xiao-Liang Qi, and Ehud Altman, Quantum error correction in scrambling dynamics and measurement-induced phase transition, Phys. Rev. Lett. **125**, 030505 (2020).

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

[16] Peter W. Shor, Scheme for reducing decoherence in quantum computer memory, Phys. Rev. A **52**, R2493 (1995).

[17] Daniel Gottesman, Stabilizer codes and quantum error correction, arXiv:quant-ph/9705052.

[18] Eric Dennis, Alexei Kitaev, Andrew Landahl, and John Preskill, Topological quantum memory, J. Math. Phys. **43**, 4452 (2002).

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

[20] Crystal Noel, Pradeep Niroula, Daiwei Zhu, Andrew Risinger, Laird Egan, Debopriyo Biswas, Marko Cetina, Alexey V. Gorshkov *et al.*, Measurement-induced quantum phases realized in a trapped-ion quantum computer, Nat. Phys. **18**, 760 (2022).

[21] J. C. Hoke, M. Ippoliti, E. Rosenberg, D. Abanin, R. Acharya, T. I. Andersen, M. Ansmann, F. Arute *et al.*, Measurement-induced entanglement and teleportation on a noisy quantum processor, Nature **622**, 481 (2023).

[22] Yimu Bao, Soonwon Choi, and Ehud Altman, Theory of the phase transition in random unitary circuits with measurements, Phys. Rev. B **101**, 104301 (2020).

