DMFT model

exp033 · 28 May 2026

Abstract

A mean-field account of exp025′s recruitment cliff. In the 4D conductance DMFT the cliff is a Hopf bifurcation of the silent fixed point. With COBANet's own LIF f-I curve as the population gain and couplings read off the biophysics (no fitted scale), the Jacobian eigenvalues place the threshold at $I^*_{\text{ext}} = 0.6$ nA, the crossing pair's imaginary part gives a gamma rhythm at $f^* \approx 27.8$ Hz set by the synaptic timescales, and a hysteresis sweep shows the onset is supercritical: continuous and reversible.

Methods

1. The 4D DMFT model

1.1 Summary of COBA model.

We start from the COBANet model (ar003 §2): conductance-based E and I membranes, a threshold-reset rule, and three exponential synapses (no E→E; I receives no inhibition):

$$C_m^E \dot{V}^E = -g_L^E(V^E - E_L) - g_e^E(V^E - E_e) - g_i^E(V^E - E_i) \quad (1)$$

$$C_m^I \dot{V}^I = -g_L^I(V^I - E_L) - g_e^I(V^I - E_e) \quad (2)$$

$$s_{t+1} = \mathbb{1}[V \geq V_{\text{th}}], \quad V \leftarrow V_{\text{reset}} \text{ if } \; s_{t+1} = 1 \;\text{ or refractory} \quad (3)$$

$$g^E_{e,t+1} = e^{-\Delta t/\tau_{\text{AMPA}}} g^E_{e,t} + W_{\text{in}} s^{\text{inp}}_t \quad (4)$$

$$g^E_{i,t+1} = e^{-\Delta t/\tau_{\text{GABA}}} g^E_{i,t} + W_{\text{ie}} s^i_t \quad (5)$$

$$g^I_{e,t+1} = e^{-\Delta t/\tau_{\text{AMPA}}} g^I_{e,t} + W_{\text{ei}} s^e_t \quad (6)$$

1.2. Continuous-time form with tonic drive.

Recast in continuous time. The synapses (4)–(6) are the exp-Euler form of first-order filters $\tau \dot{g} = -g + \tau \sum W s$, used here as ODEs. At a constant input rate g_e^E (4) settles to a steady mean, so its excitatory current into the E membrane (1) is a near-constant depolarising drive; we replace it by a tonic current I_{ext}, the swept control parameter. The E membrane then carries I_{ext} in place of g_e^E:

$$C_m^E \dot{V}^E = -g_L^E(V^E - E_L) - g_i^E(V^E - E_i) + I_{\text{ext}} \quad (7)$$

$$C_m^I \dot{V}^I = -g_L^I(V^I - E_L) - g_e^I(V^I - E_e) \quad (8)$$

Here g_i^E is the inhibition onto E and g_e^I the excitation onto I, each a continuous-time exponential filter of the presynaptic spikes:

$$\tau_{\mathrm{AMPA}}\dot{g}_e^I = -g_e^I + \tau_{\mathrm{AMPA}} W^{EI} s^E(t) \quad (9)$$

$$\tau_{\mathrm{GABA}}\dot{g}_i^E = -g_i^E + \tau_{\mathrm{GABA}} W^{IE} s^I(t) \quad (10)$$

with $s^E(t), s^I(t)$ the population spike trains and W^{EI}, W^{IE} the recurrent weight matrices; (9)–(10) are the continuous forms of (5)–(6).

1.3. Homogeneous coupling and population means.
Now resolve the populations: index E cells by $j \in \{1, ..., N_E\}$ and I cells by $k \in \{1, ..., N_I\}$. The recurrent drive in (9)–(10) is the presynaptic sum $W^{EI}s^E = \sum_j W^{EI}_{kj} s^E_j$ (and $W^{IE}s^I = \sum_k W^{IE}_{jk} s^I_k$). Replace each random weight by its population mean, $W^{EI}_{kj} \to w^{EI}$ and $W^{IE}_{jk} \to w^{IE}$; the sums become

$$\sum_j W^{EI}_{kj} s^E_j \longrightarrow w^{EI} \sum_j s^E_j = w^{EI} N_E E(t) \quad (11)$$

$$\sum_k W^{IE}_{jk} s^I_k \longrightarrow w^{IE} \sum_k s^I_k = w^{IE} N_I I(t) \quad (12)$$

introducing the *population-mean firing rates*

$$E(t) \equiv \frac{1}{N_E} \sum_{j=1}^{N_E} s^E_j(t), \quad I(t) \equiv \frac{1}{N_I} \sum_{k=1}^{N_I} s^I_k(t). \quad (13)$$

A *smooth-rate ansatz* (short-window averaging) treats $E(t), I(t)$ as continuous, dropping weight heterogeneity and finite-size noise, the shot noise $\mathrm{Var}[E(t)] \propto E(t)/N_E$ that vanishes only as $N_E, N_I \to \infty$. At finite $N_E = 1024, N_I = 256$ the residual $O\big(N^{-1/2}\big)$ fluctuations smear the Hopf onset and sustain a weak noisy gamma below threshold, both seen in the spiking simulations.

With no cell index left, every E cell sees the same g_i^E and every I cell the same g_e^I, collapsing the per-cell conductances to population means. Defining lumped couplings

$$\widetilde{W}^{EI} \equiv w^{EI} N_E, \quad \widetilde{W}^{IE} \equiv w^{IE} N_I \quad (14)$$

the conductance dynamics become

$$\tau_{\mathrm{AMPA}}\dot{g}_e^I = -g_e^I + \tau_{\mathrm{AMPA}} \widetilde{W}^{EI} E \quad (15)$$

$$\tau_{\mathrm{GABA}}\dot{g}_i^E = -g_i^E + \tau_{\mathrm{GABA}} \widetilde{W}^{IE} I \quad (16)$$

Two equations, down from $N_E + N_I$. (The fan-in scale $\widetilde{W}$ folds into W^{EI}, W^{IE} in 1.6.)

