

Supplementary Material for:

Time over noise: Asynchrony unlocks robust coordination
beyond behavioural fluctuations

Contents

S1 Supplementary Methods	2
S1.1 Beggs–Plenz Avalanche Analysis	2
S1.2 CCDF Computation	2
S1.3 Model Fitting and AIC Comparison	2
S1.4 Knee Detection (ARM)	3
S2 Supplementary Tables	4
S2.1 Table S1: SNM B&P Fitting Results	4
S2.2 Table S2: ARM B&P Fitting Results	5
S2.3 Table S3: FFT Spectral Exponent β	5
S2.4 Table S4: Knee-detection sensitivity grid	6
S3 Supplementary Figures	7
S3.1 Figure S1: SNM Avalanche CCDF with Model Fits	7
S3.2 Figure S2: ARM Avalanche CCDF with Knee Detection and Region Fits	8
S3.3 Figure S3: Topological Robustness across Network Models	9
S3.4 Figure S4: Hybrid Parameter Sweeps	10
S3.5 Figure S5: Finite-size scaling of the ARM solution plateau	10
S3.6 Figure S6: Solution rate landscape across noise parameter and time budget	12

S1 Supplementary Methods

S1.1 Beggs–Plenz Avalanche Analysis

Following the original Beggs–Plenz protocol [1], we define avalanches from the activity time series $A(t)$, where $A(t)$ is the number of nodes that changed their colour at step t (the *changed-nodes* signal). An avalanche is a maximal contiguous block of steps with $A(t) > \theta$, where the threshold $\theta = 1$. The avalanche size s is the total activity summed over the block: $s = \sum_{t \in \text{avalanche}} A(t)$. All avalanches are pooled across the 1,000 independent trials.

S1.2 CCDF Computation

Let $\{s_i\}$ be the pooled set of avalanche sizes. The complementary cumulative distribution function (CCDF) is estimated at each unique size s :

$$P(S \geq s) = \frac{|\{i : s_i \geq s\}|}{|\{s_i\}|}.$$

We retain only points with $P(S \geq s) > 10^{-6}$ (finite-sample floor). All fitting is performed in \log_{10} space.

S1.3 Model Fitting and AIC Comparison

Three distributional models are fitted to the CCDF in \log_{10} space using ordinary least-squares (OLS):

Power Law (PL).

$$\log_{10} P(S \geq s) = -(\tau - 1) \log_{10} s + c, \quad 2 \text{ free parameters } (\tau, c).$$

Truncated Power Law (TPL).

$$\log_{10} P(S \geq s) = -(\tau - 1) \log_{10} s - \frac{s}{s_c \ln 10} + c, \quad 3 \text{ free parameters } (\tau, s_c, c).$$

The exponential cut-off scale s_c is estimated by nonlinear least-squares (`scipy.optimize.curve_fit`), with bounds $\tau \in [1, 30]$, $\log_{10} s_c \in [-1, 10]$.

Exponential (EXP).

$$\log_{10} P(S \geq s) = \lambda s + c, \quad 2 \text{ free parameters } (\lambda, c),$$

where $\lambda < 0$ is fitted by OLS on $(s, \log_{10} P)$.

Model selection. Candidate models are compared using the Akaike Information Criterion computed from OLS residuals in \log_{10} space:

$$\text{AIC} = n \ln \left(\frac{\text{RSS}}{n} \right) + 2k,$$

where n is the number of data points along the CCDF grid, k the number of free parameters, and RSS the residual sum of squares in $\log_{10} P$ space.

Scope and interpretation of the AIC used here. The expression above is the Gaussian-residual (“quasi-likelihood”) approximation to AIC, obtained by treating OLS residuals in $\log_{10} P$ space as if they were drawn from an i.i.d. Gaussian. It is *not* a likelihood-based AIC on the raw avalanche-size distribution, and it does *not* directly bound the number of effective independent samples (CCDF points are serially correlated by construction). Three practical consequences follow.

