A Single Fixed Shallow Circuit for Classical Shadows of Arbitrary n-Qubit States

Yu Wang[1,*] and Xiuwu Zhu[2]

[1] *Hetao Institute of Mathematics and Interdisciplinary Sciences, Shenzhen, Guangdong 518017, China*
[2] *Beijing Key Laboratory of Topological Statistics and Applications for Complex Systems, Beijing Institute of Mathematical Sciences and Applications, Beijing 101408, China*

(Dated: September 6, 2026)

Classical shadows extract many properties of quantum states from reusable classical records, typically obtained using randomized measurement settings. Although individual settings may be shallow, switching among them can introduce control, calibration, and reconfiguration costs beyond conventional circuit and sampling metrics. Here we construct, for every n, a fixed quantum analyzer whose Born outcomes replace externally sampled settings as labels for reusable shadow snapshots. The analyzer combines a freshly prepared n-qubit fiducial register A with parallel Bell readout of A and the unknown system S, producing one $2n$-bit record per copy. The same circuit realizes a rank-one minimal informationally complete measurement for tomography: a single fixed setting replaces the 3^n local-Pauli measurement settings commonly used for complete reconstruction, while retaining the minimum d^2 outcomes with $d = 2^n$. The resulting Pauli-diagonal frame admits an analytic inverse. For fixed Hermitian observables, the Haar-averaged conditional variance is bounded by a dimension-independent multiple of the squared Hilbert–Schmidt norm of the centered observable, whereas the state-uniform bound has a dimension-dependent coefficient. For $n \geq 3$, worst-state Pauli variances remain bounded in the axial sector while growing with dimension in the mixed sector. Fiducial preparation is completed before system contact: exact unitary preparation uses $n-1$ arbitrary two-qubit gates, no work qubits beyond A, and logarithmic depth under all-to-all connectivity. The unknown system undergoes only one parallel system–ancilla entangling layer, followed by local Hadamards and readout. Measurement-assisted preparation achieves $O(1)$ adaptive quantum depth per attempt using $n+1$ extra qubits in the heralded route, or deterministic $O(1)$ adaptive quantum depth using $O(n \log n)$ extra qubits, all counted beyond A. These statistical and hardware tradeoffs show that part of the measurement-setting randomness and control complexity traditionally supplied shot by shot can instead be compiled into the preparation and circuitry of a fixed, shallow, reusable quantum analyzer.

I. INTRODUCTION

Extracting useful information from large quantum systems is a central challenge in quantum information science. Meeting it requires both statistical efficiency and measurement architectures that can be implemented and controlled at scale [1, 2]. Classical shadows provide a powerful framework: repeated measurements on independent copies produce reusable classical records from which many observables can be estimated after data acquisition [3]. In ensemble-based realizations, each record combines an externally selected setting λ with a Born outcome [4]. The setting specifies a basis, circuit instance, or evolution parameter supplied through experimental control; the outcome is generated by the quantum measurement. Physically realizing those settings can require basis switching, pulse reconfiguration, and calibration. Measurement-setting control is therefore an experimental resource not fully captured by qubit count, gate count, circuit depth, or sample complexity. Can part of this randomness and control be compiled into a fixed, shallow, reusable quantum analyzer?

On the statistical side, ensemble choice determines observable-dependent sample costs, as quantified by shadow norms for global-Clifford and local-Pauli measurements [3]. Ensemble symmetries make reconstruction and bounds tractable [5, 6], while target-aware designs optimize sampling and schedules [7, 8], respect interaction-range constraints [9], or exploit selected Pauli families [10, 11] and fermionic structure [12]. This design freedom extends beyond unitary ensembles to generalized measurements [13, 14] and their reconstruction maps [15].

Implementation costs involve quantum execution, classical reconstruction, and experimental control. Shallow circuits trade entangling depth against variance [16, 17], and approximate unitary designs give logarithmic-depth global-shadow guarantees with controlled bias and variance corrections [18]. Practical choices also depend on scrambling and gate errors [19], platform-specific pulse and readout imperfections [20], and reconstruction costs addressed by tensor networks for suitable finite-depth ensembles [21]. Hamiltonian dynamics can replace gate randomization [22, 23], although a fixed Hamiltonian can retain a sampled evolution-time label [24]. Circuit reuse reduces reconfiguration overhead by sharing settings across shots, with statistical tradeoffs [25]. Reducing sampled-circuit depth does not itself remove external setting selection and realization or their platform-dependent calibration costs.

A complementary route encodes the measurement identity directly in the outcome space. In a fixed informationally complete (IC) positive operator-valued measure

* ming-jing-happy@163.com

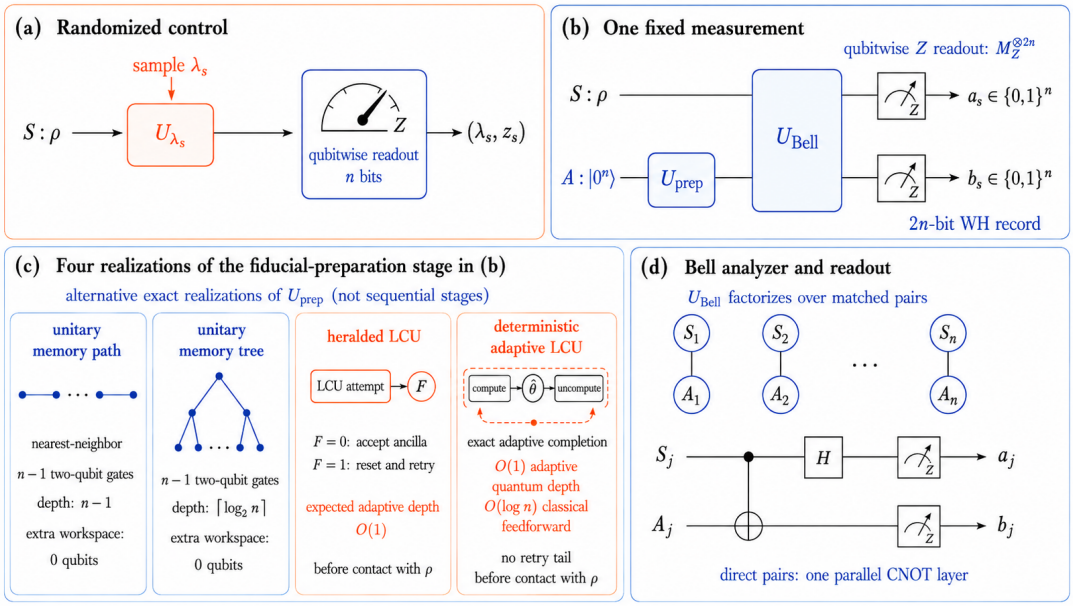

FIG. 1. Architecture and implementation hierarchy. (a) Conventional randomized shadows sample a new external setting λ_s on every shot and store it together with the n-bit Born outcome $\boldsymbol{z}_s$. Local-Pauli shadows sample from 3^n bases, while global-Clifford sampling uses an ensemble of size $2^{\Theta(n^2)}$. More generally, complete tomography of an arbitrary n-qubit state using only ancilla-free orthonormal-basis measurements requires at least 2^n+1 distinct measurement bases. (b) The fixed architecture instead repeats the same U_{prep} and U_{Bell}; qubitwise Z readout returns $\boldsymbol{a}_s$ and $\boldsymbol{b}_s$, which form one $2n$-bit WH record. (c) The unitary memory path, unitary memory tree, heralded LCU, and deterministic adaptive LCU are four alternative exact realizations of U_{prep}, not consecutive stages. Heralding retries before contact with the unknown state, whereas deterministic adaptive LCU removes the retry tail by measurement and feedforward. (d) A matched analyzer applies $\text{CNOT}_{S_j \to A_j}$, then H_{S_j} and two Z-basis readouts. Direct matched couplers place all n CNOTs in one parallel layer. Sections VII and VIII prove the induced POVM and resource claims.

(POVM), the Born outcome directly identifies the effect and its dual-frame estimator [26, 27]; fixed system–ancilla quenches already provide analog protocols for such outcome-only acquisition [28, 29]. An early circuit for single-qubit symmetric informationally complete (SIC) tomography uses two auxiliary qubits [30]. Products of single-qubit SIC measurements offer a scalable fixed implementation with local circuits and product reconstruction, demonstrated through eight qubits [31, 32]. Their Pauli second moment is 3^w at weight w, making high-weight Pauli targets costly to estimate. A global SIC instead provides an isotropic minimal-IC benchmark with d^2 rank-one outcomes in dimension d [33, 34]. Exact SICs have been established unconditionally in only finitely many dimensions; their existence for every $d = 2^n$ remains open [35–38]. Even a known fiducial need not admit shallow qubit preparation. General realization tools include Naimark dilation [39, 40] and group-covariant constructions [41, 42]; probabilistic [43] and adaptive measurement schemes [44, 45] offer alternative resource tradeoffs. A fixed analyzer therefore requires joint consideration of statistical geometry and circuit realizability. Table I summarizes representative implementation routes and their resource tradeoffs.

Here we construct a fixed analyzer for every n, using a controlled relaxation of exact SIC symmetry. An n-qubit auxiliary register A is prepared afresh for each copy in a known conjugate fiducial, before contact with the unknown system S. Fixed, fiducial-independent Bell readout produces a $2n$-bit Born outcome labeling the measurement effect and reusable shadow snapshot. With direct matched couplers, S undergoes one parallel system–ancilla CNOT layer and local Hadamards before readout, with no S–S multiqubit gates. All fiducial-dependent preparation is confined to A. Building on the general fiducial-plus-Bell interface [46], we supply explicit preparation and characterize its statistical and circuit costs. Figure 1 contrasts this interface with randomized-setting acquisition.

The analyzer uses the tensor-Pauli Weyl–Heisenberg (WH) family of Ref. [47], which established its informational completeness and projector-Gram spectrum. It gives a rank-one minimal-IC measurement with $d^2 = 4^n$ outcomes for $d = 2^n$, correlated for $n \geq 2$. We derive an analytic inverse of its Pauli-diagonal channel and exact observable-estimation variances. For any fixed Hermitian target O, writing $O_0 = O - \text{Tr}(O)\mathbb{I}/d$, we prove

$$\mathbb{E}_{\psi\sim\text{Haar}} \text{Var}_{|\psi\rangle\langle\psi|}[\widehat{O}] \leq 2(1+1/d)\,\text{Tr}(O_0^2),$$
$$\sup_{\rho} \text{Var}_{\rho}[\widehat{O}] \leq 2(d+1)\,\text{Tr}(O_0^2).$$

For $n \geq 3$, tight worst-state Pauli variances are $O(1)$

for axial strings containing only X/I or only Z/I, and $\Theta(d)$ for the remaining nonidentity strings. The bounded-cost axial sector spans every weight, while local-Pauli and product-SIC shadows retain bounded costs for fixed-weight mixed targets. Global-Clifford shadows retain a stronger dimension-independent state-uniform bound [3]. The same records also support minimal-IC tomography: one fixed setting replaces the 3^n standard local-Pauli settings, while the d^2-outcome alphabet and the sample and storage costs of dense reconstruction remain exponential.

Exact unitary fiducial preparation uses $n-1$ arbitrary two-qubit gates and no work qubits beyond A. A nearest-neighbor path has linear depth; a balanced all-to-all schedule has asymptotically optimal two-qubit depth $\lceil \log_2 n \rceil$. Measurement-assisted alternatives, completed before system contact, offer constant expected adaptive quantum depth through heralding, or deterministic constant adaptive quantum depth with additional auxiliaries (Sec. VIII). Restricted connectivity may add routing overhead.

Together, these results show how auxiliary width and preparation can be traded for reduced setting control while keeping statistical costs explicit.

II. ONE FIXED MINIMAL MEASUREMENT AND ITS REUSABLE RECORD

A. The fixed Weyl–Heisenberg minimal IC measurement

For any n-qubit system, the Hilbert-space dimension is $d = 2^n$. We label each Weyl–Heisenberg displacement by two bit strings $\boldsymbol{a} = (a_1, \ldots, a_n)$ and $\boldsymbol{b} = (b_1, \ldots, b_n)$ in $\mathbb{F}_2^n$. Up to an irrelevant Pauli phase, define

$$D_{\boldsymbol{a},\boldsymbol{b}} = Z(\boldsymbol{a})X(\boldsymbol{b}) = \bigotimes_{j=1}^{n} Z_j^{a_j} X_j^{b_j}. \tag{1}$$

The concrete fiducial used throughout is

$$|\phi_n\rangle = \frac{|+\rangle^{\otimes n} + e^{i\theta_n}|0^n\rangle}{\sqrt{\mathcal{N}_n}}, \qquad \mathcal{N}_n = 2 + 2^{1-n/2}\cos\theta_n, \tag{2}$$

$$\theta_1 = 5\pi/12, \qquad \theta_2 = 2\pi/3, \qquad \theta_n = 3\pi/4 \quad (n \geq 3).$$

Its tensor-Pauli orbit defines

$$\Pi^{(n)}_{\boldsymbol{a},\boldsymbol{b}} = D_{\boldsymbol{a},\boldsymbol{b}}|\phi_n\rangle\langle\phi_n|D^\dagger_{\boldsymbol{a},\boldsymbol{b}}, \qquad E^{(n)}_{\boldsymbol{a},\boldsymbol{b}} = \frac{1}{d}\Pi^{(n)}_{\boldsymbol{a},\boldsymbol{b}}. \tag{3}$$

The Pauli twirl gives $\sum_{\boldsymbol{a},\boldsymbol{b}} E^{(n)}_{\boldsymbol{a},\boldsymbol{b}} = \mathbb{I}$, so the positive rank-one operators in Eq. (3) form a POVM. Appendix A gives a direct Pauli-basis proof of this normalization. We denote this fixed POVM by

$$\mathsf{E}^{(n)} := \left\{ E^{(n)}_{\boldsymbol{a},\boldsymbol{b}} : \boldsymbol{a}, \boldsymbol{b} \in \mathbb{F}_2^n \right\}. \tag{4}$$

Normalization alone does not imply informational completeness. The orbit in Eq. (3) is the physical-qubit tensor-Pauli form of the characteristic-two finite-field WH family studied in Ref. [47], which establishes informational completeness and determines its projector-Gram spectrum. Since the POVM has exactly d^2 outcomes, it is minimal IC. Section III and Appendices B 1 and B 2 provide an independent tensor-Pauli derivation.

As a separate consequence, the equal-weight rank-one projectors form an operator basis, so this is also a balanced informationally complete (BIC) POVM. In the associated Bell protocol, maximal violation can certify the maximal $2\log_2 d = 2n$ device-independent private random bits [58]. The circuit developed here supplies the fixed measurement component of that protocol.

Later we analyze how this one fixed POVM performs as an alternative to the exponentially large ensembles of randomized measurement circuits used in standard classical shadows and to the 3^n local-Pauli settings used in conventional tomography.

B. One outcome, one canonical snapshot

Measuring one copy of an n-qubit state ρ with $\mathsf{E}^{(n)}$ returns a $2n$-bit outcome label $u = (\boldsymbol{a}, \boldsymbol{b})$. From now on, write $E_u := E^{(n)}_{\boldsymbol{a},\boldsymbol{b}}$ and $\Pi_u := \Pi^{(n)}_{\boldsymbol{a},\boldsymbol{b}}$. To reuse this record for either shadow estimation or tomography, we assign to it a matrix-valued estimator. The outcome occurs with Born probability $p_\rho(u) = \mathrm{Tr}(\rho E_u) = d^{-1}\,\mathrm{Tr}(\rho\Pi_u)$. If we associate the raw projector Π_u with this outcome, its average over repeated shots is

$$\mathbb{E}_\rho[\Pi_u] = \sum_u p_\rho(u)\Pi_u = \mathcal{M}(\rho),$$

where the normalized frame channel of the POVM is

$$\mathcal{M}(A) = \frac{1}{d}\sum_u \mathrm{Tr}(\Pi_u A)\Pi_u. \tag{5}$$

Thus $\mathcal{M}$ maps the input state ρ to the average raw estimator $\mathbb{E}_\rho[\Pi_u]$. Minimal informational completeness makes $\mathcal{M}$ invertible, so we correct each outcome by defining the canonical snapshot

$$\widehat{\rho}_u = \mathcal{M}^{-1}(\Pi_u). \tag{6}$$

Since $\mathcal{M}(\rho) = \sum_u p_\rho(u)\Pi_u$, the object inverted for each outcome is Π_u. Equivalently,

$$\widehat{\rho}_u = d\,\mathcal{M}^{-1}(E_u).$$

When the measurement outcome u is drawn according to the Born probabilities $p_\rho(u)$, the corrected snapshot is unbiased:

$$\mathbb{E}_{u\sim p_\rho}[\widehat{\rho}_u] = \sum_u p_\rho(u)\mathcal{M}^{-1}(\Pi_u) = \mathcal{M}^{-1}[\mathcal{M}(\rho)] = \rho. \tag{7}$$

TABLE I. Fixed-analyzer implementation tradeoffs. Here d is the dimension of one encoded qudit, whereas n counts individually addressable qubits. A whole-POVM realization returns one of all outcomes from the same per-copy device; an effect scan programs the effects separately. The routes are representative and need not be mutually exclusive. The present correlated family combines minimal IC outcomes and an analytic inverse with an explicit all-n qubit circuit and, with direct matched couplers, a single system-facing entangling layer.

implementation route	dimension / scope	map applied to each copy	circuit implication
effect-by-effect photonics [48]	SIC projectors and tomography at $d = 6, 10$	**Effect scan:** program one SIC projector at a time and combine data from d^2 settings	one encoded qudit; programmable projection measurements acquired separately and combined for reconstruction
fixed single-qudit dilations [49–55]	fixed multioutcome SICs at $d = 2, 3, 4$; quantum-walk SICs at $d = 2, 3$	**Fixed dilation:** the same loop, multiport, or walk acts on every copy; its terminal port or position is the outcome	resources are expressed in native modes, ports, or walk steps; translation to individually addressable qubits depends on the encoding
product qubit SIC [30–32, 56]	demonstrated through $n = 8$; tensor-product minimal IC for every n, with an explicit product inverse	**Fixed local product measurement:** n one-qubit tetrahedral-SIC circuits act independently and in parallel	tensoring the optimized one-qubit ancilla circuit [32] gives n disjoint CNOTs in one layer; higher-level embeddings offer an alternative [56]. Product structure gives Pauli second moment 3^w at weight w
adaptive generalized measurements [44, 45, 57]	general synthesis; SIC demonstrations on an encoded $d = 4$ qudit [44] and on up to two qubits [45]	**Adaptive realization:** ancilla readout and reset, with outcome-conditioned operations; hybrid schemes terminate in a Naimark block	a complete logical POVM with a reused auxiliary qubit; sequential control trades auxiliary width for measurement and feedback, with target- and platform-dependent gate costs
present tensor-Pauli family [47]	**every** n; rank-one minimal IC, correlated for $n \geq 2$, with exact SICs at $n = 1, 3$	**Fixed whole-POVM unitary analyzer:** prepare $\|\phi_n^*\rangle$ and apply matched Bell readout; analytic Pauli-diagonal inverse	**System interface:** one parallel CNOT layer, no S–S gates. Balanced all-to-all total for $n \geq 2$: $3n - 3$ CNOTs, CNOT depth $2\lceil \log_2 n \rceil$, before routing; at $n = 1$, one Bell CNOT

Thus every Born outcome, after the fixed correction $\mathcal{M}^{-1}$, gives an unbiased one-shot estimator of the full state, although it need not be positive semidefinite. Averaging these snapshots gives linear tomography, whereas contracting them with selected observables gives classical-shadow estimates. Section III evaluates the correction map explicitly.

C. Two estimators from the same records

a. Observable prediction. For each Hermitian target O, one evaluates only the scalar

$$\widehat{O}(u) = \mathrm{Tr}(O\widehat{\rho}_u), \qquad \widehat{O}_N = \frac{1}{N}\sum_{s=1}^{N} \widehat{O}(u_s). \tag{8}$$

Equation (7) gives $\mathbb{E}_\rho \widehat{O}_N = \mathrm{Tr}(\rho O)$. This is the classical-shadow route: many observables can be evaluated from the same records without first constructing a dense estimate of ρ.

b. Linear tomography. Let N_u be the number of times outcome u occurs and $f_u = N_u/N$ its empirical frequency. Averaging the same snapshots gives the closed-form linear-inversion estimator

$$\widehat{\rho}_N = \frac{1}{N}\sum_{s=1}^{N} \widehat{\rho}_{u_s} = \sum_u f_u \, \mathcal{M}^{-1}(\Pi_u), \qquad \mathbb{E}_\rho \widehat{\rho}_N = \rho. \tag{9}$$

The two expressions are identical: one averages record-by-record snapshots, or equivalently applies the dual frame to the empirical POVM probabilities. Because the POVM is minimal IC, this unbiased linear inverse is unique. For an exact SIC, tightness reduces each snapshot to the familiar form $\widehat{\rho}_u = (d+1)\Pi_u - \mathbb{I}$ [33, 34, 59]. Our WH POVM is generally non-tight: different Pauli components are rescaled by different inverse eigenvalues, rather than by a single universal SIC factor. Its Pauli-basis inverse and the closed-form dual for our fiducial are derived below.

III. MEASUREMENT CHANNEL AND CANONICAL RECONSTRUCTION

For the fiducial in Eq. (2), write $\mathcal{M}_{\phi_n} \equiv \mathcal{M}$ for the frame channel in Eq. (5). The Pauli-basis diagonalization below uses only tensor-Pauli covariance; the chosen fiducial enters through its Pauli expectation values. Each outcome u is converted into the canonical snapshot $\widehat{\rho}_u = \mathcal{M}_{\phi_n}^{-1}(\Pi_u)$. We evaluate this map before discussing either its statistical performance or its physical implementation.

A. Measurement channel in the Pauli basis

Let P_v, with $v = (\boldsymbol{c}, \boldsymbol{e})$, denote the Hermitian tensor-Pauli whose jth factor is I, X, Z, Y for $(c_j, e_j) = (0,0), (0,1), (1,0), (1,1)$, respectively. It differs from $D_{\boldsymbol{c},\boldsymbol{e}}$ in Eq. (1) only by a phase. These operators satisfy $P_v^\dagger = P_v$ and $\mathrm{Tr}(P_v P_w) = d\delta_{v,w}$. For each P_v, define

its expectation value in the fiducial state $|\phi_n\rangle$ by

$$\chi_v := \langle\phi_n|P_v|\phi_n\rangle. \tag{10}$$

Because P_v is Hermitian, every χ_v is real.

Pauli-diagonal reconstruction for Pauli-invariant unitary ensembles was established by Bu et al. [6]. We derive the corresponding form for the fixed Pauli-covariant POVM below.

Theorem 1 (Measurement channel and its inverse in the Pauli basis)**.** *For every tensor-Pauli operator P_v,*

$$\mathcal{M}_{\phi_n}(P_v) = |\chi_v|^2 P_v. \tag{11}$$

Consequently, for any operator A, the channel has the first form below. If every χ_v is nonzero, then $\mathcal{M}_{\phi_n}$ is invertible and its inverse is the second:

$$\begin{aligned}\mathcal{M}_{\phi_n}(A) &= \frac{1}{d}\sum_v |\chi_v|^2 \operatorname{Tr}(P_v A) P_v,\\ \mathcal{M}^{-1}_{\phi_n}(A) &= \frac{1}{d}\sum_v \frac{\operatorname{Tr}(P_v A)}{|\chi_v|^2} P_v.\end{aligned} \tag{12}$$

Pauli covariance reduces all d^2 snapshots to one reference dual. With $D_u := D_{\boldsymbol{a},\boldsymbol{b}}$ and $\Pi_0 := |\phi_n\rangle\langle\phi_n|$, define

$$F_n := \mathcal{M}^{-1}_{\phi_n}(\Pi_0) = \frac{1}{d}\sum_v \frac{1}{\chi_v} P_v, \tag{13}$$

$$\widehat{\rho}_u = \mathcal{M}^{-1}_{\phi_n}(\Pi_u) = D_u F_n D_u^\dagger. \tag{14}$$

Thus one computes the inverse only for the fiducial projector Π_0; conjugation by the observed displacement D_u gives every snapshot. Appendix B 1 proves the Pauli diagonalization, the channel and inverse formulas, and the covariance used above. Equation (12) therefore reduces the inverse to coefficientwise scalar division; no $d^2 \times d^2$ matrix is inverted.

B. Closed-form dual for the chosen fiducial

Equation (14) reduces all d^2 reconstruction rules to one reference dual F_n. For $n \geq 3$, set $r_d := \sqrt{2/d}$, $\kappa_d := 1 - r_d$, and $z_n := e^{3\pi i/4}$. Equation (2) then gives $\mathcal{N}_n = 2 - r_d$.

