Optimal Conversion from Classical to Quantum Randomness via Quantum Chaos

Wai-Keong Mok,[1] Tobias Haug,[2] Adam L. Shaw,[1,*] Manuel Endres,[1] and John Preskill[1,3]
[1]Institute for Quantum Information and Matter, California Institute of Technology, Pasadena, California 91125, USA
[2]Quantum Research Centre, Technology Innovation Institute, Abu Dhabi, United Arab Emirates
[3]AWS Center for Quantum Computing, Pasadena, California 91125, USA

(Received 18 October 2024; revised 29 January 2025; accepted 8 April 2025; published 6 May 2025)

Quantum many-body systems provide a unique platform for exploring the rich interplay between chaos, randomness, and complexity. In a recently proposed paradigm known as deep thermalization, random quantum states of system A are generated by performing projective measurements on system B following chaotic Hamiltonian evolution acting jointly on AB. In this scheme, the randomness of the projected state ensemble arises from the intrinsic randomness of the outcomes when B is measured. Here, we propose a modified scheme in which classical randomness injected during the protocol is converted by quantum chaos into quantum randomness of the resulting state ensemble. We show that for generic chaotic systems this conversion is optimal in that each bit of injected classical entropy generates as much additional quantum randomness as adding an extra qubit to B. This significantly enhances the randomness of the projected ensemble without imposing additional demands on the quantum hardware. Our scheme can be easily implemented on typical analog quantum simulators, providing a more scalable route for generating quantum randomness valuable for many applications. In particular, we demonstrate that the accuracy of a shadow tomography protocol can be substantially improved.

DOI: 10.1103/PhysRevLett.134.180403

Introduction—Preparing ensembles of random quantum states is an increasingly important task in quantum information science. Quantum randomness is theoretically interesting for its connections to quantum chaos [1–3] and quantum descriptions of black hole dynamics [4–6], and is also practically relevant for a wide range of applications such as randomized benchmarking [7,8], quantum communication [9], phase retrieval [10], shadow tomography [11,12], cryptography [13,14], and large-scale device benchmarking [15–18]. However, in general it is difficult to generate truly random quantum states—known as Haar random states—which have exponential complexity. Further, for these applications it is typically desirable to create random states on many qubits, but this task is stymied by limitations on the fidelity of near-term quantum processors.

In lieu of creating such truly random states, many applications can already be accomplished through only low-order approximations to the Haar random ensemble, known as k-designs [8,19], which are statistically indistinguishable from Haar random states for any observable involving up to k copies of the state. Thus, it is useful to characterize the randomness of a quantum state ensemble by quantifying its distance from a k-design for $k \geq 2$. It is known that approximate k-designs can be generated efficiently by random unitary circuits (RUCs) [20–23]. Effectively, such RUCs convert classical randomness (the arrangement of the circuit) into quantum randomness (the generated k-design), albeit while requiring a high degree of spatiotemporal control and low noise.

Recently, a new paradigm has emerged for realizing random states without requiring the high level of control of RUCs, but instead relying on projective measurements following a fixed chaotic Hamiltonian evolution [24,25] in a process known as deep thermalization [24–38]. In this approach, one evolves a many-qubit quantum system under a chaotic Hamiltonian (with spectral statistics described by random matrix theory), and then measures a subset of the qubits (the bath). For a wide range of physical systems, the resulting state ensemble of the unmeasured qubits (the projected ensemble) approximates a k-design [24,25]. This method for generating random states can be used for practical applications such as benchmarking analog quantum simulators [25]. Since the dynamics is fixed, unlike RUCs, the randomness stems from the intrinsic uncertainty of quantum measurements. While this approach yields exact Haar random states in the limit of infinitely many measured qubits, the convergence to successively higher order k-designs is limited by the size of the projectively measured bath, which poses a practical bottleneck for near-term quantum hardware with limited qubit number.

Here, we propose and numerically implement new strategies for circumventing this constraint via injection of classical randomness into quantum chaotic dynamics. This approach allows us to effectively generate higher order

*Present address: Department of Physics, Stanford University, Stanford, California, USA.

k-designs for a fixed system size by supplying classical entropy rather than additional bath qubits. We inject classical entropy via either randomizing the initial basis states prepared for the Hamiltonian evolution, or injecting time-independent disorder into the evolution itself. This approach allows us to sample random quantum states with significantly fewer quantum resources than would be required using RUCs or the previously proposed version of deep thermalization. In fact, we find that each bit of classical entropy can produce as much additional quantum randomness as adding an extra qubit to the bath, signifying an optimal conversion of classical to quantum randomness. The injected classical randomness can effectively double the size of the bath; hence using the same number of physical qubits we can generate state ensembles that are far closer to the Haar ensemble compared to the projected ensemble protocol proposed in [24].

Our procedures are experimentally simple to execute on modern quantum platforms, including most analog quantum simulators, and can thus enable the creation of more complex random states on near-term quantum devices. Our approach can improve many applications of quantum randomness [7–18]; we describe one in particular by demonstrating how injecting classical randomness can enhance the performance of shadow tomography [11,12].

