INFORMATION-CALIBRATED QUANTUM DIFFUSION: LINKING FORWARD INFORMATION LOSS TO REVERSE RECOVERABILITY

Anonymous authors
Paper under double-blind review

ABSTRACT

Quantum diffusion models generate distributions over quantum states by learning to reverse a prescribed sequence of noisy quantum channels. Existing methods typically organize the forward process by a scalar noise strength, such as the depolarizing strength, rather than by the information actually lost from the state ensemble. As a result, this noise parameter does not directly indicate how difficult the corresponding reverse step should be. We introduce the classical–quantum information decrement $\Delta_t = I(X{:}Q_{t-1}) - I(X{:}Q_t)$ for depolarizing channels to directly connect forward scheduling with reverse recovery. We prove that the same decrement sets an attainable common-channel recovery budget and, at fixed dimension and ensemble size, vanishes exactly when the optimal common-channel recovery error vanishes; equalizing these decrements minimizes the largest one-step information loss. We further show that local recovery alone leaves the generated distribution underdetermined, while distribution matching supplies the missing distribution-level constraint. Together, these results recast quantum diffusion as a sequence of information-controlled inverse problems, where the forward schedule determines how reverse difficulty is allocated rather than merely how noise is applied. Rather than prescribing a noise schedule and learning its inverse afterward, information calibration makes the recovery burden induced by the forward process explicit: it allocates information loss across timesteps, assigns each step an attainable common-channel recovery reference, and motivates the joint design of forward scheduling, reverse recovery, and distribution-level generation.

1 INTRODUCTION

Learning to generate quantum states with a prescribed distribution is a central task in quantum generative modeling, with applications to modeling quantum noise and structured ensembles of many-body states (Zhang et al., 2024). Inspired by classical diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020), quantum diffusion approaches this task by progressively corrupting quantum data and learning a sequence of reverse operations that transforms a simple reference ensemble into the target distribution (Zhang et al., 2024; Kwun et al., 2025). This stepwise construction offers a route to learning complex quantum ensembles through intermediate denoising tasks rather than a single generative transformation. Its design, however, involves more than choosing a neural architecture: the forward corruption and learned reverse operations must respect quantum dynamics, and the forward schedule determines the inverse problems that the learner must solve.

Recent work in classical diffusion has begun to allocate steps according to entropy or information change rather than noise strength alone (Stancevic et al., 2025; Raya et al., 2026), suggesting that a schedule should reflect the information-processing demands along the diffusion path. In quantum diffusion, this perspective raises an additional requirement: the information removed by each forward step must be related to what a common physical reverse operation can recover. A raw channel parameter such as depolarizing strength does not directly quantify either the ensemble information lost or the accuracy attainable by a common recovery channel. A noise-centric schedule can therefore allocate steps unevenly in information space, without an explicit reference for the recovery burden assigned to a finite-capacity learner. The challenge is not merely to replace one noise schedule with another, but to connect forward discretization with physically attainable reverse recovery. This leads

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

to our central question: can a single information quantity determine how the forward process should be discretized while setting an attainable recovery scale for each reverse step?

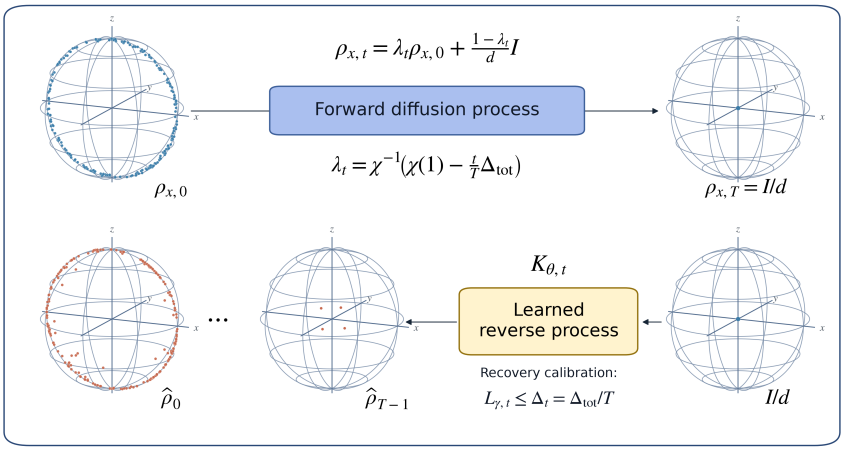


Figure 1: **Overview of the information-calibrated quantum diffusion framework.** Target states $\rho_{x,0}$ are progressively depolarized toward the maximally mixed state on an equal-information grid. The learned reverse process aims to recover target distribution. Per-step information decrement also sets an attainable common-channel recovery scale, linking forward scheduling with reverse recovery.

We answer this question using the classical–quantum information (Holevo information) retained about the ensemble label, as illustrated in Figure 1. For an ensemble $\mathcal{E} = \{p_x, \rho_x\}$, let X denote the classical label identifying which state ρ_x is drawn from the ensemble. We then define the one-step information decrement

$$\Delta_t := I(X{:}Q_{t-1}) - I(X{:}Q_t). \tag{1}$$

Along depolarization, the same decrement has a two-sided operational recovery meaning. Building on universal-recovery results (Buscemi et al., 2016; Junge et al., 2018), there exists a single label-independent completely positive trace-preserving (CPTP) recovery channel whose expected log-fidelity error is no larger than Δ_t. Conversely, at fixed Hilbert-space dimension and ensemble size, Δ_t vanishes if and only if the optimal common-channel recovery error vanishes. The same coordinate also defines the forward discretization: equalizing Δ_t yields the unique monotone grid that minimizes the largest one-step information loss.

Local recoverability, however, is not the whole generative problem. Quantum diffusion models ultimately need to recover the target distribution over quantum states, and previous work therefore relies on distribution-level objectives in addition to state-wise similarity (Zhang et al., 2024; Kwun et al., 2025). We show that this distinction remains essential after the local inverse has been physically calibrated: in a fixed noncommuting two-qubit construction, two stochastic CPTP reverse models can have identical local-fidelity statistics and both satisfy the same information-calibrated recovery budget while producing substantially different output distributions. Thus, stepwise physical recovery and generative distribution matching are complementary, non-substitutable requirements: the former controls the physical accuracy of individual reverse steps, while the latter ensures that these locally accurate steps collectively recover the target quantum-state distribution.

Using reverse architectures based on MSQuDDPM (Kwun et al., 2025), our experiments study a one-qubit mixed-state ring and four-qubit TFIM ensembles. Larger information decrements track greater attained recovery error, supporting the proposed information coordinate. The ring study distinguishes improved local recovery from endpoint generation, demonstrating the complementary role of distribution matching. Across both tasks, the effect of recovery calibration on generation depends on the forward schedule, supporting a joint approach in which information- guided forward process and reverse recovery are designed together rather than optimized in isolation.

Three main contributions of the paper can be summarized as

- **A two-sided information coordinate for recoverability.** We prove that Δ_t sets an attainable common-channel recovery budget and, at fixed Hilbert-space dimension and ensemble size, vanishes

if and only if the optimal common-channel recovery error vanishes. Equalizing this coordinate then gives the unique grid minimizing the largest one-step information loss.

- **A sharp separation between local recovery and distribution matching.** We show that in a noncommuting two-qubit construction, identical local-fidelity laws and the same budget-feasible local loss coexist with macroscopically different output distributions. This indicates that recovery calibration cannot determine stochastic generation. Distribution matching therefore provides a complementary constraint for stochastic quantum generation.
- **A recovery-constrained stochastic learner with local-to-global guarantees.** We combine distribution matching with timestep-wise recovery constraints, establish finite-sample calibration and compositional trace-Wasserstein control, and evaluate the framework through controlled comparisons using an MSQuDDPM-style reverse architecture.

2 Related work

Classical diffusion has recently moved beyond schedules defined purely by noise strength toward entropy- or information-based parameterizations (Stancevic et al., 2025; Raya et al., 2026; Ambrogioni et al., 2026; Foresti et al., 2026). These works motivate allocating diffusion steps according to information change, but they do not address physically constrained recovery by quantum channels. Quantum diffusion models, meanwhile, have developed depolarizing forward processes, stochastic or measurement-conditioned reverse dynamics, and ensemble-level generative objectives (Parigi et al., 2025; Huang et al., 2025; Zhang et al., 2024; Chen et al., 2026; Kwun et al., 2025; Liu et al., 2025; Xu et al., 2026; Bompais et al., 2026). Recent measurement-based and reverse-time constructions (Liu et al., 2025; Bompais et al., 2026) primarily address how reverse dynamics are realized for a specified forward evolution; our focus is complementary: the forward classical–quantum information decrement determines both where that evolution is discretized and an attainable recovery scale for a common CPTP reverse channel.

Relating the information lost in a forward step to the accuracy achievable by reverse recovery builds on standard results from quantum information theory, including Holevo information, relative-entropy data processing, and universal recovery (Holevo, 1973; Lindblad, 1975; Petz, 1986; Fawzi & Renner, 2015; Junge et al., 2018; Buscemi et al., 2016; Shirokov, 2017). These results are standard recovery ingredients. Our contribution is to use them to connect forward scheduling and reverse learning in quantum diffusion. The same Δ_t used to construct the forward schedule also gives an attainable error budget for a common reverse channel. Conversely, if the optimal common-channel recovery error vanishes, then the information loss must also vanish. Thus, Δ_t provides a physically meaningful reference scale for each reverse step without prescribing the dynamics that the learned reverse model must implement.

Quantum-ensemble distances explain why independent state pairing does not define ensemble distinguishability (Oreshkov & Calsamiglia, 2009). Our separation sharpens this observation quantitatively and ties it to physical recovery calibration: a perfectly covering generator and a collapsed generator can share the complete marginal fidelity law and the same budget-feasible local log-risk on a fixed noncommuting two-qubit ensemble. This connects to reconstruction–distribution tensions in classical inverse problems (Blau & Michaeli, 2018), with the local criterion here fixed by quantum recoverability.

3 Information-calibrated diffusion and recoverability

We consider a nontrivial finite ensemble $\mathcal{E} = \{p_x, \rho_x\}$ on a $d \geq 2$ dimensional Hilbert space, with $p_x > 0$. The depolarizing forward process maps

$$\rho_{x,t} = \mathcal{D}_{\lambda_t}(\rho_x) = \lambda_t \rho_x + (1 - \lambda_t)\mathbb{I}/d, \tag{2}$$

where smaller λ_t means stronger depolarization. The reverse model is stochastic and aims to recover the target distribution over quantum states, not only individual states. For two states ρ and σ, we use root fidelity $F_{\mathrm{r}}(\rho, \sigma) := \left\|\sqrt{\rho}\sqrt{\sigma}\right\|_1$ to quantify their overlap, and trace distance $d_{\mathrm{tr}}(\rho, \sigma) := \frac{1}{2}\left\|\rho - \sigma\right\|_1$ to quantify their distinguishability. Both take values in $[0, 1]$. For probability measures μ and ν on quantum states, we compare the two state distributions by the trace-cost

Wasserstein distance

$$W_{1,\mathrm{tr}}(\mu,\nu) := \inf_{\pi\in\Pi(\mu,\nu)} \int d_{\mathrm{tr}}(\rho,\sigma)\,\pi(\mathrm{d}\rho,\mathrm{d}\sigma), \tag{3}$$

where $\Pi(\mu,\nu)$ is the set of couplings with marginals μ and ν. Thus $W_{1,\mathrm{tr}}$ is the smallest expected trace distance over all pairings of the two state distributions.

The information retained about the ensemble label is measured by the classical–quantum mutual information, which for the cq ensemble state equals the Holevo information,

$$\chi(\mathcal{E}) = I(X{:}Q) = S(\bar{\rho}) - \sum_x p_x S(\rho_x), \qquad \bar{\rho} = \sum_x p_x \rho_x, \tag{4}$$

where X is the classical ensemble label, $\bar{\rho}$ is the average state, and $S(\rho) = -\operatorname{Tr}(\rho\log\rho)$ is the von Neumann entropy. Keeping the label X is important because two ensembles can have the same average density operator while defining different probability laws over states.

Our goal is to use this retained information, rather than the raw depolarizing strength, as a common coordinate for forward spacing and reverse recoverability. For a state ρ, write $\mathcal{D}_\lambda(\rho) = \lambda\rho + (1-\lambda)\mathbb{I}/d$. The quantity $\chi(\lambda) := \chi(\{p_x, \mathcal{D}_\lambda(\rho_x)\})$ is the Holevo information of the ensemble after depolarization to retention level λ, and records how much ensemble information remains along the forward path.

We fix $\lambda_0 = 1$ and $\lambda_T < 1$, with total information loss $\Delta_{\mathrm{tot}} := \chi(1) - \chi(\lambda_T)$. The first consequence concerns forward allocation: for fixed T, the following theorem gives the unique grid that equalizes this loss while minimizing the largest one-step decrement.

Theorem 1 (Equal-information depolarizing schedule)**.** *For any ensemble with at least two distinct states, $\chi(\lambda)$ is continuous and strictly increasing on $[0,1]$. For fixed integer $T \geq 1$, $\lambda_0 = 1$, and $0 \leq \lambda_T < 1$, there is a unique monotone grid satisfying*

$$\Delta_t := \chi(\lambda_{t-1}) - \chi(\lambda_t) = \Delta_{\mathrm{tot}}/T, \qquad \lambda_t = \chi^{-1}\!\left(\chi(1) - \frac{t}{T}\Delta_{\mathrm{tot}}\right). \tag{5}$$

It uniquely minimizes $\max_t \Delta_t$ and, more generally, $\sum_t \varphi(\Delta_t)$ for every strictly convex φ.

Using this schedule, the forward process is $\rho_{x,t} := \mathcal{D}_{\lambda_t}(\rho_x)$ for $t = 0,\ldots,T$.

The decrement Δ_t quantifies the information erased by the tth forward step. However, the scheduling statement alone does not characterize reverse difficulty. Since recovery cannot depend on the hidden label x, we next ask what the same Δ_t implies for a single label-independent CPTP recovery channel $\mathcal{R}_t$.

Theorem 2 (Attainable common-channel recovery budget)**.** *For every step there exists a common recovery channel $\mathcal{R}_t$, depending only on $\bar{\rho}_{t-1}$ and the known forward step, such that*

$$-2\sum_x p_x \log F_{\mathrm{r}}\left(\rho_{x,t-1}, \mathcal{R}_t(\rho_{x,t})\right) \leq \Delta_t. \tag{6}$$

Thus Δ_t gives an attainable upper bound on the expected log-fidelity error of a common recovery map, whereas λ_t only marks position along the depolarizing path. To assess how close this guarantee is to the best possible recovery, we next minimize trace-distance error over all common recovery channels.

Theorem 3 (Two-sided common-channel recoverability)**.** *For the tth forward step, define the optimal recovery error over all label-independent CPTP channels by $\varepsilon_t^\star := \inf_{\mathcal{R}\in\mathrm{CPTP}} \sum_x p_x d_{\mathrm{tr}}(\rho_{x,t-1}, \mathcal{R}(\rho_{x,t}))$. Then*

$$\Delta_t \leq \Omega_{d,m}(\varepsilon_t^\star), \qquad \varepsilon_t^\star \leq \sqrt{1 - e^{-\Delta_t}}. \tag{7}$$

where $\Omega_{d,m}$ is the Holevo-continuity modulus defined in Appendix A.2. Consequently, at fixed Hilbert-space dimension d and ensemble size m, $\Delta_t \to 0$ if and only if $\varepsilon_t^\star \to 0$.

Thus Δ_t not only gives an attainable recovery scale through Theorem 2, but also characterizes when the optimal common-channel recovery error becomes small. A pairwise lower bound is given in Appendix B.1.

4 LOCAL RECOVERY AND DISTRIBUTION MATCHING

Learning a distribution of quantum states is the central objective in quantum diffusion models. QuDDPM uses ensemble-level objectives such as maximum mean discrepancy (MMD) and Wasserstein distance (Zhang et al., 2024). Importantly, QuDDPM gives a counterexample showing that fidelity-kernel MMD can vanish for two distinct state distributions. MSQuDDPM extends the same distribution-level viewpoint to mixed-state generation (Kwun et al., 2025).

Section 3 provides a physically calibrated criterion for local recovery, with its scale set by the information decrement Δ_t. The question here is whether satisfying this criterion is enough to recover the correct distribution over quantum states. We show that it is not: a reverse model can satisfy the prescribed local recovery requirement while still assigning probability mass incorrectly across the target ensemble. Local recovery and distribution matching are therefore distinct requirements for stochastic generation.

A simple symmetric construction already shows the mechanism: a correctly covering generator and a collapsed generator can have the same distribution of local target–output scores while assigning probability mass differently across the target states. The formal statement is given in Appendix A.3. The remaining question is whether this blindness persists for the root-fidelity criterion used by our information-calibrated recovery theory while both models satisfy the same Δ-based physical recovery budget. The next theorem shows that it does, in a fixed genuinely quantum system.

Theorem 4 (Local recovery does not determine the generated distribution in a fixed two-qubit system)**.** *There exists a one-step complete-depolarization problem on two qubits with* 16 *equiprobable pairwise noncommuting target states and two physically valid stochastic CPTP reverse models with the following properties. One model reproduces the target distribution exactly, whereas the other is completely collapsed onto a single target state. Nevertheless,*

(i) the two models have identical local root-fidelity distributions,

$$F_{\mathrm{r}}(\rho_X, \widehat{\rho}^{\mathrm{good}}) \overset{\mathrm{law}}{=} F_{\mathrm{r}}(\rho_X, \widehat{\rho}^{\mathrm{col}}); \tag{8}$$

(ii) their expected local log-fidelity losses are also identical and both satisfy the exact information-calibrated recovery budget,

$$-2\mathbb{E}\log F_{\mathrm{r}}(\rho_X, \widehat{\rho}^{\mathrm{good}}) = -2\mathbb{E}\log F_{\mathrm{r}}(\rho_X, \widehat{\rho}^{\mathrm{col}}) \le \Delta, \quad \Delta = \frac{1}{2}\log 3; \tag{9}$$

(iii) their generated distributions are nevertheless macroscopically different:

$$W_{1,\mathrm{tr}}\left(\mu, \mathrm{Law}(\widehat{\rho}^{\mathrm{good}})\right) = 0, \quad W_{1,\mathrm{tr}}\left(\mu, \mathrm{Law}(\widehat{\rho}^{\mathrm{col}})\right) = \frac{\sqrt{5}}{4} > \frac{1}{2}, \tag{10}$$

where $\mu = \frac{1}{16}\sum_{x=1}^{16} \delta_{\rho_x}$ is the target ensemble distribution.

The theorem makes the distinction between the two objectives explicit inside the information-calibrated quantum diffusion setting. The Δ-based criterion answers a local physical question: whether the reverse outputs are compatible with the recoverability scale induced by the forward information loss. It does not answer the generative question of whether the stochastic reverse process reproduces the correct probability law over quantum states. In our two-qubit noncommuting construction, the latter can fail severely even when the complete local-fidelity statistics and the information-calibrated recovery budget are indistinguishable. The explicit SIC-based construction and proof are given in Appendix A.3.

This gap can become arbitrarily large in higher-dimensional systems: for every $c \in (0, 1)$, the construction can be chosen so that one model has zero trace-Wasserstein error while the other exceeds c, despite identical local-fidelity laws and budget-feasible local losses (Appendix A.3).

Thus the role of the learning objective is now separated into two parts. The information-calibrated constraint sets a physically meaningful local recovery target for each reverse step, while a distribution objective controls the probability law of the stochastic outputs.

5 INFORMATION-CALIBRATED STOCHASTIC GENERATION

Section 4 shows that local recovery and distribution matching are non-substitutable. The recovery results in Section 3 set a reference scale through a common CPTP channel; the learned stochastic

reverse process need not itself be a single CPTP map. We therefore formulate the learner as a finite-outcome quantum instrument, which includes the measurement-conditioned reverse steps used by MSQuDDPM (Kwun et al., 2025).

At timestep t, let $\{\mathcal{M}_{\theta,t,b}\}_b$ be completely positive maps with $\sum_b \mathcal{M}_{\theta,t,b}$ trace-preserving. For an input state ρ, the induced state-valued reverse kernel is

$$K_{\theta,t}(\rho,\cdot) = \sum_{b:\, p_b(\rho)>0} p_b(\rho)\, \delta_{\mathcal{M}_{\theta,t,b}(\rho)/p_b(\rho)}, \qquad p_b(\rho) := \operatorname{Tr} \mathcal{M}_{\theta,t,b}(\rho). \tag{11}$$

For a probability measure μ on input states, we let $K_{\theta,t}\mu$ denote the corresponding output law,

$$(K_{\theta,t}\mu)(B) := \int K_{\theta,t}(\rho, B)\, \mu(\mathrm{d}\rho) \tag{12}$$

for every measurable set B of states. Thus a measurement outcome is sampled with its Born probability and the normalized conditional system state is propagated. Input-independent stochastic CPTP branches are the special case $\mathcal{M}_{\theta,t,b} = \pi_t(b)\mathcal{R}_{\theta,t,b}$. A normalized measurement-conditioned branch need not itself be a linear CPTP channel. The circuit realization used by MSQuDDPM is given in Appendix B.12.