Running system, end of 1.3: conductances are now two population means; the membrane is still per-cell but sees those means:

$$\begin{aligned}
C_m^E \dot{V}_j^E &= -g_L^E\big(V_j^E - E_L\big) - g_i^E\big(V_j^E - E_i\big) + I_{\mathrm{ext}} \\
C_m^I \dot{V}_k^I &= -g_L^I\big(V_k^I - E_L\big) - g_e^I\big(V_k^I - E_e\big) \\
\tau_{\mathrm{AMPA}}\dot{g}_e^I &= -g_e^I + \tau_{\mathrm{AMPA}} \widetilde{W}^{EI} E \\
\tau_{\mathrm{GABA}}\dot{g}_i^E &= -g_i^E + \tau_{\mathrm{GABA}} \widetilde{W}^{IE} I
\end{aligned}$$

1.4. Driving-force linearisation.
The synaptic current is conductance times a *driving force*, $-g(V - E_{\text{rev}})$, a g–V product, hence nonlinear. Freeze V at rest, $V_{\text{rest}} = E_L = -65$ mV, *in the driving force only* (leak and threshold keep their full V-dependence, handled by the f-I curve in 1.5). Each driving force becomes a fixed voltage gap:

$$\Delta V_{\text{inh}} \equiv V_{\text{rest}} - E_i = -65 - (-80) = 15 \text{ mV} \quad (17)$$

$$\Delta V_{\text{exc}} \equiv V_{\text{rest}} - E_e = -65 - 0 = -65 \text{ mV}$$

The synaptic currents in (7)–(8) then lose their V-dependence and become proportional to conductance alone:

$$-g_i^E\big(V_j^E - E_i\big) \approx -g_i^E \Delta V_{\text{inh}} \quad (18)$$

$$-g_e^I\big(V_k^I - E_e\big) \approx -g_e^I \Delta V_{\text{exc}} = +g_e^I \cdot |E_e - V_{\text{rest}}| \quad (19)$$

(inhibition pulls V down, $\Delta V_{\text{inh}} = +15$ mV; excitation pushes it up, $|\Delta V_{\text{exc}}| = 65$ mV). Removing the g–V coupling reduces COBA to a current-based (CUBA) form; the cost is shunting: with V fixed we ignore that conductance also lowers the effective time constant ($\tau_{\text{eff}} = C_m/g_{\text{tot}}$).

Running system, end of 1.4: the synaptic currents are now linear in conductance (no V left in the driving force):

$$\begin{aligned}
C_m^E \dot{V}_j^E &= -g_L^E\big(V_j^E - E_L\big) - g_i^E \Delta V_{\text{inh}} + I_{\text{ext}} \\
C_m^I \dot{V}_k^I &= -g_L^I\big(V_k^I - E_L\big) + g_e^I \ |\Delta V_{\text{exc}}| \\
\tau_{\text{AMPA}} \dot{g}_e^I &= -g_e^I + \tau_{\text{AMPA}} \widetilde{W}^{EI} E \\
\tau_{\text{GABA}} \dot{g}_i^E &= -g_i^E + \tau_{\text{GABA}} \widetilde{W}^{IE} I
\end{aligned}$$

1.5. Population rate from an f-I curve.
Under (17)–(19) the membrane equations (7)–(8) read $C_m \dot{V} = -g_L(V - E_L) + I_{\text{syn}}$, LIF with a synaptic current. A LIF cell under constant net current I fires at its f-I rate $\varphi(I)$; replacing each cell's spikes by that rate gives

$$E(t) \approx \varphi_E\big(I_{\text{eff}}^E(t)\big), \quad I(t) \approx \varphi_I\big(I_{\text{eff}}^I(t)\big) \quad (20)$$

with effective input currents (from (7)–(8) with (17)–(19) substituted)

$$I_{\text{eff}}^E(t) = I_{\text{ext}}(t) - g_i^E(t) \Delta V_{\text{inh}} \quad (21)$$

$$I_{\text{eff}}^I(t) = g_e^I(t) \ |\Delta V_{\text{exc}}| \quad (22)$$

(I receives only excitation; E receives the drive minus the GABA shunt.) The instantaneous-rate replacement (20) holds only for slow inputs; in reality $E(t)$ relaxes toward the f-I fixed point on τ_E, which we encode explicitly:

$$\tau_E \dot{E} = -E + \Phi_E(I_{\text{ext}} - g_i^E \Delta V_{\text{inh}}) \quad (23)$$

$$\tau_I \dot{I} = -I + \Phi_I(g_e^I \ |\Delta V_{\text{exc}}|) \quad (24)$$

where Φ_E, Φ_I are the smooth steady-state gain functions (COBANet's LIF f-I curve in the numerics; see Locating the Hopf). Two more equations down, together with (15)–(16), four equations in (E, I, g_e^I, g_i^E).

Running system, end of 1.5: a closed 4D rate model in (E, I, g_e^I, g_i^E), constants not yet absorbed:

$$\begin{aligned}
\tau_E \dot{E} &= -E + \Phi_E(I_{\text{ext}} - g_i^E \Delta V_{\text{inh}}) \\
\tau_I \dot{I} &= -I + \Phi_I(g_e^I \, |\Delta V_{\text{exc}}|) \\
\tau_{\text{AMPA}} \dot{g}_e^I &= -g_e^I + \tau_{\text{AMPA}} \widetilde{W}^{EI} E \\
\tau_{\text{GABA}} \dot{g}_i^E &= -g_i^E + \tau_{\text{GABA}} \widetilde{W}^{IE} I
\end{aligned}$$

1.6. Absorb the driving-force constants.
The prefactors $\Delta V_{\text{inh}}, |\Delta V_{\text{exc}}|$ in (23)–(24) and the fan-in scalings in (15)–(16) are constants carrying no dynamics; fold them into the couplings:

$$W^{EI} \equiv \widetilde{W}^{EI} \cdot |\Delta V_{\text{exc}}|, \quad W^{IE} \equiv \widetilde{W}^{IE} \cdot \Delta V_{\text{inh}} \quad (25)$$

and absorb $|\Delta V_{\text{exc}}|$ into the I-cell argument by redefining $g_e^I \mapsto g_e^I \cdot |\Delta V_{\text{exc}}|$ (similarly g_i^E). The conductances now carry current units and the f-I curves take their argument directly: a change of variables, no dynamics lost.

1.7 The 4D system
After 1.1–1.6, the mean-field equations are

$$\tau_E \dot{E} = -E + \Phi_E(I_{\text{ext}} - g_i^E), \quad \tau_I \dot{I} = -I + \Phi_I(g_e^I) \quad (26)$$

$$\tau_{\text{AMPA}} \dot{g}_e^I = -g_e^I + \tau_{\text{AMPA}} W^{EI} E, \quad \tau_{\text{GABA}} \dot{g}_i^E = -g_i^E + \tau_{\text{GABA}} W^{IE} I \quad (27)$$

in state (E, I, g_e^I, g_i^E). Whether 4D is minimal, or a 2D reduction would do, is settled in the appendix, Is it really 4D?