1. **Use as a relative ranking only.** AIC values are interpreted here strictly as a model-ordering device within a fixed CCDF (i.e., a single row of Tables S1–S2). The best-supported model (lowest AIC) is reported there; absolute AIC magnitudes carry no direct probabilistic interpretation in this setting.
2. **AIC magnitudes.** In $\log_{10} P$ space, the RSS is intrinsically small (CCDF values are bounded and densely sampled), so $n \ln(\text{RSS}/n)$ can be a large negative number. Reported AIC values in Tables S1–S2 therefore span $\mathcal{O}(10^3)$ – $\mathcal{O}(10^5)$ in absolute value; what is informative is the *difference* between candidates within a row, not the raw numeric scale. A difference of $|\Delta\text{AIC}| \gtrsim 10$ within a row is treated as decisive.
3. **Distributional coverage.** The candidate set (PL, TPL, EXP) is a deliberately narrow one chosen to distinguish the qualitative questions “is a heavy-tailed scaling present?”, “is the tail cut off?”, and “is the signal compatible with a non-scaling exponential decay?”. Formal maximum-likelihood comparison against stretched-exponential, log-normal, gamma, or discrete power-law alternatives [2] is not performed here, and is flagged as a limitation in the main text.

S1.4 Knee Detection (ARM)

Because the ARM CCDF exhibits a clear two-regime structure (steep small-scale power law followed by a shallower large-scale truncated power law), we locate the crossover point (“knee”) automatically using a *direction-constrained piecewise OLS* algorithm:

1. For each candidate split point s_k (in log-size space), fit two separate OLS lines to the left ($s \leq s_k$) and right ($s > s_k$) segments, restricting the search to points with $\log_{10} P \geq -4$ (body of the distribution, excluding the deep tail).
2. Accept only splits where the right-segment slope γ_R is *shallower* (i.e., less negative) than the left-segment slope γ_L ($\Delta\gamma = \gamma_R - \gamma_L > 0$), which is the physically meaningful direction (large-scale softening).
3. Among all direction-valid splits, select the one maximising $\Delta\gamma$ (largest contrast between the two regimes).
4. If no direction-valid split exists, fall back to the minimum-RSS split within the body.

The resulting knee \hat{s}_k partitions each ARM CCDF into a *small-scale* region ($s \leq \hat{s}_k$) and a *large-scale* region ($s > \hat{s}_k$), each fitted independently.

Sensitivity analysis. Three design choices in the procedure above—the body-range cutoff $\log_{10} P \geq -4$, the direction constraint ($\Delta\gamma = \gamma_R - \gamma_L > 0$), and the min-RSS fallback rule triggered when no direction-valid split exists—each influence the detected knee in principle. To assess robustness we re-ran the procedure on the $N = 20, k = 3$ ARM dataset under a 3×2 grid of configurations: cutoff $\log_{10} P \in \{-3, -4, -5\}$ and fallback $\in \{\text{on}, \text{off}\}$. For each configuration we recomputed the large-scale exponent τ_{large} at every $p \in \{0.05, 0.15, \dots, 0.95\}$. Results are tabulated in Table S4 (`sensitivity_knee.tex`, generated by `figures/fig_supp_knee_sensitivity.py`).

