

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

Lluis Eriksson
Independent Researcher
lluiseriksson@gmail.com

December 2025

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 (references and updates):

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 (high-resolution run), we perform a comparative geometry test across three disjoint physical chains on the same backend. We find that the effective decay scale extracted from the same protocol differs significantly across chains (with $> 10\sigma$ separations), 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. Accordingly, we interpret the static–dynamic ordering agreement as conditional and geometry-specific under the present operationalization, while the geometry dependence of dynamical decay is robust.

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 experiments 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 here is operational: identify whether a practically measurable static structure metric predicts dynamical transport decay on-device.

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 below we use logical gates such as `CX` and `SWAP`; these are compiled by Qiskit to the backend basis (including `cz` and single-qubit rotations) 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):

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

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

Scale-up and stress tests (Figure 5 and Table 3). Scale-up v2 job ID: `d5ak62jht8fs73a5pjbq` ($n = 18$ 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 $n - 1$:

$$\text{SWAP}(1, 2) \text{SWAP}(2, 3) \cdots \text{SWAP}(n - 2, n - 1).$$

- Measure in Pauli bases XX , YY , and 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:

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

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.

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.

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; however, 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:

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

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.




Figure 1: High-resolution prefix-path Bell transport on `ibm_fez` (single-chain run). The figure file is `rip_final_defense.png`.

5.2 Observed geometry dependence

The extracted decay scales 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:

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

which correspond to $> 10\sigma$ differences under the reported uncertainties.

Here the quoted \pm uncertainties on μ are fit standard errors from the weighted nonlinear least-squares procedure (weights given by the per- L SEM across $n_{\text{rep}} = 3$ repetitions). These error bars do not include additional systematic/model uncertainty (e.g., alternative functional forms, asymptote-identifiability choices such as B_{fix} , or residual drift not captured by the control-referenced correction). Because the three chains are disjoint, the corresponding datasets are approximately independent at the hardware level; we therefore interpret the quoted $> 10\sigma$ separations as operational (fit-based) evidence for geometry dependence under the fixed protocol and compilation constraints.




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 $F(L) = B_{\text{fix}} + (F_0 - B_{\text{fix}})e^{-\mu L}$ with B_{fix} defined as the empirical mean of the two largest- L points per chain.

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 (timestamp: 2025-12-31 11:29 UTC).

| 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 |

6.1 Static-to-dynamic consistency (descriptive)

We emphasize that the calibration snapshot does not uniquely determine μ . This is expected: μ is a protocol-level effective parameter aggregating multiple error sources (coherent compilation effects, gate heterogeneity along the chain, crosstalk, relaxation/dephasing, and SPAM contributions). Accordingly, a monotone dependence on any single averaged calibration metric is not expected in general.

To provide an auxiliary descriptive check, we plot drift-corrected fidelities against the composite proxy

$$\Lambda_{2Q}(L) = N_{2Q}(L) \bar{\epsilon}_{2Q},$$

where $N_{2Q}(L)$ is the observed two-qubit gate count per transpiled circuit at length L (metadata), and $\bar{\epsilon}_{2Q}$ is the mean reported two-qubit gate error along the chain.




Figure 3: Descriptive static-to-dynamic consistency check: drift-corrected mean fidelity (SEM over $n_{\text{rep}} = 3$) plotted against $\Lambda_{2Q}(L) = N_{2Q}(L) \bar{\epsilon}_{2Q}$. This plot is intended as a consistency check rather than 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 of this idea 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 $+\otimes^8$, idle for a duration T , and measure in the X basis using a final layer of Hadamards. From the resulting bitstring distribution we compute nearest-neighbor covariances in the X basis:

$$C_{i,i+1}^{(X)}(T) = \langle X_i X_{i+1} \rangle_T - \langle X_i \rangle_T \langle X_{i+1} \rangle_T, \quad i = 0, \dots, 6.$$

We aggregate these into

$$S_{\text{stat}}(T) = \frac{1}{7} \sum_{i=0}^6 \left| C_{i,i+1}^{(X)}(T) \right|.$$

To remove static offsets dominated by SPAM/readout correlations, we use the corrected scalar metric

$$S_{\text{stat,corr}} = \overline{S_{\text{stat}}(T \geq 2 \mu\text{s})} - S_{\text{stat}}(0),$$

where $\overline{(\cdot)}$ denotes an average over measured $T \in \{2, 4, 6, 8\} \mu\text{s}$.

Figure 4 plots the dynamical decay scale μ extracted from Bell transport against $S_{\text{stat,corr}}$ measured on the idle chain for the curated three-chain set. The ordering agrees: $\text{Path}_1 < \text{Path}_3 < \text{Path}_2$ for both static structure and dynamical rate in this curated set.




Figure 4: Curated three-chain ordering: static structure metric $S_{\text{stat,corr}}$ (Ramsey- X nearest-neighbor covariance in idle) vs dynamical decay rate μ from Bell transport on the same three chains. Ordering agreement holds for this curated set.

