Supplemental Information:
Error correction advances the point of no return

Hikaru Wakaura[1]

[1] *QIRI (Quantum Integrated Research Institute Inc.), Minato-ku, Tokyo 107-0061, Japan*

(Dated: June 12, 2026)

This document gives the derivations, exact procedures, parameters, random seeds, and reproduction commands behind the main text. Every dataset and figure is indexed in the companion data document; all code is in an open-source software package.

S1. DERIVATION OF THE MOTIONAL-NARROWING SUPPRESSION FACTOR

Motional narrowing follows from second-order (Redfield) relaxation theory [1, 2]. Let the radical pair experience a fluctuating hyperfine field of amplitude ω_0 (in angular-frequency units) whose direction is randomized by protein tumbling on a correlation time τ_c, with autocorrelation $\langle b(t)b(0)\rangle = \omega_0^2 e^{-|t|/\tau_c}$. The Redfield dephasing rate is the zero-frequency power spectral density,

$$\frac{1}{T_2^{\mathrm{MN}}} = \omega_0^2 \int_0^\infty e^{-t/\tau_c}\, dt = \omega_0^2\, \tau_c. \tag{S1}$$

In the fast-motion limit $\omega_0\tau_c \ll 1$ this rate is suppressed relative to the static dephasing ω_0 by the factor $\omega_0\tau_c$. Discretizing the continuous averaging as N effective refocusing events per hyperfine cycle (the formal equivalence of motional narrowing to N-pulse dynamical decoupling [3] when $\omega_0\tau_c \simeq \pi/2N$) yields the closed form used in the main text,

$$\eta(\tau_c) = \frac{1}{1 + N^2(\tau_c\omega_c)^2}, \qquad \omega_c = 1/\tau_c, \tag{S2}$$

so $\tau_c\omega_c = 1$ and, for $N = 4$, $\eta = 1/(1+16) = 1/17 \approx 0.059$. With $\gamma_{\mathrm{eff}}^{\mathrm{CRY}} = 3.25$ this gives $\gamma_{\mathrm{DD}} = 0.191$; with $\gamma_{\mathrm{eff}}^{\text{MAO-A}} = 4.55$, $\gamma_{\mathrm{DD}} = 0.268$. The working-point hyperfine amplitude $A_\perp \simeq 4\,\mathrm{MHz}$ and $\tau_c \simeq 10\,\mathrm{ns}$ give $(\omega_0\tau_c)^2 \approx 1/16$, consistent with the code value $1/17$ to within the model's precision.

S2. ENTANGLEMENT-BREAKING THRESHOLD

For the channel $\mathcal{N}_\gamma$ of Eq. (1) of the main text we evaluate the coherent information on a maximally entangled purification $|\Phi\rangle$ of a maximally mixed input,

$$I_c = S\big(\mathcal{N}_\gamma(\rho)\big) - S\big[(\mathbb{I}\otimes\mathcal{N}_\gamma)\,|\Phi\rangle\langle\Phi|\,\big], \tag{S3}$$

with $S(\rho) = -\mathrm{Tr}\rho\log_2\rho$. Scanning γ, $\max_\rho I_c$ crosses zero near $\gamma \approx 0.30$, which we take as γ_{c}. We flag that γ_{c} is model-dependent (it depends on the depolarizing weight α in Eq. (1));

the value 0.30 corresponds to $\alpha = 0.3$. The qualitative conclusion $\gamma_{\text{eff}} \gg \gamma_{\text{c}}$ at biological rates is insensitive to this choice over $\alpha \in [0.2, 0.5]$ (main-text Result I).