Projected ensembles and deep thermalization—In the basic framework of the projected ensemble [24,25,39,40], a fixed initial state of N qubits, taken to be $|0\rangle^{\otimes N}$, is evolved under a fixed unitary evolution U_{AB}. The system is bipartitioned into two subsystems, delineated as A and B, of size N_A and N_B qubits, respectively. Qubits in B are then measured in the computational basis producing a bitstring $z \in \{0,1\}^{N_B}$ with probability $p(z)$, which induces a pure quantum state in A conditioned on the measurement outcome. The set of these postmeasurement states, together with the probabilities $p(z)$, collectively encode the entire state and is known as the projected ensemble.

Our scheme exploits this basic framework, but improves the convergence of the projected ensemble to k-designs via the injection of classical randomness. In the simplest case [see Fig. 1(a)], this is accomplished by preparing not a fixed initial state for the evolution, but a randomized one. Concretely, the initial state is randomly sampled from the ensemble $\mathcal{E}_{\text{init}} = \{|x\rangle, q(x)\}_{x\in\{0,1\}^N}$ of computational basis states with probability $q(x)$, and then evolved under the same fixed U_{AB}. The bath is then projectively measured, now yielding bitstrings z labeled by the choice of initial state with probabilities $p_x(z) = q(x)\|(I_A \otimes \langle z|_B) U_{AB}|x\rangle\|^2$. Accordingly, the postmeasurement states in A are labeled by both the bitstring measured in B as well as the initial state, $|\psi_x(z)\rangle \propto (I_A \otimes \langle z|_B)U_{AB}|x\rangle$. The set of these probabilities and (normalized) postmeasurement states then forms the *classically enhanced* projected ensemble $\mathcal{E} = \{q(x)p_x(z); |\psi_x(z)\rangle\}$.

In order to assess the degree of randomness of the projected ensemble $\mathcal{E}$, we compare its moments against that

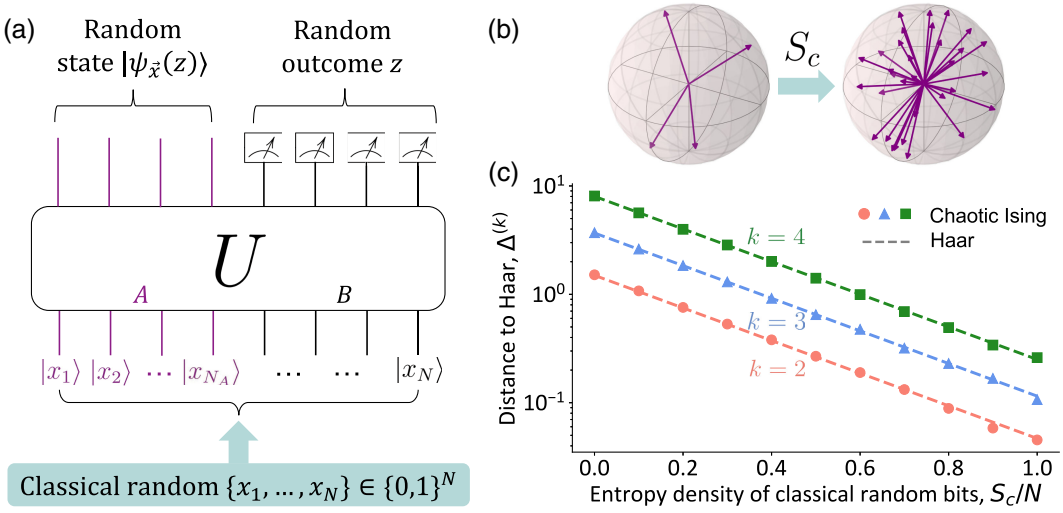

FIG. 1. Classically enhanced projected ensembles via random basis state initialization. (a) The system is initialized in a random computational basis state $\otimes_{i=1}^{N} |x_i\rangle$, where the N-bit string $x = (x_1, \ldots, x_N) \in \{0,1\}^N$ is drawn randomly from a distribution $q(x)$ with entropy S_c. The initial state evolves under a fixed unitary U, and N_B bath qubits are measured in the computational basis to yield a random outcome z. For generic chaotic dynamics, the projected states $|\psi_x(z)\rangle$ on the remaining N_A qubits form a projected ensemble $\mathcal{E}$ that approximates a k-design. (b) Classical randomness of entropy S_c increases the size of the projected ensemble $\mathcal{E}$ by a factor of up to 2^{S_c}, illustrated on the Bloch sphere. (c) Normalized Hilbert-Schmidt distance $\Delta^{(k)}$ between the kth moment of the projected ensemble $\mathcal{E}$ and the Haar ensemble, against the entropy density S_c/N of the classical distribution $q(x)$, with $N_A = 4$ and $N_B = 6$. The points are obtained numerically by evolving the initial state with the chaotic mixed-field 1D Ising Hamiltonian H_0 (4) for a time $JT = 10^3$. The dashed lines denote the analytical root-mean-square distance $\Delta^{(k)}_{\text{rms}}$ (3) when U is a Haar random unitary.

