Token–Layer Activation Event Cascades in LLMs: Rate-Matched Connectivity under Gain Scaling

Ali Uyar
Independent Researcher

Abstract

Activation patterns in large language models are often studied via static sparsity or outlier magnitudes, but less is known about how *activation events* form connected cascades across token positions and depth. We model gated-MLP activations as a token-by-layer *event lattice* by standardizing a fixed internal tensor and thresholding it into sparse binary events. This enables connected-component analysis that yields avalanche-like cascades on a two-dimensional lattice. We introduce directional branching metrics that decompose local propagation into token-direction and depth-direction components, and we evaluate a low-compute gain intervention that scales each layer's MLP residual contribution at inference time.

To isolate connectivity effects from trivial rate effects, we design two strong controls: (i) *per-layer rate-matched thresholds* that equalize marginal event rates across gains, and (ii) a *marginals-preserving raster shuffle* that permutes token order within each layer while preserving per-layer event counts exactly. We avoid equating heavy-tailed component statistics with criticality by treating tail fits as descriptive diagnostics and using multiple falsifiable signatures. Finally, we calibrate a gain $g^\star$ on Dataset A using a mechanistic criterion based on branching, then evaluate the same $g^\star$ unchanged on Dataset B and ARC multiple-choice, comparing against the baseline $g{=}1$.

In the provided 7B runs, total branching b_{tot} varies with gain under rate matching (Table 8 and Figure 3) and the null-controlled residual Δb_{tot} remains non-zero (Figure 4). This framework supports both positive and negative transfer outcomes under strong controls by reporting mechanistic signatures alongside task metrics.

1 Introduction

Large language models (LLMs) exhibit rich internal activation structure, including sparsity patterns, context-dependent gating, and occasional extreme outliers. However, many analyses treat activations as independent samples or focus on layerwise distributions, which can obscure *connectivity in token-by-depth space*. Motivated by cascade analyses in neuroscience [1], we ask a simple mechanistic question: *when we threshold a fixed internal tensor into sparse activation events, do those events form connected cascades across token position and layer depth, and how does that connectivity change under controlled perturbations?*

We propose a concrete measurement construct: a token-by-layer event raster derived from standardized gated-MLP activations. This event lattice supports connected-component analysis ("avalanches") and directional branching metrics that quantify local propagation along token and depth directions. We then study a global *gain* intervention that scales the MLP residual contribution by a scalar g at inference time. Because g trivially rescales activations and thus event *rates*, we include reviewer-proof controls that (i) match marginal event rates per layer across g and (ii) destroy within-layer temporal structure while preserving marginals exactly. We use these controls to probe quasi-critical-like regimes without claiming a phase transition or brain equivalence.

Contributions.

- **Event lattice construct.** A token-by-layer binary event raster derived from a fixed standardized gated-MLP tensor, enabling avalanche-style connected-component analysis in transformers.
- **Directional branching decomposition.** Metrics b_{time}, b_{depth}, and b_{tot} that quantify local propagation along token and depth directions, plus a marginals-controlled residual Δb.
- **Reviewer-proof controls.** Per-layer rate-matched thresholds $\tau_\ell(g)$ and a post-hoc within-layer time-permutation null that preserves marginals exactly.
- **Mechanistic calibration and transfer test.** Calibrate $g^\star$ on Dataset A by minimizing $|b_{\mathrm{tot}}(g)-1|$ over a fixed gain grid (allowing boundary solutions), then evaluate unchanged on Dataset B and ARC multiple-choice against $g{=}1$.
- **Conservative quasi-critical probing.** Multiple signatures used as falsifiable probes; heavy-tail fits are reported as descriptive only [2, 3].

2 Related work

We draw inspiration from neuronal avalanche analyses and branching measures [1] while avoiding any claim of brain equivalence. We also follow statistical cautions that power-law-like tails alone are insufficient evidence of criticality

[2, 3]. In deep learning, "edge-of-chaos" style analyses study signal propagation and criticality-like regimes in random networks and deep models [5–7]; our work instead operates at inference time on trained transformers and uses rate-matched binary event rasters. Finally, we focus on gated MLP activations (common in modern LLMs; [8]) and distinguish event connectivity from activation-magnitude outlier studies [4].

3 Method

3.1 Token-by-layer event lattice

Let a transformer with L blocks process a token sequence of length T. For each token position $t \in \{1, \ldots, T\}$ and layer $\ell \in \{1, \ldots, L\}$, we extract a fixed internal gated-MLP tensor $u_{t,\ell,i}$ (indexed by MLP hidden dimension i) at a specified hookpoint (pre-down-projection for gated MLPs). We standardize using per-layer moments (μ_ℓ, σ_ℓ) estimated on a fixed calibration slice:

$$z_{t,\ell,i} = \frac{u_{t,\ell,i} - \mu_\ell}{\sigma_\ell + \varepsilon}.$$

We define spike events in two ways:

$$s^{(+)}_{t,\ell,i} = \mathbb{1}[z_{t,\ell,i} > \tau_\ell] \quad \text{and} \quad s^{(\pm)}_{t,\ell,i} = \mathbb{1}[|z_{t,\ell,i}| > \tau_\ell].$$

We aggregate spikes into an event-count field and a binary occupancy field:

$$A_{t,\ell} = \sum_i s_{t,\ell,i}, \qquad X_{t,\ell} = \mathbb{1}[A_{t,\ell} > 0].$$

The binary raster $X \in \{0,1\}^{T \times L}$ is the event lattice used for connected components.

3.2 Rate-matched thresholds

For each gain g and layer ℓ, we choose $\tau_\ell(g)$ so that the marginal event rate matches a target $r^\star$:

$$\mathbb{E}_{t,i}\left[s_{t,\ell,i}\right] \approx r^\star.$$

Operationally we compute $\tau_\ell(g)$ as a per-layer quantile of $z_{t,\ell,i}$ on the calibration slice (separately for $s^{(+)}$ and $s^{(\pm)}$). Rate-matching success is verified by the maximum absolute rate error across layers (Figure 2).

3.3 Avalanches and connected components

We define avalanches as connected components of active sites in the token-by-layer lattice under a 4-neighborhood adjacency (time and depth moves). Each component C has size $S(C) = \sum_{(t,\ell)\in C} A_{t,\ell}$, duration in tokens (span along t), and depth span along ℓ.

3.4 Directional branching metrics

For each active site (t, ℓ) we count forward-neighbor activations in time $(t+1, \ell)$ and depth $(t, \ell+1)$, normalized by the number of possible forward neighbors. Aggregating yields:

$$b_{\text{time}}, \quad b_{\text{depth}}, \quad b_{\text{tot}} = b_{\text{time}} + b_{\text{depth}}.$$

We also compute a susceptibility proxy χ (variance-based; see Appendix) and a descriptive crackling exponent fit on the avalanche size distribution (reported with bootstrap confidence intervals). We treat these as *signatures* rather than proofs of criticality.

3.5 Marginals-preserving null and Δb

To isolate connectivity from marginals, we construct a within-layer time-permutation null: for each layer ℓ, apply a permutation π_ℓ to token indices, permuting $A_{t,\ell}$ across t while preserving each layer's multiset of counts exactly. We compute branching on this permuted raster to obtain $b_{\cdot,\text{perm}}$, then define:

$$\Delta b_\cdot = b_\cdot - b_{\cdot,\text{perm}}.$$

A non-zero Δb indicates structure beyond marginals.

3.6 Gain intervention and mechanistic $g^\star$

We modify each transformer block's residual update to scale the MLP branch:

$$h_{\ell+1} = h_\ell + \text{Attn}_\ell(h_\ell) + g \cdot \text{MLP}_\ell(h_\ell).$$

On Dataset A, for each (spike definition, target rate) condition, we select

$$g^\star = \arg\min_{g \in \mathcal{G}} |b_{\text{tot}}(g) - 1|.$$

We then evaluate the same $g^\star$ unchanged on Dataset B and ARC multiple-choice, comparing against g=1.

4 Experiments

4.1 Model and datasets

We evaluate two 7B checkpoints from the same model family (Qwen2.5-7B-Instruct and Qwen2.5-7B base) using three datasets:

- **Dataset A:** a fixed slice of Wikitext-103 validation (mechanistic calibration and signatures).
- **Dataset B:** a fixed slice of C4-en validation (transfer evaluation).
- **ARC multiple-choice:** ARC-Challenge (task metric evaluation).

All experiments are inference/analysis heavy and fit within a single-GPU budget by limiting the number of sequences and gain conditions in the published run.

4.2 Conditions and controls

We evaluate:

- Two spike definitions: one-sided $s^{(+)}$ and two-sided $s^{(\pm)}$.
- Target marginal rates: $r^\star \in \{1, 2, 4, 8\} \times 10^{-5}$.
- Gain grid $\mathcal{G} = \{0.70, 0.80, 0.85, 0.90, 0.95, 1.00, 1.05, 1.10, 1.15, 1.20, 1.30\}$.

For each g we rate-match $\tau_\ell(g)$ and evaluate two within-layer nulls: within-layer time permutation and within-layer circular shift. We report mechanistic metrics (branching, Δb, χ) and task metrics (perplexity, ARC accuracy) with bootstrap confidence intervals.

5 Results

5.1 Event rasters and rate matching

Figure 1 shows a representative token-by-layer event raster. Rate matching succeeds across all Dataset A conditions, with maximum absolute per-layer rate error below the fixed tolerance (Figure 2 and Table 8). This control is required to interpret any gain-dependent changes in branching.

5.2 Gain affects branching under rate matching

Figure 3 plots b_{time}, b_{depth}, and b_{tot} versus gain for both spike definitions at a representative target rate. Full results across all target rates are reported in Table 8. In the provided runs, b_{tot} varies with gain under rate matching (Table 8). The null-controlled residual Δb_{tot} remains non-zero under a strong within-layer permutation null (Figure 4; Table 8). This indicates connectivity structure beyond marginals even after rate matching.

5.3 Mechanistic $g^\star$ selection

Figure 5 visualizes $g^\star$ selection by minimizing $|b_{\text{tot}}(g) - 1|$ on Dataset A. The selected $g^\star$ differs across the eight (spike definition, target rate) conditions and can fall on the boundary of the gain grid (Table 1), illustrating why an explicit cross-dataset test is necessary.

5.4 Cross-dataset evaluation: negative transfer

We evaluate the mechanistically selected $g^\star$ unchanged on Dataset B and ARC multiple-choice, comparing against $g{=}1$ (Figures 6 and 7). Figures visualize one representative condition with confidence intervals; Tables 2 and 3 report all spike-definition/target-rate conditions with bootstrap confidence intervals. Replication across base vs instruct checkpoints is summarized in Appendix Table 7.