Theorem 2 (Closed-form canonical dual)**.** *For the fiducial family in Eq. (2), the two exceptional reference duals are*

$$\begin{aligned}F_1 &= 3\Pi_0 - \mathbb{I},\\ F_2 &= 6\Pi_0 + 3\Pi_0^{\mathsf{T}} - 2\mathbb{I},\end{aligned} \tag{15}$$

where the transpose is taken in the computational basis. For $n \geq 3$, the reference dual is

$$\begin{aligned}F_n = {}& \frac{\mathcal{N}_n}{r_d \kappa_d}(|+^n\rangle\langle+^n| + |0^n\rangle\langle 0^n|)\\ &+ \frac{\mathcal{N}_n}{r_d^2}\left(z_n^*|+^n\rangle\langle 0^n| + z_n|0^n\rangle\langle+^n|\right) - \frac{r_d}{\kappa_d}\mathbb{I}.\end{aligned} \tag{16}$$

For every n and every outcome u, the canonical snapshot is $\widehat{\rho}_u = D_u F_n D_u^\dagger$.

At $n = 1$ and $n = 3$, the Pauli–WH orbits are the tetrahedral and Hoggar SICs, respectively [33, 47]. In both cases the result reduces to the standard SIC rule $\widehat{\rho}_u = (d+1)\Pi_u - \mathbb{I}$, giving $3\Pi_u - \mathbb{I}$ and $9\Pi_u - \mathbb{I}$, respectively.

For $u = (\boldsymbol{a}, \boldsymbol{b})$, write $|\boldsymbol{b}\rangle := X(\boldsymbol{b})|0^n\rangle$. Tensor-Pauli conjugation acts on the two product branches as

$$D_u|+\rangle^{\otimes n} = Z(\boldsymbol{a})|+\rangle^{\otimes n}, \qquad D_u|0^n\rangle = (-1)^{\boldsymbol{a}\cdot\boldsymbol{b}}|\boldsymbol{b}\rangle. \tag{17}$$

For $n \geq 3$, Eq. (17) shows directly that tensor-Pauli conjugation preserves the four rank-one product-state terms in F_n. Hence every snapshot has the same five-term structure. Appendix B 2 obtains the fiducial-specific sectorwise inverse, and Appendix B 3 converts it into this product form and its explicit outcome-dependent version.

The product form also makes the role of coherence transparent. Apart from the identity label $(\mathbf{0}, \mathbf{0})$, the Pauli directions split into the axial and mixed sets

$$\begin{aligned}\mathsf{Ax} := {}& \underbrace{\{(\mathbf{0}, \boldsymbol{e}) : \boldsymbol{e} \neq \mathbf{0}\}}_{X/I\text{-only}}\\ &\cup \underbrace{\{(\boldsymbol{c}, \mathbf{0}) : \boldsymbol{c} \neq \mathbf{0}\}}_{Z/I\text{-only}},\\ \mathsf{Mix} := {}& \underbrace{\{(\boldsymbol{c}, \boldsymbol{e}) : \boldsymbol{c} \neq \mathbf{0},\ \boldsymbol{e} \neq \mathbf{0}\}}_{\text{mixed directions}}.\end{aligned} \tag{18}$$

Together with the identity, the axial sector contains $2d - 1$ Pauli directions, while the mixed sector contains the remaining $(d-1)^2$. The WH orbit of the equal-weight mixture $\tau_{\text{inc}} = (|+^n\rangle\langle+^n| + |0^n\rangle\langle 0^n|)/2$ spans only the identity and axial directions and therefore has frame rank $2d - 1$. The two cross outer products in the coherent fiducial projector give nonzero Pauli coefficients in the remaining $(d-1)^2$ directions, completing the IC frame. Section IV quantifies the resulting direction-dependent reconstruction cost.

IV. OBSERVABLE-ESTIMATION PERFORMANCE

A. Spectral input and standard benchmarks

For canonical classical shadows constructed from an IC POVM, inversion of the measurement frame enters the single-shot raw second moments and controls standard variance and sample-complexity bounds [13, 27]. The weakest nonidentity frame direction is therefore the relevant spectral quantity for such bounds.

For the Pauli–WH POVM generated by the fiducial $|\phi_n\rangle$ in Eq. (2), this amplification can be read directly from the projector-Gram spectrum computed in Ref. [47]. Define $G^{\Pi}_{u,u'} := \operatorname{Tr}(\Pi_u \Pi_{u'})$. Pauli covariance associates direction P_v with the projector-Gram eigenvalue $\lambda_v :=$

$d|\chi_v|^2$ (Appendix C 1). Thus Eq. (11) becomes

$$\mathcal{M}_{\phi_n}(P_v) = \frac{\lambda_v}{d} P_v,$$
$$\mathcal{M}^{-1}_{\phi_n}(P_v) = \frac{d}{\lambda_v} P_v = \frac{1}{|\chi_v|^2} P_v.$$

Hence d/λ_v is the inverse-frame amplification of the Pauli direction P_v. For a Pauli target, this factor is exactly the state-independent raw second moment of the canonical estimator; for a general target, the inverse spectrum enters through a quadratic form and controls bounds rather than determining every variance by itself. The largest nonidentity inverse factor is $d/\lambda_{\min}$, where

$$\lambda_{\min} = \min_{v\neq 0} \lambda_v, \qquad \eta_n = \frac{d+1}{d}\lambda_{\min}. \tag{19}$$

Maximizing $\lambda_{\min}$ therefore minimizes the worst Pauli-direction amplification. This gives the phase-selection principle for Eq. (2): within its equal-weight two-product-state family, the listed phases maximize $\lambda_{\min}$ over the relative phase (Appendix C 2). An exact global SIC has $\lambda_{\min} = d/(d+1)$, so $\eta = 1$ is the unit benchmark.

Theorem 3 (Exact spectral floor [47])**.** *For the family in Eq. (2),*

$$\lambda_{\min} = \begin{cases} 2/3, & n = 1, \\ 4/9, & n = 2, \\ \dfrac{2}{(2-\sqrt{2/d})^2}, & n \geq 3. \end{cases} \tag{20}$$

Consequently, $(\eta_1, \eta_2, \eta_3) = (1, 5/9, 1)$*, and for all* n*,*

$$\eta_n \geq \frac{1}{2}, \qquad \eta_n \longrightarrow \frac{1}{2}. \tag{21}$$

The exact finite-n multiplicities are listed in Appendix C 2.

Thus, although the family is not a SIC except at $n = 1, 3$, it retains at least half the SIC spectral floor in every size. Equivalently, its worst-direction inverse-frame amplification is at most twice the SIC benchmark:

measurement	$d/\lambda_{\min}$
global SIC	$d+1$
product tetrahedral SIC	3^n
present global WH	$\leq 2(d+1)$

(22)

The global-SIC and product-SIC values are standard [31–34]. The exact value of $\lambda_{\min}$ for the present construction is given in Eq. (20). The product construction is explicit for every n, but its weakest-direction amplification exceeds the global-SIC value by the factor $3^n/(d+1) \sim (3/2)^n$; equivalently,

$$\eta_{\text{prod}}(n) = \frac{d+1}{3^n} \longrightarrow 0. \tag{23}$$

In contrast, the present WH family is explicit for every n-qubit system and satisfies $\eta_n \geq 1/2$. Its worst-direction inverse-frame amplification is therefore at most $1/\eta_n \leq 2$ times the global-SIC benchmark; this ratio equals one for $n = 1, 3$ and approaches two as $n \to \infty$. The comparison uses the global-SIC value only as a benchmark and does not assume the existence of an exact SIC in every dimension $d = 2^n$ [37, 38].

B. General Hermitian observables

For a Hermitian target O, one record returns $\widehat{O}(u) = \mathrm{Tr}(O\widehat{\rho}_u)$.

Writing $O_0 = O - \mathrm{Tr}(O)\mathbb{I}/d$, define for the canonical snapshots the state-dependent squared shadow norm and its worst-state counterpart by

$$\|O_0\|^2_{\mathrm{sh},\rho} := \mathbb{E}_\rho[\widehat{O}_0^2], \qquad \|O_0\|^2_{\mathrm{sh}} := \sup_\rho \|O_0\|^2_{\mathrm{sh},\rho},$$

where the supremum is over density operators. Thus the shadow norm is a raw second moment, while $\mathrm{Var}_\rho[\widehat{O}] = \|O_0\|^2_{\mathrm{sh},\rho} - \mathrm{Tr}(\rho O_0)^2$.

Proposition 4 (General observable: average and worst–case variance)**.** *Let* O *be any Hermitian observable and let* O_0 *be its traceless part defined above. Draw* $|\psi\rangle$ *from Haar measure on pure states and set* $\rho_\psi = |\psi\rangle\langle\psi|$*. For each fixed* ρ_ψ*,* $\mathrm{Var}_{\rho_\psi}[\widehat{O}]$ *is the variance over the measurement outcome. Averaging this variance over the Haar-random choice of* $|\psi\rangle$ *gives*

$$\begin{aligned} &\mathbb{E}_{\psi\sim\mathrm{Haar}} \mathrm{Var}_{\rho_\psi}[\widehat{O}] \\ &\quad = \frac{1}{d}\mathrm{Tr}[O_0 \mathcal{M}^{-1}_{\phi_n}(O_0)] - \frac{\mathrm{Tr}(O_0^2)}{d(d+1)} \\ &\quad \leq 2\left(1 + \frac{1}{d}\right)\mathrm{Tr}(O_0^2). \end{aligned} \tag{24}$$

Uniformly over all input states, the worst-case variance obeys

$$\sup_\rho \mathrm{Var}_\rho[\widehat{O}] \leq 2(d+1)\,\mathrm{Tr}(O_0^2). \tag{25}$$

The distinction between the two bounds is already visible in the raw second moment. It is linear in the input state, so its Haar average equals its value at $\mathbb{I}/d$; for an arbitrary state, $\rho \leq \mathbb{I} = d(\mathbb{I}/d)$ bounds it by d times that value. The proof is given in Appendix C 3.

For arbitrary Hermitian targets, including fully global ones, the Haar-input average in Eq. (24) has a dimension-independent coefficient for all the global benchmarks considered here. An ideal global SIC, uniform sampling from a complete set of stabilizer MUBs, and randomized global-Clifford measurements realize the same depolarizing frame and give $(1 + 1/(d+1))\,\mathrm{Tr}(O_0^2)$ [3, 33, 34, 60, 61]. The complete stabilizer MUB uses the minimum $d+1$ Clifford measurement bases; it is an exact finite realization of this frame, not a minimal-IC POVM. The present fixed

d^2-outcome minimal-IC measurement is instead bounded by $2(1+1/d)\,\mathrm{Tr}(O_0^2)$.

State-uniformly, the comparison is different. The full randomized global-Clifford ensemble satisfies, for every Hermitian target, $\sup_\rho \mathrm{Var}_\rho[\widehat{O}_{\mathrm{Cl}}] \le 3\,\mathrm{Tr}(O_0^2)$, whereas Eq. (25) gives the upper bound $2(d+1)\,\mathrm{Tr}(O_0^2)$ for the fixed WH measurement [3].

A shallow randomized alternative trades circuit depth for a controlled statistical penalty. On a one-dimensional nearest-neighbor architecture without ancillas, a relative-error $\varepsilon_{\mathrm{des}}$-approximate unitary 3-design reduces the Clifford depth from $O(n)$ to $O(\log(n/\varepsilon_{\mathrm{des}}))$. For $O \ge 0$, the standard reconstruction has variance at most $3\,\mathrm{Tr}(O_0^2) + 10\varepsilon_{\mathrm{des}}[\mathrm{Tr}(O)]^2$ and bias at most $2\varepsilon_{\mathrm{des}}\,\mathrm{Tr}(O)$ [18].

The gap between averaged and state-uniform guarantees is also reflected in outcome count. The present POVM is minimal IC and has d^2 outcomes, whereas, among rank-one IC POVMs, a bound of the form $\sup_\rho \mathbb{E}_\rho[\widehat{O}_0^2] \le C\,\mathrm{Tr}(O_0^2)$, uniformly for all O_0 with C independent of d, requires $\Omega(d^3)$ outcomes [62].

Corollary 5 (State-dependent simultaneous estimation)**.** *Fix an input state ρ, let $O_1, \dots, O_M$ be Hermitian targets, and put*

$$B_{\mathrm{var}}(\rho) := \max_k \mathrm{Var}_\rho[\widehat{O}_k].$$

Median-of-means postprocessing of independent records drawn from ρ estimates all M expectations to additive error ϵ with failure probability at most δ using

$$N = O\left(\frac{B_{\mathrm{var}}(\rho)\log(M/\delta)}{\epsilon^2}\right). \tag{26}$$

The quantum circuit and raw records are independent of the chosen targets; Eq. (25) may be substituted for $B_{\mathrm{var}}(\rho)$ when a state-independent shot count is required.

Proposition 4 and Corollary 5 apply to arbitrary Hermitian targets. Since Pauli operators are eigenoperators of the measurement channel, their variances admit an exact evaluation.

C. Exact Pauli expectations

For $n \ge 3$, we use the Pauli-sector split displayed after Theorem 2:

$$\{I\} \;\dot\cup \underbrace{(X/I\text{-only}) \,\dot\cup\, (Z/I\text{-only})}_{\mathsf{Ax},\ 2(d-1)} \;\dot\cup \underbrace{(\text{mixed directions})}_{\mathsf{Mix},\ (d-1)^2}.$$

For a Pauli label $v = (\boldsymbol{c}, \boldsymbol{e})$, the axial class Ax has exactly one of $\boldsymbol{c}, \boldsymbol{e}$ nonzero, whereas the mixed class Mix has both nonzero. Equivalently, a mixed string contains a Y, or contains an X and a Z on different qubits.

For $u = (\boldsymbol{a}, \boldsymbol{b})$ and $v = (\boldsymbol{c}, \boldsymbol{e})$, the commutation sign and the canonical Pauli estimator are

$$\begin{aligned} &s_v(u) := (-1)^{\boldsymbol{a}\cdot\boldsymbol{e}+\boldsymbol{b}\cdot\boldsymbol{c}}, && D_u P_v D_u^\dagger = s_v(u) P_v, \\ &\widehat{P}_v(u) := \mathrm{Tr}(P_v \widehat{\rho}_u) = \frac{s_v(u)}{\chi_v}, && s_v(u) \in \{+1, -1\}, \\ &\widehat{P}_v(u)^2 = \frac{1}{|\chi_v|^2} = \frac{d}{\lambda_v} =: C(P_v). \end{aligned} \tag{27}$$

Thus the outcome changes only the sign of the estimate, not its magnitude. Consequently, $C(P_v) = \mathbb{E}_\rho[\widehat{P}_v^{\,2}]$ is independent of the input state. The sign-resolved derivation is given in Appendix C 4.

Theorem 6 (Exact classwise Pauli variances for $n \ge 3$)**.** *Define*

$$\begin{aligned} A_d &:= \left(\frac{2-\sqrt{2/d}}{1-\sqrt{2/d}}\right)^2 \in (4, 9], \quad A_d \to 4, \\ M_d &:= (\sqrt{2d}-1)^2 < 2d, \quad M_d = \Theta(d) = \Theta(2^n). \end{aligned} \tag{28}$$

For a fixed state ρ, write $m_v(\rho) := \mathrm{Tr}(\rho P_v)$. Then every nonidentity Pauli direction has the following exact variance, raw second moment, and tight worst-state variance:

$$\begin{aligned} \mathrm{Var}_\rho(\widehat{P}_v) &= C(P_v) - m_v(\rho)^2, \\ C(P_v) = \mathbb{E}_\rho[\widehat{P}_v^{\,2}] &= \sup_\sigma \mathrm{Var}_\sigma(\widehat{P}_v) \\ &= \begin{cases} A_d, & P_v \in \mathsf{Ax}, \\ M_d, & P_v \in \mathsf{Mix}. \end{cases} \end{aligned} \tag{29}$$

The supremum in each row is attained at $\sigma = I/d$; the identity has zero variance. The axial row contains the $2(d-1)$ nonidentity X/I-only and Z/I-only strings, while the mixed row contains the remaining $(d-1)^2$ directions. At $n = 3$, the two values coincide at $A_8 = M_8 = 9$; for $n > 3$, $M_d > A_d$. Consequently,

$$\sup_{\rho,\, P \ne I} \mathrm{Var}_\rho(\widehat{P}) = M_d = \Theta(2^n), \qquad n \ge 3.$$

The same variance identity also covers the two smaller cases. At $n = 1$, all three nonidentity Paulis have $C(P) = 3$. At $n = 2$, $\boldsymbol{c}\cdot\boldsymbol{e}$ is the parity of the number of Y factors: the even- and odd-Y sectors contain nine and six directions with $C(P) = 9$ and 3, respectively. In both cases, the worst-state value is attained at I/d.

The classwise law above contrasts with the weight law of local measurements. For the n-fold product tetrahedral SIC [31] and uniformly randomized local-Pauli shadows [3], the canonical estimators obey

$$\mathbb{E}_\rho[\widehat{P}^2] = 3^{w(P)}, \qquad \mathrm{Var}_\rho(\widehat{P}) = \mathbb{E}_\rho[\widehat{P}^2] - \mathrm{Tr}(\rho P)^2. \tag{30}$$

Their cost depends only on the Pauli weight $w(P)$ and is unchanged when X, Y, and Z are exchanged on a fixed support. The present WH cost depends instead on the

axial–mixed class and not on weight. Thus, for $n > 3$, $X^{\otimes n}$ and $Z^{\otimes n}$ have second moment $A_d = O(1)$ here but 3^n under the local schemes, whereas a weight-one Y_j has second moment $M_d = \Theta(2^n)$ here but only 3 locally. Neither scheme uniformly dominates the other: the fixed WH measurement favors the axial sector across all weights, while the local schemes favor low-weight mixed directions. Appendix C 5 gives a tensor-product derivation of Eq. (30).

Combining the exact moments with Corollary 5 gives, for $n \geq 3$,

$$N_{\mathsf{Ax}} = O\left(\frac{n + \log(1/\delta)}{\epsilon^2}\right),$$
$$N_{\text{all Pauli}} = O\left(\frac{2^n[n + \log(1/\delta)]}{\epsilon^2}\right).$$

The first bound estimates all $2(d-1)$ axial strings, which span every Pauli weight; each target has dimension-independent variance, and the factor n is only logarithmic in the number of targets. The second bound controls all $4^n - 1$ nonidentity Paulis: $M_d = \Theta(d)$ is the mixed-sector worst-state cost and $\log(4^n) = \Theta(n)$ controls the full target set. The same fixed-measurement records support either task, and the queried Paulis may be chosen after acquisition. By contrast, a single Pauli specified in advance can be measured directly in its eigenbasis using $O(\log(1/\delta)/\epsilon^2)$ copies.

For the complete Pauli family, global-Clifford shadows have $\mathbb{E}_\rho[\widehat{P}^2] = d + 1$ for every nonidentity P and hence the same $O(d[n + \log(1/\delta)]/\epsilon^2)$ scaling [3]. A complete stabilizer MUB realizes this channel using the minimum $d + 1$ Clifford bases [60, 61]. The statistical scaling is therefore comparable, but the architectures differ: Clifford shadows sample a global stabilizer basis on each shot, and the MUB implementation selects among $d + 1$ such bases, whereas the present scheme repeats one fixed minimal-IC d^2-outcome POVM and records one $2n$-bit outcome.

The factor d is unavoidable in the single-copy, no-quantum-memory model. Even an adaptive protocol that chooses an arbitrary POVM on each fresh copy requires $\Omega(d/\epsilon^2)$ copies to estimate the complete Pauli family with constant success probability [63, 64]. Thus the fixed-WH bound has the optimal dependence on dimension, apart from the $\Theta(n)$ target-set logarithm in the upper bound. Allowing joint two-copy measurements changes the model: data-dependent multistage protocols can remove the exponential dependence on n, with $O(n/\epsilon^4)$ copies information-theoretically and $O(n \log(n/\epsilon)/\epsilon^4)$ copies for a triply efficient Clifford implementation [64, 65].

These formulas concern individual Pauli estimators. For a Pauli sum $O = \sum_v c_v P_v$, cross-covariances also contribute, so the individual costs cannot simply be added; general observables are covered by Sec. IV. Informational completeness also permits full tomography from the same records.

V. FULL TOMOGRAPHY FROM THE SAME RECORDS

Direct fidelity estimation against a known ideal target [66, 67] and cross-platform overlap estimation using randomized local measurements [68] avoid full reconstruction. For full tomography, sample requirements depend on the reconstruction loss and measurement model [69]; here we evaluate canonical single-copy linear inversion under Hilbert–Schmidt loss.

Because the fixed POVM is informationally complete, the same records also give the unbiased linear estimator in Eq. (9). WH covariance makes all canonical snapshots unitarily equivalent and hence gives an exact finite-sample Frobenius error. We compare it below with uniformly randomized local-Pauli tomography as well as the fixed product SIC.

Corollary 7 (Exact linear-tomography error)**.** *For any density operator ρ and N independent and identically distributed (i.i.d.) outcomes obtained by applying the fixed circuit to identically prepared copies of ρ,*

$$\mathbb{E}\widehat{\rho}_N = \rho, \qquad \mathbb{E}\|\widehat{\rho}_N - \rho\|_F^2 = \frac{H_n - \mathrm{Tr}(\rho^2)}{N}, \tag{31}$$

where the squared snapshot norm is independent of the outcome and, for $n \geq 3$, has the exact value

$$H_n := \mathrm{Tr}(\widehat{\rho}_u^2) = \frac{1}{d}\left(1 + \sum_{v \neq 0} \frac{d}{\lambda_v}\right),$$
$$= \frac{1}{d}\left[1 + 2(d-1)A_d + (d-1)^2 M_d\right]. \tag{32}$$

The spectral floor gives

$$H_n \leq 2(d^2 + d - 1) - \frac{1}{d} = O(d^2). \tag{33}$$

Together with $M_d = \Theta(d)$, the exact expression shows $H_n = \Theta(d^2)$, so the Hilbert–Schmidt MSE is $\Theta(4^n/N)$.

Canonical linear inversion gives

$$H_{\mathrm{SIC}} = H_{\mathrm{Cl}} = d^2 + d - 1, \qquad H_{\mathrm{Pauli}} = H_{\mathrm{prod}} = 5^n, \tag{34}$$

for an ideal global SIC, global Clifford, uniformly randomized local-Pauli acquisition, and the n-fold product tetrahedral SIC, respectively [3, 33, 34]. The last equality follows because both local schemes use one-qubit canonical factors of the form $3\Pi - \mathbb{I}$, whose squared Hilbert–Schmidt norm is 5. Since $H_n < 2H_{\mathrm{SIC}}$ and $H_n < H_{\mathrm{prod}}$ for every $n \geq 3$, the present coefficient stays within a factor of two of the $\Theta(4^n)$ global benchmark and below both local 5^n coefficients. More explicitly,

$$\frac{H_n}{H_{\mathrm{Pauli}}} = \frac{H_n}{H_{\mathrm{prod}}} = \frac{H_n}{5^n} \sim 2\left(\frac{4}{5}\right)^n.$$

Thus full tomography remains exponential, but the canonical sample-coefficient advantage over either local baseline

grows asymptotically as $\frac{1}{2}(5/4)^n$. Table II separates this statistical comparison from measurement-context complexity. The Pauli ensemble has 3^n joint settings and 6^n possible setting–outcome records, although each shot still requires only n local basis choices. The finite-n formulas are proved in Appendix D.

TABLE II. Canonical i.i.d. linear tomography. The coefficient H enters $\mathbb{E}\|\widehat{\rho}_N - \rho\|_F^2 = (H - \mathrm{Tr}\,\rho^2)/N$. A Pauli label includes the sampled setting and its Born outcome; the two fixed minimal-IC measurements return the outcome alone.

acquisition	settings	raw labels	H
random Pauli	3^n random	6^n	5^n
product SIC	1	4^n	5^n
fixed WH	1	4^n	$H_n \sim 2 \cdot 4^n$

n	Pauli / product SIC	present H_n	ideal global SIC
1	5	5	5
2	25	25	19
3	125	71	71
4	625	317.19	271
5	3125	1482.11	1055

Explicitly forming $\widehat{\rho}_N$ requires $\Theta(d^2) = \Theta(4^n)$ storage. Each individual snapshot, however, has a compact algebraic representation that can be contracted directly with structured targets without constructing a dense matrix.