S3. PETZ RECOVERY FROM A NOISY REFERENCE

The transpose (Petz) recovery map [4] for channel $\mathcal{N}$ and reference state σ is

$$\mathcal{R}_{\sigma,\mathcal{N}}(\cdot) = \sigma^{1/2}\,\mathcal{N}^{\dagger}\Big[\mathcal{N}(\sigma)^{-1/2}(\cdot)\,\mathcal{N}(\sigma)^{-1/2}\Big]\sigma^{1/2}, \tag{S4}$$

where $\mathcal{N}^{\dagger}$ is the adjoint channel. The "noisy-reference" variant replaces σ by a dephased nuclear-spin reference $\sigma \to \mathcal{N}_{\gamma_{\text{ref}}}(\sigma)$ with $\gamma_{\text{ref}} = 0.1$. Because Eq. (S4) requires $\sigma^{1/2}$, $\mathcal{N}^{\dagger}$, and $\mathcal{N}(\sigma)^{-1/2}$—none of which has an identified biological implementation—we report the resulting concurrence $C = 0.46$ strictly as an *information-theoretic bound*, not a mechanism. (This wording corrects earlier drafts that read it as a protocol.)

S4. QEC STABILISER, DFS ENCODING, AND CONCATENATION

DFS encoding. The logical qubit is encoded in the collective-decoherence-free subspace [5] of two physical spins, $|0_L\rangle = |S\rangle = (|01\rangle - |10\rangle)/\sqrt{2}$, $|1_L\rangle = |T_0\rangle = (|01\rangle + |10\rangle)/\sqrt{2}$; collective dephasing acts trivially on this subspace, leaving the residual rate $\gamma_{\text{DFS}} = \gamma_{\text{eff}}(1 - f_{\text{coll}})$.

Stabiliser (gauging + purification). The biological stabiliser used in the QRC pipeline applies, every T_{QEC} steps: (i) gauging, an adaptive projection onto the total-S_z sector of the noiseless state; (ii) purification, $\rho \to \rho^2/\text{Tr}(\rho^2)$ applied twice. It uses only these low-level primitives (no external recovery).

Concatenation. A distance-d code concatenated to level L uses d^L physical copies; the recovered fidelity follows the standard recursion [6, 7] $p_L \approx (p_{L-1}/p_{\text{th}})^2 p_{\text{th}}$ in the sub-threshold regime. At $d = 3$, $L = 3$ (27 copies) and $\gamma_{\text{DFS}} = 0.195$ ($f_{\text{coll}} = 0.94$) the CRY fidelity is $F = 0.9957$ (main-text Table I).

S5. QUANTUM RESERVOIR: READ-OUT, LOSSES, BASELINES

Input encoding and evolution. Per step, the scalar input u_t enters through $H(u_t) = u_t \langle S_z \rangle$; the state is propagated by $U = e^{-\mathrm{i}H(u_t)\delta t}$ and subjected to $\mathcal{N}_\gamma$ [main-text Eq. (1)],

optionally followed by the stabiliser of Sec. S4. The procedure order matches the simulation code exactly and is sketched in Fig. S1.

Read-out observables (the observation method). Thirteen observables per qubit are read: the two diagonal populations; the real and imaginary parts of the off-diagonal coherences; the purity $\mathrm{Tr}\rho^2$; the two-qubit concurrence C; and the singlet yield $Y_S = \mathrm{Tr}(\Pi_S\rho)$ with $\Pi_S = |S\rangle\langle S|$—the same spin-selective yield read physically by Layer 3. These observables form the feature vector $\mathbf{x}_t$.

Losses and where used. Time-series prediction (NARMA-5) trains a ridge-regression read-out $\hat{y}_t = \mathbf{w}^\top \mathbf{x}_t$ by minimizing $\sum_t (\hat{y}_t - y_t)^2 + \beta\|\mathbf{w}\|^2$ and is scored by the normalized mean-squared error [main-text Eq. (7)], $\mathrm{NMSE} = \langle(\hat{y}_t - y_t)^2\rangle/\mathrm{Var}(y_t)$. Classification (8×8 digits) trains a support-vector machine (radial-basis kernel) and is scored by accuracy Acc (held-out fraction correct); the reported loss is $1 - \mathrm{Acc}$. The quantum Fisher information [8], $\mathrm{QFI} = 2\sum_{kl} \frac{(\lambda_k-\lambda_l)^2}{\lambda_k+\lambda_l} |\langle k|\,\partial_\theta\rho\,|l\rangle|^2$ over eigendecomposition $\rho = \sum_k \lambda_k |k\rangle\langle k|$, and the concurrence C, are computed as resource *diagnostics*, not training objectives.