[23] Yaodong Li, Yijian Zou, Paolo Glorioso, Ehud Altman, and Matthew P. A. Fisher, Cross entropy benchmark for measurement-induced phase transitions, Phys. Rev. Lett. **130**, 220404 (2023).
[24] Fergus Barratt, Utkarsh Agrawal, Andrew C. Potter, Sarang Gopalakrishnan, and Romain Vasseur, Transitions in the learnability of global charges from local measurements, Phys. Rev. Lett. **129**, 200602 (2022).
[25] Michele Dall'Arno, Giacomo Mauro D'Ariano, and Massimiliano F. Sacchi, Informational power of quantum measurements, Phys. Rev. A **83**, 062304 (2011).
[26] Matteo Ippoliti, Michael J. Gullans, Sarang Gopalakrishnan, David A. Huse, and Vedika Khemani, Entanglement phase transitions in measurement-only dynamics, Phys. Rev. X **11**, 011030 (2021).
[27] Ali Lavasani, Yahya Alavirad, and Maissam Barkeshli, Measurement-induced topological entanglement transitions in symmetric random quantum circuits, Nat. Phys. **17**, 342 (2021).
[28] Adam Nahum, Sthitadhi Roy, Brian Skinner, and Jonathan Ruhman, Measurement and entanglement phase transitions in all-to-all quantum circuits, on quantum trees, and in Landau-Ginsburg theory, PRX Quantum **2**, 010352 (2021).
[29] Yaodong Li and Matthew P. A. Fisher, Statistical mechanics of quantum error correcting codes, Phys. Rev. B **103**, 104306 (2021).
[30] Ruihua Fan, Sagar Vijay, Ashvin Vishwanath, and Yi-Zhuang You, Self-organized error correction in random unitary circuits with measurement, Phys. Rev. B **103**, 174309 (2021).
[31] Xiaozhou Feng, Brian Skinner, and Adam Nahum, Measurement-induced phase transitions on dynamical quantum trees, PRX Quantum **4**, 030333 (2023).
[32] Jin Ming Koh, Shi-Ning Sun, Mario Motta, and Austin J. Minnich, Measurement-induced entanglement phase transition on a superconducting quantum processor with mid-circuit readout, Nat. Phys. **19**, 1314 (2023).
[33] Matteo Ippoliti and Vedika Khemani, Postselection-free entanglement dynamics via spacetime duality, Phys. Rev. Lett. **126**, 060501 (2021).
[34] Matteo Ippoliti, Tibor Rakovszky, and Vedika Khemani, Fractal, logarithmic, and volume-law entangled nonthermal steady states via spacetime duality, Phys. Rev. X **12**, 011045 (2022).
[35] Tsung-Cheng Lu and Tarun Grover Spacetime duality between localization transitions and measurement-induced transitions, PRX Quantum **2**, 040319 (2021).
[36] Hossein Dehghani, Ali Lavasani, Mohammad Hafezi, and Michael J. Gullans, Neural-network decoders for measurement induced phase transitions, Nat. Commun. **14**, 2918 (2023).
[37] Samuel J. Garratt, Zack Weinstein, and Ehud Altman, Measurements conspire nonlocally to restructure critical quantum states, Phys. Rev. X **13**, 021026 (2023).
[38] Samuel J. Garratt and Ehud Altman, Probing post-measurement entanglement without post-selection, arXiv: 2305.20092.
[39] Hsin-Yuan Huang, Learning quantum states from their classical shadows, Nat. Rev. Phys. **4**, 81 (2022).
[40] G. I. Struchalin, Y. A. Zagorovskii, E. V. Kovlakov, S. S. Straupe, and S. P. Kulik, Experimental estimation of quantum state properties from classical shadows, PRX Quantum **2**, 010307 (2021).
[41] Senrui Chen, Wenjun Yu, Pei Zeng, and Steven T. Flammia, Robust shadow estimation, PRX Quantum **2**, 030348 (2021).
[42] Dax Enshan Koh and Sabee Grewal, Classical shadows with noise, Quantum **6**, 776 (2022).
[43] Hsin-Yuan Huang, Richard Kueng, and John Preskill, Efficient estimation of Pauli observables by derandomization, Phys. Rev. Lett. **127**, 030503 (2021).
[44] Atithi Acharya, Siddhartha Saha, and Anirvan M. Sengupta, Shadow tomography based on informationally complete positive operator-valued measure, Phys. Rev. A **104**, 052418 (2021).
[45] H. Chau Nguyen, Jan Lennart Bonsel, Jonathan Steinberg, and Otfried Guhne, Optimizing shadow tomography with generalized measurements, Phys. Rev. Lett. **129**, 220502 (2022).
[46] Kianna Wan, William J. Huggins, Joonho Lee, and Ryan Babbush, Matchgate shadows for fermionic quantum simulation, Commun. Math. Phys. **404**, 629 (2023).
[47] Hong-Ye Hu, Soonwon Choi, and Yi-Zhuang You, Classical shadow tomography with locally scrambled quantum dynamics, Phys. Rev. Res. **5**, 023027 (2023).
[48] Ahmed A. Akhtar, Hong-Ye Hu, and Yi-Zhuang You, Scalable and flexible classical shadow tomography with tensor networks, Quantum **7**, 1026 (2023).
[49] Christian Bertoni, Jonas Haferkamp, Marcel Hinsche, Marios Ioannou, Jens Eisert, and Hakop Pashayan, Shallow shadows: Expectation estimation using low-depth random Clifford circuits, arXiv:2209.12924.
[50] Mirko Arienzo, Markus Heinrich, Ingo Roth, and Martin Kliesch, Closed-form analytic expressions for shadow estimation with brickwork circuits, arXiv:2211.09835.
[51] Matteo Ippoliti, Yaodong Li, Tibor Rakovszky, and Vedika Khemani, Operator relaxation and the optimal depth of classical shadows, Phys. Rev. Lett. **130**, 230403 (2023).
[52] Matteo Ippoliti, Classical shadows based on locally-entangled measurements, arXiv:2305.10723.
[53] Minh C. Tran, Daniel K. Mark, Wen Wei Ho, and Soonwon Choi, Measuring arbitrary physical properties in analog quantum simulation, Phys. Rev. X **13**, 011049 (2023).
[54] Saumya Shivam, Curt W. von Keyserlingk, and Shivaji L. Sondhi, On classical and hybrid shadows of quantum states, SciPost Phys. **14**, 094 (2023).
[55] Max McGinley and Michele Fava, Shadow tomography from emergent state designs in analog quantum simulators, Phys. Rev. Lett. **131**, 160601 (2023).
[56] Michele Dall'Arno, Accessible information and informational power of quantum 2-designs, Phys. Rev. A **90**, 052311 (2014).
[57] Michele Dall'Arno, Francesco Buscemi, and Masanao Ozawa, Tight bounds on accessible information and informational power, J. Phys. A: Math. Theor. **47**, 235302 (2014).
[58] Paul Hausladen and William K. Wootters, A 'pretty good' measurement for distinguishing quantum states, J. Mod. Opt. **41**, 2385 (1994).