of the Haar ensemble (see Supplemental Material (SM) [41] for more details). The kth moment operator of the ensemble $\mathcal{E}$ is given by

$$\rho^{(k)} = \sum_{\substack{x\in\{0,1\}^N \\ z\in\{0,1\}^{N_B}}} q(x)p_x(z)(|\psi_x(z)\rangle\langle\psi_x(z)|)^{\otimes k}. \quad (1)$$

We then compute the normalized Hilbert-Schmidt distance to the kth moment of the Haar ensemble on A,

$$\Delta^{(k)} = \frac{\left\|\rho^{(k)} - \rho^{(k)}_{\text{Haar}}\right\|_2}{\left\|\rho^{(k)}_{\text{Haar}}\right\|_2}, \quad (2)$$

where $\|\cdot\|_2$ is the Hilbert-Schmidt norm. This has an intuitive entropic interpretation since $\Delta^{(k)}$ decreases when the quantum Rényi 2-entropy of $\rho^{(k)}$ increases. $\mathcal{E}$ forms an ϵ-approximate state k-design if $\Delta^{(k)} \le \epsilon$, which is consistent with the usual definition using the trace distance [41]. The classical randomness of the initial states $\mathcal{E}_{\text{init}}$ on the full system AB can be quantified by S_c, the Rényi 2-entropy of $q(x)$. The original protocol in [24,25] corresponds to $S_c = 0$.

Optimal conversion from classical to quantum randomness—While it is obvious that injecting classical randomness via $\mathcal{E}_{\text{init}}$ can only increase the quantum randomness of the projected ensemble [Fig. 1(b)], the interesting question is to what extent quantum randomness can be increased. For analytical tractability, we study the case where U_{AB} is a fixed unitary drawn from the Haar measure on the unitary group $\mathcal{U}(2^{N_A+N_B})$. The quantum randomness in the "typical" case is measured by the root-mean-square distance $\Delta^{(k)}_{\text{rms}} = \sqrt{\mathbb{E}_{U\sim\text{Haar}}(\Delta^{(k)})^2}$, obtained by averaging $(\Delta^{(k)})^2$ (which is a function of U_{AB}) over the Haar measure. Note that the trace distance between $\rho^{(k)}$ and $\rho^{(k)}_{\text{Haar}}$ can be upper bounded by $\mathcal{O}(\Delta^{(k)}_{\text{rms}})$, which we find is exponentially smaller in both N_A and k compared to the bound of Ref. [24] (see Theorem 2 in SM [41]). Our main analytical result is stated as follows.

Theorem 1 (Classically enhanced projected ensembles): Let $\mathcal{E}_{\text{init}} = \{q(x), |x\rangle\}$ be an set of orthonormal initial states $|x\rangle$ with probability distribution $q(x)$ and $S_c = -\log_2(\sum_x q(x)^2)$ is the Rényi 2-entropy of $q(x)$. For any $2 \le k < 2^{(N_A+N_B)/4}$, the root-mean-square distance $\Delta^{(k)}_{\text{rms}}$ between the kth moments of the projected ensemble $\mathcal{E}$ and the Haar ensemble approaches

$$(\Delta^{(k)}_{\text{rms}})^2 = \frac{1}{k!2^{S_c+N_B-kN_A}} \quad (3)$$

as N_A and $N_B \to \infty$.

The detailed proof of Theorem 1 using Weingarten calculus is provided in SM [41]. From Eq. (3), we can see that quantum randomness is maximized when $\mathcal{E}_{\text{init}}$ forms a 1-design, i.e., $S_c = N_A + N_B$ corresponding to the maximum amount of classical entropy in the initial state distribution. The expression for $k = 1$ is trivial and derived in SM [41]. This reveals explicitly the conversion of classical to quantum randomness. Moreover, since $\Delta^{(k)}_{\text{rms}}$ vanishes exponentially with $N_B + S_c$, Theorem 1 implies an (asymptotically) optimal conversion from classical to quantum randomness, where injecting S_c bits of classical entropy yields the same $\Delta^{(k)}_{\text{rms}}$ as adding S_c bath qubits without classical randomness. In other words, classical randomness and bath qubits are interconvertible resources from the perspective of projected ensembles.

This is particularly advantageous for experimental implementation since N_B is often limited on the quantum hardware due to noise and practical constraints. Additionally, up to N_B bits of classical entropy can be obtained naturally by reusing the bath measurement outcome z as an initial state of B in a subsequent run of the protocol. Equation (3) indicates that injecting classical randomness can reduce the distance between the projected ensemble and the Haar ensemble by a factor exponential in $N = N_A + N_B$ at a very modest cost for the quantum hardware. Thus injecting classical randomness significantly reduces the quantum resources needed to sample random quantum states on analog simulators.