Given $X \sim p$ and $\widehat{\rho}_{t-1} \sim K_{\theta,t}(\rho_{X,t}, \cdot)$, retain the population local recovery loss

$$L_{\gamma,t}(\theta) = \mathbb{E}[-2\log\max\{\gamma, F_{\mathrm{r}}(\rho_{X,t-1}, \widehat{\rho}_{t-1})\}], \qquad 0<\gamma<1. \tag{13}$$

The expectation is over the actual output law in Equation (11); equivalently, it is Born-probability weighted over the conditional outputs. A common recovery channel is a one-outcome instrument, so Theorem 2 implies that the level $L_{\gamma,t} \le \Delta_t$ is attainable over the unrestricted instrument class. This is a reference budget: a restricted circuit need not attain it.

Let $\widehat{\mathcal{J}}_{\mathrm{dist}}(\theta)$ denote an empirical distribution-matching objective. The learning problem is

$$\min_\theta \widehat{\mathcal{J}}_{\mathrm{dist}}(\theta) \quad \text{s.t.} \quad \widehat{L}_{\gamma,t}(\theta) \le \widehat{\Delta}_t, \qquad t = 1,\dots,T. \tag{14}$$

This includes MMD-based objectives as well as the ensemble optimal-transport objective used in MSQuDDPM. The recovery constraint is always evaluated on the label-paired forward states $(\rho_{x,t}, \rho_{x,t-1})$, even when the distribution objective is trained on generated reverse inputs.

Theorem 5 (Uniform finite-sample calibration)**.** *Consider the bounded-generator instrument class in Appendix A.4, with P trainable rotations and n i.i.d. training examples, and suppose the forward grid is fixed independently of these examples. If* $\max_t \left[\widehat{L}_{\gamma,t}(\widehat{\theta}) - \widehat{\Delta}_t\right]_+ \le \eta_{\mathrm{train}}$*, define*

$$\varepsilon_{\mathrm{gen}} := \frac{4\log(1/\gamma)}{\sqrt{n}} + 2\log(1/\gamma)\sqrt{\frac{P\log\left(1+\frac{6\pi Pn}{\gamma^2\log^2(1/\gamma)}\right)+\log(T/\delta)}{2n}}. \tag{15}$$

Then, with probability at least $1-\delta$, $L_{\gamma,t}(\widehat{\theta}) \le \Delta_t + \varepsilon_{\mathrm{cal}}$ for every t, where

$$\varepsilon_{\mathrm{cal}} := \eta_{\mathrm{train}} + \varepsilon_{\mathrm{gen}} + \max_t |\widehat{\Delta}_t - \Delta_t|. \tag{16}$$

No positive lower bound on the measurement-outcome probabilities is required. For input-independent CPTP branches, the original sharper constant is recovered; see Appendix A.4.

A data-estimated grid can instead be certified by an independent pilot or by the validation results in Appendix B.9. In the finite-support experiments, the corresponding population calibration quantities can also be evaluated directly.

To connect the local risks to generation, let $\mu_t = \sum_x p_x \delta_{\rho_{x,t}}$ be the target distribution and let ν_t be the distribution produced by the learned reverse chain. A local bound $L_{\gamma,t} \le a$ implies the one-step trace-Wasserstein bound

$$W_{1,\mathrm{tr}}(\mu_{t-1}, K_{\theta,t}\mu_t) \le \psi_\gamma(a), \qquad \psi_\gamma(a) = \sqrt{1-\left[e^{-a/2} - \frac{\gamma a}{2\log(1/\gamma)}\right]_+^2}. \tag{17}$$

The proof uses the actual conditional-output law, not the averaged channel. A single measurement-conditioned step need not be nonexpansive, but any finite sequence of such steps is itself a finite-outcome instrument. This gives the following endpoint guarantee without multiplying a factor at every timestep.

Theorem 6 (Compositional endpoint bound)**.** *For a fixed sequence of finite-outcome reverse instruments satisfying* $L_{\gamma,t}(\widehat{\theta}) \le \Delta_t + \varepsilon_{\mathrm{cal}}$*,*

$$W_{1,\mathrm{tr}}(\mu_0,\nu_0) \le 3W_{1,\mathrm{tr}}(\mu_T,\nu_T) + \psi_\gamma(\Delta_1+\varepsilon_{\mathrm{cal}}) + 3\sum_{t=2}^{T}\psi_\gamma(\Delta_t+\varepsilon_{\mathrm{cal}}). \tag{18}$$

For the input-independent latent-conditioned CPTP kernels used in the original formulation, the factor 3 *is replaced by* 1*, recovering the stronger bound*

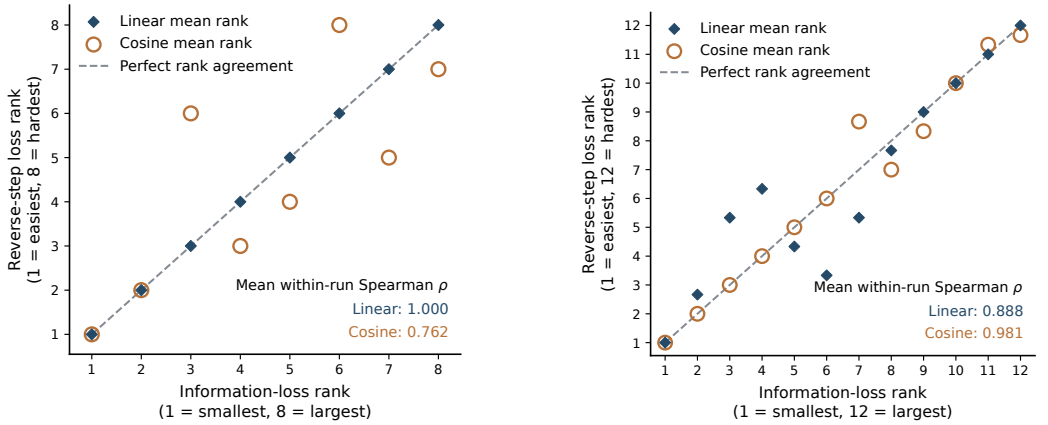
$$W_{1,\mathrm{tr}}(\mu_0,\nu_0) \le W_{1,\mathrm{tr}}(\mu_T,\nu_T) + \sum_{t=1}^{T}\psi_\gamma(\Delta_t+\varepsilon_{\mathrm{cal}}). \tag{19}$$

If $\nu_T = \delta_{\mathbb{I}/d}$*, then* $W_{1,\mathrm{tr}}(\mu_T,\nu_T) \le \lambda_T(1-1/d)$.

The complete proof is in Appendix A.5. At fixed terminal mismatch and common calibration slack, equal information minimizes the corresponding worst-step-derived bound. This is a one-sided certificate, not a claim that it minimizes the actual trained endpoint error, and distribution matching remains independently necessary by Theorem 4.

6 Experiments

Using a mixed-state ring as the primary testbed, we examine whether Holevo information tracks learned recovery difficulty, then study how scheduling, recovery calibration, and distribution matching affect generation. Supplementary four-qubit TFIM experiments probe the information–recovery relationship and the joint role of scheduling and recovery training.



(a) One-qubit ring (b) Four-qubit TFIM

Figure 2: **Information loss tracks learned recovery difficulty.** Steps are ranked within each run by Δ_t and $L_{\gamma,t}^{\mathrm{pop}}$. Diamonds and circles show mean reverse-loss ranks across 10 seeds for linear and cosine schedules, respectively; dashed lines indicate perfect rank agreement.

6.1 Information loss and recovery difficulty

To examine whether the proposed information coordinate captures reverse-step difficulty, we train independent recovery maps along linear and cosine paths on the 200-state mixed-state ring ($T = 8$), using 10,000 Adam updates per stage. A supplementary four-qubit TFIM study uses the same recovery-only objective. Neither information constraints nor distribution matching enter training, so the relationship with Δ_t is assessed without explicitly enforcing it. We exclude equal-information grids because their decrements are constant.

As shown in Figure 2, larger information losses are associated with harder recovery steps across both schedules and tasks, with strong positive correlations in every run. This agreement gives the Holevo coordinate an empirical role as a measure of learned recovery difficulty: it links information lost along the forward path to the recovery error attained by a finite-budget learner, supporting the intended connection between forward scheduling and reverse recovery.

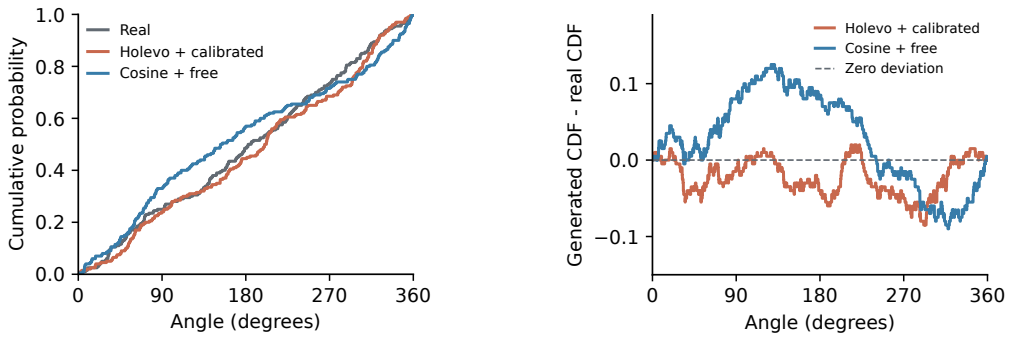
6.2 MIXED-STATE GENERATION

Using the same data, architecture, and update budget, we compare cosine and equal-information schedules under free or calibrated training, adding calibrated linear and equal-information local-only (OT-selected) baselines. Free minimizes MSQuDDPM $1 - G$ OT; calibrated adds Δ_t-based recovery constraints; local-only uses recovery gradients alone. We report endpoint $W_{1,\mathrm{tr}}$, maximum budget excess, and maximum local error; see Appendices C and C.2 for details.

Table 1: **Generation and local recovery on the mixed-state ring ($T = 8$).** The cosine/equal-information configurations examine scheduling and recovery calibration; calibrated linear extends the schedule comparison, while equal-information local-only isolates recovery without distribution matching. Free and calibrated denote distribution-only and combined training. Endpoint $W_{1,\mathrm{tr}}$ measures generation (reported with mean + s.e.); maximum excess and local trace error measure recovery. Results use 10 seeds and 200 evaluations; lower is better. Mean paired differences (row minus equal-information/calibrated) are reported with mean + 95% percentile CI.

Schedule	Training	Endpoint $W_{1,\mathrm{tr}}$	Paired difference	Max excess	Max local error
Linear	Calibrated	0.0980 ± 0.0019	$0.0223\ [0.0191, 0.0261]$	0.0010	0.8066
Cosine	Free	0.0779 ± 0.0010	$0.0022\ [0.0010, 0.0035]$	1.3588	0.9973
Cosine	Calibrated	0.0884 ± 0.0018	$0.0127\ [0.0099, 0.0154]$	0.0005	0.8075
Equal-information	Local-only	0.2747 ± 0.0054	$0.1990\ [0.1883, 0.2090]$	0.4302	0.8865
Equal-information	Free	0.0828 ± 0.0010	$0.0072\ [0.0049, 0.0096]$	1.0926	0.9921
Equal-information	Calibrated	$\mathbf{0.0757} \pm 0.0008$	—	0.0007	0.6658

Forward scheduling and reverse calibration should be considered jointly. Equal-information/calibrated attains the lowest endpoint error in Table 1 and outperforms displayed alternative on all 10 seeds. It reduces endpoint error by 14.4% compared with cosine/calibrated and by 8.6% compared with equal-information/free. The constrained forward comparison reinforces this pattern. With recovery calibration held fixed, equal-information also improves on linear retention by 22.8%. Conversely, equal-information/free does not outperform cosine/free. Thus neither the schedule nor the recovery constraint alone accounts for the best displayed result; their combination is favorable in this comparison, reinforcing the joint treatment of forward scheduling and reverse recovery underlying our theory and design.



(a) Angular cumulative distribution

(b) Deviation from the real distribution

Figure 3: **Angular distribution of generated ring states.** (a) Empirical cumulative distributions of the angles of real and generated states. (b) Generated-minus-real CDF differences; curves closer to zero indicate better agreement in cumulative angular mass.

Figure 3 illustrates one aspect of this endpoint comparison: equal-information/calibrated more closely reproduces the cumulative angular mass than cosine/free. This angular advantage need not imply radial advantage (Appendix C.3).

Better step recovery does not guarantee better endpoint generation. In Table 1, recovery constraints reduce maximum budget excess and maximum local error under both schedules, but these local gains do not guarantee better generation: cosine/calibrated has higher endpoint $W_{1,\mathrm{tr}}$ than cosine/free (0.088419 versus 0.077853). On the equal-information grid, local-only training also performs substantially worse than distribution-only or combined training. These results support the distinction in Section 4: recovery calibration improves local inversion but cannot replace distribution matching. The local-only baseline tests this distinction in practice, rather than realizing the exactly feasible theoretical counterexample.

Supplementary four-qubit generation. On 100 TFIM ground states ($g \in [0.2, 2.2)$), we use a depth-12, $T = 12$ instrument with two measured ancillas. Equal-information/calibrated gives the lowest mean endpoint error in Table 2, consistent with the joint-design perspective that forward information allocation and per-step recovery calibration should be considered together. Although substantial budget excess remains, a targeted continuation with stronger recovery emphasis reveals a tunable intermediate regime: mean maximum excess falls from 0.1640 to 0.0771 while endpoint $W_{1,\mathrm{tr}}$ rises from 0.0621 to 0.0858. Further emphasizing recovery reduces the violation at increasing generation cost (Appendix C.5).

Table 2: **Supplementary TFIM generation.** All 10 seeds and 200 evaluations are retained; endpoint entries are mean $\pm$ s.e., and recovery metrics are means. Endpoint scores average twenty 200-versus-200 evaluations per model. Lower is better.

Schedule	Training	Endpoint $W_{1,\mathrm{tr}}$	Max excess	Max local error
Cosine	Free	0.1218 ± 0.0120	0.1654	0.3251
Cosine	Calibrated	0.1187 ± 0.0097	0.1515	0.3165
Equal-information	Free	0.0971 ± 0.0103	0.1638	0.2939
Equal-information	Calibrated	$\mathbf{0.0621} \pm 0.0027$	0.1640	0.2915

7 Discussion and conclusion

Our results support a recovery-aware view of quantum diffusion: rather than prescribing a sequence of noisy quantum channels and asking a reverse model to invert them afterward, the forward process can be designed as a curriculum of inverse problems with explicit recovery references. This distinction is consequential because a diffusion schedule does more than determine how noise is applied—it allocates the recovery burden that a finite-capacity reverse model is asked to solve across timesteps. The classical–quantum information decrement Δ_t provides an operational coordinate for this allocation. Equalizing Δ_t controls how much ensemble information is removed at each forward step, while the same quantity sets an attainable accuracy scale for a common physical recovery channel. At the same time, recoverability alone does not determine a generative model: locally accurate reverse steps can still induce different distributions over quantum states. Quantum diffusion therefore requires two complementary forms of calibration—information-based control of the physical inverse problems and distribution-level matching of their collective generative outcome. Our experiments provide empirical support for this theoretical picture and the resulting recovery-aware design. They further show that the generative benefit of recovery calibration depends on the forward schedule, underscoring the need to consider forward discretization and reverse recovery training jointly.

More broadly, these results suggest that the forward process, reverse learner, and available modeling resources should be designed jointly through the inverse problems they induce. A noise schedule is also a resource-allocation mechanism: it determines where recovery capacity, circuit expressivity, measurements, and computation are required. An operational information coordinate makes this allocation recovery-aware, while a distribution objective controls the global generative law. Although established here for depolarizing dynamics, the same perspective may extend to other quantum channels: corruption should be designed according to what information must be recovered, what recovery is physically attainable, and how limited generative resources can be used effectively.

AI USE STATEMENT

Large language models were used as general-purpose research assistance for language polishing, identifying potentially relevant related work, and generating portions of the experimental code. References surfaced with AI assistance were checked by the authors against the original sources before citation. AI-generated code used in the reported experiments was manually reviewed, tested, and checked against the declared experimental protocol before execution. Reported numerical results were obtained by executing the reviewed code rather than by asking a language model to generate result values. The authors take full responsibility for the paper's claims, citations, code, and experimental results.

REPRODUCIBILITY STATEMENT.

Complete proofs and experimental protocols are provided in the appendix. The anonymized code package documents data generation, forward schedules, reverse-model implementations, training objectives, configurations and seeds, saved results, and analysis scripts, including the audited external QuDDPM reproduction and transfer evaluation.

REFERENCES

Luca Ambrogioni, Giulio Franzese, Alberto Foresti, Gabriel Raya, Bac Nguyen, Georgios Batzolis, Yuhta Takida, Naoki Murata, Chieh-Hsin Lai, and Yuki Mitsufuji. From atoms to entropy: Optimal noise allocation for diffusion training in the convex regime, 2026. URL https://arxiv.org/abs/2607.20540.

Koenraad M. R. Audenaert. A sharp continuity estimate for the von neumann entropy. *Journal of Physics A: Mathematical and Theoretical*, 40(28):8127–8136, 2007. doi: 10.1088/1751-8113/40/28/S18. URL https://doi.org/10.1088/1751-8113/40/28/S18.

Yochai Blau and Tomer Michaeli. The perception-distortion tradeoff. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6228–6237, 2018. doi: 10.1109/CVPR.2018.00652. URL https://openaccess.thecvf.com/content_cvpr_2018/html/Blau_The_Perception-Distortion_Tradeoff_CVPR_2018_paper.html.

Maël Bompais, Mădălin Guţă, and Juan P. Garrahan. Generating quantum ensembles via reverse-time quantum diffusions, 2026. URL https://arxiv.org/abs/2606.03848.

Francesco Buscemi, Siddhartha Das, and Mark M. Wilde. Approximate reversibility in the context of entropy gain, information gain, and complete positivity. *Physical Review A*, 93(6):062314, 2016. doi: 10.1103/PhysRevA.93.062314. URL https://doi.org/10.1103/PhysRevA.93.062314.

Haoyuan Cai, Qi Ye, and Dong-Ling Deng. Sample complexity of learning parametric quantum circuits. *Quantum Science and Technology*, 7(2):025014, 2022. doi: 10.1088/2058-9565/ac4f30. URL https://doi.org/10.1088/2058-9565/ac4f30.

Matthias C. Caro, Hsin-Yuan Huang, M. Cerezo, Kunal Sharma, Andrew Sornborger, Lukasz Cincio, and Patrick J. Coles. Generalization in quantum machine learning from few training data. *Nature Communications*, 13:4919, 2022. doi: 10.1038/s41467-022-32550-3. URL https://doi.org/10.1038/s41467-022-32550-3.

Chuangtao Chen, Qinglin Zhao, MengChu Zhou, Zhimin He, Zhili Sun, and Haozhen Situ. Quantum generative diffusion model: A fully quantum-mechanical model for generating quantum state ensemble. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–15, 2026. doi: 10.1109/TPAMI.2026.3718311. URL https://doi.org/10.1109/TPAMI.2026.3718311. Early Access.

Man-Duen Choi. Completely positive linear maps on complex matrices. *Linear Algebra and its Applications*, 10(3):285–290, 1975. doi: 10.1016/0024-3795(75)90075-0. URL https://doi.org/10.1016/0024-3795(75)90075-0.

Omar Fawzi and Renato Renner. Quantum conditional mutual information and approximate markov chains. *Communications in Mathematical Physics*, 340:575–611, 2015. doi: 10.1007/s00220-015-2466-x. URL https://doi.org/10.1007/s00220-015-2466-x.

Alberto Foresti, Mustapha Bounoua, Giulio Franzese, Luca Ambrogioni, and Pietro Michiardi. Improved sampling schedules for discrete diffusion models, 2026. URL https://arxiv.org/abs/2602.06849.

Christopher A. Fuchs and Jeroen van de Graaf. Cryptographic distinguishability measures for quantum-mechanical states. *IEEE Transactions on Information Theory*, 45(4):1216–1227, 1999. doi: 10.1109/18.761271. URL https://doi.org/10.1109/18.761271.

Jonathan Ho, Ajay N. Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html.

Alexander S. Holevo. Bounds for the quantity of information transmitted by a quantum communication channel. *Problems of Information Transmission*, 9(3):177–183, 1973. URL https://www.mathnet.ru/eng/ppi903.

Haitao Huang, Chuangtao Chen, and Qinglin Zhao. Continuous-variable quantum diffusion model for state generation and restoration, 2025. URL https://arxiv.org/abs/2506.19270.

Anna Jenčová. Comparison of quantum channels and statistical experiments. In *2016 IEEE International Symposium on Information Theory (ISIT)*, pp. 2249–2253, 2016. doi: 10.1109/ISIT.2016.7541699. URL https://doi.org/10.1109/ISIT.2016.7541699. Extended version: arXiv:1512.07016.

Marius Junge, Renato Renner, David Sutter, Mark M. Wilde, and Andreas Winter. Universal recovery maps and approximate sufficiency of quantum relative entropy. *Annales Henri Poincaré*, 19: 2955–2978, 2018. doi: 10.1007/s00023-018-0716-0. URL https://doi.org/10.1007/s00023-018-0716-0.