[59] Richard Jozsa, Daniel Robb, and William K. Wootters, Lower bound for accessible information in quantum mechanics, Phys. Rev. A **49**, 668 (1994).
[60] Denes Petz, Sufficient subalgebras and the relative entropy of states of a von Neumann algebra, Commun. Math. Phys. **105**, 123 (1986).
[61] H. Barnum and E. Knill, Reversing quantum dynamics with near-optimal quantum and classical fidelity, J. Math. Phys. **43**, 2097 (2002).
[62] Mark M. Wilde, Recoverability in quantum information theory, Proc. R. Soc. A: Math. Phys. Eng. Sci. **471**, 20150338 (2015).
[63] Geoff Penington, Stephen H. Shenker, Douglas Stanford, and Zhenbin Yang, Replica wormholes and the black hole interior, J. High Energy Phys. **2022**, 205 (2022).
[64] Romain Vasseur, Andrew C. Potter, Yi-Zhuang You, and Andreas W. W. Ludwig, Entanglement transitions from holographic random tensor networks, Phys. Rev. B **100**, 134203 (2019).
[65] Chao-Ming Jian, Yi-Zhuang You, Romain Vasseur, and Andreas W. W. Ludwig, Measurement-induced criticality in random quantum circuits, Phys. Rev. B **101**, 104302 (2020).
[66] Kaifeng Bu, Dax Enshan Koh, Roy J. Garcia, and Arthur Jaffe, Classical shadows with Pauli-invariant unitary ensembles, arXiv:2202.03272.
[67] Fergus Barratt, Utkarsh Agrawal, Sarang Gopalakrishnan, David A. Huse, Romain Vasseur, and Andrew C. Potter, Field theory of charge sharpening in symmetric monitored quantum circuits, Phys. Rev. Lett. **129**, 120604 (2022).
[68] M. Buchhold, Y. Minoguchi, A. Altland, and S. Diehl, Effective theory for the measurement-induced phase transition of Dirac fermions, Phys. Rev. X **11**, 041004 (2021).
[69] Ahmed A. Akhtar, Hong-Ye Hu, and Yi-Zhuang You, Measurement-induced criticality is tomographically optimal, arXiv:2308.01653.
[70] Jordan S. Cotler, Daniel K. Mark, Hsin-Yuan Huang, Felipe Hernandez, Joonhee Choi, Adam L. Shaw, Manuel Endres, and Soonwon Choi, Emergent quantum state designs from individual many-body wave functions, PRX Quantum **4**, 010311 (2023).
[71] Joonhee Choi, Adam L. Shaw, Ivaylo S. Madjarov, Xin Xie, Ran Finkelstein, Jacob P. Covey, Jordan S. Cotler, Daniel K. Mark *et al.*, Preparing random states and benchmarking with many-body quantum chaos, Nature **613**, 468 (2023).
[72] Matteo Ippoliti and Wen Wei Ho, Solvable model of deep thermalization with distinct design times, Quantum **6**, 886 (2022).
[73] Hui Khoon Ng and Prabha Mandayam, Simple approach to approximate quantum error correction based on the transpose channel, Phys. Rev. A **81**, 062342 (2010).
[74] Jordan Cotler, Patrick Hayden, Geoffrey Penington, Grant Salton, Brian Swingle, and Michael Walter, Entanglement wedge reconstruction via universal recovery channels, Phys. Rev. X **9**, 031011 (2019).
[75] Christopher A. Fuchs and Masahide Sasaki, Squeezing quantum information through a classical channel: Measuring the "quantumness" of a set of quantum states, Quantum Inf. Comput. **3**, 377 (2003).
[76] Christopher A. Fuchs, On the quantumness of a Hilbert space, Quantum Inf. Comput. **4**, 467 (2004).
[77] Georg Kostenberger, Weingarten calculus, arXiv:2101.00921.