Since S_c is at most $N_A + N_B$ bits, this gives an effective bath size of up to $N_A + 2N_B$ qubits, essentially increasing the bath size by more than a factor of 2. This allows one to generate higher-order designs with a maximal achievable k almost twice as large [41] as in the original protocol [24,25]. For example, 2-designs can be generated on A with $N_B < N_A$. In contrast, not even a 1-design can be achieved for $N_B < N_A$ if no classical randomness is injected.

We now show numerically that Theorem 1 holds even for Hamiltonian dynamics. We plot $\Delta^{(k)}$ against the classical entropy density S_c/N in Fig. 1(c) for $U_{AB} = \exp(-iH_0T)$ with quench time $JT = 10^3$, where H_0 is the mixed-field 1D Ising Hamiltonian

$$H_0 = \sum_{i=1}^{N_A+N_B} (h_xX_i + h_yY_i + JX_iX_{i+1}) \quad (4)$$

with open boundary conditions. X_i, Y_i, and Z_i are Pauli operators for qubit i. The computational basis states are eigenstates of Z_i, and thus satisfy $\langle H_0\rangle = 0$, which is necessary for deep thermalization at infinite temperature. We choose $\{h_x, h_y, J\} = \{0.8090, 0.9045, 1\}$ in the nonintegrable and chaotic regime [56,57], which was also numerically demonstrated to produce random projected ensembles [24] for $S_c = 0$. We observe that $\Delta^{(k)}$ decreases exponentially with the classical entropy S_c, in excellent agreement with our analytical formula (3) derived for Haar random dynamics. This suggests that the near-optimal

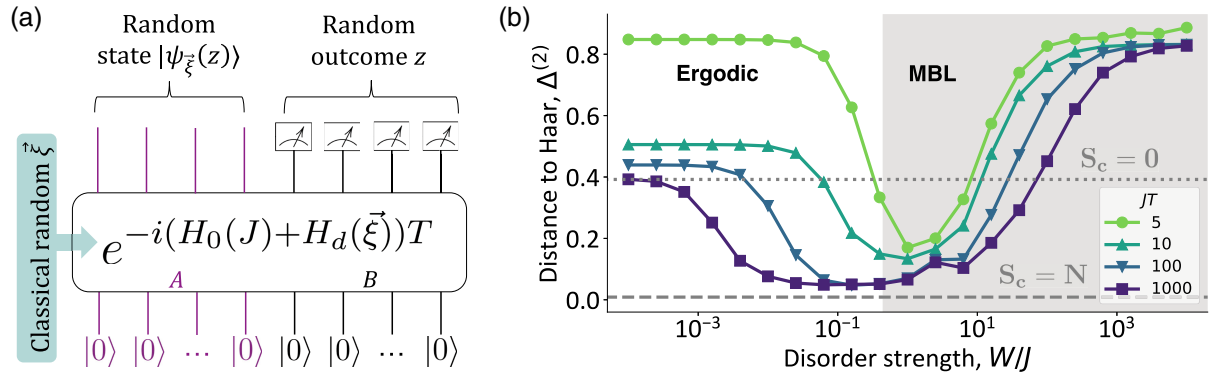

FIG. 2. (a) Classically enhanced projected ensembles via random Hamiltonian disorder. The initial state is fixed as $|0\rangle^{\otimes N}$. A spatially inhomogeneous random field $H_d(\vec{\xi})$ with fluctuations of size W is applied on N qubits on top of a fixed chaotic Hamiltonian H_0 with interaction energy J, and the state is evolved for a quench time T. (b) Normalized Hilbert-Schmidt distance $\Delta^{(2)}$ between the projected ensemble $\mathcal{E}$ and the Haar ensemble, against the disorder strength W/J. The points are obtained numerically by evolving the initial state under the chaotic (mixed-field) Ising model with disorder (5), for $N_A = 3$ and $N_B = 8$. A total of $2^{N_A+N_B}$ disorder realizations are sampled, giving a projected ensemble of size $|\mathcal{E}| = 2^{N_A+2N_B}$. Numerical results are compared against analytical benchmarks (3) of $\Delta^{(k)}_{\text{rms}}$ for Haar random unitary evolution with $S_c = 0$ and N, denoted by the dotted and dashed lines respectively. The gray shaded region marks the many-body localized (MBL) regime, with the ergodic-MBL crossover point at $W/J \approx 0.42$ [41]. For very large disorder $W/J \gg 1$, the projected ensemble converges to a 1-design, and $\Delta^{(2)}$ saturates and becomes independent of N_B.

conversion from classical to quantum randomness is a more generic feature of quantum chaos.

Classical to quantum randomness using Hamiltonian disorder—Another way to inject classical randomness is by fixing the initial state and applying a disorder Hamiltonian $H_d(\vec{\xi})$ with 2^{S_c} independent disorder realizations $\vec{\xi}$, on top of a fixed chaotic Hamiltonian H_0 [58]; see Fig. 2(a). Such a setup is motivated both by experimental considerations of analog quantum simulators that may more readily apply time-constant disordered potentials rather than localized bit-flip operations, and also for its theoretical connections to many-body localization, as we shall explore below.

