

Prefix-Path Bell Transport on IBM Quantum Hardware: High-Resolution Replication, Comparative Geometry Dependence, and Stress Tests of a Static–Dynamic Link

Version 2

Lluis Eriksson
Independent Researcher
lluiseriksson@gmail.com

December 2025 (v1); July 2026 (v2)

Editorial Status: Companion Note. This manuscript is a companion note in the coherence-maintenance program. Role: experimental protocol note: controlled geometry, comparative rate evidence, and stress tests of a static–dynamic link (RIP program context). Program index: https://ai.vixra.org/author/lluis_eriksson

Abstract

This note reports a replicated, high-resolution Bell-transport experiment on IBM Quantum superconducting hardware using a prefix-path protocol that controls spatial heterogeneity across transport lengths. A single physical qubit chain is fixed and increasing transport length L is realized via prefixes of that chain, so that L changes depth while keeping qubits nested rather than switching to different qubit subsets. We reconstruct the Bell-state fidelity $F_{\Phi^+}(L)$ from Pauli correlators E_{XX}, E_{YY}, E_{ZZ} and apply a minimal drift correction using interleaved full- Φ^+ control blocks. Beyond a single-chain sweep, we perform a comparative geometry test across three disjoint physical chains on the same backend: the effective decay scale differs significantly across chains (with $> 10\sigma$ separations under fit uncertainties), providing operational evidence that the transport decay scale is geometry-dependent under fixed compilation constraints. Motivated by the Rate Inheritance Principle (RIP) framing, we also investigate whether a phase-sensitive static correlation metric measured on idle chains can predict dynamical transport decay. A curated three-chain set exhibits an ordering agreement between a static Ramsey-X nearest-neighbor covariance metric and the transport decay scales μ measured on the same chains; however, scale-up studies over $n = 18$ randomly sampled chains and a preregistered out-of-sample prediction test do not show statistically significant monotone association under permutation testing. We interpret the static–dynamic ordering agreement as conditional and geometry-specific, while the geometry dependence of dynamical decay is robust. *v2 adds:* a fitting-choice caveat with a synthetically validated re-analysis pipeline for future variant tables, a power bound showing the negative preregistered test excludes a strong device-wide static–dynamic law (from published quantities alone), an all-lengths rate proxy specification, small-sample and estimator caveats, context references, and an offline re-analysis script with documented schema for a future data revision.

1 Scope and claim level

This manuscript is an experimental companion note. Its primary contribution is a reproducible protocol and dataset that isolates a dominant geometry confound in short-depth transport exper-

iments on superconducting QPUs. We use the phrase *operational rate inheritance* in a minimal comparative sense: changing only the physical geometry (the qubit chain) changes the effective decay scale extracted from an otherwise fixed protocol. We also report phase-sensitive static measurements and stress tests of a static–dynamic link motivated by the RIP program notes [9, 10]. We do not claim a universal microscopic law relating a static correlation envelope $f(\epsilon)$ to a Davies-rate envelope $\kappa(\epsilon)$ across all regimes; the goal is operational: identify whether a practically measurable static structure metric predicts dynamical transport decay on-device. Geometry-dependent error heterogeneity and crosstalk on superconducting devices are well documented [11, 12]; the contribution here is a controlled protocol that turns this into a quantitative comparative measurement under fixed compilation constraints.

2 Hardware, job metadata, and fixed physical geometry

All experiments were executed on IBM Quantum `ibm_fez` via IBM Quantum Runtime (plan: open). The native entangling primitive exposed on the relevant target subgraphs was `cz`. In the protocol description we use logical gates such as `CX` and `SWAP`; these are compiled by Qiskit [7] to the backend basis under the fixed compilation constraints described.

Single-chain high-resolution run (Figure 1). Job ID: `d5aic8jht8fs73a5nrv0`. Shots per circuit: 2000. Repetitions per L : $n_{\text{rep}} = 3$. Control mode: full- Φ^+ drift blocks.

Comparative geometry run (Figure 2). Job ID: `d5aipuonsj9s73b7n65g`. Shots per circuit: 2000. Repetitions per L : $n_{\text{rep}} = 3$. Three disjoint chains (length 8):

$$\begin{aligned} P_1 &= [130, 131, 138, 151, 152, 153, 154, 155], \\ P_2 &= [68, 69, 78, 89, 88, 87, 97, 107], \\ P_3 &= [0, 1, 2, 3, 4, 5, 6, 7]. \end{aligned}$$

Curated static Ramsey-X scan (Figure 4). Job ID: `d5ajj5rht8fs73a5p15g`. Shots per circuit: 2500. Idle times: $T \in \{0, 2, 4, 6, 8\} \mu\text{s}$. (Same three chains as the comparative transport run.)