VI. STRUCTURED SNAPSHOTS AND EFFICIENT POSTPROCESSING

Computational-basis sparsity provides one route to efficient qudit-shadow postprocessing when direct matrix-element access and the target trace are available [70].

Equations (14), (16), and (17) turn the algebraic inverse into an efficient per-record rule: every snapshot is a sum of four product-state outer products and the identity, without constructing a dense $d \times d$ matrix.

This representation has two immediate consequences for exact single-record postprocessing. First, the four rank-one terms are outer products between product stabilizer states. For a Pauli string P_v, $\mathrm{Tr}(P_v\widehat{\rho}_u)$ is computed directly from the binary parity in Eq. (27) in $O(n)$ time, and a K-term Pauli sum costs $O(Kn)$. For a stabilizer-state projector, the four rank-one terms require only a constant number of phase-sensitive stabilizer overlaps, while the identity contribution is fixed by the trace. Standard overlap algorithms therefore give the conservative costs $O(n^3)$ and $O(Kn^3)$ for one projector and a K-term stabilizer-projector decomposition, respectively [71, 72].

These Pauli and stabilizer-decomposition target classes also admit exact polynomial-time evaluation for global-Clifford shadows: a single-record snapshot is $(d+1)|s\rangle\langle s| - \mathbb{I}$, where $|s\rangle$ is a stabilizer state, so tableau algorithms apply [3, 71]. The costs listed below are for the present snapshots.

Second, the same representation is a sum of five product operators: the four rank-one outer products and the identity. Each term is a bond-one MPO, which yields the uniform tensor-network guarantee below.

Proposition 8 (Constant-bond snapshot evaluation)**.** *For every $n \geq 3$, the canonical snapshot has an exact MPO representation of bond dimension at most five. If O is given exactly as an open-boundary MPO of bond dimension χ, then the single-record value $\mathrm{Tr}(O\widehat{\rho}_u)$ can be computed exactly using $O(n\chi^2)$ arithmetic operations and $O(\chi)$ additional working memory beyond storage of the MPO tensors.*

Proof. Each of the five product operators is a bond-one MPO. Taking direct sums of their virtual bonds gives an exact MPO of bond dimension at most five. Contracting each term separately with the target MPO from left to right propagates a boundary vector of dimension at most χ. Each site costs $O(\chi^2)$ arithmetic operations and $O(\chi)$ working memory. Summing the five contractions therefore costs $O(n\chi^2)$ operations and $O(\chi)$ additional working memory. □

Table III summarizes both postprocessing routes. In every row the dense canonical snapshot is avoided.

TABLE III. Classical cost per record after receiving the $2n$-bit outcome. The dense canonical snapshot is never constructed.

target representation	per-record cost
Pauli string	$O(n)$
K-term Pauli sum	$O(Kn)$
stabilizer-state projector	$O(n^3)$
K-term stabilizer-projector sum	$O(Kn^3)$
bond-χ MPO	$O(n\chi^2)$

The MPO row has broader scope: it applies to every bond-χ target, including those without a short Pauli or stabilizer decomposition. General global-Clifford snapshots have no analogous uniform guarantee. Computing certain probabilities $p = |\langle\phi_1 \otimes \cdots \otimes \phi_n|G\rangle|^2$, where $|G\rangle$ is a stabilizer graph state, is #P-hard [73]. The product-state projector $O_\phi = \bigotimes_j |\phi_j\rangle\langle\phi_j|$ is a bond-one MPO. Substituting the global-Clifford snapshot $\widehat{\rho}_G = (d+1)|G\rangle\langle G| - \mathbb{I}$ gives $\mathrm{Tr}(O_\phi\widehat{\rho}_G) = (d+1)p - 1$. Thus computing this quantity exactly is #P-hard in the worst case even for bond-one targets. By contrast, Proposition 8 gives a uniform $O(n\chi^2)$ algorithm for every bond-χ MPO target here.

We now give the fixed quantum circuit that generates these records.

VII. FROM THE PAULI–WH ORBIT TO ONE FIXED CIRCUIT

Let S carry the unknown n-qubit state, and let A be an n-qubit ancilla initialized in $|0^n\rangle$. Every shot applies the same premeasurement unitary to SA, followed

by computational-basis measurement of all $2n$ qubits. The unitary has two stages. First, a fiducial-dependent preparation acting only on A maps $|0^n\rangle_A$ to the conjugate fiducial $|\phi_n^*\rangle_A$. Second, a fiducial-independent Bell-basis transform factorizes over the matched pairs (S_j, A_j). With the convention used below, the readouts of S and A are $\boldsymbol{a}$ and $\boldsymbol{b}$, respectively, so the $2n$-bit outcome directly labels the Pauli–WH effect $E^{(n)}_{\boldsymbol{a},\boldsymbol{b}}$. Thus neither the measurement setting nor the circuit changes between shots.

The common Bell construction below applies to an arbitrary normalized fiducial and realizes the corresponding Pauli–WH POVM. The only fiducial-dependent stage is the exact preparation of $|\phi_n^*\rangle$, constructed and costed in Sec. VIII.

A. Generic Bell interface

Let $|\phi\rangle \in \mathbb{C}^d$ be any normalized fiducial. The premeasurement unitary contains the following two ordered blocks.

(i) Fiducial preparation. Choose any unitary satisfying

$$U_{\mathrm{prep}}(\phi^*)\,|0^n\rangle_A = |\phi^*\rangle_A\,, \tag{35}$$

where complex conjugation is taken in the computational basis.

(ii) Bell-basis transform. Independently of $|\phi\rangle$, define

$$U_{\mathrm{Bell}} = \left(\bigotimes_{j=1}^{n} H_{S_j}\right)\left(\prod_{j=1}^{n} \mathrm{CNOT}_{S_j\to A_j}\right), \tag{36}$$

The matched CNOTs act on disjoint pairs. With direct S_j–A_j couplers, they form one parallel layer, followed by one parallel layer of Hadamards. For fixed ϕ, the complete premeasurement unitary is the ordered composition of blocks (i) and (ii):

$$U_{\mathrm{IC}}(\phi) = U_{\mathrm{Bell}}\big[\mathbb{I}_S \otimes U_{\mathrm{prep}}(\phi^*)\big]. \tag{37}$$

For the fiducial family in Eq. (2), we write $U^{(n)}_{\mathrm{IC}} := U_{\mathrm{IC}}(\phi_n)$, as in Fig. 1. The following proposition identifies the POVM induced by computational-basis readout after $U_{\mathrm{IC}}(\phi)$. It specializes the standard fiducial-state-plus-generalized-Bell construction [41, 46] to the tensor-Pauli convention used here.

Proposition 9 (Fixed-circuit realization of a Pauli–WH POVM)**.** *Let $|\phi\rangle \in \mathbb{C}^d$ be normalized. Initialize S in the state ρ and A in $|0^n\rangle$. After applying $U_{\mathrm{IC}}(\phi)$ and measuring all $2n$ qubits in the computational basis, let $\boldsymbol{a}$ and $\boldsymbol{b}$ denote the outcomes on S and A, respectively. Then*

$$p(\boldsymbol{a},\boldsymbol{b}\mid\rho) = \frac{1}{d}\,\mathrm{Tr}\Big[\rho D_{\boldsymbol{a},\boldsymbol{b}}|\phi\rangle\langle\phi|D^\dagger_{\boldsymbol{a},\boldsymbol{b}}\Big]\,. \tag{38}$$

For $\phi = \phi_n$, Eq. (3) reduces this expression to $p(\boldsymbol{a},\boldsymbol{b}\mid\rho) = \mathrm{Tr}[\rho E^{(n)}_{\boldsymbol{a},\boldsymbol{b}}]$. Thus a single fixed premeasurement unitary followed by computational-basis readout implements the corresponding $d^2 = 4^n$-outcome Pauli–WH POVM.

Appendix E explicitly maps the two-register readout $(\boldsymbol{a},\boldsymbol{b})$ to the WH displacement $D_{\boldsymbol{a},\boldsymbol{b}}$ and proves Eq. (38) for an arbitrary input state ρ. Since U_{Bell} is fiducial independent, any exact circuit preparing $|\phi^*\rangle$ realizes the same POVM and reconstruction map; different choices affect only the preparation resources, such as gate count, depth, connectivity, and ancillary workspace.

VIII. EXACT FIDUCIAL PREPARATION AND RESOURCE TRADEOFFS

It remains to prepare the known ancilla state $|\phi_n^*\rangle$. For the two unitary routes, the construction is easiest to read in three layers: (i) encode the two product branches in one root qubit, (ii) expand that one-qubit memory along a path or a tree, and (iii) attach the common Bell analyzer. The first two layers prepare $|\phi_n\rangle$. Complex conjugation only sends $\theta_n \mapsto -\theta_n$, so the same topology and resource counts prepare $|\phi_n^*\rangle$; below we therefore analyze $|\phi_n\rangle$ and postpone the already-fixed Bell block.

There is only one parameterized local rule in the expansion: a two-qubit split turns one logical memory into two smaller memories while preserving the coherent branch label. Each memory occupies one physical output qubit; its label (m) records the number of output factors still encoded, not the number of qubits occupied by the memory. A nearest-neighbor path applies this rule serially, whereas a balanced tree applies disjoint splits in parallel. Figure 2 puts the rule and both schedules side by side. The optional measurement-assisted LCU route uses a different control model and admits both heralded and deterministic adaptive realizations. Table IV compares these resource points; each is followed by the same Bell-analysis block.

A. Exact preparation by binary memory splitting

1. Nearest-neighbor path

The path construction introduces no separate memory qubit. Before step j, qubits $1,\dots,j-1$ already have their final values in each coherent branch, qubit j of A carries the encoding of the two product tails that remain, and all qubits to its right are in $|0\rangle$. The adjacent unitary on $(j, j+1)$ fixes qubit j as the next $|+\rangle$ or $|0\rangle$ output and transfers the shortened encoding to qubit $j+1$, which is also an output qubit of A. Linearity preserves the relative coefficient $e^{i\theta_n}$. Thus all intermediate information stays within A, and no work qubit is used.

Theorem 10 (Ancilla-free nearest-neighbor preparation)**.** *For every $n \geq 2$, there exist an explicitly preparable one-*

qubit state $|\eta_n\rangle$ and explicitly constructible adjacent two-qubit unitaries G_j, acting on qubits $(j, j+1)$, such that

$$G_{n-1}\cdots G_2 G_1\big(|\eta_n\rangle_1 \left|0^{n-1}\right\rangle_{2,\dots,n}\big) = |\phi_n\rangle\,. \tag{39}$$

The circuit uses one single-qubit initialization and $n-1$ adjacent two-qubit gates, with two-qubit depth $n-1$ and no work qubits, measurements, or feedforward. Compiling the initialization and first split jointly gives $1+2(n-2) = 2n-3$ CNOTs and the same CNOT depth.

Proof. Memory encoding and one path step. For $m \geq 1$, define the one-qubit memory states representing the two unfinished m-factor product tails by

$$\begin{aligned} c_m &:= 2^{-m/2}, \qquad r_m := \sqrt{1-c_m^2}, \\ \left|\mu_0^{(m)}\right\rangle &:= |0\rangle\,, \qquad \left|\mu_+^{(m)}\right\rangle := c_m\,|0\rangle + r_m\,|1\rangle\,. \end{aligned} \tag{40}$$

Here m is the number of output factors still to be generated. The two encoding states have overlap c_m, exactly matching the overlap of the product tails $|0^{\otimes m}\rangle$ and $|+^{\otimes m}\rangle$. First prepare on qubit 1

$$|\eta_n\rangle = \frac{\left|\mu_+^{(n)}\right\rangle + e^{i\theta_n}\left|\mu_0^{(n)}\right\rangle}{\sqrt{\mathcal{N}_n}}. \tag{41}$$

By the definition of $\mathcal{N}_n$ in Eq. (2), $|\eta_n\rangle$ is normalized.

For $m \geq 2$, prescribe one shortening step by

$$\begin{aligned} U_m \left|\mu_0^{(m)}\right\rangle |0\rangle &= \left|\mu_0^{(1)}\right\rangle \left|\mu_0^{(m-1)}\right\rangle, \\ U_m \left|\mu_+^{(m)}\right\rangle |0\rangle &= \left|\mu_+^{(1)}\right\rangle \left|\mu_+^{(m-1)}\right\rangle. \end{aligned} \tag{42}$$

Since $\left|\mu_0^{(1)}\right\rangle = |0\rangle$ and $\left|\mu_+^{(1)}\right\rangle = |+\rangle$, this step writes one physical output factor and passes the shortened encoding to the next qubit of A. The input overlap is c_m, and the output overlap is $2^{-1/2}c_{m-1} = c_m$. The two pairs therefore have the same Gram matrix, so the prescribed map is an isometry on their span and extends to a two-qubit unitary. Appendix F gives an explicit unitary completion. At chronological step j, the unfinished tail contains $m = n-j+1$ output factors; define the physical gate on sites $(j, j+1)$ by

$$G_j := U_{n-j+1}^{(j,j+1)}, \qquad j = 1, \dots, n-1.$$

For the path, applying $G_1 = U_n^{(1,2)}$ produces the root state below; we suppress the untouched factor $\left|0^{n-2}\right\rangle$ on qubits $3, \dots, n$:

$$\begin{aligned} \left|\xi_{1,n-1}^{(n)}\right\rangle &:= G_1\big(|\eta_n\rangle_1 |0\rangle_2\big) \\ &= \frac{1}{\sqrt{\mathcal{N}_n}}\Big(|+\rangle_1 \left|\mu_+^{(n-1)}\right\rangle_2 \\ &\quad + e^{i\theta_n}|0\rangle_1 \left|\mu_0^{(n-1)}\right\rangle_2\Big). \end{aligned} \tag{43}$$

The first output factor is now fixed in both coherent branches, while qubit 2 carries the encoding of the remaining $n-1$ factors.

Loop invariant and termination. Iterating the same rule gives, after k gates and for $1 \leq k \leq n-1$,

$$\begin{aligned} |\psi_k\rangle = \frac{1}{\sqrt{\mathcal{N}_n}}\Big(&|+\rangle^{\otimes k} \left|\mu_+^{(n-k)}\right\rangle \\ &+ e^{i\theta_n}|0\rangle^{\otimes k} \left|\mu_0^{(n-k)}\right\rangle\Big) |0\rangle^{\otimes(n-k-1)}\,. \end{aligned} \tag{44}$$

The three displayed factors occupy, from left to right, qubits $1{:}k$, $k+1$, and $k+2{:}n$; the final factor is absent when $k = n-1$. Equation (43) is the case $k = 1$. Applying $G_{k+1} = U_{n-k}^{(k+1,k+2)}$ and Eq. (42) proves the step from k to $k+1$ for $1 \leq k \leq n-2$. For $k = n-1$, $\left|\mu_+^{(1)}\right\rangle = |+\rangle$ and $\left|\mu_0^{(1)}\right\rangle = |0\rangle$. The two completed components are therefore $|+\rangle^{\otimes n}$ and $e^{i\theta_n}|0\rangle^{\otimes n}$, and their coherent sum is Eq. (2), proving Eq. (39). The one-CNOT root block in Appendix F directly prepares Eq. (43) from $|00\rangle$; it replaces the separate initialization and G_1. The same appendix gives a two-CNOT completion of each later split, yielding the CNOT count stated in the theorem. □

2. *Balanced interaction tree*

Here m labels the unresolved block size, whereas $\left|\mu_s^{(m)}\right\rangle$ occupies only one qubit. The left-to-right path repeatedly partitions this block size as $m = 1 + (m-1)$. To expose parallelism, choose instead $m = a + b$ with $a, b \geq 1$ and apply $U_{a,b}$ to the memory and an unused output qubit. The result is two single-qubit memories associated with child blocks of sizes a and b. Explicitly,

$$\begin{aligned} U_{a,b} \left|\mu_0^{(m)}\right\rangle |0\rangle &= \left|\mu_0^{(a)}\right\rangle \left|\mu_0^{(b)}\right\rangle, \\ U_{a,b} \left|\mu_+^{(m)}\right\rangle |0\rangle &= \left|\mu_+^{(a)}\right\rangle \left|\mu_+^{(b)}\right\rangle. \end{aligned} \tag{45}$$

Indeed, the input overlap is c_m, while the output overlap is $c_a c_b = c_{a+b} = c_m$; the two prescribed pairs have the same Gram matrix and therefore define an isometry that extends to a two-qubit unitary. The path is the special choice $(a, b) = (1, m-1)$, which recovers Eq. (42) with $U_m = U_{1,m-1}$. Choosing a and b as nearly equal as possible gives the parallel schedule below.

Theorem 11 (Ancilla-free balanced-tree preparation)**.** *With arbitrary one- and two-qubit gates and balanced-tree connectivity, $|\phi_n\rangle$ can be prepared exactly on its n-qubit register using no work qubits. Initialize one of the n output qubits in $|\eta_n\rangle$ from Eq. (41) and the other $n-1$ output qubits in $|0\rangle$. Whenever a qubit carries a one-qubit encoding for $m > 1$ output factors not yet written, choose*

$$a = \lfloor m/2 \rfloor, \qquad b = \lceil m/2 \rceil. \tag{46}$$

Apply $U_{a,b}$ to that qubit and one output qubit of A still in $|0\rangle$. After the split, the original wire carries the one-qubit

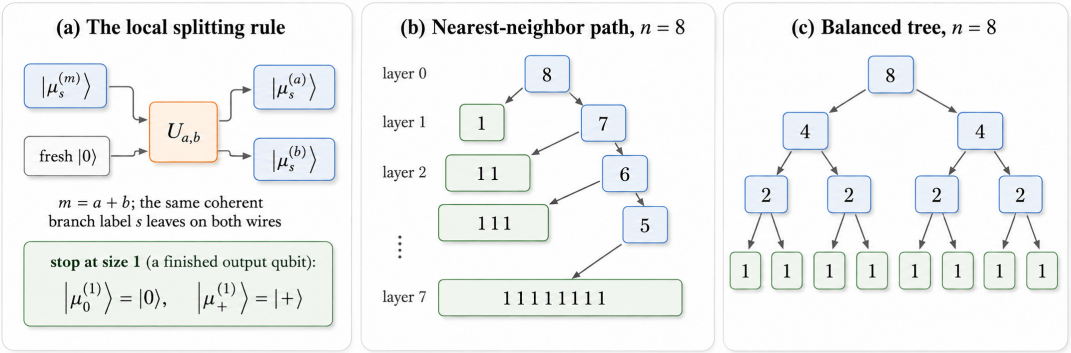

FIG. 2. Binary-memory splitting for $n = 8$. Blue boxes carry one-qubit encodings $\left|\mu_s^{(k)}\right\rangle$; green boxes labeled 1 are completed output qubits. (a) $U_{a,b}$ preserves s while splitting $m = a + b$; "fresh" is an unused output qubit of A. (b) Repeated $(1, m-1)$ splits form a nearest-neighbor path; each green bar collects the outputs completed so far. (c) Balanced splits on disjoint pairs run in parallel. Both schedules use only the n output qubits.

memory for the size-a child block, while the previously unused output wire carries the corresponding memory for the size-b child block. Perform all splits on disjoint pairs in parallel and repeat until every unresolved block has size one. The circuit uses $n-1$ two-qubit gates and has two-qubit depth $\lceil \log_2 n \rceil$.

Proof. Recursive rule and correctness. We prove correctness by strong induction on the number m of output factors still to be generated. Start with one qubit carrying the memory for an unresolved block of size m and $m-1$ output qubits in $|0\rangle$. We claim that the recursive splits implement the two branchwise maps

$$\begin{aligned} \left|\mu_0^{(m)}\right\rangle |0^{m-1}\rangle &\longmapsto |0^m\rangle\,, \\ \left|\mu_+^{(m)}\right\rangle |0^{m-1}\rangle &\longmapsto |+\rangle^{\otimes m}\,. \end{aligned} \tag{47}$$

For $m = 1$, these are just $\left|\mu_0^{(1)}\right\rangle = |0\rangle$ and $\left|\mu_+^{(1)}\right\rangle = |+\rangle$. For $m > 1$, choose a and b from Eq. (46). Equation (45) maps that memory to two single-qubit memories for child blocks of sizes a and b. Assign $a-1$ of the remaining zero qubits to the first child and $b-1$ to the second. The induction hypothesis expands these two disjoint blocks into a and b completed factors, respectively, and the two expansions can run in parallel. This proves both maps above.

Equation (47) describes the result after the entire recursion has terminated. Starting from $|\eta_n\rangle\left|0^{n-1}\right\rangle$, the $\left|\mu_+^{(n)}\right\rangle$ component becomes $|+\rangle^{\otimes n}$, while the $\left|\mu_0^{(n)}\right\rangle$ component becomes $|0^n\rangle$. Linearity preserves their relative coefficient, so the final output is

$$|\eta_n\rangle\left|0^{n-1}\right\rangle \xrightarrow{\text{all recursive splits}} \frac{|+\rangle^{\otimes n} + e^{i\theta_n}|0^n\rangle}{\sqrt{\mathcal{N}_n}} = |\phi_n\rangle\,. \tag{48}$$

Gate count and depth. Each split replaces one encoding by two and uses one previously untouched output qubit. Starting with one unresolved block and ending with n size-one blocks therefore requires $n-1$ splits. All of these qubits belong to A, so no work qubit is used.

Let $L_{\text{tree}}(m)$ be the number of two-qubit layers needed to resolve a block of size m. The two child blocks are disjoint and can be expanded in parallel, so

$$L_{\text{tree}}(1) = 0, \qquad L_{\text{tree}}(m) = 1 + \max\{L_{\text{tree}}(a), L_{\text{tree}}(b)\}, \tag{49}$$

for $m > 1$. Since $b = \lceil m/2 \rceil \geq a$, strong induction gives

$$L_{\text{tree}}(m) = 1 + \lceil \log_2 b \rceil = \lceil \log_2 m \rceil. \tag{50}$$

Here we used $1 + \lceil \log_2 \lceil m/2 \rceil \rceil = \lceil \log_2 m \rceil$, which also holds when m is not a power of two. □

Eight-qubit illustration. To make the recursion concrete, we now spell out the $n = 8$ instance step by step. Label the eight output qubits of A as $1, \ldots, 8$. All eight are present from the start. An "unused" qubit below means one of these output qubits still in $|0\rangle$, not an additional work qubit.

Initialization. This is the $n = 8$ instance of the one-qubit initialization in Eq. (41). Starting from $\left|0^8\right\rangle_A$, prepare qubit 1 in

$$|\eta_8\rangle = \frac{(1/16 + e^{i\theta_8})\,|0\rangle + (\sqrt{255}/16)\,|1\rangle}{\sqrt{\mathcal{N}_8}}.$$

A single-qubit gate suffices because $|\eta_8\rangle$ is normalized. Leave the other seven qubits in $|0\rangle$. The initial state for the recursive splits is therefore

$$\begin{aligned} |\Psi_0\rangle &:= |\eta_8\rangle_1\,|0\rangle_2 \cdots |0\rangle_8 \\ &= \frac{1}{\sqrt{\mathcal{N}_8}}\left(\left|\mu_+^{(8)}\right\rangle_1 + e^{i\theta_8}\left|\mu_0^{(8)}\right\rangle_1\right) \\ &\quad \otimes |0\rangle_2 \cdots |0\rangle_8\,. \end{aligned} \tag{51}$$

Here both $\left|\mu_+^{(8)}\right\rangle$ and $\left|\mu_0^{(8)}\right\rangle$ are *one-qubit* memory states representing unresolved eight-factor product tails; neither is an eight-qubit register.

