

Research Article

Exact Purification Rates for Alternating Qubit Measurements and Their Blind Spot at Right Angles

Agus Mulia Bakti¹

1. Universitas Pendidikan Mandalika, Indonesia

The usual way to quantify measurement incompatibility is noise robustness: how much noise must be mixed in before some joint measurement becomes possible. That is an extremal, worst-case quantity, computed by optimization, and it says little about what a real experiment would actually record if it just kept running. We ask a plainer question. A qubit is measured, over and over, in two projective bases separated by a Bloch angle θ . Once the resulting chain settles into steady state, how far does the state typically move between one measurement and the next? For sharp measurements this rate, $r(\theta)$, turns out to have the closed form $\frac{1}{2}\sin\theta\sqrt{1+\sin\theta}$. It agrees with the standard robustness threshold $\nu(\theta) = 1 - 1/\sqrt{1+\sin\theta}$ to first order near $\theta = 0$, then departs from it everywhere else. For unsharp (POVM) measurements in the weak-measurement limit $\lambda \rightarrow 0$, full-rank Kraus operators preserve purity exactly, and this collapses the problem onto a one-dimensional diffusion on the Bloch circle. The resulting rate is $c(\theta) = \sin(\theta/2)/[\sqrt{2}E(m)]$ with $m = 1 - \tan^2(\theta/2)$ and E the complete elliptic integral of the second kind, giving the clean value $c(\pi/2) = 1/\pi$ at the orthogonal point. A blind spot shows up along the way: at $\theta = \pi/2$, $r(\theta)$ stops responding to noise entirely, for any Pauli-type or amplitude-damping channel, simply because $\sin 45^\circ = \cos 45^\circ$. That is unfortunate news for anyone hoping to use $r(\theta)$ as a calibration check at the angle experimentalists tend to prefer. Purification under repeated or simultaneous incompatible measurement has by now a sizable literature, but in every case we examined the angle θ itself is held fixed and something else, measurement strength or the number of axes, is varied instead. Treating θ as the free parameter and solving for it in closed form is the gap this paper tries to close. Past the qubit, the same construction becomes a random walk on $\mathbb{CP}^{d-1}$ rather than on a circle, and its closed form we leave open.

Corresponding author: Agus Mulia Bakti, agusmuliabakti02@gmail.com

I. Introduction

Non-commuting observables are the oldest source of measurement incompatibility in quantum mechanics, and for a single pair of measurements, checked once, the consequences for joint measurability are by now textbook material; the recent *Reviews of Modern Physics* colloquium by Gühne, Haapasalo, Kraft, Pellonpää and Uola gives a broad survey of the field ^[1], and the closely related notion of quantum steering is reviewed by Uola, Costa, Nguyen and Gühne ^[2]. The tool everyone reaches for is *noise robustness*: how much depolarizing (or other) noise has to be mixed in before a joint, compatible measurement exists ^{[3][4][5]}. For two rank-one projective qubit measurements at Bloch angle θ apart, this threshold is well known,

$$\lambda_c(\theta) = \frac{1}{\cos(\theta/2) + \sin(\theta/2)} = \frac{1}{\sqrt{1 + \sin \theta}}, \quad (1)$$

and it gives back the familiar $1/\sqrt{2}$ at $\theta = \pi/2$, the anticommuting case of σ_z and σ_x .

But λ_c is an extremal, static number. It answers a worst-case question about a single shot, and computing it in general means solving an optimization, a semidefinite program once one leaves the qubit. None of that says what happens if the same two bases are actually measured over and over, in alternation, on one qubit, with the Born rule updating the state after every shot. That is a different situation, and unlike λ_c , it is something one could watch happen in a single run of the experiment.

Repeated and continuous measurement of non-commuting observables on a qubit is, of course, not a new subject. Most of it is written in the language of continuous weak measurement and quantum trajectories, that is, conditioned state evolution, filtering, and purification under monitoring, for which Wiseman and Milburn remains the standard reference ^[6]. Chantasri, Dressel and Jordan built a stochastic path-integral formulation that reaches much of the same territory from a different direction ^[7]. We do not attempt a reformulation of the alternating-basis problem in that language here, though it would be a natural next step. On the experimental side, continuous weak monitoring of a single superconducting qubit was demonstrated directly by Murch, Weber, Macklin and Siddiqi ^[8], and Weber, Murch, Schwartz, Roch and Siddiqi give a broader review of quantum trajectories in these circuit-QED systems ^[9]; this is exactly the kind of platform where the alternating-basis chain considered below could in principle be run. Measurement-induced phase transitions in many-body random circuits form a related but genuinely separate body of work, where purification of a monitored system is again the central object ^{[10][11]}. That setting is high-dimensional and many-body, and what gets varied there is the measurement rate across

an extensive circuit rather than the angle between two fixed single-qubit bases, so its closed-form results simply do not transfer to the problem considered here. Narrowing to qubit-specific weak measurement, Wei and Nazarov looked at monitoring and purification under simultaneous weak measurement of non-commuting variables [12]. Ruskov, Korotkov and Mølmer, and later Ruskov, Combes, Mølmer and Wiseman, worked out how much faster purification goes when two or three complementary observables are measured simultaneously [13][14]. Hacothen-Gourgy *et al.* actually built the experiment: simultaneous continuous measurement of two non-commuting observables on a superconducting qubit, with the diffusive dynamics observed directly [15]. Atalaya *et al.* derived the matching analytical correlators and checked them against data [16]. Closest to what we do here, Walls and Ford used quantum state diffusion to study *alternating* measurements of S_z and S_x , working out the resulting memory effects and purification cascades [17]. What all of this shares is a fixed relative angle between the two axes, usually exact orthogonality, with measurement strength or the number of simultaneous axes doing the varying. Nobody, as far as we can tell, has let the inter-basis angle θ itself be the free parameter and solved for the closed-form dependence on it.

That is the gap this paper fills. For an infinite alternating chain of measurements at Bloch angle θ , define the steady-state average single-step trace distance

$$r(\theta) := \lim_{n \rightarrow \infty} \mathbb{E} [D(\rho_n, \rho_{n-1})], \quad (2)$$

for sharp (projective) measurements, with an unsharp (POVM) counterpart $c(\theta, \lambda) = \lim_{\lambda \rightarrow 0} r(\theta, \lambda)/\lambda$ for weak measurements of sharpness λ . Neither quantity requires tomography or an optimization to estimate: a single measured trajectory is enough, and both come out as closed-form functions of θ . One terminological point is worth making explicit. $r(\theta)$ and $c(\theta)$ are operational distinguishability and purification rates, not incompatibility measures in the technical sense that $\nu(\theta)$ is. We call θ the “incompatibility angle” because it is the same Bloch-sphere parameter entering the standard joint-measurability threshold $\lambda_c(\theta)$, and because r, c agree with ν to first order near $\theta = 0$ (Sec. B), not because we think $r(\theta)$ or $c(\theta)$ qualifies as an incompatibility measure in the Busch, Heinosaari and Designolle sense. What follows derives these closed forms, checks them symbolically and numerically, compares them against the static robustness $\nu(\theta) = 1 - \lambda_c(\theta)$, and works out their sensitivity to physical noise, turning up an exact blind spot at $\theta = \pi/2$ along the way. Every claim below is tagged with its epistemic status, proved, symbolically verified, numerically verified, or open; a result that has not been derived in closed form is not presented as though it had been.

