Monitored feedback versus open-loop control for Kuramoto–XY synchronisation on IBM dynamic circuits

Miroslav Šotek[1,2,*]

[1] *ANULUM, Marbach SG, Switzerland*

[2] *ORCID: 0009-0009-3560-0851*

(Dated: May 20, 2026)

(Nodo DOI: 10.5281/zenodo.20382099)

Abstract

Dynamic circuits make it possible to test monitored quantum-control policies on gate-model hardware rather than only static circuit families. We report a preregistered four-qubit Kuramoto–XY feedback experiment on `ibm_kingston`. The payload uses three system qubits, one monitor qubit, three dynamic rounds, and paired hardware arms: a monitored-feedback arm and a matched open-loop control. The preregistered binary synchronisation endpoint is negative: the feedback arm gives mean target error 0.2148958333, while the matched control gives 0.2116406250, so the feedback policy does not improve the selected target. We then replicate the full S1–S1f stack on `ibm_fez` and `ibm_marrakesh`. The primary binary endpoint is positive on those two backends, with target-error improvements 0.0927777778 and 0.0678710938, respectively. Direct-observable extensions show that this is not a simple backend-independent controller promotion: Fez gives mostly negative `YYI` extension structure, Marrakesh gives consistently positive `YYI` and positive primary response, and the original Kingston run remains a negative primary falsifier. The contribution is therefore a hardware-control boundary result. A dedicated follow-up lane adds full three-bit readout calibration and local gate-folding ZNE on Fez and Marrakesh for S1b–S1f. Readout mitigation does not erase the feedback/control separations, and ZNE generally preserves or enlarges them except for the Fez S1f quadrature check. Paired dynamic-circuit feedback tests and direct Pauli readouts can reveal backend- and calibration-window sensitive response modes, but the tested controller is not supported as a backend-general or robust synchronisation-improving policy. A same-layout Marrakesh repeat remains positive, but with a smaller target-error improvement, making drift robustness a measured constraint rather than an assumption.

* protoscience@anulum.li

I. INTRODUCTION

Kuramoto synchronisation is a canonical model for collective phase dynamics and has a long history in nonlinear science [1, 2]. Gate-model quantum circuits can encode small XY-type phase-coupled systems, but noisy intermediate-scale quantum hardware requires conservative claim boundaries [3, 4]. In this setting, the scientific question is not whether a broad quantum advantage has been shown. The useful question is narrower: can a small monitored feedback policy improve a preregistered hardware observable relative to a matched open-loop arm under the same backend, layout, shot budget, repetitions, and calibration window?

This work is the dynamic-circuit control branch of the same Kuramoto–XY hardware programme that produced the DLA-parity leakage observation, the FIM hardware falsification, and the reduced-Pauli mechanism-separation study [5–7]. Its role is not to promote a controller after one positive backend response. It asks whether monitored feedback has a reproducible paired-arm signature once direct observables, backend replication, readout calibration, and local-folding ZNE are all treated as part of the claim boundary.

Measurement-conditioned feedback is a standard quantum-control primitive [8]. Dynamic circuits provide the required gate-model hardware surface: mid-circuit measurement, classical conditioning, and subsequent gates within one submitted payload [9–11]. Mid-circuit readout has also become a practical resource for near-term quantum simulation [12], and adaptive circuits can express tasks unavailable to fixed-depth unitary circuits at the same depth [13]. They also create a stricter experimental standard. A feedback arm alone is not interpretable, because backend drift, layout selection, readout bias, and ordinary shot noise can all masquerade as control response. The S1 experiment therefore uses a paired-arm design: every feedback run is compared with a matched open-loop control generated from the same circuit body and executed in the same hardware window.

This paper reports the completed S1 hardware programme. The primary preregistered binary proxy is negative. The same-paper extensions then ask whether the negative proxy hides smaller channel-level responses in direct XY-sector observables. The final result is scientifically useful, but conservative: the tested controller is not promoted. Instead, the paired-arm method identifies a structured failure and sensitivity boundary for future redesigned feedback policies.