1.8 4D Jacobian
At a fixed point $(E^*, I^*, g_e^{I*}, g_i^{E*})$:

$$J = \begin{pmatrix} -1/\tau_E & 0 & 0 & -\Phi_E'/\tau_E \\ 0 & -1/\tau_I & \Phi_I'/\tau_I & 0 \\ W^{EI} & 0 & -1/\tau_{\text{AMPA}} & 0 \\ 0 & W^{IE} & 0 & -1/\tau_{\text{GABA}} \end{pmatrix} \quad (28)$$

with Φ_E', Φ_I' evaluated at the fixed-point arguments.

2. Locating the Hopf

2.1 The Hopf condition and sweep
Below the cliff the network is silent: nudge it and the perturbation dies away. Above the cliff the same nudge grows into a sustained rhythm. The switch is a Hopf bifurcation of the silent fixed point: the smallest drive I_{ext}^* at which the fixed point stops damping oscillations and starts amplifying them.

Linear stability makes this precise. Each mode of the linearised system evolves as $e^{\lambda t}$, where λ is an eigenvalue of the Jacobian (28): a negative real part decays, a positive one grows. The Hopf is the instant an oscillating mode (a complex-conjugate pair) crosses from the left half-plane to the right ($\mathrm{Re}\,\lambda = 0$, $\mathrm{Im}\,\lambda \neq 0$), so the silent state gives way to a growing oscillation rather than a static shift. To find it we sweep I_{ext}, track the fixed point, and diagonalise J at each step; I^*_{ext} is the first drive where the leading pair reaches the axis (Figure 2). The local analysis holds only if a single pair crosses while the rest stay damped (a simple Hopf); two pairs crossing at once (a double-Hopf) would seed tori the linearisation cannot see. Timescales: $\tau_E = 20$, $\tau_I = 5$, $\tau_{\mathrm{AMPA}} = 2$, $\tau_{\mathrm{GABA}} = 6$ ms.

2.2 Gain and couplings from COBANet

The gain Φ and the couplings come from COBANet, not from a fit. For Φ we use the LIF f-I curve: a cell under white-noise input of mean μ and standard deviation σ_V fires at the Siegert rate

$$\varphi(\mu) = \left[\tau_{\mathrm{ref}} + \tau_m\sqrt{\pi}\int_{(V_{\mathrm{reset}}-\mu_V)/\sigma_V}^{(V_{\mathrm{th}}-\mu_V)/\sigma_V} e^{u^2}(1+\mathrm{erf}\,u)\,\mathrm{d}u\right]^{-1}, \quad \mu_V = E_L + \mu/g_L \qquad (29)$$

with COBANet's $\tau_m, g_L, V_{\mathrm{th}}, V_{\mathrm{reset}}, \tau_{\mathrm{ref}}$ per population. The couplings are $W^{EI} = w^{EI}N_E\,|\Delta V_{\mathrm{exc}}|$ and $W^{IE} = w^{IE}N_I\Delta V_{\mathrm{inh}}$ (eqs 14, 25): the fan-in normalisation ($w \mapsto w/N$) makes $w^{EI}N_E, w^{IE}N_I$ the ei-strength values s and rs ($\approx$ 1 and 2 µS), times the driving forces (17). With Φ in physical units I_{ext} is a current in nanoamps: the absolute scale is fixed by the biophysics, not chosen. The membrane-noise std σ_V is the one free parameter, set to 4 mV; the located Hopf is insensitive to it (f^* moves under 1 Hz over $\sigma_V \in [3, 6]$ mV).

3. Calculating its frequency

At the crossing $\lambda = \pm i\omega^*$, so $f^* = \omega^*/2\pi$ is read from the Jacobian at threshold. To test whether inhibitory decay sets the period, we re-find the Hopf at each τ_{GABA} of exp041′s retrained sweep and compare f^* to the spiking f_γ (Figure 3).

4. Classifying sub- or supercritical

4.1 The hysteresis sweep

Above the Hopf a limit cycle exists; whether it appears gently or abruptly decides whether exp025′s recruitment cliff is the graded, reversible onset that picture assumes. We test it with a hysteresis sweep: step I_{ext} quasi-statically up through I^*_{ext} and back down, at each step integrating (26)–(27) to steady state from the previous step's end state (so any coexisting cycle is carried along), and record the peak-to-peak amplitude $A = \max_t E - \min_t E$.

A *stable* cycle born at threshold makes the rising and falling branches coincide, with A returning to zero at I^*_{ext}: a reversible, supercritical onset. An *unstable* cycle sitting below threshold instead acts as a basin boundary: the silent state jumps to a distant large-amplitude cycle that survives as the drive is lowered back, so the branches split into a hysteresis loop with a bistable window: subcritical (Figure 4).

4.2 The cycle waveform

The same integration gives the cycle's shape above onset: E and I are near-sinusoidal, with the E burst leading the I burst by the round-trip synaptic delay: E recruits I, I shunts E a few milliseconds later (Figure 5).

Results

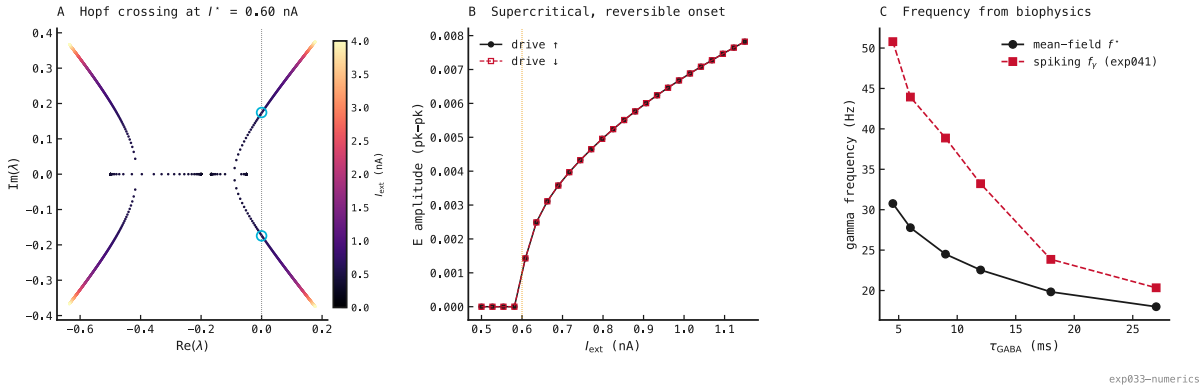

Figure 1: Summary of the analysis; the detailed panels follow. **A**, the Hopf: one eigenvalue pair crosses into the right half-plane at $I^* = 0.6$ nA, fixing a gamma rhythm at $f^* \approx 27.8$ Hz (full plot, Figure 2). **B**, the onset is supercritical and reversible, the up/down branches coinciding (Figure 4). **C**, the predicted frequency falls with τ_{GABA}, tracking the spiking measurement (Figure 3). exp025′s recruitment cliff is a supercritical Hopf, set by the synaptic time constants, derived below from COBANet's f-I curve and biophysical couplings with no fitted scale.