II. Model and definitions

Take a qubit and two orthonormal projective bases on the Bloch sphere, $A = \{|+n_A\rangle, |-n_A\rangle\}$ and $B = \{|+n_B\rangle, |-n_B\rangle\}$, with unit vectors n_A, n_B separated by angle θ . Measure in the sequence $A, B, A, B, \dots$, updating the state after each shot via the Lüders rule and whatever outcome actually occurs; not an average over outcomes, and not an adversarial choice. Figure 1 shows the geometry.

alternating chain: $A, B, A, B, \dots$




Figure 1. Schematic of the alternating measurement chain. The qubit is measured in the sequence $A, B, A, B, \dots$, where the two projective bases are separated by Bloch angle θ . Each step lands on the “near” pole of the other basis with probability $\cos^2(\theta/2)$ or the “far” pole with probability $\sin^2(\theta/2)$.

Lemma 1 (i.i.d. reduction). *Because each basis has two antipodal poles, every step, regardless of which pole the qubit currently occupies, lands on the “near” pole of the other basis (angular separation θ , probability $\cos^2(\theta/2)$) or the “far” pole (angular separation $\pi - \theta$, probability $\sin^2(\theta/2)$), independently of the history. The chain of per-step outcomes is therefore i.i.d.*

It is this antipodal symmetry that makes the sharp case solvable outright, with no Markov-chain machinery needed, just a single-step average.

For two pure qubit states φ apart on the Bloch sphere, the trace distance is $D = \sin(\varphi/2)$. So each step of the sharp chain contributes

$$D_{\text{near}}(\theta) = \sin(\theta/2), D_{\text{far}}(\theta) = \cos(\theta/2),$$

with probabilities $\cos^2(\theta/2)$ and $\sin^2(\theta/2)$ respectively.

III. Exact rate for sharp projective measurement chains

A. Closed form

From Lemma 1,

$$\begin{aligned} r(\theta) &= \cos^2(\theta/2)\sin(\theta/2) + \sin^2(\theta/2)\cos(\theta/2) \\ &= \sin(\theta/2)\cos(\theta/2)[\sin(\theta/2) + \cos(\theta/2)]. \end{aligned} \quad (3)$$

Using $\sin(\theta/2)\cos(\theta/2) = \frac{1}{2}\sin\theta$ and $[\sin(\theta/2) + \cos(\theta/2)]^2 = 1 + \sin\theta$,

$$\boxed{r(\theta) = \frac{1}{2}\sin\theta\sqrt{1 + \sin\theta}}. \quad (4)$$

Checks. $r(0) = 0$, and there is no cliff edge here: r climbs smoothly from zero rather than jumping like the binary yes/no of whether some infinite deterministic chain exists. At $\theta = \pi/2$, $r = \frac{1}{2} \cdot 1 \cdot \sqrt{2} = \sqrt{2}/2 \approx 0.7071$. Near the origin,

$$r(\theta) = \frac{\theta}{2} + O(\theta^2), \quad (5)$$

so the onset is plain linear, slope 1/2 in radians, with nothing fractional or anomalous going on at the compatible point.

B. Comparison with static noise robustness

Write the robustness itself, from Eq. (1), as $\nu(\theta) = 1 - \lambda_c(\theta) = 1 - 1/\sqrt{1 + \sin\theta}$. Both r and ν vanish at $\theta = 0$ and expand as $\theta/2 + O(\theta^2)$ near the origin, but that agreement is just the shared local Bloch geometry near the compatible point asserting itself, not a sign that the two functions are secretly the same thing. Away from the origin they are not proportional: $r/\nu \approx 1.67$ at $\theta = 30^\circ$, and ≈ 2.41 at $\theta = 90^\circ$. Substituting $u = 1/\lambda_c = \sqrt{1 + \sin\theta}$ gives an exact relation between them,

$$r = \frac{1 - \lambda_c^2}{2\lambda_c^3}, \quad (6)$$

manifestly non-linear, which is really just confirmation that r carries information ν alone does not have. The conceptual difference is worth stating plainly. $\nu(\theta)$ is extremal and adversarial, the noise needed for *some* joint measurement to exist, while $r(\theta)$ is an average over what an actual repeated Born-rule process does. They agree to first order and then go their separate ways. Figure 2 plots both over $\theta \in [0, \pi]$, with the ratio r/ν in the inset.




Figure 2. The operational rate $r(\theta)$ [Eq. (4)] and the static noise-robustness threshold $\nu(\theta) = 1 - \lambda_c(\theta)$ [Eq. (1)]. The two curves agree to first order near $\theta = 0$ but diverge globally; the inset shows the ratio r/ν , which grows monotonically from 1 toward $\theta = \pi/2$.

IV. Generalization to unsharp POVM chains

A. Model and the purification principle

Now swap each sharp projector for an unsharp POVM of sharpness $\lambda \in [0, 1]$ along direction n : $E_{\pm} = (I \pm \lambda n \cdot \sigma)/2$. The Lüders update splits into components parallel ($s_{\parallel}$) and perpendicular ($s_{\perp}$) to

n , and is exactly the partial-collapse map

$$s'_{\parallel} = \frac{s_{\parallel} \pm \lambda}{1 \pm \lambda s_{\parallel}}, \quad s'_{\perp} = \frac{s_{\perp} \sqrt{1 - \lambda^2}}{1 \pm \lambda s_{\parallel}}, \quad (7)$$

with outcomes $\pm$ occurring at probability $p_{\pm} = (1 \pm \lambda s_{\parallel})/2$. Something nice happens if this is written in terms of rapidity, $\eta = \text{artanh}(s_{\parallel})$ and $\mu = \text{artanh}(\lambda)$: the update becomes $\eta \rightarrow \eta + \mu$, plain rapidity addition, exactly like a boost in relativity. The shrinkage of the perpendicular component is just the matching aberration factor.

One fact worth flagging is exact for every λ , not an approximation valid only near $\lambda = 0$: the non-selective average obeys

$$\mathbb{E}[s'_{\parallel}] = s_{\parallel}, \mathbb{E}[s'_{\perp}] = s_{\perp} \sqrt{1 - \lambda^2}, \quad (8)$$

which means the non-selective Lüders channel is, exactly, a phase-damping channel with damping parameter $\gamma = \sqrt{1 - \lambda^2}$.

Proposition 1 (Exact purification of conditioned trajectories, pure initial state). *For any POVM with invertible (full-rank) Kraus operators K_i , a pure state $|\psi\rangle$ conditioned on outcome i maps to $K_i|\psi\rangle$, which remains a (normalized) pure state. Hence a trajectory conditioned on its actual sequence of outcomes converges to the boundary of the Bloch ball (a pure state) even though the unsharp POVM used at each step is not projective.*

Full-rank Kraus operators alone do *not* extend this to a mixed initial state: Kraus operators proportional to a fixed unitary are full rank, map every pure state to a pure state, and yet leave the purity of every mixed state unchanged, since they act only as a rotation of the Bloch ball. What is needed instead is an argument specific to the alternating Lüders chain considered here.