Protocol field	Value
System qubits	3
Monitor qubits	1
Classical bits	6
Dynamic rounds	3
Hardware arms	feedback, matched open-loop control
Primary circuits	2
Shots per circuit	1024
Repetitions per arm	12
Target synchronisation proxy	0.72

TABLE I. Primary S1 preregistered dynamic-circuit payload.

II. PREREGISTERED S1 PROTOCOL

The primary S1 payload uses three system qubits and one monitor qubit. The feedback arm runs three dynamic rounds. In each round, the monitor qubit is coupled to the system, measured, unconditionally reset, and used to condition a small correction on the system. The matched open-loop arm preserves the same body structure without using the measured monitor outcome as an active feedback signal. The readiness artefact is generated by the committed S1 readiness script, and the canonical preregistration package is the JSON file dated May 6, 2026 under `data/s1_feedback_loop/`. The live payload used unconditional monitor reset and conditional correction only, because the submitted provider path rejected reset operations inside conditional blocks. This implementation detail is part of the experimental boundary: the circuit tests provider-supported dynamic-circuit feedback, not arbitrary low-latency real-time control.

The primary endpoint is the absolute target-order-parameter error,

$$\epsilon_{\mathrm{target}} = |R_{\mathrm{live}} - R_{\mathrm{target}}|, \tag{1}$$

where $R_{\mathrm{target}} = 0.72$. The feedback arm is useful under the preregistered decision rule only if it reduces this error relative to the matched open-loop control.

III. DIRECT-OBSERVABLE EXTENSIONS

The primary binary proxy saturates near one on the selected backend. To avoid over-interpreting a single nearly saturated readout, S1 includes a series of same-paper direct-observable extensions. These are not separate papers and do not change the primary endpoint. They localise the failure mode by measuring XY-sector and control observables from matched feedback/control arms.

S1b preserves the original dynamic feedback body and measures final `IXX`, `IYY`, `XXI`, and `YYI` correlators. S1c reduces the body to one dynamic round, correction angle 0.06, and base gain 0.4. S1d compares three preregistered one-round policy variants: `current_shallow_positive`, `polarity_flipped`, and `weak_positive`. S1e repeats that policy sweep at five repetitions per arm. S1f keeps the S1d/S1e-favoured shallow policy and measures `XXI`, `XYI`, `YXI`, `YYI`, and `ZZI`. The `XYI`/`YXI` rows test cross-quadrature leakage, while `ZZI` is a population-sector control.

Figure 1 gives the extension-level summary across `ibm_kingston`, `ibm_fez`, and `ibm_marrakesh`. The figure is intended to expose the same logic as the tables: the direct observable response is structured, but its sign and localisation are backend-dependent.

For an observable P, each arm expectation is reduced from counts as

$$\langle P \rangle = \frac{1}{N} \sum_{b} n_b \prod_{i \in \mathrm{supp}(P)} (-1)^{b_i}, \tag{2}$$

with n_b the observed count for bitstring b and N the total shots. The reported extension quantity is

$$\Delta_P = \langle P \rangle_{\mathrm{feedback}} - \langle P \rangle_{\mathrm{control}}. \tag{3}$$

IV. PRIMARY S1 RESULT AND REPLICATION

The primary paired-arm execution first completed on `ibm_kingston`. The feedback job was `d86qn3lg7okc73elg2eg`; the matched open-loop control job was `d86qn65g7okc73elg2hg`. Each arm used 12 repetitions and 12,288 total shots. The transpiled depths were 717 for the feedback arm and 684 for the matched control. This first execution was negative under the preregistered decision tree.

After review, the same primary endpoint was replicated on `ibm_fez` and `ibm_marrakesh`. The Fez jobs were `d86sl45g7okc73elimdg` and `d86slcp789is73904bng`. The Marrakesh