Baselines. Prediction baseline: a hyperparameter-optimized classical echo-state network [9] with the same ridge read-out form (reservoir size scanned $N_{\mathrm{ens}} = 1$–16); best NMSE ≈ 0.017. Classification baseline: an SVM on the raw 64 pixel intensities. Both baselines use identical train/test splits and the same `scikit-learn` estimators [10].

S6. DECISION MODEL: PARAMETERS AND PROCEDURE

The drift-to-bound model [main-text Eqs. (8)–(9)] uses a two-qubit decision register with choice states $|L\rangle = |01\rangle$, $|R\rangle = |10\rangle$ and Hamiltonian $H = -t_{\mathrm{tun}}(|L\rangle\langle R| + \mathrm{h.c.}) + \frac{b}{2}(|L\rangle\langle L| - |R\rangle\langle R|) + J\,\sigma_z\otimes\sigma_z$. The committed-decision parameters are: drift bias $b = 2.0$, tunnelling $t_{\mathrm{tun}} = 0.45$, exchange $J = 0.3$, accumulator leak $\lambda = 0$, readiness-potential threshold $A_{\mathrm{rp}} = 0.03$, motor bound $A_{\mathrm{bound}} = 0.60$, aversive (stop) strength 3.0 and duration 40 steps, step $\delta t = 0.05$, up to 400 steps; QEC every 20 steps. **Veto criterion (corrected).** A veto succeeds iff a stop was applied and the action ($A \geq A_{\mathrm{bound}}$) does not fire within a fixed, duration-independent horizon $T_{\mathrm{eval}} = 80$ steps after the stop, with the go-drive resuming after the pulse. This replaces an earlier criterion that terminated the trial at the (variable) end of the aversive pulse and so spuriously rewarded short pulses (Sec. S7). Ensembles use 40 stochastic trials per delay (small bias jitter $\mathcal{N}(0, 0.02)$); seeds are deterministic (a fixed

