Perturbations: gamma gates rates not just mean inhibition

exp042 · 2 June 2026

Abstract

exp025′s rate gap admits a cheap reading: PING fires less because the I-loop delivers more inhibition. This entry forecloses that reading and identifies *what* about the I-stream is doing the suppressing: qualitatively (rhythm vs mean), and quantitatively (which temporal precision is required). Scaffolded by ar009 §Leg 1 item 2.

Methods

Pure inference on the trained exp025 PING baseline (seed 42, $\theta_u = \text{off}$). For each batch the I-population spike tensor $s^I_{\text{base}} \in \{0,1\}^{T\times B\times N_I}$ is recorded from a baseline forward pass, then an override tensor replaces it in a second pass via the exp037 hidden-perturbation hook. The E-population experiences only the override I-stream through W^{IE}; the readout consumes the perturbed E spikes.

Every perturbation only *moves* spikes in time (or, for Poisson, redraws at the matched count) — none adds or removes them — so the mean per-cell I rate is matched to baseline by construction: exactly for phase-shuffle, and to within ≈ 3% for the jitter families across the range where each result is read. The one exception is cycle-coherent jitter at the largest σ: a Gaussian block offset with $\sigma = 100$ ms displaces part of each burst past the ends of the fixed presentation window, where it is clamped and lost, so the *realised* I rate falls to 40.5 Hz (24% below the 53.2 Hz baseline). Realised I is therefore plotted on every sweep, and the strict same-mean-inhibition comparison (the compound figure below) is anchored at $\sigma = 14$ ms, where realised I is still within 3% of baseline on both arms.

Five perturbation families:

1. **Baseline**: no override; trained PING dynamics.
2. **Cycle-coherent jitter**: partition the trial into blocks of length $1/f_\gamma$ (≈ 22.8 ms at the trained operating point from exp041). For each (trial, block), draw a single Gaussian offset $\Delta \sim \mathcal{N}(0, \sigma^2)$ and shift every I-spike in that block by Δ. Within-burst cross-cell synchrony is preserved exactly; only the *placement* of each burst is perturbed. Sweep $\sigma \in \{0, 1, 3, 7, 14, 21, 28, 42, 60, 100\}$ ms.
3. **Per-I-cell jitter**: draw an independent Gaussian offset $\Delta_{b,n,k} \sim \mathcal{N}(0, \sigma^2)$ for *every* I-spike and shift each spike by its own offset. Destroys within-burst cross-cell synchrony while keeping the mean per-cell rate; the within-burst counterpart of cycle-coherent jitter. Sweep $\sigma \in \{0, 0.5, 1, 2, 5, 9, 14, 21, 50\}$ ms.

4. **Phase-shuffle**: per-trial permutation π_b of the time axis applied to all I-cells together: $s^I_{\text{shuf}[t,b,n]} = s^I_{\text{base}[\pi_b(t),b,n]}$. Preserves cross-cell co-firing within a timestep; destroys all phase structure.
5. **Rate-matched Poisson**: per-(trial, cell) Bernoulli with $p = \text{count}_{b,n}/T$. Destroys both temporal and cross-cell structure; tests the g_i variance limit.