As an example, we consider the Hamiltonian

$$H = H_0 + H_d(\vec{\xi}) = H_0 + \sum_{i=1}^{N_A+N_B} \xi_i X_i, \tag{5}$$

where H_0 is given by Eq. (4) and disorder $\xi_i \overset{\text{i.i.d.}}{\sim} \text{Uniform}[-W, W]$ with disorder strength W/J (i.i.d., independent identically distributed). Figure 2(b) shows the behavior of $\Delta^{(2)}$ against disorder strength W/J (similar behavior is observed for other $k > 2$), with $2^{N_A+N_B}$ disorder realizations. The analytical values of $\Delta^{(k)}_{\text{rms}}$ in Eq. (3) for $S_c = 0$ and $S_c = N$ are indicated by dashed lines as a benchmark. At very weak disorder strengths $W/J \to 0$, $\Delta^{(k)}$ converges for large JT near the benchmark value with $S_c = 0$, consistent with previous results [24]. As the disorder strength increases, classical randomness gets converted into quantum randomness, causing $\Delta^{(k)}$ to decrease. At sufficiently long evolution times, $\Delta^{(k)}$ can become close to the benchmark value with $S_c = N$. This signifies a near-optimal conversion from classical to quantum randomness. Our analytical formula (3) works well here even though classical randomness is injected via the dynamics instead of the initial state, which demonstrates the generality of our protocol.

At strong disorder strengths $W/J \gg 1$, $\Delta^{(k)}$ increases and saturates. We attribute this behavior to the fact that for strong disorder, the projected ensemble behaves like an ensemble of random product states (at best a low-randomness 1-design), due to many-body localization effects that become relevant when $W/J \gtrsim 1$. We note that a more experimentally accessible scheme of adding a random global detuning $H_d = \xi \sum_i X_i$, where $\xi \sim \text{Uniform}[-W, W]$ (instead of spatially inhomogeneous disorder), also yields qualitatively similar results (not shown).

The crossover between the benchmark values at $S_c = 0$ and $S_c = N$ as disorder strength increases can be roughly estimated via a simple argument. Firstly, we need a quench time $JT \gtrsim N_A$ to get volume-law entanglement between the system and bath qubits. To contribute appreciably to the randomness of the projected ensemble, the effects of disorder must be integrated over a time T such that $WT \gtrsim 1$. On the other hand, for efficient conversion of classical to quantum randomness, we must avoid the many-body localized regime $W/J \gtrsim 1$ [41]. Therefore, for $JT \gg N_A$ we expect nearly maximal conversion of classical to quantum randomness for $1/JT \lesssim W/J \lesssim 1$. The behavior of $\Delta^{(2)}$ shown in Fig. 2(b) is consistent with this expectation. An interesting future direction would be to characterize the many-body localization transition via the projected ensemble.

Application: Classical shadow tomography—A practical application of our protocol is classical shadow tomography

[11] for learning expectation values of observables in unknown states. The state is scrambled with a unitary, followed by measurements in the computational basis. From the outcomes and the inverted scrambling dynamics, one can construct a classical representation of the unknown state, which can be used to accurately estimate the expectation values of many observables. In Refs. [12,59], it was proposed to use projected ensembles to generate the scrambling dynamics for shadow tomography. However, the estimation accuracy crucially depends on the quantum randomness of the projected ensemble [12]. Now we show that one can gain an exponential increase in accuracy by adding classical randomness, without incurring extra cost on the quantum computer.

The initial state is $\rho_A \otimes |x\rangle\langle x|_B$, where ρ_A is the unknown state to be learned. Classical randomness is injected by randomly initializing the B subsystem in computational basis states $|x\rangle$, similar to the setup in Fig. 1(a), up to a maximum of $S_c = 2^{N_B}$ bits. A unitary U is then applied on the full system. The projected ensemble $\mathcal{E}$ is constructed by measuring subsystem B in the computational basis $\{z_B\}$. The states in $\mathcal{E}$ are then measured in the computational basis $\{z_A\}$. The measurement outcomes (z_A, z_B) occur with probability $p_x(z_A, z_B) = \langle z_A, z_B|U(\rho_A \otimes |x\rangle\langle x|)U^\dagger|z_A, z_B\rangle$. In the classical postprocessing, we construct the shadow estimator as

$$\hat{\rho}_{A,x} = \frac{(2^{N_A}+1)\langle x|U^\dagger|z_A, z_B\rangle\langle z_A, z_B|U|x\rangle}{\mathrm{Tr}_A(\langle x|U^\dagger|z_A, z_B\rangle\langle z_A, z_B|U|x\rangle)} - I_A, \quad (6)$$