Spike def	Target rate	$g^\star$	$b_{\text{tot}}(g{=}1)$	$b_{\text{tot}}(g^\star)$	$\Delta b_{\text{tot}}(g{=}1)$	$\Delta b_{\text{tot}}(g^\star)$
one-sided (+)	1e-05	0.700	0.593	0.599	0.263	0.266
one-sided (+)	2e-05	0.700	0.823	0.833	0.260	0.271
one-sided (+)	4e-05	1.300	1.102	1.089	0.226	0.210
one-sided (+)	8e-05	1.300	1.395	1.387	0.158	0.145
two-sided ($|\cdot|$)	1e-05	0.850	0.627	0.629	0.295	0.300
two-sided ($|\cdot|$)	2e-05	1.000	0.853	0.853	0.305	0.305
two-sided ($|\cdot|$)	4e-05	1.300	1.113	1.103	0.267	0.253
two-sided ($|\cdot|$)	8e-05	1.300	1.386	1.378	0.197	0.182

Table 1: Mechanistic gain calibration on Dataset A: selected $g^\star$ by minimizing $|b_{\text{tot}}(g) - 1|$ under rate-matched thresholds, and corresponding branching statistics at $g = 1$ and $g^\star$. Values from Table T01 and `gstar.json`.

Spike def	Target rate	$g^\star$	PPL($g{=}1$) [95% CI]	PPL($g^\star$) [95% CI]	ΔPPL
one-sided (+)	1e-05	0.700	15.173 [13.473, 17.130]	17.627 [15.698, 19.816]	2.453
one-sided (+)	2e-05	0.700	15.173 [13.473, 17.130]	17.627 [15.698, 19.816]	2.453
one-sided (+)	4e-05	1.300	15.173 [13.473, 17.130]	16.300 [14.522, 18.308]	1.127
one-sided (+)	8e-05	1.300	15.173 [13.473, 17.130]	16.300 [14.522, 18.308]	1.127
two-sided ($|\cdot|$)	1e-05	0.850	15.173 [13.473, 17.130]	15.746 [13.973, 17.799]	0.573
two-sided ($|\cdot|$)	2e-05	1.000	15.173 [13.473, 17.130]	15.173 [13.473, 17.130]	0.000
two-sided ($|\cdot|$)	4e-05	1.300	15.173 [13.473, 17.130]	16.300 [14.522, 18.308]	1.127
two-sided ($|\cdot|$)	8e-05	1.300	15.173 [13.473, 17.130]	16.300 [14.522, 18.308]	1.127

Table 2: Dataset B evaluation: perplexity at $g = 1$ and at the mechanistically calibrated $g^\star$ (selected on Dataset A), with bootstrap 95% confidence intervals from sequence resampling.

5.5 Robustness across spike definitions

Figure 8 summarizes qualitative robustness across spike definitions, reducing the risk that results depend on a particular thresholding convention.

5.6 Claim boundaries

We do not infer criticality from tail shapes alone. Tail fits are reported as descriptive diagnostics [2, 3], and the signature suite is treated as a set of falsifiable probes that can yield negative results under the same controls.

6 Discussion, limitations, and ethics

Our results support two practical takeaways. First, gain scaling changes local event connectivity even when marginal rates are matched and marginals are controlled by a strong null. Second, a mechanistic calibration based on a branching target does not necessarily generalize to task metrics; negative transfer should be expected and reported.

Limitations. The included runs use deterministic dataset slices to fit within a single-GPU budget (Dataset A: 64 calibration sequences of length 256 for mechanistic metrics; plus a separate 128-sequence slice for raster/null extraction; Dataset B: 96 sequences of length 256; ARC-Challenge test split). Uncertainty estimates for mechanistic metrics can still tighten with additional sampling. We evaluate two checkpoints (instruction-tuned and base) within one model family and a single gain intervention family.

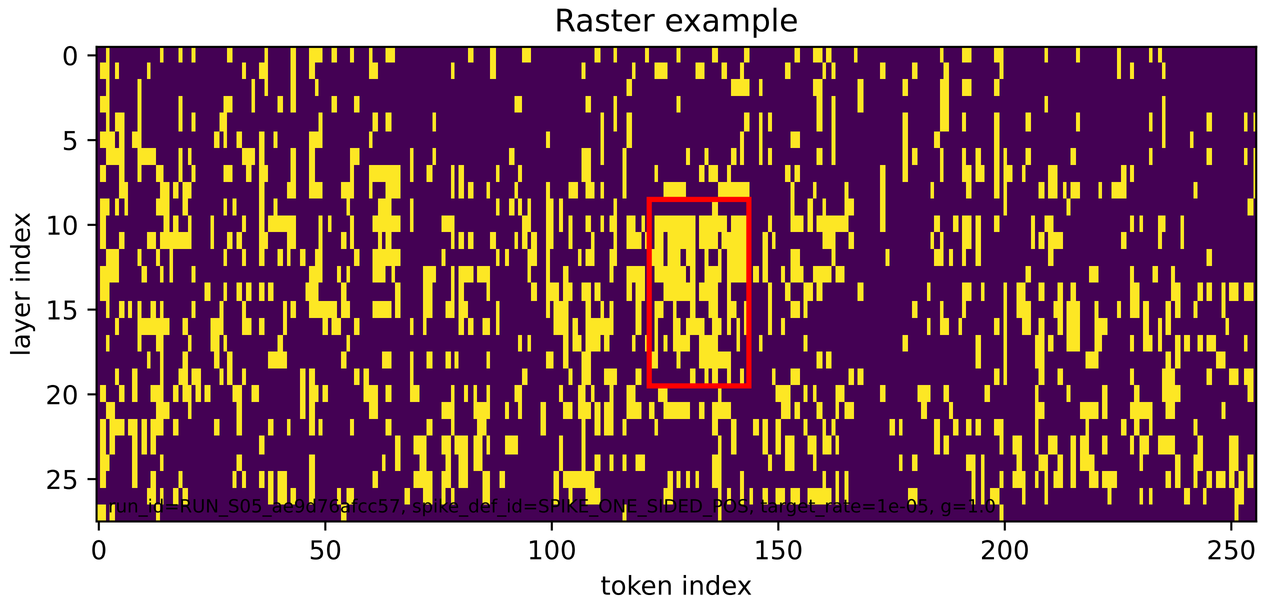

Figure 1: F01: Token-by-layer event raster example (Dataset A). Active sites are thresholded standardized MLP-gate activations; connected components correspond to avalanche-like event cascades on the token-by-layer lattice.

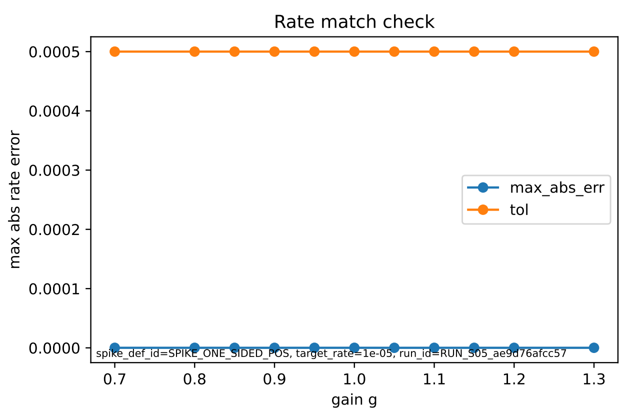

Figure 2: F02: Rate-matching verification across gains (representative condition). Achieved marginal spike rates per layer match the target rate within tolerance across the full gain grid.

Spike def	Target rate	$g^\star$	Acc(g=1) [95% CI]	Acc($g^\star$) [95% CI]	ΔAcc
one-sided (+)	1e-05	0.700	0.885 [0.866, 0.902]	0.873 [0.852, 0.890]	-0.012
one-sided (+)	2e-05	0.700	0.885 [0.866, 0.902]	0.873 [0.852, 0.890]	-0.012
one-sided (+)	4e-05	1.300	0.885 [0.866, 0.902]	0.857 [0.837, 0.875]	-0.028
one-sided (+)	8e-05	1.300	0.885 [0.866, 0.902]	0.857 [0.837, 0.875]	-0.028
two-sided ($\lvert\cdot\rvert$)	1e-05	0.850	0.885 [0.866, 0.902]	0.886 [0.866, 0.903]	0.001
two-sided ($\lvert\cdot\rvert$)	2e-05	1.000	0.885 [0.866, 0.902]	0.885 [0.866, 0.902]	0.000
two-sided ($\lvert\cdot\rvert$)	4e-05	1.300	0.885 [0.866, 0.902]	0.857 [0.837, 0.875]	-0.028
two-sided ($\lvert\cdot\rvert$)	8e-05	1.300	0.885 [0.866, 0.902]	0.857 [0.837, 0.875]	-0.028

Table 3: ARC-Challenge multiple-choice evaluation: accuracy at $g = 1$ vs $g^\star$ (selected on Dataset A), with bootstrap 95% confidence intervals from question resampling.

Ethics. This work analyzes internal activations of open models and does not introduce new training data or deployment. Interpretability results should not be overinterpreted as cognitive equivalence.

Reproducibility statement

All figures and tables referenced in this paper are included as immutable artifacts in the accompanying bundle. Each producing run contains `run_record.json` and `config_resolved.yaml` describing model, data slices, conditions, and artifact hashes. Appendix A lists the run identifiers used for each main figure and table.