The key robustness quantity here is not the precise numerical value of τ_{large} —which depends on how the body region is delimited (cutoff) and where the algorithm is allowed to place the knee—but whether the qualitative two-regime, SOC-like structure is recovered at all. Across the 3×2 grid in Table S4, three signals address this directly. First, the direction-constrained branch ($\Delta\gamma > 0$) fires in all 28/28 plateau configurations whenever the cutoff is restricted to $\log_{10} P \in \{-3, -4\}$, and in 38/42 $\approx 90\%$ overall, so the existence of a knee separating a steeper small- s slope from a heavier large- s slope is itself robust. Second, the min-RSS fallback rule returns numerically identical estimates whenever the direction-valid branch succeeds (F=T and F=F columns coincide for cutoff $\in \{-3, -4\}$ throughout the plateau), so this rule does not drive the headline result. Third, every direction-valid plateau configuration yields τ_{large} in the narrow range [1.23, 1.50], with a plateau-wide mean of $\bar{\tau}_{\text{large}} = 1.36$, sitting squarely within the canonical SOC range and confirming the universality of the large-scale exponent across design choices. The point-estimate $\bar{\tau}_{\text{large}} = 1.376 \pm 0.081$ reported in the main text and in Table S2 comes from an independent full-tail TPL fit (Methods); together with the knee-detection sweep above it provides two orthogonal lines of evidence that the plateau is in a critical-like state. Outside the plateau ($p \geq 0.75$) the direction-constrained branch fails increasingly often and the only-permissive cutoff $\log_{10} P \geq -5$ returns spurious knees deep in the noisy tail (cells with $\tau_{\text{large}} \gtrsim 5$ in Table S4), exactly as expected for a regime that is no longer critical-like. The sensitivity grid is reported in Supplementary Table S4.

S2 Supplementary Tables

S2.1 Table S1: SNM B&P Fitting Results

Table S1: Avalanche size CCDF fitting results for the Synchronous Noise Model (SNM). “Best” indicates the AIC-best of the three candidate models PL, TPL, EXP (see Methods, “Model Fitting and AIC Comparison”). “ τ / λ ” reports the leading parameter of the AIC-best model: τ (the power-law exponent) when Best $\in \{\text{PL}, \text{TPL}\}$, λ (the exponential decay rate) when Best = EXP. The AIC-best entry in each row is set in boldface.

ε	Best	τ / λ	AIC _{PL}	AIC _{TPL}	AIC _{EXP}	R_{PL}^2	R_{TPL}^2	R_{EXP}^2
0.05	TPL	2.289	-4610	-8701	-4756	0.900	0.984	0.906
0.15	TPL	1.373	-12759	-37381	-30604	0.781	0.993	0.982
0.25	TPL	1.075	-49647	-149559	-136883	0.670	0.998	0.996
0.35	EXP	0.000 [†]	-14750	-39984	-39986	0.614	0.996	0.996
0.45	EXP	0.000 [†]	-3664	-4415	-4417	0.397	0.616	0.616
0.55	EXP	0.000 [†]	-1930	-2012	-2014	0.103	0.171	0.171
0.65	EXP	0.000 [†]	-1755	-1785	-1787	0.052	0.081	0.081
0.75	EXP	0.000 [†]	-1663	-1684	-1686	0.019	0.041	0.041
0.85	EXP	0.000 [†]	-1574	-1587	-1589	0.002	0.018	0.018
0.95	EXP	0.001 [†]	-2991	—	-2995	0.866	—	0.867

[†] Reported $\lambda = 0.000$ – 0.001 values are exact fit outputs rounded to three decimals; they indicate that the AIC-best exponential model has a decay rate at or below this precision floor, corresponding to an avalanche size distribution close to uniform on the fitted log-size grid (i.e., short, isolated bursts with no detectable characteristic scale at the resolution of the analysis). The EXP model nonetheless wins on AIC because PL and TPL introduce additional curvature that the data do not support. “—” in the AIC_{TPL} or R_{TPL}^2 column indicates that the TPL nonlinear fit (`scipy.optimize.curve_fit` with the bounds specified in Methods) failed to converge or returned a non-finite AIC; that row’s AIC comparison was performed across the remaining (PL, EXP) candidates.