Proposition 2 (Exact purification from an arbitrary initial state). *For every $\lambda \in (0, 1)$ and any initial state ρ_0 (pure or mixed), the alternating Lüders chain of this section, with axes alternating between n_A and n_B at Bloch angle $\theta \in [0, \pi]$ (including the degenerate single-axis case $\theta = 0$), satisfies $r_n^2 \rightarrow 1$ almost surely, where $r_n = |s_n|$ is the Bloch-vector radius of the conditioned state after n steps.*

Proof. Write $g_n = 1 - r_n^2$ and let $s_{\parallel,n}$ denote the Bloch component along whichever axis is measured at step n . Squaring the exact single-step update above and averaging over the two outcomes with probabilities $p_{\pm} = (1 \pm \lambda s_{\parallel,n})/2$ gives, after simplification (verified symbolically),

$$\mathbb{E}[g_{n+1} \mid \mathcal{F}_n] = \kappa_n g_n, \kappa_n = \frac{1 - \lambda^2}{1 - \lambda^2 s_{\parallel,n}^2} \in [1 - \lambda^2, 1]. \quad (9)$$

Since $\kappa_n \leq 1$, (g_n) is a bounded nonnegative supermartingale, hence converges almost surely to some $G_\infty \geq 0$. Writing $\xi_n = g_n(1 - \kappa_n) \geq 0$ and telescoping $\mathbb{E}[g_{N+1}] = \mathbb{E}[g_0] - \sum_{n \leq N} \mathbb{E}[\xi_n]$ gives $\sum_n \mathbb{E}[\xi_n] \leq \mathbb{E}[g_0] < \infty$, hence $\sum_n \xi_n < \infty$ almost surely. Since $1 - \kappa_n \geq \lambda^2(1 - s_{\parallel,n}^2)$, this forces $\sum_n g_n(1 - s_{\parallel,n}^2) < \infty$ almost surely. On the event $\{G_\infty = L > 0\}$, eventually $g_n > L/2$, so the sum condition forces $s_{\parallel,n}^2 \rightarrow 1$ along that trajectory. But $s_{\parallel,n}^2 \leq r_n^2 = 1 - g_n \rightarrow 1 - L < 1$, a contradiction. Hence $L = 0$ almost surely. $\square$

This result is stronger than what the reduction below strictly needs: it holds exactly for the discrete chain, with no $\lambda \rightarrow 0$ limit required, and for every $\theta \in [0, \pi]$. Proposition 1 is recovered as the immediate corollary $g_0 = 0$. We verified Eq. (9) symbolically (exact algebraic cancellation) and confirmed $r_n^2 \rightarrow 1$ by direct Monte Carlo simulation of the discrete chain, for 200 independent mixed initial states at each of several θ and λ combinations, including $\theta = 0$; see the supplementary code.

Since $E_\pm = (I \pm \lambda n \cdot \sigma)/2$ has full rank whenever $\lambda < 1$, and Proposition 2 shows purification from an arbitrary initial state: the conditioned chain lives on the one-dimensional *boundary circle* of the disk spanned by n_A, n_B , not somewhere in its two-dimensional interior, regardless of whether the trajectory started pure or mixed.¹

B. Reduction to a one-dimensional diffusion

Let ψ track position on this circle, measured as an angle from the A -axis. Expand the exact Lüders update to $O(\lambda^2)$ for one measurement centered at angular position 0, and symbolically you get

$$\mathbb{E}[\Delta\psi] = -\lambda^2 \frac{\sin 2\psi}{4}, \text{Var}[\Delta\psi] = \lambda^2 \sin^2 \psi. \quad (10)$$

Add up the contributions from measurements centered alternately at $\psi = 0$ (A) and $\psi = \theta$ (B), and the drift comes out to

$$b(\psi) = -\frac{1}{2} \cos \theta \sin(2\psi - \theta), \quad (11)$$

with a diffusion coefficient (per unit λ^2 -time) that is simply $D(\psi) = 1$ at $\theta = \pi/2$, since the two contributions' $\sin^2 \psi + \cos^2 \psi$ add up to unity there, and more generally

$$D(\psi) = 1 - \cos \theta \cos(2\psi - \theta). \quad (12)$$

Equations (11)–(12) admit a natural reading as competition between deterministic phase-locking and measurement-induced diffusion. The drift $b(\psi)$ attracts ψ toward $\psi = \theta \pmod{\pi}$, mathematically analogous to noisy phase-locking models familiar from synchronization theory, while $D(\psi)$ supplies the

accompanying state-dependent diffusion. At $\theta = \pi/2$ the drift vanishes identically ($\cos \theta = 0$), so the conditioned dynamics reduce to pure uniform diffusion on the circle with no preferred phase, the unlocked regime of a noisy phase-oscillator model, consistent with the flat stationary density derived in Sec. C below. Moving θ away from $\pi/2$ restores the drift and tilts the stationary density $\pi_\theta(\psi)$ [Eq. (14)] increasingly toward $\psi = \theta$ (Fig. 3), the measurement-chain analogue of the locking transition.

It is worth being upfront about what this buys us and what it does not. The drift and diffusion above come from a formal, leading-order-in- λ expansion of the exact discrete Lüders map, followed by the usual heuristic step of replacing the discrete process with its continuum Fokker–Planck limit as $\lambda \rightarrow 0$. We have not proved a rigorous limit theorem here: there is no functional convergence result with a controlled bound on the $O(\lambda^3)$ remainder. What we do have is a symbolic derivation of the $O(\lambda^2)$ coefficients themselves, done exactly, plus numerical agreement to six significant digits between the resulting closed form and direct high-precision Monte Carlo on the original discrete chain (Table 1). So every boxed result in this section that rests on the Fokker–Planck reduction, Eq. (14), $c(\pi/2) = 1/\pi$, and the general closed form Eq. (17), should be read as proved within that formal diffusion limit, checked symbolically and numerically, rather than derived rigorously from the discrete chain itself. Table 4 labels them accordingly.

C. Exact result at $\theta = \pi/2$

At $\theta = \pi/2$, $\cos \theta$ drops out, so $b(\psi) \equiv 0$ and $D(\psi) \equiv 1$ identically: what is left is pure Brownian motion on the circle, and its stationary distribution is just uniform. That gives

$$c\left(\frac{\pi}{2}\right) = \frac{1}{2} \mathbb{E}_{\text{uniform}} [|\sin \psi|] = \frac{1}{2} \cdot \frac{2}{\pi} = \boxed{\frac{1}{\pi}}, \quad (13)$$

which follows exactly from the diffusion-limit picture above, and is not something read off a numerical fit.

D. General θ : stationary density and closed form

Lemma 2 (Zero probability current). *The drift and diffusion satisfy the exact structural identity $b(\psi) = -D'(\psi)/4$ for all θ . This forces the stationary Fokker–Planck solution to have the form $\pi(\psi) \propto [C - 2J \int \sqrt{D} d\psi] / D^{3/2}$. Since $\int_0^{2\pi} \sqrt{D} d\psi > 0$ always, periodicity of the domain forces $J = 0$.*

Hence, for all θ ,

$$\pi_\theta(\psi) \propto [1 - \cos \theta \cos(2\psi - \theta)]^{-3/2}. \quad (14)$$

Figure 3 shows this density over the full (θ, ψ) domain at once, rather than at a single angle. The flat red slice at $\theta = \pi/2$ is the uniform distribution of the previous subsection; move away from that angle and the density tilts increasingly toward $\psi = \theta$, which is exactly the skew that the closed form Eq. (17) below has to account for.