Dennis Kretschmann, Dirk Schlingemann, and Reinhard F. Werner. A continuity theorem for stinespring's dilation. *Journal of Functional Analysis*, 255(8):1889–1904, 2008. doi: 10.1016/j.jfa.2008.07.023. URL https://doi.org/10.1016/j.jfa.2008.07.023.

Gino Kwun, Bingzhi Zhang, and Quntao Zhuang. Mixed-state quantum denoising diffusion probabilistic model. *Physical Review A*, 111(3):032610, 2025. doi: 10.1103/PhysRevA.111.032610. URL https://doi.org/10.1103/PhysRevA.111.032610.

Göran Lindblad. Completely positive maps and entropy inequalities. *Communications in Mathematical Physics*, 40(2):147–151, 1975. doi: 10.1007/BF01609396. URL https://doi.org/10.1007/BF01609396.

Xinyu Liu, Jingze Zhuang, Wanda Hou, and Yi-Zhuang You. Measurement-based quantum diffusion models, 2025. URL https://arxiv.org/abs/2508.08799.

Ognyan Oreshkov and John Calsamiglia. Distinguishability measures between ensembles of quantum states. *Physical Review A*, 79(3):032336, 2009. doi: 10.1103/PhysRevA.79.032336. URL https://doi.org/10.1103/PhysRevA.79.032336.

Marco Parigi, Stefano Martina, and Filippo Caruso. Quantum-noise-driven generative diffusion models. *Advanced Quantum Technologies*, 8(12):2300401, 2025. doi: https://doi.org/10.1002/qute.202300401. URL https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/qute.202300401.

Dénes Petz. Sufficient subalgebras and the relative entropy of states of a von neumann algebra. *Communications in Mathematical Physics*, 105(1):123–131, 1986. doi: 10.1007/BF01212345. URL https://doi.org/10.1007/BF01212345.

Gabriel Raya, Bac Nguyen, Georgios Batzolis, Yuhta Takida, Dejan Stancevic, Naoki Murata, Chieh-Hsin Lai, Yuki Mitsufuji, and Luca Ambrogioni. Noise scheduling as information-guided allocation in diffusion training, 2026. URL https://arxiv.org/abs/2602.18647.

Joseph M. Renes, Robin Blume-Kohout, A. J. Scott, and Carlton M. Caves. Symmetric informationally complete quantum measurements. *Journal of Mathematical Physics*, 45(6):2171–2180, 2004. doi: 10.1063/1.1737053. URL https://doi.org/10.1063/1.1737053.

M. E. Shirokov. Tight uniform continuity bounds for the quantum conditional mutual information, for the Holevo quantity, and for capacities of quantum channels. *Journal of Mathematical Physics*, 58(10):102202, 2017. doi: 10.1063/1.4987135. URL https://doi.org/10.1063/1.4987135.

Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 2256–2265. PMLR, 2015. URL https://proceedings.mlr.press/v37/sohl-dickstein15.html.

Dejan Stancevic, Florian Handke, and Luca Ambrogioni. Entropic time schedulers for generative diffusion models. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-1474. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/3ed10616ecdd776be283c0a45cf9332d-Abstract-Conference.html.

John Watrous. Semidefinite programs for completely bounded norms. *Theory of Computing*, 5(11): 217–238, 2009. doi: 10.4086/toc.2009.v005a011. URL https://theoryofcomputing.org/articles/v005a011/.

Jian Xu, Wei Chen, Shigui Li, Chao Li, Delu Zeng, John Paisley, and Qibin Zhao. Local-time riemannian score matching on the quantum pure-state manifold, 2026. URL https://arxiv.org/abs/2605.03573.

Bingzhi Zhang, Peng Xu, Xiaohui Chen, and Quntao Zhuang. Generative quantum machine learning via denoising diffusion probabilistic models. *Physical Review Letters*, 132(10):100602, 2024. doi: 10.1103/PhysRevLett.132.100602. URL https://doi.org/10.1103/PhysRevLett.132.100602.

Huangjun Zhu, Yong Siah Teo, and Berthold-Georg Englert. Two-qubit symmetric informationally complete positive-operator-valued measures. *Physical Review A*, 82(4):042308, 2010. doi: 10.1103/PhysRevA.82.042308. URL https://doi.org/10.1103/PhysRevA.82.042308.

A PROOFS OF MAIN RESULTS

This appendix section collects the proofs of the results stated in the main text, in the same order as the main narrative: the equal-information schedule, common-channel recoverability, recoverability–matching separation, finite-sample calibration, and the local-to-global endpoint certificate. Additional refinements, implementation certificates, and experimental details are organized separately below.

A.1 CQ IDENTITIES AND THE FORWARD EQUAL-INFORMATION SCHEDULE

The main text uses two properties of the forward coordinate: the cq identity and invertibility of the depolarizing curve. This section proves both before any recovery argument is invoked.

A.1.1 THE CQ MUTUAL INFORMATION EQUALS THE HOLEVO QUANTITY

For $\omega^{XQ} = \sum_x p_x |x\rangle\langle x| \otimes \rho_x$, block diagonalization gives

$$S(\omega^{XQ}) = H(p) + \sum_x p_x S(\rho_x), \tag{20}$$

where $H(p) = -\sum_x p_x \log p_x$ is the Shannon entropy. Moreover, $S(\omega^X) = H(p)$ and $S(\omega^Q) = S(\bar{\rho})$. Substitution into the definition of quantum mutual information in the introduction proves Equation (4).

A.1.2 PROOF OF THE FORWARD-SCHEDULE THEOREM

Write $\rho_x(\lambda) := \mathcal{D}_\lambda(\rho_x)$ and $\bar{\rho}(\lambda) := \sum_x p_x \rho_x(\lambda)$. Continuity follows from finite-dimensional entropy continuity. For $0 < \lambda < 1$, all $\rho_x(\lambda)$ and $\bar{\rho}(\lambda)$ are positive definite, so analyticity follows from analyticity of the matrix logarithm on the positive cone.

For $\lambda = \alpha\lambda_1 + (1-\alpha)\lambda_2$, both arguments in the definition of $\chi(\lambda)$ are the corresponding convex combinations. Joint convexity of quantum relative entropy therefore yields

$$\chi(\lambda) \le \alpha\chi(\lambda_1) + (1-\alpha)\chi(\lambda_2). \tag{21}$$

For $a \in [0,1]$,

$$(\rho_x(a\lambda), \bar{\rho}(a\lambda)) = a(\rho_x(\lambda), \bar{\rho}(\lambda)) + (1-a)(\mathbb{I}/d, \mathbb{I}/d). \tag{22}$$

Joint convexity therefore gives the contraction relation

$$\chi(a\lambda) \le a\chi(\lambda). \tag{23}$$

For $\lambda > 0$, the channel $\mathcal{D}_\lambda$ is injective on trace-one operators. A nontrivial ensemble therefore remains nontrivial, so $\chi(\lambda) > 0$. If $0 < \mu < \lambda$, apply Equation (23) with $a = \mu/\lambda$:

$$\chi(\mu) \le \frac{\mu}{\lambda}\chi(\lambda) < \chi(\lambda). \tag{24}$$

Together with $\chi(0) = 0$, this proves strict increase. The inverse grid consequently exists and is unique.

For any monotone grid with fixed endpoints,

$$\sum_{t=1}^{T} \Delta_t = \Delta_{\mathrm{tot}}, \tag{25}$$

so $\max_t \Delta_t \ge \Delta_{\mathrm{tot}}/T$, with equality if and only if all decrements are equal. Strict monotonicity of χ then gives uniqueness of the grid. Jensen's inequality gives

$$\sum_{t=1}^{T} \varphi(\Delta_t) \ge T\varphi(\Delta_{\mathrm{tot}}/T), \tag{26}$$

with equality only at the equal vector when φ is strictly convex. □

A.2 Learner-facing and two-sided common-channel recovery

We use the following notation for the recovery proofs. Write $\chi_t = \chi(\lambda_t)$, $\bar{\rho}_t = \sum_x p_x \rho_{x,t}$, and, whenever $\lambda_{t-1} > 0$, $\mathcal{N}_t = \mathcal{D}_{\lambda_t/\lambda_{t-1}}$. Also write

$$\mathcal{L}_t^{\log}(\mathcal{R}) := -2 \sum_x p_x \log F_{\mathrm{r}}(\rho_{x,t-1}, \mathcal{R}(\rho_{x,t})), \tag{27}$$

$$\mathcal{L}_t^{\mathrm{cq}}(\mathcal{R}) := -2 \log F_{\mathrm{r}}\Big(\omega_{t-1}^{XQ}, (\mathrm{id}_X \otimes \mathcal{R})(\omega_t^{XQ})\Big), \tag{28}$$

where $\omega_t^{XQ} = \sum_x p_x |x\rangle\langle x| \otimes \rho_{x,t}$ and

$$F_{\mathrm{r}}\Big(\omega_{t-1}^{XQ}, (\mathrm{id}_X \otimes \mathcal{R})(\omega_t^{XQ})\Big) = \sum_x p_x F_{\mathrm{r}}(\rho_{x,t-1}, \mathcal{R}(\rho_{x,t})). \tag{29}$$

Let $m = |\mathcal{X}|$. With $h_2(r) = -r \log r - (1-r)\log(1-r)$ and $0 \log 0 := 0$, define

$$g(r) = (1+r)\log(1+r) - r \log r, \tag{30}$$

$$\mathfrak{f}_d(r) = h_2(\min\{r, 1-1/d\}) + \min\{r, 1-1/d\} \log(d-1), \tag{31}$$

and the same-prior Holevo continuity modulus

$$\Omega_{d,m}(r) = \min\{2\mathfrak{f}_d(r),\, r \log \min\{d,m\} + g(r),\, \log \min\{d,m\}\}, \tag{32}$$

with generalized inverse $\Omega_{d,m}^{\leftarrow}(a) = \inf\{r \in [0,1] : \Omega_{d,m}(r) \geq a\}$. For two ensembles on the same index set, Shirokov's distance is

$$D_{\mathrm{ens}}(\{p_x, \rho_x\}, \{q_x, \sigma_x\}) = \frac{1}{2} \sum_x \|p_x \rho_x - q_x \sigma_x\|_1. \tag{33}$$

This appendix separates the ingredients of the local recovery theory. Ensemble-level Holevo-loss recoverability is established in prior work (Buscemi et al., 2016), and the universal recovery theorem supplies a map depending only on the reference state and forward channel (Junge et al., 2018). We first show that this same map controls the stricter average-log risk used by our learner. We then combine data processing with established Holevo-continuity bounds and, separately, derive the pairwise-geometric converse used in the strengthened lower bracket. A decision-value witness from the statistical-comparison viewpoint also yields an exact symmetric orthogonal benchmark.

A.2.1 Proof of the learner-facing common-channel budget

The universal recovery theorem states that, for states ρ, σ and a quantum channel $\mathcal{N}$, there is a recovery channel $\mathcal{R}_{\sigma,\mathcal{N}}^{\mathrm{U}}$ depending only on $(\sigma, \mathcal{N})$ such that

$$D(\rho\|\sigma) - D(\mathcal{N}(\rho)\|\mathcal{N}(\sigma)) \geq -2 \log F_{\mathrm{r}}\big(\rho, (\mathcal{R}_{\sigma,\mathcal{N}}^{\mathrm{U}} \circ \mathcal{N})(\rho)\big), \tag{34}$$

and $(\mathcal{R}_{\sigma,\mathcal{N}}^{\mathrm{U}} \circ \mathcal{N})(\sigma) = \sigma$ (Junge et al., 2018).

Because $p_x > 0$ and $\bar{\rho}_{t-1} = \sum_x p_x \rho_{x,t-1}$, one has $\operatorname{supp} \rho_{x,t-1} \subseteq \operatorname{supp} \bar{\rho}_{t-1}$ for every x. Apply Equation (34) to every pair $(\rho_{x,t-1}, \bar{\rho}_{t-1})$ with the same channel $\mathcal{N}_t$. Universality means that the same recovery channel, depending only on $(\bar{\rho}_{t-1}, \mathcal{N}_t)$, works for all x. Moreover,

$$\Delta_t = \sum_x p_x \Big[D(\rho_{x,t-1}\|\bar{\rho}_{t-1}) - D(\rho_{x,t}\|\bar{\rho}_t) \Big]. \tag{35}$$

Let

$$F_x := F_{\mathrm{r}}\big(\rho_{x,t-1}, \mathcal{R}_t^{\mathrm{U}}(\rho_{x,t})\big). \tag{36}$$

Summing with weights p_x gives

$$\Delta_t \geq -2 \sum_x p_x \log F_x = \mathcal{L}_t^{\log}(\mathcal{R}_t^{\mathrm{U}}). \tag{37}$$

Since $-\log$ is convex,

$$-2 \log \sum_x p_x F_x \leq -2 \sum_x p_x \log F_x. \tag{38}$$

Combining this inequality with Equation (6) and exponentiating gives the appendix-local consequence

$$\sum_x p_x F_x \geq e^{-\Delta_t/2}. \tag{39}$$

By Equation (29), the same weighted sum is the root fidelity of the corresponding cq block states. □

A.2.2 PROOF OF THE TWO-SIDED INTRINSIC RECOVERABILITY THEOREM

Fix a timestep and an arbitrary label-independent CPTP recovery channel $\mathcal{R}$. Write

$$\sigma_x := \mathcal{R}(\rho_{x,t}), \qquad \mathcal{E}_{\mathcal{R}} := \{p_x, \sigma_x\}, \qquad \varepsilon(\mathcal{R}) := \sum_x p_x d_{\mathrm{tr}}(\rho_{x,t-1}, \sigma_x). \tag{40}$$

Because $\mathcal{E}_t$ is mapped to $\mathcal{E}_{\mathcal{R}}$ by $\mathcal{R}$, data processing for the Holevo quantity gives $\chi(\mathcal{E}_{\mathcal{R}}) \le \chi_t$. Hence

$$\Delta_t = \chi_{t-1} - \chi_t \le \chi_{t-1} - \chi(\mathcal{E}_{\mathcal{R}}). \tag{41}$$

The two ensembles on the right have the same probabilities. Their Shirokov ensemble distance D_{ens} from Equation (33) is exactly $\varepsilon(\mathcal{R})$. Proposition 5, Eqs. (60)–(61), of Shirokov (2017) gives separate bounds with coefficients $\log d$ and $\log m$; the same proposition states that when the priors coincide, the generic $2g(D_{\mathrm{ens}})$ remainder can be replaced by $g(D_{\mathrm{ens}})$. Hence

$$|\chi_{t-1} - \chi(\mathcal{E}_{\mathcal{R}})| \le \varepsilon(\mathcal{R}) \log \min\{d, m\} + g(\varepsilon(\mathcal{R})). \tag{42}$$

An independent entropy-continuity route yields the other branch in Equation (32). Let $\bar{\sigma} = \sum_x p_x \sigma_x$ and $\varepsilon_x = d_{\mathrm{tr}}(\rho_{x,t-1}, \sigma_x)$. Convexity of trace distance gives $d_{\mathrm{tr}}(\bar{\rho}_{t-1}, \bar{\sigma}) \le \varepsilon(\mathcal{R})$. Audenaert continuity and concavity of the monotone modulus $\mathfrak{f}_d$ then give

$$|\chi_{t-1} - \chi(\mathcal{E}_{\mathcal{R}})| \le \mathfrak{f}_d(\varepsilon(\mathcal{R})) + \sum_x p_x \mathfrak{f}_d(\varepsilon_x) \tag{43}$$

$$\le 2\mathfrak{f}_d(\varepsilon(\mathcal{R})). \tag{44}$$

Finally, every Holevo quantity lies in $[0, \log \min\{d, m\}]$. Combining these three valid bounds with Equation (41) proves, for every $\mathcal{R}$,

$$\Delta_t \le \Omega_{d,m}(\varepsilon(\mathcal{R})). \tag{45}$$

In finite dimension the set of CPTP maps is compact in the Choi representation, and the objective defining $\varepsilon_t^\star$ is continuous, so the infimum is attained by some $\mathcal{R}_t^\star$. Applying Equation (45) to $\mathcal{R}_t^\star$ gives the main-text converse Equation (7).

For the upper bound, use the universal channel $\mathcal{R}_t^{\mathrm{U}}$ from Theorem 2. With $F_x := F_{\mathrm{r}}(\rho_{x,t-1}, \mathcal{R}_t^{\mathrm{U}}(\rho_{x,t}))$, Equation (39), the upper Fuchs–van de Graaf inequality Fuchs & van de Graaf (1999), and concavity of $f \mapsto \sqrt{1-f^2}$ give

$$\sum_x p_x d_{\mathrm{tr}}(\rho_{x,t-1}, \mathcal{R}_t^{\mathrm{U}}(\rho_{x,t})) \le \sum_x p_x \sqrt{1 - F_x^2} \tag{46}$$

$$\le \sqrt{1 - \left(\sum_x p_x F_x\right)^2} \tag{47}$$

$$\le \sqrt{1 - e^{-\Delta_t}}. \tag{48}$$

Taking the infimum over common channels proves Theorem 3.

The final equivalence in Theorem 3 now follows directly. If $\Delta_t \to 0$, the upper bound gives $\varepsilon_t^\star \to 0$. Conversely, Equation (7) and continuity of $\Omega_{d,m}$ at the origin, with $\Omega_{d,m}(0) = 0$, imply that $\varepsilon_t^\star \to 0$ forces $\Delta_t \to 0$. □

A.3 PROOFS FOR RECOVERABILITY–MATCHING SEPARATION

We first state the auxiliary symmetric observation used to motivate the quantum-specific separation.

Proposition 7 (Local scores can miss distributional differences)**.** *Let the target distribution be uniform over $\{\rho_1, \ldots, \rho_m\}$. Suppose a local score S satisfies $S(\rho_i, \rho_i) = s_{\mathrm{self}}$ and $S(\rho_i, \rho_j) = s_{\mathrm{off}}$ for all $i \ne j$. Let X and Z be independent and uniform on $\{1, \ldots, m\}$. For the correctly covering generator $\widehat{\rho}^{\mathrm{good}} = \rho_Z$ and the collapsed generator $\widehat{\rho}^{\mathrm{col}} = \rho_{j_0}$, the local score has the same distribution: it equals s_{self} with probability $1/m$ and s_{off} with probability $1 - 1/m$. Hence the complete local-score distribution cannot distinguish the two generators even though their output distributions differ.*

A.3.1 ASSIGNMENT BLINDNESS OF PRODUCT-COUPLED SCORES

For the good model in Prop 7, conditioning on $Z = j$ gives the same two-point score law stated in the proposition. Therefore

$$\mathrm{Law}(S(\rho_X, \rho_Z)) = \frac{1}{m}\sum_{j=1}^{m} \nu_S = \nu_S = \mathrm{Law}(S(\rho_X, \rho_{j_0}))\,. \tag{49}$$

The good output law is μ, so its Wasserstein error is zero. The collapsed output law is the Dirac measure $\delta_{\rho_{j_0}}$; its coupling with μ is unique, giving

$$W_{1,\mathrm{tr}}(\mu, \delta_{\rho_{j_0}}) = \frac{1}{m}\sum_{x} d_{\mathrm{tr}}(\rho_x, \rho_{j_0}). \tag{50}$$

This proves Prop 7. □

A.3.2 FIXED TWO-QUBIT NONCOMMUTING SEPARATION

Four-dimensional SIC-POVMs exist and admit an explicit two-qubit realization (Renes et al., 2004; Zhu et al., 2010). Let $P_j = |\psi_j\rangle\langle\psi_j|$, $j = 1, \ldots, 16$, denote its normalized rank-one projectors. The SIC identities are

$$\frac{1}{16}\sum_{j=1}^{16} P_j = \frac{\mathbb{I}}{4}, \qquad \mathrm{Tr}(P_i P_j) = \frac{1}{5} \quad (i \neq j). \tag{51}$$

Identifying the four-dimensional Hilbert space with two qubits, define

$$\rho_j := \frac{2}{3}P_j + \frac{1}{12}\mathbb{I}, \qquad \mu_{\mathrm{SIC}} := \frac{1}{16}\sum_{j=1}^{16}\delta_{\rho_j}, \qquad \tau := \frac{\mathbb{I}}{4}. \tag{52}$$

The average state is τ, and complete depolarization maps every ρ_j to τ. Moreover, for $i \neq j$,

$$[\rho_i, \rho_j] = \frac{4}{9}[P_i, P_j] \neq 0. \tag{53}$$

Indeed, two distinct rank-one projectors commute only if their ranges are orthogonal, whereas Equation (51) gives nonzero, nonunit overlap. Thus the target ensemble is pairwise noncommuting.

Let X, Z be independent and uniform on $\{1, \ldots, 16\}$. Define the good and collapsed stochastic reverse models by replacer branches,

$$\mathcal{R}_z^{\mathrm{good}}(A) = \mathrm{Tr}(A)\rho_z, \qquad \mathcal{R}_z^{\mathrm{col}}(A) = \mathrm{Tr}(A)\rho_1. \tag{54}$$

The good output law is exactly μ_{SIC}, while the collapsed output law is δ_{ρ_1}.