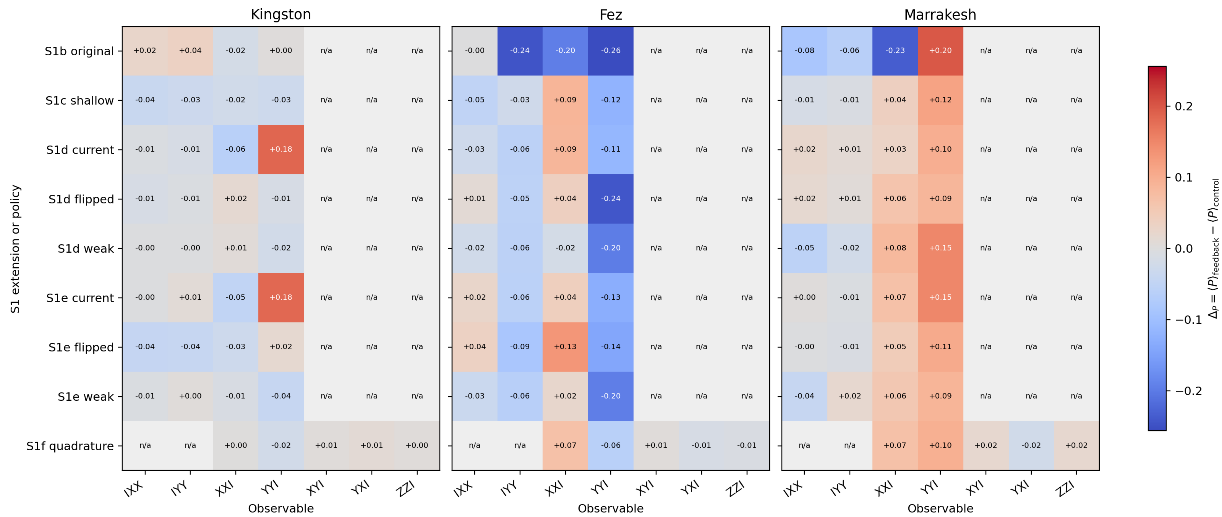

FIG. 1. Feedback-minus-control Pauli expectations across the S1 direct observable extensions on three IBM Heron backends. Grey cells mark observables not included in that extension. The original Kingston `YYI` feature is not backend-universal: Fez is mostly negative in the `YYI` column, while Marrakesh is positive through the shallow policy and quadrature checks.

Backend	Feedback error	Control error	Improvement	Decision
`ibm_kingston`	0.2148958333	0.2116406250	-0.0032552083	negative
`ibm_fez`	0.0355338542	0.1283116319	0.0927777778	positive
`ibm_marrakesh`	0.0440559896	0.1119270833	0.0678710938	positive
`ibm_marrakesh` repeat	0.0232183160	0.0269118924	0.0036935764	positive

TABLE II. Primary S1 binary-proxy hardware result across three backends. The improvement column is control target error minus feedback target error, so positive values favour feedback. The final row is a same-layout Marrakesh repeat on physical layout [7,17,6,8].

jobs were `d86sr9h789is73904jm0` and `d86src2s46sc73f7edm0`. Both same-day replication runs were positive on the binary target-error endpoint, as shown in Table II. A later Marrakesh same-layout repeat pinned to system qubits [7,17,6] and monitor qubit 8 also remained positive. Its jobs were `d86tk8op0eas73dlf4pg` and `d86tkbp789is73905hkg`. The repeat reduced the effect size from 0.0678710938 to 0.0036935764, so it supports the sign of the Marrakesh primary response but not calibration-window invariance.

The primary result is therefore not a universal controller failure, but it is also not a

Observable	S1b Δ_P	S1c Δ_P
`IXX`	0.0247395833	-0.0364583333
`IYY`	0.0358072917	-0.0319010417
`XXI`	-0.0175781250	-0.0221354167
`YYI`	0.0039062500	-0.0292968750
Mean $\lvert\Delta_P\rvert$	0.0205078125	0.0299479167

TABLE III. S1b direct XY-sector extension and S1c shallow/gain-tuned extension. S1b is mixed-sign; S1c is negative in all four channels.

backend-general controller success. The same preregistered three-round payload is negative on Kingston and positive on Fez and Marrakesh. This makes backend and calibration-window dependence part of the main result rather than a footnote. The direct-observable stack below is needed to determine whether the positive primary endpoint corresponds to a stable mechanism.

V. XY-SECTOR LOCALISATION

S1b asks whether the negative binary endpoint hides channel-level differences in direct XY-sector observables. The original dynamic feedback body is preserved, but the final basis is rotated to the selected Pauli correlator. The results are shown in Table III.

