Shadow Tomography from Emergent State Designs in Analog Quantum Simulators

Max McGinley[1,2] and Michele Fava[1,3]
[1]*Rudolf Peierls Centre for Theoretical Physics, Clarendon Laboratory, Parks Road, Oxford OX1 3PU, United Kingdom*
[2]*T.C.M. Group, Cavendish Laboratory, JJ Thomson Avenue, Cambridge CB3 0HE, United Kingdom*
[3]*Philippe Meyer Institute, Physics Department, École Normale Supérieure (ENS), Université PSL, 24 rue Lhomond, F-75231 Paris, France*

(Received 12 January 2023; accepted 11 September 2023; published 18 October 2023)

We introduce a method that allows one to infer many properties of a quantum state—including nonlinear functions such as Rényi entropies—using only global control over the constituent degrees of freedom. In this protocol, the state of interest is first entangled with a set of ancillas under a fixed global unitary, before projective measurements are made. We show that when the unitary is sufficiently entangling, a universal relationship between the statistics of the measurement outcomes and properties of the state emerges, which can be connected to the recently discovered phenomeonon of emergent quantum state designs in chaotic systems. Thanks to this relationship, arbitrary observables can be reconstructed using the same number of experimental repetitions that would be required in classical shadow tomography [Huang *et al.*, Nat. Phys. **16**, 1050 (2020)]. Unlike previous approaches to shadow tomography, our protocol can be implemented using only global Hamiltonian evolution, as opposed to qubit-selective logic gates, which makes it particularly well suited to analog quantum simulators, including ultracold atoms in optical lattices and arrays of Rydberg atoms.

DOI: 10.1103/PhysRevLett.131.160601

Introduction.—The ability to control interactions in a many-body quantum system allows one to simulate and study other complex quantum systems of interest [1,2]. In a universal quantum computer, where logical gates can be selectively applied to a few qubits at a time, one can, in principle, mimic the dynamics of any Hamiltonian [3]; however, at present such devices are limited by their size and noisiness [4]. In contrast, analog quantum simulators—such as ultracold atoms in optical lattices [5,6] and arrays of Rydberg atoms [7–10]—typically possess global rather than site-specific control and, as such, are more tailored to synthesizing specific classes of Hamiltonians. Despite their limitations in terms of programmability, such platforms are often more scalable and less noisy than computationally universal devices and have already been used to shed light on a wide variety of many-body quantum phenomena [11–20].

In any such experiment, a key task is to infer the properties of some many-body state once it has been prepared. In computationally universal devices, a particularly powerful technique known as shadow tomography can be employed for this purpose [21–23], wherein random unitary rotations are applied before projective measurements of each qubit are made (see also [24–26]). Using this scheme, many properties of the state can be estimated simultaneously, and nonlinear properties such as Rényi entropies can also be accessed. However, measurement strategies of this kind currently involve the application of spatially inhomogeneous sequences of site-selective gates. While these operations are natural in digital devices, they are not available in analog quantum simulators, wherein all degrees of freedom evolve simultaneously under some global uniform Hamiltonian. Accordingly, the set of observables that can be directly accessed therein (efficiently or otherwise) is at present much more limited.

In this Letter, we bridge this gap by introducing a new protocol that allows one to simultaneously infer many properties of a state (including Rényi entropies, etc.) without needing to address each degree of freedom individually. Rather than applying inhomogeneous unitaries drawn randomly and compiled from few-qubit gates, we propose to apply some fixed deterministic global unitary U to the system together with a set of ancillas, followed by measurements in the computational basis [see Fig. 1(a)]. The unitary need not be fine-tuned, and so can be native to the system in question, making our protocol particularly well suited to analog quantum simulators. Importantly, our scheme offers the same performance guarantees as classical shadow tomography [23], meaning that the number of

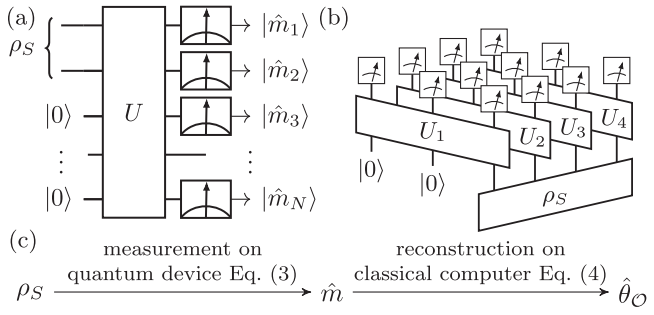

FIG. 1. (a) In our protocol, the target state ρ_S is evolved together with a set of ancillas under a fixed unitary U, before projective measurements are made in the computational basis. (b) For a many-body state, qubits can be subdivided into "blocks", each of which interact with a separate set of ancillas. (c) Schematic of the process for constructing estimators $\hat{\theta}_{\mathcal{O}}$ of expectation values $\langle\mathcal{O}\rangle$.

measurements needed to estimate a wide range of expectation values does not grow with system size.