Layer 1: $8 \to 4+4$. Use qubit 5, which is still in $|0\rangle$, and apply $U_{4,4}$ to qubits $(1,5)$. Equation (45) acts on the two branches as

$$\left|\mu_0^{(8)}\right\rangle_1 |0\rangle_5 \longmapsto \left|\mu_0^{(4)}\right\rangle_1 \left|\mu_0^{(4)}\right\rangle_5,$$
$$\left|\mu_+^{(8)}\right\rangle_1 |0\rangle_5 \longmapsto \left|\mu_+^{(4)}\right\rangle_1 \left|\mu_+^{(4)}\right\rangle_5.$$

The other six output qubits remain in $|0\rangle$ and are suppressed below. Applying these two rules to $|\Psi_0\rangle$ gives

$$|\Psi_1\rangle = \frac{1}{\sqrt{\mathcal{N}_8}}\Big(\left|\mu_+^{(4)}\right\rangle_1 \left|\mu_+^{(4)}\right\rangle_5 + e^{i\theta_8} \left|\mu_0^{(4)}\right\rangle_1 \left|\mu_0^{(4)}\right\rangle_5\Big). \tag{52}$$

Layer 2: each $4 \to 2+2$. Apply $U_{2,2}$ simultaneously to qubit pairs $(1,3)$ and $(5,7)$. Thus each size-four memory is split using a distinct output qubit still in $|0\rangle$, while the other four output qubits remain in $|0\rangle$ and are suppressed. The result is

$$|\Psi_2\rangle = \frac{1}{\sqrt{\mathcal{N}_8}}\Big(\bigotimes_{j\in\{1,3,5,7\}} \left|\mu_+^{(2)}\right\rangle_j + e^{i\theta_8} \bigotimes_{j\in\{1,3,5,7\}} \left|\mu_0^{(2)}\right\rangle_j\Big). \tag{53}$$

Layer 3: each $2 \to 1+1$. Finally apply $U_{1,1}$ in parallel to qubit pairs $(1,2)$, $(3,4)$, $(5,6)$, and $(7,8)$. Every qubit now contains one completed factor:

$$|\Psi_3\rangle = \frac{1}{\sqrt{\mathcal{N}_8}}\left(\bigotimes_{j=1}^{8} \left|\mu_+^{(1)}\right\rangle_j + e^{i\theta_8} \bigotimes_{j=1}^{8} \left|\mu_0^{(1)}\right\rangle_j\right)$$
$$= \frac{|+\rangle^{\otimes 8} + e^{i\theta_8}|0\rangle^{\otimes 8}}{\sqrt{\mathcal{N}_8}} = |\phi_8\rangle. \tag{54}$$

The block sizes have therefore evolved as $8 \to 4+4 \to 2+2+2+2 \to 1+\cdots+1$. Each layer acts linearly on the two coherent terms, so their normalization and relative phase are preserved. The preparation consists of one single-qubit initialization followed by $1+2+4=7$ two-qubit gates in three two-qubit layers. Thus its two-qubit depth is three, and it uses no qubit outside A. This is precisely the $n=8$ instance of the recursive construction proved above.

3. Compilation and common CNOT ledger

The path and balanced-tree constructions above act entirely within the n-qubit output register A. Neither invokes a fanout primitive, and neither uses clean or dirty work qubits beyond A. At each split, the input qubit in $|0\rangle$ is one of the final output qubits of A that has not yet been written; after the split it remains part of the output. Thus the tree hides no ancillary workspace. It remains to compile the abstract two-qubit splits into CNOTs and one-qubit gates.

At the abstract level, either schedule contains $n-1$ split gates. Compiling all of them independently as one-to-two-qubit isometries would use $2(n-1)$ CNOTs [74]. Appendix F gives one explicit two-CNOT template for every $U_{a,b}$, including balanced splits with $a,b>1$. The root is cheaper: before it, the input is one known product state and is not yet correlated with any completed output. We therefore fuse the initialization and first split and compile their fixed two-qubit output state, rather than the full root unitary. Let $(a,b)=(1,n-1)$ for the path and $(\lfloor n/2\rfloor, \lceil n/2\rceil)$ for the balanced tree. In either case, write their joint target as

$$\left|\xi_{a,b}^{(n)}\right\rangle := \frac{\left|\mu_+^{(a)}\right\rangle \left|\mu_+^{(b)}\right\rangle + e^{i\theta_n}|00\rangle}{\sqrt{\mathcal{N}_n}}. \tag{55}$$

Its coefficient-matrix determinant is $e^{i\theta_n} r_a r_b/\mathcal{N}_n \neq 0$, so $\left|\xi_{a,b}^{(n)}\right\rangle$ is entangled. Appendix F constructs its direct preparation from $|00\rangle$ using exactly one CNOT and one-qubit gates. This root block replaces jointly—it does not follow—the initialization and first split. For $n \geq 2$, the remaining $n-2$ nonroot splits use two CNOTs each, giving

$$C_{\text{prep}} = \underbrace{1}_{\substack{\text{initialization and}\\ \text{root split}}} + \underbrace{2(n-2)}_{\substack{n-2 \text{ nonroot splits,}\\ \text{two CNOTs each}}} = 2n-3. \tag{56}$$

For the path, these CNOTs are sequential, so the CNOT depth is $2n-3$. For the balanced tree, let $L=\lceil \log_2 n\rceil$. The joint first block has CNOT depth one, and each of the remaining $L-1$ tree levels becomes two parallel CNOT layers. Its CNOT depth is therefore $1+2(L-1)=2L-1$. For $n=1$, no preparation CNOT is needed.

4. Why the unitary schedules are optimal

Proposition 12 (Optimal two-qubit-gate count)**.** *For $n \geq 2$, every unitary circuit that starts from a product state and prepares $|\phi_n\rangle$ exactly requires at least $n-1$ two-qubit gates, even with any number of additional qubits initialized in a product state. The two-qubit-gate count in Theorem 11 is therefore optimal.*

Proof. Across every nontrivial bipartition $S \mid \bar{S}$, the two branch vectors $|+\rangle_S^{\otimes|S|}, |0\rangle_S^{\otimes|S|}$ are linearly independent, as are their counterparts on $\bar{S}$. Hence $|\phi_n\rangle$ has Schmidt rank two across every such cut. In the interaction graph whose edges are the two-qubit gates, all n output qubits must therefore lie in one connected component; otherwise the output would factor across a nontrivial cut. A connected graph containing n specified vertices has at least $n-1$ edges. □

Proposition 13 (Logarithmic unitary-depth lower bound)**.** *For $n \geq 2$, consider circuits starting from a product state on the output and work registers, with arbitrary one-qubit gates, layers of disjoint two-qubit gates,*

all-to-all connectivity, and any number of work qubits. Exact preparation of $|\phi_n\rangle$ requires two-qubit depth

$$D \geq \frac{1}{2}\log_2 n. \tag{57}$$

The balanced tree is therefore asymptotically depth optimal.

Proof. The key observation is that the target $|\phi_n\rangle$ has nonzero connected Z correlation between every pair of output qubits. To verify this, fix two distinct outputs i and j, let Z_i denote Pauli Z on output i, and set $x_n = 2c_n \cos\theta_n$ and $\nu_n = (1+x_n)/(2+x_n)$. Direct evaluation gives

$$\begin{aligned} \langle Z_i\rangle_{\phi_n} &= \langle Z_i Z_j\rangle_{\phi_n} = \nu_n, \\ \mathrm{Cov}_{\phi_n}(Z_i, Z_j) &= \nu_n(1-\nu_n) > 0. \end{aligned} \tag{58}$$

For $n \geq 2$, the phases in Eq. (2) give $x_n \in [-1/2, 0)$ and hence $\nu_n \in [1/3, 1/2)$. We now turn this all-pairs correlation into a light-cone bound. Let $\mathcal{B}_i$ be the backward input light cone of output i. At depth D, $|\mathcal{B}_i| \leq 2^D$. If $\mathcal{B}_i$ and $\mathcal{B}_j$ were disjoint, the product input would give zero connected correlation; hence $\mathcal{B}_i \cap \mathcal{B}_j \neq \varnothing$ for every pair. Let $\mathcal{F}_a$ denote the forward output light cone of input wire a. Fix i. Every output j shares some ancestor $a \in \mathcal{B}_i$ with i, so $j \in \mathcal{F}_a$. Since $|\mathcal{F}_a| \leq 2^D$, all n outputs obey

$$n \leq \sum_{a\in\mathcal{B}_i} |\mathcal{F}_a| \leq 4^D. \tag{59}$$

This proves Eq. (57). □

5. *Small-system endpoints*

For $n = 1$, the complete tetrahedral-SIC circuit uses one arbitrary preparation gate $U_{\mathrm{prep}}(\phi_1^*)$, one Hadamard, and one CNOT. One CNOT is minimal for a deterministic implementation of a qubit SIC POVM using one ancilla qubit and computational-basis readout of both qubits [32].

For $n = 3$, the ancilla-free CNOT complexity of preparing the Hoggar fiducial is exactly three, assuming arbitrary one-qubit gates and all-to-all CNOT connectivity. Figure 6 gives a three-CNOT construction, and Appendix G proves that no circuit with at most two CNOTs can prepare the state. Adding one Bell-analysis CNOT for each of the three system–ancilla pairs gives a six-CNOT implementation of the three-qubit Hoggar SIC POVM.

B. Optional measurement-assisted LCU

The balanced tree already gives the strongest measurement-free resource point used in this work. Mid-circuit measurement and feedforward have also been used for dynamic generalized measurements and constant-depth long-range entangling gates [45, 75]. Related measurement-assisted constructions prepare arbitrary states at constant quantum depth with exponential auxiliary resources [76], or exploit permutation symmetry [77] and suitable matrix-product-state structure [78, 79]. Here they allow the two product branches to be combined by a short heralded LCU procedure. Following the two-term LCU idea [80], a branch-label qubit coherently selects $H^{\otimes n}$ or $e^{i\theta_n}\mathbb{I}$; interference produces the target in one branch:

$$\frac{H^{\otimes n} + e^{i\theta_n}\mathbb{I}}{2}|0^n\rangle = \frac{\sqrt{\mathcal{N}_n}}{2}|\phi_n\rangle = \sqrt{p_n}\,|\phi_n\rangle\,. \tag{60}$$

Here $p_n = \mathcal{N}_n/4$ is the probability of that branch.

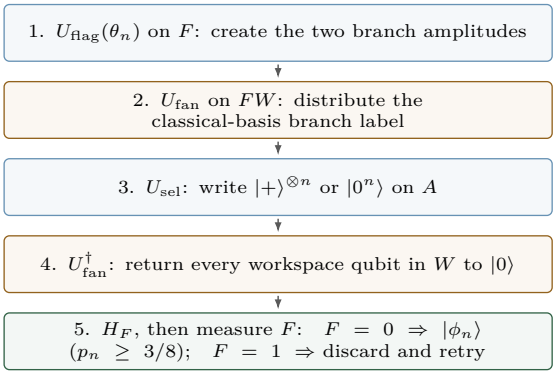

FIG. 3. Coherent reference decomposition of the optional heralded LCU preparation, read from top to bottom. The adaptive realization below compiles steps 2–4 into one exact dynamic common-control primitive. In the coherent decomposition, the fanout workspace carries the branch label only between steps 2 and 4 and is reset before the flag is measured. The unknown state has not yet interacted with A, so failed ancilla preparations can be discarded without consuming a copy of ρ.

For the coherent reference implementation, let A be the n-qubit output register, F a single flag qubit carrying the branch label, and W an $(n-1)$-qubit fanout register. Write $|\mathbf{b}\rangle_W := |b\rangle^{\otimes(n-1)}$ for $b \in \{0,1\}$. For $n = 1$, W is empty and the fanout unitary below is the identity.

Lemma 14 (Constant-success branch synthesis)**.** *For every $n \geq 1$, with arbitrary one-qubit gates and all-to-all CNOT connectivity, there is an explicit unitary V_n on FAW such that*

$$\begin{aligned} V_n|0^{2n}\rangle = \frac{1}{2}\Big[&|0\rangle_F \left(|+\rangle^{\otimes n} + e^{i\theta_n}|0^n\rangle\right)_A \\ &+ |1\rangle_F \left(|+\rangle^{\otimes n} - e^{i\theta_n}|0^n\rangle\right)_A\Big] |\mathbf{0}\rangle_W\,. \end{aligned} \tag{61}$$

Consequently, measuring F yields $F = 0$ and leaves A exactly in $|\phi_n\rangle$ with probability $p_n = \mathcal{N}_n/4 \geq 3/8$; W is reset in either branch. The circuit uses n additional initialized work qubits—the flag qubit F and the $(n-1)$-qubit register W—together with $3n-2$ CNOTs, $O(n)$ one-qubit gates, and CNOT depth $2\lceil\log_2 n\rceil + 1$.

Proof. Figure 3 gives the register-level roadmap for the coherent reference implementation. Initialize FAW in

$|0\rangle_F |0^n\rangle_A |\mathbf{0}\rangle_W$; the five steps below are listed in physical time order.

(1) *Create the branch superposition.* Apply $U_{\text{flag}}(\theta_n)$ to F, where

$$U_{\text{flag}}(\theta) := \text{diag}(1, e^{i\theta})_F H_F.$$

It produces

$$U_{\text{flag}}(\theta_n) |0\rangle_F = \frac{|0\rangle_F + e^{i\theta_n} |1\rangle_F}{\sqrt{2}},$$

with A and W still in $|0^n\rangle_A |\mathbf{0}\rangle_W$.

(2) *Fan out the branch label.* For $b \in \{0,1\}$, the relevant action of the coherent fanout on its initialized workspace is

$$U_{\text{fan}} |b\rangle_F |\mathbf{0}\rangle_W = |b\rangle_F |\mathbf{b}\rangle_W .$$

This is the standard binary-tree cat-state (repetition-code) encoding primitive: it coherently encodes the computational-basis branch label rather than cloning an arbitrary input state. Under all-to-all connectivity, a binary tree of $n-1$ CNOTs and depth $\lceil \log_2 n \rceil$ implements U_{fan} [81], giving

$$\frac{|0\rangle_F |0^n\rangle_A |\mathbf{0}\rangle_W + e^{i\theta_n} |1\rangle_F |0^n\rangle_A |\mathbf{1}\rangle_W}{\sqrt{2}}.$$

(3) *Select the two product branches.* Let F control A_1 and W_j control A_{j+1} for $1 \le j \le n-1$. Denote the parallel product of the n zero-controlled Hadamards by U_{sel}. Since each pair obeys $|0\rangle |0\rangle \mapsto |0\rangle |+\rangle$ and $|1\rangle |0\rangle \mapsto |1\rangle |0\rangle$, this step gives

$$\frac{|0\rangle_F |+\rangle_A^{\otimes n} |\mathbf{0}\rangle_W + e^{i\theta_n} |1\rangle_F |0^n\rangle_A |\mathbf{1}\rangle_W}{\sqrt{2}}.$$

(4) *Erase the copied label.* Applying $U_{\text{fan}}^\dagger$ resets W without undoing the two branches already written to A:

$$\frac{|0\rangle_F |+\rangle_A^{\otimes n} + e^{i\theta_n} |1\rangle_F |0^n\rangle_A}{\sqrt{2}} |\mathbf{0}\rangle_W .$$

(5) *Interfere the branches.* A final H_F interferes the branch amplitudes, routing their sum—the LCU numerator in Eq. (60)—to $F=0$ and their difference to $F=1$.

The complete branch-synthesis unitary is therefore

$$V_n := H_F U_{\text{fan}}^\dagger U_{\text{sel}} U_{\text{fan}} U_{\text{flag}}(\theta_n), \tag{62}$$

where the rightmost factor acts first and identities on untouched registers are implicit. Applying the final H_F to the state in step (4) gives Eq. (61). Its $F=0$ component is $\sqrt{\mathcal{N}_n} |\phi_n\rangle /2$ and therefore has probability

$$p_n = \frac{\mathcal{N}_n}{4} \ge \frac{3}{8}. \tag{63}$$

Explicitly, $p_1 = (3+\sqrt{3})/8$, $p_2 = 3/8$, and $p_n = 1/2 - 2^{-(n+3)/2}$ for $n \ge 3$.

The forward and reverse fanout trees contribute $2(n-1)$ CNOTs and depth $2\lceil \log_2 n \rceil$; the parallel layer U_{sel} contributes n further CNOTs and one layer. Indeed, $H = R_y(\pi/4) Z R_y(-\pi/4)$, so a controlled Hadamard is locally equivalent to a controlled Z, and hence to one CNOT; changing the control from one to zero requires only local X gates. This proves the stated resources. □

1. *Constant-depth heralded realization*

Here adaptive quantum depth counts only the quantum layers between measurement and feedforward rounds, excluding classical processing, reset, and queuing latency. The flag outcome F heralds exact ancilla preparation: failed attempts are discarded and repeated before A interacts with ρ, and only a successful ancilla is coupled to ρ. Because acceptance is independent of ρ, heralding leaves both the logical POVM and its reconstruction map unchanged and cannot bias the Born distribution. The internal measurement outcomes are implementation records, not externally sampled shadow settings.

Proposition 15 (Exact heralded constant adaptive quantum depth)**.** *Our construction uses the one-round dynamic-fanout primitive of Ref. [75] on its alternating 1D layout. In each attempt, conditional on obtaining $F=0$, it prepares $|\phi_n\rangle$ exactly. Each attempt has $O(n)$ quantum size, $O(1)$ adaptive quantum depth, one round of mid-circuit measurements followed by the final measurement of F, and $O(\log n)$ classical parity depth. Repeating with a freshly reinitialized ancilla register until success gives*

$$\mathbb{E}[N_{\text{attempts}}] = p_n^{-1} \le \frac{8}{3}. \tag{64}$$

Thus expected adaptive quantum depth and total gate count are $O(1)$ and $O(n)$, respectively.

In Eq. (62), the only blocks whose depth grows with n are U_{fan} and $U_{\text{fan}}^\dagger$. On the initialized workspace, the former coherently broadcasts the branch bit,

$$(\alpha |0\rangle + \beta |1\rangle)_F \left|0^{n-1}\right\rangle_W \longmapsto \alpha |0\rangle_F \left|0^{n-1}\right\rangle_W + \beta |1\rangle_F \left|1^{n-1}\right\rangle_W ,$$

which is a GHZ-type encoding; the inverse erases the copied label. The measurement-and-feedforward construction of Ref. [75] implements this fanout exactly at constant CNOT depth using measured auxiliaries and Pauli feedforward, while U_{sel} is already one parallel layer. A literal replacement in Eq. (62) would therefore use the dynamic-fanout primitive twice.

For the present Hadamard targets, however, one call suffices. When W starts in $|\mathbf{0}\rangle$, the full sandwich obeys

$$U_{\text{fan}}^\dagger U_{\text{sel}} U_{\text{fan}} |b\rangle_F |\psi\rangle_A |\mathbf{0}\rangle_W = |b\rangle_F \left(H_A^{\otimes n}\right)^{1-b} |\psi\rangle_A |\mathbf{0}\rangle_W , \quad b \in \{0,1\}.$$

TABLE IV. Preparation resources for the two unitary schedules and two adaptive LCU realizations. Extra qubits are counted beyond the n-qubit fiducial register A, so total preparation width is n plus the second-column entry. The heralded value is $n+1$, rather than the coherent reference's n, because dynamic fanout uses n measured auxiliaries in place of the $(n-1)$-qubit workspace W; this extra qubit is a space-for-depth cost. The deterministic adaptive bound is a conservative exact-primitives composition [75, 82, 83]. The unitary counts apply for $n \geq 2$, and the common n-pair Bell-analysis layer is excluded. NN denotes nearest-neighbor.

route and model	extra qubits beyond A	gate-count bound	depth bound	success and measurements
nearest-neighbor path (Thm. 10) measurement-free; NN A rail	0	$n-1$ two-qubit gates; $2n-3$ CNOTs	$n-1$ two-qubit; $2n-3$ CNOT	deterministic; no mid-circuit measurement
balanced tree (Thm. 11) measurement-free; tree connectivity	0	$n-1$ two-qubit gates; $2n-3$ CNOTs	$\lceil \log_2 n \rceil$ two-qubit; $2\lceil \log_2 n \rceil - 1$ CNOT	deterministic; no mid-circuit measurement
heralded LCU (Prop. 15) adaptive; alternating 1D [75]	$n+1$	$3n-1$ CNOTs/attempt; expected at most $\frac{8}{3}(3n-1)$	5 CNOT layers/attempt; $O(1)$ expected adaptive; $O(\log n)$ classical	nondeterministic; $p_n \geq 3/8$; $\mathbb{E}[N_{\text{attempts}}] \leq 8/3$; 2 measurement rounds/attempt
deterministic adaptive LCU (Prop. 16) adaptive; NN grid	$O(n \log n)$	$O(n \log n)$ quantum gates	$O(1)$ adaptive; $O(\log n)$ classical	deterministic; $p_{\text{succ}} = 1$; constant measurement rounds; no retry

Thus it is a zero-controlled $H^{\otimes n}$. To expose its fanout form, set $L = R_y(\pi/4)H$ and $R = HR_y(-\pi/4)$. A Hadamard on A_j controlled by $F = 1$ obeys $C_{F=1}(H_{A_j}) = (I_F \otimes L_{A_j})\text{CNOT}_{F\to A_j}(I_F \otimes R_{A_j})$. Our control is on $F = 0$, which is obtained by flipping F before and after the standard one-controlled operation: $C_{F=0}(U) = X_F C_{F=1}(U) X_F$. Hence

$$C_{F=0}(H_A^{\otimes n}) = X_F(I_F \otimes L_A^{\otimes n}) \times \left(\prod_{j=1}^{n} \text{CNOT}_{F\to A_j}\right)(I_F \otimes R_A^{\otimes n})X_F.$$

The single-qubit gates in each outer layer act in parallel, and the middle product is exactly one common-control fanout. The one-round primitive of Ref. [75] then uses n measured auxiliaries, $3n-1$ CNOTs, and five CNOT layers on its alternating 1D layout. A balanced XOR tree computes the Pauli corrections in $O(\log n)$ classical depth. After feedforward, H_F and the measurement of F herald the LCU branch. Hence one attempt has peak width $2n+1$, one mid-circuit measurement round followed by the final flag measurement, and expected CNOT count $(3n-1)/p_n \leq \frac{8}{3}(3n-1)$. The corrected channel, Eq. (61), and p_n are unchanged.

2. *Deterministic completion*

The heralded procedure succeeds after a random number of attempts, with constant expected adaptive quantum depth. Exact amplitude amplification makes the preparation deterministic. With additional auxiliary qubits, measurement and feedforward also keep the completed preparation at constant adaptive quantum depth.

Proposition 16 (Deterministic preparation at constant adaptive quantum depth)**.** *For every $n \geq 2$, the local alternating quantum–classical computation (LAQCC) model of Ref. [82], with exact arbitrary one-qubit rotations, admits an exact deterministic preparation of $|\phi_n\rangle$ with $O(1)$ adaptive quantum depth. The construction uses $O(n \log n)$ initialized auxiliary qubits beyond A and $O(n \log n)$ quantum gates, a constant number of measurement rounds, and $O(\log n)$ classical feedforward depth.*

The construction combines the two LCU branches coherently by one exact amplification step. Its branch-control operations and all-zero selective phase admit constant-depth implementations using measurement and feedforward [75, 82, 83]. Appendix H proves the amplification identity and gives the selective-phase construction and resource count. Preparation finishes before A contacts S, so the same Bell interface and reconstruction apply. Relative to the measurement-free balanced tree, this route trades additional auxiliary qubits and classical feedback for lower quantum depth; the two resource points are compared in Table IV.

C. Comparison and end-to-end resources

Table IV places four resource points on one ledger: exact two-qubit and CNOT resources for the strict-unitary schedules, asymptotic resources for the heralded and deterministic adaptive LCU realizations, and the universal Bell readout kept outside every preparation count.

For $n \geq 2$ and before topology-dependent routing, either unitary preparation followed by the common Bell layer contains $n-1$ preparation isometries and n Bell-analysis CNOTs, hence $2n-1$ two-qubit gates in total. At readout, the fiducial-assisted analyzer acts on exactly the $2n$ measured qubits SA; no additional online workspace remains. The optimized root isometry uses one CNOT, the remaining $n-2$ preparation isometries use two each, and the Bell layer contributes n. The resulting pre-routing

count is therefore

$$C_{\text{full}} = \underbrace{1 + 2(n-2)}_{\text{preparation}} + \underbrace{n}_{\text{Bell}} = 3n - 3 \tag{65}$$

CNOTs. With direct matched S_j–A_j couplers, the CNOT depth is $2n-2$ for the path and $2\lceil \log_2 n \rceil$ for the tree. At $n = 1$, the complete tetrahedral-SIC measurement uses one Bell-analysis CNOT.

The two schedules expose a connectivity tradeoff without a width tradeoff: the path uses only neighboring couplings within A, whereas the tree requires long-range pairs. On a single interleaved line, additional routing may still be needed to combine either preparation with the matched Bell layer. The adaptive implementations use preparation measurements internally, but only an exactly prepared ancilla is coupled to the unknown state; any transport needed for that coupling lies outside the preparation-only resources in Table IV. Every implementation yields the same 4^n external outcomes $(\boldsymbol{a}, \boldsymbol{b})$ and uses the same reconstruction map.

IX. FROM-CIRCUIT NUMERICAL VALIDATION AND FINITE-SHOT BENCHMARKS

The numerics ask three questions not assumed by the analytic derivation: whether coherence completes the frame, whether raw records show the predicted finite-sample behavior, and whether the reconstructed costs exhibit the class–weight tradeoff. Panels (a) and (c) are deterministic frame diagnostics; panel (b) continues through Born sampling and canonical sample means. We choose $n = 4$, the first anisotropic case after the $n = 3$ Hoggar SIC and a size at which direct dense frame inversion remains practical.