Stationary density $\pi_\theta(\psi) \propto [1 - \cos \theta \cos(2\psi - \theta)]^{-3/2}$




Figure 3. The stationary density $\pi_\theta(\psi) \propto [1 - \cos \theta \cos(2\psi - \theta)]^{-3/2}$ [Eq. (14)], over the full (θ, ψ) domain. The red curve marks $\theta = \pi/2$, where Lemma 2 together with $b \equiv 0$ forces the density to be exactly uniform, $1/2\pi$. Moving away from $\theta = \pi/2$ in either direction skews the density toward $\psi = \theta$.

Normalization. With $k = \cos \theta$, $m = 2k/(1 + k)$,

$$Z(\theta) = \int_0^{2\pi} \pi_\theta(\psi) d\psi = \frac{4E(m)}{(1 - k)\sqrt{1 + k}}, \quad (15)$$

with E the complete elliptic integral of the second kind, derived by hand, checked symbolically, and matching direct numerical quadrature to six digits.

Numerator. Define $N(\theta) = \int_0^{2\pi} \pi_\theta(\psi) |\sin \psi| d\psi$. Using the π -periodicity of both π_θ and $|\sin \psi|$, substituting $u = 2\psi - \theta$, and splitting $\sin\left(\frac{u+\theta}{2}\right)$ into two elementary pieces, each piece integrates via an

ordinary $(A + Bs^2)^{-3/2}$ antiderivative. No elliptic integral of the third kind is needed, which is not what we expected going in. The result is the plain closed form

$$N(\theta) = \frac{4}{\sin \theta}. \quad (16)$$

Final result. Combine $Z(\theta)$ and $N(\theta)$, simplify with half-angle identities, and

$$c(\theta) = \frac{\sin(\theta/2)}{\sqrt{2}E(m)}, m = 1 - \tan^2(\theta/2). \quad (17)$$

Corollary 1. $c(\theta) = c(\pi - \theta)$, which follows from Eq. (17) via the imaginary-modulus (reciprocal-modulus) transformation of the complete elliptic integral of the second kind, and serves as an independent consistency check.

Proof. Write $m = 1 - \tan^2(\theta/2)$ and $m' = 1 - \tan^2\left(\frac{\pi-\theta}{2}\right)$. Because $\tan\left(\frac{\pi-\theta}{2}\right) = \cot(\theta/2) = 1/\tan(\theta/2)$, this gives $m' = -m/(1-m)$, the standard reciprocal-modulus substitution. The imaginary-modulus transformation for E ^[18] gives $E(m') = E(m)/\sqrt{1-m}$, which we also checked directly to 30 significant digits. Since $1-m = \tan^2(\theta/2)$, we have $\sqrt{1-m} = \tan(\theta/2)$ for $\theta \in (0, \pi)$, and

$$c(\pi - \theta) = \frac{\cos(\theta/2)}{\sqrt{2}E(m')} = \frac{\cos(\theta/2)\tan(\theta/2)}{\sqrt{2}E(m)} = \frac{\sin(\theta/2)}{\sqrt{2}E(m)} = c(\theta). \blacksquare$$

Small-angle limit. As $\theta \rightarrow 0$, Eq. (17) gives $c(\theta) \rightarrow \theta/(2\sqrt{2})$, slope $1/(2\sqrt{2}) \approx 0.35355$. That is not the same slope as $r(\theta)$ or $\nu(\theta)$, both $1/2$, so the unsharp, weak-measurement limit does not simply reduce to the sharp-case geometry as $\lambda \rightarrow 0$. Figure 4 puts the two slopes side by side.




Figure 4. Small-angle behavior of $r(\theta)$, $\nu(\theta)$, and $c(\theta)$, with their leading-order linear asymptotes (dotted). $r(\theta)$ and $\nu(\theta)$ share slope $1/2$, while $c(\theta)$ rises with the shallower slope $1/(2\sqrt{2})$, showing that the unsharp weak-measurement limit is not a simple continuation of the sharp-case geometry.

Table 1 sets Eq. (17) next to high-precision Monte Carlo estimates from direct simulation of the Lüders chain at small λ , with burn-in scaled as $O(1/\lambda^2)$, the correct relaxation-time scaling for this chain.²

θ	$c(\theta)$ exact	Monte Carlo
10°	0.06085	0.0620
30°	0.16934	0.1714
50°	0.25110	0.2524
70°	0.30138	0.3015
90°	$0.31831 (= 1/\pi)$	0.3186

Table 1. Exact closed form, Eq. (17), versus high-precision Monte Carlo (burn-in $\sim 6/\lambda^2$ steps).

Figure 5 puts the exact curve of Eq. (17) on the same plot as the Monte Carlo points from Table 1.




Figure 5. The exact closed form $c(\theta) = \sin(\theta/2)/[\sqrt{2}E(m)]$, $m = 1 - \tan^2(\theta/2)$ [Eq. (17)], compared with high-precision Monte Carlo estimates (Table 1). The dotted line marks the exact value $c(\pi/2) = 1/\pi$.

V. Noise sensitivity and the $\theta = \pi/2$ blind spot

A. A model-independent identity

Put any single-qubit noise channel between measurements in the alternating sharp chain (it need not be isotropic) and the per-step trace distance still obeys

$$r(\theta) = D_{\text{far}}(\theta) + p_{\text{near}} [D_{\text{near}}(\theta) - D_{\text{far}}(\theta)], \quad (18)$$

exactly, where $D_{\text{near}} = \sin(\theta/2)$ and $D_{\text{far}} = \cos(\theta/2)$ are purely geometric and carry no noise-dependence at all; everything about the noise sits in p_{near} .

Proposition 3 (Exact blind spot at $\theta = \pi/2$). At $\theta = \pi/2$, $\sin(\pi/4) = \cos(\pi/4)$ exactly, so the bracket in Eq. (18) vanishes identically, and $r(\pi/2) = \sqrt{2}/2$ regardless of what p_{near} is, that is, for any noise channel

whatsoever. The argument only needs the coincidence $D_{\text{near}} = D_{\text{far}}$, nothing about how p_{near} arises.

We checked this against two physically different channels. For depolarizing noise with parameter p applied after each collapse, $p_{\text{near}}(\theta, p) = \frac{1}{2}[1 + (1 - p)\cos\theta]$, which gives $r(\theta, p) = p_{\text{near}}\sin(\theta/2) + (1 - p_{\text{near}})\cos(\theta/2)$ and a Fisher information for estimating p from near/far statistics of

$$I(p = 0, \theta) = \cot^2 \theta. \quad (19)$$

This blows up as $\theta \rightarrow 0, \pi$, very sensitive, but in the least interesting, near-compatible regime, and hits exactly zero at $\theta = \pi/2$ (Table 2). For amplitude damping with parameter γ , a strongly anisotropic channel with a fixed point at the north pole and the natural choice for superconducting qubits, building and diagonalizing the resulting four-state ($A^\pm, B^\pm$) Markov chain gives $dr/d\gamma$ curves that track the depolarizing case almost exactly, again hitting zero at $\theta = \pi/2$ to machine precision.

θ	depolarizing, dr/dp	amplitude damping, $dr/d\gamma$
5°	0.4759	0.4637
45°	0.1913	0.1913
90°	0	0
150°	0.3062	0.3060

Table 2. Sensitivity $dr/d(\text{noise})$ at noise = 0, for two physically distinct channels.