We show that, for generic choices of U, a universal relationship between properties of the target state and the distribution of measurement outcomes emerges. Specifically, the procedure becomes equivalent to making measurements of the state in bases drawn randomly from the Haar ensemble. This equivalence is made precise later through our introduction of a construction called the "tomographic ensemble": a probability distribution of wave functions that describes the overall measurement process [Eqs. (1) and (2)]. For sufficiently scrambling U, integer moments of this ensemble agree closely with the Haar ensemble, i.e., an approximate quantum state design (QSD) is formed [27,28]. Consequently, properties of the system density matrix can be reconstructed through appropriate postprocessing of the measurement outcomes. This can be achieved with moderate resources, while allowing low errors in observables' estimates ($\lesssim 1\%$).

The emergence of QSDs from a single global unitary (as opposed to random sequences of local gates [29]) can be related to the recently introduced concept of "deep thermalization," where QSDs appear in the projected ensemble of many-body quantum states [30–35]. By adapting analytical arguments developed in that context, we rigorously establish the existence of QSDs for particular representative cases. We supplement this with numerical evidence for generic choices of U, which allows us to benchmark the full tomography procedure and understand the effect of symmetries.

Protocol.—Our aim is to measure properties of some state of interest ρ_S, which is prepared at the beginning of each run of the experiment in some register of qubits S labeled $i = 1, \ldots, N_S$ (similar considerations also apply beyond qubit systems). We assume that projective measurements of all qubits can be made in some computational basis $\{|m\rangle\}$ and fix a basis where $|m\rangle$ are Z_i diagonal, where (X_i, Y_i, Z_i) are Pauli operators.

Projective measurements in the fixed basis $\{|m\rangle\}$ give us access to expectation values of diagonal observables, e.g., $\langle Z_i Z_j\rangle$. To learn off-diagonal observables, one can apply an appropriate unitary to the system qubits before measurement. For instance, if we rotate every qubit by $e^{-i\pi\sum_i Y_i/4}$, then observables such as $\langle X_i X_j\rangle$ can be learned. However, in analog quantum simulators, where we have only global control, observables such as $\langle X_i Y_j\rangle$ cannot be measured in this way, since different unitaries would have to be applied to qubits i and j separately—an operation that we assume to be unavailable. (See also Refs. [36,37] for a discussion.)

To overcome this limitation, we propose a protocol that employs a set of ancilla qubits A initialized in some predetermined state, which for convenience we assume to be a pure product state $|0^{\otimes N_A}\rangle$ (this assumption is not strictly necessary). The system and ancilla qubits are jointly evolved using some fixed global unitary U, which is generated by a (possibly time-dependent) Hamiltonian that can be readily simulated on the platform in question. We refer to all such unitaries as "native." Finally, all qubits are measured in the computational basis $\{|m\rangle\}$. This is repeated M times, resulting in a collection of M bit strings $m^{(r)} = m_1^{(r)} m_2^{(r)} \ldots m_N^{(r)}$, for $r = 1, \ldots, M$, each of length $N = N_S + N_A$. This protocol is illustrated in Fig. 1(a).

Our claim is that if a native unitary U is sufficiently entangling (in a sense soon to be made precise), then *any* observable can be inferred from the distribution of measurement outcomes m, and—crucially—that the number of experimental repetitions M and amount of classical computation required to estimate most observables of interest can be bounded, in the same spirit as classical shadow tomography [23]. Remarkably, this is possible using just a single, fixed choice of U each run (although later we will show that quantitative performance improvements can be obtained by sampling U from an ensemble of native unitaries in each run).

The above claim can be more precisely specified using the formalism of positive operator-valued measures (POVMs). Any measurement scheme on a state ρ_S with possible outcomes $\{m\}$ can be described by a set of positive Hermitian operators $F_m \geq 0$, known as a POVM, chosen such that the probability of obtaining outcome m is $\mathbb{P}(m|\rho_S) = \mathrm{Tr}[F_m \rho_S]$. Conservation of probability implies $\sum_m F_m = \mathbb{I}_S$. For our protocol, the POVM operators are given by

$$F_m = (\mathbb{I}_S \otimes \langle 0|_A) U^\dagger |m\rangle \langle m| U (\mathbb{I}_S \otimes |0\rangle_A). \tag{1}$$

We have assumed that the initial state of the ancilla is pure and the evolution is unitary. Therefore, F_m are proportional to rank-1 projectors $F_m = d_S q_m |\phi_m\rangle\langle\phi_m|$, where $q_m = \mathbb{P}(m|\mathbb{I}_S/d_S) = \mathrm{Tr}[F_m]/d_S$ are the outcome probabilities for the maximally mixed state $\mathbb{I}_S/d_S$, $|\phi_m\rangle$ are normalized wave functions, and $d_S = 2^{N_S}$. Since $q_m \geq 0$ and

$\sum_m q_m = 1$, formally we can define a probability distribution over pure states on S, where the normalized wave function $|\phi_m\rangle$ occurs with probability q_m. We refer to this distribution, which contains complete information about the POVM, as the tomographic ensemble.