S1b is not a positive controller result. Three channels move in the feedback direction and one moves against it, with mean absolute separation only 0.0205078125. It is, however, more informative than the saturated binary proxy: the feedback/control difference is channel-structured rather than uniformly absent.

S1c then tests whether the depth and gain of the original dynamic body are the problem. The shallower, lower-gain policy reduces the feedback depths to about 237, but all four XY-sector channels move negative. The direct-observable response therefore cannot be rescued by this shallow/gain-tuned policy on Kingston. The Fez and Marrakesh replications show a different boundary: the same direct-observable lanes have larger aggregate separations than Kingston, but their channel signs disagree. Fez is dominated by negative `YYI`/`IYY` structure, whereas Marrakesh has a persistent positive `YYI` column.

Run	Variant	IXX	IYY	XXI	YYI	Mean signed	Mean abs.
S1d	current_shallow_positive	-0.007813	-0.013672	-0.060547	0.184896	0.025716	0.066732
S1d	polarity_flipped	-0.013021	-0.006510	0.015625	-0.012370	-0.004069	0.011882
S1d	weak_positive	-0.002604	-0.003906	0.006510	-0.020182	-0.005046	0.008301
S1e	current_shallow_positive	-0.003516	0.008203	-0.048828	0.182422	0.034570	0.060742
S1e	polarity_flipped	-0.042969	-0.040625	-0.028125	0.016797	-0.023730	0.032129
S1e	weak_positive	-0.007422	0.004297	-0.010547	-0.040625	-0.013574	0.015723

TABLE IV. S1d policy-direction sweep and S1e higher-repetition repeat. Values are feedback-minus-control Pauli expectations.

VI. POLICY-DIRECTION SWEEP AND REPEAT

S1d compares three one-round policy variants. S1e repeats the same policy sweep at five repetitions per arm. The purpose is to distinguish a genuine policy-direction response from a single noisy calibration window. The results are shown in Table IV. The artefact timestamps place S1d at 13:46:14Z and S1e at 14:09:19Z on May 20, 2026, a separation of about 23 minutes under the same backend family and day.

The same narrow feature appears in both S1d and S1e: `current_shallow_positive` has a large positive `YYI` separation, 0.1848958333 in S1d and 0.1824218750 in S1e. The alternative policies remain close to zero or negative by mean signed separation. This is a real policy-sensitive signal candidate, but it is too narrow for controller promotion: `IXX` and `XXI` remain negative for the same policy, and the favourable aggregate is driven by one channel.

VII. QUADRATURE MECHANISM CHECK

S1f is the final same-paper discriminator. It keeps `current_shallow_positive` and tests whether the repeated `YYI` response belongs to a stable quadrature structure. It measures the original `XXI`/`YYI` pair, the cross-quadrature controls `XYI`/`YXI`, and a population control `ZZI`. The result is shown in Table V. Its artefact timestamp is 14:22:13Z, about 13 minutes after S1e and about 103 minutes after the primary S1 paired-arm timestamp. This is therefore a later same-day calibration-window check, not a multi-day stability claim.

The S1f result blocks the strongest possible positive interpretation of S1d and S1e on Kingston. The `YYI` separation changes sign and the mean absolute feedback–control separation falls to 0.0082812500. The cross-quadrature controls are small, and the population

Observable	Feedback	Control	Δ_P
XXI	0.2187500000	0.2175781250	0.0011718750
XYI	0.1617187500	0.1566406250	0.0050781250
YXI	0.2242187500	0.2136718750	0.0105468750
YYI	-0.0445312500	-0.0242187500	-0.0203125000
ZZI	0.6523437500	0.6480468750	0.0042968750
Mean $\|\Delta_P\|$			0.0082812500

TABLE V. S1f quadrature and population-control check. The large positive YYI response seen in S1d/S1e does not persist in this later window.