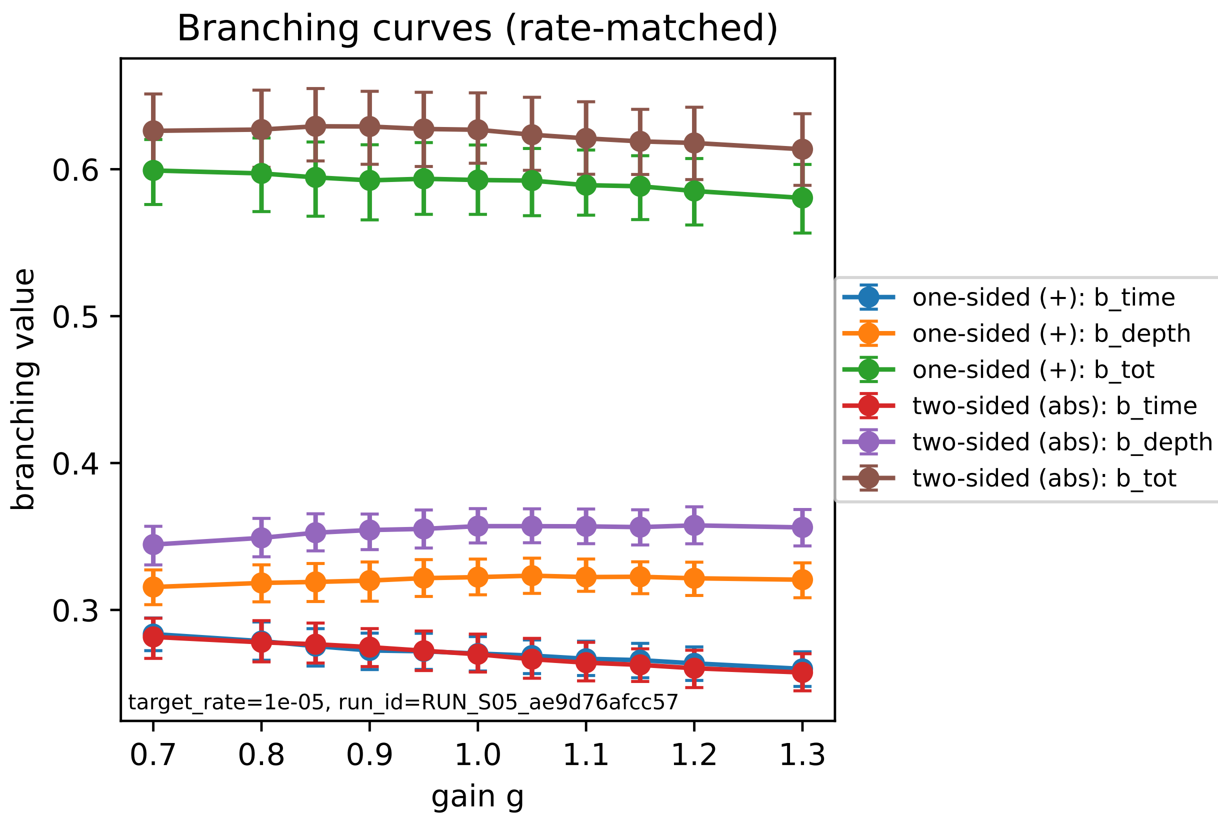

Figure 3: F03: Directional branching metrics versus gain under rate matching (representative target rate). Total branching b_{tot} decomposes into token-direction and depth-direction components.

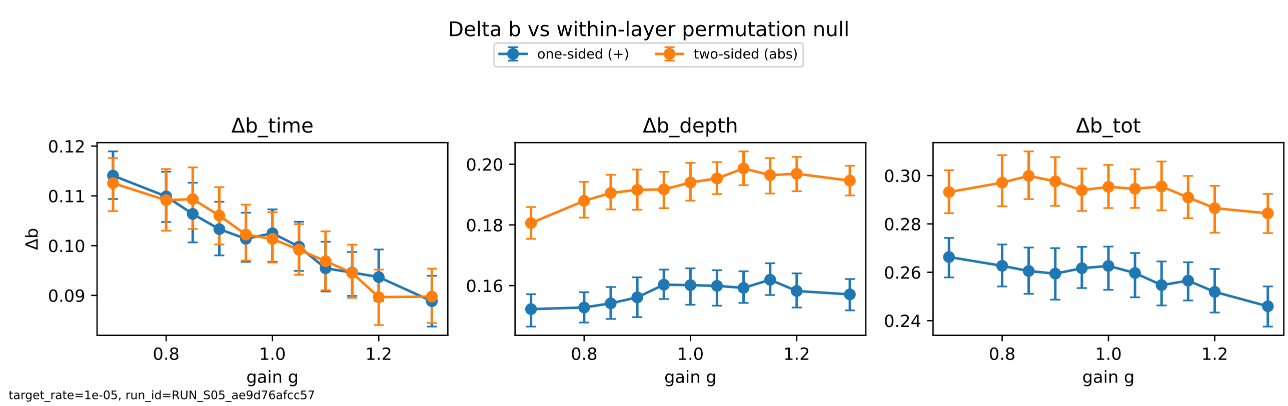

Figure 4: F04: Null-controlled residual connectivity Δb relative to a within-layer time-permutation null that preserves per-layer event-count marginals exactly (representative target rate).

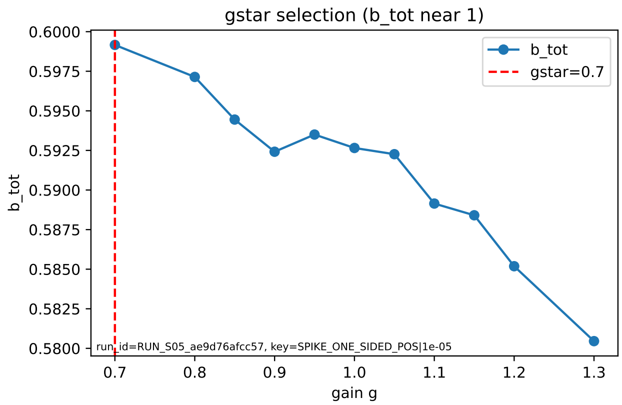

Figure 5: F05: Mechanistic $g^{\star}$ selection by minimizing $|b_{\mathrm{tot}}(g)-1|$ on Dataset A (no performance signals used).

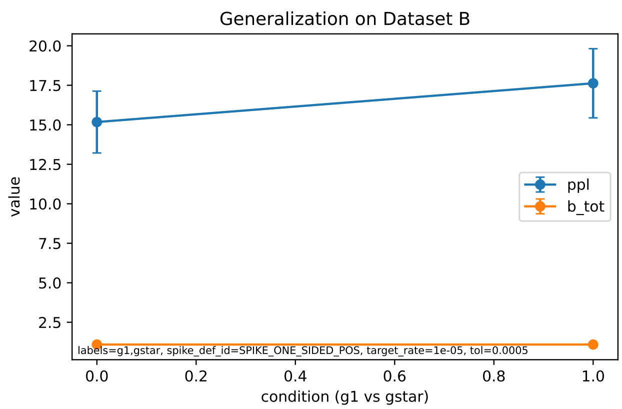

Figure 6: F06: Dataset B evaluation at $g=1$ vs $g^{\star}$ (selected on Dataset A), representative condition with confidence intervals. Tables 2 and 8 report all conditions.

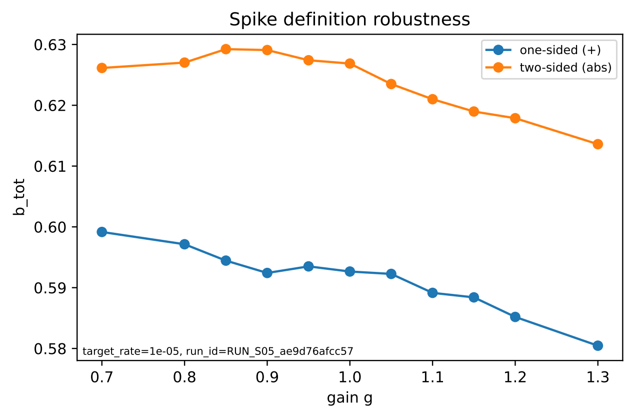

Figure 8: F08: Robustness across spike definitions (representative target rate). Qualitative gain-dependent patterns persist under one-sided and two-sided event definitions.

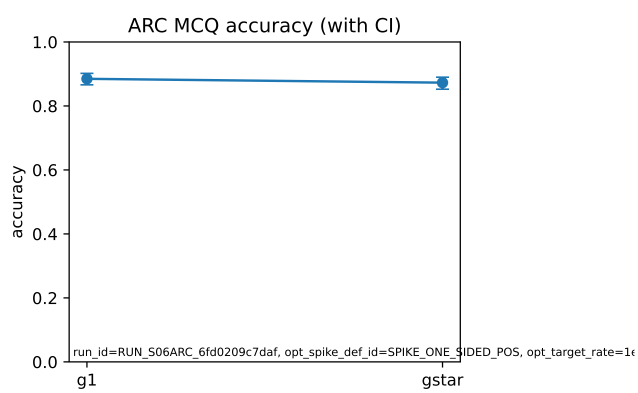

Figure 7: F07: ARC multiple-choice accuracy at $g=1$ vs $g^{\star}$ with confidence intervals (representative condition). Table 3 reports all conditions.

References

[1] John M. Beggs and Dietmar Plenz. Neuronal avalanches in neocortical circuits. *Journal of Neuroscience*, 23(35):11167–11177, 2003. doi:10.1523/JNEUROSCI.23-35-11167.2003.

[2] Aaron Clauset, Cosma Rohilla Shalizi, and Mark E. J. Newman. Power-law distributions in empirical data. *SIAM Review*, 51(4):661–703, 2009. doi:10.1137/070710111.

[3] Jonathan Touboul and Alain Destexhe. Can power-law scaling and neuronal avalanches arise from stochastic dynamics? *PLOS ONE*, 5(2):e8982, 2010. doi:10.1371/journal.pone.0008982.

[4] Mingjie Sun, Xinlei Chen, J. Zico Kolter, and Zhuang Liu. Massive activations in large language models. arXiv preprint arXiv:2402.17762, 2024. doi:10.48550/arXiv.2402.17762.

[5] Samuel S. Schoenholz, Justin Gilmer, Surya Ganguli, and Jascha Sohl-Dickstein. Deep information propagation. arXiv preprint arXiv:1611.01232, 2016. doi:10.48550/arXiv.1611.01232.

[6] Ben Poole, Subhaneil Lahiri, Maithra Raghu, Jascha Sohl-Dickstein, and Surya Ganguli. Exponential expressivity in deep neural networks through transient chaos. arXiv preprint arXiv:1606.05340, 2016. doi:10.48550/arXiv.1606.05340.

[7] Jeffrey Pennington, Samuel S. Schoenholz, and Surya Ganguli. Resurrecting the sigmoid in deep learning through dynamical isometry: theory and practice. arXiv preprint arXiv:1711.04735, 2017. doi:10.48550/arXiv.1711.04735.

[8] Noam Shazeer. GLU variants improve transformer. arXiv preprint arXiv:2002.05202, 2020. doi:10.48550/arXiv.2002.05202.