We first compute the off-diagonal root fidelity. More generally, for

$$\rho_P = v\mathbb{I} + \lambda P, \qquad \rho_Q = v\mathbb{I} + \lambda Q, \qquad u := v + \lambda, \qquad s := \mathrm{Tr}(PQ), \tag{55}$$

with rank-one projectors P, Q in dimension d, the orthogonal complement of $\mathrm{span}\{P, Q\}$ contributes $(d-2)v$ to the root fidelity. On the two-dimensional span, the positive matrix $\sqrt{\rho_P}\rho_Q\sqrt{\rho_P}$ has trace $2uv + \lambda^2 s$ and determinant $(uv)^2$. Hence

$$F_{\mathrm{r}}(\rho_P, \rho_Q) = (d-2)v + \sqrt{4uv + \lambda^2 s}. \tag{56}$$

For Equation (52), $d = 4$, $\lambda = 2/3$, $v = 1/12$, $u = 3/4$, and $s = 1/5$ off the diagonal. Therefore

$$f_{\mathrm{SIC}} := F_{\mathrm{r}}(\rho_i, \rho_j) = \frac{1}{6} + \sqrt{\frac{61}{180}} = \frac{5 + \sqrt{305}}{30} \qquad (i \neq j). \tag{57}$$

The fidelity equals one on the diagonal. In the good model $X = Z$ with probability $1/16$; in the collapsed model $X = 1$ with the same probability. Thus both complete scalar fidelity laws are

$$F = \begin{cases} 1, & \text{with probability } 1/16, \\ f_{\mathrm{SIC}}, & \text{with probability } 15/16, \end{cases} \tag{58}$$

proving Equation (8). Their common unclipped local log-fidelity risk is

$$L_{\mathrm{SIC}}^{\log} = -\frac{15}{8}\log f_{\mathrm{SIC}}. \tag{59}$$

Each ρ_j has spectrum $\{3/4, 1/12, 1/12, 1/12\}$. Since the average state is τ, the complete-depolarization cq-information decrement is

$$\Delta_{\mathrm{SIC}} = \log 4 - S(\rho_j) = \frac{1}{2}\log 3. \tag{60}$$

The budget inequality is strict. To verify it without a floating-point comparison, note first that $\sqrt{305} > 87/5$, hence

$$f_{\mathrm{SIC}} > \frac{56}{75}. \tag{61}$$

Moreover,

$$\left(\frac{56}{75}\right)^3 > \frac{27}{65} \quad \text{and} \quad \left(\frac{27}{65}\right)^5 > \frac{1}{81}; \tag{62}$$

the two cross-multiplied numerator differences are respectively 24415 and 1970842. Consequently $f_{\mathrm{SIC}}^{15} > 1/81$, which is equivalent to

$$-\frac{15}{8}\log f_{\mathrm{SIC}} < \frac{1}{2}\log 3 = \Delta_{\mathrm{SIC}}. \tag{63}$$

This proves the exact budget-feasibility claim in Equation (9).

Finally, for distinct SIC projectors,

$$d_{\mathrm{tr}}(P_i, P_j) = \sqrt{1 - \mathrm{Tr}(P_i P_j)} = \frac{2}{\sqrt{5}}, \tag{64}$$

and homogeneity under the common depolarizing mixture gives

$$d_{\mathrm{tr}}(\rho_i, \rho_j) = \frac{4}{3\sqrt{5}} \qquad (i \neq j). \tag{65}$$

The good model has zero endpoint error. For the collapsed model the second Wasserstein marginal is a Dirac measure, so the coupling is unique and

$$W_{1,\mathrm{tr}}(\mu_{\mathrm{SIC}}, \delta_{\rho_1}) = \frac{1}{16}\sum_{j=1}^{16} d_{\mathrm{tr}}(\rho_j, \rho_1) \tag{66}$$

$$= \frac{15}{16}\frac{4}{3\sqrt{5}} = \frac{\sqrt{5}}{4} > \frac{1}{2}. \tag{67}$$

This proves Theorem 4. □

The separation can be amplified arbitrarily in higher dimension.

Corollary 8 (Arbitrarily severe budget-feasible separation)**.** *For every $c \in (0, 1)$ there exists a finite-dimensional complete-depolarization problem and two stochastic CPTP reverse models with identical local-fidelity laws and identical expected local log-fidelity losses, both within the exact information-calibrated recovery budget, while one model has zero trace-Wasserstein error and the other has error larger than c. The dimension may be chosen as a power of two.*

A.3.3 Amplification to an arbitrary constant below one

We prove Corollary 8 by the symmetric commuting family used to amplify the separation. Fix $d \geq 2$ and $\lambda \in (0, 1)$, let $\tau = \mathbb{I}/d$, and define

$$\rho_x^\lambda = \lambda|x\rangle\langle x| + (1-\lambda)\tau, \qquad \mu_\lambda^{(d)} = \frac{1}{d}\sum_{x=1}^{d} \delta_{\rho_x^\lambda}. \tag{68}$$

The complete-depolarization channel maps every ρ_x^λ to τ, so the information decrement of this one-step forward problem is

$$\Delta_d(\lambda) = \chi(\mathcal{E}_\lambda^{(d)}). \tag{69}$$

Let Z be uniform on $\{1, \dots, d\}$ and independent of the target label X. Define the good branch family and the collapsed branch family by the CPTP replacer channels

$$\mathcal{R}_z^{\mathrm{good}}(A) = \mathrm{Tr}(A)\rho_z^\lambda, \tag{70}$$

$$\mathcal{R}_z^{\mathrm{col}}(A) = \mathrm{Tr}(A)\rho_1^\lambda \qquad \text{for every } z. \tag{71}$$

The good output law is exactly $\mu_\lambda^{(d)}$, whereas the collapsed output law is $\delta_{\rho_1^\lambda}$.

Set

$$u_\lambda := \frac{1 + (d-1)\lambda}{d}, \qquad v_\lambda := \frac{1-\lambda}{d}. \tag{72}$$

The states commute in the computational basis. For $x \neq z$, their root fidelity is

$$f_d(\lambda) = 2\sqrt{u_\lambda v_\lambda} + (d-2)v_\lambda, \tag{73}$$

while the fidelity is one when $x = z$. In the good model, $X = Z$ with probability $1/d$; in the collapsed model, $X = 1$ with probability $1/d$. Hence both scalar fidelity random variables have the same law,

$$F = \begin{cases} 1, & \text{with probability } 1/d, \\ f_d(\lambda), & \text{with probability } (d-1)/d, \end{cases} \tag{74}$$

and their common unclipped branch-averaged log risk is

$$\mathcal{L}_d^{\log,\mathrm{br}}(\lambda) = -2\frac{d-1}{d}\log f_d(\lambda). \tag{75}$$

The good model reproduces $\mu_\lambda^{(d)}$, so its Wasserstein error is zero. For the collapsed model the second marginal is a Dirac measure, hence the coupling is unique and

$$W_{1,\mathrm{tr}}\big(\mu_\lambda^{(d)}, \delta_{\rho_1^\lambda}\big) = \frac{1}{d}\sum_x d_{\mathrm{tr}}(\rho_x^\lambda, \rho_1^\lambda) \tag{76}$$

$$= \lambda\left(1 - \frac{1}{d}\right), \tag{77}$$

because $d_{\mathrm{tr}}(\rho_x^\lambda, \rho_y^\lambda) = \lambda$ for $x \neq y$.

The average state of $\mathcal{E}_\lambda^{(d)}$ is τ, so

$$\Delta_d(\lambda) = \log d - S(\rho_x^\lambda) = \log d + u_\lambda \log u_\lambda + (d-1)v_\lambda \log v_\lambda. \tag{78}$$

For every fixed $\lambda \in (0,1)$, as $d \to \infty$,

$$\Delta_d(\lambda) = \lambda \log d - h_2(\lambda) + o(1), \tag{79}$$

whereas Equation (73) gives

$$f_d(\lambda) \longrightarrow 1 - \lambda, \qquad \mathcal{L}_d^{\log,\mathrm{br}}(\lambda) \longrightarrow -2\log(1-\lambda) < \infty. \tag{80}$$

Thus for every fixed $\lambda \in (0,1)$ there is a finite $d_0(\lambda)$ such that

$$\mathcal{L}_d^{\log,\mathrm{br}}(\lambda) \leq \Delta_d(\lambda) \qquad \text{for all } d \geq d_0(\lambda). \tag{81}$$

Given $c < 1$, choose $\lambda \in (c, 1)$. Equation 77 exceeds c for all sufficiently large d, while the budget inequality also holds for sufficiently large d. We may choose such a d along the subsequence $d = 2^n$, proving the corollary for qubit Hilbert spaces. □

Small-noise complement. The amplification family is not purely a large-budget construction. For fixed $d > 2$ and small λ, direct Taylor expansion gives

$$\chi(\mathcal{E}_\lambda^{(d)}) = \frac{d-1}{2}\lambda^2 - \frac{(d-1)(d-2)}{6}\lambda^3 + O(\lambda^4), \tag{82}$$

$$\mathcal{L}_d^{\log,\mathrm{br}}(\lambda) = \frac{d-1}{2}\lambda^2 - \frac{(d-1)(d-2)}{4}\lambda^3 + O(\lambda^4). \tag{83}$$

Hence

$$\chi(\mathcal{E}_\lambda^{(d)}) - \mathcal{L}_d^{\log,\mathrm{br}}(\lambda) = \frac{(d-1)(d-2)}{12}\lambda^3 + O(\lambda^4) > 0 \tag{84}$$

for all sufficiently small positive λ.

A.4 Uniform finite-sample calibration

We prove Theorem 5 for the actual outcome-weighted risk, rather than for an averaged output state. Fix the forward grid, the ancillary preparations, and all non-trainable conditioning before drawing the risk sample; they may be fixed conditionally on an independent pilot. At each timestep, the instrument is obtained from

$$U_{\theta,t} = W_{P,t}e^{-i\theta_P H_{P,t}/2}\cdots W_{1,t}e^{-i\theta_1 H_{1,t}/2}W_{0,t}, \qquad \theta \in [-\pi,\pi]^P, \tag{85}$$

followed by a fixed measurement and, if needed, a partial trace. Here $\|H_{j,t}\|_\infty \le 1$ and the $W_{j,t}$ are fixed unitaries. Each coordinate occurs at most once in a given timestep; setting its generator to zero accommodates unused coordinates and disjoint timestep parameter blocks. Coordinates may be shared across timesteps. Repeated gates controlled by the same coordinate require counting their occurrences in the parameter-continuity bound. A separately trainable classical conditioner is not included. The parameter dependence is consistent with gate-based QML bounds (Caro et al., 2022; Cai et al., 2022).

For an outcome of positive probability, use the sample loss

$$\ell_{\gamma,t}(\theta; x, b) = -2\log\max\left\{\gamma, F_{\mathrm{r}}\left(\rho_{x,t-1}, \frac{\mathcal{M}_{\theta,t,b}(\rho_{x,t})}{\operatorname{Tr}\mathcal{M}_{\theta,t,b}(\rho_{x,t})}\right)\right\}. \tag{86}$$

Set it arbitrarily in $[0, 2\log(1/\gamma)]$ at a zero-probability outcome. For i.i.d. labels $X_1,\dots,X_n \sim p$, the empirical risk is

$$\widehat{L}_{\gamma,t}(\theta) = \frac{1}{n}\sum_{i=1}^{n}\sum_b \operatorname{Tr}\mathcal{M}_{\theta,t,b}(\rho_{X_i,t})\ell_{\gamma,t}(\theta; X_i, b). \tag{87}$$

Thus n counts independent labels, not enumerated outcomes or optimization updates. The same indexed labels can be used across timesteps. For external CPTP branches, $b = z$ and the weight is $\pi_t(z)$, recovering the branch loss used by the finite-latent implementation. Exact summation over the prescribed finite ensemble gives $L_{\gamma,t}$ itself instead of a sampling estimate.

A.4.1 Stability of conditional output laws

The following estimate handles small outcome probabilities without dividing a continuity bound by a minimum probability.

Lemma 9 (Subnormalized outputs and instrument stability)**.** *Let $A_b, B_b \succeq 0$ satisfy $\sum_b \operatorname{Tr} A_b = \sum_b \operatorname{Tr} B_b = 1$. Then*

$$W_{1,\mathrm{tr}}\left(\sum_b \operatorname{Tr} A_b\, \delta_{A_b/\operatorname{Tr} A_b}, \sum_b \operatorname{Tr} B_b\, \delta_{B_b/\operatorname{Tr} B_b}\right) \le \frac{3}{2}\sum_b \|A_b - B_b\|_1 . \tag{88}$$

Zero-mass terms are omitted. Consequently, for any finite-outcome instrument,

$$W_{1,\mathrm{tr}}\big(K(\rho,\cdot), K(\sigma,\cdot)\big) \le 3d_{\mathrm{tr}}(\rho,\sigma). \tag{89}$$

The constant is independent of the number of outcomes.

Proof. Couple the outcome weights maximally. For a matched outcome with both traces positive,

$$\min\{\operatorname{Tr} A_b, \operatorname{Tr} B_b\} \left\| \frac{A_b}{\operatorname{Tr} A_b} - \frac{B_b}{\operatorname{Tr} B_b} \right\|_1 \leq \|A_b - B_b\|_1 + |\operatorname{Tr} A_b - \operatorname{Tr} B_b|. \tag{90}$$

For example, if $\operatorname{Tr} A_b \leq \operatorname{Tr} B_b$, apply the triangle inequality to $A_b - (\operatorname{Tr} A_b / \operatorname{Tr} B_b) B_b$. A zero trace gives zero matched mass. The unmatched mass is $\frac{1}{2}\sum_b |\operatorname{Tr} A_b - \operatorname{Tr} B_b|$ and costs at most that amount, since state-space trace distance has diameter one. Combining matched and unmatched costs gives

$$\frac{1}{2}\sum_b \|A_b - B_b\|_1 + \sum_b |\operatorname{Tr} A_b - \operatorname{Tr} B_b| \leq \frac{3}{2}\sum_b \|A_b - B_b\|_1 .$$

This proves Equation (88). For $A_b = \mathcal{M}_b(\rho)$ and $B_b = \mathcal{M}_b(\sigma)$, the channel retaining the classical measurement record, $\rho \mapsto \sum_b |b\rangle\langle b| \otimes \mathcal{M}_b(\rho)$, is CPTP. Its trace-distance contractivity gives $\sum_b \|A_b - B_b\|_1 \leq \|\rho - \sigma\|_1$, proving Equation (89). □

A.4.2 Parameter continuity of the Born-weighted risk

A telescoping expansion of Equation (85) gives

$$\|U_{\theta,t} - U_{\theta',t}\|_\infty \leq \frac{1}{2}\|\theta - \theta'\|_1 \leq \frac{\sqrt{P}}{2}\|\theta - \theta'\|_2 . \tag{91}$$

Unitary conjugation followed by the fixed measurement channel implies, for every input ρ,

$$\frac{1}{2}\sum_b \|\mathcal{M}_{\theta,t,b}(\rho) - \mathcal{M}_{\theta',t,b}(\rho)\|_1 \leq \|U_{\theta,t} - U_{\theta',t}\|_\infty . \tag{92}$$

Applying Lemma 9 bounds the output-law distance by $3\|U_{\theta,t} - U_{\theta',t}\|_\infty$.

For a fixed target state, Equation (155) and the $2/\gamma$-Lipschitz property of $f \mapsto -2\log\max\{\gamma, f\}$ bound the difference of two output losses by $2\sqrt{2d_{\mathrm{tr}}(\sigma,\sigma')}/\gamma$. Integrating over a Wasserstein-optimal coupling and applying Jensen's inequality to the square root therefore gives, uniformly in x, t,

$$\begin{aligned}\left|\sum_b \operatorname{Tr}\mathcal{M}_{\theta,t,b}(\rho_{x,t})\ell_{\gamma,t}(\theta;x,b) - \sum_b \operatorname{Tr}\mathcal{M}_{\theta',t,b}(\rho_{x,t})\ell_{\gamma,t}(\theta';x,b)\right| \\ \leq \frac{2\sqrt{6}}{\gamma}\sqrt{\|U_{\theta,t} - U_{\theta',t}\|_\infty} \leq \frac{2\sqrt{3}\,P^{1/4}}{\gamma}\|\theta - \theta'\|_2^{1/2} .\end{aligned} \tag{93}$$

This controls the full Born-weighted risk, including parameter-dependent probability masses. Individual normalized branches need not obey this estimate. No assumption of a positive minimum outcome probability has been used.

A.4.3 Covering argument and population calibration

A Euclidean parameter net of radius $r = \gamma^2\varepsilon^2/(12\sqrt{P})$ makes the outcome-averaged losses uniformly ε-close. The cube $[-\pi,\pi]^P$ lies in a ball of radius $\pi\sqrt{P}$, so it admits a net of cardinality

$$N_r \leq \left(1 + \frac{2\pi\sqrt{P}}{r}\right)^P = \left(1 + \frac{24\pi P}{\gamma^2\varepsilon^2}\right)^P . \tag{94}$$

Each outcome-averaged sample loss lies in $[0, 2\log(1/\gamma)]$. Hoeffding's inequality and a union bound over TN_r timestep/net-point pairs yield, with probability at least $1-\delta$,

$$L_{\gamma,t}(\theta_j) \leq \widehat{L}_{\gamma,t}(\theta_j) + 2\log(1/\gamma)\sqrt{\frac{\log(TN_r/\delta)}{2n}}$$

for every net point and timestep. Approximating an arbitrary θ changes its population and empirical risks by at most ε each. Setting $\varepsilon = 2\log(1/\gamma)/\sqrt{n}$ gives

$$\begin{aligned}\varepsilon_{\mathrm{gen}} := \frac{4\log(1/\gamma)}{\sqrt{n}} \\ + 2\log(1/\gamma)\sqrt{\frac{P\log\left(1 + \frac{6\pi P n}{\gamma^2\log^2(1/\gamma)}\right) + \log(T/\delta)}{2n}}.\end{aligned} \tag{95}$$

On this uniform event, for any data-selected $\widehat{\theta}$,

$$L_{\gamma,t}(\widehat{\theta}) - \Delta_t \le \widehat{L}_{\gamma,t}(\widehat{\theta}) - \widehat{\Delta}_t + \varepsilon_{\mathrm{gen}} + |\widehat{\Delta}_t - \Delta_t|. \tag{96}$$

The empirical residual assumption proves Theorem 5. The clock-error term is deterministic on the selected grid; a separate probabilistic clock bound is combined with this event by a union bound. Conditioning on a grid chosen from these same n labels does not justify the fixed-grid argument.

For input-independent external CPTP branches, couple outputs with the same latent value. The factor 3 in the parameter-law estimate becomes 1, yielding the sharper loss modulus

$$\frac{2P^{1/4}}{\gamma}\,\|\theta - \theta'\|_2^{1/2}.$$

The corresponding bound is Equation (95) with 6π replaced by 2π. It applies either to exact latent averaging with i.i.d. labels or to i.i.d. label/latent pairs. This is the non-measurement specialization; Born-sampled outcomes drawn from one parameter's law cannot be treated as a parameter-independent sample for the uniform instrument proof.

A.4.4 Clipped risk implies average fidelity

Lemma 10 (Clipped risk implies average fidelity)**.** *For $X \sim p$ and $\widehat{\rho}_{t-1} \sim K_{\theta,t}(\rho_{X,t}, \cdot)$, if $L_{\gamma,t}(\theta) \le a$ with $a \ge 0$, then*

$$\mathbb{E}F_{\mathrm{r}}(\rho_{X,t-1}, \widehat{\rho}_{t-1}) \ge \left[e^{-a/2} - \frac{\gamma a}{2\log(1/\gamma)}\right]_+. \tag{97}$$

Proof. For the scalar random fidelity $F = F_{\mathrm{r}}(\rho_{X,t-1}, \widehat{\rho}_{t-1})$, concavity of the logarithm gives

$$\mathbb{E}\max\{\gamma, F\} \ge \exp(\mathbb{E}\log\max\{\gamma, F\}) \ge e^{-a/2}.$$

On $\{F < \gamma\}$ the clipped loss equals $2\log(1/\gamma)$, so $\Pr(F < \gamma) \le a/[2\log(1/\gamma)]$. Consequently,

$$\mathbb{E}F = \mathbb{E}\max\{\gamma, F\} - \mathbb{E}(\gamma - F)_+ \ge e^{-a/2} - \frac{\gamma a}{2\log(1/\gamma)}.$$

Nonnegativity of fidelity proves Equation (97). This argument applies to the actual output law and does not require input-independent outcome probabilities. □

A.5 Local-to-global distributional bounds

Fix the parameters θ throughout the composition and endpoint arguments. We use the action of $K_{\theta,t}$ on probability measures defined in Section 5. The target law is $\mu_t = \sum_x p_x \delta_{\rho_{x,t}}$, and the learned reverse chain satisfies $\nu_{t-1} = K_{\theta,t}\nu_t$. We use the transfer function ψ_γ from Equation (17).