We argue that for generic choices of U generated by local interactions without conservation laws, the tomographic ensemble exhibits a useful universal property, namely, that it forms an approximate QSD [27,28]. This means that, for small enough integers k, the kth moments of the ensemble

$$\mathcal{E}^{(k)} = \sum_m q_m (|\phi_m\rangle\langle\phi_m|)^{\otimes k} \equiv \frac{1}{d_S} \sum_m \mathrm{Tr}[F_m] \tilde{F}_m^{\otimes k} \quad (2)$$

agree with the kth moments of the Haar ensemble $\mathcal{E}^{(k)}_{\mathrm{Haar}}$ up to some small error. (Here, $\tilde{F}_m = F_m/\mathrm{Tr}[F_m]$ are unit-trace positive operators.) Intuitively, closeness of a given ensemble to the Haar measure [as quantified by the moments (2)] implies that the probability distribution covers the space of states approximately uniformly. If the dynamics U respects some symmetry, then $\mathcal{E}^{(k)}$ will instead tend toward an alternative ensemble, where within each symmetry charge sector a k-design is formed; we discuss this case in the Supplemental Material [38].

We first provide evidence justifying the above claim and then describe how this property can be leveraged to perform shadow tomography of target states ρ_S.

Emergent quantum state designs.—The formation of QSDs in the tomographic ensemble is reminiscent of the concept of deep thermalization. In the latter, a bipartite wave function $|\Psi^{SA}\rangle$ is prepared by applying a unitary U to a product state, and the qubits on A are measured projectively, therefore producing an ensemble of states on S. Deep thermalization is achieved if this ensemble reproduces the Haar ensemble up to the kth moment for some $k > 1$. While deep thermalization and QSDs in the tomographic ensemble are distinct concepts, they bear many similarities. This connection is particularly fruitful since there are examples [30,32,33] where the emergence of deep thermalization can be rigorously established. We have adapted these proofs to show that the tomographic ensemble forms an (approximate) QSD when U is drawn from the Haar ensemble or is a dual-unitary circuit evolved for a time $t \geq N_S$ [38].

These two cases are illustrative, albeit contrived, examples where rigorous results can be obtained. For more practical purposes, we employ numerical simulations to illustrate that the same occurs for generic unitaries that arise in analog quantum simulators. As figure of merit, following Ref. [30], we use the trace distance $\Delta^{(k)} := \frac{1}{2}\|\mathcal{E}^{(k)} - \mathcal{E}^{(k)}_{\mathrm{Haar}}\|_1$, which quantifies how far the tomographic ensemble is from being a k-design ($\|C\|_1 = \mathrm{Tr}[\sqrt{C^\dagger C}]$ is the trace norm). We study dynamics under the Hamiltonian $H(t) = \sum_j X_j X_{j+1} + h^x(t) X_j + h^y(t) Y_j + h^z(t) Z_j$, which approximates the native dynamics of Rydberg atom quantum simulators [17] and in certain parameter regimes is known to exhibit fast scrambling of information [53–56]. Furthermore, when the fields $h^{x,y,z}(t)$ are time dependent, there are no conserved quantities, and we find that this encourages a rapid approach to k-design. We find that Floquet evolution works well, with $h_x = 0.8$, $h_z = 0$, and $h_y(t)$ toggling periodically between 0.9 for $t \in [n, n + 0.5)$ and 1.8 for $t \in [n - 0.5, n)$, with $n \in \mathbb{Z}$. In the following, the system qubits are located at the center of a chain with open boundary conditions.

The behavior of the trace distance for $k = 2$ as a function of time is shown in Fig. 2 for various different N_A. We see approximately exponential decay with time, until a plateau is reached. The value of this plateau is close to the average trace distance that one obtains by replacing $|\phi_m\rangle$ with 2^N independently sampled Haar-random wave functions, indicating that the states making up the tomographic ensemble are effectively quasirandom. Accordingly, the plateau trace distance scales as $\sim 1/\sqrt{2^N}$. This behavior is qualitatively similar behavior to that seen in the projected ensemble of wave functions generated from non-energy-conserving dynamics [34].

Extracting properties of the state.—Having established that the POVMs generated from our protocol generically form QSDs, we now describe how this property can be leveraged to efficiently learn properties of ρ_S. While 2-designs are known to be optimal for full reconstruction of the system density matrix [57] or process tomography [58,59], here we describe an explicitly shadow tomographic scheme for extracting information about ρ_S, which in comparison keeps the sample complexity and classical computational cost bounded [21–23].