Because every route in Table IV implements the same POVM, it suffices to simulate one. We independently compile $|\phi_4^*\rangle$ as a nearest-neighbor bond-dimension-two circuit on the four qubits of A, with no work qubits, and append the Bell layer. For each computational-basis input $|x\rangle$, the circuit gives

$$A_{u,x} = \langle u| U_{\text{IC}} |x, 0^n\rangle .$$

The row $A_u := (A_{u,x})_x$ determines $E_u = A_u^\dagger A_u$ and hence $\mathcal{M}_{\text{num}}(A) = \sum_u \text{Tr}(E_u A) E_u / \text{Tr}\, E_u$. We invert this frame numerically and apply the same procedure to independently constructed tensor-SIC and complete-MUB effects. The MUB has the same depolarizing frame, and hence the same Haar benchmark below, as uniform global Clifford [31, 60, 61]. No analytic WH spectrum, snapshot, or variance formula enters a numerical marker or RMSE band. The ideal global SIC is only a tight-frame benchmark; no dimension-16 SIC fiducial or circuit is inserted [33, 34].

Panel (a) repeats the reconstruction for $|0^n\rangle$, $|+^n\rangle$, their equal-weight incoherent mixture, and the coherent fiducial. Their ranks are $16, 16, 31, 256$: the mixture retains only the two axial sectors, whereas the coherent cross terms supply the 225 missing mixed directions. These are diagnostic controls, not alternative implementations.

Panel (b) fixes the dense observable before sampling. The same 256 independent Haar states are used for all measurements, with a separate 100,000-outcome Born stream for each state and measurement; each plotted N is a nested prefix. Markers are RMSEs over these 256 end-to-end trials, and bands are pointwise percentile-bootstrap intervals from 1000 resamples. Dashed curves use the numerically reconstructed maximally-mixed second moment and $\mathbb{E}_\psi \langle\psi|O|\psi\rangle^2 = \text{Tr}(O^2)/[d(d+1)]$, not a fit to the markers. The complete-MUB variance is $18/17$, shared here by uniform global Clifford and an ideal global SIC. The raw-record RMSE is consistent with $N^{-1/2}$ scaling.

Panel (c) reads the single-shot moments directly from the numerical frames at I/d. For all 30 axial strings, the present value $A_{16} \simeq 6.49$ is weight independent, compared with 3^w for the tensor SIC and 17 for the complete MUB. All 225 mixed WH strings give $M_{16} = (\sqrt{32}-1)^2 \simeq 21.69$, including each weight-one Y_j. Uniform global Clifford and an ideal global SIC give 17 at I/d. Thus the present measurement beats the tensor SIC from weight two onward in the axial sector, but pays a larger mixed-sector cost; this is a classwise tradeoff, not uniform dominance.

These comparisons place the present measurement on a broader hardware–statistical frontier. It is attractive when shot-to-shot setting control is costly and the target family contains axial or structured global observables. Product SICs remain natural for low-weight mixed targets, whereas global-Clifford shadows provide the stronger dimension-independent state-uniform guarantee at the price of a setting ensemble. Locally entangled shadows and target-set-specific strategies under bounded interaction range occupy other points on this frontier [9, 10].

X. DISCUSSION AND OUTLOOK

Precompiled analyzer and system-facing interface. In the present construction, no shot-dependent measurement setting is applied to the unknown system S: every copy encounters the same system-facing analyzer, and the Born outcome itself labels the shadow snapshot. The fiducial register A is freshly prepared before interacting with S, with all fiducial-dependent processing confined to the analyzer side. Once A is ready, each system qubit S_j undergoes only the matched Bell-analysis operations $\text{CNOT}_{S_j \to A_j}$ and H_{S_j}, followed by terminal Z-basis measurement. With direct matched couplers and ignoring routing overhead, the system-facing entangling depth is one, and no S–S gate is required. The nonlocal structure of the POVM is supplied by the correlations prepared in A together with the fixed pairwise Bell interface, rather than by a global entangling circuit acting directly on S.

Randomized local-Pauli and global-Clifford shadows sample from large, exponentially growing measurement

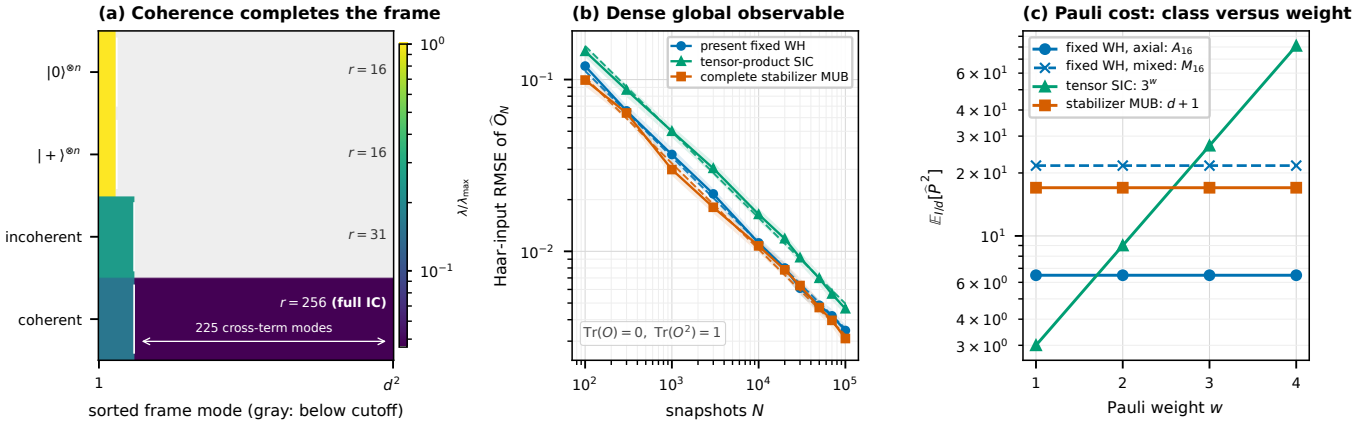

FIG. 4. From-circuit numerical validation at $n = 4$. (a) Common-normalized numerical frame spectra for the two product branches, their incoherent mixture, and the coherent fiducial. The rank ladder $16, 16, 31, 256$ shows that coherent cross terms add 225 directions and complete the frame; gray entries lie below the $10^{-10}\lambda_{\max}$ rank cutoff. (b) Haar-input root-mean-square error (RMSE) for one fixed centered dense target with $\mathrm{Tr}(O^2) = 1$. The three measurements use the same 256 Haar-random input states and an independent outcome stream for every state and measurement. Markers are empirical RMSEs of the canonical sample means, bands are pointwise 95% percentile-bootstrap intervals over the 256 input–record trials, and dashed curves are numerical-frame predictions rather than fits. For the complete stabilizer MUB, this Haar prediction equals those of uniform global-Clifford shadows and an ideal global SIC. (c) Maximally-mixed single-shot second moments grouped by Pauli weight. Filled markers summarize the weight-binned means over all 30 axial strings for the present WH measurement, tensor SIC, and complete MUB; blue crosses give the corresponding means over all 225 mixed WH strings. The dashed blue line is the independent closed-form value $M_{16} = (\sqrt{32} - 1)^2$, while the axial WH markers equal A_{16}. Uniform global Clifford and an ideal global SIC share the complete-MUB value $d + 1$ at I/d.

ensembles. Local Pauli measurements have 3^n possible settings, each requiring only local basis changes. Even when each setting is shallow, a long acquisition may invoke many distinct physical control sequences; global-Clifford settings generally also involve multiqubit gates on S [3]. Precomputing the setting labels does not remove the need to execute the corresponding physical configurations, although caching and batching may reduce the associated overhead. Repeating the fixed system-facing interface eliminates shot-dependent reconfiguration of the system-facing measurement setting during acquisition, at the cost of an additional n-qubit analyzer register and per-shot analyzer preparation.

The analyzer architecture and system-facing control configuration remain fixed across shots, while A is reset and freshly prepared before each interaction with S. A simple platform-dependent measure of the operational cost is the time required to reach a target precision,

$$T_\epsilon^{(x)} \simeq N_\epsilon^{(x)} \left(\tau_{\mathrm{shot}}^{(x)} + \bar{\tau}_{\mathrm{update}}^{(x)}\right) + T_{\mathrm{cal}}^{(x)}, \qquad x \in \{\mathrm{rand}, \mathrm{fix}\}.$$

Here N_ϵ is the required number of shots, $\bar{\tau}_{\mathrm{update}}$ denotes the average setting-update overhead, and $T_{\mathrm{cal}}^{(x)}$ accounts for the calibration cost of the corresponding acquisition scheme. For the fixed analyzer, $\bar{\tau}_{\mathrm{update}}^{(\mathrm{fix})} = 0$ after initialization, while $\tau_{\mathrm{shot}}^{(\mathrm{fix})}$ includes fresh preparation of A, the fixed Bell-analysis layer, and readout.

Small-system tomography. The same fixed circuit provides informationally complete tomography with a single measurement setting, replacing the 3^n local-Pauli settings conventionally used for n-qubit tomography while retaining the minimal $d^2 = 4^n$ outcomes of a rank-one IC measurement. Beyond this reduction in setting complexity, the Hilbert–Schmidt MSE coefficient is also exponentially smaller under the canonical i.i.d. linear inversion of Sec. V. Randomized local-Pauli tomography and a fixed tensor-product qubit SIC measurement both have Hilbert–Schmidt MSE coefficient 5^n, whereas the present measurement has coefficient H_n, with

$$H_n < 5^n \quad (n \geq 3), \qquad H_n \sim 2 \cdot 4^n.$$

Equivalently, $H_n/5^n \sim 2(4/5)^n$, so the sample requirement at fixed Hilbert–Schmidt MSE is smaller by an asymptotic factor of order $(5/4)^n$ than for these two benchmarks. For $n = 3$, the Hoggar SIC gives $H_3 = 71$, compared with $5^3 = 125$ for the local-Pauli and product-SIC benchmarks. Table II summarizes the finite-size comparison.

Tensor-product qubit SICs already provide scalable fixed minimal-IC measurements, but their Pauli statistics mirror those of randomized local-Pauli shadows: the second moment grows as 3^w with operator weight w. The present correlated measurement instead keeps axial-Pauli costs bounded at every weight and, for arbitrary Hermitian targets, has the same dimension-independent scaling as global-Clifford shadows in the Haar-averaged sense. These gains require correlated fiducial preparation and come with larger costs for low-weight mixed-Pauli targets.

From minimal IC to shallow SIC measurements. A

SIC provides a rank-one minimal-IC benchmark with equal pairwise overlaps between distinct measurement projectors, but its existence has not been established in every power-of-two dimension [37, 38]. Existence and implementation are separate questions: knowing a SIC fiducial does not by itself provide a shallow circuit for preparing it on an n-qubit device.

The present construction gives an explicit minimal-IC measurement for every n, with shallow fiducial preparation and exact SIC symmetry at $n = 1, 3$. For other n, the projector overlaps are not all equal, but the statistical bounds established above still apply. An interesting question is whether WH-covariant measurements with SIC symmetry, or uniformly near-SIC projector overlaps and comparable statistical performance, can be realized for all n with shallow qubit-compatible preparation. This asks for more than SIC existence alone.

Beyond minimality. For general targets, the present measurement has a dimension-independent Haar-input coefficient, matching global-Clifford shadow scaling, but an $O(d)$ state-uniform coefficient. Relaxing minimality could improve the latter while retaining a fixed analyzer. For rank-one IC measurements, Ref. [62] shows that a dimension-independent uniform raw-second-moment bound for all traceless, Hilbert–Schmidt-normalized observables requires $\Omega(d^3)$ outcomes, and achieves this scaling using a collection of $\Theta(d^2)$ measurement bases.

Here we consider a different implementation question: whether the same outcome scaling can be realized within the WH–Bell architecture. A natural d^3-outcome candidate uses d WH fiducials coherently indexed by an n-qubit label register, together with the n-qubit fiducial and system registers, giving a $3n$-qubit fixed analyzer. If the resulting combined orbit forms an equal-weight complex projective 3-design, the measurement yields a dimension-independent, state-uniform variance bound for Hilbert–Schmidt-normalized targets [3]. Whether such an all-n WH family exists, and whether its coherent fiducial preparation can be compiled into a shallow $3n$-qubit circuit with computationally efficient reconstruction, are interesting questions for further study.

These results highlight measurement-setting control as an experimental resource complementary to conventional metrics such as circuit width, gate count, depth, and sample complexity. Randomized protocols realize part of their measurement complexity through shot-dependent setting selection and implementation, whereas a fixed analyzer shifts part of that complexity into precompiled measurement geometry, analyzer-state preparation, and a fixed system-facing interface reused across shots. More broadly, this raises the question of where the complexity of quantum measurement should reside. How much randomness and control must be supplied anew on each shot, and how much can instead be compiled once into the geometry and circuitry of the measuring device itself?

ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China under Grants No. 42330707 and No. 42530108.

Appendix A: Normalization of the WH orbit

For an n-qubit system, $d = 2^n$. This appendix proves the WH-orbit normalization for an arbitrary normalized fiducial $|\phi\rangle$; the fixed POVM in Eq. (3) is obtained by setting $|\phi\rangle = |\phi_n\rangle$. We first make the two kinds of labels explicit. Write

$$u = (\boldsymbol{a}, \boldsymbol{b}), \qquad v = (\boldsymbol{c}, \boldsymbol{e}),$$

where all four bold symbols belong to $\mathbb{F}_2^n$. The label u identifies an orbit element, so we abbreviate $D_u = D_{\boldsymbol{a},\boldsymbol{b}}$, $\Pi_u = \Pi_{\boldsymbol{a},\boldsymbol{b}}$, and $E_u = E_{\boldsymbol{a},\boldsymbol{b}}$. The independent label v will identify a Pauli-basis element. Each label has $d^2 = 4^n$ possible values.

Every effect $E_u = d^{-1}\Pi_u$ is positive because Π_u is a rank-one projector. To show that the effects form a POVM, it therefore remains only to prove

$$\sum_u E_u = \mathbb{I}.$$

Step 1: Expand an operator in the Pauli basis. Let $P_v = P_{\boldsymbol{c},\boldsymbol{e}}$ be the Hermitian tensor Pauli proportional to $Z(\boldsymbol{c})X(\boldsymbol{e})$. Its irrelevant phase is chosen so that P_v is Hermitian. The d^2 operators $\{P_v\}$ form an orthogonal basis, with $P_0 = \mathbb{I}$ and

$$\mathrm{Tr}(P_v P_w) = d\delta_{vw}.$$

Every operator A can therefore be written as

$$A = \frac{1}{d}\sum_v \mathrm{Tr}(P_v A) P_v. \tag{A1}$$

Step 2: Average one Pauli component over all displacements. Recall that $u = (\boldsymbol{a}, \boldsymbol{b})$ labels the conjugating displacement and $v = (\boldsymbol{c}, \boldsymbol{e})$ labels the Pauli being conjugated. The single-qubit relation $ZX = -XZ$ gives

$$D_u P_v D_u^\dagger = (-1)^{\boldsymbol{a}\cdot\boldsymbol{e}+\boldsymbol{b}\cdot\boldsymbol{c}} P_v =: s_v(u) P_v.$$

Thus $s_v(u) = +1$ when D_u and P_v commute and $s_v(u) = -1$ when they anticommute. The notation $s_v(u)$ is only shorthand for the explicit binary phase displayed above.

For a fixed $\boldsymbol{e}$, the sum $\sum_{\boldsymbol{a}}(-1)^{\boldsymbol{a}\cdot\boldsymbol{e}}$ equals d if $\boldsymbol{e} = \boldsymbol{0}$. If $\boldsymbol{e} \neq \boldsymbol{0}$, exactly half of the bit strings $\boldsymbol{a}$ give $\boldsymbol{a}\cdot\boldsymbol{e} = 0$ and half give $\boldsymbol{a}\cdot\boldsymbol{e} = 1$, so the signs cancel and the sum is zero. The same argument applies to the sum over $\boldsymbol{b}$. Consequently,

$$\begin{aligned}\sum_u s_v(u) &= \sum_{\boldsymbol{a},\boldsymbol{b}} (-1)^{\boldsymbol{a}\cdot\boldsymbol{e}+\boldsymbol{b}\cdot\boldsymbol{c}} \\ &= \left[\sum_{\boldsymbol{a}} (-1)^{\boldsymbol{a}\cdot\boldsymbol{e}}\right]\left[\sum_{\boldsymbol{b}} (-1)^{\boldsymbol{b}\cdot\boldsymbol{c}}\right] \\ &= d^2\delta_{v0}.\end{aligned} \tag{A2}$$

Here $\delta_{v0} = 1$ exactly when $\boldsymbol{c} = \boldsymbol{e} = \boldsymbol{0}$. Thus every nonidentity Pauli component cancels in the displacement average; only $P_0 = \mathbb{I}$ survives.

Step 3: Average the full operator. Inside each conjugated operator $D_u A D_u^\dagger$, replace A by its expansion in Eq. (A1). Then apply Eq. (A2):

$$\begin{aligned}\frac{1}{d^2}\sum_u D_u A D_u^\dagger &= \frac{1}{d^3}\sum_v \mathrm{Tr}(P_v A) P_v \sum_u s_v(u) \\ &= \frac{\mathrm{Tr}\, A}{d}\mathbb{I}.\end{aligned} \tag{A3}$$

The last equality keeps only $v = 0$, for which $P_0 = \mathbb{I}$ and $\mathrm{Tr}(P_0 A) = \mathrm{Tr}\, A$.

Step 4: Apply the twirl to the fiducial projector. Take $A = |\phi\rangle\langle\phi|$. Because $|\phi\rangle$ is normalized, $\mathrm{Tr}\, A = 1$, and its conjugates are precisely the orbit projectors $\Pi_u = D_u|\phi\rangle\langle\phi|D_u^\dagger$. Equation (A3) therefore gives

$$\frac{1}{d^2}\sum_u \Pi_u = \frac{1}{d}\mathbb{I}, \qquad \sum_u \Pi_u = d\mathbb{I}, \qquad \sum_u E_u = \mathbb{I}. \tag{A4}$$

Thus the factor $1/d$ is exactly the physical POVM normalization. This argument uses only normalization of $|\phi\rangle$ and Pauli covariance; it does not by itself imply informational completeness. For the family generated by $|\phi_n\rangle$, Ref. [47] establishes both informational completeness and uniform spectral stability.

Appendix B: Measurement-channel and reconstruction proofs

1. Pauli diagonalization and covariance

We first prove Theorem 1 and the covariance formulas in Eqs. (13) and (14). We use the conjugation sign $s_v(u) \in \{\pm 1\}$ introduced in Appendix A, so that $D_u P_v D_u^\dagger = s_v(u) P_v$. From the explicit sign formula there, the exponents add modulo two, so $s_v(u)s_w(u) = s_{v\oplus w}(u)$. Equation (A2) then gives

$$\sum_u s_v(u) s_w(u) = d^2\delta_{v,w}. \tag{B1}$$

By the Pauli expansion in Eq. (A1) and the definition of χ_v,

$$\begin{aligned}\Pi_0 &= \frac{1}{d}\sum_v \mathrm{Tr}(P_v\Pi_0) P_v = \frac{1}{d}\sum_v \chi_v P_v, \\ \Pi_u &= D_u \Pi_0 D_u^\dagger = \frac{1}{d}\sum_v s_v(u)\chi_v P_v.\end{aligned} \tag{B2}$$

Taking the Hilbert–Schmidt inner product with P_w and using $\mathrm{Tr}(P_v P_w) = d\delta_{v,w}$ then gives

$$\mathrm{Tr}(\Pi_u P_w) = s_w(u)\chi_w. \tag{B3}$$

Because P_v is Hermitian, $\chi_v = \mathrm{Tr}(\Pi_0 P_v)$ is real. Substituting Eqs. (B2) and (B3) directly into the definition of

$\mathcal{M}_{\phi_n}$, and using Eq. (B1) gives

$$\begin{aligned}\mathcal{M}_{\phi_n}(P_v) &= \frac{1}{d}\sum_u \mathrm{Tr}(\Pi_u P_v)\Pi_u \\ &= \frac{\chi_v}{d^2}\sum_{u,w} s_v(u)s_w(u)\chi_w P_w = \chi_v^2 P_v = |\chi_v|^2 P_v.\end{aligned} \tag{B4}$$

Writing $A = \sum_v a_v P_v$, with $a_v = \mathrm{Tr}(P_v A)/d$, Eq. (B4) acts on each Pauli coefficient separately:

$$\begin{aligned}\mathcal{M}_{\phi_n}(A) &= \sum_v |\chi_v|^2 a_v P_v, \\ \mathcal{M}_{\phi_n}^{-1}(A) &= \sum_v \frac{a_v}{|\chi_v|^2} P_v.\end{aligned}$$

These are the two formulas in Eq. (12). The inverse exists exactly when every χ_v is nonzero; equivalently, no Pauli direction is erased, and the orbit is IC.

Assume henceforth that every χ_v is nonzero. For $A = \Pi_0$, Eq. (B2) shows that the Pauli coefficient is $a_v = \chi_v/d$. Since χ_v is real, $\chi_v/|\chi_v|^2 = 1/\chi_v$, including when $\chi_v < 0$. Substituting these coefficients into the inverse gives

$$\begin{aligned}F_n = \mathcal{M}_{\phi_n}^{-1}(\Pi_0) &= \frac{1}{d}\sum_v \frac{\chi_v}{|\chi_v|^2} P_v \\ &= \frac{1}{d}\sum_v \frac{1}{\chi_v} P_v.\end{aligned}$$

It remains to pass from the reference outcome to an arbitrary u. For $A = \sum_v a_v P_v$, Pauli conjugation gives $D_u A D_u^\dagger = \sum_v s_v(u) a_v P_v$. Applying the coefficientwise inverse to this expansion proves covariance:

$$\begin{aligned}\mathcal{M}_{\phi_n}^{-1}(D_u A D_u^\dagger) &= \sum_v \frac{s_v(u)a_v}{|\chi_v|^2} P_v \\ &= D_u \mathcal{M}_{\phi_n}^{-1}(A) D_u^\dagger.\end{aligned} \tag{B5}$$

Finally, the orbit definition gives $\Pi_u = D_u \Pi_0 D_u^\dagger$. Applying Eq. (B5) with $A = \Pi_0$ yields

$$\begin{aligned}\widehat{\rho}_u &= \mathcal{M}_{\phi_n}^{-1}(\Pi_u) \\ &= \mathcal{M}_{\phi_n}^{-1}(D_u \Pi_0 D_u^\dagger) \\ &= D_u \mathcal{M}_{\phi_n}^{-1}(\Pi_0) D_u^\dagger = D_u F_n D_u^\dagger.\end{aligned}$$

This proves Eqs. (13) and (14).

2. Fiducial overlaps and channel eigenvalues

It remains to evaluate these overlaps. Write $P_{\boldsymbol{c},\boldsymbol{e}} = \gamma_{\boldsymbol{c},\boldsymbol{e}} D_{\boldsymbol{c},\boldsymbol{e}}$, where $D_{\boldsymbol{c},\boldsymbol{e}} = Z(\boldsymbol{c})X(\boldsymbol{e})$ is phase free and $|\gamma_{\boldsymbol{c},\boldsymbol{e}}| = 1$. The corresponding phase-free overlap $\widetilde{\chi}_{\boldsymbol{c},\boldsymbol{e}} := \langle\phi_n|D_{\boldsymbol{c},\boldsymbol{e}}|\phi_n\rangle$ satisfies $\chi_{\boldsymbol{c},\boldsymbol{e}} = \gamma_{\boldsymbol{c},\boldsymbol{e}}\widetilde{\chi}_{\boldsymbol{c},\boldsymbol{e}}$, so $|\chi_{\boldsymbol{c},\boldsymbol{e}}|^2 = |\widetilde{\chi}_{\boldsymbol{c},\boldsymbol{e}}|^2$. With $z_n = e^{i\theta_n}$, direct expansion of Eq. (2) gives

$$\begin{aligned}\widetilde{\chi}_{\boldsymbol{c},\boldsymbol{e}} &= \langle\phi_n|D_{\boldsymbol{c},\boldsymbol{e}}|\phi_n\rangle \\ &= \frac{\delta_{\boldsymbol{c},\boldsymbol{0}} + \delta_{\boldsymbol{e},\boldsymbol{0}} + d^{-1/2}[z_n(-1)^{\boldsymbol{c}\cdot\boldsymbol{e}} + z_n^*]}{\mathcal{N}_n}.\end{aligned} \tag{B6}$$

For $n = 1$, direct substitution gives $|\widetilde{\chi}_v|^2 = 1/3$ for all $v \neq 0$. Hence $\mathcal{M}_{\phi_1}$ is the identity on the scalar component and multiplies the traceless component by $1/3$. Therefore

$$\begin{aligned}\mathcal{M}_{\phi_1}(A) &= \frac{A + \mathrm{Tr}(A)\mathbb{I}}{3}, \\ \mathcal{M}_{\phi_1}^{-1}(A) &= 3A - \mathrm{Tr}(A)\mathbb{I}.\end{aligned} \tag{B7}$$

Taking $A = \Pi_0$ gives $F_1 = 3\Pi_0 - \mathbb{I}$.