Results

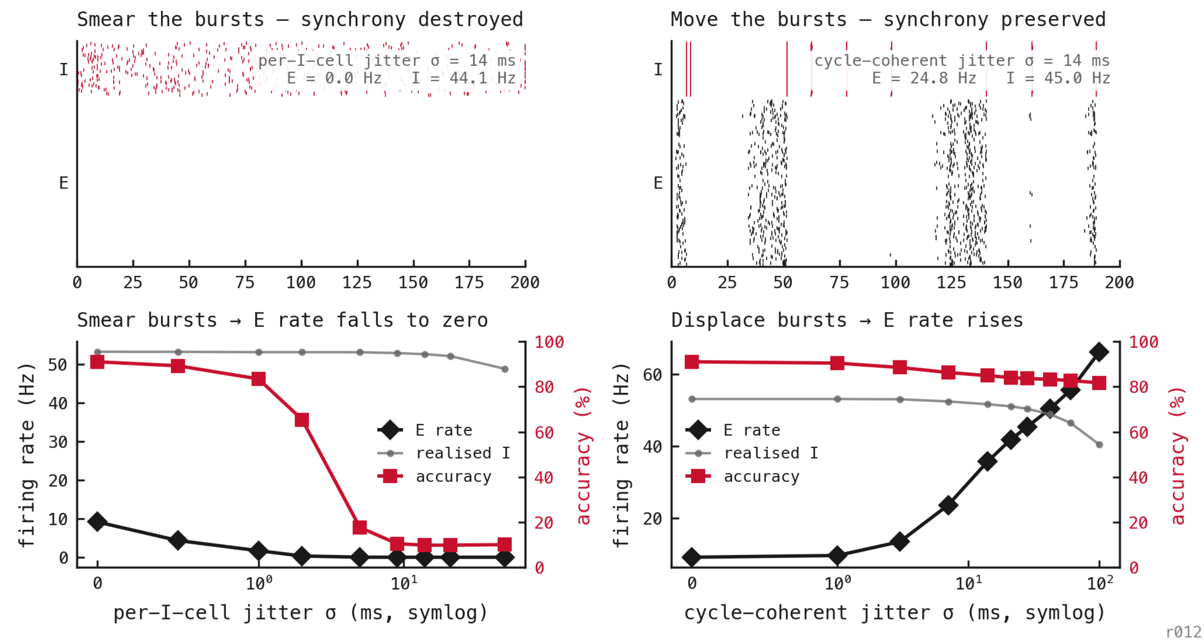

Figure 1: Two inference-time perturbations of the trained-PING I-stream, **both holding the mean per-cell I rate fixed at ≈ 53.2 Hz**, push the E rate in *opposite* directions, which a mean-inhibition account cannot produce. **Both columns use the same jitter magnitude, σ = 14 ms** — only the *kind* of jitter differs. **Left column, smear the bursts:** per-I-cell jitter (realised I 52.5 Hz) scatters the spikes within each burst, destroying synchrony while leaving the mean untouched; the burst dissolves into a continuous shunt, the E rate falls to zero, and accuracy collapses toward chance (9.9% at the Poisson limit). **Right column, move the bursts:** cycle-coherent jitter (realised I 51.7 Hz — within 3% of the left column) displaces each gamma burst bodily but keeps its within-burst synchrony; the I-stream opens gaps and the E rate *rises* from 9.1 Hz at baseline to 35.8 Hz, accuracy holding near 84.9%. Same jitter magnitude, same mean inhibition (both ≈ 53.2 Hz), opposite outcome: what gates the E rate is the *timing* of inhibition, the rhythm, not its average level. The cycle-coherent rise continues past the full phase-shuffle level to 66.3 Hz by σ = 100 ms (bottom-right sweep), but there the finite trial window truncates the most-displaced bursts and realised I falls 24%, so the strict rate-matched reading is taken at σ = 14 ms; see Methods.