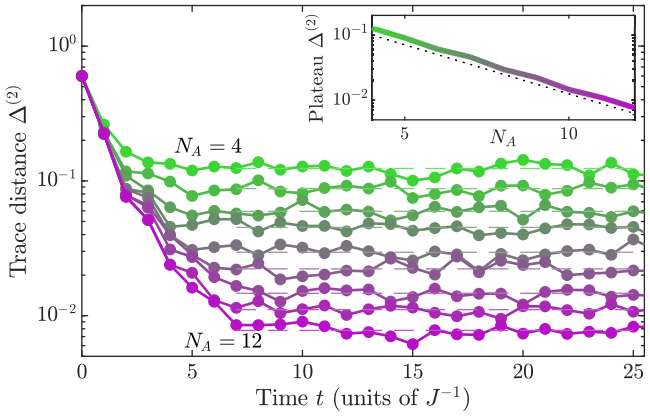

FIG. 2. Trace distance between the moments of the tomographic ensemble (2) and the Haar ensemble for a unitary $U = U_F^t$, with $N_S = 2$ and N_A increasing from 4 (green) to 12 (purple). The Floquet unitary is $U_F = e^{-iH_2/2} e^{-iH_1/2}$, with Hamiltonians $H_{1,2}$ describing the tilted-field Ising model with different field values; see main text. Inset: the plateau values (dashed lines in main plot) scale approximately as $\sim 2^{-N/2}$ (dotted line).

For a fixed unitary U, the distribution of measurement outcomes p_m depends on the state ρ_S through the POVM operators (1). It will be useful to treat operators on S as vectors, denoted using double angled brackets $|\mathcal{O}\rangle\!\rangle$ and equipped with the inner product $\langle\!\langle \mathcal{O}|\mathcal{O}'\rangle\!\rangle = \mathrm{Tr}[\mathcal{O}^\dagger \mathcal{O}']$. Similarly, the outcome distribution can be written as a 2^N-dimensional vector $|p) = \sum_m p_m |m)$, where $|m)$ is an orthonormal basis for $\mathbb{R}^{2^N}$, i.e., $(m|m') = \delta_{m,m'}$. One can then define a completely positive linear map, which we call the POVM channel,

$$\mathcal{F} = \sum_m |m)\langle\!\langle F_m|. \tag{3}$$

The observed experimental outcomes $\{\hat{m}^{(r)}\}$ $(r = 1, \ldots, M)$ are evidently distributed according to the probability vector $|p) = \mathcal{F}|\rho_S\rangle\!\rangle$.

The inverse problem of learning properties of ρ_S from experimental data $\{\hat{m}^{(r)}\}$ can be solved by finding a map $\mathcal{G}$ satisfying $\mathcal{G}\mathcal{F} = \mathrm{id}$, where id is the identity superoperator. This allows us to construct an unbiased estimator $\hat{\theta}_\mathcal{O}$ for any expectation value $\langle\mathcal{O}\rangle = \mathrm{Tr}[\mathcal{O}\rho_S]$ according to $\hat{\theta}_\mathcal{O} = M^{-1}\sum_{r=1}^M \langle\!\langle \mathcal{O}|\mathcal{G}|\hat{m}^{(r)})$. In the spirit of shadow tomography [21–23], this estimator can be computed without needing to reconstruct the full density matrix ρ_S, which would be sample inefficient. Such an inverse $\mathcal{G}$ only exists when $\mathcal{F}$ has full row rank: a condition known as "informational completeness," which is guaranteed when the tomographic ensemble forms a 2-design [38].

While $\mathcal{G}$ is nonunique, in general, to minimize sample complexity we choose the inverse that minimizes the (average-case) variance $\mathrm{Var}\hat{\theta}_\mathcal{O} = \mathbb{E}_{\hat{m}}[\hat{\theta}_\mathcal{O}^2] - \mathbb{E}_{\hat{m}}[\hat{\theta}_\mathcal{O}]^2$, namely [38],

$$\mathcal{G}^* := \mathcal{M}^{-1}\tilde{\mathcal{F}}^\dagger, \quad \text{where } \mathcal{M} := \sum_m \mathrm{Tr}[F_m]|\tilde{F}_m\rangle\!\rangle\langle\!\langle \tilde{F}_m|, \tag{4}$$

where we defined the normalized channel $\tilde{\mathcal{F}} = \sum_m |m)\langle\!\langle \tilde{F}_m|$. The map $\mathcal{M}$ is a superoperator mapping the space of operators on S to itself. It has full rank whenever $\mathcal{F}$ is informationally complete and therefore has a unique inverse.