A Appendix: Additional tables, figures, and provenance

A.1 Susceptibility curves

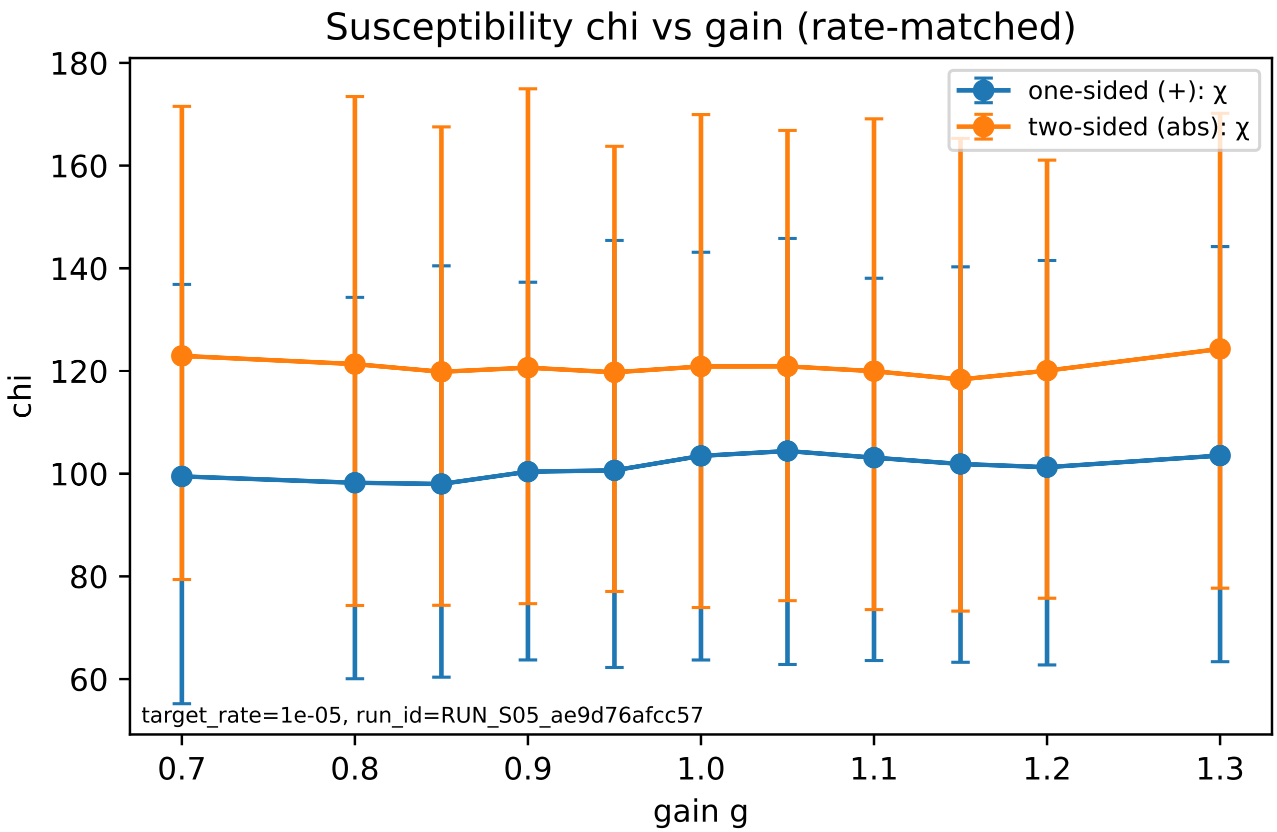

Figure 9: F09: Susceptibility proxy $\chi(g)$ with uncertainty across spike definitions (representative target rate). Full results are in Table T01.

A.2 Null comparison

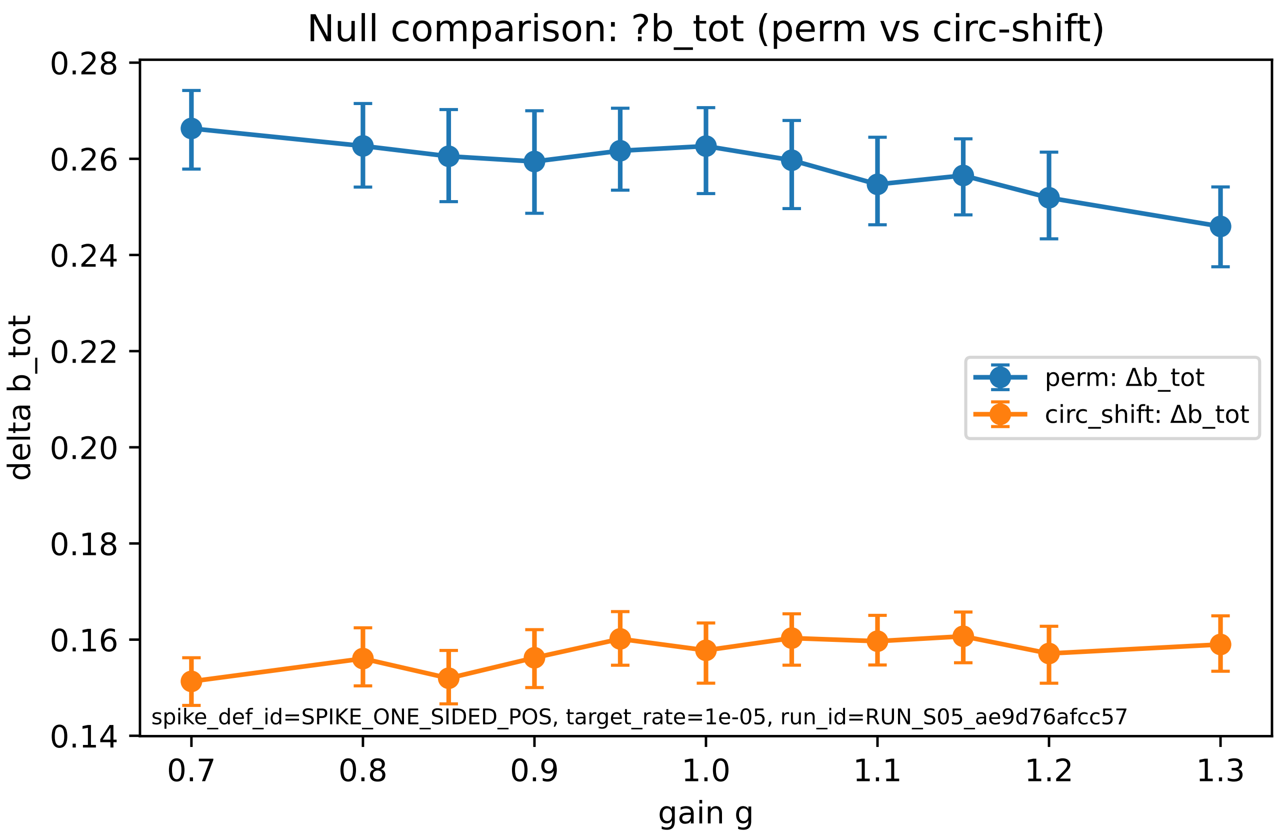

Figure 10: F10: $\Delta b_{\mathrm{tot}}(g)$ under multiple nulls, including a structure-preserving circular-shift null.

A.3 Ablations

Table 4: Ablation comparison of gain interventions at matched rates.

Intervention	Spike def	Target rate	g	Δb_{tot}	max—rate err—
ATTN_GLOBAL	one-sided (+)	1e-05	0.700	0.239	9.81e-08
ATTN_GLOBAL	one-sided (+)	2e-05	0.700	0.242	2.53e-07
ATTN_GLOBAL	one-sided (+)	4e-05	1.300	0.235	3.86e-07
ATTN_GLOBAL	one-sided (+)	8e-05	1.300	0.164	9.53e-07
ATTN_GLOBAL	two-sided ($|\cdot|$)	1e-05	0.850	0.287	4.34e-08
ATTN_GLOBAL	two-sided ($|\cdot|$)	2e-05	1.000	0.301	6.21e-08
ATTN_GLOBAL	two-sided ($|\cdot|$)	4e-05	1.300	0.273	2.06e-07
ATTN_GLOBAL	two-sided ($|\cdot|$)	8e-05	1.300	0.201	4.36e-07
MLP_BAND_EARLY	one-sided (+)	1e-05	0.700	0.256	1.03e-07
MLP_BAND_EARLY	one-sided (+)	2e-05	0.700	0.254	2.12e-07
MLP_BAND_EARLY	one-sided (+)	4e-05	1.300	0.220	4.85e-07
MLP_BAND_EARLY	one-sided (+)	8e-05	1.300	0.152	1.07e-06
MLP_BAND_EARLY	two-sided ($|\cdot|$)	1e-05	0.850	0.296	4.36e-08
MLP_BAND_EARLY	two-sided ($|\cdot|$)	2e-05	1.000	0.306	6.21e-08
MLP_BAND_EARLY	two-sided ($|\cdot|$)	4e-05	1.300	0.262	1.93e-07
MLP_BAND_EARLY	two-sided ($|\cdot|$)	8e-05	1.300	0.188	4.27e-07
MLP_BAND_LATE	one-sided (+)	1e-05	0.700	0.262	1.19e-07
MLP_BAND_LATE	one-sided (+)	2e-05	0.700	0.260	1.93e-07
MLP_BAND_LATE	one-sided (+)	4e-05	1.300	0.223	3.66e-07
MLP_BAND_LATE	one-sided (+)	8e-05	1.300	0.156	8.43e-07
MLP_BAND_LATE	two-sided ($|\cdot|$)	1e-05	0.850	0.298	3.65e-08
MLP_BAND_LATE	two-sided ($|\cdot|$)	2e-05	1.000	0.303	6.21e-08
MLP_BAND_LATE	two-sided ($|\cdot|$)	4e-05	1.300	0.261	2.05e-07
MLP_BAND_LATE	two-sided ($|\cdot|$)	8e-05	1.300	0.192	4.44e-07
MLP_BAND_MID	one-sided (+)	1e-05	0.700	0.263	9.37e-08

Continued on next page

Table 4 (continued)