For $n = 2$, Eq. (B6) gives

$$|\widetilde{\chi}_{\boldsymbol{c},\boldsymbol{e}}|^2 = \begin{cases} 1/9, & (\boldsymbol{c},\boldsymbol{e}) \neq (\boldsymbol{0},\boldsymbol{0}),\ \boldsymbol{c}\cdot\boldsymbol{e} = 0, \\ 1/3, & \boldsymbol{c}\cdot\boldsymbol{e} = 1. \end{cases} \tag{B8}$$

The Hermitian Pauli convention satisfies $P_{\boldsymbol{c},\boldsymbol{e}}^{\mathsf{T}} = (-1)^{\boldsymbol{c}\cdot\boldsymbol{e}} P_{\boldsymbol{c},\boldsymbol{e}}$. Thus the symmetric traceless sector has channel eigenvalue $1/9$, while the antisymmetric sector has eigenvalue $1/3$. Writing

$$A_\pm = \frac{A \pm A^{\mathsf{T}}}{2}, \qquad A_0 = \frac{\mathrm{Tr}(A)}{4}\mathbb{I},$$

gives

$$\begin{aligned}\mathcal{M}_{\phi_2}(A) &= \frac{2A - A^{\mathsf{T}} + 2\,\mathrm{Tr}(A)\mathbb{I}}{9}, \\ \mathcal{M}_{\phi_2}^{-1}(A) &= 6A + 3A^{\mathsf{T}} - 2\,\mathrm{Tr}(A)\mathbb{I}.\end{aligned} \tag{B9}$$

Taking $A = \Pi_0$ gives $F_2 = 6\Pi_0 + 3\Pi_0^{\mathsf{T}} - 2\mathbb{I}$.

For $n \geq 3$, we use the axial and mixed sets defined in Eq. (18). For any operator A, define its scalar, axial, and mixed Pauli components by

$$A_0 := \frac{\mathrm{Tr}\,A}{d}\mathbb{I}, \qquad c_v(A) := \frac{1}{d}\mathrm{Tr}(P_v A), \tag{B10}$$

$$A_{\mathsf{Ax}} := \sum_{v\in\mathsf{Ax}} c_v(A) P_v, \qquad A_{\mathsf{Mix}} := A - A_0 - A_{\mathsf{Ax}}. \tag{B11}$$

To evaluate the channel on these sectors, set

$$\begin{aligned}&r_d := \sqrt{\frac{2}{d}}, \quad \mathcal{N}_n = 2 - r_d, \quad \kappa_d := 1 - r_d, \\ &\alpha_d := \frac{\kappa_d^2}{\mathcal{N}_n^2}, \quad \beta_d := \frac{2}{d\mathcal{N}_n^2}.\end{aligned} \tag{B12}$$

Now $z_n = e^{3\pi i/4}$. On Ax, exactly one Kronecker delta in Eq. (B6) equals one and $\boldsymbol{c}\cdot\boldsymbol{e} = 0$, giving $\widetilde{\chi}_v = (1 - r_d)/\mathcal{N}_n$. On Mix, both deltas vanish, and the remaining numerator has modulus r_d for either value of $\boldsymbol{c}\cdot\boldsymbol{e}$. Hence the channel eigenvalue is 1 in the identity direction, α_d on Ax, and β_d on Mix. Grouping the Pauli expansion into these three sectors gives

$$\begin{aligned}\mathcal{M}_{\phi_n}(A) &= A_0 + \alpha_d A_{\mathsf{Ax}} + \beta_d A_{\mathsf{Mix}}, \\ \mathcal{M}_{\phi_n}^{-1}(A) &= A_0 + \alpha_d^{-1} A_{\mathsf{Ax}} + \beta_d^{-1} A_{\mathsf{Mix}}.\end{aligned} \tag{B13}$$

Taking $A = \Pi_u$ gives

$$\widehat{\rho}_u = \frac{\mathbb{I}}{d} + \alpha_d^{-1}(\Pi_u)_{\mathsf{Ax}} + \beta_d^{-1}(\Pi_u)_{\mathsf{Mix}}. \tag{B14}$$

At $n=3$, $d=8$ and the two nonidentity values coincide at $1/9 = 1/(d+1)$. Since $A_{\mathsf{Ax}} + A_{\mathsf{Mix}} = A - A_0$ and $dA_0 = \mathrm{Tr}(A)\mathbb{I}$, the sector formulas reduce directly to

$$\mathcal{M}_{\phi_3}(A) = A_0 + \frac{A - A_0}{d+1} = \frac{A + \mathrm{Tr}(A)\mathbb{I}}{d+1},$$
$$\mathcal{M}^{-1}_{\phi_3}(A) = A_0 + (d+1)(A - A_0) = (d+1)A - \mathrm{Tr}(A)\mathbb{I}.$$

These are exactly the standard SIC channel and inverse. In particular, $\mathrm{Tr}(\Pi_u) = 1$ gives $\mathcal{M}^{-1}_{\phi_3}(\Pi_u) = 9\Pi_u - \mathbb{I}$. All listed overlaps are nonzero, completing the direct proof of minimal informational completeness.

3. Closed-form dual and outcome-dependent snapshots

We now convert the sector inverse in Eq. (B14) into the product form stated in Theorem 2. For $n \geq 3$ and outcome $u = (\boldsymbol{a}, \boldsymbol{b})$, set

$$z_n = e^{i\theta_n}, \qquad |x_{\boldsymbol{a}}\rangle = Z(\boldsymbol{a})\,|+\rangle^{\otimes n},$$
$$|\boldsymbol{b}\rangle = X(\boldsymbol{b})\,|0^n\rangle, \qquad \omega_{\boldsymbol{a},\boldsymbol{b}} = (-1)^{\boldsymbol{a}\cdot\boldsymbol{b}}.$$

First, direct expansion of the orbit projector gives

$$\Pi_{\boldsymbol{a},\boldsymbol{b}} = \frac{|x_{\boldsymbol{a}}\rangle\langle x_{\boldsymbol{a}}| + |\boldsymbol{b}\rangle\langle\boldsymbol{b}| + z_n^*\omega_{\boldsymbol{a},\boldsymbol{b}}|x_{\boldsymbol{a}}\rangle\langle\boldsymbol{b}| + z_n\omega_{\boldsymbol{a},\boldsymbol{b}}|\boldsymbol{b}\rangle\langle x_{\boldsymbol{a}}|}{\mathcal{N}_n}. \tag{B15}$$

Second, we isolate its axial and mixed Pauli components. Define complete dephasing in the Z and X bases by

$$\Delta_Z(A) := \sum_{\boldsymbol{y}\in\{0,1\}^n} |\boldsymbol{y}\rangle\langle\boldsymbol{y}|A|\boldsymbol{y}\rangle\langle\boldsymbol{y}|,$$
$$\Delta_X(A) := H^{\otimes n}\Delta_Z\big(H^{\otimes n}AH^{\otimes n}\big)\,H^{\otimes n}. \tag{B16}$$

Equivalently, Δ_X is complete dephasing in the X-eigenbasis $\{|x_{\boldsymbol{y}}\rangle\}_{\boldsymbol{y}}$. The two maps retain, respectively, the identity together with the Z-only and X-only Pauli components. Hence

$$A_{\mathsf{Ax}} = \Delta_Z(A) + \Delta_X(A) - 2\frac{\mathrm{Tr}\,A}{d}\mathbb{I},$$
$$A_{\mathsf{Mix}} = A - \Delta_Z(A) - \Delta_X(A) + \frac{\mathrm{Tr}\,A}{d}\mathbb{I}. \tag{B17}$$

With $\kappa_d = 1 - r_d$, direct dephasing gives

$$\Delta_Z(\Pi_{\boldsymbol{a},\boldsymbol{b}}) = \frac{\mathbb{I}/d + \kappa_d|\boldsymbol{b}\rangle\langle\boldsymbol{b}|}{\mathcal{N}_n},$$
$$\Delta_X(\Pi_{\boldsymbol{a},\boldsymbol{b}}) = \frac{\mathbb{I}/d + \kappa_d|x_{\boldsymbol{a}}\rangle\langle x_{\boldsymbol{a}}|}{\mathcal{N}_n}. \tag{B18}$$

Equations (B17) and (B18) therefore determine $(\Pi_{\boldsymbol{a},\boldsymbol{b}})_{\mathsf{Ax}}$ and $(\Pi_{\boldsymbol{a},\boldsymbol{b}})_{\mathsf{Mix}}$ explicitly. Finally, substituting them into Eq. (B14) gives

$$\widehat{\rho}_{\boldsymbol{a},\boldsymbol{b}} = \frac{\mathcal{N}_n}{r_d\kappa_d}(|x_{\boldsymbol{a}}\rangle\langle x_{\boldsymbol{a}}| + |\boldsymbol{b}\rangle\langle\boldsymbol{b}|) + \frac{\mathcal{N}_n}{r_d^2}\left(z_n^*\omega_{\boldsymbol{a},\boldsymbol{b}}|x_{\boldsymbol{a}}\rangle\langle\boldsymbol{b}| + z_n\omega_{\boldsymbol{a},\boldsymbol{b}}|\boldsymbol{b}\rangle\langle x_{\boldsymbol{a}}|\right) - \frac{r_d}{\kappa_d}\mathbb{I}. \tag{B19}$$

Setting $\boldsymbol{a} = \boldsymbol{b} = \boldsymbol{0}$ recovers the reference dual in Eq. (16). Together with the $n = 1, 2$ calculations in Appendix B 2, this proves Theorem 2.

Appendix C: Gram spectrum and observable-estimation proofs

1. Projector-Gram connection

The expectation-value calculation above diagonalizes the measurement channel. To use the projector-Gram spectrum in the performance comparison, we now show that the same Pauli sign patterns also diagonalize the projector-Gram matrix. Fix one Pauli operator P_v, and let the displacement—equivalently, outcome—label u run over all d^2 values. The conjugation rule

$$D_u P_v D_u^\dagger = s_v(u) P_v$$

assigns one sign to each u: $s_v(u) = +1$ when D_u commutes with P_v, and $s_v(u) = -1$ when they anticommute. Collecting these signs produces the d^2-component column vector

$$\mathbf{s}_v := (s_v(u))_u \in \{+1, -1\}^{d^2}.$$

As v runs over all d^2 Pauli labels, the vectors $\mathbf{s}_v$ form the orthogonal character basis in Eq. (B1); each has norm d, so $\mathbf{s}_v/d$ is normalized.

The matrix G^Π records the pairwise overlaps of the orbit projectors. Using Eq. (B2) gives

$$G^\Pi_{u,u'} = \mathrm{Tr}(\Pi_u\Pi_{u'}) = \frac{1}{d}\sum_w s_w(u)s_w(u')|\chi_w|^2,$$
$$\big(G^\Pi\mathbf{s}_v\big)_u = \sum_{u'} G^\Pi_{u,u'}s_v(u') = d|\chi_v|^2 s_v(u). \tag{C1}$$

The second line is ordinary matrix-vector multiplication. After inserting the first line, character orthogonality gives $\sum_{u'} s_w(u')s_v(u') = d^2\delta_{w,v}$: every term with $w \neq v$ cancels, leaving only $w = v$. Therefore, for each Pauli label v, the projector-Gram eigenpair is

$$G^\Pi\mathbf{s}_v = \lambda_v\mathbf{s}_v, \qquad \mathbf{s}_v = (s_v(u))_u, \qquad \lambda_v = d|\chi_v|^2.$$

Here G^Π is the $d^2 \times d^2$ matrix being diagonalized, $\mathbf{s}_v$ is its d^2-component eigenvector, and λ_v is the corresponding scalar eigenvalue. The Pauli label v indexes these eigenpairs. This is the bridge between the fiducial expectation values that control reconstruction and the projector-Gram spectrum used in Sec. IV A.

2. Phase selection and full finite-size projector-Gram spectrum

a. Phase selection. We use the axial and mixed label sets of Eq. (18) in Appendix B 2. To explain the phases in

Eq. (2), allow an arbitrary relative phase θ while keeping the two branch amplitudes equal. Set $\mathcal{N}_\theta = 2+2\cos\theta/\sqrt{d}$. Replacing z_n by $e^{i\theta}$ in Eq. (B6) and using the Gram relation in Appendix C 1 gives the three nonidentity eigenvalue branches

$$\lambda_{\mathrm{Ax}}(\theta) = \frac{(\sqrt{d}+2\cos\theta)^2}{\mathcal{N}_\theta^2},$$
$$\lambda_{\mathrm{even}}(\theta) = \frac{4\cos^2\theta}{\mathcal{N}_\theta^2}, \qquad \lambda_{\mathrm{odd}}(\theta) = \frac{4\sin^2\theta}{\mathcal{N}_\theta^2}.$$

The last two branches are the mixed directions with even and odd numbers of Y factors; both are present for $d \geq 4$. For $d \geq 8$, put $x = |\cos\theta|$. Their smaller eigenvalue obeys

$$\min\{\lambda_{\mathrm{even}}, \lambda_{\mathrm{odd}}\} \leq f_d(x) := \frac{\min\{x^2, 1-x^2\}}{(1-x/\sqrt{d})^2},$$

with equality when $\cos\theta = -x$. The function f_d increases up to $x = 1/\sqrt{2}$ and decreases thereafter: on the second interval its derivative has the sign of $1/\sqrt{d}-x < 0$. Hence the mixed-sector upper bound is attained at $\theta = 3\pi/4$. At this phase,

$$\frac{\lambda_{\mathrm{Ax}}}{\lambda_{\mathrm{even}}} = \frac{(\sqrt{d}-\sqrt{2})^2}{2} \geq 1.$$

The axial branch does not lower the floor, proving phase optimality for $d \geq 8$.

For $d = 4$ and $\cos\theta = -x \leq 0$, the smaller of the axial and even mixed branches is

$$\frac{\min\{(1-x)^2, x^2\}}{(1-x/2)^2} \leq \frac{4}{9}.$$

It is maximized at $x = 1/2$; there the odd branch is $4/3$, so $\theta = 2\pi/3$ attains the optimal floor $4/9$. For $\cos\theta \geq 0$, the mixed-sector floor is at most $2/(2+1/\sqrt{2})^2 < 4/9$. Finally, for $d = 2$ the even mixed branch is absent, and $\theta = 5\pi/12$ makes all three nonidentity eigenvalues $2/3$. This reaches the SIC upper bound: the nonidentity eigenvalues of any pure-state Pauli orbit sum to d^2-d, so their minimum cannot exceed $(d^2-d)/(d^2-1) = d/(d+1)$. These optimizations concern the relative phase in the stated equal-weight family, not unrestricted fiducials.

b. Spectrum at the selected phases. Substituting these phases gives the finite-size spectrum of Theorem 8 in Ref. [47]. Its nonidentity part is

$$\begin{cases} \{(2/3)^{(3)}\}, & d = 2, \\ \{(4/9)^{(9)}, (4/3)^{(6)}\}, & d = 4, \\ \{(8/9)^{(63)}\}, & d = 8, \\ \{\mu_d^{((d-1)^2)}, \nu_d^{(2(d-1))}\}, & d > 8, \end{cases} \tag{C2}$$

where

$$\mu_d = \frac{2}{(2-\sqrt{2/d})^2}, \qquad \nu_d = \frac{d(1-\sqrt{2/d})^2}{(2-\sqrt{2/d})^2}. \tag{C3}$$

For $d \geq 8$, Eq. (B6) assigns ν_d to the $2(d-1)$ nonidentity X- and Z-axis labels and μ_d to the $(d-1)^2$ mixed labels; the two values coincide at $d = 8$. At $d = 4$, Eq. (B8) instead gives the even- and odd-Y parity sectors with multiplicities nine and six.

For $d \geq 8$, the mixed sector attains the smallest eigenvalue, so

$$\lambda_{\min} = d\beta_d = \frac{2}{\mathcal{N}_n^2},$$
$$\eta_n = \frac{d+1}{d}\lambda_{\min} = \frac{2(d+1)}{d\mathcal{N}_n^2} \longrightarrow \frac{1}{2}.$$

Together with $(\eta_1, \eta_2, \eta_3) = (1, 5/9, 1)$, these expressions give Eqs. (20) and (21). The spectrum is not asymptotically flat: $\nu_d/\mu_d = (\sqrt{d}-\sqrt{2})^2/2$ diverges. The main-text claim is a uniform lower-edge statement, not geometric convergence to a SIC.

3. General-observable variance bounds

We prove Proposition 4. At $\rho = \mathbb{I}/d$, every rank-one effect has probability $1/d^2$. By self-adjointness of the inverse and the frame identity,

$$\mathbb{E}_{\mathbb{I}/d}\widehat{O}^2 = \frac{1}{d^2}\sum_u [\mathrm{Tr}(\Pi_u \mathcal{M}_{\phi_n}^{-1}(O))]^2$$
$$= \frac{1}{d}\,\mathrm{Tr}[O\mathcal{M}_{\phi_n}^{-1}(O)]. \tag{C4}$$

Since $\mathcal{M}^{-1}(\mathbb{I}) = \mathbb{I}$, applying Eq. (C4) to O_0 gives its maximally mixed second moment. Its mean vanishes at $\mathbb{I}/d$, so it is also the variance there. In the normalized Pauli basis, the nonidentity inverse eigenvalues are d/λ_v, so this quantity is at most $\mathrm{Tr}(O_0^2)/\lambda_{\min}$. Equations (19) and (21) give the stated upper bound.

The conditional second moment is linear in the input state, so its Haar average equals its value at $\mathbb{I}/d$, whereas

$$\mathbb{E}_{\psi\sim\mathrm{Haar}}\langle\psi|O_0|\psi\rangle^2 = \frac{\mathrm{Tr}(O_0^2)}{d(d+1)}. \tag{C5}$$

Subtracting this squared mean proves Eq. (24). Finally, rank-one positivity gives $p_\rho(u) = \mathrm{Tr}(\rho\Pi_u)/d \leq 1/d = d\,p_{\mathbb{I}/d}(u)$. Hence

$$\mathbb{E}_\rho[\widehat{O}_0^2] \leq d\mathbb{E}_{\mathbb{I}/d}[\widehat{O}_0^2]$$
$$\leq 2(d+1)\,\mathrm{Tr}(O_0^2). \tag{C6}$$

Since variance is at most the second moment, taking the supremum over ρ proves Eq. (25). Applying a standard median-of-means estimate with $B_{\mathrm{var}}(\rho)$ from Corollary 5 to each target and taking a union bound proves Eq. (26).

4. Exact Pauli second moments

We derive Eq. (27) and prove Theorem 6. Pauli conjugation gives $\mathrm{Tr}(P_v\Pi_u) = s_v(u)\chi_v$, while $\mathcal{M}_{\phi_n}^{-1}(P_v) =$

$P_v/|\chi_v|^2$. Since P_v is Hermitian, χ_v is real. The Pauli-diagonal inverse is self-adjoint, and therefore

$$\widehat{P}_v(u) = \mathrm{Tr}[P_v \mathcal{M}_{\phi_n}^{-1}(\Pi_u)] = \mathrm{Tr}[\mathcal{M}_{\phi_n}^{-1}(P_v)\Pi_u] = \frac{s_v(u)}{\chi_v}. \tag{C7}$$

Its square is $1/|\chi_v|^2 = d/\lambda_v$ for every outcome and hence for every input state. Subtracting the squared mean proves Eq. (29). The $n = 1$ expectation value and Eq. (B8) give the low-dimensional values stated after the theorem. For $n \geq 3$, inserting α_d and β_d gives Eq. (28). Here $r_d \leq 1/2$, so the mixed value dominates the axial value, with equality at $n = 3$; the former is exactly $(\sqrt{2d}-1)^2$, while the latter is at most 9 and tends to 4. Together with the exceptional cases, these observations complete the proof.

5. Product-SIC baselines

For one qubit the tetrahedral canonical snapshot is $3\Pi_r - \mathbb{I}$. Every tetrahedral Bloch component is $\pm 1/\sqrt{3}$, so a nonidentity Pauli estimator is $\pm\sqrt{3}$. The same snapshot has eigenvalues $2, -1$ and therefore squared Hilbert–Schmidt norm 5. Tensorization gives $3^{w(P)}$ for a weight-$w(P)$ Pauli and 5^n for the squared snapshot norm. Born averaging and unbiasedness give Eq. (30); inserting $H_{\mathrm{prod}} = 5^n$ from Eq. (34) into the iid calculation in Appendix D gives the product-SIC tomography error.

Appendix D: Canonical tomography error

We prove Corollary 7. Unbiasedness is Eq. (7). Independence and zero mean of $\widehat{\rho}_{u_s} - \rho$ remove all cross terms, leaving

$$\mathbb{E}\|\widehat{\rho}_N - \rho\|_F^2 = \frac{1}{N}\left[\mathbb{E}\,\mathrm{Tr}(\widehat{\rho}_u^2) - \mathrm{Tr}(\rho^2)\right]. \tag{D1}$$

WH covariance makes all snapshots unitarily conjugate and hence equal in norm. In the normalized Pauli basis, the squared coefficient of Π_u in direction $v \neq 0$ is λ_v/d^2, while inversion multiplies that coefficient by d/λ_v. Parseval therefore gives

$$\mathrm{Tr}(\widehat{\rho}_u^2) = \frac{1}{d}\left(1 + \sum_{v\neq 0}\frac{d}{\lambda_v}\right) = H_n. \tag{D2}$$

For the low-dimensional cases this yields

$$\begin{aligned} H_1 &= \frac{1+3\cdot 3}{2} = 5, \\ H_2 &= \frac{1+9\cdot 9+6\cdot 3}{4} = 25, \\ H_3 &= \frac{1+63\cdot 9}{8} = 71. \end{aligned} \tag{D3}$$

For $n \geq 3$, counting $2(d-1)$ axis and $(d-1)^2$ mixed directions gives the second line of Eq. (32). Bounding every nonidentity term by $1/\lambda_{\min}$ and using Eq. (21) yields Eq. (33); substituting into Eq. (D1) proves the corollary.

For completeness, the global benchmarks in Eq. (34) follow from the standard canonical snapshot

$$R = (d+1)\Pi - \mathbb{I}, \qquad \Pi^2 = \Pi, \quad \mathrm{Tr}\,\Pi = 1. \tag{D4}$$

For a global SIC, Π is the outcome projector; the corresponding linear-tomography MSE is also given explicitly in Eq. (103) of Ref. [34]. For uniformly randomized global-Clifford measurements, $\Pi = U^\dagger|b\rangle\langle b|U$, and the same snapshot formula is Eq. (S16) of Ref. [3]. In both cases,

$$\begin{aligned} R^2 &= (d^2-1)\Pi + \mathbb{I}, \\ \mathrm{Tr}(R^2) &= d^2 + d - 1. \end{aligned} \tag{D5}$$

Since this norm is independent of the measurement record, unbiasedness and Eq. (D1) give

$$\mathbb{E}\|\widehat{\rho}_N - \rho\|_F^2 = \frac{d^2 + d - 1 - \mathrm{Tr}(\rho^2)}{N}. \tag{D6}$$

Thus $H_{\mathrm{SIC}} = H_{\mathrm{Cl}} = d^2 + d - 1 = 4^n + 2^n - 1$ for $d = 2^n$, under the same i.i.d. canonical linear-inversion convention.

For the local baselines, the product-SIC snapshot norm is 5^n, as shown in Appendix C 5. Uniformly randomized local-Pauli acquisition has the same norm because its single-qubit canonical factors also have the form $3\Pi - \mathbb{I}$. We can now compare the coefficients. Equation (33) gives $H_n \leq 2H_{\mathrm{SIC}} - 1/d < 2H_{\mathrm{SIC}}$. For the local comparison, $H_3 = 71 < 5^3$. For $n \geq 4$, the same bound gives $H_n < 2d^2 + 2d \leq 5^n$, since $2(4/5)^n + 2(2/5)^n$ is less than one at $n = 4$ and decreases with n. Hence $H_n < H_{\mathrm{Pauli}} = H_{\mathrm{prod}}$ for every $n \geq 3$.