Scale-up and stress tests (Figure 5 and Table 3). Scale-up v2 job ID: `d5ak62jht8fs73a5pjbg` ($n = 18$ random chains). Edge-level static scan job ID: `d5ak8vbht8fs73a5pm5g` (per-edge covariances). Preregistered screening job ID: `d5akfnrht8fs73a5psm0` (select top/bottom groups by static metric). Preregistered dynamic test job ID: `d5akj33ht8fs73a5pvt0` (out-of-sample prediction test on selected chains).

3 Protocol

3.1 Circuit family

For each transport length L , we set $n = L + 2$ and construct the following circuit on n qubits: prepare Φ^+ on sites $(0, 1)$ via `H` on qubit 0 followed by `CX(0, 1)`; apply a `SWAP` chain transporting the second qubit from site 1 to site $n-1$: `SWAP(1, 2) SWAP(2, 3) ... SWAP(n-2, n-1)`; measure in Pauli bases XX , YY , ZZ by applying standard basis-change rotations on qubits 0 and $n-1$, followed by computational measurement on those two qubits only. The prefix-path protocol fixes

the physical mapping by choosing a physical chain P_{phys} and using its prefix of length n as the initial layout. Routing is disabled, and compilation uses optimization level 0 to avoid structural rewriting.

3.2 Fidelity estimator

For each length L , we estimate correlators E_{XX}, E_{YY}, E_{ZZ} from parity outcomes and compute

$$F_{\Phi^+}(L) = \frac{1 + E_{XX}(L) - E_{YY}(L) + E_{ZZ}(L)}{4}. \quad (1)$$

3.3 Drift controls and correction

To reduce slow drift, we interleave full- Φ^+ control blocks at $L = 0$ (three Pauli bases) and use a minimal control-referenced correction:

$$F_{\text{corr}} = F_{\text{raw}} - (F_{\text{ctrl,near}} - \bar{F}_{\text{ctrl}}).$$

This corrects coarse drift at the level of the control schedule; it is not intended as a full time-series model. (*Added in v2*) F_{corr} is an affine-corrected estimator and is not constrained to $[0, 1]$; no clipping is applied, so that averages remain unbiased.

3.4 Uncertainty reporting

For each L , we run $n_{\text{rep}} = 3$ repetitions and report the mean and standard error of the mean (SEM) across repetitions. (*Added in v2*) With $n_{\text{rep}} = 3$ the SEM itself carries large relative uncertainty (2 degrees of freedom); confidence intervals built from these SEMs should use Student- t factors, and the fit-based uncertainties quoted below inherit this caveat. The scale-up analyses instead use multinomial bootstrapping of shot outcomes, which does not suffer from this small- n_{rep} issue.

4 Results I: single-chain high-resolution sweep

Figure 1 shows the replicated single-chain sweep (prefix-path) with error bars. A descriptive fit of the form $F(L) = Ae^{-\mu L} + B$ can be used as a compact summary; fitted parameters are model-dependent and should be treated as descriptive rather than as direct observables.

5 Results II: comparative geometry test (operational rate inheritance)

To test whether the effective decay scale depends on physical geometry rather than being a universal constant of the protocol, we repeated the sweep on three disjoint physical chains under identical compilation constraints.

5.1 Fit protocol for the comparative test

For the comparative test we summarize each chain by an effective decay scale μ using

$$F(L) = B_{\text{fix}} + (F_0 - B_{\text{fix}})e^{-\mu L}. \quad (2)$$




Figure 1: High-resolution prefix-path Bell transport on `ibm_fez` (single-chain run).

Here B_{fix} is defined operationally for each chain as the empirical mean of the two largest- L fidelity points available on that chain; we then fit (F_0, μ) by weighted nonlinear least squares (weights $1/\sigma_{\text{SEM}}(L)^2$). This stabilizes inference when the asymptote is not identifiable in the sampled L range.