Intervention	Spike def	Target rate	g	Δb_{tot}	max—rate err—
MLP_BAND_MID	one-sided (+)	2e-05	0.700	0.263	2.62e-07
MLP_BAND_MID	one-sided (+)	4e-05	1.300	0.218	8.40e-07
MLP_BAND_MID	one-sided (+)	8e-05	1.300	0.154	9.60e-07
MLP_BAND_MID	two-sided ($\lvert\cdot\rvert$)	1e-05	0.850	0.297	5.15e-08
MLP_BAND_MID	two-sided ($\lvert\cdot\rvert$)	2e-05	1.000	0.304	6.21e-08
MLP_BAND_MID	two-sided ($\lvert\cdot\rvert$)	4e-05	1.300	0.262	2.62e-07
MLP_BAND_MID	two-sided ($\lvert\cdot\rvert$)	8e-05	1.300	0.191	4.34e-07
MLP_GLOBAL	one-sided (+)	1e-05	1.000	0.263	1.12e-07
MLP_GLOBAL	one-sided (+)	1e-05	0.700	0.266	9.63e-08
MLP_GLOBAL	one-sided (+)	2e-05	1.000	0.260	1.71e-07
MLP_GLOBAL	one-sided (+)	2e-05	0.700	0.271	1.78e-07
MLP_GLOBAL	one-sided (+)	4e-05	1.000	0.226	3.66e-07
MLP_GLOBAL	one-sided (+)	4e-05	1.300	0.210	4.65e-07
MLP_GLOBAL	one-sided (+)	8e-05	1.000	0.158	8.98e-07
MLP_GLOBAL	one-sided (+)	8e-05	1.300	0.145	1.12e-06
MLP_GLOBAL	two-sided ($\lvert\cdot\rvert$)	1e-05	1.000	0.295	3.65e-08
MLP_GLOBAL	two-sided ($\lvert\cdot\rvert$)	1e-05	0.850	0.300	4.30e-08
MLP_GLOBAL	two-sided ($\lvert\cdot\rvert$)	2e-05	1.000	0.305	6.21e-08
MLP_GLOBAL	two-sided ($\lvert\cdot\rvert$)	2e-05	1.000	0.305	6.21e-08
MLP_GLOBAL	two-sided ($\lvert\cdot\rvert$)	4e-05	1.000	0.267	2.05e-07
MLP_GLOBAL	two-sided ($\lvert\cdot\rvert$)	4e-05	1.300	0.253	2.54e-07
MLP_GLOBAL	two-sided ($\lvert\cdot\rvert$)	8e-05	1.000	0.197	4.44e-07
MLP_GLOBAL	two-sided ($\lvert\cdot\rvert$)	8e-05	1.300	0.182	6.07e-07

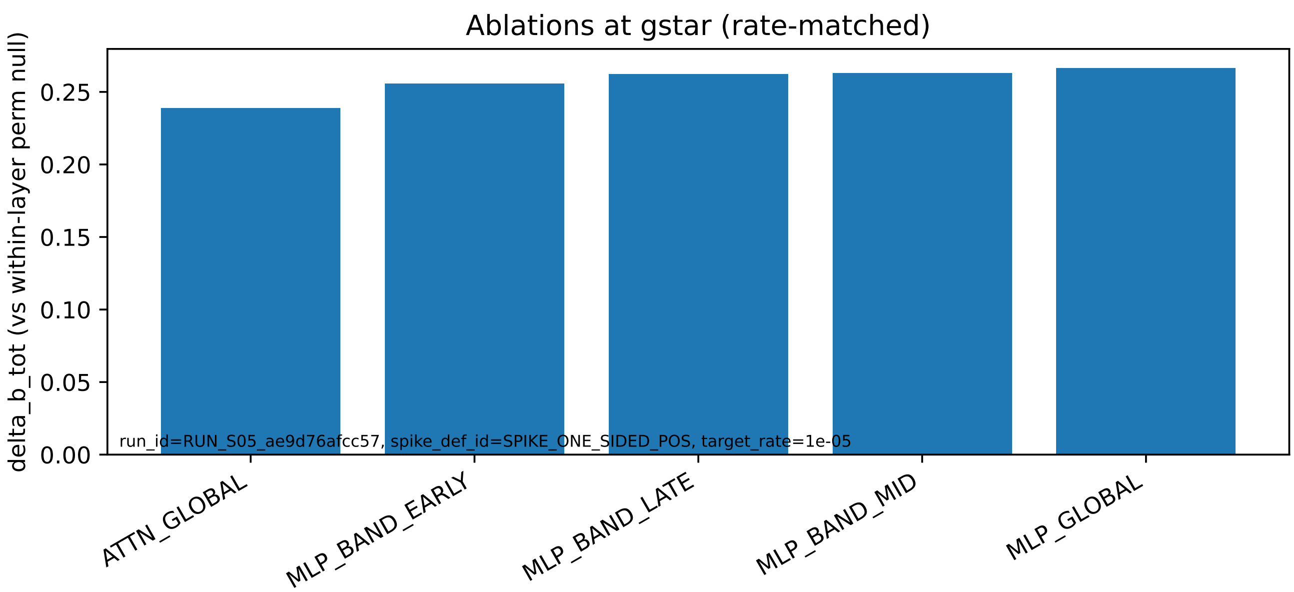

Figure 11: F11: Ablation comparison of gain interventions (MLP vs attention and layer-banded variants) at matched rates.

A.4 Tail fits and crackling diagnostics (descriptive)

Table 5: Tail-fit diagnostics on avalanche size (descriptive only). Continuous-approximation fits on the upper tail defined by a fixed percentile.

Spike def	Target rate	g	$x_{\min}$	n_{tail}	α (PL)	LLR(PL–LN)	LLR(PL–EXP)
one-sided (+)	1e-05	0.700	3.000	7948	2.669	4056.992	2297.358
one-sided (+)	1e-05	0.800	3.000	7976	2.668	3987.043	2252.031
one-sided (+)	1e-05	0.850	3.000	8000	2.676	4029.989	2270.047
one-sided (+)	1e-05	0.900	3.000	7969	2.675	4032.787	2269.391
one-sided (+)	1e-05	0.950	3.000	8000	2.682	4075.005	2319.080
one-sided (+)	1e-05	1.000	3.000	7981	2.685	4090.435	2348.173
one-sided (+)	1e-05	1.050	3.000	7948	2.676	4071.773	2345.749
one-sided (+)	1e-05	1.100	3.000	7933	2.683	4089.890	2371.667
one-sided (+)	1e-05	1.150	3.000	7979	2.696	4147.541	2421.596

Continued on next page

Table 5 (continued)

Spike def	Target rate	g	$x_{\min}$	n_{tail}	α (PL)	LLR(PL–LN)	LLR(PL–EXP)
one-sided (+)	1e-05	1.200	3.000	7932	2.705	4178.151	2508.271
one-sided (+)	1e-05	1.300	3.000	7886	2.711	4181.969	2579.957
one-sided (+)	2e-05	0.700	5.000	7017	2.203	2814.018	2419.748
one-sided (+)	2e-05	0.800	5.000	7036	2.206	2809.992	2411.345
one-sided (+)	2e-05	0.850	5.000	7036	2.208	2806.333	2434.081
one-sided (+)	2e-05	0.900	5.000	7030	2.207	2802.220	2425.993
one-sided (+)	2e-05	0.950	5.000	7054	2.205	2773.071	2361.295
one-sided (+)	2e-05	1.000	5.000	7052	2.198	2729.209	2298.310
one-sided (+)	2e-05	1.050	5.000	7045	2.198	2732.065	2316.180
one-sided (+)	2e-05	1.100	5.000	7057	2.205	2788.389	2366.974
one-sided (+)	2e-05	1.150	5.000	7107	2.210	2804.680	2369.985
one-sided (+)	2e-05	1.200	5.000	7066	2.206	2760.866	2339.463
one-sided (+)	2e-05	1.300	5.000	7058	2.218	2821.391	2445.851
one-sided (+)	4e-05	0.700	7.000	5438	1.853	2053.212	3350.284
one-sided (+)	4e-05	0.800	7.000	5480	1.851	2036.804	3298.712
one-sided (+)	4e-05	0.850	7.000	5468	1.850	2036.417	3291.294
one-sided (+)	4e-05	0.900	7.000	5498	1.853	2049.012	3321.405
one-sided (+)	4e-05	0.950	7.000	5532	1.856	2090.317	3336.432
one-sided (+)	4e-05	1.000	7.000	5485	1.856	2091.606	3366.843
one-sided (+)	4e-05	1.050	7.000	5464	1.853	2073.021	3330.642
one-sided (+)	4e-05	1.100	7.000	5450	1.848	2008.264	3248.724
one-sided (+)	4e-05	1.150	7.000	5512	1.855	2065.282	3309.023
one-sided (+)	4e-05	1.200	7.000	5487	1.851	2038.951	3257.426
one-sided (+)	4e-05	1.300	7.000	5508	1.847	1992.441	3196.796
one-sided (+)	8e-05	0.700	8.000	2480	1.667	1320.508	3931.708
one-sided (+)	8e-05	0.800	8.000	2599	1.658	1349.754	3908.835
one-sided (+)	8e-05	0.850	8.000	2615	1.651	1329.981	3847.509
one-sided (+)	8e-05	0.900	8.000	2577	1.647	1313.601	3788.598
one-sided (+)	8e-05	0.950	8.000	2576	1.643	1293.005	3745.451
one-sided (+)	8e-05	1.000	8.000	2621	1.660	1373.614	3933.414
one-sided (+)	8e-05	1.050	8.000	2597	1.665	1378.840	3976.578
one-sided (+)	8e-05	1.100	8.000	2618	1.667	1398.571	3997.533
one-sided (+)	8e-05	1.150	8.000	2644	1.681	1451.381	4156.312
one-sided (+)	8e-05	1.200	8.000	2641	1.683	1464.244	4170.795
one-sided (+)	8e-05	1.300	7.000	2924	1.699	1678.102	4884.305
two-sided ($\|\cdot\|$)	1e-05	0.700	3.000	7602	2.588	3954.727	2594.214
two-sided ($\|\cdot\|$)	1e-05	0.800	3.000	7594	2.585	3967.263	2579.633
two-sided ($\|\cdot\|$)	1e-05	0.850	3.000	7595	2.576	3948.046	2542.628
two-sided ($\|\cdot\|$)	1e-05	0.900	3.000	7621	2.569	3877.639	2462.573
two-sided ($\|\cdot\|$)	1e-05	0.950	3.000	7615	2.560	3810.091	2383.591
two-sided ($\|\cdot\|$)	1e-05	1.000	3.000	7641	2.562	3816.369	2371.430
two-sided ($\|\cdot\|$)	1e-05	1.050	3.000	7685	2.571	3871.749	2377.783
two-sided ($\|\cdot\|$)	1e-05	1.100	3.000	7685	2.575	3876.614	2385.655
two-sided ($\|\cdot\|$)	1e-05	1.150	3.000	7713	2.577	3884.445	2372.329
two-sided ($\|\cdot\|$)	1e-05	1.200	3.000	7644	2.566	3819.534	2368.780
two-sided ($\|\cdot\|$)	1e-05	1.300	3.000	7582	2.560	3774.944	2392.893
two-sided ($\|\cdot\|$)	2e-05	0.700	5.000	6804	2.172	2729.655	2542.164
two-sided ($\|\cdot\|$)	2e-05	0.800	5.000	6704	2.157	2648.830	2464.156
two-sided ($\|\cdot\|$)	2e-05	0.850	5.000	6786	2.166	2719.956	2514.505
two-sided ($\|\cdot\|$)	2e-05	0.900	5.000	6751	2.162	2677.362	2487.523
two-sided ($\|\cdot\|$)	2e-05	0.950	5.000	6777	2.169	2738.187	2553.888
two-sided ($\|\cdot\|$)	2e-05	1.000	5.000	6717	2.157	2653.084	2481.295
two-sided ($\|\cdot\|$)	2e-05	1.050	5.000	6653	2.151	2597.878	2466.324
two-sided ($\|\cdot\|$)	2e-05	1.100	5.000	6691	2.150	2612.198	2434.546
two-sided ($\|\cdot\|$)	2e-05	1.150	5.000	6666	2.145	2578.680	2394.654
two-sided ($\|\cdot\|$)	2e-05	1.200	5.000	6681	2.151	2581.616	2427.482
two-sided ($\|\cdot\|$)	2e-05	1.300	5.000	6628	2.151	2568.994	2460.674
two-sided ($\|\cdot\|$)	4e-05	0.700	7.000	5408	1.831	1934.127	3126.533
two-sided ($\|\cdot\|$)	4e-05	0.800	7.000	5424	1.830	1909.250	3094.482
two-sided ($\|\cdot\|$)	4e-05	0.850	7.000	5480	1.833	1920.907	3109.686
two-sided ($\|\cdot\|$)	4e-05	0.900	7.000	5432	1.836	1941.444	3178.971
two-sided ($\|\cdot\|$)	4e-05	0.950	7.000	5445	1.839	1957.922	3213.232
two-sided ($\|\cdot\|$)	4e-05	1.000	7.000	5402	1.833	1924.979	3140.572
two-sided ($\|\cdot\|$)	4e-05	1.050	7.000	5390	1.835	1923.012	3170.906
two-sided ($\|\cdot\|$)	4e-05	1.100	7.000	5366	1.830	1908.787	3124.674
two-sided ($\|\cdot\|$)	4e-05	1.150	7.000	5341	1.831	1921.037	3146.645
two-sided ($\|\cdot\|$)	4e-05	1.200	7.000	5321	1.822	1841.185	3016.642
two-sided ($\|\cdot\|$)	4e-05	1.300	7.000	5400	1.836	1930.342	3160.118
two-sided ($\|\cdot\|$)	8e-05	0.700	9.000	2601	1.610	1127.714	3086.452
two-sided ($\|\cdot\|$)	8e-05	0.800	9.000	2661	1.607	1128.619	3053.789
two-sided ($\|\cdot\|$)	8e-05	0.850	9.000	2686	1.610	1155.508	3096.750
two-sided ($\|\cdot\|$)	8e-05	0.900	9.000	2684	1.606	1132.866	3047.846
two-sided ($\|\cdot\|$)	8e-05	0.950	9.000	2705	1.608	1145.213	3070.833
two-sided ($\|\cdot\|$)	8e-05	1.000	9.000	2686	1.608	1147.021	3075.422
two-sided ($\|\cdot\|$)	8e-05	1.050	9.000	2712	1.615	1190.946	3158.379
two-sided ($\|\cdot\|$)	8e-05	1.100	9.000	2711	1.614	1188.181	3142.813
two-sided ($\|\cdot\|$)	8e-05	1.150	9.000	2716	1.622	1206.134	3234.563