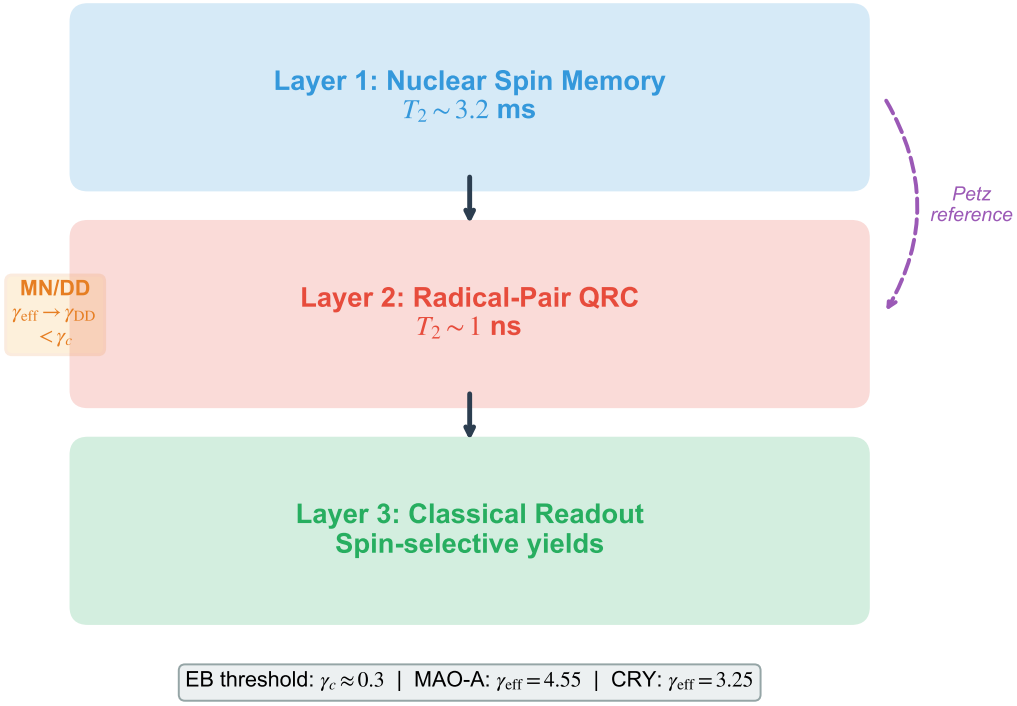

FIG. S1. **Three-layer architecture and the reservoir procedure.** Layer 1 nuclear-spin memory, Layer 2 radical-pair reservoir, Layer 3 spin-selective read-out, with the per-step procedure of Sec. S5.

base offset plus $1000d + s$ for delay d and trial s) for exact reproducibility.

S7. VETO-CRITERION AUDIT AND THE WINDOW-EXTENSION SEARCH

Figure S2 shows the veto-success sigmoids [main-text Eq. (10)] in both reference frames.

The early-termination artifact. An earlier veto criterion declared a success the moment the aversive pulse *ended* with the accumulator still below the bound, terminating the trial and never letting the go-drive resume. A shorter pulse then terminates sooner, giving the accumulator less time to cross—so shorter pulses were spuriously rewarded, producing a non-physical "shorter stop → longer window" trend and inflating the apparent window. We replaced it with the duration-independent, run-to-completion criterion of Sec. S6.

Audit results (corrected criterion). With the fixed horizon and no early termination: (i) the duration artifact vanishes (the window no longer grows for shorter pulses); (ii) the

TABLE S1. Vetoability across regimes (CRY). $\langle a \rangle$ is the mean decision signal; $A(N)$ the accumulator after N steps ($A_{\text{bound}} = 0.60$).

Regime	γ	$\langle a \rangle$	$A(400)$	$A(2000)$	veto window
no MN ($\gamma_{\text{eff}} > \gamma_c$)	3.25	0.0004	0.014	0.014	none (frozen)
classical	50	0.000	0.002	0.001	none (frozen)
MN ($\gamma_{\text{DD}} < \gamma_c$)	0.191	0.102	2.11	2.21	marginal

apparent window is small and *horizon-sensitive*—for the no-QEC baseline the fitted window is $\sim$0.5–2.4 ms at $T_{\text{eval}} = 40$–80 steps and is *undefined* (no $\geq 50\%$ veto) for $T_{\text{eval}} \geq 160$ steps, because a transient stop only *delays* the action; (iii) under a strict "permanent reversal" criterion (accumulator quenched below A_{rp} by the horizon) *no* condition—none, continuous QEC, or stop-triggered QEC—achieves a durable veto; (iv) error correction only shrinks or eliminates the window; with QEC no stop delay reaches a $\geq 50\%$ veto. Hence no operation extends the window, and the robust content is the halving of T_{move} (main text). Figure S3 shows the sweep under the *previous* (early-termination) criterion and is retained only to illustrate the artifact it produced.

Vetoability across the three regimes (no MN / MN / MN+QEC). To answer the natural question "what is the veto window with no motional narrowing?" we run the corrected experiment at the bare biological point ($\gamma_{\text{eff}} > \gamma_c$), the rescued point ($\gamma_{\text{DD}} < \gamma_c$), and the classical limit ($\gamma = 50$); Table S1. Without MN the mean decision signal $\langle a \rangle = \langle P_R - P_L \rangle$ is suppressed $\sim$250-fold ($0.102 \to 0.0004$) and the accumulator saturates at $A = 0.014 \ll A_{\text{bound}} = 0.60$, unchanged between 400 and 2000 steps: the decision is genuinely *frozen* (not merely slow) and numerically identical to the classical limit. Hence the veto window is undefined at the bare point—there is no committed decision to veto—whereas under QEC it is undefined for the opposite reason (the decision forms too fast). Vetoability is therefore non-monotone in coherence, nonzero only at the MN operating point.