Frequency

Sweeping I_{ext} and diagonalising J at each step locates the Hopf at

$$I^*_{\mathrm{ext}} = 0.6 \text{ nA}, \quad \omega^* = 0.175 \text{ rad/ms}, \quad f^* = \frac{\omega^*}{2\pi} \approx 27.8 \text{ Hz},$$

in the PING gamma band, from COBANet's f-I curve and biophysical couplings with no fitted scale. One complex pair crosses the imaginary axis while the others stay damped, a simple Hopf (Figure 2). The predictive content is the *dependence* on the synaptic timescales: across a τ_{GABA} sweep f^* tracks the spiking network's gamma qualitatively, not quantitatively (Figure 3).

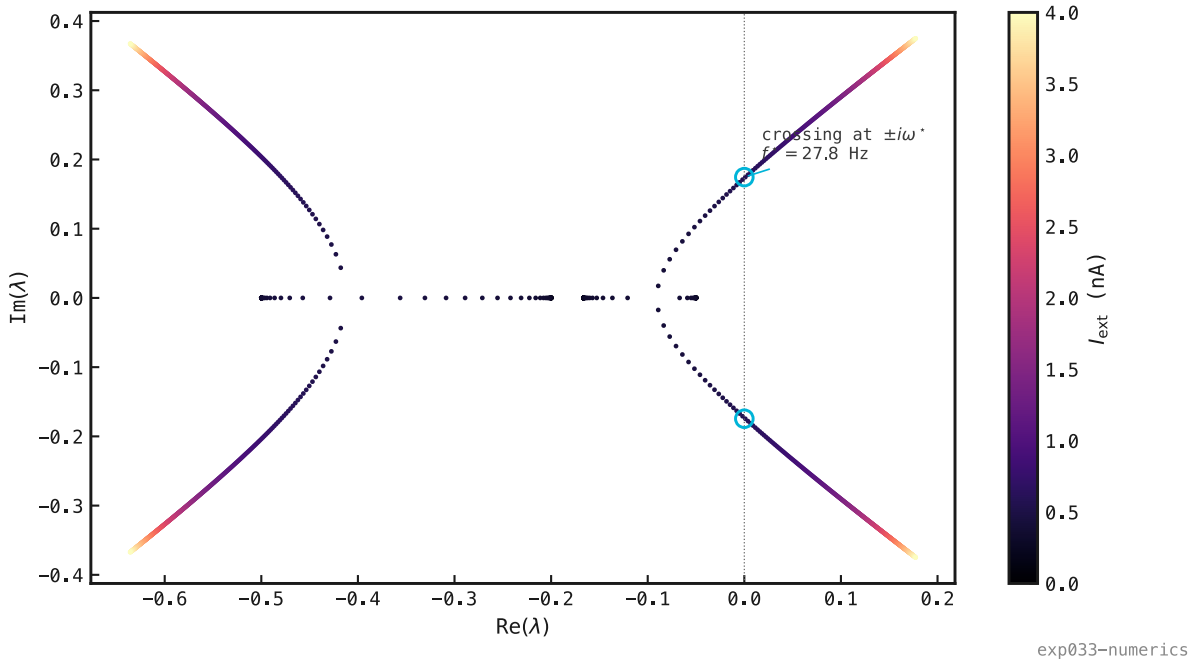

Figure 2: **How to read it.** Each dot is one eigenvalue λ of the Jacobian (28). Its horizontal position is the growth rate Re λ (negative means that mode decays, positive means it grows), and the dotted line at Re $\lambda = 0$ is the stability boundary: everything to its left is damped. Its vertical position is the oscillation frequency Im λ (zero = a non-oscillating mode, on the real axis). The Jacobian is 4×4, so each drive contributes four dots; sweeping I_{ext} (colour, dark→bright) drags them along the curves shown. Follow the upper and lower arms: one complex-conjugate pair lifts off the real axis and marches rightward, touching the boundary at $\pm i\omega^*$ (cyan circles) when $I^*_{\text{ext}} = 0.6$ nA. That crossing is the Hopf, where the network tips from damping to amplifying, and the height of the crossing sets the rhythm, $f^* = \omega^*/2\pi \approx 27.8$ Hz. The other two eigenvalues stay left of the line throughout, so a single oscillation goes unstable: a simple Hopf. (A second pair reaches the axis only at far higher drive, above the recruitment cliff.)