Backend	Primary improvement	S1b mean abs.	S1c mean abs.	S1d current YYI	S1e current YYI	S1f YYI	S1f mean abs.
ibm_kingston	-0.003255	0.020508	0.029948	0.184896	0.182422	-0.020313	0.008281
ibm_fez	0.092778	0.173991	0.071777	-0.114583	-0.128906	-0.060547	0.031484
ibm_marrakesh	0.067871	0.145182	0.044596	0.099609	0.148047	0.096875	0.045313

TABLE VI. Compact cross-backend S1–S1f summary. The primary improvement is control target error minus feedback target error. Direct-observable entries are feedback-minus-control Pauli expectations or mean absolute separations.

control is also small. On Fez, S1f remains negative in YYI with mean absolute separation 0.0314843750. On Marrakesh, S1f remains positive in YYI with mean absolute separation 0.0453125000. The conservative interpretation is not a stable backend-general quadrature-localised feedback mechanism; it is a backend-dependent response boundary.

VIII. READOUT MITIGATION AND ZNE STRESS LANE

The first S1b–S1f direct-observable executions did not include dedicated readout calibration circuits. We therefore do not retrofit mitigation onto those artefacts. Instead, we execute a new preregistered readout/ZNE lane on Fez and Marrakesh. Each backend/lane submission contains the folded main circuits at local gate-folding noise scales 1, 3, and 5, plus all eight computational-basis calibration circuits for the three final system-readout bits. The calibration and main circuits are pinned to the same selected system/monitor physical layout. The folding is local rather than global unitary folding because the S1 feedback

Backend	Lane	Scale-1 mean abs.	Linear ZNE mean abs.	Readout-mitigated linear ZNE mean abs.	Readout matrix cond.
`ibm_fez`	S1b	0.056315	0.063192	0.063868	1.078215
`ibm_fez`	S1c	0.031738	0.038425	0.039430	1.059059
`ibm_fez`	S1d	0.038032	0.043936	0.045297	1.059970
`ibm_fez`	S1e	0.037435	0.045204	0.046707	1.061207
`ibm_fez`	S1f	0.054141	0.041439	0.042975	1.066406
`ibm_marrakesh`	S1b	0.030273	0.040256	0.040437	1.036386
`ibm_marrakesh`	S1c	0.113932	0.120402	0.122239	1.029654
`ibm_marrakesh`	S1d	0.095486	0.101951	0.102965	1.031348
`ibm_marrakesh`	S1e	0.104590	0.113810	0.115997	1.038547
`ibm_marrakesh`	S1f	0.069531	0.076432	0.077617	1.040666

TABLE VII. Dedicated readout-mitigation and local-folding ZNE lane for the S1 direct-observable stack. The table reports mean absolute feedback-minus-control separation across each lane's preregistered Pauli channels or policy-channel rows. The readout confusion matrices are well-conditioned in this lane; mitigation does not remove the observed separations.

arm contains mid-circuit measurement, reset, and classical conditioning; the reducer folds invertible quantum gates while preserving the dynamic-circuit control-flow operations.

Table VII changes the limitation boundary. The dominant direct-observable separations are not explained away by a poorly conditioned readout inversion: the full-basis readout matrices have condition numbers near one, and the readout-mitigated ZNE values remain close to the unmitigated ZNE values. The ZNE stress test also does not convert the result into a controller-promotion claim. On Marrakesh, all S1b–S1f mean absolute separations increase under the linear ZNE fit. On Fez, S1b–S1e also increase, while S1f decreases. The conservative interpretation is that the feedback/control response is not a simple removable readout artefact, but the sign and localisation remain backend- and lane-dependent.

IX. PROVIDER CONSTRAINTS AND CONTROLLER REDESIGN

The revised S1 analysis records provider constraints as experimental variables. Both Fez and Marrakesh reported 4 ns values for `dt` and `dtm` in the runtime backend snapshot. The target timing constraints exposed by the Qiskit target were acquire alignment 1, pulse alignment 1, granularity 1, and minimum pulse length 2. The provider API used for this manuscript did not expose a per-branch classical-conditioning latency or reset fidelity value. We therefore do not quote a fabricated latency or reset number. Instead, the measured

Backend	Mean readout error	Mean retention	Max leakage	Multi-bit proxy
`ibm_fez`	0.010895	0.974121	0.040430	0.000293
`ibm_marrakesh`	0.006378	0.985718	0.023242	0.000073