Proposition 11 (Local distributional recovery)**.** *For any reverse output kernel and its risk in Equation* (13)*, $L_{\gamma,t}(\theta) \le a$ implies*

$$W_{1,\mathrm{tr}}(\mu_{t-1}, K_{\theta,t}\mu_t) \le \psi_\gamma(a). \tag{98}$$

Under $X \sim p$ and $\widehat{\rho}_{t-1} \sim K_{\theta,t}(\rho_{X,t}, \cdot)$, the directly evaluated average fidelity also gives

$$W_{1,\mathrm{tr}}(\mu_{t-1}, K_{\theta,t}\mu_t) \le \sqrt{1 - \left(\mathbb{E}F_{\mathrm{r}}(\rho_{X,t-1}, \widehat{\rho}_{t-1})\right)^2}. \tag{99}$$

Proof. Sample the target label X and then the actual reverse output at $\rho_{X,t}$. The joint law of $(\rho_{X,t-1}, \widehat{\rho}_{t-1})$ couples μ_{t-1} and $K_{\theta,t}\mu_t$. The upper Fuchs–van de Graaf inequality (Fuchs & van de Graaf, 1999) and concavity of $f \mapsto \sqrt{1-f^2}$ give

$$W_{1,\mathrm{tr}}(\mu_{t-1}, K_{\theta,t}\mu_t) \le \mathbb{E}d_{\mathrm{tr}}(\rho_{X,t-1}, \widehat{\rho}_{t-1}) \tag{100}$$

$$\le \mathbb{E}\sqrt{1 - F_{\mathrm{r}}(\rho_{X,t-1}, \widehat{\rho}_{t-1})^2} \tag{101}$$

$$\le \sqrt{1 - \left(\mathbb{E}F_{\mathrm{r}}(\rho_{X,t-1}, \widehat{\rho}_{t-1})\right)^2}. \tag{102}$$

This proves Equation (99). Applying Lemma 10 proves Equation (98). Independence between the target label and the measurement outcome is not needed: the conditional outcome probabilities are exactly those of the learned instrument. □

A.5.1 Composition preserves the instrument structure

Lemma 12 (Instrument sequences have a uniform stability factor)**.** *For a fixed sequence of finite-outcome instruments, every nonempty consecutive composition of their state-valued kernels obeys*

$$W_{1,\mathrm{tr}}\big((K_{\theta,s}\circ\cdots\circ K_{\theta,t})(\rho,\cdot),(K_{\theta,s}\circ\cdots\circ K_{\theta,t})(\sigma,\cdot)\big)\le 3d_{\mathrm{tr}}(\rho,\sigma),\qquad s\le t. \tag{103}$$

Here $K_{\theta,t}$ acts first. The factor does not grow with the length of the sequence.

Proof. A measurement history $(b_t,\dots,b_s)$ corresponds to the completely positive map

$$\mathcal{M}_{\theta,s,b_s}\circ\cdots\circ\mathcal{M}_{\theta,t,b_t}.$$

Summing these maps over all histories is a composition of trace-preserving channels, hence is trace-preserving. The product of successive conditional Born probabilities telescopes to

$$\mathrm{Tr}[(\mathcal{M}_{\theta,s,b_s}\circ\cdots\circ\mathcal{M}_{\theta,t,b_t})(\rho)]\,,$$

and the final conditional state is the same composed subnormalized output divided by this trace. Histories of zero probability have zero final mass. The sequence therefore induces the output law of a single finite-outcome instrument. Applying Lemma 9 to this instrument proves Equation (103). □

Lemma 13 (Wasserstein propagation through a kernel)**.** *Let K be a Borel kernel. If*

$$W_{1,\mathrm{tr}}(K(\rho,\cdot),K(\sigma,\cdot))\le\kappa d_{\mathrm{tr}}(\rho,\sigma)$$

for every pair of states, then

$$W_{1,\mathrm{tr}}(K\mu,K\nu)\le\kappa W_{1,\mathrm{tr}}(\mu,\nu). \tag{104}$$

Proof. For every real-valued 1-Lipschitz function f on the compact state space, the function $\rho\mapsto\int f(\sigma)K(\rho,\mathrm{d}\sigma)$ is κ-Lipschitz by Kantorovich–Rubinstein duality. Applying the same duality to μ,ν proves Equation (104). □

A.5.2 Proof of the compositional endpoint bound

Insert the intermediate laws $\mu_0, K_{\theta,1}\mu_1, (K_{\theta,1}\circ K_{\theta,2})\mu_2,\dots,(K_{\theta,1}\circ\cdots\circ K_{\theta,T})\mu_T,\nu_0$. The triangle inequality gives

$$\begin{aligned}W_{1,\mathrm{tr}}(\mu_0,\nu_0)\le{}& W_{1,\mathrm{tr}}(\mu_0,K_{\theta,1}\mu_1)\\&+\sum_{t=2}^{T}W_{1,\mathrm{tr}}((K_{\theta,1}\circ\cdots\circ K_{\theta,t-1})\mu_{t-1},(K_{\theta,1}\circ\cdots\circ K_{\theta,t})\mu_t)\\&+W_{1,\mathrm{tr}}((K_{\theta,1}\circ\cdots\circ K_{\theta,T})\mu_T,(K_{\theta,1}\circ\cdots\circ K_{\theta,T})\nu_T)\,.\end{aligned} \tag{105}$$

Apply Lemmas 12 and 13 to each entire downstream sequence, rather than multiplying per-step estimates. This yields

$$\begin{aligned}W_{1,\mathrm{tr}}(\mu_0,\nu_0)\le{}& 3W_{1,\mathrm{tr}}(\mu_T,\nu_T)+W_{1,\mathrm{tr}}(\mu_0,K_{\theta,1}\mu_1)\\&+3\sum_{t=2}^{T}W_{1,\mathrm{tr}}(\mu_{t-1},K_{\theta,t}\mu_t).\end{aligned} \tag{106}$$

Substitute Equation (98) with the actual risks, or with any certified upper bounds on them. In particular, taking $\theta=\widehat{\theta}$ and $L_{\gamma,t}(\widehat{\theta})\le\Delta_t+\varepsilon_{\mathrm{cal}}$ proves Equation (18). The terminal estimate is proved in Appendix B.8.2. This argument also covers $T=1$, with an empty sum over $t\ge 2$.

A.5.3 The nonexpansive latent-channel specialization

If the branch is a CPTP channel $\mathcal{R}_{\theta,t,z}$ selected under an input-independent probability law π_t, its kernel is

$$K_{\theta,t}(\rho,B)=\int\mathbf{1}_B(\mathcal{R}_{\theta,t,z}(\rho))\,\pi_t(\mathrm{d}z). \tag{107}$$

Couple two outputs with the same latent draw. Trace-distance contractivity gives

$$W_{1,\mathrm{tr}}(K_{\theta,t}(\rho,\cdot),K_{\theta,t}(\sigma,\cdot)) \le \int d_{\mathrm{tr}}(\mathcal{R}_{\theta,t,z}(\rho),\mathcal{R}_{\theta,t,z}(\sigma))\pi_t(\mathrm{d}z) \tag{108}$$

$$\le d_{\mathrm{tr}}(\rho,\sigma). \tag{109}$$

A sequence with fresh independent draws has the same property by coupling the entire latent history. The preceding endpoint proof consequently gives

$$W_{1,\mathrm{tr}}(\mu_0,\nu_0) \le W_{1,\mathrm{tr}}(\mu_T,\nu_T) + \sum_{t=1}^{T} \psi_\gamma(\Delta_t + \varepsilon_{\mathrm{cal}}). \tag{110}$$

This is the stronger certificate for the input-independent CPTP models; it is not assumed for a general measurement-conditioned instrument.

A.5.4 Worst-step schedule optimality

Since ψ_γ is nondecreasing, Equation (18) implies

$$W_{1,\mathrm{tr}}(\mu_0,\nu_0) \le 3W_{1,\mathrm{tr}}(\mu_T,\nu_T) + (3T-2)\psi_\gamma(\max_t \Delta_t + \varepsilon_{\mathrm{cal}}). \tag{111}$$

For the latent-channel specialization, the corresponding expression is

$$W_{1,\mathrm{tr}}(\mu_0,\nu_0) \le W_{1,\mathrm{tr}}(\mu_T,\nu_T) + T\psi_\gamma(\max_t \Delta_t + \varepsilon_{\mathrm{cal}}). \tag{112}$$

At fixed endpoints, T, terminal mismatch and common calibration slack, Theorem 1 minimizes both displayed right-hand sides. Where $e^{-a/2} - \gamma a/[2\log(1/\gamma)] > 0$, $\psi_\gamma(a)$ is strictly increasing, so uniqueness follows from uniqueness of the equal-information grid. If the positive part vanishes, the bound saturates and ties are possible. This statement concerns the worst-step relaxation, not optimization of the weighted sum in Equation (18) or of the actual learned endpoint error. Calibration slack and circuit approximation error may themselves change with the chosen grid.

B ADDITIONAL THEORETICAL REFINEMENTS AND CERTIFICATES

This section contains supporting results that sharpen or extend the main theory but are not needed to follow the main proof chain: pairwise-geometric recovery witnesses, decision-theoretic benchmarks, explicit Petz representatives, approximate-schedule statements, finite-shot validation, restricted-channel certificates, and instrument stability.

B.1 PROOF OF THE PAIRWISE-GEOMETRIC CONVERSE

For each pair $x < y$, define the distinguishability erased by the tth forward step,

$$\delta_{xy}^{(t)} := d_{\mathrm{tr}}(\rho_{x,t-1}, \rho_{y,t-1}) - d_{\mathrm{tr}}(\rho_{x,t}, \rho_{y,t}) \geq 0, \tag{113}$$

where nonnegativity follows from contractivity of trace distance under the forward CPTP map. Define the geometry-only relaxation

$$\varepsilon_t^{\mathrm{geo}} := \min_{r_x \geq 0} \left\{ \sum_x p_x r_x : \; r_x + r_y \geq \delta_{xy}^{(t)} \; \forall x < y \right\}. \tag{114}$$

Fix a label-independent CPTP recovery channel $\mathcal{R}$ and set

$$e_x := d_{\mathrm{tr}}\big(\rho_{x,t-1}, \mathcal{R}(\rho_{x,t})\big). \tag{115}$$

For any pair $x < y$, the triangle inequality gives

$$d_{\mathrm{tr}}(\rho_{x,t-1}, \rho_{y,t-1}) \leq e_x + d_{\mathrm{tr}}\big(\mathcal{R}(\rho_{x,t}), \mathcal{R}(\rho_{y,t})\big) + e_y \tag{116}$$

$$\leq e_x + d_{\mathrm{tr}}(\rho_{x,t}, \rho_{y,t}) + e_y, \tag{117}$$

where the second line again uses contractivity under CPTP maps. Hence

$$e_x + e_y \geq \delta_{xy}^{(t)} \tag{118}$$

for every pair. The channel-induced vector $(e_x)_x$ is therefore feasible for Equation (114), so

$$\sum_x p_x e_x \geq \varepsilon_t^{\mathrm{geo}}. \tag{119}$$

Taking the infimum over $\mathcal{R}$ yields the channel-independent converse

$$\varepsilon_t^\star \geq \varepsilon_t^{\mathrm{geo}}. \tag{120}$$

For the depolarizing path,

$$\rho_{x,t} - \rho_{y,t} = \lambda_t(\rho_x - \rho_y), \tag{121}$$

so homogeneity of the trace norm yields

$$d_{\mathrm{tr}}(\rho_{x,t}, \rho_{y,t}) = \lambda_t d_{\mathrm{tr}}(\rho_x, \rho_y). \tag{122}$$

Therefore

$$\delta_{xy}^{(t)} = (\lambda_{t-1} - \lambda_t) d_{\mathrm{tr}}(\rho_x, \rho_y). \tag{123}$$

Define the ensemble-geometric LP

$$\mathfrak{G}(\mathcal{E}) := \min_{r_x \geq 0} \left\{ \sum_x p_x r_x : \; r_x + r_y \geq d_{\mathrm{tr}}(\rho_x, \rho_y) \; \forall x < y \right\}. \tag{124}$$

Positive homogeneity gives

$$\varepsilon_t^{\mathrm{geo}} = (\lambda_{t-1} - \lambda_t)\mathfrak{G}(\mathcal{E}). \tag{125}$$

Finally, writing $d_{xy} = d_{\mathrm{tr}}(\rho_x, \rho_y)$, standard LP duality gives

$$\mathfrak{G}(\mathcal{E}) = \max_{\alpha_{xy} \geq 0} \left\{ \sum_{x<y} \alpha_{xy} d_{xy} : \; \sum_{y \neq x} \alpha_{xy} \leq p_x \; \forall x \right\}, \tag{126}$$

with $\alpha_{xy} = \alpha_{yx}$ understood on unordered pairs. Feasibility and boundedness are immediate, so strong duality applies.

For later use, monotonicity of $\Omega_{d,m}$ together with Equation (45) gives $\varepsilon_t^\star \geq \Omega_{d,m}^{\leftarrow}(\Delta_t)$. Combining this with Equation (120), define

$$b_t := \max\{\Omega_{d,m}^{\leftarrow}(\Delta_t), \varepsilon_t^{\mathrm{geo}}\}, \qquad \varepsilon_t^\star \geq b_t. \tag{127}$$

This appendix-only strengthening also yields a risk-level refinement. For an arbitrary $\mathcal{R}$, let $d_x = d_{\mathrm{tr}}(\rho_{x,t-1}, \mathcal{R}(\rho_{x,t}))$ and $F_x = F_{\mathrm{r}}(\rho_{x,t-1}, \mathcal{R}(\rho_{x,t}))$. The upper Fuchs–van de Graaf inequality and concavity of $r \mapsto \sqrt{1-r^2}$ give

$$\sum_x p_x F_x \leq \sqrt{1 - \left(\sum_x p_x d_x\right)^2} \leq \sqrt{1 - b_t^2}. \tag{128}$$

Hence

$$-\log(1 - b_t^2) \leq \inf_{\mathcal{R}} \mathcal{L}_t^{\mathrm{cq}}(\mathcal{R}) \leq \inf_{\mathcal{R}} \mathcal{L}_t^{\log}(\mathcal{R}) \leq \Delta_t. \tag{129}$$

□

B.2 Decision-value witness and the exact orthogonal benchmark

Corollary 14 (Exact orthogonal depolarizing benchmark)**.** *For the uniform complete orthogonal ensemble*

$$\mathcal{E}_\lambda^{(d)} = \left\{\frac{1}{d}, \rho_x^\lambda = \lambda|x\rangle\langle x| + (1-\lambda)\frac{\mathbb{I}}{d}\right\}_{x=1}^{d},$$

the optimal common-channel average trace error for recovering retention a from retention $b < a$ is

$$\varepsilon^\star(a \leftarrow b) = (a-b)\left(1 - \frac{1}{d}\right). \tag{130}$$

A simple operational witness follows from the randomization/deficiency viewpoint for quantum statistical experiments (Jenčová, 2016). Let $\mathcal{A}$ be a finite action set and let $u(x,a) \in [0,1]$ be a bounded utility. Define the Bayes value

$$V_u(\mathcal{E}) := \max_{\{M_a\}} \sum_{x,a} p_x u(x,a) \operatorname{Tr}(M_a \rho_x), \tag{131}$$

where the maximum ranges over POVMs. Then every common recovery channel obeys

$$\sum_x p_x d_{\mathrm{tr}}\big(\rho_{x,t-1}, \mathcal{R}(\rho_{x,t})\big) \geq \left[V_u(\mathcal{E}_{t-1}) - V_u(\mathcal{E}_t)\right]_+. \tag{132}$$

Indeed, let $\sigma_x = \mathcal{R}(\rho_{x,t})$ and choose a POVM $\{M_a\}$ optimal for $\mathcal{E}_{t-1}$. For each x, the operator $A_x := \sum_a u(x,a) M_a$ satisfies $0 \leq A_x \leq \mathbb{I}$, hence

$$\operatorname{Tr}[A_x(\rho_{x,t-1} - \sigma_x)] \leq d_{\mathrm{tr}}(\rho_{x,t-1}, \sigma_x). \tag{133}$$

Summing with p_x gives

$$V_u(\mathcal{E}_{t-1}) - V_u(\{p_x, \sigma_x\}) \leq \sum_x p_x d_{\mathrm{tr}}(\rho_{x,t-1}, \sigma_x). \tag{134}$$

Since $\{p_x, \sigma_x\}$ is a post-processing of $\mathcal{E}_t$, its Bayes value is at most $V_u(\mathcal{E}_t)$. This proves Equation (132). Choosing $u(x,a) = \mathbf{1}\{a = x\}$ yields

$$\varepsilon_t^\star \geq P_{\mathrm{guess}}(\mathcal{E}_{t-1}) - P_{\mathrm{guess}}(\mathcal{E}_t). \tag{135}$$

We now prove Theorem 14. For the uniform complete orthogonal ensemble $\mathcal{E}_\lambda^{(d)}$, all states commute in the computational basis. Dephasing any POVM in this common eigenbasis leaves its outcome

probabilities unchanged, reducing minimum-error discrimination to the corresponding classical problem; the maximum-likelihood basis decision is therefore optimal. Hence

$$P_{\text{guess}}(\mathcal{E}_\lambda^{(d)}) = \lambda + \frac{1-\lambda}{d}. \tag{136}$$

Therefore Equation (135) gives

$$\varepsilon^\star(a \leftarrow b) \geq (a-b)\left(1 - \frac{1}{d}\right). \tag{137}$$

For the matching upper bound, use the identity channel. Since

$$\rho_x^a - \rho_x^b = (a-b)\left(|x\rangle\langle x| - \frac{\mathbb{I}}{d}\right), \tag{138}$$

and $|x\rangle\langle x| - \mathbb{I}/d$ has one eigenvalue $1 - 1/d$ and $d-1$ eigenvalues $-1/d$,

$$d_{\text{tr}}(\rho_x^a, \rho_x^b) = (a-b)\left(1 - \frac{1}{d}\right). \tag{139}$$

Averaging over x matches the lower bound and proves the corollary. □

Corollary 15 (Stochastic-branch converse)**.** *If $Z \sim \pi_t$ is input-independent and every $\mathcal{R}_{t,z}$ is label-independent CPTP, then*

$$\mathbb{E}_{X,Z} d_{\text{tr}}(\rho_{X,t-1}, \mathcal{R}_{t,Z}(\rho_{X,t})) \geq \varepsilon_t^\star \geq \max\{\Omega_{d,m}^{\leftarrow}(\Delta_t), \varepsilon_t^{\text{geo}}\}. \tag{140}$$

B.3 Proof of the stochastic-branch converse

For every fixed latent value z, the branch $\mathcal{R}_{t,z}$ is itself a label-independent CPTP channel and is therefore feasible in the infimum defining $\varepsilon_t^\star$. Hence

$$\sum_x p_x d_{\text{tr}}(\rho_{x,t-1}, \mathcal{R}_{t,z}(\rho_{x,t})) \geq \varepsilon_t^\star$$

for every z. Averaging with respect to the input-independent latent law π_t proves the first inequality in Equation (140); the second is the strengthened appendix bound Equation (127). □

Corollary 16 (Instrument recovery and the common-channel optimum)**.** *Let $\{\mathcal{M}_{t,b}\}_b$ be any label-independent finite-outcome instrument, and let $\widehat{\rho}_{t-1}$ have the conditional-output law Equation* (11) *at $\rho_{X,t}$. Then*

$$\mathbb{E} d_{\text{tr}}(\rho_{X,t-1}, \widehat{\rho}_{t-1}) \geq \varepsilon_t^\star. \tag{141}$$

Moreover, the infimum of the left-hand side over all such instruments equals $\varepsilon_t^\star$.