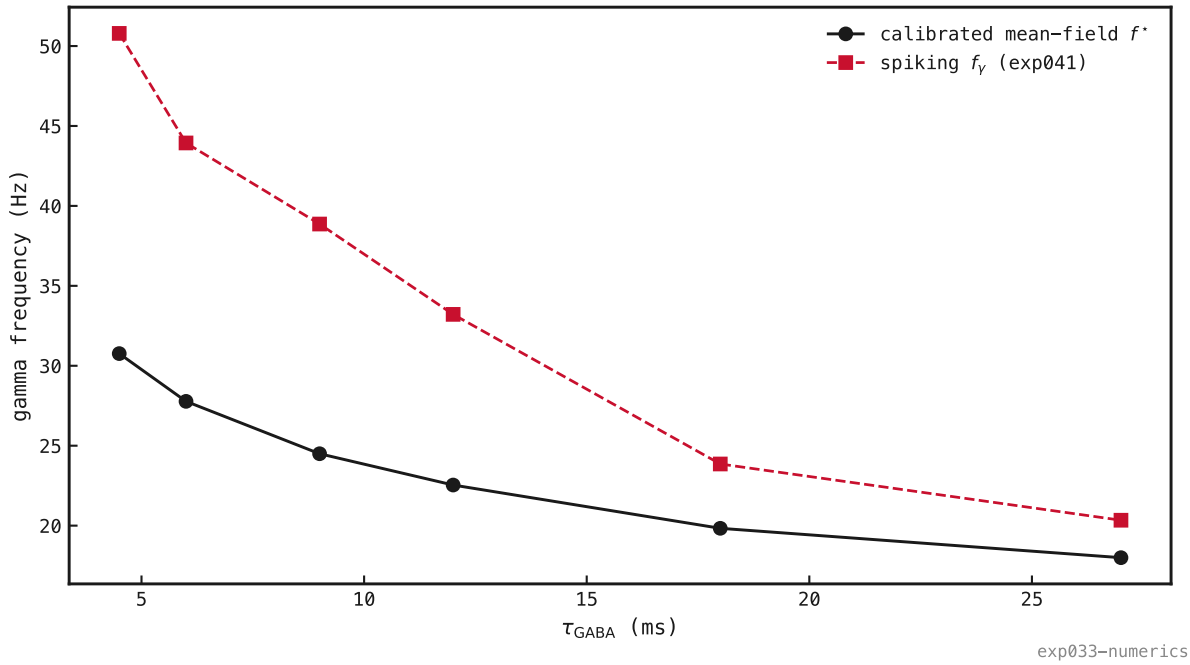

Figure 3: Both fall monotonically with τ_{GABA}: the inhibitory decay is the clock in both. The match is qualitative, not quantitative. The calibrated mean-field is *flatter* and runs below the spiking network across the sweep, the gap largest at short τ_{GABA} (27.8 vs 43.9 Hz at the canonical 6 ms) and narrowing as τ_{GABA} grows. The reduction captures the mechanism but not the full sensitivity, expected since it drops the spike-synchrony of the E-volley that sharpens the cycle most at short τ_{GABA}.

Super or subcritical

The rising and falling sweeps coincide (Figure 4): the amplitude grows continuously from zero at I^*_{ext} and retraces exactly on the way down, with no hysteresis (loop width 0 nA, branch gap 2×10^{-6}) and no cycle coexisting with the silent state. This is a supercritical Hopf, the recruitment cliff of exp025. The rising branch obeys the predicted $A^2 \propto (I_{\mathrm{ext}} - I^*_{\mathrm{ext}})$ (slope 1.1×10^{-4}, $R^2 = 0.999$), so $A \propto \sqrt{I_{\mathrm{ext}} - I^*_{\mathrm{ext}}}$ and $\mathrm{d}A/\,\mathrm{d}I_{\mathrm{ext}}$ diverges at threshold: most of the amplitude appears in a narrow band of drive above I^*_{ext}.

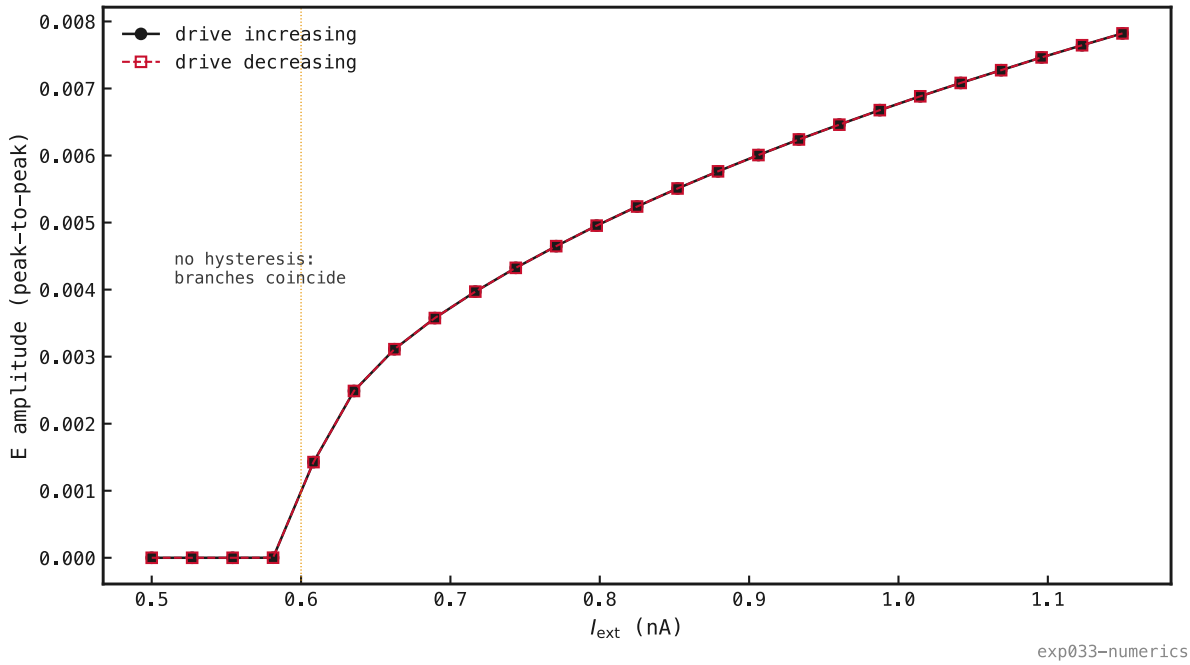

Figure 4: Quasi-static up/down ramp of the drive across I^*_{ext}. The rising branch (black, drive increasing) and falling branch (red, drive decreasing) coincide: the cycle turns on and off at the same drive, with amplitude growing continuously from zero. No loop, no bistable window: the reversible onset of a supercritical Hopf. (A subcritical onset would show the two branches separating into a hysteresis loop.)

Above onset, E leads I by ≈ 5.3 ms, the loop delay the 2D reduction omits.