At this point, we notice that the superoperator $\mathcal{M}$ is equivalent to the second moment of the tomographic ensemble $\mathcal{E}^{(2)}$, Eq. (2) [38]. Having established that QSDs generically appear in our protocol, we can replace $\mathcal{M}$ with its universal 2-design form $\mathcal{M} = (\mathrm{id} + |\mathbb{I}\rangle\!\rangle\langle\!\langle \mathbb{I}|)/(d_S + 1)$, which has an inverse

$$\mathcal{M}^{-1}[\mathcal{O}] = (d_S + 1)\mathcal{O} - \mathrm{Tr}[\mathcal{O}]\mathbb{I}. \tag{5}$$

Using the fact that a 2-design is formed, we circumvent having to explicitly compute $\mathcal{M}^{-1}$, which keeps the classical computational cost bounded.

Using $\langle\!\langle \mathcal{O}|\mathcal{G}^*|m) = \langle\!\langle \mathcal{M}^{-1}[O]|\tilde{F}_m\rangle\!\rangle$, we can express the variance of the estimator (4) as

$$\mathrm{Var}\,\hat{\theta}_\mathcal{O} = \frac{1}{M}\mathrm{Tr}\left[(\rho_S \otimes \mathcal{M}^{-1}[\mathcal{O}]^{\otimes 2})\left(\sum_m F_m \otimes \tilde{F}_m^{\otimes 2}\right)\right]. \tag{6}$$

The factor in rounded brackets we identify as the third moment, $\mathcal{E}^{(3)}$ in Eq. (2). Therefore, if the tomographic ensemble forms a 3-design, as we expect for generic unitaries U, then the variance (6) will be the same as for any other 3-design POVM. One such POVM arises in classical shadow tomography with random global Clifford unitaries [38]. Therefore, we can conclude that our scheme can be used to estimate expectation values of ρ_S using the same number of repetitions M as one would need for ordinary classical shadow tomography. The dependence of the variance on the observable in question is well characterized in Ref. [23]: observables with bounded norm $\|\mathcal{O}\|_2 = \sqrt{\mathrm{Tr}[\mathcal{O}^\dagger\mathcal{O}]}$ can be efficiently estimated for any system size N_S. The procedure can be generalized in the same way as classical shadow tomography to estimate nonlinear observables, e.g., Rényi entropies [38].

To summarize, we have shown how the formation of QSDs in the tomographic ensemble facilitates efficient reconstruction of expectation values, since the map $\mathcal{M}^{-1}$ appearing in Eq. (4) can be replaced by its universal form (5). The deviation from the 2-design will govern the systematic error, since $|\mathbb{E}_m\hat{\theta}_\mathcal{O} - \langle\mathcal{O}\rangle| \leq (d_S + 1)\Delta^{(2)}\|O\|_\infty$, where $\Delta^{(2)}$ is the trace distance, while the $k = 3$ moments $\mathcal{E}^{(3)}$ determine the variance via (6). It is evidently favorable to have the tomographic ensemble as close to a 2- and 3-design as possible, which occurs for generic chaotic evolution as we saw above.

Benchmarking the protocol.—We now provide numerical simulations of our full protocol, including the joint evolution of the system and ancillas, sampling of measurement outcomes, and reconstruction of observables. We test our measurement scheme on a family of 2-qubit target states $\rho_S(\alpha) = \alpha|\mathrm{EPR}\rangle\langle\mathrm{EPR}| + (1-\alpha)[|00\rangle\langle 00| + |11\rangle\langle 11|]/2$, where $|\mathrm{EPR}\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$ is an Einstein-Podolsky-Rosen state. The coherence parameter $\alpha \in [0, 1]$ allows us to interpolate between fully dephased ($\alpha = 0$) and pure ($\alpha = 1$) EPR pairs. For the purpose of demonstration, the observables we choose to reconstruct are the fidelity with the EPR state $\mathrm{Tr}[\rho_S|\mathrm{EPR}\rangle\langle\mathrm{EPR}|]$ and the purity $\mathrm{Tr}[\rho_S^2]$.

In one set of simulations, we generate U from Floquet evolution using the tilted-field Ising model as a generating Hamiltonian, as before. In a second set, we also add some randomness to U—that is, for each repetition r we generate a distinct $U^{(r)}$ by selecting random magnetic fields. Then, $U^{(r)}$ is used in the joint system-ancilla evolution and in the construction of estimators. This helps to bring the tomographic ensemble closer to a 2-design, therefore further reducing systematic errors [38]. To construct random