S8. PHYSICAL-TIME CALIBRATION AND THE COHERENCE CARRIER

One model step equals $T_2^{\text{carrier}}/\tau_{\text{coh}}$, where τ_{coh} is the $1/e$ coherence-decay time measured *in the simulation* under pure dephasing at the operating γ, and T_2^{carrier} is the coherence time

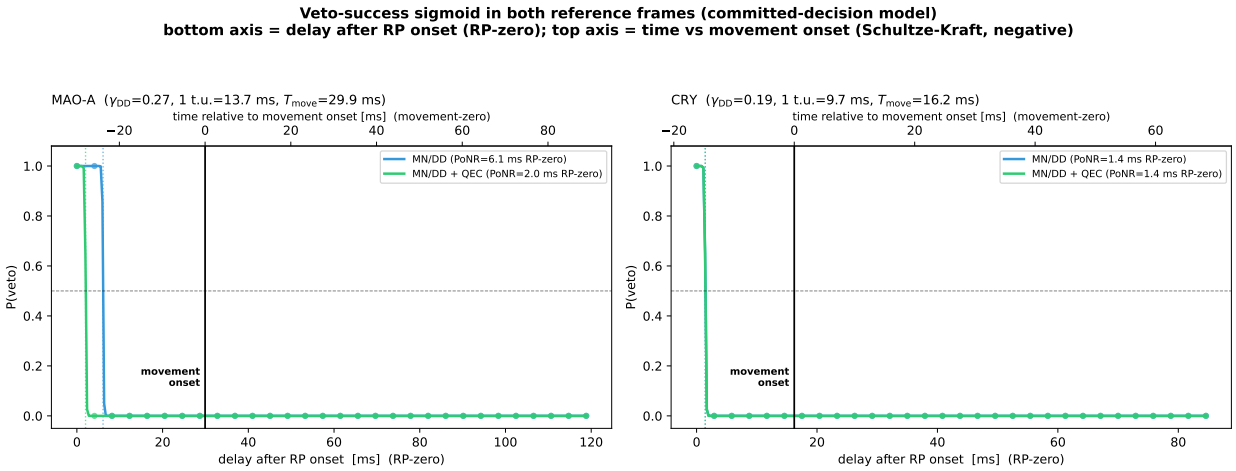

FIG. S2. **Veto-success sigmoids in two frames.** Bottom axis: delay after readiness-potential onset. Top axis: time relative to movement onset (Schultze-Kraft). Error correction shifts the half-success point earlier because it shortens the onset-to-movement latency.

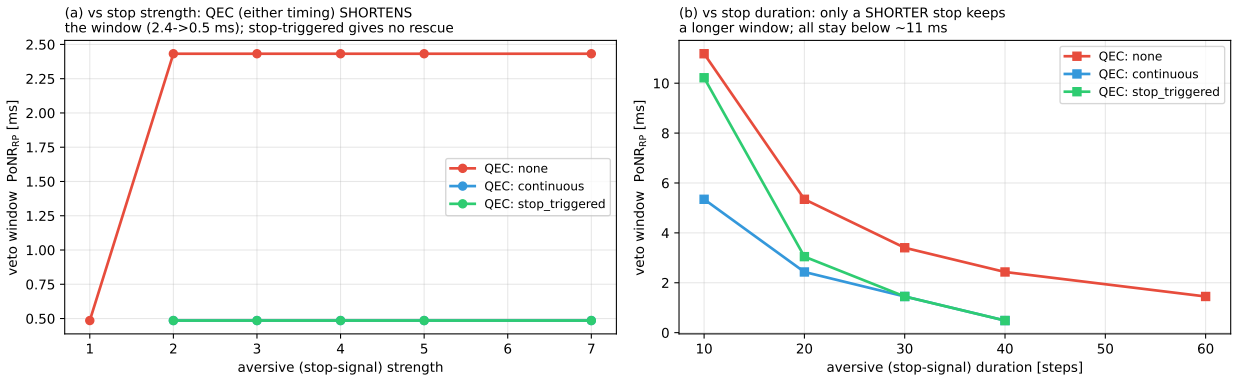

FIG. S3. **Window-extension search under the superseded (early-termination) criterion**—retained to illustrate the artifact. Window (readiness-potential frame) versus stop strength (left) and stop duration (right) for three QEC timings. The non-physical "shorter stop → longer window" trend (right panel) is the early-termination artifact corrected in Sec. S7; the duration-fair audit removes it. Error correction still only shortens the window.