Per-I-cell jitter

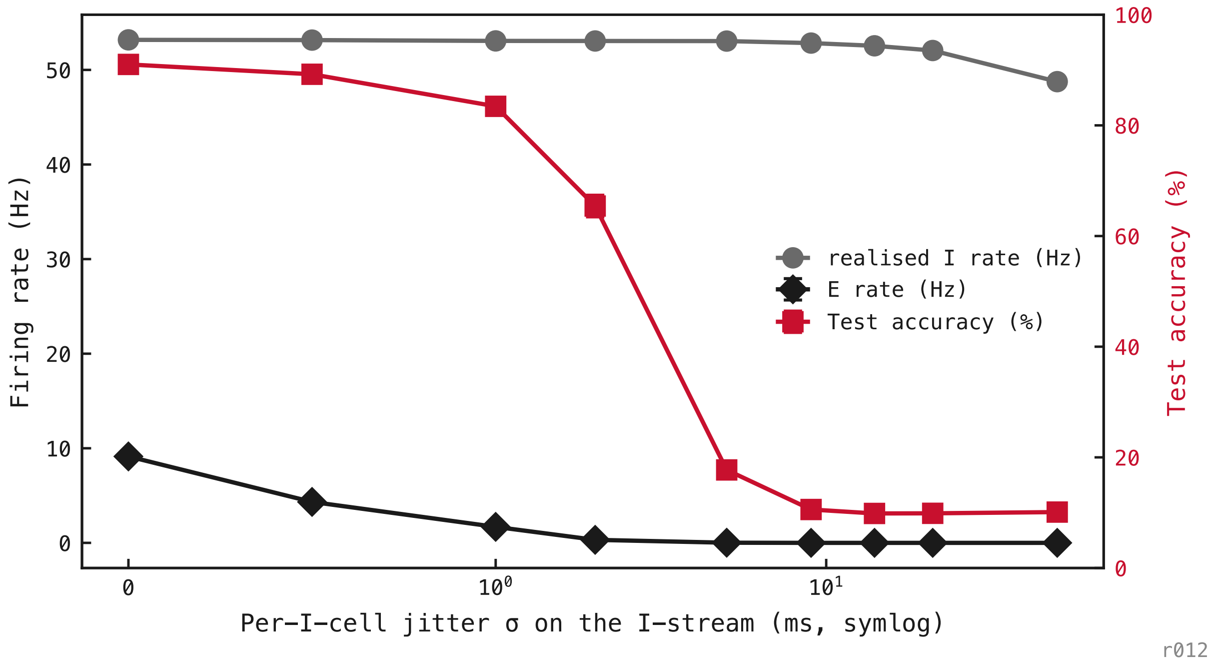

Figure 2: Per-I-cell jitter sweep, three seeds. Each spike receives an independent Gaussian offset; mean per-cell I rate is preserved exactly. **E rate (black diamonds, left axis) falls monotonically from baseline** (9.1 Hz), already more than halved to 4.3 Hz by $\sigma = 0.5$ ms and essentially zero (0 Hz) by $\sigma \approx 5$ ms, below $\tau_{\text{GABA}} = 6$ ms. **Accuracy (red squares, right axis) holds at ≈ 89.2% up to $\boldsymbol{\sigma = 0.5}$ ms**, then collapses through 83.4% (σ = 1), 65.4% (σ = 2) and 17.7% (σ = 5), bottoming at chance (10.6%) by $\sigma = 9$ ms. The grey trace is the realised mean I rate, held flat near 53.2 Hz across the sweep: the E collapse happens under matched inhibition. The $\sigma \to \infty$ asymptote is the rate-matched Poisson regime: E silent, accuracy at chance.