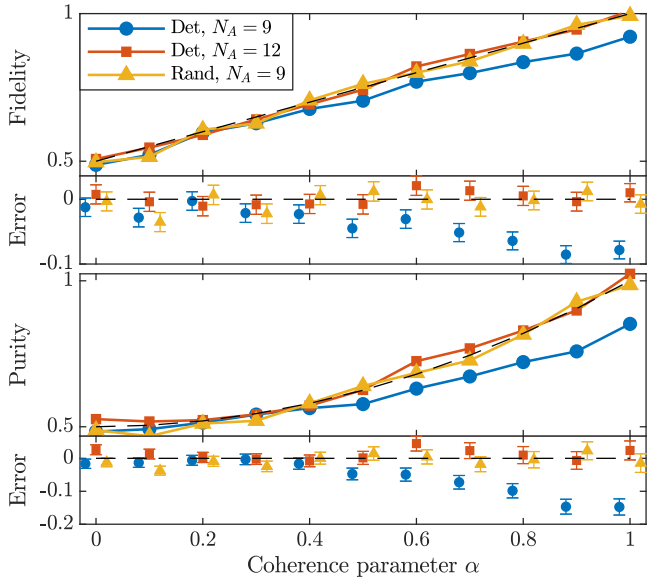

FIG. 3. Estimations of the fidelity $\mathrm{Tr}[\rho_S|\mathrm{EPR}\rangle\langle\mathrm{EPR}|]$ (top panels) and purity $\mathrm{Tr}[\rho_S^2]$ (bottom panels) using the deterministic protocol (Det; fixed U), and the semi-randomized protocol (Rand), using the state $|\mathrm{EPR}\rangle$ after dephasing with strength $1-\alpha$ as a target, see main text. The total evolution time is $t = 10$, and $M = 5 \times 10^3$ repetitions are used for all data points. Errors relative to the true data are shown, with data artificially shifted horizontally for readability.

unitaries $U^{(r)}$, for each time interval of length $\tau = 1$, we sample each field component $h^{x,y,z}$ independently from a normal distribution with zero mean and standard deviation $\sqrt{2}$.

Figure 3 shows estimations of the fidelity and purity for various different α and tomography schemes, using $M = 5 \times 10^3$ repetitions each and evolving for a total time $t = 10$. We see closer agreement with the true fidelity as N_A is increased and when randomness is introduced.

Classical computations.—As in classical shadow tomography, the estimation of expectation values from experimental data requires a certain amount of classical postprocessing, the complexity of which we wish to bound. Specifically, when an outcome m is observed we must evaluate $\langle\langle\mathcal{O}|\mathcal{G}_0|m\rangle$, which requires computation of the backward time evolution $U^\dagger|m\rangle$.

When the number of system qubits N_S is $O(1)$, the evolution time required to obtain an approximate QSD is also $O(1)$, and hence efficient matrix product state techniques can be used even for large N_A. For tomography of many-body states, the present strategy must be modified, since the time of evolution required to reach a QSD grows with N_S. Instead of evolving all system qubits with a single collection of ancillas, one can instead block the system into n groups of $N_S/n = O(1)$ qubits and evolve each block jointly with a separate collection of ancillas A_j under a unitary U_j, where $j = 1, \ldots, n$. This scheme, illustrated in Fig. 1(b), yields POVM operators $F_{m_1,\ldots,m_n} = \bigotimes_{j=1}^n F_{m_j}^{(j)}$, where each $F_{m_j}^{(j)}$ is of the form (1). The tomographic ensemble for each separate block reaches an approximate 3-design in a $O(1)$ time, allowing $F_{m_1,\ldots,m_n}$ to be evaluated efficiently using matrix product methods as before. The trade-off is that $\mathcal{M}^{-1}$ must be replaced by an n-fold tensor product of (5), and this will affect how the estimator variance (6) depends on the observable $\mathcal{O}$. By analogy to shadow tomography with random local Pauli measurements [23], observables with support on a small number of blocks will still be accessible using a reasonable number of repetitions M, regardless of how big S is. (We demonstrate this rigorously in the Supplemental Material [38].)

Note that one could, in principle, compute the map $\mathcal{M}^{-1}$ without using the universal 2-design form (5), which would eliminate any systematic error in estimation. However, this is only feasible for a small number of ancillas N_A, since $2^{N_S+N_A}$ separate terms must be summed to construct $\mathcal{M}$.

Note added.—Recently, we became aware of a complementary study, which appeared in the same arXiv posting, where a similar measurement scheme is presented [37]. The protocol introduced in that work follows the same steps as ours, where the state is first entangled with ancillas, before measurements in the computational basis are made, data from which are postprocessed classically to infer properties of the state. In contrast to our proposal, no assumption is made about the formation of a QSD; instead, the inverse map $\mathcal{M}^{-1}$ needs to be explicitly computed.

M. M. thanks Shivaji Sondhi for helpful discussions. We are especially grateful to Sounak Biswas for insight throughout the completion of this work. We acknowledge support from UK Engineering and Physical Sciences Research Council Grant No. EP/S020527/1 and the European Research Council under the European Union Horizon 2020 Research and Innovation Programme, Grant Agreement No. 804213-TMCS (M. F.).

