

Scalar tracking under burst noise

A reproducible comparison of state-dependent and fixed feedback

Aboudy Keddeh
Systems Architect

Keddeh Systems · Computational control research
Research draft · Version 1.1 · 14 October 2026

Abstract

This study evaluates whether a specified state- and noise-dependent feedback gain improves scalar set-point tracking during a temporary increase in additive process noise. It preserves an existing discrete scalar experiment and formalises its comparison, statistical unit, outcome and limitations. Two policies receive the same Gaussian innovations over 1,200 samples. The primary outcome is the mean paired difference in full-run root mean square error (RMSE) across 200 seeds. In the legacy block, adaptive minus fixed RMSE is -0.012533 , with a paired 95% Student-t interval of $[-0.016957, -0.008109]$. The adaptive policy wins only 87 of 200 pairs: lower mean loss does not imply more frequent wins. A separate seed block gives -0.020948 $[-0.027281, -0.014615]$. The direction reverses without a noise burst and against a stronger fixed baseline. The contribution is a reproducible, bounded comparison and an executable decision aid for further controller research. It is not evidence of universal adaptive superiority or a deployed synchronisation mechanism.

Interactive companion

The accompanying [study.html](#) provides protocol execution, exploratory controls and CSV exports. This PDF presents the formal study and static seed-297 figures. Model, protocol, seed blocks and numerical results are unchanged in this editorial revision.

© 2026 Aboudy Keddeh. Published by Keddeh Systems.

All rights reserved. KEDDEH-RESERVED-1.0. No licence is granted by this document. Reuse, redistribution and implementation require separate written permission, subject to applicable law. Third-party material retains its own rights.

1. Research question and useful decision

Does the supplied adaptive gain reduce expected full-run scalar tracking RMSE relative to fixed gain 0.18 s^{-1} in this specified burst-noise workload? An engineer can use this experiment to decide whether this particular gain law deserves further study, which disturbance conditions expose failure, and whether a claimed improvement survives a stronger comparator. A negative answer remains a useful result. Application to clock recovery, distributed synchronisation or a physical controller requires a separate plant, measurement and actuation model; those sectors are not validated here.

Let $D_r = \text{RMSE}(\text{adaptive}, r) - \text{RMSE}(\text{fixed}, r)$, and $\Delta = E[D_r]$. The directional hypothesis is $H_1: \Delta < 0$, against $H_0: \Delta \geq 0$. A two-sided 95% interval entirely below zero is evidence favouring the adaptive law in this model; entirely above zero favours fixed gain; crossing zero is inconclusive. This criterion was documented after the legacy results existed. It is a retrospective formalisation, not preregistration. No minimum practically important difference has been justified, so the criterion cannot authorise deployment.

2. Scope, objects and units

- x is a real-valued, dimensionless state; $r = 0.297$ is the inherited numerical set-point; $x_0 = 0.45$. Neither x nor r is assigned radians or a physical measurement unit.
- t is model time in seconds; $\Delta t = 0.01 \text{ s}$. Samples are $t_i = i\Delta t$ for $i = 0, \dots, 1199$: the final sample is 11.99 s , with $1,199$ updates. “12-second window” refers to $[0, 12)$, not a sample at 12 .
- Feedback gain k has units s^{-1} . Noise amplitude σ has units $\text{state}/\sqrt{\text{s}}$; z is a standard normal variate. RMSE has state units. The target and initial state are fixed across comparisons.
- The legacy score $C = 1 - |x - r|$ is a dimensionless proximity score that may be negative. Its 0.75 line is equivalent to $|x - r| = 0.25$. It measures neither circular phase coherence nor agreement among oscillators.
- The native Keddeh definition of 0.297 as a jitter function is a separate research object. Its implementation is not supplied by assigning the same number to this scalar target.

3. Mechanism and justified abstraction

The fixed policy is linear mean reversion toward a constant set-point. The adaptive policy changes restoring strength as a function of the prior error and a scenario-supplied noise amplitude. This isolates the feedback-law comparison from differences in disturbances. It is a minimal computational control model, not an identified physical plant.

$$x_i = x_{i-1} - k_i(x_{i-1} - r)\Delta t + \sigma_i/\Delta t z_i$$

$$k_i = 0.18 \text{ s}^{-1}$$

$$k a_i = \text{clip}(0.12 + 1.2|x_{i-1} - r|^2 - 0.25\sigma_i, 0.04, 0.80) \text{ s}^{-1}$$

The numerical coefficients inherit the normalised state convention; dimensionally, the quadratic coefficient is 1.2 s^{-1} per squared state unit and the noise coefficient is $0.25/(\text{state} \cdot \sqrt{\text{s}})$. $\sigma_i = 0.25$ for destination indices $400 \leq i < 800$ and 0.05 otherwise. This is the preserved discrete convention: the update ending at 4.00 s uses burst noise. A different continuous-time boundary convention would be a model revision.