Continued on next page

Table 5 (continued)

Spike def	Target rate	g	$x_{\min}$	n_{tail}	α (PL)	LLR(PL–LN)	LLR(PL–EXP)
two-sided ($\lvert \cdot \rvert$)	8e-05	1.200	9.000	2742	1.619	1205.001	3199.957
two-sided ($\lvert \cdot \rvert$)	8e-05	1.300	8.000	2949	1.642	1402.427	3859.723

Table 6: Crackling fit diagnostics (descriptive) with a fail-closed gate.

Spike def	Target rate	g	γ	CI low	CI high	CI width	n_{avals}	R^2	gate
one-sided (+)	1e-05	0.700	1.745	1.631	1.877	0.246	2969	0.970	pass
one-sided (+)	1e-05	0.800	1.718	1.630	1.896	0.266	2911	0.943	pass
one-sided (+)	1e-05	0.850	1.814	1.663	1.912	0.249	2856	0.970	pass
one-sided (+)	1e-05	0.900	1.793	1.602	1.953	0.352	2801	0.930	pass
one-sided (+)	1e-05	0.950	1.792	1.632	1.903	0.272	2756	0.970	pass
one-sided (+)	1e-05	1.000	1.686	1.568	1.897	0.329	2719	0.942	pass
one-sided (+)	1e-05	1.050	1.682	1.559	1.912	0.352	2701	0.957	pass
one-sided (+)	1e-05	1.100	1.632	1.471	1.859	0.387	2685	0.949	pass
one-sided (+)	1e-05	1.150	1.702	1.479	1.834	0.355	2707	0.970	pass
one-sided (+)	1e-05	1.200	1.730	1.591	1.858	0.267	2653	0.975	pass
one-sided (+)	1e-05	1.300	1.702	1.574	1.847	0.273	2606	0.977	pass
one-sided (+)	2e-05	0.700	1.726	1.649	1.835	0.186	5095	0.981	pass
one-sided (+)	2e-05	0.800	1.729	1.650	1.818	0.168	5059	0.972	pass
one-sided (+)	2e-05	0.850	1.751	1.686	1.817	0.130	5064	0.981	pass
one-sided (+)	2e-05	0.900	1.708	1.635	1.806	0.171	5030	0.943	pass
one-sided (+)	2e-05	0.950	1.598	1.533	1.806	0.274	5031	0.906	pass
one-sided (+)	2e-05	1.000	1.613	1.526	1.832	0.306	4995	0.902	pass
one-sided (+)	2e-05	1.050	1.670	1.589	1.846	0.257	4986	0.932	pass
one-sided (+)	2e-05	1.100	1.686	1.609	1.845	0.236	4978	0.932	pass
one-sided (+)	2e-05	1.150	1.729	1.624	1.820	0.196	5005	0.971	pass
one-sided (+)	2e-05	1.200	1.565	1.503	1.762	0.259	4994	0.913	pass
one-sided (+)	2e-05	1.300	1.685	1.608	1.757	0.149	4927	0.972	pass
one-sided (+)	4e-05	0.700	1.652	1.616	1.711	0.095	5081	0.967	pass
one-sided (+)	4e-05	0.800	1.684	1.650	1.744	0.094	5066	0.984	pass
one-sided (+)	4e-05	0.850	1.675	1.653	1.738	0.084	5049	0.981	pass
one-sided (+)	4e-05	0.900	1.683	1.657	1.737	0.081	5092	0.975	pass
one-sided (+)	4e-05	0.950	1.639	1.616	1.722	0.107	5114	0.966	pass
one-sided (+)	4e-05	1.000	1.669	1.625	1.730	0.105	5096	0.981	pass
one-sided (+)	4e-05	1.050	1.650	1.616	1.713	0.096	5096	0.966	pass
one-sided (+)	4e-05	1.100	1.602	1.579	1.693	0.114	5109	0.960	pass
one-sided (+)	4e-05	1.150	1.624	1.599	1.693	0.094	5180	0.973	pass
one-sided (+)	4e-05	1.200	1.645	1.625	1.699	0.074	5146	0.976	pass
one-sided (+)	4e-05	1.300	1.647	1.622	1.712	0.090	5226	0.984	pass
one-sided (+)	8e-05	0.700	1.685	1.654	1.739	0.085	2272	0.956	pass
one-sided (+)	8e-05	0.800	1.736	1.709	1.788	0.079	2371	0.969	pass
one-sided (+)	8e-05	0.850	1.703	1.670	1.753	0.083	2397	0.966	pass
one-sided (+)	8e-05	0.900	1.725	1.690	1.773	0.082	2386	0.963	pass
one-sided (+)	8e-05	0.950	1.705	1.685	1.764	0.079	2398	0.960	pass
one-sided (+)	8e-05	1.000	1.717	1.688	1.774	0.085	2420	0.963	pass
one-sided (+)	8e-05	1.050	1.712	1.683	1.765	0.082	2416	0.964	pass
one-sided (+)	8e-05	1.100	1.709	1.678	1.754	0.076	2466	0.961	pass
one-sided (+)	8e-05	1.150	1.698	1.674	1.761	0.087	2498	0.958	pass
one-sided (+)	8e-05	1.200	1.708	1.673	1.756	0.083	2544	0.955	pass
one-sided (+)	8e-05	1.300	1.727	1.694	1.791	0.097	2548	0.962	pass
two-sided ($\lvert \cdot \rvert$)	1e-05	0.700	1.736	1.637	1.834	0.197	2587	0.945	pass
two-sided ($\lvert \cdot \rvert$)	1e-05	0.800	1.719	1.585	1.879	0.294	2508	0.954	pass
two-sided ($\lvert \cdot \rvert$)	1e-05	0.850	1.827	1.653	1.959	0.306	2485	0.986	pass
two-sided ($\lvert \cdot \rvert$)	1e-05	0.900	1.646	1.551	1.947	0.397	2467	0.942	pass
two-sided ($\lvert \cdot \rvert$)	1e-05	0.950	1.748	1.594	1.869	0.274	2461	0.974	pass
two-sided ($\lvert \cdot \rvert$)	1e-05	1.000	1.823	1.662	1.913	0.251	2453	0.990	pass
two-sided ($\lvert \cdot \rvert$)	1e-05	1.050	1.813	1.646	1.928	0.282	2442	0.986	pass
two-sided ($\lvert \cdot \rvert$)	1e-05	1.100	1.769	1.662	1.979	0.318	2441	0.975	pass
two-sided ($\lvert \cdot \rvert$)	1e-05	1.150	1.700	1.632	2.017	0.385	2446	0.938	pass
two-sided ($\lvert \cdot \rvert$)	1e-05	1.200	1.809	1.693	2.026	0.333	2414	0.958	pass
two-sided ($\lvert \cdot \rvert$)	1e-05	1.300	1.622	1.445	1.878	0.433	2377	0.989	pass
two-sided ($\lvert \cdot \rvert$)	2e-05	0.700	1.703	1.640	1.833	0.192	4653	0.926	pass
two-sided ($\lvert \cdot \rvert$)	2e-05	0.800	1.760	1.705	1.858	0.153	4568	0.982	pass
two-sided ($\lvert \cdot \rvert$)	2e-05	0.850	1.760	1.719	1.854	0.135	4512	0.984	pass
two-sided ($\lvert \cdot \rvert$)	2e-05	0.900	1.761	1.707	1.859	0.152	4461	0.982	pass
two-sided ($\lvert \cdot \rvert$)	2e-05	0.950	1.787	1.726	1.876	0.150	4455	0.978	pass
two-sided ($\lvert \cdot \rvert$)	2e-05	1.000	1.714	1.666	1.840	0.174	4449	0.970	pass
two-sided ($\lvert \cdot \rvert$)	2e-05	1.050	1.667	1.626	1.799	0.174	4414	0.969	pass
two-sided ($\lvert \cdot \rvert$)	2e-05	1.100	1.701	1.659	1.791	0.132	4423	0.979	pass
two-sided ($\lvert \cdot \rvert$)	2e-05	1.150	1.688	1.641	1.776	0.135	4405	0.976	pass
two-sided ($\lvert \cdot \rvert$)	2e-05	1.200	1.655	1.623	1.776	0.153	4354	0.978	pass
two-sided ($\lvert \cdot \rvert$)	2e-05	1.300	1.637	1.619	1.730	0.110	4328	0.984	pass
two-sided ($\lvert \cdot \rvert$)	4e-05	0.700	1.646	1.605	1.717	0.112	4868	0.968	pass
two-sided ($\lvert \cdot \rvert$)	4e-05	0.800	1.713	1.673	1.770	0.097	4791	0.975	pass