[1] Y. Manin, *Computable and Uncomputable* (Sovetskoye Radio Press, Moscow, 1980) in Russian.
[2] R. P. Feynman, Simulating physics with computers, Int. J. Theor. Phys. **21**, 467 (1982).
[3] S. Lloyd, Universal quantum simulators, Science **273**, 1073 (1996).
[4] J. Preskill, Quantum computing in the NISQ era and beyond, Quantum **2**, 79 (2018).
[5] I. Bloch, J. Dalibard, and S. Nascimbene, Quantum simulations with ultracold quantum gases, Nat. Phys. **8**, 267 (2012).
[6] C. Gross and I. Bloch, Quantum simulations with ultracold atoms in optical lattices, Science **357**, 995 (2017).

[7] H. Weimer, M. Müller, I. Lesanovsky, P. Zoller, and H. P. Büchler, A Rydberg quantum simulator, Nat. Phys. **6**, 382 (2010).

[8] D. Barredo, S. de Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys, An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays, Science **354**, 1021 (2016).

[9] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner, and M. D. Lukin, Atom-by-atom assembly of defect-free one-dimensional cold atom arrays, Science **354**, 1024 (2016).

[10] A. Browaeys and T. Lahaye, Many-body physics with individually controlled Rydberg atoms, Nat. Phys. **16**, 132 (2020).

[11] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms, Nature (London) **415**, 39 (2002).

[12] B. Paredes, A. Widera, V. Murg, O. Mandel, S. Fölling, I. Cirac, G. V. Shlyapnikov, T. W. Hänsch, and I. Bloch, Tonks–Girardeau gas of ultracold atoms in an optical lattice, Nature (London) **429**, 277 (2004).

[13] M. Aidelsburger, M. Atala, M. Lohse, J. T. Barreiro, B. Paredes, and I. Bloch, Realization of the Hofstadter Hamiltonian with Ultracold Atoms in Optical Lattices, Phys. Rev. Lett. **111**, 185301 (2013).

[14] M. Schreiber, S. S. Hodgman, P. Bordia, H. P. Lüschen, M. H. Fischer, R. Vosk, E. Altman, U. Schneider, and I. Bloch, Observation of many-body localization of interacting fermions in a quasirandom optical lattice, Science **349**, 842 (2015).

[15] J. yoon Choi, S. Hild, J. Zeiher, P. Schauß, A. Rubio-Abadal, T. Yefsah, V. Khemani, D. A. Huse, I. Bloch, and C. Gross, Exploring the many-body localization transition in two dimensions, Science **352**, 1547 (2016).

[16] J. Smith, A. Lee, P. Richerme, B. Neyenhuis, P. W. Hess, P. Hauke, M. Heyl, D. A. Huse, and C. Monroe, Many-body localization in a quantum simulator with programmable random disorder, Nat. Phys. **12**, 907 (2016).

[17] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner *et al.*, Probing many-body dynamics on a 51-atom quantum simulator, Nature (London) **551**, 579 (2017).

[18] S. de Léséleuc, V. Lienhard, P. Scholl, D. Barredo, S. Weber, N. Lang, H. P. Büchler, T. Lahaye, and A. Browaeys, Observation of a symmetry-protected topological phase of interacting bosons with Rydberg atoms, Science **365**, 775 (2019).

[19] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho *et al.*, Quantum phases of matter on a 256-atom programmable quantum simulator, Nature (London) **595**, 227 (2021).

[20] P. N. Jepsen, Y. K. Lee, H. Lin, I. Dimitrova, Y. Margalit, W. W. Ho, and W. Ketterle, Long-lived phantom helix states in Heisenberg quantum magnets, Nat. Phys. **18**, 899 (2022).

[21] S. Aaronson, Shadow tomography of quantum states, in *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018* (Association for Computing Machinery, New York, 2018), pp. 325–338.

[22] S. Aaronson and G. N. Rothblum, Gentle measurement of quantum states and differential privacy, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019* (Association for Computing Machinery, New York, 2019), pp. 322–333.

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

[24] A. Elben, B. Vermersch, M. Dalmonte, J. I. Cirac, and P. Zoller, Rényi Entropies from Random Quenches in Atomic Hubbard and Spin Models, Phys. Rev. Lett. **120**, 050406 (2018).

[25] T. Brydges, A. Elben, P. Jurcevic, B. Vermersch, C. Maier, B. P. Lanyon, P. Zoller, R. Blatt, and C. F. Roos, Probing Rényi entanglement entropy via randomized measurements, Science **364**, 260 (2019).

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

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

[28] A. Ambainis and J. Emerson, Quantum *t*-designs: *t*-wise independence in the quantum world, arXiv:quant-ph/0701126.

[29] A. W. Harrow and R. A. Low, Random quantum circuits are approximate 2-designs, Commun. Math. Phys. **291**, 257 (2009).