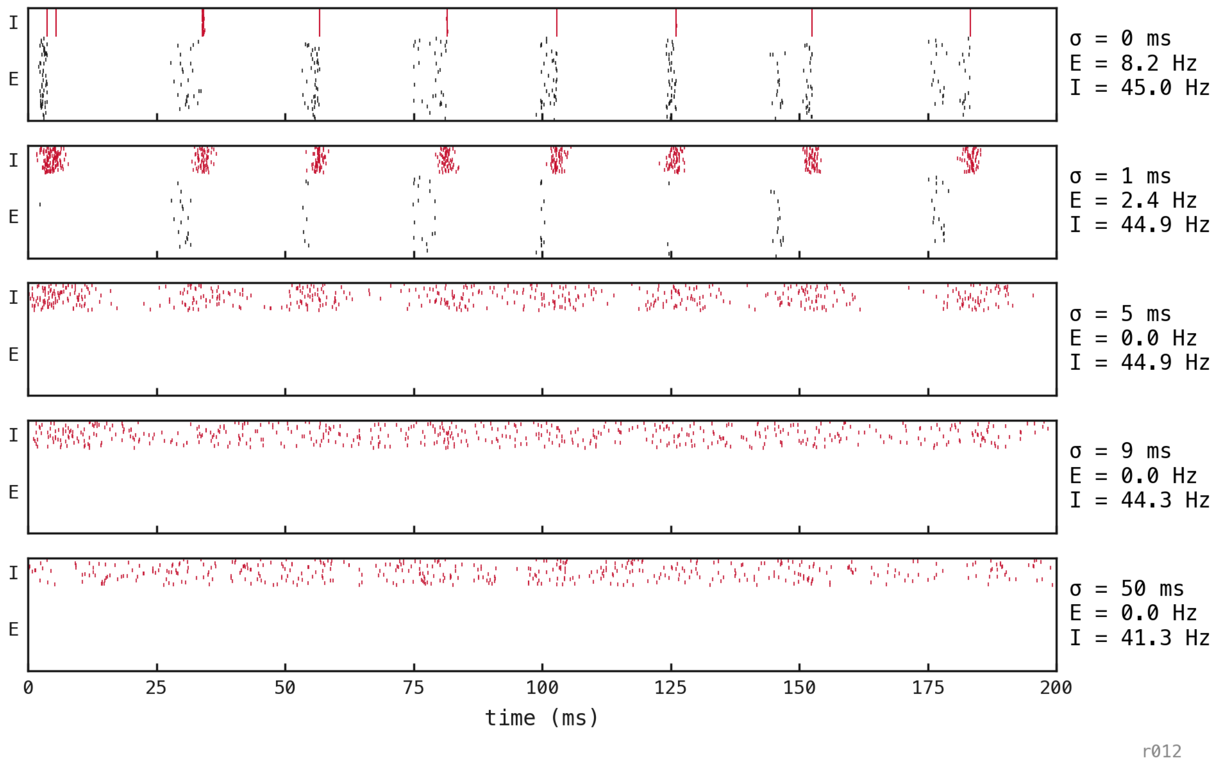

Figure 3: Single trial replayed at five per-cell jitter levels. At $\sigma = 0$ the I-bursts are crisp vertical bands. **At $\boldsymbol{\sigma = 1}$ ms** the bursts visibly smear into a few-ms-wide cluster and E firing already collapses to the low single digits (per-trial E annotated on each panel; sweep mean 1.7 Hz). **At $\boldsymbol{\sigma \geq 5}$ ms** the I-stream looks indistinguishable from a continuous low-variance shunt, and E is silenced. Per-cell jitter doesn't release E; it destroys the bursty structure that gave E its recovery troughs in the first place.

Cycle-coherent jitter

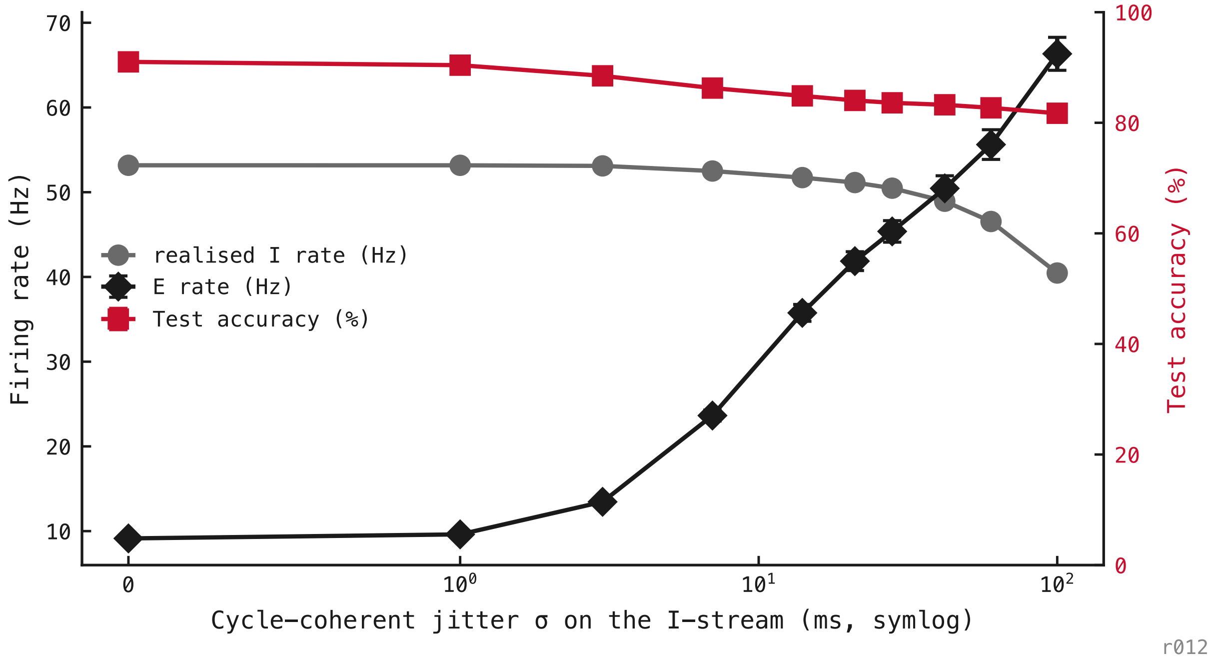

Figure 4: E rate (black diamonds) and accuracy (red squares) vs cycle-coherent jitter σ, three seeds. As σ grows the displaced bursts open wider gaps and the E rate climbs from baseline (9.1 Hz) *past* the full phase-shuffle level (25.9 Hz, the reference with within-burst structure destroyed) to 66.3 Hz by $\sigma = 100$ ms, with the sharpest rise near the predicted transition timescale $\sigma = 1/f_\gamma \approx 22.8$ ms; accuracy declines only gently, holding near 81.7%. The grey trace is the realised mean I rate: it holds within 3% of baseline through $\sigma \approx 14$ ms — where the E rate has already risen to 35.8 Hz — then droops to 40.5 Hz (24% below baseline) by $\sigma = 100$ ms, as the finite trial window truncates the most-displaced bursts. The strict same-mean-inhibition comparison is read at the smaller σ, where the rate is matched and the E rise is already unambiguous.