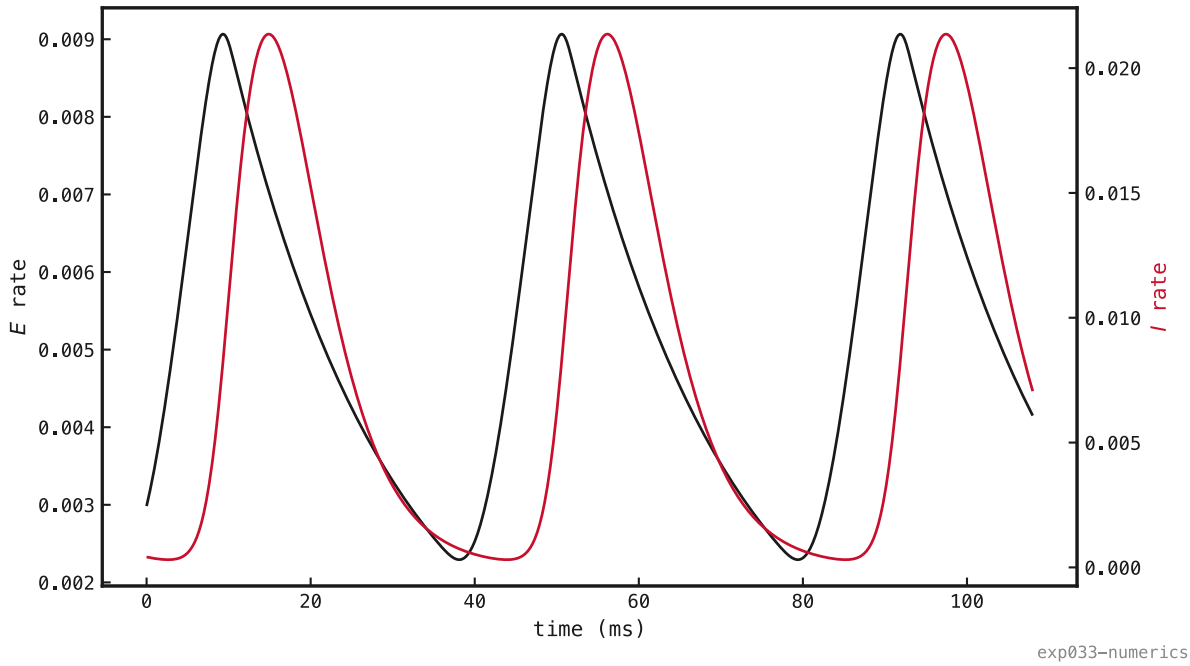

Figure 5: E (black) and I (red) over the limit cycle at $I_{\text{ext}} = I^* + 0.4$ nA. The E burst leads the I burst by ≈ 5.3 ms: E recruits I, I shunts E a few milliseconds later, and the cycle repeats, the loop delay the 2D reduction omits.

Appendix

Is it really 4D?

A Hopf gives a closed-loop trajectory, so the long-run motion is planar, so there should be a 2D description. This appendix collects the attempts. Verdict: a 2D *description* exists (a centre manifold), but the *vector field* does not reduce below three.

1. **A Hopf gives planar (two-dimensional) dynamics, so 4D looks suspect.** A limit cycle is a closed loop, traversed with two coordinates (amplitude and phase), so the motion lies in a plane. If it is 2D, a 2D model should exist; steps 3–5 hunt for it.

2. **The geometry confirms the motion is 2D (Figures 6–7).** The four variables cycle in loop order (E, then g_e^I, then I, then g_i^E, each lagging the last, Figure 6). Projected onto every variable pair (Figure 7), the cycle is one closed loop on a thin 2D sheet, a *centre manifold* (the loop is a 1D curve on it). Only (E, g_e^I) nearly collapses to a line (AMPA trails the E rate); the rest enclose real area, so no variable pair can stand in for the sheet.

3. **Route A: the textbook Wilson-Cowan model (slave the conductances).** The standard 2D tool is two rates with instantaneous coupling, the 4D model with infinitely fast synapses. Slave each conductance to its filter's steady value (15)–(16), $g_e^I = \tau_{\text{AMPA}} W^{EI} E$ and $g_i^E = \tau_{\text{GABA}} W^{IE} I$, and substitute into (26):

$$\tau_E \dot{E} = -E + \Phi_E(I_{\text{ext}} - \tau_{\text{GABA}} W^{IE} I), \quad (30)$$

$$\tau_I \dot{I} = -I + \Phi_I(\tau_{\text{AMPA}} W^{EI} E). \quad (31)$$

   Its divergence (the Jacobian trace),

$$\frac{\partial \dot{E}}{\partial E} + \frac{\partial \dot{I}}{\partial I} = -\frac{1}{\tau_E} - \frac{1}{\tau_I} < 0, \quad (32)$$

   is a negative constant, so Bendixson–Dulac forbids a periodic orbit: no Hopf, for any drive or coupling. It has dropped the round-trip delay (E excites I after τ_{AMPA}, I shunts E after τ_{GABA}): zero-lag inhibition damps but cannot overshoot. The reduction is valid only if synapses are fast relative to the rhythm, which they are not ($\tau_{\text{GABA}} \approx 6$ ms is the gamma period's order).

4. **Route B: quasi-steady-state the rates instead.** The dual move: slave the rates, $E = \Phi_E(I_{\text{ext}} - g_i^E)$ and $I = \Phi_I(g_e^I)$, into the conductance equations (27), giving a 2D system in (g_e^I, g_i^E):

$$\tau_{\text{AMPA}} \dot{g}_e^I = -g_e^I + \tau_{\text{AMPA}} W^{EI} \Phi_E(I_{\text{ext}} - g_i^E), \quad (33)$$

$$\tau_{\text{GABA}} \dot{g}_i^E = -g_i^E + \tau_{\text{GABA}} W^{IE} \Phi_I(g_e^I), \quad (34)$$

   with the same negative-constant divergence,

$$\frac{\partial \dot{g}_e^I}{\partial g_e^I} + \frac{\partial \dot{g}_i^E}{\partial g_i^E} = -\frac{1}{\tau_{\text{AMPA}}} - \frac{1}{\tau_{\text{GABA}}} < 0, \quad (35)$$

   so no cycle: it rings down (Figure 8). (These rates are the membrane variables, already reduced to an f-I rate.)

5. **Route C: lump into fast and slow timescales.** Slave the two fastest variables, the AMPA conductance ($\tau_{\text{AMPA}} = 2$ ms) and the I rate ($\tau_I = 5$ ms), keeping the two slowest $\{E, g_i^E\}$ ($\tau_{\text{GABA}} = 6$, $\tau_E = 20$ ms):

$$\tau_E \dot{E} = -E + \Phi_E(I_{\text{ext}} - g_i^E), \quad (36)$$

$$\tau_{\text{GABA}} \dot{g}_i^E = -g_i^E + \tau_{\text{GABA}} W^{IE} \Phi_I(\tau_{\text{AMPA}} W^{EI} E). \quad (37)$$

   Trace $-1/\tau_E - 1/\tau_{\text{GABA}} < 0$: no cycle. The split is forced anyway: the constants interleave, $\tau_{\text{AMPA}} = 2 < \tau_I = 5 < \tau_{\text{GABA}} = 6 < \tau_E = 20$ ms, so "fast" and "slow" each mix a conductance with a rate.