Continued on next page

Table 6 (continued)

Spike def	Target rate	g	γ	CI low	CI high	CI width	n_{avals}	R^2	gate
two-sided ($\lvert\cdot\rvert$)	4e-05	0.850	1.665	1.650	1.736	0.086	4831	0.977	pass
two-sided ($\lvert\cdot\rvert$)	4e-05	0.900	1.664	1.628	1.728	0.100	4827	0.980	pass
two-sided ($\lvert\cdot\rvert$)	4e-05	0.950	1.642	1.626	1.712	0.085	4860	0.976	pass
two-sided ($\lvert\cdot\rvert$)	4e-05	1.000	1.687	1.652	1.744	0.092	4846	0.976	pass
two-sided ($\lvert\cdot\rvert$)	4e-05	1.050	1.660	1.636	1.721	0.085	4865	0.980	pass
two-sided ($\lvert\cdot\rvert$)	4e-05	1.100	1.669	1.632	1.733	0.100	4797	0.977	pass
two-sided ($\lvert\cdot\rvert$)	4e-05	1.150	1.655	1.620	1.712	0.093	4784	0.976	pass
two-sided ($\lvert\cdot\rvert$)	4e-05	1.200	1.658	1.630	1.725	0.095	4848	0.970	pass
two-sided ($\lvert\cdot\rvert$)	4e-05	1.300	1.622	1.600	1.688	0.087	4912	0.977	pass
two-sided ($\lvert\cdot\rvert$)	8e-05	0.700	1.625	1.601	1.691	0.090	2537	0.953	pass
two-sided ($\lvert\cdot\rvert$)	8e-05	0.800	1.625	1.604	1.687	0.082	2540	0.959	pass
two-sided ($\lvert\cdot\rvert$)	8e-05	0.850	1.625	1.607	1.683	0.076	2543	0.957	pass
two-sided ($\lvert\cdot\rvert$)	8e-05	0.900	1.605	1.583	1.666	0.083	2611	0.949	pass
two-sided ($\lvert\cdot\rvert$)	8e-05	0.950	1.636	1.595	1.690	0.095	2609	0.958	pass
two-sided ($\lvert\cdot\rvert$)	8e-05	1.000	1.639	1.607	1.696	0.089	2627	0.953	pass
two-sided ($\lvert\cdot\rvert$)	8e-05	1.050	1.662	1.627	1.712	0.085	2706	0.961	pass
two-sided ($\lvert\cdot\rvert$)	8e-05	1.100	1.654	1.617	1.703	0.086	2735	0.957	pass
two-sided ($\lvert\cdot\rvert$)	8e-05	1.150	1.646	1.612	1.705	0.093	2751	0.958	pass
two-sided ($\lvert\cdot\rvert$)	8e-05	1.200	1.676	1.638	1.716	0.078	2810	0.967	pass
two-sided ($\lvert\cdot\rvert$)	8e-05	1.300	1.663	1.627	1.718	0.091	2838	0.961	pass

A.5 Replication summary (base vs instruct)

Spike def	Target rate	$g^\star_{\mathrm{base}}$	$g^\star_{\mathrm{inst}}$	Δb_{tot} (base)	Δb_{tot} (inst)	χ (base)	χ (inst)
one-sided (+)	1e-05	0.700	0.700	0.274	0.266	154.531	99.469
one-sided (+)	2e-05	0.700	0.700	0.276	0.271	251.222	160.923
one-sided (+)	4e-05	1.300	1.300	0.217	0.210	361.316	242.203
one-sided (+)	8e-05	1.300	1.300	0.148	0.145	562.081	338.812
two-sided ($\lvert\cdot\rvert$)	1e-05	0.800	0.850	0.304	0.300	196.479	119.844
two-sided ($\lvert\cdot\rvert$)	2e-05	0.800	1.000	0.313	0.305	316.137	201.021
two-sided ($\lvert\cdot\rvert$)	4e-05	1.300	1.300	0.264	0.253	458.077	329.981
two-sided ($\lvert\cdot\rvert$)	8e-05	1.300	1.300	0.188	0.182	726.533	491.691

Table 7: Replication summary comparing base vs instruction-tuned checkpoints at their respective $g^\star$ (appendix).

A.6 Selected Dataset A condition table

Table 8: Selected columns from Table T01 (Dataset A): rate-matching error, branching, null-controlled residual, susceptibility proxy, crackling exponent estimate, and component summary statistics across all conditions.

Spike def	Target rate	g	max—rate err—	b_{tot}	Δb_{tot}	χ	γ	#avals	mean size	mean span (tok)	mean span (layers)
one-sided (+)	1e-05	0.700	9.63e-08	0.599	0.266	99.469	1.745	30584	2.833	1.442	1.529
one-sided (+)	1e-05	0.800	9.95e-08	0.597	0.263	98.227	1.718	30662	2.825	1.431	1.539
one-sided (+)	1e-05	0.850	8.44e-08	0.594	0.261	97.997	1.814	30786	2.816	1.424	1.539
one-sided (+)	1e-05	0.900	7.62e-08	0.592	0.259	100.388	1.793	30882	2.806	1.418	1.541
one-sided (+)	1e-05	0.950	9.63e-08	0.594	0.262	100.642	1.792	30771	2.815	1.417	1.548
one-sided (+)	1e-05	1.000	1.12e-07	0.593	0.263	103.465	1.686	30794	2.813	1.414	1.549
one-sided (+)	1e-05	1.050	7.76e-08	0.592	0.260	104.412	1.682	30710	2.822	1.412	1.552
one-sided (+)	1e-05	1.100	9.67e-08	0.589	0.255	103.114	1.632	30817	2.810	1.408	1.548
one-sided (+)	1e-05	1.150	8.92e-08	0.588	0.257	101.886	1.702	30820	2.810	1.407	1.548
one-sided (+)	1e-05	1.200	8.28e-08	0.585	0.252	101.261	1.730	30903	2.804	1.404	1.545
one-sided (+)	1e-05	1.300	9.95e-08	0.580	0.246	103.526	1.702	31031	2.792	1.395	1.540
one-sided (+)	2e-05	0.700	1.78e-07	0.833	0.271	160.923	1.726	33555	5.163	1.713	1.913
one-sided (+)	2e-05	0.800	1.62e-07	0.830	0.270	163.233	1.729	33686	5.144	1.700	1.922
one-sided (+)	2e-05	0.850	2.18e-07	0.829	0.267	162.730	1.751	33711	5.136	1.693	1.921
one-sided (+)	2e-05	0.900	2.01e-07	0.828	0.266	163.570	1.708	33741	5.131	1.690	1.928
one-sided (+)	2e-05	0.950	2.34e-07	0.824	0.263	163.832	1.598	33920	5.104	1.686	1.932
one-sided (+)	2e-05	1.000	1.71e-07	0.823	0.260	165.146	1.613	33974	5.098	1.682	1.931
one-sided (+)	2e-05	1.050	2.04e-07	0.822	0.259	164.680	1.670	33948	5.102	1.680	1.930
one-sided (+)	2e-05	1.100	2.36e-07	0.820	0.257	162.738	1.686	34009	5.096	1.679	1.931
one-sided (+)	2e-05	1.150	1.91e-07	0.818	0.258	158.867	1.729	34110	5.077	1.679	1.928
one-sided (+)	2e-05	1.200	2.47e-07	0.814	0.252	159.493	1.565	34292	5.050	1.672	1.921
one-sided (+)	2e-05	1.300	2.41e-07	0.810	0.251	165.867	1.685	34493	5.018	1.662	1.911
one-sided (+)	4e-05	0.700	4.58e-07	1.109	0.233	255.178	1.652	25770	13.436	2.144	2.500

Continued on next page

Table 8 (continued)