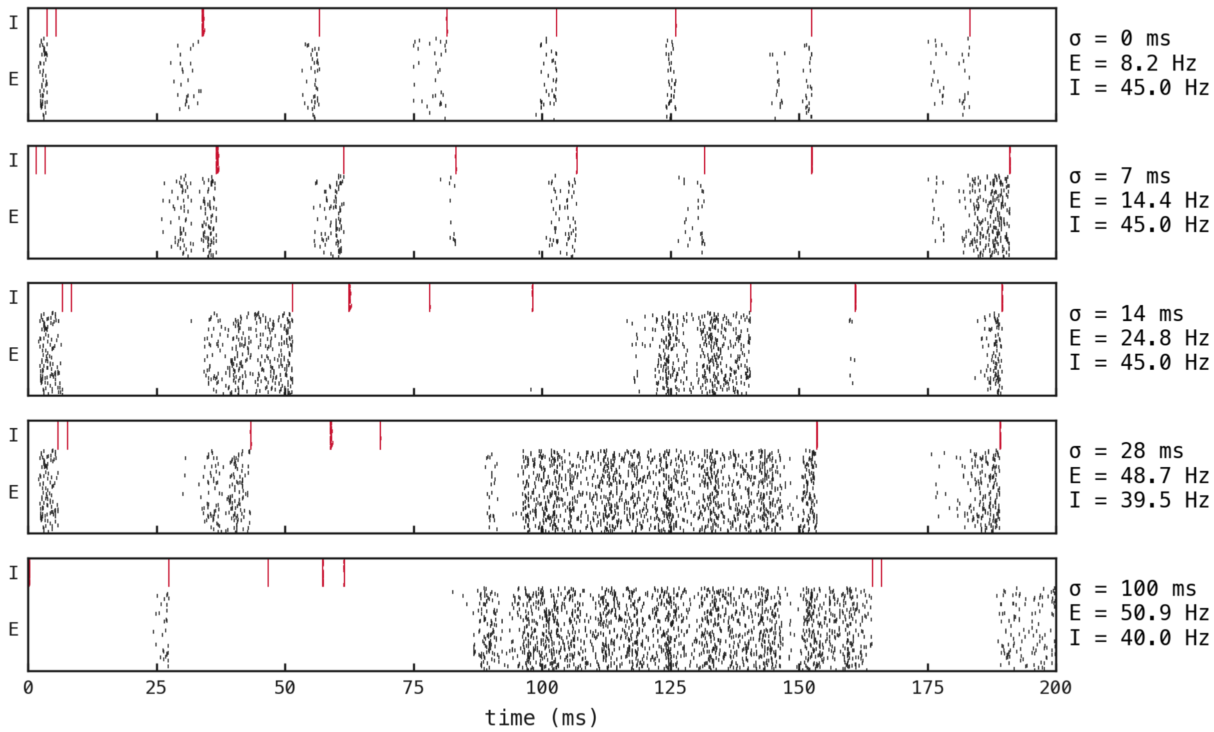

Figure 5: Single trial replayed at five jitter levels ($\sigma = 0, 7, 14, 28, 100$ ms; seed 42, MNIST digit 0 sample 0). Per-trial E rate annotated on each panel. **The I-bands stay vertical and crisp at every $\boldsymbol{\sigma}$:** within-burst synchrony is preserved exactly. What changes is *where* each burst lands: at larger σ the bursts are displaced bodily from their phase-locked positions, opening longer gaps in the I-stream that E fires through, and the E rate climbs accordingly.

Next steps

Toroidal (wrapped) jitter, to extend the rate-matched range and disentangle release from truncation. The strict same-mean-inhibition claim is now anchored at $\sigma = 14$ ms, where realised I holds within 3% of baseline on both arms and the E rate has already risen to 35.8 Hz — the qualitative result stands on rate-matched ground. Beyond $\sigma \approx 30$ ms, though, the cycle-coherent E-rate rise and the realised-I droop become confounded: some of the extra E firing is genuine gap-opening, and some is simply less inhibition delivered, because the finite window clamps and loses the most-displaced bursts. Wrapping each block offset modulo the trial length would restore exact spike-count preservation at every σ and separate the two, at the cost of re-injecting a wrapped burst at the opposite trial edge — a phase artifact of its own, so the wrapped sweep is a robustness check, not a replacement. The prediction: if the wrapped E rate still climbs past the phase-shuffle level (25.9 Hz), the release at large σ is real; if it flattens there, part of the 66.3 Hz overshoot at $\sigma = 100$ ms was the truncation. Either way the anchored rhythm-vs-mean conclusion is unaffected.