5.2 Observed geometry dependence

The extracted decay scales (baseline variant, as in v1) are

$$\mu_1 = 0.2224 \pm 0.0182, \quad \mu_2 = 0.5398 \pm 0.0186, \quad \mu_3 = 0.3014 \pm 0.0068.$$

Key separations:

$$\mu_2 - \mu_1 = 0.3174 \pm 0.0261, \quad \mu_2 - \mu_3 = 0.2384 \pm 0.0198,$$

corresponding to $> 10\sigma$ differences under the reported uncertainties. The quoted $\pm$ are fit standard errors from the weighted nonlinear least-squares procedure; they do not include systematic/model uncertainty. Because the three chains are disjoint, the corresponding datasets are approximately independent at the hardware level.

5.3 Fitting-choice caveat (added in v2)

B_{fix} (asymptote from the two largest- L points) is a modeling choice that can bias the absolute μ when the asymptote is not reached in the sampled range. Synthetic-data validation of the re-analysis pipeline shipped with this note (`verification/0105/reanalyze_0105.py`, `--demo` mode;




Figure 2: Comparative geometry test on `ibm_fez`. Error bars are SEM over $n_{\text{rep}} = 3$ repetitions per L . Curves show the weighted fit Eq. (2) with the v1 baseline B_{fix} .

Section 8) confirms both effects: on synthetic chains with known ground truth, the v1 B_{fix} variant overestimates μ by $\sim 25\%$, while the *ordering* of well-separated decay scales is preserved across the v1 choice, a free- B three-parameter fit, and $B_{\text{fix}} = 1/4$ (the depolarizing floor), with $\geq 7\sigma$ minimal separation. For the on-device dataset the corresponding variant table requires the archived per- L fidelities and is deferred to a data revision. The headline geometry claim of this note therefore rests on the v1 fit uncertainties quoted above, read as *ordering and separation* rather than absolute μ values, with the explicit caveat that quoted errors exclude model-choice systematics.

6 Calibration snapshot and descriptive static-to-dynamic check

We queried a backend-reported calibration snapshot for the three chains (timestamp: 2025-12-31 11:29 UTC). Reported values are simple averages over qubits (T_1 , readout error) and over adjacent edges (mean two-qubit gate error) along each chain.

Table 1: Backend-reported calibration snapshot for the three disjoint chains.

Path	μ	T_1 (μs)	Avg 2Q err.	Readout err.
1	0.2224 ± 0.0182	133.5	0.0026	0.0181
2	0.5398 ± 0.0186	116.0	0.0029	0.0177
3	0.3014 ± 0.0068	182.2	0.0078	0.0158

The calibration snapshot does not uniquely determine μ ; this is expected, since μ is a protocol-level effective parameter aggregating multiple error sources. As an auxiliary descriptive check, Figure 3 plots drift-corrected fidelities against the composite proxy $\Lambda_{2Q}(L) = N_{2Q}(L) \epsilon_{2Q}$, where $N_{2Q}(L)$ is the observed two-qubit gate count per transpiled circuit and ϵ_{2Q} the mean reported two-qubit gate error along the chain.




Figure 3: Descriptive static-to-dynamic consistency check: drift-corrected mean fidelity vs $\Lambda_{2Q}(L)$. A consistency check, not a universal predictive model.

7 Static measurements and a candidate static–dynamic ordering link

7.1 Motivation (RIP bridge)

In the RIP framing [9], a central hinge is whether a practically measurable static structure observable can predict a dynamical decay penalty extracted from an active protocol. Prior notes emphasize that inheritance is conditional and can fail in specific regimes [10]. Here we test an operational version on-device using a phase-sensitive static probe (Ramsey-X).

7.2 Curated three-chain static–dynamic ordering agreement

On each 8-qubit chain we prepare $|+\rangle^{\otimes 8}$, idle for a duration T , and measure in the X basis. From the resulting bitstring distribution we compute nearest-neighbor covariances $C_{i,i+1}^{(X)}(T) = \langle X_i X_{i+1} \rangle_T - \langle X_i \rangle_T \langle X_{i+1} \rangle_T$, aggregated as $S_{\text{stat}}(T) = \frac{1}{7} \sum_{i=0}^6 |C_{i,i+1}^{(X)}(T)|$, and use the SPAM-corrected scalar $S_{\text{stat,corr}} = \overline{S_{\text{stat}}(T \geq 2\mu\text{s})} - S_{\text{stat}}(0)$. Figure 4 plots μ against $S_{\text{stat,corr}}$ for the curated three-chain set; the ordering agrees (Path1 < Path3 < Path2 in both).

Table 2: Curated three-chain summary on `ibm_fez`.