S2.2 Table S2: ARM B&P Fitting Results

Table S2: Avalanche size CCDF fitting results for the Asynchronous Random Model (ARM), split at the automatically detected knee \hat{s}_k (Methods, “Knee Detection (ARM)”). The CCDF for each p is partitioned into a *small-scale* region ($s \leq \hat{s}_k$) and a *large-scale* region ($s > \hat{s}_k$), each fit independently with PL, TPL, and EXP. Best_s and Best_l name the AIC-best model on the respective region. τ_{small} and τ_{large} are the PL/TPL exponents (or, when Best is EXP, the exponential decay rate λ); $s_{c,\text{small}}$ and $s_{c,\text{large}}$ are the exponential cut-off scales returned by the TPL fit. “—” in an s_c column indicates that the AIC-best model for that region was not TPL (i.e., PL or EXP), so a TPL cut-off scale is not defined.

p	$\log_{10} \hat{s}_k$	Best _s	τ_{small}	$s_{c,\text{small}}$	Best _l	τ_{large}	$s_{c,\text{large}}$
0.05	2.389	TPL	2.461	277	TPL	1.434	4990
0.15	2.093	TPL	2.478	96	TPL	1.431	1850
0.25	1.875	TPL	2.537	54	TPL	1.410	775
0.35	1.643	TPL	2.556	27	TPL	1.291	282
0.45	1.431	TPL	2.620	16	TPL	1.294	116
0.55	1.255	TPL	2.744	11	TPL	1.277	47
0.65	1.041	TPL	2.973	7	TPL	1.499	20
0.75	1.000	PL	4.175	—	EXP	0.140	—
0.85	1.000	PL	4.841	—	EXP	0.354	—
0.95	0.812	PL	7.036	—	PL	7.958	—

Best_s: best model for small-scale region. Best_l: best model for large-scale region.

S2.3 Table S3: FFT Spectral Exponent β

Table S3: Spectral exponent β from linear regression of $\log_{10} |\hat{C}(f)|$ on $\log_{10} f$ in the range $f \in [10^{-3.5}, 10^{-1.0}]$. Dashes indicate $R^2 < 0.50$ (SNM) or $R^2 < 0.90$ (ARM), where the power-law fit is deemed unreliable.

p / ε	β_{SNM}	R^2_{SNM}	β_{ARM}	R^2_{ARM}
0.05	0.288	0.960	0.599	0.984
0.15	0.091	0.739	0.565	0.983
0.25	0.036	0.535	0.557	0.985
0.35	—	0.325	0.551	0.984
0.45	—	0.177	0.542	0.987
0.55	—	0.064	0.548	0.988
0.65	—	0.013	0.575	0.992
0.75	—	0.002	0.587	0.994
0.85	—	0.001	0.652	0.996
0.95	—	0.000	0.809	0.993

ARM mean: $\beta = 0.599 \pm 0.077$ (all 10 values).

S2.4 Table S4: Knee-detection sensitivity grid

Table S4: **Sensitivity of the large-scale exponent τ_{large} to knee-detection design choices.** For each p and each (cutoff, fallback) configuration, the direction-constrained piecewise OLS is re-run to detect a knee, after which the right (large-scale) region is fit independently with PL, TPL, and EXP and the AIC-best model is selected (mirroring the main pipeline in `fig4_criticality.py`). Each cell reports the resulting τ_{large} . “F=T enables the min-RSS fallback when no direction-valid split exists; “F=F disables it. Right-most columns report the mean and standard deviation of τ_{large} across all six configurations at each p . “— denotes either (i) no direction-valid split with fallback disabled, or (ii) AIC-best on the right region was EXP, in which case τ_{large} is undefined.

p	$\log_{10} P \geq -3$		$\log_{10} P \geq -4$		$\log_{10} P \geq -5$		mean	s.d.
	F=T	F=F	F=T	F=F	F=T	F=F		
0.05	1.435	1.435	1.434	1.434	1.441	1.441	1.437	0.004
0.15	1.433	1.433	1.431	1.431	1.434	1.434	1.433	0.001
0.25	1.429	1.429	1.410	1.410	1.415	1.415	1.418	0.009
0.35	1.272	1.272	1.291	1.291	1.300	1.300	1.288	0.013
0.45	1.266	1.266	1.294	1.294	1.310	1.310	1.290	0.020
0.55	1.235	1.235	1.277	1.277	—	—	1.256	0.024
0.65	1.426	1.426	1.499	1.499	3.262	—	1.822	0.806
0.75	—	—	—	—	—	—	—	—
0.85	—	—	—	—	9.260	—	9.260	—
0.95	—	—	—	—	—	—	—	—