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

[31] J. Choi, A. L. Shaw, I. S. Madjarov, X. Xie, R. Finkelstein, J. P. Covey, J. S. Cotler, D. K. Mark, H.-Y. Huang, A. Kale, H. Pichler, F. G. S. L. Brandão, S. Choi, and M. Endres, Emergent quantum randomness and benchmarking from Hamiltonian many-body dynamics, Nature (London) **613**, 468 (2023).

[32] W. W. Ho and S. Choi, Exact Emergent Quantum State Designs from Quantum Chaotic Dynamics, Phys. Rev. Lett. **128**, 060601 (2022).

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

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

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

[36] K. V. Kirk, J. Cotler, H.-Y. Huang, and M. D. Lukin, Hardware-efficient learning of quantum many-body states, arXiv:2212.06084.

[37] M. C. Tran, D. K. Mark, W.-W. Ho, and S. Choi, Measuring Arbitrary Physical Properties in Analog Quantum Simulation, Phys. Rev. X **13**, 011049 (2023).

[38] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.131.160601 for proofs of the existence of *k*-designs in the tomographic ensemble

for both Haar-random unitaries and dual-unitary circuits, a more in-depth discussion of the semirandomized protocol, a treatment of symmetry-respecting dynamics, and details on classical postprocessing, which includes Refs. [39–52].

[39] S. Gammelmark, B. Julsgaard, and K. Mølmer, Past Quantum States of a Monitored System, Phys. Rev. Lett. **111**, 160401 (2013).

[40] M. Ledoux, *The Concentration of Measure Phenomenon*, Mathematical Surveys and Monographs (American Mathematical Society, Providence, 2001).

[41] R. A. Low, Large deviation bounds for k-designs, Proc. R. Soc. A **465**, 3289 (2009).

[42] B. Bertini, P. Kos, and T. Prosen, Exact Spectral Form Factor in a Minimal Model of Many-Body Quantum Chaos, Phys. Rev. Lett. **121**, 264101 (2018).

[43] B. Bertini, P. Kos, and T. Prosen, Entanglement Spreading in a Minimal Model of Maximal Many-Body Quantum Chaos, Phys. Rev. X **9**, 021033 (2019).

[44] S. Gopalakrishnan and A. Lamacraft, Unitary circuits of finite depth and infinite width from quantum channels, Phys. Rev. B **100**, 064309 (2019).

[45] B. Bertini, P. Kos, and T. Prosen, Exact Correlation Functions for Dual-Unitary Lattice Models in $1+1$ Dimensions, Phys. Rev. Lett. **123**, 210601 (2019).

[46] B. Buča and T. Prosen, A note on symmetry reductions of the Lindblad equation: Transport in constrained open spin chains, New J. Phys. **14**, 073007 (2012).

[47] Z. Webb, The clifford group forms a unitary 3-design, Quantum Inf. Comput. **16**, 1379 (2016).

[48] G. Vidal and R. F. Werner, Computable measure of entanglement, Phys. Rev. A **65**, 032314 (2002).

[49] M. B. Plenio, Logarithmic Negativity: A Full Entanglement Monotone That is Not Convex, Phys. Rev. Lett. **95**, 090503 (2005).

[50] A. Elben, R. Kueng, H.-Y. R. Huang, R. van Bijnen, C. Kokail, M. Dalmonte, P. Calabrese, B. Kraus, J. Preskill, P. Zoller, and B. Vermersch, Mixed-State Entanglement from Local Randomized Measurements, Phys. Rev. Lett. **125**, 200501 (2020).

[51] T. S. Ferguson, U-statistics, lecture notes for statistics 200B, *UCLA* (2003), accessed August 2022, https://www.math.ucla.edu/~tom/Stat200C/Ustat.pdf.

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

[53] M. C. Bañuls, J. I. Cirac, and M. B. Hastings, Strong and Weak Thermalization of Infinite Nonintegrable Quantum Systems, Phys. Rev. Lett. **106**, 050405 (2011).

[54] H. Kim and D. A. Huse, Ballistic Spreading of Entanglement in a Diffusive Nonintegrable System, Phys. Rev. Lett. **111**, 127205 (2013).

[55] H. Kim, T. N. Ikeda, and D. A. Huse, Testing whether all eigenstates obey the eigenstate thermalization hypothesis, Phys. Rev. E **90**, 052105 (2014).

[56] P. Hosur, X.-L. Qi, D. A. Roberts, and B. Yoshida, Chaos in quantum channels, J. High Energy Phys. 02 (2016) 004.

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

[58] J. Emerson, R. Alicki, and K. Życzkowski, Scalable noise estimation with random unitary operators, J. Opt. B **7**, S347 (2005).

[59] C. Dankert, R. Cleve, J. Emerson, and E. Livine, Exact and approximate unitary 2-designs and their application to fidelity estimation, Phys. Rev. A **80**, 012304 (2009).