Path	$S_{\text{stat,corr}}$	μ
1	0.012543	0.222437 ± 0.018202
3	0.016221	0.301424 ± 0.006824
2	0.023694	0.539802 ± 0.018637




Figure 4: Curated three-chain ordering: $S_{\text{stat,corr}}$ vs μ .

7.3 Scale-up stress tests and preregistered prediction test (negative)

To evaluate whether the static metric predicts dynamical decay more broadly, we performed stress tests on randomly sampled chains and a preregistered out-of-sample prediction test.

Dynamic rate proxies. To avoid sensitivity to long-length saturation, v1 defined the short-length proxy

$$\mu_{04} := \frac{1}{4} \ln \left(\frac{1 - F_{\Phi^+}(4)}{1 - F_{\Phi^+}(0)} \right). \quad (3)$$

(Added in v2) Since $L \in \{0, 2, 4, 6\}$ was measured, an all-lengths proxy μ_{all} – the slope of a least-squares fit of $\ln(1 - F)$ over all four lengths – is specified in the re-analysis script as a cross-check on the two-point choice; its evaluation on the archived per- L fidelities is deferred to a data revision.

Scale-up v2 ($n = 18$ random chains). A Spearman rank test with a permutation p -value yields no statistically significant association (job d5ak62jht8fs73a5pjb9): $\rho \approx 0.23$, two-sided permutation $p \approx 0.35$ (Figure 5).

Edge-restricted static metric (negative). Restricting the static metric to the SWAP edges used at $L = 4$ does not improve association (job d5ak8vbht8fs73a5pm5g): $\rho \approx 0.18$, $p \approx 0.46$.

Preregistered out-of-sample prediction test (negative). Using a static screening run on $n = 30$ random chains, we selected $k = 5$ chains with the smallest and largest values of the simplified metric $S_{\text{stat}}(6\mu\text{s}) - S_{\text{stat}}(0)$ (screening job d5akfnrht8fs73a5psm0), then measured μ_{04} on these 10 chains with $n_{\text{rep}} = 3$ (dynamic job d5akj33ht8fs73a5pvt0). The preregistered statistic $\Delta := \overline{\mu_{04}}|_{\text{HIGH}} - \overline{\mu_{04}}|_{\text{LOW}}$ was observed to be $\Delta \approx -0.0014$ with permutation $p_{1\text{-sided}} \approx 0.50$ (two-sided $p = 1.0$).




Figure 5: Scale-up v2 stress test ($n = 18$ random chains): bootstrap means with 16–84% intervals, $S_{\text{stat,corr}}$ vs μ_{04} .

Power of the negative test (added in v2). A negative result is informative only relative to the effects it could have detected. For the preregistered design ($k = 5$ per group, one-sided $\alpha = 0.05$, power 0.80) the minimum detectable effect is

$$\Delta_{\text{MDE}} = (z_{0.95} + z_{0.80}) \sigma_{\text{chain}} \sqrt{2/k} \approx 1.57 \sigma_{\text{chain}},$$

where σ_{chain} is the between-chain dispersion of μ_{04} (numerical evaluation deferred to the data revision). A bound follows from published quantities alone: for an effect of the size suggested by the curated three-chain set (spread ≈ 0.3 in μ) to have escaped detection, one would need $\sigma_{\text{chain}} \gtrsim 0.19$ – roughly 60% of the mean $\overline{\mu_{04}} \approx 0.33$, implausibly large for chains on a single backend. The design was therefore adequately powered against a *strong* device-wide static–dynamic law; the null result rules that out, while remaining compatible with a weak or mechanism-specific association below Δ_{MDE} .

Table 3: Preregistered out-of-sample prediction test on `ibm_fez`.

Group	$\overline{\mu_{04}}$	Notes
LOW	0.334723	$k = 5$ chains, smallest static metric
HIGH	0.333317	$k = 5$ chains, largest static metric

Interpretation. The curated three-chain set exhibits an ordering agreement between the static Ramsey-X metric and the dynamical decay scales. However, the scale-up and preregistered tests do not support a device-wide monotone static–dynamic law under the present metrics and sampling, within the power quantified above. Either (i) the dominant mechanisms determining transport decay on generic chains are not captured by the Ramsey-X nearest-neighbor covariance metric, or

(ii) the relevant static predictor is more mechanism-specific (e.g., two-qubit ZZ -type static probes on transport edges), which is left for future work.