which satisfies the normalization $\mathrm{Tr}(\hat{\rho}_{A,x}) = 1$ (see SM [41] for a detailed explanation of the protocol). Intuitively, $\hat{\rho}_{A,x}$ is constructed to "undo" the scrambling of ρ_A, in order to estimate ρ_A. This gives the estimator $\hat{O} = \mathrm{Tr}(O\hat{\rho}_{A,x}) \approx \mathrm{Tr}(O\rho_A)$ for the observable O, averaged over L measurement shots. The shadow estimator can be analogously defined if classical randomness is instead injected by adding random disorder such as in Eq. (5) to the dynamics for a fixed initial state. The estimation error is $\delta O = |\hat{O} - \mathrm{Tr}(O\rho_A)|$, with the bias (i.e., systematic) error given by δO as $L \to \infty$. $\hat{\rho}_A$ is an unbiased estimator of ρ_A if the projected ensemble forms an exact 2-design [12]. For approximate 2-designs, the estimation incurs a bias error that grows with the distance $\Delta^{(2)}$ from a 2-design. Thus, from Theorem 1, we expect the bias error to be exponentially reduced by increasing S_c.

This is demonstrated in Fig. 3(a), which shows the average bias error against the classical entropy S_c injected into the projected ensemble, with randomly chosen Pauli operators O. We numerically simulate two different settings of classical randomness: (i) bath qubits randomly initialized from the uniform distribution over S_c classical bits, with U being a fixed Haar random unitary; (ii) bath qubits fixed in the $|0\rangle^{\otimes N_B}$ state and evolved under the disordered Hamiltonian (5) sampled from a set of 2^{S_c} randomly chosen disorder realizations [12].

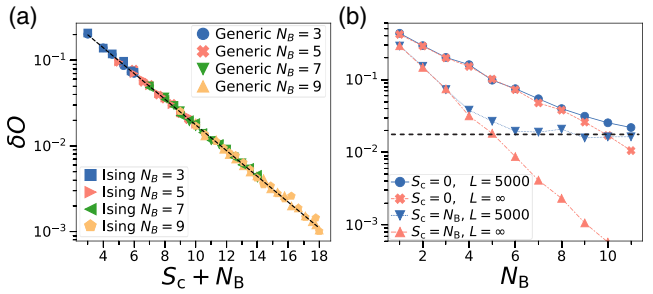

FIG. 3. (a) Average bias error δO (with $L \to \infty$ measurement shots) against classical entropy S_c plus number of bath qubits N_B. For "Generic," we randomize the initial state of the N_B bath qubits and evolve with a fixed Haar random unitary. For "Ising," we evolve with (5) for $JT = 100$, where the initial state is fixed and the disorder of the Hamiltonian is randomized for 2^{S_c} realizations with $W/J = 0.5$. Dashed line depicts $\delta O \propto 2^{-(S_c+N_B)/2}$. Error is nearly independent of whether randomness originates from measuring bath qubits, adding random initialization, or disorder in the time evolution. (b) Error δO against N_B for different L and S_c. Black dashed line shows the error for $L = 5000$ of an unbiased classical shadow protocol [11]. In both plots, we have $N_A = 1$ and average over 100 randomly chosen instances of unitaries to suppress statistical fluctuations (see SM [41] for details).

Similar to our analytical prediction for $\Delta^{(k)}_{\mathrm{rms}}$ in Theorem 1, we observe that the bias error decreases exponentially as $\sim 2^{-(N_B+S_c)/2}$ in both settings of classical randomness. Injecting S_c bits of entropy yields nearly the same reduction in bias error as having additional S_c bath qubits without classical randomness, arising from the near-optimal conversion from classical to quantum randomness. In Fig. 3(b), we plot δO against N_B, with classical randomness in the setting (i). We find that for a finite L, the error converges at large N_B to that expected of an unbiased shadow estimator (black dashed line) [11]. Increasing S_c causes the error to decrease significantly. Our protocol remains computationally efficient and scalable at large N_A, using the "patched quench" setup [59]; see SM [41] for details.

Discussion—We have shown how quantum chaos can convert classical randomness into quantum randomness by injecting classical entropy into the deep thermalization framework. For generic chaotic dynamics, each bit of classical entropy generates nearly as much quantum randomness as an additional bath qubit. From a practical viewpoint, injecting classical entropy allows one to improve the generation of approximate k-designs, with the maximum achievable k almost doubled, at only a small additional cost to the quantum hardware. This enhancement in quantum randomness directly translates to better performance for applications that utilize random quantum dynamics as a resource, as we have demonstrated through the example of shadow tomography. Our scheme is easier

to implement experimentally on many analog quantum simulators, as compared to well-established schemes using RUCs, by leveraging the inherent complexity of many-body quantum dynamics. From a theoretical perspective, our work raises interesting questions about whether injecting classical randomness improves the convergence of other related protocols such as finite-temperature projected ensembles and temporal ensembles to their respective maximum entropy ensembles [35], about the crossover between ergodic and localized behavior, and about the universal features of quantum chaos.