Table 2: Curated three-chain summary on `ibm_fez`. Static job: `d5ajj5rht8fs73a5p15g`. Dynamic job: `d5aipuonsj9s73b7n65g`. Static metric is $S_{\text{stat,corr}} = S_{\text{stat}}(T \geq 2 \mu\text{s}) - S_{\text{stat}}(0)$.

| 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 |

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

The curated ordering agreement above involves only three chains. To evaluate whether the same static metric predicts dynamical decay more broadly on the device, we performed stress tests on randomly sampled chains and a preregistered out-of-sample prediction test.

Dynamic rate proxy used in stress tests. To avoid sensitivity to long-length saturation points, we define a robust short-length dynamical rate proxy using $L = 0$ and $L = 4$:

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

Uncertainties for the static metric $S_{\text{stat,corr}}$ and for μ_{04} in the scale-up analysis are estimated by multinomial bootstrapping of the observed outcome distributions.

Scale-up v2 ($n = 18$ random chains). We sampled $n = 18$ random length-8 chains on `ibm_fez` and measured $S_{\text{stat,corr}}$ (Ramsey- X scan) and μ_{04} (Bell transport at $L \in \{0, 2, 4, 6\}$). A Spearman rank test with a permutation p -value yields no statistically significant association (job `d5ak62jht8fs73a5pjbg`): Spearman $\rho \approx 0.23$ with two-sided permutation $p \approx 0.35$. Figure 5 shows the corresponding scatter (with bootstrap intervals).




Figure 5: Scale-up v2 stress test on `ibm_fez` ($n = 18$ random chains): bootstrap means with 16–84% intervals for static metric $S_{\text{stat,corr}}$ vs dynamical proxy μ_{04} . A rank-based permutation test shows no statistically significant monotone association under this operationalization.

Edge-restricted static metric (negative). To better match the transport support, we also computed an edge-restricted static metric using only the nearest-neighbor covariance terms corresponding to the SWAP edges used at $L = 4$. This does not improve association in the tested dataset (job `d5ak8vbt8fs73a5pm5g`): the edge-restricted metric yields Spearman $\rho \approx 0.18$ with permutation $p \approx 0.46$.

Preregistered out-of-sample prediction test (negative). We preregistered a prediction test as follows: using a static screening run on $n = 30$ random chains, we selected $k = 5$ chains with the smallest and largest values of a simplified static metric $S_{\text{stat,corr}} = S_{\text{stat}}(6 \mu s) - S_{\text{stat}}(0)$ (screening job `d5akfnrht8fs73a5psm0`). We then measured μ_{04} on these 10 selected chains with $n_{\text{rep}} = 3$ repeats each (dynamic test job `d5akj33ht8fs73a5pvt0`). The preregistered group statistic

$$\Delta := \overline{\mu_{04} \mid \text{HIGH}} - \overline{\mu_{04} \mid \text{LOW}}$$

was observed to be $\Delta \approx -0.0014$ with permutation $p_{1\text{-sided}} \approx 0.50$ (two-sided $p = 1.0$), i.e., no support for the preregistered directional prediction.

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 prereg-

Table 3: Preregistered out-of-sample prediction test on `ibm_fez`. LOW/HIGH groups ($k = 5$ each) selected by static screening ($S_{\text{stat}}(6\mu\text{s}) - S_{\text{stat}}(0)$) before running dynamics. Dynamic metric is $\mu_{04} = \frac{1}{4} \ln\left(\frac{1-F(4)}{1-F(0)}\right)$. Reported are group means over per-chain repeat means ($n_{\text{rep}} = 3$).

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

istered prediction tests do not support a device-wide monotone static–dynamic law under the present metrics and sampling procedure. This suggests that, on `ibm_fez`, 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.

Reproducibility artifacts. All primary figures in this note are generated from IBM Quantum Runtime job outputs identified by the Job IDs listed in Section 2, together with archived analysis scripts and JSON artifacts (e.g., `rip_comparative_fit_v3_fixB.json`, `rip_strong_scaleup_v2.json`, and `rip_predictive_dynamic.json`). The figures referenced in this manuscript correspond to the files `rip_final_defense.png`, `rip_comparative_plot_v3_fixB.png`, `rip_strong_link.png`, and `rip_strong_scaleup_v2.png`.

8 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. The latter yields $> 10\sigma$ differences in an effective decay scale μ extracted under fixed protocol and compilation constraints, providing operational evidence that the transport decay scale is geometry-dependent.

We also investigated a candidate static–dynamic link motivated by RIP-style framing using a phase-sensitive static Ramsey- X metric. A curated three-chain set exhibits an ordering agreement between the static metric and the measured dynamical decay scales. However, scale-up analyses and a preregistered out-of-sample prediction test do not show statistically significant association under the present operationalization on `ibm_fez`. We therefore 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.

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 (author profile)*. 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). <https://ai.vixra.org/abs/2512.0072>
- [10] L. Eriksson, *Stress Testing the Rate Inheritance Principle: Spectral Decoherence Rates and an Operational Resource Horizon*, ai.viXra:2512.0070 (2025). <https://ai.vixra.org/abs/2512.0070>