TABLE VIII. Provider and calibration constraints for the S1 readout/ZNE layouts. The mean readout error is taken from the live backend properties on the selected four-qubit layout. Mean retention, max leakage, and multi-bit proxy come from the dedicated final-readout calibration circuits, averaged across S1b–S1f. The multi-bit proxy is the observed probability of a calibration outcome differing from the prepared state in at least two final-readout bits; it is not mid-circuit crosstalk tomography.

dynamic-circuit cost boundary is the transpiled control-flow structure: on the same-layout Marrakesh repeat, the feedback arm has maximum transpiled depth 715 with three `if_else` operations, six measurements, and three resets; the matched control has maximum transpiled depth 696 with six measurements and three resets. On the Fez primary replication, the corresponding depths are 717 and 684.

These constraints motivate a concrete next controller rather than a generic "try again" loop. The next S1 controller should replace the fixed correction angle with a monitor-conditioned signed gain: the correction amplitude should depend on the monitor outcome and on a preregistered confidence map derived from the calibration lane. It should add a two-monitor variant that separately observes the edge associated with the `YYI` response and an orthogonal control edge, so that a single monitor cannot ambiguously couple policy signal and readout pathway. Finally, it should test an explicit topology variant that targets the `YYI`-sensitive edge seen on Marrakesh while carrying the same polarity, cross-quadrature, and population controls used here. The acceptance gate for such a controller should require the primary binary proxy, the preregistered `YYI` direct observable, and a same-layout repeat to agree in sign before any controller-promotion claim is made.

X. REPRODUCIBILITY ARTEFACTS

The hardware evidence is committed as raw-count and analysis artefacts under `data/s1_feedback_loop/`. Table IX lists the primary analysis artefacts and SHA256 digests used for this manuscript. The exact hardware job identifiers are retained in the JSON files and

Run	Analysis artefact	SHA256
S1	`s1_feedback_analysis_summary_ibm_kingston_20260520T123941Z.json`	`dcfae0016806ac3f705ce9ecad62f1ab47d0d3986775170a7f15b380f38fcf24`
S1b	`s1b_xy_observable_analysis_ibm_kingston_20260520T130238Z.json`	`91fdedc8043d25c68f95a57c21be7ed828c7e83f34af48f11010fae686e4b152`
S1c	`s1c_xy_observable_analysis_ibm_kingston_20260520T132455Z.json`	`ade9a407dec9759ed4fe789f738d4ff3560f2ecf73561f6f07435c4bab9f50dd`
S1d	`s1d_xy_observable_analysis_ibm_kingston_20260520T134614Z.json`	`6e6bc2f513f23633c72756016392b58cbf140d89c6f9cc7fdd07880c66b61c76`
S1e	`s1e_xy_observable_analysis_ibm_kingston_20260520T140919Z.json`	`4f7e106e2af8e42aa74c9c07d66032bc76a78a936c499370c56fcef8c32ad522`
S1f	`s1f_xy_observable_analysis_ibm_kingston_20260520T142213Z.json`	`17b37753460137c796a5fb8c52784ce694d7ce4ffb63e55c92853abe8f91313e`

TABLE IX. Committed analysis artefacts supporting the S1–S1f manuscript.

in the Markdown draft notes.

The Fez and Marrakesh replication artefacts follow the same naming scheme under `data/s1_feedback_loop/`, with backend names `ibm_fez` and `ibm_marrakesh` and timestamps 20260520T145159Z–20260520T150249Z for Fez and 20260520T150509Z–20260520T151450Z for Marrakesh. Their analysis SHA256 digests are recorded in the JSON sidecars and in the repository history. The dedicated readout/ZNE lane uses timestamps 20260520T153231Z–20260520T153820Z for Fez and 20260520T153231Z–20260520T153847Z for Marrakesh. Its job IDs are `d86t84p789is739052r0`, `d86t8ftg7okc73eljf10`, `d86t8qp789is739053og`, `d86t9hop0eas73dleo20`, and `d86tarlg7okc73elji20` on

Fez, and d86t84op0eas73dlem4g, d86t8gtg7okc73eljf20, d86t8stg7okc73eljfig, d86t9ptg7okc73eljgo0, and d86tb2dg7okc73eljiag on Marrakesh. The same-layout Marrakesh repeat analysis artefact is s1_feedback_analysis_summary_ibm_marrakesh_20260520T155826Z.json; the provider-constraint artefact is s1_dynamic_circuit_constraints_20260520T160234Z.json.