Acknowledgments—We thank Abhishek Anand, Jielun Chen, Wen Wei Ho, Daniel Mark, and Richard Tsai for helpful discussions. We acknowledge support from the DARPA ONISQ program (W911NF2010021), the DOE (DE-SC0021951), the Army Research Office MURI program (W911NF2010136), the NSF CAREER Grant (1753386), the Institute for Quantum Information and Matter, an NSF Physics Frontiers Center (NSF Grant PHY-1733907), and the Technology Innovation Institute (TII). J. P. also acknowledges support from the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research (DE-NA0003525, DE-SC0020290), and the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator.

[1] J. Cotler, N. Hunter-Jones, J. Liu, and B. Yoshida, Chaos, complexity, and random matrices, J. High Energy Phys. 11 (2017) 048.
[2] P. Hosur, X.-L. Qi, D. A. Roberts, and B. Yoshida, Chaos in quantum channels, J. High Energy Phys. 02 (2016) 004.
[3] D. A. Roberts and B. Yoshida, Chaos and complexity by design, J. High Energy Phys. 04 (2017) 121.
[4] D. N. Page, Information in black hole radiation, Phys. Rev. Lett. **71**, 3743 (1993).
[5] D. N. Page, Time dependence of hawking radiation entropy, J. Cosmol. Astropart. Phys. 09 (2013) 028.
[6] P. Hayden and J. Preskill, Black holes as mirrors: Quantum information in random subsystems, J. High Energy Phys. 09 (2007) 120.
[7] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Randomized benchmarking of quantum gates, Phys. Rev. A **77**, 012307 (2008).
[8] 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).
[9] P. Hayden, M. Horodecki, A. Winter, and J. Yard, A decoupling approach to the quantum capacity, Open Syst. Inf. Dyn. **15**, 7 (2008).
[10] S. Kimmel and Y.-K. Liu, Phase retrieval using unitary 2-designs, in *2017 International Conference on Sampling Theory and Applications (SampTA)* (2017), pp. 345–349, arXiv:2017.80244.
[11] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, Nat. Phys. **16**, 1050 (2020).
[12] M. McGinley and M. Fava, Shadow tomography from emergent state designs in analog quantum simulators, Phys. Rev. Lett. **131**, 160601 (2023).
[13] D. DiVincenzo, D. Leung, and B. Terhal, Quantum data hiding, IEEE Trans. Inf. Theory **48**, 580 (2002).
[14] P. Ananth, F. Kaleoglu, and H. Yuen, Simultaneous haar indistinguishability with applications to unclonable cryptography, arXiv:2405.10274.
[15] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell *et al.*, Quantum supremacy using a programmable superconducting processor, Nature (London) **574**, 505 (2019).
[16] A. Morvan, B. Villalonga, X. Mi, S. Mandrà, A. Bengtsson, P. V. Klimov, Z. Chen, S. Hong, C. Erickson, I. K. Drozdov *et al.*, Phase transition in random circuit sampling, arXiv:2304.11119.
[17] Q. Zhu, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan, M. Gong *et al.*, Quantum computational advantage via 60-qubit 24-cycle random circuit sampling, Sci. Bull. **67**, 240 (2022).
[18] A. L. Shaw, Z. Chen, J. Choi, D. K. Mark, P. Scholl, R. Finkelstein, A. Elben, S. Choi, and M. Endres, Benchmarking highly entangled states on a 60-atom analogue quantum simulator, Nature (London) **628**, 71 (2024).
[19] D. Gross, K. Audenaert, and J. Eisert, Evenly distributed unitaries: On the structure of unitary designs, J. Math. Phys. (N.Y.) **48**, 052104 (2007).
[20] F. G. S. L. Brandão, A. W. Harrow, and M. Horodecki, Local random quantum circuits are approximate polynomial-designs, Commun. Math. Phys. **346**, 397 (2016).
[21] J. Haferkamp, Random quantum circuits are approximate unitary t-designs in depth $O(nt^{5+o(1)})$, Quantum **6**, 795 (2022).
[22] A. W. Harrow and S. Mehraban, Approximate unitary t-designs by short random quantum circuits using nearest-neighbor and long-range gates, Commun. Math. Phys. **401**, 1531 (2023).
[23] T. Schuster, J. Haferkamp, and H.-Y. Huang, Random unitaries in extremely low depth, arXiv:2407.07754.
[24] J. S. Cotler, D. K. Mark, H.-Y. Huang, F. Hernández, J. Choi, A. L. Shaw, M. Endres, and S. Choi, Emergent quantum state designs from individual many-body wave functions, PRX Quantum **4**, 010311 (2023).
[25] 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, Preparing random states and benchmarking with many-body quantum chaos, Nature (London) **613**, 468 (2023).
[26] M. Ippoliti and W. W. Ho, Solvable model of deep thermalization with distinct design times, Quantum **6**, 886 (2022).
[27] W. W. Ho and S. Choi, Exact emergent quantum state designs from quantum chaotic dynamics, Phys. Rev. Lett. **128**, 060601 (2022).
[28] H. Wilming and I. Roth, High-temperature thermalization implies the emergence of quantum state designs, arXiv:2202.01669.