Proof. The averaged map $\sum_b \mathcal{M}_{t,b}$ is label-independent and CPTP. Convexity of trace distance gives, for each x,

$$\sum_b \operatorname{Tr} \mathcal{M}_{t,b}(\rho_{x,t}) d_{\text{tr}}\left(\rho_{x,t-1}, \frac{\mathcal{M}_{t,b}(\rho_{x,t})}{\operatorname{Tr} \mathcal{M}_{t,b}(\rho_{x,t})}\right) \geq d_{\text{tr}}\left(\rho_{x,t-1}, \sum_b \mathcal{M}_{t,b}(\rho_{x,t})\right),$$

with zero-mass outcomes omitted. Averaging over x proves Equation (141). Conversely, every CPTP channel is a one-outcome instrument, so an optimal common channel attains the same unrestricted infimum. □

Corollary 17 (Blockwise common-channel recovery)**.** *For $0 \leq s < t \leq T$ with $\lambda_s > 0$, let $\Delta_{s:t} = \chi_s - \chi_t$ and define $\varepsilon_{s:t}^\star$ analogously. Then*

$$\max\left\{\Omega_{d,m}^{\leftarrow}(\Delta_{s:t}),\ (\lambda_s - \lambda_t)\mathfrak{G}(\mathcal{E})\right\} \leq \varepsilon_{s:t}^\star \leq \sqrt{1 - e^{-\Delta_{s:t}}}. \tag{142}$$

Moreover, a common channel $\mathcal{R}_{s\leftarrow t}$ satisfies

$$\sum_x p_x F_{\text{r}}(\rho_{x,s}, \mathcal{R}_{s\leftarrow t}(\rho_{x,t})) \geq e^{-\Delta_{s:t}/2}, \quad \sum_x p_x d_{\text{tr}}(\rho_{x,s}, \mathcal{R}_{s\leftarrow t}(\rho_{x,t})) \leq \sqrt{1 - e^{-\Delta_{s:t}}}. \tag{143}$$

B.4 Proof of the blockwise recovery corollary

For $\lambda_s > 0$, the depolarizing semigroup identity gives

$$\rho_{x,t} = \mathcal{D}_{\lambda_t/\lambda_s}(\rho_{x,s}), \qquad \bar{\rho}_t = \mathcal{D}_{\lambda_t/\lambda_s}(\bar{\rho}_s). \tag{144}$$

Apply the universal recovery theorem to the reference state $\bar{\rho}_s$ and the cumulative channel $\mathcal{D}_{\lambda_t/\lambda_s}$. The support argument used in the one-step proof applies unchanged, so one channel works for all labels. The Holevo-continuity argument from Theorem 3 applies verbatim to this cumulative channel. The pairwise argument also applies to the block endpoints; depolarizing homogeneity gives pairwise distinguishability loss $(\lambda_s - \lambda_t) d_{\mathrm{tr}}(\rho_x, \rho_y)$, and the same LP-scaling argument as in Equation (125) contributes $(\lambda_s - \lambda_t)\mathfrak{G}(\mathcal{E})$. Together these give the lower side of Equation (142). If

$$\delta^{\mathrm{rel}}_{x;s:t} := D(\rho_{x,s}\|\bar{\rho}_s) - D(\rho_{x,t}\|\bar{\rho}_t), \tag{145}$$

then $\delta^{\mathrm{rel}}_{x;s:t} \geq 0$, $\sum_x p_x \delta^{\mathrm{rel}}_{x;s:t} = \Delta_{s:t}$. Define the blockwise fidelity

$$F_{x;s:t} := F_{\mathrm{r}}\big(\rho_{x,s}, \mathcal{R}_{s\leftarrow t}(\rho_{x,t})\big). \tag{146}$$

Universal recoverability gives $F_{x;s:t} \geq e^{-\delta^{\mathrm{rel}}_{x;s:t}/2}$. Averaging and applying Jensen to the convex function $r \mapsto e^{-r/2}$ proves Equation (143). For the trace-distance statement, apply the upper Fuchs–van de Graaf inequality and then Jensen to the concave function $f \mapsto \sqrt{1-f^2}$:

$$\sum_x p_x d_{\mathrm{tr}}\big(\rho_{x,s}, \mathcal{R}_{s\leftarrow t}(\rho_{x,t})\big) \leq \sum_x p_x \sqrt{1 - F^2_{x;s:t}} \tag{147}$$

$$\leq \sqrt{1 - \left(\sum_x p_x F_{x;s:t}\right)^2} \tag{148}$$

$$\leq \sqrt{1 - e^{-\Delta_{s:t}}}. \tag{149}$$

The same channel proves the upper side of Equation (142), completing Theorem 17. □

Corollary 18 (Clipped-loss attainability)**.** *Viewing $\mathcal{R}^{\mathrm{U}}_t$ as a degenerate stochastic kernel, for every $\gamma \in (0,1)$,*

$$L_{\gamma,t}(\mathcal{R}^{\mathrm{U}}_t) \leq \mathcal{L}^{\log}_t(\mathcal{R}^{\mathrm{U}}_t) \leq \Delta_t. \tag{150}$$

Thus the same decrement is an attainable budget for the clipped loss used in training.

B.5 Proof of clipped-budget attainability

For every fidelity value $F \in [0,1]$,

$$-2\log\max\{\gamma, F\} \leq -2\log F, \tag{151}$$

with the right-hand side interpreted as $+\infty$ at $F = 0$. Apply this pointwise to the deterministic universal recovery channel and average over x. The second inequality is Theorem 2. □

Proposition 19 (Diamond-norm degradation of the budget)**.** *Let $\mathcal{R}_{\theta,t}$ be a deterministic CPTP channel and define*

$$\varepsilon_{\diamond,t}(\theta) = \tfrac{1}{2}\left\|\mathcal{R}_{\theta,t} - \mathcal{R}^{\mathrm{U}}_t\right\|_\diamond. \tag{152}$$

Then

$$\sum_x p_x F_{\mathrm{r}}\big(\rho_{x,t-1}, \mathcal{R}_{\theta,t}(\rho_{x,t})\big) \geq \left[e^{-\Delta_t/2} - \sqrt{2\varepsilon_{\diamond,t}(\theta)}\right]_+. \tag{153}$$

Whenever the positive part is nonzero, the cq risk is at most minus twice its logarithm. No analogous statewise average-log bound follows without an individual fidelity lower bound.

B.6 Proof of the channel-approximation proposition

By the definition of diamond norm,

$$d_{\mathrm{tr}}\big(\mathcal{R}_{\theta,t}(\rho), \mathcal{R}_t^{\mathrm{U}}(\rho)\big) \le \varepsilon_{\diamond,t}(\theta) \tag{154}$$

for every input state. The Powers–Størmer inequality and Hölder's inequality imply the standard fidelity continuity estimate

$$|F_{\mathrm{r}}(\sigma,\rho) - F_{\mathrm{r}}(\sigma,\rho')| \le \sqrt{\|\rho-\rho'\|_1} = \sqrt{2d_{\mathrm{tr}}(\rho,\rho')}. \tag{155}$$

Hence every statewise fidelity is reduced by at most $\sqrt{2\varepsilon_{\diamond,t}(\theta)}$. Averaging and applying the appendix-local fidelity consequence Equation (39) proves Equation (153). □

B.7 Petz and averaged rotated-Petz maps

The recovery theorem is used as an existence result, but an explicit representative is useful for teacher construction. Let σ be a state and $\mathcal{N}$ a channel. The Hilbert–Schmidt adjoint $\mathcal{N}^\dagger$ is defined by $\mathrm{Tr}[Y\mathcal{N}(X)] = \mathrm{Tr}[\mathcal{N}^\dagger(Y)X]$. On $\mathrm{supp}\,\mathcal{N}(\sigma)$, the Petz map is (Petz, 1986; Fawzi & Renner, 2015; Junge et al., 2018)

$$\mathcal{P}_{\sigma,\mathcal{N}}(A) = \sigma^{1/2}\mathcal{N}^\dagger\Big(\mathcal{N}(\sigma)^{-1/2}A\mathcal{N}(\sigma)^{-1/2}\Big)\sigma^{1/2}, \tag{156}$$

where inverse powers are generalized inverses on their supports. Following the convention of Junge et al. (2018), define the rotated map

$$\mathcal{R}^s_{\sigma,\mathcal{N}}(A) = \sigma^{-is}\mathcal{P}_{\sigma,\mathcal{N}}\big(\mathcal{N}(\sigma)^{is}A\mathcal{N}(\sigma)^{-is}\big)\sigma^{is}. \tag{157}$$

Each map is CPTP on $\mathrm{supp}\,\mathcal{N}(\sigma)$. Let Π project onto this support and choose a state τ_0 supported on $\mathrm{supp}\,\sigma$. The completion

$$\widetilde{\mathcal{T}}(A) = \mathcal{T}(\Pi A\Pi) + \mathrm{Tr}[(\mathbb{I}-\Pi)A]\tau_0 \tag{158}$$

is CPTP on the full output space and agrees with $\mathcal{T}$ on every relevant output $\mathcal{N}(\rho)$ with $\mathrm{supp}\,\rho \subseteq \mathrm{supp}\,\sigma$. Indeed, $\rho \le c\sigma$ for some finite c, so $\mathcal{N}(\rho) \le c\mathcal{N}(\sigma)$ and the output support is contained in $\mathrm{supp}\,\mathcal{N}(\sigma)$.

A universal recovery channel satisfying the inequality used in Equation (34) is

$$\mathcal{R}^{\mathrm{U}}_{\sigma,\mathcal{N}} = \int_{\mathbb{R}} \beta_0(s)\mathcal{R}^{s/2}_{\sigma,\mathcal{N}}\,\mathrm{d}s, \qquad \beta_0(s) = \frac{\pi}{2\left[\cosh(\pi s)+1\right]}. \tag{159}$$

The density β_0 integrates to one. The ordinary Petz map is the zero-rotation member of this family; it is not assumed to satisfy the universal fidelity remainder by itself for every ensemble.

B.8 Terminal randomness and approximate schedules

The endpoint theorem allows a general terminal prior. The first proposition isolates the role of stochasticity at complete depolarization; the second relates an estimated information grid to the true decrements on that selected grid.

Proposition 20 (Complete depolarization and stochastic generation)**.** *If $\mu_T = \delta_{\mathbb{I}/d}$ and every reverse step is deterministic, then every generated measure is a Dirac measure. Conversely, any finite ensemble $\sum_z q_z\delta_{\sigma_z}$ can be generated from $\delta_{\mathbb{I}/d}$ by sampling $Z \sim q$ and applying the replacer channel $\mathcal{R}_z(A) = \mathrm{Tr}(A)\sigma_z$.*

B.8.1 Proof of the terminal-randomness proposition

A deterministic channel maps a Dirac measure to a Dirac measure, and compositions preserve this property. Conversely, every map $\mathcal{R}_z(A) = \mathrm{Tr}(A)\sigma_z$ is CPTP. Sampling $Z \sim q$ therefore produces the output law $\sum_z q_z\delta_{\sigma_z}$ from any fixed input state, including $\mathbb{I}/d$. □

B.8.2 Proof of the truncated-terminal mismatch proposition

Set $\tau = \mathbb{I}/d$, the maximally mixed terminal prior used in the main text. Because the second marginal is a Dirac measure, its coupling with μ_T is unique. Hence

$$W_{1,\mathrm{tr}}(\mu_T, \delta_\tau) = \sum_x p_x d_{\mathrm{tr}}(\mathcal{D}_{\lambda_T}(\rho_x), \tau) \tag{160}$$

$$= \lambda_T \sum_x p_x d_{\mathrm{tr}}(\rho_x, \tau), \tag{161}$$

where the second equality uses $\tau = \mathbb{I}/d$ and $\mathcal{D}_{\lambda_T}(\rho_x) - \tau = \lambda_T(\rho_x - \tau)$. The trace distance from the maximally mixed state is convex in the state and therefore maximized on a pure state. For a pure state, $\rho - \tau$ has eigenvalues $1 - 1/d$ and $-1/d$ with multiplicity $d - 1$, so $d_{\mathrm{tr}}(\rho, \tau) \le 1 - 1/d$. This proves the terminal mismatch bound in Theorem 6. □

Proposition 21 (Approximate equal-information schedule)**.** *Suppose* $\sup_{\lambda\in[\lambda_T,1]} |\widehat{\chi}(\lambda) - \chi(\lambda)| \le \varepsilon_\chi$. *Assume the estimated endpoint drop* $\widehat{\Delta}_{\mathrm{tot}} := \widehat{\chi}(1) - \widehat{\chi}(\lambda_T)$ *is positive and that an exact equalizing grid exists; let* $\widehat{\lambda}_t$ *denote such a grid between the fixed endpoints. Define*

$$\Delta_t^{\mathrm{true}} := \chi(\widehat{\lambda}_{t-1}) - \chi(\widehat{\lambda}_t), \qquad \widehat{\Delta}_t := \frac{\widehat{\chi}(1) - \widehat{\chi}(\lambda_T)}{T}. \tag{162}$$

Then

$$\max_t |\Delta_t^{\mathrm{true}} - \widehat{\Delta}_t| \le 2\varepsilon_\chi, \tag{163}$$

$$\max_t |\Delta_t^{\mathrm{true}} - \Delta_{\mathrm{tot}}/T| \le 2\varepsilon_\chi \left(1 + \frac{1}{T}\right). \tag{164}$$

Thus under exact inversion the training and endpoint results may take $\varepsilon_{\mathrm{clk}} = 2\varepsilon_\chi$*, with* Δ_t *interpreted as the true decrement on the selected pilot grid.*

B.8.3 Proof of the schedule-estimation proposition

Let $\widehat{\Delta}_{\mathrm{tot}} = \widehat{\chi}(1) - \widehat{\chi}(\lambda_T)$. Uniform estimation error implies

$$\left|\widehat{\Delta}_{\mathrm{tot}} - \Delta_{\mathrm{tot}}\right| \le 2\varepsilon_\chi. \tag{165}$$

Let $\widehat{\lambda}_t$ denote the estimated grid and set $\Delta_t^{\mathrm{true}} := \chi(\widehat{\lambda}_{t-1}) - \chi(\widehat{\lambda}_t)$. The estimated decrement is $\widehat{\Delta}_t = \widehat{\Delta}_{\mathrm{tot}}/T$, while each endpoint evaluation differs from the true value by at most ε_χ. Hence

$$\left|\Delta_t^{\mathrm{true}} - \widehat{\Delta}_t\right| \le 2\varepsilon_\chi, \tag{166}$$

which proves Equation (163). The triangle inequality gives $|\Delta_t^{\mathrm{true}} - \Delta_{\mathrm{tot}}/T| \le \left|\Delta_t^{\mathrm{true}} - \widehat{\Delta}_t\right| + \left|\widehat{\Delta}_{\mathrm{tot}} - \Delta_{\mathrm{tot}}\right|/T \le 2\varepsilon_\chi(1 + 1/T)$, proving Equation (164). □

Theorem 22 (Estimated schedule with inexact numerical inversion)**.** *Assume the uniform event of Theorem 21. Let*

$$\widehat{y}_t = (1 - t/T)\widehat{\chi}(1) + (t/T)\widehat{\chi}(\lambda_T), \tag{167}$$

and suppose the numerical solver returns a nonincreasing grid $1 = \widetilde{\lambda}_0 \ge \cdots \ge \widetilde{\lambda}_T = \lambda_T \ge 0$*, with* $\lambda_T < 1$*, satisfying*

$$\left|\widehat{\chi}(\widetilde{\lambda}_t) - \widehat{y}_t\right| \le \varepsilon_{\mathrm{inv}} \qquad (t = 0, \ldots, T), \tag{168}$$

Thus all returned points lie in $[\lambda_T, 1]$*, and the fixed endpoints are enforced exactly. Define*

$$\widetilde{\Delta}_t^{\mathrm{true}} := \chi(\widetilde{\lambda}_{t-1}) - \chi(\widetilde{\lambda}_t), \qquad \widehat{\Delta}_t := \frac{\widehat{\chi}(1) - \widehat{\chi}(\lambda_T)}{T}. \tag{169}$$

Then

$$\max_t \left|\widetilde{\Delta}_t^{\mathrm{true}} - \widehat{\Delta}_t\right| \le 2\varepsilon_\chi + 2\varepsilon_{\mathrm{inv}}, \tag{170}$$

$$\max_t \left|\widetilde{\Delta}_t^{\mathrm{true}} - \Delta_{\mathrm{tot}}/T\right| \le 2\varepsilon_\chi \left(1 + \frac{1}{T}\right) + 2\varepsilon_{\mathrm{inv}}. \tag{171}$$

Let $\lambda_t^\star$ be the exact true-information grid point satisfying $\chi(\lambda_t^\star) = (1 - t/T)\chi(1) + (t/T)\chi(\lambda_T)$. If $\lambda_t^\star$ and $\widetilde{\lambda}_t$ lie in an interval on which $\chi'(\lambda) \ge m_* > 0$, then $|\widetilde{\lambda}_t - \lambda_t^\star| \le (2\varepsilon_\chi + \varepsilon_{\mathrm{inv}})/m_*$ for every t.

Proof. The residual and uniform curve error imply

$$\left|\chi(\widetilde{\lambda}_t) - \widehat{y}_t\right| \le \varepsilon_\chi + \varepsilon_{\mathrm{inv}}. \tag{172}$$

Since $\widehat{y}_{t-1} - \widehat{y}_t = \widehat{\Delta}_t$, subtracting two adjacent relations proves Equation (170). Moreover, $\left|\widehat{\Delta}_t - \Delta_{\mathrm{tot}}/T\right| \le 2\varepsilon_\chi/T$, which gives Equation (171). Finally, with $y_t^\star := (1 - t/T)\chi(1) + (t/T)\chi(\lambda_T)$, the endpoint interpolation gives $|\widehat{y}_t - y_t^\star| \le \varepsilon_\chi$. Combining this with $|\chi(\widetilde{\lambda}_t) - \widehat{y}_t| \le \varepsilon_\chi + \varepsilon_{\mathrm{inv}}$ yields $|\chi(\widetilde{\lambda}_t) - \chi(\lambda_t^\star)| \le 2\varepsilon_\chi + \varepsilon_{\mathrm{inv}}$. The mean-value theorem on the stated interval then gives the coordinate bound. □

B.9 Finite-shot calibration

The main learning theorem assumes exact loss evaluation. The results in this section provide an independent hardware-facing route based on pilot tomography and held-out validation.

Recall the monotone Audenaert entropy-continuity modulus $\mathfrak{f}_d$ defined in Equation (31). We reuse it here to control the information-schedule error from state tomography.

Theorem 23 (Finite-shot, path-uniform equal-information schedule)**.** *Assume the probabilities p_x are known. Let $\mathsf{M} = \{M_y\}_{y\in\mathcal{Y}}$ be informationally complete with Hermitian dual frame $\{F_y\}_{y\in\mathcal{Y}}$ and $\Lambda_{\mathsf{M}} = \sum_{y\in\mathcal{Y}} \|F_y\|_1$. Measure N_χ copies of each ρ_x, let $\widehat{q}_{xy}$ be the empirical frequencies, and define*

$$\widetilde{\rho}_x = \sum_{y\in\mathcal{Y}} \widehat{q}_{xy} F_y, \qquad \widehat{\rho}_x \in \arg\min_{\sigma\in\mathsf{S}(Q)} \|\sigma - \widetilde{\rho}_x\|_1 . \tag{173}$$

Propagate the physical estimates through the known forward channel,

$$\widehat{\rho}_x(\lambda) = \mathcal{D}_\lambda(\widehat{\rho}_x), \qquad \widehat{\chi}(\lambda) = S\left(\sum_x p_x \widehat{\rho}_x(\lambda)\right) - \sum_x p_x S(\widehat{\rho}_x(\lambda)). \tag{174}$$

For $\delta_\chi \in (0,1)$, set

$$\varepsilon_\rho = \min\left\{1, \Lambda_{\mathsf{M}} \sqrt{\frac{\log(2m|\mathcal{Y}|/\delta_\chi)}{2N_\chi}}\right\}. \tag{175}$$

With probability at least $1 - \delta_\chi$,

$$\max_x d_{\mathrm{tr}}(\widehat{\rho}_x, \rho_x) \le \varepsilon_\rho \tag{176}$$

and, simultaneously for all $\lambda \in [0,1]$,

$$|\widehat{\chi}(\lambda) - \chi(\lambda)| \le 2\mathfrak{f}_d(\varepsilon_\rho). \tag{177}$$

If the estimated endpoint drop $\widehat{\Delta}_{\mathrm{tot}} := \widehat{\chi}(1) - \widehat{\chi}(\lambda_T)$ is positive, the estimated equal-information grid exists uniquely and its budget error obeys

$$\max_t |\widehat{\Delta}_t - \Delta_t^{\mathrm{true}}| \le 4\mathfrak{f}_d(\varepsilon_\rho). \tag{178}$$

A sufficient condition on the same event is $\Delta_{\mathrm{tot}} > 4\mathfrak{f}_d(\varepsilon_\rho)$.

Theorem 24 (Independent finite-shot validation of recovery risk)**.** *Condition on a trained parameter $\widehat{\theta}$ and on a pilot grid, both fixed before validation. At each timestep, draw n_{v} iid validation observations; the same indexed trajectory may be reused across timesteps, so only iid sampling across the validation index is required. For every target/output pair, assume that both density operators can be prepared repeatedly. For a stochastic reverse step this means that the recorded latent value or measurement*

branch can be conditioned on during validation. Tomograph each state from N_{v} copies using the same POVM M and the physical estimator in Equation (173)*. Let $\widetilde{L}_{\gamma,t}(\widehat{\theta})$ be the empirical clipped loss computed from the tomographic estimates. Define*

$$\varepsilon_{\rho,\mathrm{v}} := \min\left\{1, \Lambda_{\mathsf{M}}\sqrt{\frac{\log(4n_{\mathrm{v}}T|\mathcal{Y}|/\delta_{\mathrm{tom}})}{2N_{\mathrm{v}}}}\right\}, \tag{179}$$

$$\varepsilon_{\mathrm{tom}} := \min\left\{2\log(1/\gamma), \frac{4\sqrt{2\varepsilon_{\rho,\mathrm{v}}}}{\gamma}\right\}, \tag{180}$$

$$\varepsilon_{\mathrm{samp}} := 2\log(1/\gamma)\sqrt{\frac{\log(T/\delta_{\mathrm{v}})}{2n_{\mathrm{v}}}}. \tag{181}$$

Then, with probability at least $1-\delta_{\mathrm{v}}-\delta_{\mathrm{tom}}$, simultaneously for all t,

$$L_{\gamma,t}(\widehat{\theta}) \le \widetilde{L}_{\gamma,t}(\widehat{\theta}) + \varepsilon_{\mathrm{samp}} + \varepsilon_{\mathrm{tom}}. \tag{182}$$

If

$$\max_t \left[\widetilde{L}_{\gamma,t}(\widehat{\theta}) - \widehat{\Delta}_t\right]_+ \le \eta_{\mathrm{v}}, \tag{183}$$

then, on the additional pilot event in Theorem 23, let $\varepsilon_{\mathrm{clk}}$ be any simultaneous upper bound on $\max_t|\widehat{\Delta}_t - \Delta_t^{\mathrm{true}}|$; under exact inversion one may take $\varepsilon_{\mathrm{clk}} = 4\mathfrak{f}_d(\varepsilon_\rho)$, while Theorem 22 adds $2\varepsilon_{\mathrm{inv}}$ for inexact inversion. Then

$$\max_t \left[L_{\gamma,t}(\widehat{\theta}) - \Delta_t^{\mathrm{true}}\right]_+ \le \eta_{\mathrm{v}} + \varepsilon_{\mathrm{samp}} + \varepsilon_{\mathrm{tom}} + \varepsilon_{\mathrm{clk}}. \tag{184}$$

The joint confidence is at least $1-\delta_\chi-\delta_{\mathrm{v}}-\delta_{\mathrm{tom}}$.