Figure 6 shows the exact depolarizing curve $dr/dp = \frac{1}{2}\cos\theta[\cos(\theta/2) - \sin(\theta/2)]$ alongside the amplitude-damping data from Table 2; both hit zero at $\theta = \pi/2$.




Figure 6. Noise sensitivity $dr/d(\text{noise})$ at zero noise, for a depolarizing channel (exact curve) and an amplitude-damping channel (points, Table 2). Both vanish exactly at $\theta = \pi/2$, the model-independent blind spot of Proposition 3.

B. Practical implication

Equation (18) makes clear that this blind spot is a plain trigonometric coincidence at $\theta = \pi/2$, not a feature of one particular noise model, and it holds for any single-qubit channel that can be written through p_{near} , whatever its microscopic origin. In practice, that rules out $r(\theta)$ as a calibration check at exactly the angle experiments tend to favor, since $\theta = \pi/2$ is usually called “maximally incompatible.” Away from that angle, though, $r(\theta)$ and $c(\theta, \lambda)$ can still be read off directly from raw sequential-measurement statistics: a few hundred shots at one fixed pair of bases (Sec. IV) is enough to catch a percent-level departure from the ideal curve.

It is worth saying clearly what this diagnostic is not. It is not a replacement for standard randomized benchmarking (RB) [19][20][21], and we checked that it is not a more shot-efficient alternative to it either. RB piles up a small per-gate error over m Clifford gates before a single readout, amplifying the signal by a factor that grows with m . Per elementary operation, a well-chosen RB sequence length gives at least an order of magnitude more Fisher information about a depolarizing rate p than one step of the alternating-

basis chain does, across the whole realistic range $p \lesssim 0.1$. These two protocols are answering different questions and should not be conflated: RB estimates a gate-averaged error rate under a randomized (twirled) noise model, while $r(\theta)$ and $c(\theta)$ describe the steady-state statistics of one specific, un-randomized, repeated Born-rule process at a fixed angle θ . The rigorous group-theoretic underpinnings of RB, and just how far its exponential-decay picture can be pushed, are laid out in the general framework of Helsen *et al.* [21].

Where $r(\theta)$ does hold a real, narrower advantage is in settings where RB's prerequisites are simply not available. RB needs a calibrated universal (or at least Clifford-generating) gate set, and it relies on twirling to make gate-dependent errors look approximately depolarizing [19][20]. $r(\theta)$ needs none of that: no gates beyond the two measurement bases already being alternated, and no twirling assumption. This matters for the continuous- and sequential-weak-measurement experiments motivating this work in the first place (e.g. Refs. [8][15]), where an alternating-basis chain may already be exactly what is being studied, and bolting on a random-gate RB sequence would be an extra, non-native step rather than something that falls out of the experiment for free. So the value here is niche, not general-purpose: a low-overhead residual check obtained for free from a protocol already running, not a competitor to RB for gate calibration in general.

There is a further limitation we would rather state outright than bury: an identifiability gap. Equation (18) shows noise entering only through p_{near} , but that is exactly the same channel through which a systematic miscalibration of the intended angle ($\theta \rightarrow \theta + \delta$) would enter, since $p_{\text{near}} = \cos^2(\theta/2)$ in the noiseless case. So a single measured value of $r(\theta)$ cannot on its own tell decoherence apart from a small angular miscalibration; separating the two needs either an independent calibration of θ from gate-level characterization, or measurements at several θ values, since the two effects diverge in their θ -dependence away from $\theta = \pi/2$. Nor does a single value of $r(\theta)$ identify the noise mechanism: Table 2 shows depolarizing and amplitude-damping channels producing $dr/d(\text{noise})$ curves that are nearly impossible to tell apart at this precision. So the diagnostic can tell you *that* noise is present, and roughly *how much*, but not *what kind*. These are open limitations of the single-angle protocol, not solved problems. At the time of writing, the only piece of this section checked against a real device is the blind-spot prediction of Proposition 3 itself (Sec. C); the Fisher-information curves, the RB throughput comparison, and the general noise-sensitivity results of Table 2 remain untested on hardware. See Table 4.

C. Preliminary validation on superconducting hardware

Proposition 3 needs only the coincidence $D_{\text{near}} = D_{\text{far}}$ at $\theta = \pi/2$; it does not assume a noise model, so it should hold on a real device regardless of what that device's actual noise looks like. We tested this on the IBM Quantum superconducting processor `ibm_fez` (a Heron r2 device), using a single physical qubit. Each alternating-basis step was realized in one continuous circuit rather than as separate reset-and-reprepare shots: basis- B steps apply $R_y(-\theta)$, mid-circuit-measure, then $R_y(+\theta)$ to return to the computational frame, so a single n_{steps} -measurement circuit executed for a given shot count yields one full bitstring per shot.³ Five angles ($\theta = 30^\circ, 45^\circ, 60^\circ, 75^\circ, 90^\circ$) were run at $n_{\text{steps}} = 20$, burn-in = 5, across shot counts from 250 to 8000, for 40 jobs, 1.3825×10^6 pooled transitions in total, and three independent repeats of the $\theta = 90^\circ$ baseline submitted as separate jobs (different calibration cycles) to check run-to-run stability.

θ	$p_{\text{near}}^{\text{theory}}$	$p_{\text{near}}^{\text{hw}}$	$\Delta p_{\text{near}} (\%)$	z	r^{theory}	r^{hw}	$\Delta r (\%)$
30°	0.9330	0.9275	0.586	-7.06	0.3062	0.3101	1.263
45°	0.8536	0.8488	0.551	-4.40	0.4619	0.4645	0.551
60°	0.7500	0.7469	0.407	-2.35	0.5915	0.5926	0.189
75°	0.6294	0.6312	0.277	+1.21	0.6772	0.6768	0.047
90°	0.5000	0.4995	0.091	-0.30	0.7071	0.7071	$\sim 10^{-13}$

Table 3. Theory versus `ibm_fez` hardware, single qubit, shots= 8000 (the largest-shot run, hence smallest shot noise). Δ columns are relative error against the noiseless closed form, in percent; z is the p_{near} deviation in units of its predicted shot-noise standard error, $z = (p_{\text{near}}^{\text{hw}} - p_{\text{near}}^{\text{theory}})/\sigma_{p_{\text{near}}}$.




Figure 7. p_{near} on ibm_fez (shots= 8000, error bars are the predicted shot-noise σ) against the noiseless closed form $p_{\text{near}}(\theta) = \cos^2(\theta/2)$. The hardware point sits visibly below theory at small θ and converges onto the curve near $\theta = \pi/2$.

The last two columns of Table 3 reproduce the point made in Sec. A: r^{hw} tracks r^{theory} to machine precision at $\theta = 90^\circ$ while missing it by over a percent at $\theta = 30^\circ$, exactly the blind-spot signature. Figure 7 shows this directly: the raw $p_{\text{near}}^{\text{hw}}$ points sit visibly below the noiseless curve at small θ and converge onto it as $\theta \rightarrow \pi/2$. But the more informative column is z , not Δ . A relative error of 0.586% at $\theta = 30^\circ$ sounds unremarkable until it is measured against the 0.077% shot-noise floor at shots= 8000: that is a 7σ discrepancy, not noise. Figure 8 makes this explicit by tracking z against shots at each angle. Two regimes are visible. At $\theta = 30^\circ, 45^\circ, 60^\circ$, z drifts steadily out of the $|z| < 2$ band as shots increase, from ~ -1 at 250 shots to $-7.06, -4.40, -2.35$ respectively at 8000 shots. That growth with shots is the diagnostic signature of a genuine, shot-count-independent *bias*: the absolute deviation $p_{\text{near}}^{\text{theory}} - p_{\text{near}}^{\text{hw}}$ stays essentially flat (within ~ 0.001) across the 32-fold increase in shots, while its statistical uncertainty shrinks as $1/\sqrt{\text{shots}}$, so the same fixed bias becomes more and more significant. Pure shot noise cannot produce that pattern; it would keep z scattered around zero at every shot count.