[29] P. W. Claeys and A. Lamacraft, Emergent quantum state designs and biunitarity in dual-unitary circuit dynamics, Quantum **6**, 738 (2022).
[30] H. Shrotriya and W. W. Ho, Nonlocality of deep thermalization, arXiv:2305.08437.
[31] C. Liu, Q. C. Huang, and W. W. Ho, Deep thermalization in continuous-variable quantum systems, arXiv:2405.05470.
[32] T. Bhore, J.-Y. Desaules, and Z. Papić, Deep thermalization in constrained quantum systems, Phys. Rev. B **108**, 104317 (2023).
[33] M. Ippoliti and W. W. Ho, Dynamical purification and the emergence of quantum state designs from the projected ensemble, PRX Quantum **4**, 030322 (2023).
[34] M. Lucas, L. Piroli, J. De Nardis, and A. De Luca, Generalized deep thermalization for free fermions, Phys. Rev. A **107**, 032215 (2023).
[35] D. K. Mark, F. Surace, A. Elben, A. L. Shaw, J. Choi, G. Refael, M. Endres, and S. Choi, A maximum entropy principle in deep thermalization and in Hilbert-space ergodicity, arXiv:2403.11970.
[36] A. Chan and A. D. Luca, Projected state ensemble of a generic model of many-body quantum chaos, arXiv:2402.16939.
[37] R.-A. Chang, H. Shrotriya, W. W. Ho, and M. Ippoliti, Deep thermalization under charge-conserving quantum dynamics, arXiv:2408.15325.
[38] N. D. Varikuti and S. Bandyopadhyay, Unraveling the emergence of quantum state designs in systems with symmetry, Quantum **8**, 1456 (2024).
[39] S. Goldstein, J. L. Lebowitz, R. Tumulka, and N. Zanghì, On the distribution of the wave function for systems in thermal equilibrium, J. Stat. Phys. **125**, 1193 (2006).
[40] S. Goldstein, J. L. Lebowitz, C. Mastrodonato, R. Tumulka, and N. Zanghì, Universal probability distribution for the wave function of a quantum system entangled with its environment, Commun. Math. Phys. **342**, 965 (2016).
[41] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.134.180403 for additional calculation details, which includes Refs. [42–55].
[42] A. W. Harrow, The church of the symmetric subspace, arXiv:1308.6595.
[43] A. Holevo, Statistical decision theory for quantum systems, J. Multivariate Anal. **3**, 337 (1973).
[44] C. W. Helstrom, Quantum detection and estimation theory, J. Stat. Phys. **1**, 231 (1969).
[45] D. Aharonov, J. Cotler, and X.-L. Qi, Quantum algorithmic measurement, Nat. Commun. **13**, 887 (2022).
[46] B. Collins and S. Matsumoto, Weingarten calculus via orthogonality relations: New applications, Lat. Am. J. Probab. Math. Stat. **14**, 631 (2017).
[47] G. Köstenberger, Weingarten calculus, arXiv:2101.00921.
[48] M. Cramer, Thermalization under randomized local Hamiltonians, New J. Phys. **14**, 053051 (2012).
[49] K. M. Audenaert, A sharp continuity estimate for the von Neumann entropy, J. Phys. A **40**, 8127 (2007).
[50] V. Oganesyan and D. A. Huse, Localization of interacting fermions at high temperature, Phys. Rev. B **75**, 155111 (2007).
[51] Y. Y. Atas, E. Bogomolny, O. Giraud, and G. Roux, Distribution of the ratio of consecutive level spacings in random matrix ensembles, Phys. Rev. Lett. **110**, 084101 (2013).
[52] B. Collins, S. Matsumoto, and J. Novak, The Weingarten calculus, Not. Am. Math. Soc. **69**, 1 (2022).
[53] T. Gorin, T. Prosen, T. H. Seligman, and M. Žnidarič, Dynamics of loschmidt echoes and fidelity decay, Phys. Rep. **435**, 33 (2006).
[54] L. D'Alessio, Y. Kafri, A. Polkovnikov, and M. Rigol, From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics, Adv. Phys. **65**, 239 (2016).
[55] L. Schatzki, Random real valued and complex valued states cannot be efficiently distinguished, arXiv:2410.17213.
[56] H. Kim and D. A. Huse, Ballistic spreading of entanglement in a diffusive nonintegrable system, Phys. Rev. Lett. **111**, 127205 (2013).
[57] H. Kim, T. N. Ikeda, and D. A. Huse, Testing whether all eigenstates obey the eigenstate thermalization hypothesis, Phys. Rev. E **90**, 052105 (2014).
[58] H. Perrin, T. Scoquart, A. I. Pavlov, and N. V. Gnezdilov, Dynamic thermalization on noisy quantum hardware, arXiv:2407.04770.
[59] 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).