Corollary 25 (End-to-end finite-resource endpoint bound)**.** *Under Theorems 23 and 24, let $K_t := K_{\widehat{\theta},t}$ be the fixed reverse output-law kernel at timestep t. Write $K_t\rho$ for its output probability law on input state ρ, and define*

$$D_t := W_{1,\mathrm{tr}}(\mu_t,\nu_t), \qquad \kappa_t := \sup_{\rho\neq\sigma}\frac{W_{1,\mathrm{tr}}(K_t\rho, K_t\sigma)}{d_{\mathrm{tr}}(\rho,\sigma)}.$$

Set

$$\varepsilon_{\mathrm{FR}} := \eta_{\mathrm{v}} + \varepsilon_{\mathrm{samp}} + \varepsilon_{\mathrm{tom}} + \varepsilon_{\mathrm{clk}}. \tag{185}$$

With joint probability at least $1-\delta_\chi-\delta_{\mathrm{v}}-\delta_{\mathrm{tom}}$, the local recovery bound and kernel stability give, simultaneously for all t,

$$D_{t-1} \le \kappa_t D_t + \psi_\gamma(\Delta_t^{\mathrm{true}} + \varepsilon_{\mathrm{FR}}).$$

Iterating this recursion yields

$$D_0 \le \left(\prod_{j=1}^{T}\kappa_j\right)D_T + \sum_{t=1}^{T}\left(\prod_{j=1}^{t-1}\kappa_j\right)\psi_\gamma(\Delta_t^{\mathrm{true}} + \varepsilon_{\mathrm{FR}}). \tag{186}$$

For input-independent latent-conditioned CPTP kernels, $\kappa_t \le 1$, giving the corresponding worst-step bound with $\varepsilon_{\mathrm{FR}}$ in place of $\varepsilon_{\mathrm{cal}}$. For a finite-outcome instrument chain, the stronger whole-chain composition argument of Appendix A.5 gives

$$D_0 \le 3D_T + \psi_\gamma(\Delta_1^{\mathrm{true}} + \varepsilon_{\mathrm{FR}}) + 3\sum_{t=2}^{T}\psi_\gamma(\Delta_t^{\mathrm{true}} + \varepsilon_{\mathrm{FR}}), \tag{187}$$

so the factor 3 is not multiplied across timesteps.

B.9.1 Proof of the finite-shot schedule theorem

For fixed x, y, Hoeffding's inequality gives

$$\Pr(|\widehat{q}_{xy} - q_{xy}| > a) \le 2e^{-2N_\chi a^2}. \tag{188}$$

A union bound over $m|\mathcal{Y}|$ pairs, with

$$a = \sqrt{\frac{\log(2m|\mathcal{Y}|/\delta_\chi)}{2N_\chi}}, \tag{189}$$

shows that all frequency errors are at most a with probability $1 - \delta_\chi$. On this event,

$$\|\widetilde{\rho}_x - \rho_x\|_1 \le \sum_{y\in\mathcal{Y}} |\widehat{q}_{xy} - q_{xy}| \, \|F_y\|_1 \le \Lambda_{\mathsf{M}} a. \tag{190}$$

Because $\widehat{\rho}_x$ is a trace-norm projection onto the compact state space and the true state is feasible,

$$\|\widehat{\rho}_x - \rho_x\|_1 \le \|\widehat{\rho}_x - \widetilde{\rho}_x\|_1 + \|\widetilde{\rho}_x - \rho_x\|_1 \tag{191}$$

$$\le 2\,\|\widetilde{\rho}_x - \rho_x\|_1 . \tag{192}$$

Thus $d_{\mathrm{tr}}(\widehat{\rho}_x, \rho_x) \le \Lambda_{\mathsf{M}} a$, capped by one. Contractivity of trace distance under $\mathcal{D}_\lambda$ gives the same bound for all λ. Convexity of the trace norm gives

$$d_{\mathrm{tr}}\left(\sum_x p_x \widehat{\rho}_x(\lambda), \sum_x p_x \rho_x(\lambda)\right) \le \sum_x p_x d_{\mathrm{tr}}(\widehat{\rho}_x(\lambda), \rho_x(\lambda)) \le \varepsilon_\rho. \tag{193}$$

The sharp entropy-continuity inequality of Audenaert (2007), maximized over trace distances no larger than ε_ρ, gives

$$|S(\widehat{\rho}) - S(\rho)| \le \mathfrak{f}_d(\varepsilon_\rho). \tag{194}$$

Applying it once to the average state and once to every ensemble member yields Equation (177). Finally, $|\widehat{\Delta}_{\mathrm{tot}} - \Delta_{\mathrm{tot}}| \le 4\mathfrak{f}_d(\varepsilon_\rho)$, so $\Delta_{\mathrm{tot}} > 4\mathfrak{f}_d(\varepsilon_\rho)$ implies $\widehat{\Delta}_{\mathrm{tot}} > 0$. The forward-schedule theorem applied to the reconstructed nontrivial ensemble then gives a unique estimated grid, and Theorem 21 with $\varepsilon_\chi = 2\mathfrak{f}_d(\varepsilon_\rho)$ gives Equation (178). □

B.9.2 Proof of the finite-shot validation theorem

There are at most $2n_{\mathrm{v}}T$ states to reconstruct. Repeating the preceding Hoeffding–projection argument with a union bound over $2n_{\mathrm{v}}T|\mathcal{Y}|$ state–outcome pairs gives, with probability at least $1 - \delta_{\mathrm{tom}}$, trace-distance error at most $\varepsilon_{\rho,\mathrm{v}}$ for every reconstructed state. The fidelity continuity estimate Equation (155), applied to each argument in turn, gives

$$|F_{\mathrm{r}}(\rho, \sigma) - F_{\mathrm{r}}(\widehat{\rho}, \widehat{\sigma})| \le 2\sqrt{2\varepsilon_{\rho,\mathrm{v}}}. \tag{195}$$

Because $f \mapsto -2\log\max\{\gamma, f\}$ is $2/\gamma$-Lipschitz and has range $[0, 2\log(1/\gamma)]$, every tomographic loss differs from its ideal counterpart by at most $\varepsilon_{\mathrm{tom}}$. Conditional on the fixed trained model, Hoeffding's inequality and a union bound over timesteps give

$$L_{\gamma,t}(\widehat{\theta}) \le \widehat{L}^{\mathrm{ideal}}_{\gamma,t}(\widehat{\theta}) + \varepsilon_{\mathrm{samp}} \tag{196}$$

simultaneously for all t, with probability $1 - \delta_{\mathrm{v}}$. Combining the two events gives Equation (182). Adding the pilot budget error and taking positive parts proves Equation (184). □

B.9.3 Proof of the finite-resource endpoint corollary

On the joint pilot and validation event, Equation (184) gives

$$L_{\gamma,t}(\widehat{\theta}) \le \Delta_t^{\mathrm{true}} + \varepsilon_{\mathrm{FR}} \tag{197}$$

for every t. The proof of Theorem 6 then applies verbatim. □

B.10 RESTRICTED-CHANNEL AND CHOI CERTIFICATES

Universal recoverability is an unrestricted-channel statement. This section quantifies the loss incurred by a restricted Stinespring ansatz and gives a post-training channel certificate.

For a stochastic model, the corresponding barycentric (average-output) channel is

$$\overline{\mathcal{R}}_{\theta,t}(\rho) := \int_{\mathcal{Z}} \mathcal{R}_{\theta,t,z}(\rho)\pi_t(\mathrm{d}z). \tag{198}$$

It is CPTP but does not determine the full state-valued law $K_{\theta,t}(\rho,\cdot)$. For a finite-outcome instrument, the corresponding barycentric channel is instead $\overline{\mathcal{R}}_{\theta,t} = \sum_b \mathcal{M}_{\theta,t,b}$; the same distinction between an averaged channel and the conditional-output law remains.

Theorem 26 (Restricted-class attainability from Stinespring approximation)**.** *Assume $V_{\theta,t}$ and V_t^{U} are Stinespring isometries into a common dilation space $Q \otimes E$, and let $\mathsf{U}(E)$ be the unitary group on the environment. For any θ and timestep t, set*

$$\varepsilon_{\mathrm{St},t}(\theta) := \inf_{W\in\mathsf{U}(E)} \left\|(\mathbb{I}_Q \otimes W)V_{\theta,t} - V_t^{\mathrm{U}}\right\|_\infty, \qquad \varepsilon_{\mathrm{app}} := \inf_{\theta\in[-\pi,\pi]^P} \max_{1\le t\le T} \varepsilon_{\mathrm{St},t}(\theta). \tag{199}$$

Then

$$\tfrac{1}{2}\left\|\mathcal{R}_{\theta,t} - \mathcal{R}_t^{\mathrm{U}}\right\|_\diamond \le \varepsilon_{\mathrm{St},t}(\theta). \tag{200}$$

Let

$$\delta_{x,t}^{\mathrm{rel}} := D(\rho_{x,t-1}\|\bar\rho_{t-1}) - D(\rho_{x,t}\|\bar\rho_t), \qquad \Delta_t = \sum_x p_x \delta_{x,t}^{\mathrm{rel}}. \tag{201}$$

For a deterministic model, interpret $L_{\gamma,t}$ using a degenerate latent law concentrated on the single branch. Then

$$F_{\mathrm{r}}(\rho_{x,t-1}, \mathcal{R}_{\theta,t}(\rho_{x,t})) \ge \left[e^{-\delta_{x,t}^{\mathrm{rel}}/2} - \sqrt{2\varepsilon_{\mathrm{St},t}(\theta)}\right]_+, \tag{202}$$

$$L_{\gamma,t}(\theta) \le \mathcal{A}_{\gamma,t}(\varepsilon_{\mathrm{St},t}(\theta)), \tag{203}$$

where

$$\mathcal{A}_{\gamma,t}(\varepsilon) := \sum_x p_x \left[-2\log\max\left\{\gamma, \left[e^{-\delta_{x,t}^{\mathrm{rel}}/2} - \sqrt{2\varepsilon}\right]_+\right\}\right]. \tag{204}$$

If the infimum defining $\varepsilon_{\mathrm{app}}$ in Equation (199) *is attained, a minimizing shared parameter satisfies $\varepsilon_{\mathrm{St},t}(\theta) \le \varepsilon_{\mathrm{app}}$ for every t. In general, for every $\upsilon > 0$ there exists θ_υ with $\max_t \varepsilon_{\mathrm{St},t}(\theta_\upsilon) \le \varepsilon_{\mathrm{app}} + \upsilon$. At $\varepsilon = 0$, $\mathcal{A}_{\gamma,t}(0) \le \Delta_t$.*

For a single timestep, an unrestricted Stinespring unitary with ancilla dimension at least d^2 can realize $\mathcal{R}_t^{\mathrm{U}}$ exactly. A shared conditional model has zero approximation radius whenever it realizes the finite family $\{V_t^{\mathrm{U}}\}_{t=1}^T$ at the corresponding conditions.

For a linear map $\Phi : Q \to Q$, $\|\Phi\|_\diamond$ denotes the standard completely bounded trace norm. Define its normalized Choi operator

$$J(\Phi) := (\mathrm{id}_R \otimes \Phi)(|\Omega\rangle\langle\Omega|), \qquad |\Omega\rangle := \frac{1}{\sqrt{d}}\sum_{i=1}^d |i\rangle^R |i\rangle^Q. \tag{205}$$

For a channel, $J(\Phi)$ is a density operator; this is the normalized Choi–Jamiołkowski representation (Choi, 1975; Watrous, 2009).

Theorem 27 (Finite-shot Choi certificate for teacher approximation)**.** *Let $\widetilde{\mathcal{R}}_t^{\mathrm{U}}$ be a numerically implemented teacher satisfying*

$$d_{\mathrm{tr}}(J(\widetilde{\mathcal{R}}_t^{\mathrm{U}}), J(\mathcal{R}_t^{\mathrm{U}})) \le \varepsilon_{\mathrm{teach},t}. \tag{206}$$

After training a deterministic channel $\mathcal{R}_{\widehat\theta,t}$, perform tomography on N_J copies of the learned normalized Choi state at each timestep using an informationally complete POVM $\mathsf{M}_J = \{M_{J,y}\}_{y\in\mathcal{Y}_{\mathrm{Choi}}}$ on dimension d^2, with dual-frame condition number Λ_J. Let $\widehat{J}_t$ be the physical estimate and define

$$\varepsilon_J := \min\left\{1, \Lambda_J\sqrt{\frac{\log(2T|\mathcal{Y}_{\mathrm{Choi}}|/\delta_J)}{2N_J}}\right\}. \tag{207}$$

Then, with probability at least $1-\delta_J$*, simultaneously for all* t*,*

$$\tfrac{1}{2}\left\|\mathcal{R}_{\widehat{\theta},t}-\mathcal{R}_t^{\mathrm{U}}\right\|_\diamond \le \varepsilon_{\diamond,t}^{\mathrm{Choi}}, \tag{208}$$

where

$$\varepsilon_{\diamond,t}^{\mathrm{Choi}} := \min\Big\{1, d\left[d_{\mathrm{tr}}(\widehat{J}_t, J(\widetilde{\mathcal{R}}_t^{\mathrm{U}})) + \varepsilon_J + \varepsilon_{\mathrm{teach},t}\right]\Big\}. \tag{209}$$

Substituting $\varepsilon_{\diamond,t}^{\mathrm{Choi}}$ *into Theorem 19 gives a post-training average-fidelity certificate, and the same half-diamond control yields a clipped-risk certificate through Equation* (204)*. For a stochastic model, applying the construction to the barycentric channel* $\overline{\mathcal{R}}_{\widehat{\theta},t}$ *certifies the averaged channel, while branchwise expected loss is validated directly by Theorem 24.*

B.10.1 Proof of the Stinespring-approximation theorem

Environment unitaries do not change the induced channel, so fix one attaining the infimum up to an arbitrarily small error and write the aligned isometries as V and W. Because the difference of two channels is Hermiticity preserving, its diamond norm can be optimized on a pure input with a reference system of input dimension. For such an input, the trace distance between the two pure dilated outputs is at most the Euclidean distance of their state vectors, hence at most $\|V-W\|_\infty$. Partial trace cannot increase trace distance. Taking the supremum gives

$$\tfrac{1}{2}\left\|\mathcal{R}_{\theta,t}-\mathcal{R}_t^{\mathrm{U}}\right\|_\diamond \le \|V-W\|_\infty, \tag{210}$$

consistent with the continuity theorem for Stinespring representations (Kretschmann et al., 2008). Applying Equation (34) to the individual pair indexed by x gives

$$F_{\mathrm{r}}(\rho_{x,t-1}, \mathcal{R}_t^{\mathrm{U}}(\rho_{x,t})) \ge e^{-\delta_{x,t}^{\mathrm{rel}}/2}. \tag{211}$$

The fidelity-continuity argument in Theorem 19 then proves Equation (202). Applying the clipped loss, averaging over x, and using its monotonic decrease in fidelity gives Equation (203). At $\varepsilon=0$, each summand equals $\min\{\delta_{x,t}^{\mathrm{rel}}, 2\log(1/\gamma)\}$, whose average is at most Δ_t. □

B.10.2 A normalized-Choi norm lemma

For a Hermiticity-preserving map $\Phi: Q\to Q$ and the normalized Choi operator in Equation (205),

$$\|J(\Phi)\|_1 \le \|\Phi\|_\diamond \le d\,\|J(\Phi)\|_1. \tag{212}$$

The lower bound follows by evaluating the diamond norm on $|\Omega\rangle\langle\Omega|$. For the upper bound, any normalized pure reference-assisted input can be written as $(A\otimes\mathbb{I})|\Omega\rangle$ with $\|A\|_2^2 = d$. Therefore

$$\|(\mathrm{id}\otimes\Phi)(|\psi\rangle\langle\psi|)\|_1 = \left\|(A\otimes\mathbb{I})J(\Phi)(A^\dagger\otimes\mathbb{I})\right\|_1 \tag{213}$$

$$\le \|A\|_\infty^2\,\|J(\Phi)\|_1 \le d\,\|J(\Phi)\|_1. \tag{214}$$

For Hermiticity-preserving maps the diamond norm may be optimized on a pure input with a reference of input dimension, yielding the upper bound; see also Watrous (2009).

B.10.3 Proof of the Choi-validation theorem

The IC-tomography argument of Theorem 23, now in dimension d^2 and union-bounded over T Choi states, gives

$$d_{\mathrm{tr}}(\widehat{J}_t, J(\mathcal{R}_{\widehat{\theta},t})) \le \varepsilon_J \tag{215}$$

for all t with probability $1-\delta_J$. The triangle inequality then gives

$$d_{\mathrm{tr}}(J(\mathcal{R}_{\widehat{\theta},t}), J(\mathcal{R}_t^{\mathrm{U}})) \le d_{\mathrm{tr}}(\widehat{J}_t, J(\widetilde{\mathcal{R}}_t^{\mathrm{U}})) + \varepsilon_J + \varepsilon_{\mathrm{teach},t}. \tag{216}$$

Applying the upper bound in Equation (212) to the difference of the channels proves Equation (208); the cap at one uses the fact that half the diamond distance between channels is at most one. □

B.11 Stability of finite-outcome instruments

The nonexpansive result in the main text uses an external latent variable. A measurement-conditioned branch can amplify perturbations through normalization, so a general instrument need not be nonexpansive. The following finite-factor estimate records the distinction.

Proposition 28 (Finite-outcome instrument bound)**.** *Let $\{\mathcal{M}_y\}_{y=1}^r$ be a quantum instrument, so each $\mathcal{M}_y$ is completely positive and $\sum_y \mathcal{M}_y$ is trace-preserving. Let $K_{\mathcal{M}}(\rho)$ be the probability law of the normalized post-measurement state $\mathcal{M}_y(\rho)/\operatorname{Tr}\mathcal{M}_y(\rho)$, with outcome probability $\operatorname{Tr}\mathcal{M}_y(\rho)$; branches of zero probability may be assigned an arbitrary normalized state because they carry no mass. Then the Dobrushin coefficient in trace-Wasserstein distance satisfies $\kappa(K_{\mathcal{M}}) \le 3$.*

Proof. Set $A_y = \mathcal{M}_y(\rho)$, $B_y = \mathcal{M}_y(\sigma)$, $p_y = \operatorname{Tr} A_y$, and $q_y = \operatorname{Tr} B_y$. Couple the outcome distributions maximally. The unmatched mass is their total variation distance, which is at most $d_{\mathrm{tr}}(\rho,\sigma)$ by data processing. The matched mass at outcome y is $r_y = \min(p_y, q_y)$. For outcomes with $p_y q_y > 0$,

$$r_y \left\| A_y/p_y - B_y/q_y \right\|_1 \le \left\| A_y - B_y \right\|_1 + |p_y - q_y| \,. \tag{217}$$

If $p_y q_y = 0$, then $r_y = 0$ and the matched contribution of that outcome is zero, so no normalized-state quotient is needed. The cq output channel $\rho \mapsto \sum_y |y\rangle\langle y| \otimes \mathcal{M}_y(\rho)$ is trace-distance contractive, so

$$\sum_y \left\| A_y - B_y \right\|_1 \le \left\| \rho - \sigma \right\|_1 , \qquad \sum_y |p_y - q_y| \le \left\| \rho - \sigma \right\|_1 . \tag{218}$$

The matched contribution is therefore at most $2d_{\mathrm{tr}}(\rho,\sigma)$, and the unmatched contribution is at most $d_{\mathrm{tr}}(\rho,\sigma)$ because the diameter of state space under trace distance is one. Summing gives the factor three. □

B.12 Measurement-conditioned implementation

The MSQuDDPM reverse step prepares fresh ancillas, applies a trainable joint unitary, measures the ancillas, and propagates the normalized conditional system state (Kwun et al., 2025). For a fixed pure ancillary input $|0\rangle_E$ and readout basis $\{|b\rangle_E\}_b$, define

$$A_{\theta,t,b} = (\langle b|_E \otimes \mathbb{I}_Q) U_{\theta,t} (|0\rangle_E \otimes \mathbb{I}_Q), \qquad \mathcal{M}_{\theta,t,b}(\rho) = A_{\theta,t,b} \rho A_{\theta,t,b}^\dagger . \tag{219}$$

Then $\sum_b A_{\theta,t,b}^\dagger A_{\theta,t,b} = \mathbb{I}_Q$, and Equation (11) is exactly the Born-sampled output law. This does not make the normalized map $\mathcal{M}_{\theta,t,b}(\rho)/\operatorname{Tr}\mathcal{M}_{\theta,t,b}(\rho)$ a linear CPTP channel.