8 Reproducibility artifacts

All primary figures are generated from IBM Quantum Runtime job outputs identified by the Job IDs in Section 2 and from the analysis scripts used in the original runs. The processed JSON artifacts from v1 (`rip_comparative_fit_v3_fixB.json`, `rip_strong_scaleup_v2.json`, `rip_predictive_dynamic.json`) were *not* preserved outside the provider account; therefore this release does not claim fully offline reproduction of the on-device numbers. The Job IDs remain the primary data reference, and job results remain retrievable through the IBM Quantum account records (e.g. `service.job(<id>).result()` in `qiskit-ibm-runtime`) while retained by the provider. The series repository ships `verification/0105/reanalyze` with a documented canonical JSON schema and a synthetic-data validation mode (`--demo`), covering the fit-variant table, scale-up rank tests, the preregistered statistic and its permutation p -values, the power analysis, and the μ_{all} proxy. A future data revision will populate `results/0105/` – ideally with raw per-circuit counts (chain id, L , basis, repetition, timestamp) rather than processed summaries – enabling full offline re-analysis. Unlike the other verification suites in this series, this script *re-analyzes archived measurements*; it cannot regenerate them.

9 Conclusion

We reported (i) a replicated, high-resolution prefix-path Bell-transport dataset on `ibm_fez` and (ii) a comparative geometry test across three disjoint physical chains, yielding $> 10\sigma$ differences in an effective decay scale μ under the stated fit protocol (ordering read as the robust claim; see the fitting-choice caveat). We also investigated a candidate static–dynamic link: a curated three-chain set exhibits an ordering agreement, but scale-up analyses and a preregistered out-of-sample prediction test (adequately powered against a strong device-wide law) do not show statistically significant association under the present operationalization. We present the static–dynamic ordering agreement as conditional and geometry-specific, and treat identification of a more mechanism-aligned static predictor as an open experimental direction.

Note added in v2

No experimental data changed; all Job IDs, figures and v1 numbers are unchanged. v2 adds analysis-robustness and reproducibility layers: (1) fitting-choice caveat (Section 5.3): the B_{fix} bias and the ordering-stability of well-separated scales are demonstrated on synthetic ground truth; the on-device variant table is deferred to a data revision; (2) power bound for the negative preregistered test (Section 7.3): the design excludes a strong device-wide static–dynamic law from published quantities alone, with the numerical MDE deferred; (3) all-lengths rate proxy μ_{all} specified; (4) small-sample caveats (t -factors for $n_{\text{rep}} = 3$ SEMs; no clipping of the affine drift-corrected estimator); (5) offline re-analysis script with documented schema and synthetic validation (`reanalyze_0105.py`); (6) context references on device error heterogeneity and crosstalk [11, 12].

References

- [1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press (2000).
- [2] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems*, Oxford University Press (2002).
- [3] E. B. Davies, *Markovian master equations*, Commun. Math. Phys. **39** (1974) 91–110.
- [4] D. Petz, *Sufficient subalgebras and the relative entropy of states of a von Neumann algebra*, Commun. Math. Phys. **105** (1986) 123–131.
- [5] O. Fawzi and R. Renner, *Quantum conditional mutual information and approximate Markov chains*, Commun. Math. Phys. **340** (2015) 575–611.
- [6] E. Magesan, J. M. Gambetta, and J. Emerson, *Scalable and robust randomized benchmarking of quantum processes*, Phys. Rev. Lett. **106** (2011) 180504.
- [7] Qiskit contributors, *Qiskit: An Open-source Framework for Quantum Computing*, <https://qiskit.org>.
- [8] L. Eriksson, *Coherence-maintenance program index*, https://ai.vixra.org/author/lluis_eriksson.
- [9] L. Eriksson, *The Rate Inheritance Principle: From Static Correlations to Dynamical Decoherence Rates*, ai.viXra:2512.0072 (2025).
- [10] L. Eriksson, *Stress Testing the Rate Inheritance Principle: Spectral Decoherence Rates and an Operational Resource Horizon*, ai.viXra:2512.0070 (2025).
- [11] M. Sarovar, T. Proctor, K. Rudinger, K. Young, E. Nielsen, and R. Blume-Kohout, *Detecting crosstalk errors in quantum information processors*, Quantum **4** (2020) 321.
- [12] V. Tripathi, H. Chen, M. Khezri, K.-W. Yip, E. M. Levenson-Falk, and D. A. Lidar, *Suppression of crosstalk in superconducting qubits using dynamical decoupling*, Phys. Rev. Applied **18** (2022) 024068.