At $\theta = 75^\circ, 90^\circ$, by contrast, z stays inside or just outside $|z| < 2$ with no shots-dependent trend, consistent with statistical fluctuation rather than bias.




Figure 8. Significance of the p_{near} deviation from theory, $z = (p_{\text{near}}^{\text{hw}} - p_{\text{near}}^{\text{theory}}) / \sigma_{p_{\text{near}}}$, versus shots, for each angle. Growth of $|z|$ with shots at $\theta = 30^\circ, 45^\circ, 60^\circ$ signals a fixed bias rather than shot noise; the flat, near-zero trend at $\theta = 75^\circ, 90^\circ$ is consistent with noise alone.

This is exactly the kind of departure Sec. A's depolarizing model was built to describe, and it is worth checking whether that specific model, rather than noise in the abstract, actually fits. The noiseless-versus-depolarizing formula already given there, $p_{\text{near}}(\theta, p) = \frac{1}{2}[1 + (1 - p)\cos\theta]$, predicts a bias $p_{\text{near}}^{\text{theory}} - p_{\text{near}}(\theta, p) = \frac{1}{2}p\cos\theta$: proportional to $\cos\theta$, hence largest at small θ and vanishing at $\theta = \pi/2$, the opposite θ -dependence from what a systematic angular miscalibration $\theta \rightarrow \theta + \delta$ would produce (bias $\propto \sin\theta$, largest at $\theta = \pi/2$, vanishing at $\theta \rightarrow 0$). The observed pattern, largest deviation at $\theta = 30^\circ$ and statistically indistinguishable from zero at $\theta = 90^\circ$, favors depolarizing-type noise over angular miscalibration, which is itself a useful, if informal, resolution of the identifiability concern raised in Sec. B for this particular dataset. Fitting a single effective depolarizing rate p by weighted least squares to the four informative angles ($\theta = 90^\circ$ has zero predictor weight, $\cos 90^\circ = 0$) at shots= 8000 gives per-

angle implied values 0.0126, 0.0133, 0.0122 at $\theta = 30^\circ, 45^\circ, 60^\circ$ respectively, with a fourth-angle outlier at $\theta = 75^\circ$ (-0.0135 , a 2.3σ residual from the joint fit, likely a statistical fluctuation given its small predictor weight $\cos 75^\circ \approx 0.26$). Pooling all 40 jobs (all shot levels and repeats, inverse-variance weighted) gives $p_{\text{eff}} = 0.0131$ with $\chi^2/\text{dof} = 1.17$, an acceptable fit. Repeating the fit on each of the three independent $\theta = 90^\circ$ -baseline job submissions separately gives $p_{\text{eff}} = 0.0151, 0.0105, 0.0156$, consistent with each other within the expected shot-noise scatter and with no sign of the large job-to-job drift that would indicate an unstable calibration.

We stress what this fit is and is not. It is a single-parameter, single-model consistency check using the data already in hand, not an independent noise characterization: we did not run randomized benchmarking or gate-set tomography on the same qubit on the same day to cross-validate p_{eff} against a ground truth, so the numerical agreement with a depolarizing model is evidence in that model's favor, not proof that the physical error is literally depolarizing rather than some other channel with similar $\cos \theta$ -scaling (amplitude damping, for instance, was shown in Sec. A to track the depolarizing sensitivity curve closely). What the fit does establish, on data internal to this dataset alone, is that (i) the deviation from the noiseless curve at $\theta < 90^\circ$ is a real, reproducible bias and not shot noise, (ii) its θ -dependence is far better explained by a noise channel than by an angle-calibration error, and (iii) a single free parameter accounts for it across three angles and 40 independently submitted jobs to within statistical error. This is a single-device, single-qubit, single-noise-environment check, with no independent error mitigation and no cross-platform replication, and it does not touch the RB comparison, the Fisher-information curves, or the amplitude-damping results of Sec. A on their own terms, none of which have been tested on hardware. Within that narrow scope, though, the qualitative signature Proposition 3 predicts, real, statistically significant hardware noise fully visible in p_{near} away from $\theta = \pi/2$ but exactly invisible to r at $\theta = \pi/2$, is exactly what the device shows. One further scope limitation is worth stating plainly: the three independent repeats across separate calibration cycles (Sec. C) were run only at $\theta = 90^\circ$, the angle predicted *not* to show bias. The angles that do show the bias, $30^\circ, 45^\circ, 60^\circ$, were each submitted once, so run-to-run stability of the biased p_{near} values across separate calibration cycles has not yet been checked; ruling out a single-calibration-cycle artifact at those angles is the most direct follow-up.

VI. Discussion

A. Relation to prior work

A literature search, systematic though we would not call it exhaustive, turned up Wei and Nazarov ^[12], Ruskov *et al.* ^{[13][14]}, Hacothen-Gourgy *et al.* ^[15], Atalaya *et al.* ^[16], and, closest of all, Walls and Ford ^[17], whose quantum-state-diffusion treatment of alternating S_z, S_x measurements shares the purification phenomenology at the heart of Section IV. The experimental groundwork for this line of work, direct observation of a single qubit's conditioned trajectory under continuous weak measurement, goes back to Murch *et al.* ^[8], reviewed more broadly by Weber *et al.* ^[9]. But in every one of these papers, the angle between measurement axes is fixed, almost always exact orthogonality, and what gets varied is measurement strength, the number of simultaneous axes, or detector efficiency, never θ itself. Ruskov *et al.* ^{[13][14]} specifically vary the number and arrangement of *simultaneous* complementary axes (two or three, always orthogonal to each other) and the resulting purification speed-up, which is a different question from the relative angle between two *alternating* axes.

Closer in structure, and more recent, is the path-integral treatment of monitored-qubit purification by Martins and Lima ^[22], which appeared while this manuscript was already being written. Their model monitors one fixed axis continuously through a collisional (ancilla) coupling, with coupling strength per collision as the free parameter, so there is no second axis, and hence no θ , in their formulation at all. Their exact stationary-purity distribution is consequently a different object from $c(\theta)$ here: theirs solves a one-axis continuous-monitoring problem, ours a two-axis alternating one. We think of the two as complementary rather than overlapping.

Lema *et al.* ^[23] ask a related but different question: how much information about the initial qubit configuration survives in a sequential weak-measurement record when the scheme is informationally incomplete, or when the qubit's own dynamics scrambles the state. Their quantity is a mutual-information bound on retrodicting the initial state from an imperfect readout, not the steady-state trace-distance rate between consecutive post-measurement states as a function of θ studied here. Again, the two are complementary: theirs bounds what an imperfect record can tell you about where you started, while $r(\theta)$ and $c(\theta)$ describe the steady-state statistics of the process itself.