6. **All three fail for one structural reason.** The network is a pure ring: the single loop $E \to g_e^I \to I \to g_i^E \to E$, no recurrent E→E or I→I and no self-drive, so each variable's only diagonal Jacobian term is its own decay and every gain Φ' sits off-diagonal. Eliminate *any* two variables and the 2D trace is $-1/\tau_a - 1/\tau_b < 0$; Bendixson–Dulac then rules out a cycle. Routes A–C are three of the $\binom{4}{2} = 6$ ways to pick the kept pair; the runner sweeps all six and none crosses. A 2D model that *does* oscillate has added a destabiliser PING lacks: recurrent excitation or a cubic self-gain, i.e. a positive diagonal term. We do not add one: PING has no W^{EE} to supply it, and bolting one on would make a self-excitation oscillator (van der Pol / FitzHugh–Nagumo), not PING: a different mechanism for the gamma, not this network's.

7. **Three dimensions survive.** Slave only the fastest lag, the AMPA conductance $g_e^I = \tau_{\text{AMPA}} W^{EI} E$ (step 2′s near-degenerate pair), leaving a three-lag ring:

$$\tau_E \dot{E} = -E + \Phi_E(I_{\text{ext}} - g_i^E), \quad (38)$$

$$\tau_I \dot{I} = -I + \Phi_I(\tau_{\text{AMPA}} W^{EI} E), \quad (39)$$

$$\tau_{\text{GABA}} \dot{g}_i^E = -g_i^E + \tau_{\text{GABA}} W^{IE} I, \quad (40)$$

   This still Hopfs. Located like the 4D bifurcation (sweep I_{ext}, diagonalise the 3×3 Jacobian, find the complex-pair crossing), it gives $I_{\text{ext}}^* = 0.7$ nA and $f^* = 37$ Hz, both above the 4D values, since dropping the AMPA lag stiffens the loop. The same sweep finds no crossing for either 2D reduction. Figure 8: 4D and 3D sustain the rhythm; all three 2D reductions ring down.

8. **Resolution: the 2D description is the centre manifold, not a coordinate pair.** The 2D description that exists is the centre manifold, the plane of the two critical, oscillatory eigenvectors, tilted across all of (E, I, g_e^I, g_i^E): a plane in the *eigenbasis*, not any coordinate pair. So it is genuinely 2D yet unreachable by dropping two physical variables: a coordinate projection lands on the wrong plane, where the gains go off-diagonal and Bendixson–Dulac kills the cycle (steps 3–6). The *attractor* is a 1D loop on a 2D manifold; the *vector field* does not reduce below three, since a Hopf needs three first-order lags in the $E \to I \to E$ ring to build the destabilising phase. 4D is the natural model, 3D the floor, and the 2D centre manifold where the cycle lives, not a model in the original variables.

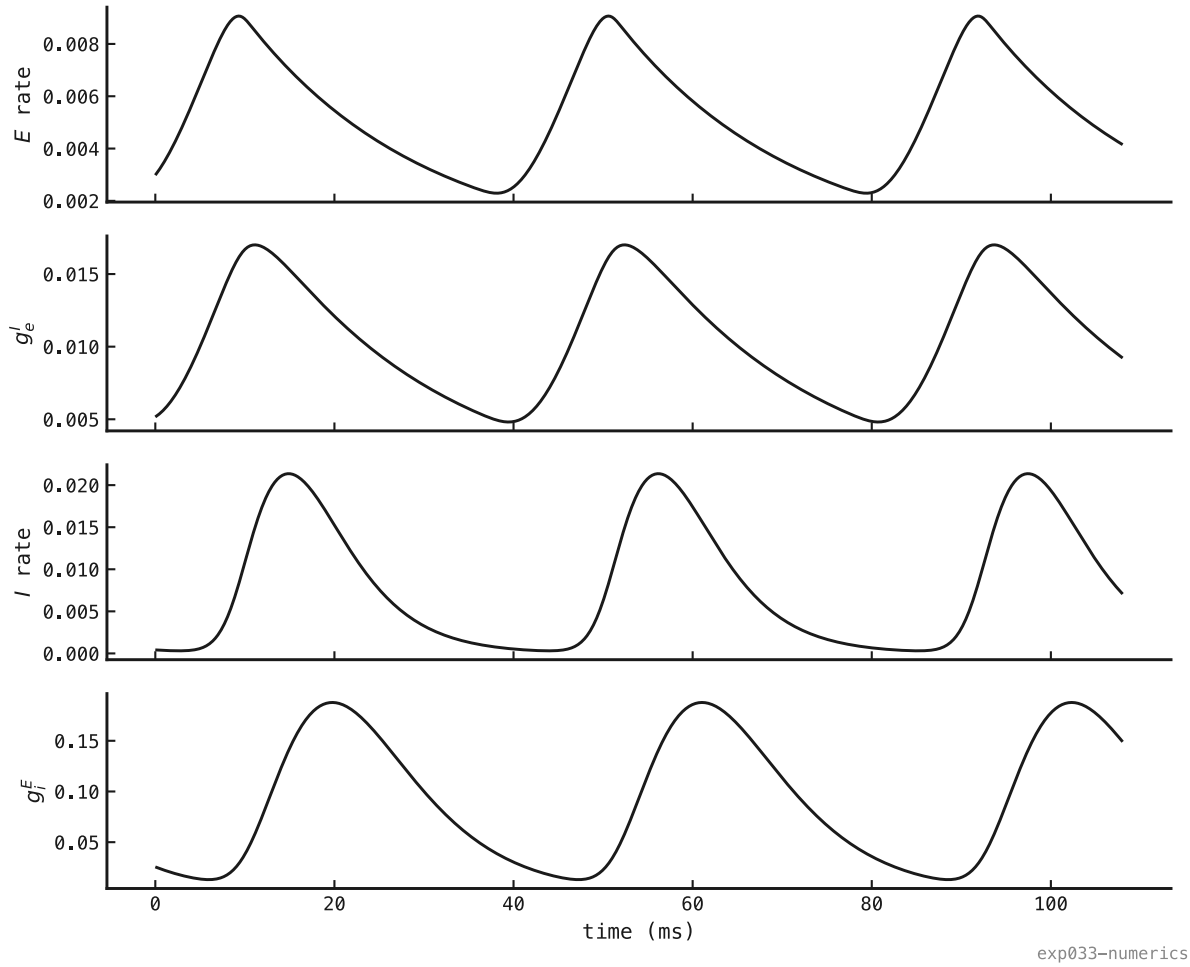

Figure 6: The four state variables over the limit cycle ($I_{\text{ext}} = I^* + 0.4$ nA), in loop order $E \to g_e^I \to I \to g_i^E$. Each peaks after the one above it: g_e^I tracks the E rate almost rigidly (the near-degenerate pair of step 2), then drives I, which fills g_i^E, which shunts E: one round trip of the ring per gamma cycle.