Euler–Maruyama contributes the $\sqrt{\Delta t}$ scaling of the Gaussian increment [1]. Both policies use the exact same z_i within a trial; different trial seeds provide separate pseudo-random streams. The browser uses Mulberry32 and Box–Muller, not NumPy or the Python standard-library stream. The adaptive law knows σ exactly. It does not estimate noise from measurements, and reducing k does not reduce the additive noise term. The coefficient 0.297 is retained for historical reproducibility; this experiment contains no test establishing it as uniquely optimal.

For the fixed policy, $a = 1 - k\Delta t$ and error $e = x - r$ yield exact discrete moments $m_i = a m_{i-1}$ and $v_i = a^2 v_{i-1} + \sigma^2 \Delta t$, from $m_0 = 0.153$, $v_0 = 0$. The default expected mean squared error is 0.05936275688459807; its square root is 0.24364473498230588. This is $\sqrt{E[\text{MSE}]}$, not $E[\text{RMSE}]$. It provides an analytic check independent of individual random paths. The fixed recurrence has $|a| < 1$ throughout the allowed gain range. This observation is not a proof of every property of the nonlinear adaptive process.

4. Reproducible evaluation protocol

1. Freeze the published model, protocol version and defaults. The primary fixed gain is 0.18, baseline $\sigma = 0.05$ and burst $\sigma = 0.25$. Do not select a favourable seed after inspecting its plot.
2. Run 200 paired trials at integer seeds 297–496 for the historical replication. Separately run 1000297–1000496 for the additional block. Seeds are reproducibility identifiers, not observations sampled from a real-world population.
3. Calculate $\text{RMSE} = \sqrt{[(1/1200) \sum_i (x_i - r)^2]}$ for each policy. Include the initial sample. The unit of statistical replication is the trajectory pair; the 1,200 temporally correlated samples are not independent replicates.
4. Calculate the sample mean of D , its sample standard deviation with divisor 199, and standard error $\text{SD}/\sqrt{200}$. Use $\text{mean}(D) \pm 1.971956544 \cdot \text{SE}$ for the paired two-sided 95% t interval, $\text{df} = 199$ [2]. Coverage is approximate for these nonlinear loss distributions and relies on adequate pseudo-random streams.
5. Report every pair, both noise blocks, all window errors and all six scenario arms. No failed or unfavourable seed is removed. Invalid parameters stop a run with an error.
6. Run the declared sensitivity arms: no process noise; steady $\sigma = 0.05$; stronger burst $\sigma = 0.50$; fixed comparator gain 0.10; and fixed comparator gain 0.40. Each changes only the listed factors. Their intervals are pointwise and exploratory, not multiplicity-adjusted evidence of a global best law.
7. Keep edited-parameter sandbox runs visibly separate from the frozen protocol. Export the protocol and results together so conclusions cannot be detached from settings.

Secondary diagnostics are the [0,4), [4,8), [8,12) window RMSE; the frequency of lower adaptive RMSE; integrated squared feedback $\sum [k_i(x_{i-1} - r)]^2 \Delta t$, in state^2/s ; and sampled time outside $|\text{error}| > 0.25$, in seconds, counted on updates 1–1199. The latter is a discrete occupancy estimate, not continuous crossing time or a process-spawning count. These diagnostics do not replace the primary outcome. The no-noise arm is deterministic; repeated seeds contribute no genuine sampling uncertainty there.

5. Executed results and interpretation

Table 1. Mean paired RMSE differences by scenario. Negative values favour adaptive feedback; positive values favour fixed feedback. Units: state. Each block contains 200 trajectory pairs.

Scenario	Legacy mean Δ	Replication mean Δ
Primary burst	-0.012533	-0.020948
No process noise	+0.010127	+0.010127
Steady noise, no burst	+0.009722	+0.009687
Stronger burst	-0.091810	-0.115550
Fixed gain 0.10	-0.047566	-0.059767
Fixed gain 0.40	+0.039522	+0.037915

The adaptive law reduces mean RMSE in the primary scenario but loses more individual trials than it wins in the legacy block. These are different summaries: larger improvements in some trajectories can outweigh smaller degradations in many others. The result must not be rewritten as “adaptive is usually better.” The no-noise and steady-noise controls expose slower return under the lower adaptive gain near the target. Raising the fixed gain to 0.40 reverses the mean advantage. Because this model imposes no feedback cost, actuator limit, delay or measurement noise, stronger restoring action can benefit tracking without paying realistic penalties. A fair deployment-oriented optimisation would need those costs and constraints explicitly.

Recorded numbers are local computational observations, not measurements of deployed hardware. The JSON evidence records the execution environment and source digests. No independent assessor has reproduced or approved the study.

6. Limitations and next falsifiable experiments

- **Baseline dependence:** compare against a gain selected using a separate training seed set, then evaluate once on untouched seeds under a stated feedback-cost budget. Current 0.40 sensitivity already defeats a universal superiority claim.
- **Known noise:** define a causal estimator using only available measurements, then test estimation lag and bias against the present oracle- σ condition.
- **Discretisation:** this release claims results for $\Delta t = 0.01$ only. A future refinement study must couple Brownian increments across 0.01, 0.005 and 0.0025 grids and preserve physical horizon and event boundaries before claiming continuous-time accuracy.
- **Physical transfer:** identify a plant and measurement units; add sensor noise, delay, saturation and feedback cost; then compare against measured traces. No hardware or native-runtime benefit follows from the present simulation alone.
- **Phase and jitter:** a phase study requires circular phase definitions, oscillator dynamics and phase-lock outcomes; a native 0.297 study requires its executable jitter operator and contract. Neither can be inferred from this scalar score.