The recovery risk in Equation (13) therefore becomes the Born-weighted branch loss

$$L_{\gamma,t}(\theta) = \sum_x p_x \sum_b \operatorname{Tr}\mathcal{M}_{\theta,t,b}(\rho_{x,t}) \left[-2 \log \max \left\{ \gamma, F_{\mathrm{r}}\left(\rho_{x,t-1}, \frac{\mathcal{M}_{\theta,t,b}(\rho_{x,t})}{\operatorname{Tr}\mathcal{M}_{\theta,t,b}(\rho_{x,t})} \right) \right\} \right], \tag{220}$$

with zero-probability outcomes omitted. The average channel $\sum_b \mathcal{M}_{\theta,t,b}$ can be well calibrated while the conditional-output risk is not, so Equation (220)—not a fidelity computed after averaging the branches—is the quantity constrained in an MSQuDDPM-style implementation.

If an additional random ancillary state or other input-independent classical choice is sampled before the circuit, condition on that draw and include it as part of the instrument outcome. Fresh independent choices at successive timesteps preserve the Markov construction used in Appendix A.5.

B.13 Scope and limitations

Our scheduling result concerns a fixed depolarizing path: equal information minimizes the largest information decrement and the corresponding worst-step-derived bound, not necessarily the trained endpoint error. The recovery budget is an unrestricted common-channel reference; a finite-depth learner need not attain it, and the approximation results do not guarantee successful optimization. The endpoint certificate is one-sided and does not imply improvement in every aspect of the generated distribution. Finally, the finite-sample and tomography guarantees are worst-case uniform bounds and can be numerically conservative at the circuit sizes considered here. Our finite-support experiments therefore evaluate the corresponding finite-ensemble calibration quantities directly rather than relying on a numerically tight value of the uniform generalization bound.

C EXPERIMENTAL PROTOCOL AND DIAGNOSTIC DETAILS

C.1 TASK, OBJECTIVES, AND EVALUATION

Task and reverse model. We use the mixed-state ring task reconstructed from MSQuDDPM (Kwun et al., 2025). A state is drawn as

$$\rho = (1-q_0)|\psi(\varphi)\rangle\langle\psi(\varphi)| + q_0 I/2, \qquad |\psi(\varphi)\rangle = \cos(\varphi/2)|0\rangle + \sin(\varphi/2)|1\rangle,$$

where $\varphi \sim U[0, 2\pi)$ and $q_0 \sim U[0, 0.04)$ are independent. All models share the same 200 training states (dataset seed 202609071). Each of the eight reverse stages uses eight layers of R_x/R_y rotations and fixed nearest-neighbor CZ gates on one system qubit and two ancillas initialized to $|00\rangle$. The stages do not share parameters: there are 48 parameters per stage and 384 in total. Starting from $I/2$, generation samples the actual input-dependent Born outcome at every stage and propagates its normalized conditional output, rather than the outcome-averaged state. This is the instrument setting of Equations (219) and (220).

Paired forward schedules and objectives. Both schedules use global depolarization, $\rho_{x,t} = \lambda_t \rho_{x,0} + (1-\lambda_t) I/2$, with $\lambda_0 = 1$ and $\lambda_8 = 0$. For the implemented MS cosine schedule,

$$f_t = \cos^2\left(\frac{t/T + 0.008}{1.008}\frac{\pi}{2}\right), \quad a_t = f_t/f_0, \quad q_t = (1 - a_t/a_{t-1})^2, \quad \lambda_t = \prod_{s=1}^{t}(1-q_s),$$

with $f_T = 0$ set exactly. Equal-information scheduling instead solves $\chi(\lambda_t) = (1-t/T)\chi(1)$ using the actual equiprobable training ensemble and natural logarithms. Each schedule has its own $\Delta_t = \chi(\lambda_{t-1}) - \chi(\lambda_t)$.

The four groups cross these schedules with free or calibrated reverse training. Both use the original unregularized ensemble OT training objective with cost $1 - G(\rho, \sigma)$, where

$$G(\rho, \sigma) = \mathrm{Tr}(\rho\sigma) + \sqrt{(1 - \mathrm{Tr}\,\rho^2)(1 - \mathrm{Tr}\,\sigma^2)}.$$

The implementation floors each impurity factor at 10^{-15} for numerical stability. At each stage, the distributional objective uses 200 samples from the frozen learned reverse prefix and marginalizes the current four Born outcomes; its gradient includes both output-state and probability-mass terms. Calibrated training adds $\alpha_t(L^{\mathrm{pop}}_{\gamma,t} - \Delta_t)$ with $\gamma = 0.01$ and the projected update $\alpha_t \leftarrow [\alpha_t + 0.20(L^{\mathrm{pop}}_{\gamma,t} - \Delta_t)]_+$. There is no additional $1/T$ factor in this stagewise objective. The recovery risk is evaluated on the true forward pairs, not the generated-prefix inputs:

$$L^{\mathrm{pop}}_{\gamma,t} = \frac{1}{200}\sum_{x=1}^{200}\sum_{b=1}^{4} p_{\theta,t,b}(\rho_{x,t})\left[-2\log\max\{\gamma, F_{\mathrm{r}}(\rho_{x,t-1}, \widehat{\rho}_{\theta,t-1,b}(x))\}\right].$$

Here *population* denotes the exact finite training ensemble and all Born outcomes, not an expectation over unseen ring states. Accordingly, the reported Δ_t budgets and $L^{\mathrm{pop}}_{\gamma,t}$ are exact quantities for this fixed empirical ensemble. We do not invoke the fixed-grid finite-sample guarantee of Theorem 5 for the underlying continuous ring distribution; out-of-sample target draws are used for endpoint generation evaluation, while formal certification of a data-estimated grid is provided by the independent-pilot route in Appendix B.9.

Optimization and endpoint evaluation. Each stage receives 10,000 Adam updates, for 80,000 updates per model, trained in reverse order $8, \ldots, 1$. The learning rate is $10^{-3}10^{-j/4999}$ for updates indexed $j = 0, \ldots, 4999$, then 10^{-4} thereafter. In the original four-group experiment and the constrained-schedule comparison, each stage uses its lowest original training-OT iterate among iterates $1, \ldots, 10{,}000$, including in calibrated training. These checkpoints are also retained for the reused controls in Section 6.2; the local-only baseline uses the same training-OT selection criterion, as specified below. The four groups share 10 initialization seeds and matched sampling uniforms; common uniforms do not force identical Born outcomes. For each independently trained seed, endpoint evaluation uses twenty saved 200-versus-200 batches, each containing 200 raw generated states and 200 development target states. The twenty batch-level endpoint values are first averaged within each trained model, yielding 200 batch-level evaluations in total across the 10 seeds; these batches are sampling repetitions, not independently trained models.

The primary endpoint metric is $W_{1,\mathrm{tr}}$, with transport cost $d_{\mathrm{tr}}(\rho,\sigma) = \frac{1}{2}\|\rho-\sigma\|_1$ and no entropy regularization. It is distinct from the $1-G$ OT training objective. We also report $V_{\max}^{\mathrm{pop}} = \max_t [L_{\gamma,t}^{\mathrm{pop}} - \Delta_t]_+$ and the maximum local trace error

$$E_{\mathrm{local}}^{\max} = \max_{1\le t\le 8} \max_{\substack{1\le x\le 200,\ 1\le b\le 4 \\ p_{\theta,t,b}(\rho_{x,t})>0}} d_{\mathrm{tr}}(\rho_{x,t-1}, \widehat{\rho}_{\theta,t-1,b}(x)).$$

No average over labels or outcomes precedes this maximum.

C.2 Frozen schedule and objective ablations

The constrained-schedule comparison uses the same $T = 8$ architecture, 10 initialization seeds, 10,000-update learning-rate rule, $\gamma = 0.01$, dual learning rate 0.20, and endpoint evaluator. The additional linear schedule is linear in cumulative retention, $\lambda_t = 1 - t/8$, with its own finite-ensemble information decrements. The existing cosine/calibrated and equal-information/calibrated runs are reused without retraining or changing their checkpoints.

The training-objective comparison holds the equal-information grid fixed. Distribution-only and combined training reuse the original equal-information/free and equal-information/calibrated models, including their minimum-training-OT checkpoints. The local-only model directly minimizes $L_{\gamma,t}^{\mathrm{pop}}$ on the full true-forward training ensemble and all four Born outcomes. It uses neither an OT gradient nor dual updates. Its gradient inputs are the true forward pairs; generated-prefix inputs enter only checkpoint scoring, as described next.

All three methods select the lowest original stage-training $1-G$ OT value among iterates $1,\dots,10{,}000$, retaining the first iterate in an exact tie. At each stage, this score uses 200 samples from the model's own previously selected reverse prefix, the original sampling uniforms, and the same target ensemble, with the current four Born outcomes marginalized exactly. The local-only baseline evaluates this OT score without differentiating it; every parameter update uses only the recovery gradient. The primary endpoint trace-OT is never used for checkpoint selection.

The underlying runs use 10 seeds and 200 evaluations, matched initialization and sampling uniforms, and the unchanged learning-rate schedule and endpoint evaluator. Section 6.2 reuses the reference entries from Section 6.2 and reports the local-only mean across its three models. No additional time grid, seed, or hyperparameter search is included.

C.3 Angular allocation and radial fidelity

The direction-only diagnostic associated with Figure 3 normalizes each nonzero Bloch vector to unit length and computes empirical OT with cost $\frac{1}{2}\|\mathbf{u}-\mathbf{v}\|_2$, separately within each saved evaluation batch. This normalization is used only for the diagnostic, never for training, the primary endpoint metric, or the Bloch plots below. The plotted CDFs instead use the angle in the xz-plane; their deviation is not an exact decomposition of the three-dimensional trace-Wasserstein cost and depends on the fixed angular cut. Across all twenty batches of the two selected models, direction-only OT averages 0.066959 for equal-information/calibrated and 0.079007 for cosine/free, with the former lower in $14/20$ batches.

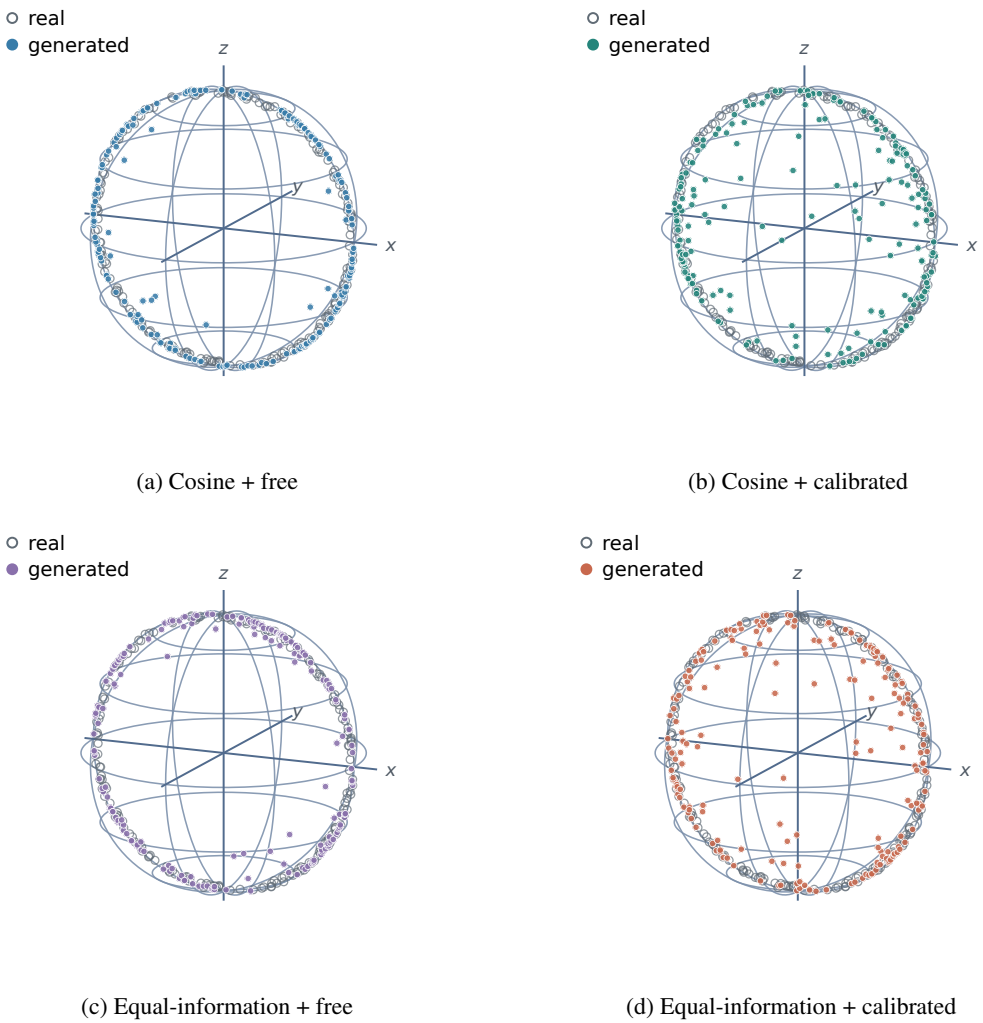
The free models place more generated states near the pure-state boundary, producing a thinner ring. Calibrated models retain the ring's angular coverage but include more interior states. As illustrated in Figure 3, reasonably accurate angular allocation can therefore coexist with imperfect radial fidelity: a lower trace-Wasserstein error does not necessarily imply a sharper ring.

C.4 Four-qubit TFIM protocol

Under the same network structure as in the ring experiment, we supplement the experiments with 4-qubit TFIM states. The states are defined using the open-boundary transverse-field Ising Hamiltonian

$$H(g) = -\sum_{i=1}^{3} Z_i Z_{i+1} - g\sum_{i=1}^{4} X_i.$$

We independently sample 100 fields $g \sim U[0.2, 2.2)$ (dataset seed 202609091) and diagonalize $H(g)$ to obtain the ground-state projectors $\rho_g = |\psi_0(g)\rangle\langle\psi_0(g)|$, forming a shared 100-state training



(a) Cosine + free

(b) Cosine + calibrated

(c) Equal-information + free

(d) Equal-information + calibrated

Figure 4: **Ring support is not fully captured by endpoint ranking.** Each panel overlays the same 200 real states (gray open circles) with all 200 generated states (colored points) from saved batch 0.

ensemble. Both TFIM studies use $T = 12$ reverse stages and circuit depth 12. The generation study uses four system qubits and two measured ancillas. The independent-recovery study uses four system qubits, two measured ancillas, and two additional hidden ancillas that are traced out.

C.5 TFIM RECOVERY–GENERATION DIAGNOSTIC

To investigate the remaining TFIM budget excess, we continue the original equal-information/calibrated checkpoints on the same previously selected ten-seed cohort. We retain the data, schedule, fixed spectral basis, and four-system/two-measured-ancilla construction. Writing $L_t = L^{\mathrm{pop}}_{\gamma,t}$, the joint continuation uses

$$J_\beta = J_{\mathrm{OT}} + \beta[L_t - \Delta_t]_+,$$

where J_{OT} is the stagewise trace-OT objective and β is fixed within each run, unlike the ring's primal–dual update. The reverse stages are visited in order $12 \to 1$; after updating a stage, the corresponding higher-stage prefix is used to regenerate the inputs for the next stage with the same sampling uniforms. Thus each β defines a training-consistent full-chain continuation from the same Original checkpoints, conditional on the previously selected cohort. We retain final iterates after at most 200 L-BFGS-B iterations per stage. Endpoint evaluation reuses identical targets, sampling uniforms, and the original evaluator. Recovery-only (R: L_t) and penalty-only (P: $0.003[L_t - \Delta_t]_+$) continuations are retained as recovery-focused references; R uses at most 2,000 iterations per stage. Depth-18 models extend depth 12 with function-preserving initialization. This is a conditional diagnostic, not independent confirmation.

Table 3: **TFIM recovery–generation diagnostic.** Mean ± s.e. across the same ten selected matched seeds; each endpoint first averages the same twenty evaluation batches. Joint rows vary the fixed recovery-penalty weight β from the same Original checkpoints. P and R use different recovery-focused objectives and are shown as references. Lower is better.

Setting	Depth	Max excess	Endpoint $W_{1,\mathrm{tr}}$
Original	12	0.164031 ± 0.011224	0.062123 ± 0.002675
Joint, $\beta = 0.003$	12	0.167981 ± 0.009410	0.064204 ± 0.004955
Joint, $\beta = 0.03$	12	0.139602 ± 0.007385	0.066005 ± 0.003944
Joint, $\beta = 0.3$	12	0.077117 ± 0.002909	0.085810 ± 0.007308
Joint, $\beta = 3$	12	0.023482 ± 0.001292	0.204143 ± 0.014698
P	12	0.017445 ± 0.001693	0.580027 ± 0.013632
R	12	0.017202 ± 0.001699	0.690155 ± 0.003869
Joint, $\beta = 0.003$	18	0.164951 ± 0.009861	0.066341 ± 0.006204
P	18	0.015744 ± 0.000948	0.573060 ± 0.013582
R	18	0.014483 ± 0.000890	0.683615 ± 0.001879

Increasing the joint recovery weight reveals an empirical operating curve rather than an immediate collapse of generation. At $\beta = 0.03$, the mean maximum excess decreases by 14.9% relative to Original while endpoint $W_{1,\mathrm{tr}}$ increases by 6.3%. At $\beta = 0.3$, the mean maximum excess decreases by 53.0%, from 0.164031 to 0.077117, while endpoint $W_{1,\mathrm{tr}}$ increases from 0.062123 to 0.085810; all ten matched seeds have lower excess and higher endpoint error than Original. Stronger recovery pressure at $\beta = 3$ reduces the mean maximum excess by 85.7% but incurs a substantially larger endpoint cost. Thus substantial recovery improvement need not immediately produce the severe generation degradation seen under recovery-only or penalty-only optimization.

None of the tested joint settings satisfies all twelve timestep budgets simultaneously. The observed curve is therefore an empirical recovery–generation operating frontier under this continuation protocol, not a Pareto-optimal frontier of the model class and not evidence of an intrinsic incompatibility between recovery and generation. It nevertheless shows that the residual violation of the original calibrated model is responsive to stronger recovery weighting, while further reduction increasingly competes with endpoint distribution matching. The depth-18 diagnostic changes the recovery-focused values only modestly and does not resolve the joint-model excess.

All R/P models retain their largest excess at $t = 12$. For this step, the label-independent replacement channel $\mathcal{R}_{\mathrm{mean}}(\rho) = \mathrm{Tr}(\rho)\sum_x p_x \rho_{x,11}$ attains $L^{\mathrm{pop}}_{\gamma,12} = 0.009205731 < \Delta_{12} = 0.020347712$.

Thus the unrestricted terminal budget is feasible, but realization by the tested circuit is not established. The finite learner's residual violation does not identify an expressivity lower bound or resolve optimization versus capacity.

C.6 External QuDDPM reproduction and cross-task transfer

We first reproduce the original four-qubit TFIM setting of QuDDPM (Zhang et al., 2024), $g \in [0.2, 0.4]$, using released data and the notebook-compatible historical official implementation. Released-weight replay matches saved trajectories to maximum state infidelity 2.4×10^{-7}. One complete retraining retains the original training function, $T = 30$, depth 12, and 3,001 Adam updates per stage. Below, both models use twenty matched batches of 100 fresh terminal inputs. These are sampling repetitions, not independent training runs. The native metrics are biased fidelity-kernel MMD^2 and computational-basis absolute magnetization $M = \mathbb{E}[|\sum_i Z_i|/4]$.

Original-setting evaluation	Magnetization M	MMD^2
Released weights	0.873524	0.024549
Retrained implementation	0.870924	0.023716

Comparable native metrics and successful replay support a basic implementation sanity check, not bitwise training reproduction. We retain the historical repeated measurement key; the authors' training minibatch sequence is unavailable.

We transfer this implementation to our 100-state, wide-field TFIM ensemble (Section C.4), retaining its scrambling forward process, reverse architecture, native objective, and final iterates. Three training replicas, each evaluated on twenty 200-state batches, attain endpoint $W_{1,\mathrm{tr}} = 0.410176 \pm 0.004385$ (mean $\pm$ s.e. across replicas), higher than the MSQuDDPM-style results in Table 2. This is a cross-model/cross-task transfer diagnostic, not evidence that QuDDPM is intrinsically inferior. The learners differ in forward corruption and reverse-model construction, with the MSQuDDPM-style learner specifically supporting mixed-state generation; budgets and model cohorts are also unmatched. Poorer transfer may thus reflect architecture/task mismatch. We report it for completeness and external implementation auditing, not as a primary performance comparison.