Appendix E: Detailed proof of the fixed-circuit realization

a. Conventions and readout labels. Throughout, $\boldsymbol{a}$ denotes the final S-register readout after the Hadamards, while $\boldsymbol{b}$ denotes the A-register readout. For $\boldsymbol{x}, \boldsymbol{a}, \boldsymbol{b} \in \mathbb{F}_2^n$, Eq. (1) gives

$$\begin{aligned} D_{\boldsymbol{a},\boldsymbol{b}} &= Z(\boldsymbol{a})X(\boldsymbol{b}), \\ Z(\boldsymbol{a})\,|\boldsymbol{x}\rangle &= (-1)^{\boldsymbol{a}\cdot\boldsymbol{x}}\,|\boldsymbol{x}\rangle, \quad X(\boldsymbol{b})\,|\boldsymbol{x}\rangle = |\boldsymbol{x}\oplus\boldsymbol{b}\rangle. \end{aligned} \tag{E1}$$

We use this representative for the exact amplitude identity below; multiplying $D_{\boldsymbol{a},\boldsymbol{b}}$ by a unit phase leaves its orbit projector and probability unchanged.

b. Pure-state amplitude through the circuit. For a pure system input, the state after conjugate-fiducial preparation is

$$|\psi\rangle_S\,|\phi^*\rangle_A = \sum_{\boldsymbol{x},\boldsymbol{y}} \psi_{\boldsymbol{x}}\phi^*_{\boldsymbol{y}}\,|\boldsymbol{x}\rangle_S\,|\boldsymbol{y}\rangle_A. \tag{E2}$$

The conjugated resource supplies the coefficients $\phi^*_{\boldsymbol{y}}$ appearing in the bra $\langle\phi|$.

Each $\mathrm{CNOT}_{S_j\to A_j}$ maps $|x_j, y_j\rangle$ to $|x_j, y_j \oplus x_j\rangle$. Applying the product of all n matched CNOTs therefore gives

$$\sum_{\boldsymbol{x},\boldsymbol{y}} \psi_{\boldsymbol{x}}\phi^*_{\boldsymbol{y}} \,|\boldsymbol{x}\rangle_S \,|\boldsymbol{y}\oplus\boldsymbol{x}\rangle_A \,. \tag{E3}$$

For the final A-register label $\boldsymbol{b}$, only terms satisfying $\boldsymbol{y}\oplus\boldsymbol{x} = \boldsymbol{b}$ contribute, equivalently $\boldsymbol{y} = \boldsymbol{x}\oplus\boldsymbol{b}$. The corresponding unnormalized system component is

$$\sum_{\boldsymbol{x}} \psi_{\boldsymbol{x}}\phi^*_{\boldsymbol{x}\oplus\boldsymbol{b}} \,|\boldsymbol{x}\rangle_S \,. \tag{E4}$$

The Walsh–Hadamard identity is

$$H^{\otimes n}\,|\boldsymbol{x}\rangle = d^{-1/2}\sum_{\boldsymbol{a}} (-1)^{\boldsymbol{a}\cdot\boldsymbol{x}}\,|\boldsymbol{a}\rangle \,. \tag{E5}$$

Applying these Hadamards and selecting the final S-register label $\boldsymbol{a}$ gives the full-circuit amplitude

$$\begin{aligned}
&\langle\boldsymbol{a},\boldsymbol{b}|\, U_{\mathrm{IC}}(\phi)(|\psi\rangle_S \otimes |0^n\rangle_A) \\
&\quad = \langle\boldsymbol{a},\boldsymbol{b}|\, U_{\mathrm{Bell}}(|\psi\rangle_S \otimes |\phi^*\rangle_A) \\
&\quad = d^{-1/2}\sum_{\boldsymbol{x}} (-1)^{\boldsymbol{a}\cdot\boldsymbol{x}}\psi_{\boldsymbol{x}}\phi^*_{\boldsymbol{x}\oplus\boldsymbol{b}}.
\end{aligned} \tag{E6}$$

c. Identify the WH displacement. For the convention above, taking the adjoint reverses the order of the Hermitian Pauli factors:

$$D^\dagger_{\boldsymbol{a},\boldsymbol{b}} = X(\boldsymbol{b})Z(\boldsymbol{a}), \qquad D^\dagger_{\boldsymbol{a},\boldsymbol{b}}\,|\boldsymbol{x}\rangle = (-1)^{\boldsymbol{a}\cdot\boldsymbol{x}}\,|\boldsymbol{x}\oplus\boldsymbol{b}\rangle \,. \tag{E7}$$

Therefore

$$\langle\phi|\, D^\dagger_{\boldsymbol{a},\boldsymbol{b}}\,|\psi\rangle = \sum_{\boldsymbol{x}} (-1)^{\boldsymbol{a}\cdot\boldsymbol{x}}\psi_{\boldsymbol{x}}\phi^*_{\boldsymbol{x}\oplus\boldsymbol{b}}. \tag{E8}$$

Comparing with Eq. (E6) proves the required amplitude identity.

d. Born probability and mixed inputs. For the pure state $\rho_\psi = |\psi\rangle\langle\psi|$, Born's rule gives

$$\begin{aligned}
p(\boldsymbol{a},\boldsymbol{b}\mid\psi) &= \frac{1}{d}\left|\langle\phi|\, D^\dagger_{\boldsymbol{a},\boldsymbol{b}}\,|\psi\rangle\right|^2 \\
&= \frac{1}{d}\,\mathrm{Tr}\Big[\rho_\psi D_{\boldsymbol{a},\boldsymbol{b}}|\phi\rangle\langle\phi|D^\dagger_{\boldsymbol{a},\boldsymbol{b}}\Big] \,.
\end{aligned} \tag{E9}$$

A mixed state has a decomposition $\rho = \sum_k r_k|\psi_k\rangle\langle\psi_k|$ with $r_k \ge 0$ and $\sum_k r_k = 1$. Both the Born probability and the trace expression are affine in ρ, so averaging the pure-state identity gives

$$p(\boldsymbol{a},\boldsymbol{b}\mid\rho) = \frac{1}{d}\,\mathrm{Tr}\Big[\rho D_{\boldsymbol{a},\boldsymbol{b}}|\phi\rangle\langle\phi|D^\dagger_{\boldsymbol{a},\boldsymbol{b}}\Big] \,. \tag{E10}$$

This is Eq. (38). For $\phi = \phi_n$, Eq. (3) identifies the operator inside the Born rule with $E^{(n)}_{\boldsymbol{a},\boldsymbol{b}}$, completing the proof of Proposition 9.

Appendix F: Explicit CNOT compilation of the root and binary splits

Root block: one CNOT. The optimization is a state preparation, not a one-CNOT implementation of the full first-split unitary. Let $a+b=n$, with $(a,b)=(1,n-1)$ for the path and $(a,b)=(\lfloor n/2\rfloor,\lceil n/2\rceil)$ for the balanced tree. Writing the state in Eq. (55) as $\left|\xi^{(n)}_{a,b}\right\rangle = \sum_{j,k=0}^{1}(X_{a,b})_{jk}\,|j,k\rangle$ gives the known coefficient matrix

$$X_{a,b} = \frac{1}{\sqrt{\mathcal{N}_n}}\begin{pmatrix} e^{i\theta_n}+c_ac_b & c_ar_b \\ r_ac_b & r_ar_b \end{pmatrix}. \tag{F1}$$

It obeys $\mathrm{Tr}(X^\dagger_{a,b}X_{a,b}) = 1$ and

$$\det X_{a,b} = \frac{e^{i\theta_n}r_ar_b}{\mathcal{N}_n}, \qquad \Delta_{a,b} := |\det X_{a,b}| = \frac{r_ar_b}{\mathcal{N}_n} > 0. \tag{F2}$$

Take a singular-value decomposition

$$X_{a,b} = L_{a,b}\,\mathrm{diag}(\sigma_0,\sigma_1)R^\dagger_{a,b}, \qquad \sigma_0 \ge \sigma_1 > 0. \tag{F3}$$

Normalization and Eq. (F2) give $\sigma^2_{0,1} = (1 \pm \sqrt{1-4\Delta^2_{a,b}})/2$. Define

$$\begin{aligned}
\gamma_{a,b} &:= \arcsin(2\Delta_{a,b}) \\
&= \arcsin\left(\frac{2r_ar_b}{\mathcal{N}_n}\right), \qquad 0 < \gamma_{a,b} \le \frac{\pi}{2}.
\end{aligned} \tag{F4}$$

Then $\sigma_0 = \cos(\gamma_{a,b}/2)$ and $\sigma_1 = \sin(\gamma_{a,b}/2)$. With $R_y(\vartheta) = e^{-i\vartheta Y/2}$, an explicit root circuit is

$$\left|\xi^{(n)}_{a,b}\right\rangle = (L_{a,b}\otimes R^*_{a,b})\,\mathrm{CNOT}_{1\to2}[R_y(\gamma_{a,b})\otimes I]\,|00\rangle \,. \tag{F5}$$

Here $L_{a,b}$ and $R^*_{a,b}$ are single-qubit unitaries obtained from the SVD; the displayed CNOT is the sole two-qubit gate. Indeed, the rotation and CNOT first produce $\sigma_0\,|00\rangle + \sigma_1\,|11\rangle$. Under local gates $A\otimes B$, a coefficient matrix transforms as $X \mapsto AXB^T$, so the final local gates produce $L_{a,b}\,\mathrm{diag}(\sigma_0,\sigma_1)R^\dagger_{a,b}$. The complex conjugate on $R_{a,b}$ is therefore essential. Since $\det X_{a,b} \neq 0$, the target is entangled and cannot be prepared using only one-qubit gates; Eq. (F5) is thus an exact and optimal one-CNOT preparation. For the path, it directly prepares Eq. (43) from $|00\rangle$, replacing the separate preparation of $|\eta_n\rangle$ and application of G_1.

All nonroot splits: one two-CNOT template. Let $a,b \ge 1$ and $m = a+b$. The two branch inputs in Eq. (45) span $\{|00\rangle, |10\rangle\}$. By linearity, the required action on these orthonormal inputs is

$$\begin{aligned}
U_{a,b}\,|00\rangle &= |00\rangle \,, \\
U_{a,b}\,|10\rangle &= \frac{1}{r_m}\big(c_ar_b\,|01\rangle + r_ac_b\,|10\rangle + r_ar_b\,|11\rangle\big).
\end{aligned} \tag{F6}$$

The second output is normalized since $r_b^2 + r_a^2 c_b^2 = r_m^2$ and is orthogonal to the first. Define the two rotation angles

$$\alpha_a := \arccos c_a, \qquad \delta_{a,b} := \arctan\left(\frac{r_a c_b}{r_b}\right). \tag{F7}$$

Thus $\cos\alpha_a = c_a$, $\sin\alpha_a = r_a$, $\sin\delta_{a,b} = r_a c_b / r_m$, and $\cos\delta_{a,b} = r_b/r_m$. An explicit unitary completion is

$$\begin{aligned} U_{a,b} = {} & [R_y(\alpha_a) \otimes I]\ \mathrm{CNOT}_{2\to1} \\ & \times [R_y(-\alpha_a) \otimes R_y(-\delta_{a,b})] \\ & \times \mathrm{CNOT}_{1\to2}[I \otimes R_y(\delta_{a,b})]. \end{aligned} \tag{F8}$$

Figure 5 shows the gates in physical time order.

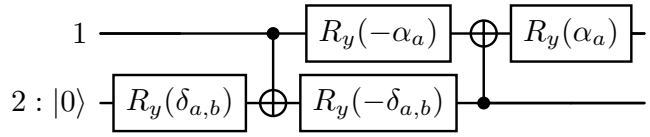

FIG. 5. Two-CNOT circuit for the split $m = a + b$. Wire 1 carries the one-qubit memory; wire 2 is an unused output qubit. The same circuit implements both the path ($a = 1$) and balanced-tree splits; only the rotation angles change.

To verify the circuit, the target rotation, first CNOT, and inverse target rotation leave $|00\rangle$ unchanged and send $|10\rangle$ to $\sin\delta_{a,b}|10\rangle + \cos\delta_{a,b}|11\rangle$. The remaining operation is the second CNOT conjugated by $R_y(\alpha_a)$ on wire 1. When wire 2 is one, it maps $|1\rangle$ on wire 1 to $c_a|0\rangle + r_a|1\rangle$; when wire 2 is zero, it acts as the identity. Substitution gives Eq. (F6). In particular,

$$U_{a,b}(c_m|00\rangle + r_m|10\rangle) = (c_a|0\rangle + r_a|1\rangle) \otimes (c_b|0\rangle + r_b|1\rangle),$$

which is the required coherent branch map.

For the path, $a = 1$ gives $\alpha_1 = \pi/4$ and $\delta_{1,m-1} = \arctan(c_m/r_{m-1})$, recovering $U_m = U_{1,m-1}$. Equation (F8) therefore supplies two CNOTs and four one-qubit rotations for every nonroot split in either schedule. The first split still uses the separate one-CNOT state preparation above, so the preparation count $1+2(n-2) = 2n-3$ and balanced CNOT depth $2\lceil\log_2 n\rceil - 1$ are unchanged.

Two CNOTs are also necessary for this full coherent action on the reachable input subspace. Up to local gates, a one-CNOT circuit receives two orthogonal first-qubit states with the same second-qubit factor. Expanding either control direction shows that the two output coefficient determinants vanish simultaneously. Here $U_{a,b}|00\rangle$ is a product state, whereas the coefficient determinant of $U_{a,b}|10\rangle$ is $-c_a r_a r_b c_b / r_m^2 \neq 0$. One CNOT cannot implement both prescribed columns. This is a specialized real-isometry compilation [74, 84, 85].

Appendix G: Hoggar fiducial: preparation lower bound and numerical circuit

At $d = 8$, $\mathcal{N}_3 = 3/2$ and

$$|\phi_3\rangle = \frac{1}{\sqrt{12}}(-1+2i, 1, 1, 1, 1, 1, 1, 1)^T. \tag{G1}$$

This is the standard Pauli-covariant Hoggar fiducial [47]. In the ancilla-free three-qubit unitary model with arbitrary one-qubit gates and arbitrary CNOT connectivity, it requires exactly three CNOTs to prepare from a product state. Three suffice for any three-qubit pure state [86]. For the lower bound, the unnumbered section "Classification with respect to $|000\rangle$" of Ref. [86] gives, for its exact two-CNOT class up to local unitaries and a qubit permutation, with $\langle a|a^\perp\rangle = 0$, the form

$$\cos\alpha\,|a\rangle\,|b\rangle\,|c\rangle + \sin\alpha\,|a^\perp\rangle\,|b'\rangle\,|c'\rangle\,. \tag{G2}$$

The zero-CNOT class is the degenerate product case. The one-CNOT class is biseparable; Schmidt-decomposing its entangled pair and, if necessary, relabeling the qubits gives the same orthogonal-factor form. Hence every state reachable with at most two CNOTs has a representative of this form. Tracing out the orthogonal-factor qubit gives

$$\rho_{23} = \cos^2\alpha|bc\rangle\langle bc| + \sin^2\alpha|b'c'\rangle\langle b'c'|, \tag{G3}$$

a separable mixture. Thus at least one two-qubit reduced concurrence is zero. Direct reduction of the Hoggar vector gives

$$C(\rho_{12}) = C(\rho_{13}) = C(\rho_{23}) = \frac{\sqrt{2}}{3} > 0. \tag{G4}$$

For an analytic check, tracing out one qubit gives the two subnormalized conditional vectors

$$\begin{aligned} v_0 = (a, b, b, b), \quad v_1 = (b, b, b, b), \\ a = \frac{-1+2i}{\sqrt{12}}, \qquad b = \frac{1}{\sqrt{12}}. \end{aligned} \tag{G5}$$

Their Wootters matrix $\tau_{rs} = \langle v_r|(\sigma_y \otimes \sigma_y)|v_s^*\rangle$ is, up to entrywise conjugation, $c\left(\begin{smallmatrix} 2 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ with $|c| = \sqrt{2}/6$. Its

TABLE V. One numerical Euler-angle realization of the ancilla-free CNOT-optimal Hoggar fiducial preparation, rounded to six decimals for display. The numerical calculations use the angles at full precision.

ℓ	qubit j	$(a_{\ell j}, b_{\ell j}, c_{\ell j})$
1	1	(−0.846651, 1.045827, −3.103113)
1	2	(−1.258829, −2.482732, −2.932261)
1	3	(2.289262, −0.126380, −0.101227)
2	1	(−2.266624, −2.170051, 0.240365)
2	2	(2.092672, 2.799899, −2.845832)
2	3	(−0.549941, −0.809641, 2.049816)
3	1	(−2.987084, −2.685355, −1.727123)
3	2	(−3.079891, −0.795702, 2.443830)
3	3	(1.842244, −0.142857, 2.655990)
4	1	(1.360092, −0.627407, −0.853500)
4	2	(−3.058065, −0.762020, 0.350864)
4	3	(−3.004442, 0.230627, 1.255897)

two singular values differ by $2|c| = \sqrt{2}/3$. Permutation symmetry of the target gives the same result for all three pairs.

Concurrence is invariant under local unitaries, and a qubit permutation only relabels the three reduced pairs, so the preceding obstruction applies to every representative in the cited at-most-two-CNOT classes. The positive values therefore exclude all zero-, one-, and two-CNOT classes. This proves CNOT-optimal preparation of the fiducial, not global optimality of a Naimark unitary. A separate question is whether the same three-qubit Hoggar SIC POVM admits an exact six-qubit, strict-unitary Naimark realization with fewer than six CNOTs under the same gate model.

An analytic three-CNOT preparation follows from the $n = 3$ instance of the general memory-splitting construction in Appendix F: one CNOT prepares the root state $\left|\xi^{(3)}_{1,2}\right\rangle$, and two implement the final $U_{1,1}$ split. Figure 6 gives a separate numerical three-CNOT compilation of the same fiducial, rather than an expansion of that analytic circuit.

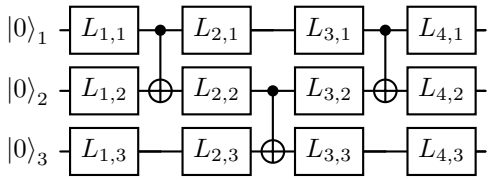

FIG. 6. Numerical ancilla-free three-CNOT compilation of the Hoggar fiducial. The preceding analytic argument proves that the CNOT count is optimal. Here $L_{\ell,j} = R_z(a_{\ell j})R_y(b_{\ell j})R_z(c_{\ell j})$, with $R_z(\vartheta) = e^{-i\vartheta Z/2}$; Table V lists the angles rounded to six decimals.

Appendix H: Alternative coherent and adaptive LCU completions

The heralded branch synthesis of Lemma 14 can be made deterministic by one exact amplitude-amplification step. We give two implementations of the same identity: a fully unitary circuit and a measurement-and-feedforward circuit. We first prove the logical identity, then compile its two selective phases in the two models.

Logical amplitude amplification.— Lemma 14 gives

$$\begin{aligned} |G\rangle &:= |0\rangle_F |\phi_n\rangle_A |\mathbf{0}\rangle_W , \\ |v\rangle &:= V_n \left|0^{2n}\right\rangle = \sqrt{p_n}\,|G\rangle + \sqrt{1-p_n}\,|B\rangle , \end{aligned} \tag{H1}$$

where $|B\rangle$ is the normalized $F = 1$ branch and is orthogonal to $|G\rangle$. The task is to remove its $|B\rangle$ amplitude coherently. Figure 7(a) previews the five operations below in their physical time order.

First phase the good branch. Because the two branches are distinguished by F, define

$$S_G(\varphi) := \mathbb{I}_{FAW} + (e^{i\varphi} - 1)(|0\rangle\langle 0|_F \otimes \mathbb{I}_{AW}). \tag{H2}$$

For $p = p_n$ and $z = e^{i\varphi}$, the state becomes $S_G(\varphi)|v\rangle = z\sqrt{p}\,|G\rangle + \sqrt{1-p}\,|B\rangle$; its overlap with $|v\rangle$ is $1 + p(z-1)$.

Next undo the preparation with $V_n^\dagger$, phase only its all-zero input, and reapply V_n. The required input phase is

$$S_0^{FAW}(\varphi) := \mathbb{I}_{FAW} + (e^{i\varphi} - 1)|0^{2n}\rangle\langle 0^{2n}|. \tag{H3}$$

Conjugating the input phase by V_n gives

$$R_v(\varphi) := V_n S_0^{FAW}(\varphi) V_n^\dagger = \mathbb{I}_{FAW} + (z-1)|v\rangle\langle v|. \tag{H4}$$

Thus the last three physical operations phase the prepared direction $|v\rangle$. Applying them to the state above, the remaining bad-branch coefficient is

$$\begin{aligned} c_B &:= \langle B|\, R_v(\varphi) S_G(\varphi)\, |v\rangle \\ &= \sqrt{1-p}\{1 + (z-1)[1 + p(z-1)]\} \\ &= \sqrt{1-p}\, z[1 - 4p\sin^2(\varphi/2)]. \end{aligned} \tag{H5}$$

The bad branch therefore vanishes if

$$\varphi_n := 2\arcsin\frac{1}{2\sqrt{p_n}}, \qquad 4p_n \sin^2(\varphi_n/2) = 1. \tag{H6}$$

This phase is real whenever $p_n \geq 1/4$, in particular under our stronger bound $p_n \geq 3/8$. At $\varphi = \pi$, S_G and S_0 are the two usual Grover sign flips; the matched non-π phase in Eq. (H6) makes the single step exact [87]. Since $c_B = 0$ and the circuit is unitary, only a normalized good component remains. Hence, up to an irrelevant global phase,

$$\begin{aligned} &\left(V_n S_0^{FAW}(\varphi_n) V_n^\dagger\right) S_G(\varphi_n) V_n \left|0^{2n}\right\rangle \\ &\qquad = e^{i\gamma_n}\,|G\rangle . \end{aligned} \tag{H7}$$

Unitary compilation of the selective phases.— The two selective phases have sharply different circuit costs. The good-branch phase S_G is local because the branch label is stored in F. Let $P(\varphi) := \mathrm{diag}(1, e^{i\varphi})$. This ordinary phase gate phases $|1\rangle$, whereas S_G must phase $F = 0$. Since X exchanges $|0\rangle$ and $|1\rangle$,

$$\begin{aligned} XP(\varphi)X &= \mathrm{diag}(e^{i\varphi}, 1) = \mathbb{I} + (e^{i\varphi} - 1)|0\rangle\langle 0|, \\ S_G(\varphi) &= [X_F P_F(\varphi) X_F] \otimes \mathbb{I}_{AW}. \end{aligned} \tag{H8}$$

Thus Eq. (H2) requires only local gates on F and no entangling gate.

By contrast, S_0^{FAW} must coherently phase one joint basis state, $\left|0^{2n}\right\rangle$. Every V_n-type block begins and ends with $W = |\mathbf{0}\rangle_W$, so immediately before S_0 it is enough to recognize the all-zero string on FA.

Following the standard generalized multi-controlled-unitary notation [88], $\Lambda_R[U_T]$ denotes the gate that applies U to a target T exactly when all $r = |R|$ control qubits are one:

$$\begin{aligned} \Lambda_R[U_T] := &\left(\mathbb{I}_R - |1^r\rangle\langle 1^r|_R\right) \otimes \mathbb{I}_T \\ &+ |1^r\rangle\langle 1^r|_R \otimes U_T, \qquad r = |R|. \end{aligned} \tag{H9}$$

For $U = X$, this is the generalized Toffoli gate [89, 90]. With $X_{FA} := X_F \otimes X_A^{\otimes n}$, the required all-zero phase is therefore

$$S_0^{FA}(\varphi) = X_{FA}\Lambda_A[P_F(\varphi)]X_{FA} = \mathbb{I}_{FA} + (e^{i\varphi} - 1)|0^{n+1}\rangle\langle 0^{n+1}|. \tag{H10}$$

From right to left, the first X_{FA} converts the zero controls to one controls, $\Lambda_A[P_F(\varphi)]$ phases the component for which all FA qubits are one, and the final X_{FA} restores the basis labels. Unlike the local S_G, this is an n-controlled phase, i.e., a general multicontrolled one-qubit unitary. Exact general constructions attain $O(n)$ size and $O(\log n)$ depth with one clean ancilla [89], while recent gate-count-oriented compilations provide linear-size alternatives for general multicontrolled $U(2)$ gates [91]. Here the $n-1$ clean qubits of W are already available, so we use the balanced-AND realization below; it keeps both the logarithmic depth and an explicit CNOT ledger. The circuit computes the joint predicate, applies the phase, and uncomputes the predicate. Measuring the predicate would instead collapse the superposition needed for the cancellation in Eq. (H7).