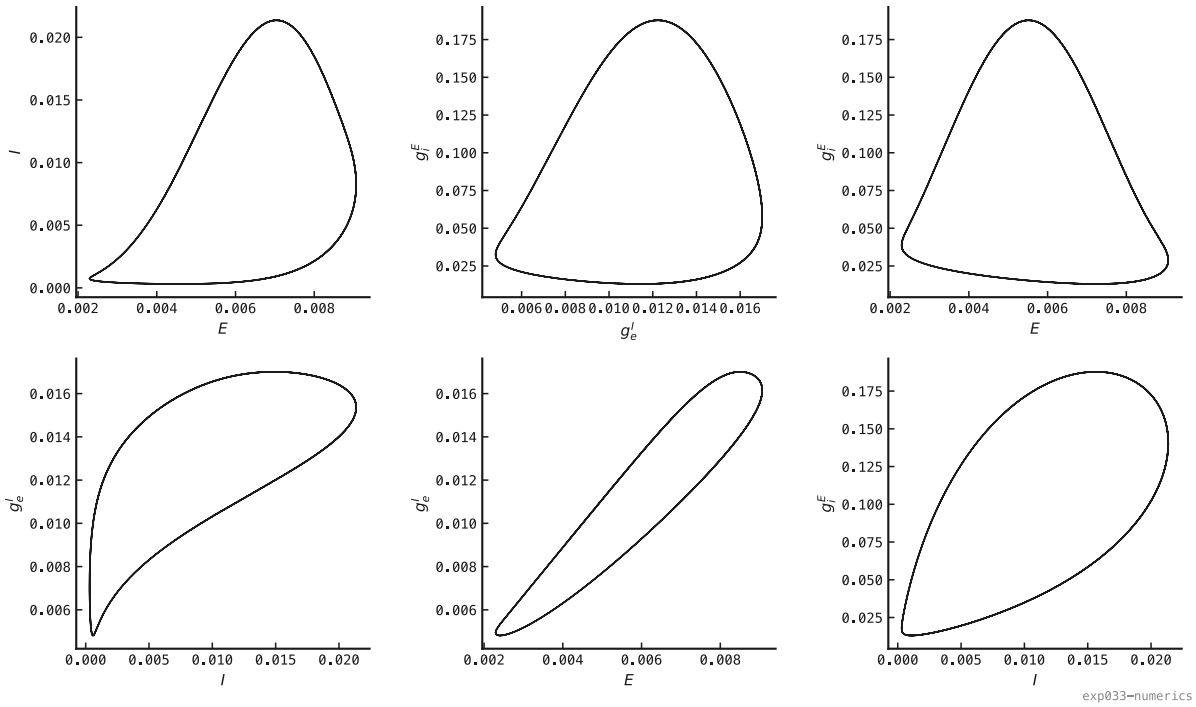

Figure 7: The 4D limit cycle ($I_{\text{ext}} = I^* + 0.4$ nA) projected onto all six variable pairs. Every projection is one closed loop: the trajectory lives on a 2D centre manifold. The (E, g_e^I) panel collapses almost to a line: the fast AMPA conductance tracks the E rate, which is why slaving it (the 3D reduction) preserves the dynamics, while the other pairs enclose genuine area and cannot be collapsed.

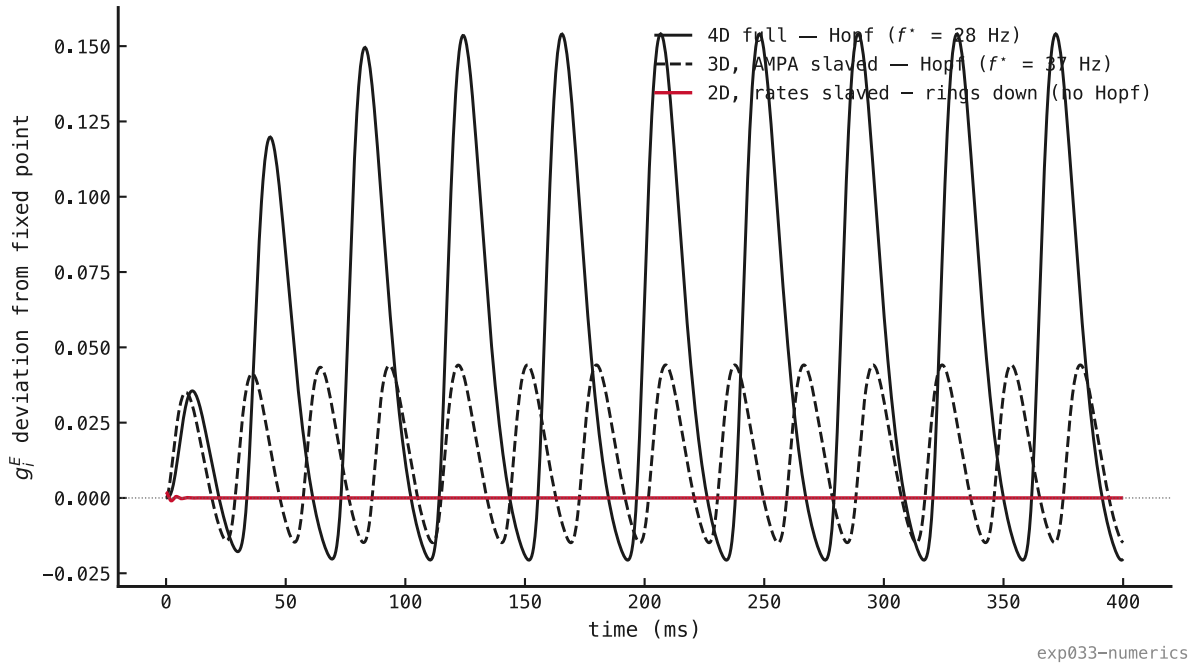

Figure 8: The reductions tested by simulation: g_i^E after a small kick at a common drive ($I_{\text{ext}} = 1$ nA, above every threshold). The 4D model (black, $f^* = 28$ Hz) and the 3D AMPA-slaved reduction (black dashed, $f^* = 37$ Hz) both sustain the limit cycle; the 2D rate-slaved reduction (red) rings down to its fixed point: no Hopf, as eq (35) requires.