of the layer carrying the decision. The self-consistency check $t_{\mathrm{hf}}/T_{2e} = 3.247 = \gamma_{\mathrm{eff}}$ validates the radical-pair clock. Table S2 lists, purely to illustrate the *calibration scaling* (physical time $\propto T_2$), the candidate carriers and the corresponding model-time-to-physical factors;

TABLE S2. Carrier-resolved calibration scaling (CRY, MN; the "window" is the non-committed-model timescale, shown for calibration only; see caption text).

Carrier (source)	T_2	veto window
electron T_{2e} (single RP episode)	1.54 ns	1.7 ns
^{31}P paramagnetic (open RP)	160 µs	0.18 ms
^{31}P-ATP, in vivo [11] (primary)	12 ms	13.4 ms
^{31}P phosphocreatine (upper bound)	100 ms	111 ms

only the literature ^{31}P-ATP value [11] and the phosphocreatine upper bound are used in the main text, and the previously quoted 40 ms to 52 ms uncited estimate is deprecated. The "window" column is the marginal MN window of the simpler (non-committed, population-flip) veto model used for calibration; it is shown only to set the timescale and is *not* a robust quantity—under the committed drift-to-bound model with the corrected criterion the window collapses (Sec. S7). The robust volition result is the halving of T_{move} (main text), not any window value.

S9. REPRODUCIBILITY

The full study regenerates from the released package and analysis scripts; the exact commands, parameters, and random seeds are listed in the companion data document. Tolerances follow the numpy convention $|a-b| \le \mathrm{atol} + \mathrm{rtol}\,|b|$ with rtol = 2%, atol = 10^{-6}. Software: NumPy, SciPy, QuTiP, scikit-learn [10], Matplotlib. Every number-to-source mapping is in the companion SOURCES file.

S10. HARDWARE DEMONSTRATION: CROSS-DEVICE REPRODUCIBILITY AND THE SPARSE-DD NULL

This section supports the main-text hardware section ("Hardware demonstration of the rescue"). **Protocol.** A Bell state is prepared on two qubits and idled for τ either freely or under a dense XY8 sequence (inter-pulse spacing ≈ 1 µs); the two-qubit state is reconstructed by nine-setting Pauli tomography and scored by fidelity to $|\Phi^+\rangle$ and concurrence (the analysis

TABLE S3. Bell-state fidelity at $\tau = 40\,\mu s$, free vs dense-XY8 DD, for four qubit pairs on two IBM devices.

Device	pair (T_2, μs)	free F	DD F
`ibm_marrakesh`	[2,3] (500/222)	0.52	0.77
`ibm_marrakesh`	[51,58] (161/185)	0.45	0.60
`ibm_fez`	[143,144] (170/217)	0.26	0.68
`ibm_fez`	[83,96] (179/160)	0.49	0.57

pipeline is first validated on a noiseless simulator). **What it tests.** As stressed in the main text this is an analog of the quantum-information rescue mechanism, not the biology; and because the concurrence is T_1-limited it is the *fidelity*—not the concurrence/entanglement-breaking witness—that exhibits the rescue. That dynamical decoupling improves fidelity on superconducting qubits is itself established [12].

The sparse-DD null (run #1). A single, sparse XY4 block over the whole idle (inter-pulse spacing up to $\sim 30\,\mu s$) on transpiler-default qubits showed *no* rescue (free $\approx$ DD): such sparse decoupling only refocuses sub-kHz noise, and the default qubits were not dephasing-limited. The rescue appears only with dense decoupling on dephasing-limited high-T_2 qubits, as expected physically.