No new tree register is needed: reuse the $n-1$ clean qubits of W as the internal nodes of a balanced AND tree with leaves $A_1, \ldots, A_n$. Only the root $w_\star$ enters the final controlled phase; the other $n-2$ work qubits store intermediate conjunctions that permit logarithmic-depth computation and uncomputation. Let $U_\wedge$ denote this reversible AND-tree computation. Then

$$X_{FA}U_\wedge^\dagger \,\mathrm{CP}_{w_\star,F}(\varphi)U_\wedge X_{FA}\,|\psi\rangle_{FA}\,|\mathbf{0}\rangle_W = S_0^{FA}(\varphi)\,|\psi\rangle_{FA}\,|\mathbf{0}\rangle_W\,, \tag{H11}$$

Here $\mathrm{CP}_{w_\star,F}(\varphi)$ applies $P_F(\varphi)$ when $w_\star = 1$; equivalently it is $\mathrm{diag}(1,1,1,e^{i\varphi})_{w_\star,F}$. This gate is symmetric under exchanging its two qubits; we write the order $(w_\star, F)$ to emphasize that $w_\star$ stores the AND result and controls the phase $P_F(\varphi)$ on F. Figure 7(b) shows $U_\wedge$ computing $w_\star = A_1 \wedge \cdots \wedge A_n$, the controlled phase acting on $(w_\star, F)$, and $U_\wedge^\dagger$ clearing W. Thus the phase is applied exactly when all flipped FA qubits are one.

For the resource ledger, let $L = \lceil \log_2 n \rceil$. The three V_n-type calls contribute $3(3n-2)$ CNOTs and depth $3(2L+1)$; S_G contributes no CNOT. The AND tree and its inverse contain $2(n-1)$ Toffolis, and the controlled phase contributes two CNOTs. Using an exact six-CNOT Toffoli gives

$$C_{\mathrm{prep}} = 3(3n-2) + 6[2(n-1)] + 2 = 21n - 16, \qquad D_{\mathrm{prep}} = 3(2L+1) + 2(6L) + 2 = 18L + 5. \tag{H12}$$

The initialized auxiliary register is only FW, of size $1 + (n-1) = n$; including A, the preparation width is $2n$. These are transparent standard-library upper bounds, not optimized counts.

Deterministic adaptive compilation.— We now prove Proposition 16. In the local alternating quantum–classical computation (LAQCC) model [82], we trade initialized auxiliary qubits for depth: measurement and feedforward compile the same exact five-block identity, Eq. (H7), from $O(\log n)$ coherent depth to $O(1)$ adaptive quantum depth. This conservative construction uses $O(n \log n)$ initialized auxiliary qubits and $O(n \log n)$ gates. No logical block is changed, and no measurement outcome is postselected.

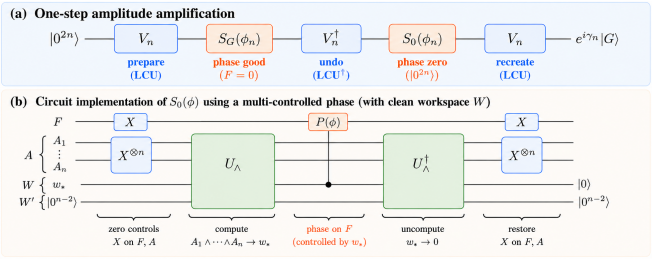

FIG. 7. Circuit view of exact amplitude amplification and its only nonlocal selective phase. Panel (a) shows Eq. (H7) in temporal order. Panel (b) expands S_0 on the invariant subspace with clean W: $U_\wedge$ stores the conjunction in $w_\star$, which controls a phase on F, and $U_\wedge^\dagger$ restores all of W to zero.

Among the five blocks, only the three V_n-type occurrences—V_n, $V_n^\dagger$, and V_n—and the phase S_0 require nonlocal coherent control; S_G is local. Each V_n-type block reduces to one exact dynamic common-control fanout [75], together with parallel one-qubit gates, and therefore has constant adaptive quantum depth. After this replacement, the only remaining nonlocal operation is S_0.

As in the unitary compilation above, W is in $|\mathbf{0}\rangle_W$ whenever S_0^{FAW} is applied. The full-register phase therefore reduces on this invariant subspace to the all-zero phase on FA. We implement it without measuring whether FA is zero: coherently compute this predicate, apply a phase, and erase the computation. Introduce a clean marker qubit T and define

$$U_{\mathrm{zero}} := X_{FA}\Lambda_{FA}[X_T]X_{FA}. \tag{H13}$$

This is an all-zero-controlled generalized Toffoli: for $x \in \{0,1\}^{n+1}$,

$$U_{\mathrm{zero}}\,|x\rangle_{FA}\,|t\rangle_T = |x\rangle_{FA}\,|t \oplus \mathbf{1}_{x=0}\rangle_T\,. \tag{H14}$$

An exact measurement-and-feedback formulation of the multicontrolled-X gate is summarized in Theorem 3 of Ref. [76], following Ref. [83]. Since $\mathbf{1}_{x=0} = 1 \oplus \mathrm{OR}(x)$, implementing U_{zero} is equivalent, up to a local X_T, to computing a reversible OR marker. The OR primitive is only the physical compiler for the same logical operation U_{zero}. The exact quantum OR construction of Ref. [83] and its local adaptive compilation in Ref. [82] provide such a measurement-and-feedforward implementation. Each marker call has $O(1)$ adaptive quantum depth, a constant number of measurement rounds, and $O(\log n)$ classical feedforward depth, using $O(n \log n)$ initialized internal auxiliaries and gates. This internal workspace—not the single logical marker T—accounts for the quoted width.

Only the internal auxiliary qubits are measured: T, and hence the value of the all-zero predicate, is kept coherent. Feedforward applies the outcome-dependent corrections, so the net operation on FAT is exactly U_{zero}.

The selective phase is now obtained in three steps: compute the marker with U_{zero}, apply $P_T(\varphi) = \text{diag}(1, e^{i\varphi})_T$, and uncompute with $U_{\text{zero}}^\dagger = U_{\text{zero}}$. Explicitly,

$$\begin{aligned} &U_{\text{zero}}^\dagger\big(\mathbb{I}_{FA} \otimes P_T(\varphi)\big)U_{\text{zero}} \left|\psi\right\rangle_{FA} \left|0\right\rangle_T \\ &\quad = \left[\mathbb{I}_{FA} + (e^{i\varphi} - 1)|0^{n+1}\rangle\langle 0^{n+1}|\right] \left|\psi\right\rangle_{FA} \left|0\right\rangle_T . \end{aligned} \tag{H15}$$

Thus T is returned to zero, while the all-zero component of FA alone acquires $e^{i\varphi}$. This is exactly S_0 on the invariant subspace; the predicate is neither observed nor postselected.

The full circuit therefore contains three constant-depth dynamic V_n-type calls, the local phase S_G, and two constant-depth marker calls surrounding one local phase. The marker construction dominates the conservative resource bound. Because the constant number of calls reuse its internal workspace, the complete preparation has $O(n \log n)$ peak initialized width and gate count, $O(1)$ adaptive quantum depth, a constant number of measurement rounds, and $O(\log n)$ classical feedforward depth. Because every corrected block is exact, the cancellation in Eq. (H7) still gives $|G\rangle$ with unit probability: there is no heralding measurement and no retry tail. This is the deterministic adaptive entry in Table IV. All auxiliary-side processing is completed before coupling the prepared fiducial to the unknown state.

[1] J. Eisert, D. Hangleiter, N. Walk, I. Roth, D. Markham, R. Parekh, U. Chabaud, and E. Kashefi, Quantum certification and benchmarking, Nat. Rev. Phys. **2**, 382–390 (2020).

[2] M. Kliesch and I. Roth, Theory of quantum system certification, PRX Quantum **2**, 010201 (2021).

[3] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, Nat. Phys. **16**, 1050–1057 (2020).

[4] 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**, 9–24 (2023).

[5] H.-Y. Hu, S. Choi, and Y.-Z. You, Classical shadow tomography with locally scrambled quantum dynamics, Phys. Rev. Research **5**, 023027 (2023).

[6] K. Bu, D. E. Koh, R. J. Garcia, and A. Jaffe, Classical shadows with Pauli-invariant unitary ensembles, npj Quantum Inf. **10**, 6 (2024).

[7] C. Hadfield, S. Bravyi, R. Raymond, and A. Mezzacapo, Measurements of quantum Hamiltonians with locally-biased classical shadows, Commun. Math. Phys. **391**, 951–967 (2022).

[8] H.-Y. Huang, R. Kueng, and J. Preskill, Efficient estimation of Pauli observables by derandomization, Phys. Rev. Lett. **127**, 030503 (2021).

[9] Z.-P. Xu, Q. Ye, J. Wang, G. Koßmann, A. Winter, and R. Schwonnek, Optimal strategies for shadow tomography with limited resources, arXiv:2608.08429 (2026).

[10] M. Ippoliti, Classical shadows based on locally-entangled measurements, Quantum **8**, 1293 (2024).

[11] Y. Wu, C. Wang, J. Yao, H. Zhai, Y.-Z. You, and P. Zhang, Contractive unitary and classical shadow tomography, npj Quantum Inf. **12**, 86 (2026).

[12] K. Wan, W. J. Huggins, J. Lee, and R. Babbush, Matchgate shadows for fermionic quantum simulation, Commun. Math. Phys. **404**, 629–700 (2023).

[13] H. C. Nguyen, J. L. Bönsel, J. Steinberg, and O. Gühne, Optimizing shadow tomography with generalized measurements, Phys. Rev. Lett. **129**, 220502 (2022).

[14] G. García-Pérez, M. A. C. Rossi, B. Sokolov, F. Tacchino, P. K. Barkoutsos, G. Mazzola, I. Tavernelli, and S. Maniscalco, Learning to measure: Adaptive informationally complete generalized measurements for quantum algorithms, PRX Quantum **2**, 040342 (2021).

[15] A. Caprotti, J. Morris, and B. Dakić, Optimizing quantum tomography via shadow inversion, Phys. Rev. Research **6**, 033301 (2024).

[16] M. Ippoliti, Y. Li, T. Rakovszky, and V. Khemani, Operator relaxation and the optimal depth of classical shadows, Phys. Rev. Lett. **130**, 230403 (2023).

[17] C. Bertoni, J. Haferkamp, M. Hinsche, M. Ioannou, J. Eisert, and H. Pashayan, Shallow shadows: Expectation estimation using low-depth random Clifford circuits, Phys. Rev. Lett. **133**, 020602 (2024).

[18] T. Schuster, J. Haferkamp, and H.-Y. Huang, Random unitaries in extremely low depth, Science **389**, 92–96 (2025).

[19] H.-Y. Hu, A. Gu, S. Majumder, H. Ren, Y. Zhang, D. S. Wang, Y.-Z. You, Z. Minev, S. F. Yelin, and A. Seif, Demonstration of robust and efficient quantum property learning with shallow shadows, Nat. Commun. **16**, 2943 (2025).

[20] S. Notarnicola, A. Elben, T. Lahaye, A. Browaeys, S. Montangero, and B. Vermersch, A randomized measurement toolbox for an interacting Rydberg-atom quantum simulator, New J. Phys. **25**, 103006 (2023).

[21] A. A. Akhtar, H.-Y. Hu, and Y.-Z. You, Scalable and flexible classical shadow tomography with tensor networks, Quantum **7**, 1026 (2023).

[22] H.-Y. Hu and Y.-Z. You, Hamiltonian-driven shadow tomography of quantum states, Phys. Rev. Research **4**, 013054 (2022).

[23] T.-G. Zhou and P. Zhang, Efficient classical shadow tomography through many-body localization dynamics, Quantum **8**, 1467 (2024).

[24] Z. Liu, Z. Hao, and H.-Y. Hu, Predicting arbitrary state properties from single Hamiltonian quench dynamics, Phys. Rev. Research **6**, 043118 (2024).

[25] J. Helsen and M. Walter, Thrifty shadow estimation: Reusing quantum circuits and bounding tails, Phys. Rev. Lett. **131**, 240602 (2023).

[26] A. Acharya, S. Saha, and A. M. Sengupta, Shadow tomography based on informationally complete positive operator-valued measure, Phys. Rev. A **104**, 052418

(2021).
[27] L. Innocenti, S. Lorenzo, I. Palmisano, F. Albarelli, A. Ferraro, M. Paternostro, and G. M. Palma, Shadow tomography on general measurement frames, PRX Quantum **4**, 040328 (2023).
[28] 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).
[29] M. McGinley and M. Fava, Shadow tomography from emergent state designs in analog quantum simulators, Phys. Rev. Lett. **131**, 160601 (2023).
[30] J. Řeháček, B.-G. Englert, and D. Kaszlikowski, Minimal qubit tomography, Phys. Rev. A **70**, 052321 (2004).
[31] R. Stricker, M. Meth, L. Postler, C. Edmunds, C. Ferrie, R. Blatt, P. Schindler, T. Monz, R. Kueng, and M. Ringbauer, Experimental single-setting quantum state tomography, PRX Quantum **3**, 040310 (2022).
[32] Z. You, Q. Liu, and Y. Zhou, Circuit optimization of informationally complete positive operator-valued qubit measurements for shadow estimation, Phys. Rev. Applied **23**, 014021 (2025).
[33] J. M. Renes, R. Blume-Kohout, A. J. Scott, and C. M. Caves, Symmetric informationally complete quantum measurements, J. Math. Phys. **45**, 2171–2180 (2004).
[34] A. J. Scott, Tight informationally complete quantum measurements, J. Phys. A: Math. Gen. **39**, 13507–13530 (2006).
[35] A. J. Scott and M. Grassl, Symmetric informationally complete positive-operator-valued measures: A new computer study, J. Math. Phys. **51**, 042203 (2010).
[36] C. A. Fuchs, M. C. Hoang, and B. C. Stacey, The SIC question: History and state of play, Axioms **6**, 21 (2017).
[37] P. Horodecki, Ł. Rudnicki, and K. Życzkowski, Five open problems in quantum information theory, PRX Quantum **3**, 010101 (2022).
[38] M. Appleby, S. T. Flammia, and G. S. Kopp, A constructive approach to Zauner's conjecture via the Stark conjectures, arXiv:2501.03970 (2025).
[39] M. Neumark, Spectral functions of a symmetric operator, Izv. Akad. Nauk SSSR Ser. Mat. **4**, 277–318 (1940).
[40] A. Peres, Neumark's theorem and quantum inseparability, Found. Phys. **20**, 1441–1453 (1990).
[41] G. M. D'Ariano, P. Perinotti, and M. F. Sacchi, Informationally complete measurements and group representation, J. Opt. B: Quantum Semiclass. Opt. **6**, S487–S491 (2004).
[42] T. Decker, D. Janzing, and M. Rötteler, Implementation of group-covariant positive operator valued measures by orthogonal measurements, J. Math. Phys. **46**, 012104 (2005).
[43] T. Singal, F. B. Maciejewski, and M. Oszmaniec, Implementation of quantum measurements using classical resources and only a single ancillary qubit, npj Quantum Inf. **8**, 82 (2022).
[44] W. Cai, J. Han, L. Hu, Y. Ma, X. Mu, W. Wang, Y. Xu, Z. Hua, H. Wang, Y. P. Song, J.-N. Zhang, C.-L. Zou, and L. Sun, High-efficiency arbitrary quantum operation on a high-dimensional quantum system, Phys. Rev. Lett. **127**, 090504 (2021).
[45] P. Ivashkov, G. Uchehara, L. Jiang, D. S. Wang, and A. Seif, High-fidelity, multiqubit generalized measurements with dynamic circuits, PRX Quantum **5**, 030315 (2024).
[46] S. Gupta and M. B. Weiss, A simple realization of Weyl–Heisenberg covariant measurements, arXiv:2512.22111 (2025).
[47] X. Zhu and Y. Wang, Uniformly stable minimal Weyl–Heisenberg measurements approaching the SIC benchmark, arXiv:2608.11850 (2026).
[48] N. Bent, H. Qassim, A. A. Tahir, D. Sych, G. Leuchs, L. L. Sánchez-Soto, E. Karimi, and R. W. Boyd, Experimental realization of quantum tomography of photonic qudits via symmetric informationally complete positive operator-valued measures, Phys. Rev. X **5**, 041006 (2015).
[49] G. N. M. Tabia, Experimental scheme for qubit and qutrit symmetric informationally complete positive operator-valued measurements using multiport devices, Phys. Rev. A **86**, 062107 (2012).
[50] Z. E. D. Medendorp, F. A. Torres-Ruiz, L. K. Shalm, G. N. M. Tabia, C. A. Fuchs, and A. M. Steinberg, Experimental characterization of qutrits using symmetric informationally complete positive operator-valued measurements, Phys. Rev. A **83**, 051801(R) (2011).
[51] L.-T. Feng, X.-M. Hu, M. Zhang, Y.-J. Cheng, C. Zhang, Y. Guo, Y.-Y. Ding, Z. Hou, F.-W. Sun, G.-C. Guo, D.-X. Dai, A. Tavakoli, X.-F. Ren, and B.-H. Liu, Higher-dimensional symmetric informationally complete measurement via programmable photonic integrated optics, Optica **12**, 1014–1019 (2025).
[52] W.-Z. Yan, L.-T. Feng, Z. Hou, Y.-Y. Zhao, C. Roch i Carceller, A. Tavakoli, H. Zhu, G.-C. Guo, X.-F. Ren, and G.-Y. Xiang, A single programmable photonic circuit for universal quantum measurements, arXiv:2603.20085 (2026).
[53] Y.-Y. Zhao, N.-K. Yu, P. Kurzyński, G.-Y. Xiang, C.-F. Li, and G.-C. Guo, Experimental realization of generalized qubit measurements based on quantum walks, Phys. Rev. A **91**, 042101 (2015).
[54] Z. Bian, J. Li, H. Qin, X. Zhan, R. Zhang, B. C. Sanders, and P. Xue, Realization of single-qubit positive-operator-valued measurement via a one-dimensional photonic quantum walk, Phys. Rev. Lett. **114**, 203602 (2015).
[55] X. Wang, X. Zhan, Y. Li, L. Xiao, G. Zhu, D. Qu, Q. Lin, Y. Yu, and P. Xue, Generalized quantum measurements on a higher-dimensional system via quantum walks, Phys. Rev. Lett. **131**, 150803 (2023).
[56] L. E. Fischer, D. Miller, F. Tacchino, P. K. Barkoutsos, D. J. Egger, and I. Tavernelli, Ancilla-free implementation of generalized measurements for qubits embedded in a qudit space, Phys. Rev. Research **4**, 033027 (2022).
[57] E. Andersson and D. K. L. Oi, Binary search trees for generalized measurements, Phys. Rev. A **77**, 052104 (2008).
[58] M. Farkas, J. Volčič, S. A. L. Storgaard, R. Chen, and L. Mančinska, Maximal device-independent randomness in every dimension, Nat. Phys. **22**, 319–324 (2026).
[59] H. Zhu and B.-G. Englert, Quantum state tomography with fully symmetric measurements and product measurements, Phys. Rev. A **84**, 022327 (2011).
[60] Q. Zhang, Q. Liu, and Y. Zhou, Minimal-Clifford shadow estimation by mutually unbiased bases, Phys. Rev. Applied **21**, 064001 (2024).
[61] Y. Wang and W. Cui, Classical shadow tomography with mutually unbiased bases, Phys. Rev. A **109**, 062406 (2024).
[62] Z. Yang, D. Chen, and H. Zhu, Optimal shadow estimation with minimal measurement settings, arXiv:2606.20003 (2026).
[63] S. Chen, J. Cotler, H.-Y. Huang, and J. Li, Exponential separations between learning with and without quantum memory, in *2021 IEEE 62nd Annual Symposium on*

Foundations of Computer Science (FOCS) (IEEE, 2022), pp. 574–585.

[64] S. Chen, W. Gong, and Q. Ye, Optimal tradeoffs for estimating Pauli observables, in *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, 2024), pp. 1086–1105.

[65] R. King, D. Gosset, R. Kothari, and R. Babbush, Triply efficient shadow tomography, PRX Quantum **6**, 010336 (2025).

[66] S. T. Flammia and Y.-K. Liu, Direct fidelity estimation from few Pauli measurements, Phys. Rev. Lett. **106**, 230501 (2011).

[67] M. P. da Silva, O. Landon-Cardinal, and D. Poulin, Practical characterization of quantum devices without tomography, Phys. Rev. Lett. **107**, 210404 (2011).

[68] A. Elben, B. Vermersch, R. van Bijnen, C. Kokail, T. Brydges, C. Maier, M. K. Joshi, R. Blatt, C. F. Roos, and P. Zoller, Cross-platform verification of intermediate scale quantum devices, Phys. Rev. Lett. **124**, 010504 (2020).

[69] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, Sample-optimal tomography of quantum states, IEEE Trans. Inf. Theory **63**, 5628–5641 (2017).

[70] Y. Wang, Quantum advantage via efficient postprocessing on qudit classical shadow tomography, Phys. Rev. Lett. **135**, 200601 (2025).

[71] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, Phys. Rev. A **70**, 052328 (2004).

[72] H. J. García, I. L. Markov, and A. W. Cross, On the geometry of stabilizer states, Quantum Inf. Comput. **14**, 683–720 (2014).

[73] S. Ghosh, A. Deshpande, D. Hangleiter, A. V. Gorshkov, and B. Fefferman, Complexity phase transitions generated by entanglement, Phys. Rev. Lett. **131**, 030601 (2023).

[74] R. Iten, R. Colbeck, I. Kukuljan, J. Home, and M. Christandl, Quantum circuits for isometries, Phys. Rev. A **93**, 032318 (2016).

[75] E. Bäumer and S. Woerner, Measurement-based long-range entangling gates in constant depth, Phys. Rev. Research **7**, 023120 (2025).

[76] W. Zi, J. Nie, and X. Sun, Constant-depth quantum circuits for arbitrary quantum state preparation via measurement and feedback, arXiv:2503.16208 (2025).

[77] R. Luo, J. Chen, and X. Ma, Constant-depth adaptive preparation of Dicke and symmetric states, arXiv:2608.01144 (2026).

[78] K. C. Smith, A. Khan, B. K. Clark, S. M. Girvin, and T.-C. Wei, Constant-depth preparation of matrix product states with adaptive quantum circuits, PRX Quantum **5**, 030344 (2024).

[79] Y. Li, A. Mitra, and T.-C. Lu, State preparation via measurement and feedback: Pushing relations, state structures, and non-invertible symmetries, arXiv:2608.08821 (2026).

[80] A. M. Childs and N. Wiebe, Hamiltonian simulation using linear combinations of unitary operations, Quantum Inf. Comput. **12**, 901–924 (2012).

[81] C. Moore and M. Nilsson, Parallel quantum computation and quantum codes, SIAM J. Comput. **31**, 799–815 (2001).

[82] H. Buhrman, M. Folkertsma, B. Loff, and N. M. P. Neumann, State preparation by shallow circuits using feed forward, Quantum **8**, 1552 (2024).

[83] Y. Takahashi and S. Tani, Collapse of the hierarchy of constant-depth exact quantum circuits, Comput. Complexity **25**, 849–881 (2016).

[84] F. Vatan and C. Williams, Optimal quantum circuits for general two-qubit gates, Phys. Rev. A **69**, 032315 (2004).

[85] H.-R. Wei and Y.-M. Di, Decomposition of orthogonal matrix and synthesis of two-qubit and three-qubit orthogonal gates, Quantum Inf. Comput. **12**, 262–270 (2012).

[86] M. Žnidarič, O. Giraud, and B. Georgeot, Optimal number of controlled-NOT gates to generate a three-qubit state, Phys. Rev. A **77**, 032320 (2008).

[87] P. Høyer, On arbitrary phases in quantum amplitude amplification, Phys. Rev. A **62**, 052304 (2000).

[88] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, Phys. Rev. A **52**, 3457–3467 (1995).

[89] J. Nie, W. Zi, and X. Sun, Quantum circuit for multi-qubit Toffoli gate with optimal resource, arXiv:2402.05053 (2024).

[90] T. Khattar and C. Gidney, Rise of conditionally clean ancillae for efficient quantum circuit constructions, Quantum **9**, 1752 (2025).

[91] B. Zindorf and S. Bose, Efficient implementation of multi-controlled quantum gates, Phys. Rev. Applied **24**, 044030 (2025).