7. Source lineage, references and review

The retained predecessor is SERVERspace mining-engine/simulate_phase_noise_sde.py, cell fd48343a, SHA-256 89cdf9faf4dff4e56f5516877094bd4728812b0597a59ff2f3ea38778933241e. It compares fixed-gain scalar noise tracks and labels some results as spawning without executing a spawn. Its gain, sample interval, horizon and random stream differ from this study; it is lineage, not numerical equivalence. The browser model source SHA-256 a1eeaa7fc879a8b1e2a7da5e439c0ff25a30b3679c0b3971e25240ba36b1eb91 is preserved unchanged. The adaptive law is inherited from the inspected browser source. Those inspected files establish the computational formula and predecessor; this study has not assessed derivation evidence elsewhere in the research corpus.

[1] D. J. Higham (2001), “An Algorithmic Introduction to Numerical Simulation of Stochastic Differential Equations,” SIAM Review 43(3), 525–546. DOI: [10.1137/S0036144500378302](https://doi.org/10.1137/S0036144500378302). Supports the numerical method, not this gain law or the choice 0.297.

<https://epubs.siam.org/doi/10.1137/S0036144500378302>

[2] NIST/SEMATECH e-Handbook of Statistical Methods, [Analysis of paired observations](#) and [Confidence intervals for differences between means](#), accessed 14 October 2026. Supports paired-difference inference; it does not certify these results.

<https://www.itl.nist.gov/div898/handbook/prc/section3/prc311.htm>

<https://www.itl.nist.gov/div898/handbook/prc/section3/prc312.htm>


[3] Keddeh Research and Formal Publishing Standard, KEDDEH-STD-PUBLISHING-001 v1.0, and Document Template, KEDDEH-TPL-PUBLISHING-001 v1.0. Consulted for document structure, attribution, rights and review fields. No independently certified template compliance is asserted.

Appendix A Static trajectory figures

Default exploratory record: seed 297, fixed gain 0.18 s^{-1} , baseline $\sigma = 0.05$, burst $\sigma = 0.25$. Shading marks 4–8 s. These model outputs illustrate one trajectory pair; inference is based on the paired protocol, not this single path.


P1 Scalar state trajectories

Fixed · Adaptive | State units; dashed reference = 0.297.




P2 Legacy proximity score

Fixed · Adaptive | Dimensionless proximity score, not phase coherence; dashed reference = 0.75.



P3 Adaptive feedback gain

Adaptive gain | Gain in s^{-1} ; dashed reference = fixed gain 0.18 s^{-1} .



Publication record and provenance

Publication record	Value
Document identifier	KEDDEH-RES-SCALAR-TRACKING-001
Content version / issue date	1.1 / 14 October 2026
Classification / status	Public research / Draft
Research direction	Aboudy Keddeh, Systems Architect
Publisher / rights holder	Keddeh Systems / Aboudy Keddeh
Licence profile	KEDDEH-RESERVED-1.0
Independent assessment / approval	Not recorded / Not recorded

Contributions and review

Research direction: Aboudy Keddeh. Formalisation, implementation, computational evaluation and publication layout: AI-assisted with OpenAI. Independent research assessment and approval have not been recorded. Layout verification is not scientific review or certification.

Corporate context and rights

The Layna Company Pty Limited is recorded in the supplied publishing standard as parent-company context; that record does not establish copyright assignment or legal publishing authority. The publisher and rights holder are recorded separately above. Authorship does not itself assign copyright or grant patent, trademark or software licences.

Revision history

Version 1.0 formalised the preserved browser model, paired t interval, additional seed block, sensitivity arms, secondary diagnostics and analytic fixed-policy check. Version 1.1 limits the derivation statement to inspected sources and revises the static publication presentation. Model behaviour, random-number generator, protocol and numerical results are unchanged.

Figure provenance

Figure records: P1, P2 and P3. Creator: AI-assisted Keddeh plotting implementation. Role: computational model output. Source: preserved model.mjs, seed-297 trajectory.json; exact values accompany this carrier. Rights basis: © 2026 Aboudy Keddeh, KEDDEH-RESERVED-1.0; no third-party artwork. Figure version: 1.0. Relationship: explanatory trajectory projection of KEDDEH-RES-SCALAR-TRACKING-001 v1.1.

Carrier record

A4 PDF with a repeating masthead and automatic Page X of Y. Typography uses Liberation Sans, an Arial-metric-compatible substitute, with DejaVu Sans for scientific glyphs. The text-only masthead follows the retained publishing standard; original template artwork was not available for this carrier. Exact Arial typography and artwork reproduction are therefore not asserted. The accompanying manifest records source identities and separate carrier hashes.