As far as we can tell, nobody has previously derived $r(\theta)$ or $c(\theta)$ as explicit closed-form functions of the incompatibility angle. That is the specific, narrow claim to originality being made here: closed-form

dependence on the angle itself, bridged to the standard robustness measure $\nu(\theta)$, not a claim that purification under repeated incompatible measurement is itself a new phenomenon. It plainly is not; it is well established, including in the many-body measurement-induced-phase-transition setting [10] [11] mentioned in the Introduction, and in the broader incompatibility literature reviewed in Refs. [1][2].

B. Qudit generalization: an open problem

The purification argument in Proposition 2 does not care about the specific two-axis qubit structure of its proof in an obvious way, but its numerical signature does extend: for a d -outcome unsharp POVM $E_i = \frac{1-\lambda}{d}I + \lambda|i\rangle\langle i|$, every Kraus operator is invertible for $\lambda < 1$, so Proposition 1 already guarantees purification of *pure* conditioned trajectories no matter what d is. We checked purification from *mixed* initial states numerically for a qutrit ($d = 3$) alternating between the computational and discrete Fourier bases, and purity $\text{Tr}(\rho^2)$ converges cleanly to 1 within a few hundred steps, consistent with a qudit generalization of Proposition 2 — but we have not proved such a generalization; the supermartingale argument used for the qubit case leans on the explicit parallel/perpendicular Bloch-vector decomposition, which does not carry over directly to $\mathbb{CP}^{d-1}$. But the qubit closed form leans on something dimension-specific: axial symmetry around the measurement axis collapsing the pure-state space to a one-dimensional circle. For $d > 2$, the pure-state space is $\mathbb{CP}^{d-1}$, real dimension $2(d-1) \geq 4$, and there is no single axial symmetry that brings it down to one parameter, except along special sub-families that just reduce back to the qubit problem. So the qudit closed form is a genuinely different open problem, a random walk on a Kähler manifold rather than a circle, and not something that falls out of the results above by simple algebra.

C. Status summary

Table 4 pulls together the epistemic status of every claim made in this paper, using the same labels applied throughout the derivation above.

Claim	Status
$r(\theta) = \frac{1}{2} \sin \theta \sqrt{1 + \sin \theta}$ (sharp)	Proved
r, ν agree to first order, differ globally	Proved
Exact purification of conditioned trajectories, pure initial state	Proved
Exact purification from arbitrary (pure or mixed) initial state	Proved
$\pi_\theta(\psi) \propto [1 - \cos \theta \cos(2\psi - \theta)]^{-3/2}$	Proved within $\lambda \rightarrow 0$ diffusion limit; symbolic + numerical (6 digits)
$c(\pi/2) = 1/\pi$	Proved within $\lambda \rightarrow 0$ diffusion limit
$c(\theta) = \sin(\theta/2)/[\sqrt{2}E(m)]$, general θ	Proved within $\lambda \rightarrow 0$ diffusion limit; symbolic + numerical (6 digits)
Blind spot at $\theta = \pi/2$, arbitrary noise channel	Proved
Blind spot at $\theta = \pi/2$ on real hardware (ibm_fez, single qubit, Sec. C)	Preliminary hardware check; single device, not independently replicated across platforms
Originality relative to cited literature	Best-effort, non-exhaustive search
Qudit closed form for $c_d(\theta)$	Open
Practical diagnostic utility (niche, low-overhead residual check; not a substitute for RB)	Untested, hypothesis

Table 4. Status of claims.

VII. Conclusion

We set out to find the steady-state incompatibility rate of an alternating sequential qubit measurement chain, and we now have exact closed forms for both the sharp case, $r(\theta)$, and the unsharp case, $c(\theta)$. The latter arrives as a complete elliptic integral, once an initially wrong two-dimensional diffusion picture was corrected by the exact purification principle of Proposition 2, which holds for arbitrary (pure or mixed) initial states. Both quantities are operational and directly observable from a single trajectory.

They agree with the standard static noise-robustness measure to first order near the compatible point, part ways from it everywhere else, and carry an exact, model-independent blind spot to noise at $\theta = \pi/2$. The qubit problem is closed analytically at this point. The qudit generalization is guaranteed, structurally, to show the same purification phenomenon, but working it out needs new tools for random walks on $\mathbb{CP}^{d-1}$, and we leave that open.

Appendix A. Methodological history: rejected intermediate hypotheses

Two missteps got corrected along the way to the results above; we record both here, for transparency, and both are cross-referenced from footnotes in Sec. IV.

- i. *The two-dimensional diffusion picture.* The first pass at the unsharp POVM chain modeled the conditioned dynamics as diffusion inside the two-dimensional Bloch-sphere disk spanned by n_A, n_B . That is wrong: Proposition 2 shows trajectories conditioned on their actual outcomes purify exactly, from any initial state, so the real state space is the one-dimensional boundary circle, not the disk's interior. Sec. B gives the corrected reduction.
- ii. *The rejected conjecture $c(\theta) = \sin \theta/\pi$.* Before Eq. (17) was in hand, Monte Carlo runs at small λ with too little burn-in (the chain's relaxation time scales as $O(1/\lambda^2)$, and the early runs did not sample long enough to reflect that) produced statistics that looked consistent with $c(\theta) = \sin \theta/\pi$. Once the burn-in was corrected to the right $O(1/\lambda^2)$ scaling, that conjecture turned out to be wrong by as much as 12% near $\theta = 10^\circ$, so it is rejected in favor of Eq. (17), which matches high-precision quadrature to at least six significant digits at every angle tested.

Appendix B. Symbolic verification notes

Every boxed closed-form result in Secs. III and IV was checked two separate ways. First, direct symbolic expansion and simplification (sympy) of the exact Lüders update maps, confirming both the drift/diffusion coefficients and the martingale identity $\mathbb{E}[\Delta s_{\parallel}] = 0$. Second, high-precision numerical quadrature of the stationary density Eq. (14) against Eq. (17), agreeing to at least six significant digits across $\theta \in \{5^\circ, 10^\circ, \dots, 170^\circ\}$. The zero-current argument in Lemma 2 was checked by substituting $\pi_\theta \propto D^{-3/2}$ directly into the stationary Fokker–Planck current; the residual current came out consistent with zero at machine precision ($\lesssim 10^{-11}$) at every angle tested.

Statements and Declarations

Funding

No specific funding was received for this work.

Potential Competing Interests

No potential competing interests to declare.

Author Contributions

Agus Mulia Bakti: Conceptualization; Formal analysis; Investigation; Methodology; Software; Validation; Writing, original draft; Writing, review & editing.

Data Availability

The symbolic derivation notebook (sympy) and Monte Carlo simulation code that support the findings of this study are openly available under the MIT License in a public repository at <https://doi.org/10.5281/zenodo.21219550>. The Qiskit code used to construct the alternating-basis hardware circuits, and the raw measurement data (bitstrings/counts) from all 40 IBM Quantum `ibmq_16melbourne` jobs underlying Table 3 and Figs. 7–8, are available at <https://doi.org/10.5281/zenodo.22657072>.

Acknowledgements

The author thanks colleagues at the Department of Physics Education, Universitas Pendidikan Mandalika, for discussion, and acknowledges the use of symbolic computation (sympy) for verification of the algebraic derivations reported here.