XI. CLAIM BOUNDARY

The supported claims are deliberately narrow. First, the S1 monitored-feedback payload and matched open-loop control executed on three IBM Heron backends, with a negative primary result on Kingston and positive primary results on Fez and Marrakesh. Second, direct XY-sector extensions show that feedback/control differences can be channel-structured and policy-sensitive. Third, the S1d/S1e/S1f YYI story is backend-dependent: Kingston loses the positive YYI response in S1f, Fez is mostly negative in that column, and Marrakesh remains positive. The controller is therefore not promoted as backend-general.

The paper does not claim quantum advantage, backend-general feedback control, full real-time control beyond provider-supported dynamic circuits, scalable control, or larger-system synchronisation. It also does not claim that the observed response is isolated from calibration drift, layout-dependent coherent error, readout effects, or finite-shot variation. The result is a bounded hardware-control measurement and a falsifier for the tested feedback law as a universal policy.

At the time of execution, authenticated live QPU access for this campaign was limited to IBM Quantum hardware and to the three Heron backends reported here. The framework contains provider-neutral execution and packaging surfaces for other gate-model, trapped-ion, neutral-atom, photonic, annealing, analogue, and simulator targets, but those software routes are not evidence of live non-IBM execution. We plan to rerun the same preregistered S1 payloads on additional hardware families when compute allocations and provider access become available; until those reruns exist, no cross-provider or cross-technology controller claim is made.

The original S1b–S1f artefacts remain unmitigated raw hardware evidence. The readout/ZNE lane is a separate preregistered execution, not a post hoc correction of those earlier jobs. Mid-circuit measurement and feedforward readout mitigation is an active technical area [14]; the present reducer only mitigates the final three-bit readout distribution and only stress-tests gate noise through local folding of invertible quantum operations.

XII. DISCUSSION

The central scientific value of S1 is methodological rather than promotional. The primary endpoint fails on Kingston but succeeds on Fez and Marrakesh, so a single-backend story would be misleading in either direction. The extensions then add mechanism information that a binary endpoint alone would hide. S1b and S1c show backend-dependent XY-sector response; S1d and S1e show that the policy sweep can repeat a `YYI` feature inside a backend window; and S1f shows that the same feature is not stable across backends. The dedicated readout/ZNE lane strengthens the negative claim boundary: the direct-observable separations survive full final-readout calibration and do not behave like a uniform monotone gate-noise artefact.

This sequence is useful for future control design because it converts a simple negative result into design constraints. Any next-generation S1 policy must preregister direct Pauli observables rather than rely only on the binary proxy. It must include a matched open-loop arm, policy-polarity controls, cross-quadrature controls, and a population control. It should also reserve a same-layout repeat in a later calibration window before promoting a controller. It should additionally treat provider constraints as first-class experimental variables: mid-circuit measurements and resets add latency, classical conditioning changes the transpiled control-flow path, reset fidelity can couple the monitor qubit back into subsequent rounds, and measurement crosstalk can affect both the monitored branch and the final readout. Deterministic error-suppression and calibration-aware workflows can improve near-term algorithm performance [15], but S1 shows that the feedback law itself still needs a matched-arm falsification gate. Under these standards, S1 is not archive clutter: it is a claim-bound falsification and localisation paper for dynamic-circuit feedback on current IBM hardware.

Monitor layer	Mid-circuit measurement with reset-fidelity and latency metadata recorded before execution.
Adaptive policy	Outcome-conditioned gain and polarity selected from a preregistered finite policy table.
Control arm	Matched open-loop body with the same routed circuit family and shot budget.
Observable gate	Binary proxy plus direct `XX`/`YY`, cross-quadrature, and population controls.
Promotion gate	Same-layout repeat and readout/ZNE lane before any backend-specific controller promotion.