Spike def	Target rate	g	max—rate err—	$b_{\mathtt{tot}}$	$\Delta b_{\mathtt{tot}}$	χ	γ	#avals	mean size	mean span (tok)	mean span (layers)
one-sided (+)	4e-05	0.800	4.28e-07	1.108	0.233	255.887	1.684	25745	13.439	2.136	2.519
one-sided (+)	4e-05	0.850	4.31e-07	1.108	0.230	254.088	1.675	25728	13.459	2.132	2.530
one-sided (+)	4e-05	0.900	4.17e-07	1.106	0.232	253.071	1.683	25908	13.366	2.124	2.523
one-sided (+)	4e-05	0.950	4.08e-07	1.104	0.230	252.461	1.639	26049	13.294	2.117	2.524
one-sided (+)	4e-05	1.000	3.66e-07	1.102	0.226	256.167	1.669	26230	13.200	2.102	2.506
one-sided (+)	4e-05	1.050	4.06e-07	1.100	0.224	257.204	1.650	26239	13.190	2.104	2.506
one-sided (+)	4e-05	1.100	5.23e-07	1.098	0.223	252.135	1.602	26387	13.110	2.106	2.494
one-sided (+)	4e-05	1.150	4.60e-07	1.095	0.219	249.300	1.624	26563	13.025	2.104	2.491
one-sided (+)	4e-05	1.200	5.20e-07	1.093	0.218	244.250	1.645	26786	12.911	2.095	2.473
one-sided (+)	4e-05	1.300	4.65e-07	1.089	0.210	242.203	1.647	27007	12.821	2.094	2.458
one-sided (+)	8e-05	0.700	1.01e-06	1.402	0.163	397.668	1.685	12101	57.153	2.945	2.740
one-sided (+)	8e-05	0.800	1.06e-06	1.399	0.164	395.740	1.736	12377	55.863	2.921	2.799
one-sided (+)	8e-05	0.850	9.29e-07	1.399	0.164	393.898	1.703	12370	55.913	2.926	2.825
one-sided (+)	8e-05	0.900	1.12e-06	1.398	0.162	392.063	1.725	12412	55.701	2.931	2.812
one-sided (+)	8e-05	0.950	9.36e-07	1.397	0.160	383.967	1.705	12536	55.174	2.912	2.801
one-sided (+)	8e-05	1.000	8.98e-07	1.395	0.158	381.054	1.717	12629	54.784	2.908	2.781
one-sided (+)	8e-05	1.050	9.75e-07	1.395	0.156	374.196	1.712	12635	54.778	2.928	2.725
one-sided (+)	8e-05	1.100	1.06e-06	1.392	0.154	366.487	1.709	12811	53.966	2.912	2.712
one-sided (+)	8e-05	1.150	8.73e-07	1.392	0.152	353.791	1.698	12788	54.079	2.924	2.681
one-sided (+)	8e-05	1.200	9.91e-07	1.390	0.149	344.150	1.708	12939	53.458	2.918	2.653
one-sided (+)	8e-05	1.300	1.12e-06	1.387	0.145	338.812	1.727	13127	52.655	2.904	2.594
two-sided ($\lvert \cdot \rvert$)	1e-05	0.700	7.60e-08	0.626	0.293	122.941	1.736	28576	3.035	1.426	1.595
two-sided ($\lvert \cdot \rvert$)	1e-05	0.800	5.91e-08	0.627	0.297	121.347	1.719	28573	3.035	1.413	1.606
two-sided ($\lvert \cdot \rvert$)	1e-05	0.850	4.30e-08	0.629	0.300	119.844	1.827	28483	3.046	1.412	1.616
two-sided ($\lvert \cdot \rvert$)	1e-05	0.900	4.66e-08	0.629	0.298	120.649	1.646	28480	3.046	1.410	1.622
two-sided ($\lvert \cdot \rvert$)	1e-05	0.950	4.70e-08	0.627	0.294	119.745	1.748	28530	3.040	1.408	1.626
two-sided ($\lvert \cdot \rvert$)	1e-05	1.000	3.65e-08	0.627	0.295	120.883	1.823	28561	3.039	1.405	1.632
two-sided ($\lvert \cdot \rvert$)	1e-05	1.050	4.66e-08	0.623	0.295	120.902	1.813	28622	3.031	1.400	1.632
two-sided ($\lvert \cdot \rvert$)	1e-05	1.100	5.71e-08	0.621	0.295	119.962	1.769	28698	3.022	1.399	1.633
two-sided ($\lvert \cdot \rvert$)	1e-05	1.150	4.73e-08	0.619	0.291	118.353	1.700	28748	3.018	1.397	1.633
two-sided ($\lvert \cdot \rvert$)	1e-05	1.200	5.33e-08	0.618	0.287	120.059	1.809	28703	3.024	1.393	1.636
two-sided ($\lvert \cdot \rvert$)	1e-05	1.300	6.62e-08	0.614	0.284	124.299	1.622	28686	3.024	1.388	1.630
two-sided ($\lvert \cdot \rvert$)	2e-05	0.700	1.63e-07	0.851	0.306	197.402	1.703	31375	5.529	1.684	1.997
two-sided ($\lvert \cdot \rvert$)	2e-05	0.800	1.23e-07	0.851	0.301	194.726	1.760	31338	5.536	1.673	2.013
two-sided ($\lvert \cdot \rvert$)	2e-05	0.850	9.25e-08	0.852	0.304	193.456	1.760	31302	5.545	1.668	2.023
two-sided ($\lvert \cdot \rvert$)	2e-05	0.900	9.08e-08	0.852	0.304	196.842	1.761	31306	5.541	1.664	2.024
two-sided ($\lvert \cdot \rvert$)	2e-05	0.950	1.00e-07	0.853	0.302	196.673	1.787	31235	5.555	1.662	2.034
two-sided ($\lvert \cdot \rvert$)	2e-05	1.000	6.21e-08	0.853	0.305	201.021	1.714	31188	5.566	1.660	2.036
two-sided ($\lvert \cdot \rvert$)	2e-05	1.050	9.85e-08	0.852	0.303	202.030	1.667	31240	5.555	1.655	2.031
two-sided ($\lvert \cdot \rvert$)	2e-05	1.100	7.44e-08	0.850	0.299	200.626	1.701	31241	5.556	1.654	2.039
two-sided ($\lvert \cdot \rvert$)	2e-05	1.150	1.21e-07	0.848	0.299	197.394	1.688	31283	5.547	1.652	2.038
two-sided ($\lvert \cdot \rvert$)	2e-05	1.200	1.23e-07	0.844	0.298	198.463	1.655	31381	5.529	1.647	2.034
two-sided ($\lvert \cdot \rvert$)	2e-05	1.300	1.07e-07	0.840	0.289	204.731	1.637	31581	5.494	1.639	2.020
two-sided ($\lvert \cdot \rvert$)	4e-05	0.700	2.33e-07	1.118	0.272	324.517	1.646	24599	14.106	2.094	2.637
two-sided ($\lvert \cdot \rvert$)	4e-05	0.800	2.29e-07	1.118	0.271	325.905	1.713	24603	14.094	2.078	2.646
two-sided ($\lvert \cdot \rvert$)	4e-05	0.850	2.49e-07	1.118	0.271	323.725	1.665	24604	14.096	2.077	2.657
two-sided ($\lvert \cdot \rvert$)	4e-05	0.900	1.94e-07	1.117	0.270	326.791	1.664	24661	14.071	2.072	2.647
two-sided ($\lvert \cdot \rvert$)	4e-05	0.950	1.89e-07	1.114	0.266	325.458	1.642	24815	13.984	2.065	2.642
two-sided ($\lvert \cdot \rvert$)	4e-05	1.000	2.05e-07	1.113	0.267	330.934	1.687	24957	13.902	2.057	2.637
two-sided ($\lvert \cdot \rvert$)	4e-05	1.050	1.94e-07	1.113	0.267	329.652	1.660	24983	13.880	2.060	2.624
two-sided ($\lvert \cdot \rvert$)	4e-05	1.100	2.12e-07	1.112	0.263	324.618	1.669	24968	13.899	2.063	2.622
two-sided ($\lvert \cdot \rvert$)	4e-05	1.150	2.10e-07	1.111	0.263	318.179	1.655	24982	13.883	2.063	2.619
two-sided ($\lvert \cdot \rvert$)	4e-05	1.200	2.07e-07	1.107	0.258	319.883	1.658	25226	13.758	2.058	2.607
two-sided ($\lvert \cdot \rvert$)	4e-05	1.300	2.54e-07	1.103	0.253	329.981	1.622	25543	13.583	2.055	2.583
two-sided ($\lvert \cdot \rvert$)	8e-05	0.700	4.33e-07	1.388	0.199	522.920	1.625	12658	54.793	2.815	3.152
two-sided ($\lvert \cdot \rvert$)	8e-05	0.800	4.64e-07	1.389	0.199	516.944	1.625	12695	54.648	2.795	3.187
two-sided ($\lvert \cdot \rvert$)	8e-05	0.850	5.76e-07	1.389	0.199	511.292	1.625	12756	54.394	2.788	3.188
two-sided ($\lvert \cdot \rvert$)	8e-05	0.900	4.55e-07	1.388	0.201	508.814	1.605	12846	53.996	2.793	3.199
two-sided ($\lvert \cdot \rvert$)	8e-05	0.950	4.49e-07	1.387	0.199	510.891	1.636	12878	53.891	2.787	3.202
two-sided ($\lvert \cdot \rvert$)	8e-05	1.000	4.44e-07	1.386	0.197	516.272	1.639	13016	53.331	2.778	3.150
two-sided ($\lvert \cdot \rvert$)	8e-05	1.050	5.57e-07	1.384	0.196	517.214	1.662	13115	52.903	2.786	3.133
two-sided ($\lvert \cdot \rvert$)	8e-05	1.100	5.28e-07	1.382	0.191	509.224	1.654	13248	52.353	2.788	3.104
two-sided ($\lvert \cdot \rvert$)	8e-05	1.150	5.94e-07	1.382	0.190	502.037	1.646	13313	52.082	2.791	3.058
two-sided ($\lvert \cdot \rvert$)	8e-05	1.200	4.34e-07	1.381	0.188	495.254	1.676	13496	51.433	2.790	3.047
two-sided ($\lvert \cdot \rvert$)	8e-05	1.300	6.07e-07	1.378	0.182	491.691	1.663	13581	51.068	2.815	2.931

A.7 Artifact provenance

Artifact	Run	Config hash (prefix)
F01_RASTER_EXAMPLE	RUN_S05_ae9d76afcc57	d7513f37dfef...
F02_RATE_MATCH_CHECK	RUN_S05_ae9d76afcc57	d7513f37dfef...
F03_BRANCHING_CURVES	RUN_S05_ae9d76afcc57	d7513f37dfef...
F04_NULL_DELTAB	RUN_S05_ae9d76afcc57	d7513f37dfef...
F05_GSTAR_SELECTION	RUN_S05_ae9d76afcc57	d7513f37dfef...
F06_GENERALIZATION_B	RUN_S06B_1791d65dc967	d7513f37dfef...
F07_ARC_MCQ	RUN_S06ARC_6fd0209c7daf	d7513f37dfef...
F08_SPIKEDEF_ROBUST	RUN_S05_ae9d76afcc57	d7513f37dfef...
F09_CHI_CURVES	RUN_S05_ae9d76afcc57	d7513f37dfef...
F10_NULL_COMPARE	RUN_S05_ae9d76afcc57	d7513f37dfef...
F11_ABLATIONS	RUN_S05_ae9d76afcc57	d7513f37dfef...
T01_SUMMARY	RUN_S05_ae9d76afcc57	d7513f37dfef...
T02_GENERALIZATION	RUN_S06B_1791d65dc967	d7513f37dfef...
T03_ARC	RUN_S06ARC_6fd0209c7daf	d7513f37dfef...
T04_TAIL_FITS	RUN_S05_ae9d76afcc57	d7513f37dfef...
T05_CRACKLING_DIAGNOSTICS	RUN_S05_ae9d76afcc57	d7513f37dfef...
T06_ABLATIONS	RUN_S05_ae9d76afcc57	d7513f37dfef...
T07_REPLICATION_SUMMARY	RUN_S07_72cab7151a1f	d7513f37dfef...

Table 9: Artifact provenance (run identifiers and resolved config hashes) for exported figures and tables. Full hashes and checksums are recorded in each run's `run_record.json` and in `MANIFEST.sha256`.