Cross-device reproducibility. Table S3 and Fig. S4 collect the fidelity at $\tau = 40\,\mu s$ for four qubit pairs across two devices (`ibm_marrakesh`, `ibm_fez`). Decoupling preserves the fidelity above the free value in every case; on one `ibm_fez` pair the free fidelity is strongly non-monotone (drift) while the decoupled fidelity is smooth, so decoupling also stabilizes the outcome.

[1] A. G. Redfield, IBM Journal of Research and Development **1**, 19 (1957).

[2] A. Abragam, *The Principles of Nuclear Magnetism* (Oxford University Press, Oxford, 1961).

[3] L. Viola, E. Knill, and S. Lloyd, Physical Review Letters **82**, 2417 (1999).

[4] D. Petz, Communications in Mathematical Physics **105**, 123 (1986).

[5] D. A. Lidar, I. L. Chuang, and K. B. Whaley, Physical Review Letters **81**, 2594 (1998).

[6] E. Knill and R. Laflamme, Physical Review A **55**, 900 (1997).

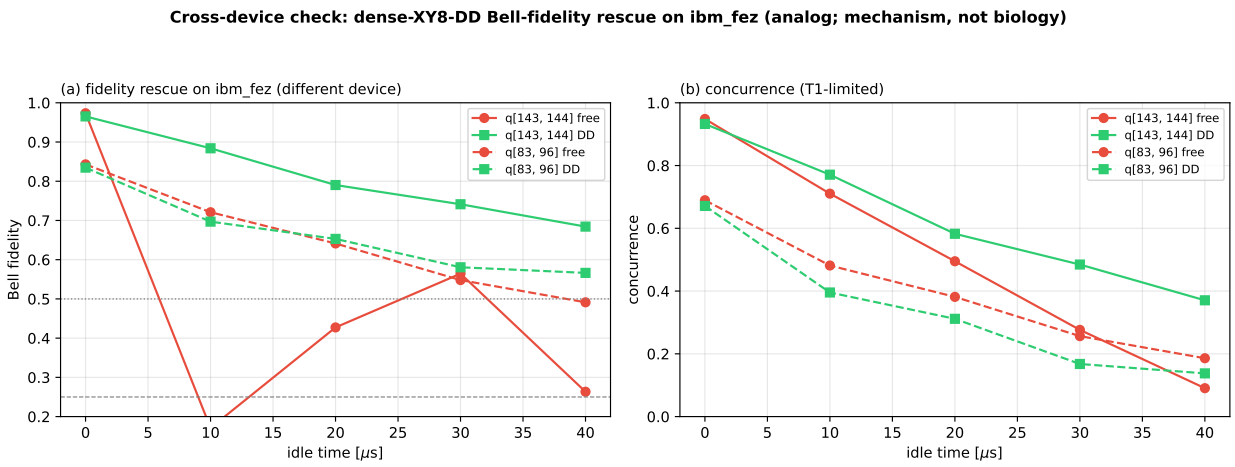

FIG. S4. **Cross-device reproduction (`ibm_fez`).** (a) Bell-state fidelity: dense XY8 decoupling preserves the state where free idling decays toward separable, on both `fez` pairs (one free curve is strongly non-monotone—drift—while the decoupled curve is smooth). (b) Concurrence is T_1-limited (free $\approx$ DD), as on `marrakesh`.

[7] B. M. Terhal, Reviews of Modern Physics **87**, 307 (2015).

[8] S. L. Braunstein and C. M. Caves, Physical Review Letters **72**, 3439 (1994).

[9] H. Jaeger and H. Haas, Science **304**, 78 (2004).

[10] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, Journal of Machine Learning Research **12**, 2825 (2011).

[11] C. Rémy, J. P. Albrand, A. L. Benábid, M. Décorps, and P. Jacrot, Magn. Reson. Med. **4**, 144 (1987).

[12] B. Pokharel, N. Anand, B. Fortman, and D. A. Lidar, Physical Review Letters **121**, 220502 (2018).