FIG. 2. Concrete redesign target implied by S1. The next controller should make provider constraints and direct observables part of the preregistered policy, rather than treating them as post-run diagnostics.

XIII. CONCLUSION

The S1 programme tests monitored feedback at the smallest scale where IBM dynamic circuits can implement an active monitor-and-correct loop with a matched open-loop control. The preregistered binary synchronisation endpoint is negative on Kingston but positive on Fez and Marrakesh. Direct-observable extensions reveal structured, policy-sensitive feedback/control differences, and a dedicated Fez/Marrakesh readout/ZNE lane shows that those differences are not erased by full final-readout mitigation or by local-folding extrapolation. The final quadrature checks still do not support a backend-general positive controller mechanism. The appropriate publication claim is therefore precise: S1 establishes a reproducible paired-arm method and a hardware-control boundary for Kuramoto–XY feedback, while rejecting backend-general promotion of the tested controller. It also establishes that monitored feedback on IBM dynamic circuits can exhibit backend-dependent response modes that are not captured by a saturated binary synchronisation proxy, which is a concrete methodological constraint for future controller design.

---

[1] Y. Kuramoto, in *International Symposium on Mathematical Problems in Theoretical Physics* (Springer, 1975) pp. 420–422.

[2] J. A. Acebrón, L. L. Bonilla, C. J. P. Vicente, F. Ritort, and R. Spigler, Reviews of Modern Physics **77**, 137 (2005).

[3] J. Preskill, Quantum **2**, 79 (2018), arXiv:1801.00862 [quant-ph].

[4] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th ed. (Cambridge University Press, Cambridge, 2010).

[5] M. Sotek, Backend-sensitive hardware observation of parity-sector and excitation-number correlated decoherence asymmetry in the xy hamiltonian (2026), preprint manuscript and public repository artefacts.

[6] M. Sotek, Fisher-information-inspired collective feedback in kuramoto–xy hamiltonians: exact small-system structure and a hardware falsification (2026), preprint manuscript and public repository artefacts.

[7] M. Sotek, Reduced-pauli entanglement checks for dla-sector leakage mechanisms on ibm heron hardware (2026), preprint manuscript and public repository artefacts.

[8] H. M. Wiseman and G. J. Milburn, *Quantum Measurement and Control* (Cambridge University Press, Cambridge, 2010).

[9] A. D. Córcoles, M. Takita, K. Inoue, S. Lekuch, Z. K. Minev, J. M. Chow, and J. M. Gambetta, Physical Review Letters **127**, 100501 (2021), arXiv:2102.01682 [quant-ph].

[10] IBM Quantum, Ibm quantum platform (2026), `https://quantum.ibm.com/`.

[11] A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon, S. Martiel, P. D. Nation, L. S. Bishop, A. W. Cross, B. R. Johnson, and J. M. Gambetta, Quantum computing with Qiskit (2024), arXiv:2405.08810 [quant-ph].

[12] J. M. Koh, S.-N. Sun, M. Motta, and A. J. Minnich, Nature Physics **19**, 1314 (2023).

[13] M. Foss-Feig, A. Tikku, T.-C. Lu, K. Mayer, M. Iqbal, T. M. Gatterman, J. A. Gerber, K. Gilmore, D. Gresh, A. Hankin, N. Hewitt, C. V. Horst, M. Matheny, T. Mengle, B. Neyenhuis, H. Dreyer, D. Hayes, T. H. Hsieh, and I. H. Kim, Experimental demonstration of the advantage of adaptive quantum circuits (2023), arXiv:2302.03029 [quant-ph].

[14] J. M. Koh, D. E. Koh, and J. Thompson, PRX Quantum **7**, 010317 (2026), arXiv:2406.07611 [quant-ph].

[15] P. S. Mundada, A. Barbosa, S. Maity, Y. Wang, T. M. Stace, T. Merkh, F. Nielson, A. R. R. Carvalho, M. Hush, M. J. Biercuk, and Y. Baum, Physical Review Applied **20**, 024034 (2023), arXiv:2209.06864 [quant-ph].