Footnotes

¹ An earlier pass at this treated the conditioned dynamics as a genuine two-dimensional diffusion inside the disk. Proposition 2 rules that out for arbitrary initial states; Appendix A gives a short account of the mistake, kept here for transparency.

² An earlier run with insufficient burn-in looked consistent with a simpler conjecture, $c(\theta) = \sin \theta / \pi$. Fixing the burn-in scaling showed that conjecture off by up to 12% near $\theta = 10^\circ$, so it is rejected. See

Appendix A.

³ Transitions from the first burn_in steps of each shot are discarded before pooling the $(i, i + 1)$ near/far outcomes into the p_{near} estimator of Eq. (18), to absorb any hardware-specific transient before the steady-state chain of Sec. II is what gets measured.

References

1. ^a, ^bGühne O, Haapasalo E, Kraft T, Pellonpää J, Uola R (2023). "Colloquium: Incompatible Measurements in Quantum Information Science." *Rev Mod Phys.* **95**:011003. doi:[10.1103/RevModPhys.95.011003](https://doi.org/10.1103/RevModPhys.95.011003).
2. ^a, ^bUola R, Costa A, Nguyen H, Gühne O (2020). "Quantum Steering." *Rev Mod Phys.* **92**:015001. doi:[10.1103/RevModPhys.92.015001](https://doi.org/10.1103/RevModPhys.92.015001).
3. ^ΔBusch P (1986). "Unsharp Reality and Joint Measurements for Spin Observables." *Phys Rev D.* **33**:2253–2261. doi:[10.1103/PhysRevD.33.2253](https://doi.org/10.1103/PhysRevD.33.2253).
4. ^ΔHeinosaari T, Kiukas J, Reitzner D (2015). "Noise Robustness of the Incompatibility of Quantum Measurements." *Phys Rev A.* **92**:022115. doi:[10.1103/PhysRevA.92.022115](https://doi.org/10.1103/PhysRevA.92.022115).
5. ^ΔDesignolle S, Farkas M, Kaniewski J (2019). "Incompatibility Robustness of Quantum Measurements: A Unified Framework." *New J Phys.* **21**:113053. doi:[10.1088/1367-2630/ab5020](https://doi.org/10.1088/1367-2630/ab5020).
6. ^ΔWiseman H, Milburn G (2010). *Quantum Measurement and Control*. Cambridge University Press.
7. ^ΔChantasri A, Dressel J, Jordan A (2013). "Action Principle for Continuous Quantum Measurement." *Phys Rev A.* **88**:042110. doi:[10.1103/PhysRevA.88.042110](https://doi.org/10.1103/PhysRevA.88.042110).
8. ^a, ^b, ^ΔMurch K, Weber S, Macklin C, Siddiqi I (2013). "Observing Single Quantum Trajectories of a Superconducting Qubit." *Nature.* **502**:211–214. doi:[10.1038/nature12539](https://doi.org/10.1038/nature12539).
9. ^a, ^bWeber S, Murch K, Schwartz M, Roch N, Siddiqi I (2016). "Quantum Trajectories of Superconducting Qubits." *C R Phys.* **17**:703–719. doi:[10.1016/j.crhy.2016.07.004](https://doi.org/10.1016/j.crhy.2016.07.004).
10. ^a, ^bGullans M, Huse D (2020). "Dynamical Purification Phase Transition Induced by Quantum Measurements." *Phys Rev X.* **10**:041020. doi:[10.1103/PhysRevX.10.041020](https://doi.org/10.1103/PhysRevX.10.041020).
11. ^a, ^bFisher M, Khemani V, Nahum A, Vijay S (2023). "Random Quantum Circuits." *Annu Rev Condens Matter Phys.* **14**:335–379. doi:[10.1146/annurev-conmatphys-031720-030658](https://doi.org/10.1146/annurev-conmatphys-031720-030658).
12. ^a, ^bWei H, Nazarov Y (2008). "Statistics of Measurement of Noncommuting Quantum Variables: Monitoring and Purification of a Qubit." *Phys Rev B.* **78**:045308. doi:[10.1103/PhysRevB.78.045308](https://doi.org/10.1103/PhysRevB.78.045308).

13. ^{a, b, c}Ruskov R, Korotkov A, Mølmer K (2010). "Qubit State Monitoring by Measurement of Three Complementary Observables." *Phys Rev Lett.* **105**:100506. doi:[10.1103/PhysRevLett.105.100506](https://doi.org/10.1103/PhysRevLett.105.100506).
14. ^{a, b, c}Ruskov R, Combes J, Mølmer K, Wiseman H (2012). "Qubit Purification Speed-Up for Three Complementary Continuous Measurements." *Phil Trans R Soc A.* **370**:5291–5307. doi:[10.1098/rsta.2011.0516](https://doi.org/10.1098/rsta.2011.0516).
15. ^{a, b, c}Hacohen-Gourgy S, Martin L, Flurin E, Ramasesh V, Whaley K, Siddiqi I (2016). "Quantum Dynamics of Simultaneously Measured Non-Commuting Observables." *Nature.* **538**:491–494. doi:[10.1038/nature19762](https://doi.org/10.1038/nature19762).
16. ^{a, b}Atalaya J, Hacohen-Gourgy S, Martin L, Siddiqi I, Korotkov A (2018). "Correlators in Simultaneous Measurement of Non-Commuting Qubit Observables." *npj Quantum Inf.* **4**:41. doi:[10.1038/s41534-018-0091-1](https://doi.org/10.1038/s41534-018-0091-1).
17. ^{a, b}Walls S, Ford I (2024). "Memory Effects in a Sequence of Measurements of Non-Commuting Observables." *Phys Rev A.* **110**:032432. doi:[10.1103/PhysRevA.110.032432](https://doi.org/10.1103/PhysRevA.110.032432).
18. ^aOlver F, Olde Daalhuis A, Lozier D, Schneider B, Boisvert R, Clark C, Miller B, Saunders B, Cohl H, McClain M (2024). "NIST Digital Library of Mathematical Functions." NIST Digital Library of Mathematical Functions. §19.7, Release 1.2.0 of 2024–03–15.
19. ^{a, b}Magesan E, Gambetta J, Emerson J (2011). "Scalable and Robust Randomized Benchmarking of Quantum Processes." *Phys Rev Lett.* **106**:180504. doi:[10.1103/PhysRevLett.106.180504](https://doi.org/10.1103/PhysRevLett.106.180504).
20. ^{a, b}Wallman J, Flammia S (2014). "Randomized Benchmarking with Confidence." *New J Phys.* **16**:103032. doi:[10.1088/1367-2630/16/10/103032](https://doi.org/10.1088/1367-2630/16/10/103032).
21. ^{a, b}Helsen J, Roth I, Onorati E, Werner A, Eisert J (2022). "General Framework for Randomized Benchmarking." *PRX Quantum.* **3**:020357. doi:[10.1103/PRXQuantum.3.020357](https://doi.org/10.1103/PRXQuantum.3.020357).
22. ^aPoltronieri Martins M, Santos Lima H (2026). "Purification of a Monitored Qubit: Exact Path-Integral Solution."
23. ^aLema C, Bou-Comas A, Acharya A, Oganessian V, Sengupta A (2025). "Reading Qubits with Sequential Weak Measurements: Limits of Information Extraction."

Supplementary data: available at <https://doi.org/10.32388/YU8RS0.2>

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.