S3 Supplementary Figures

S3.1 Figure S1: SNM Avalanche CCDF with Model Fits

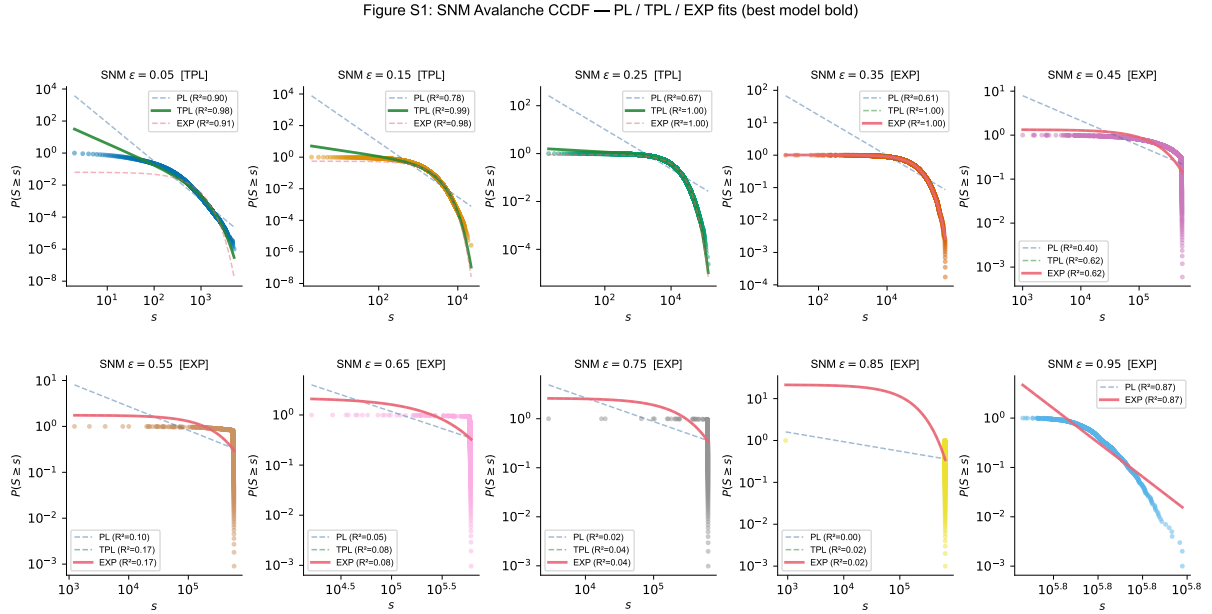


Figure S1: SNM avalanche statistics fail to exhibit scale-free behaviour at any tested noise intensity ε . Across the full sweep $\varepsilon \in \{0.05, 0.15, \dots, 0.95\}$, the AIC-best model is a truncated power-law (TPL) for $\varepsilon \leq 0.25$ and an exponential (EXP) for $\varepsilon \geq 0.35$; no parameter window produces a clean unbounded power-law tail. *Detail:* Each panel shows the empirical CCDF (circles) with three fitted curves overlaid: PL (blue), TPL (green), EXP (red); the AIC-winning model is drawn as a solid thick line, the others as dashed thin lines, and the winning model is reported in brackets in the panel title. TPL dominance at low ε reflects a soft power-law with an exponential cut-off; the EXP dominance at higher ε is consistent with the absence of scale-free avalanche statistics.

S3.2 Figure S2: ARM Avalanche CCDF with Knee Detection and Region Fits

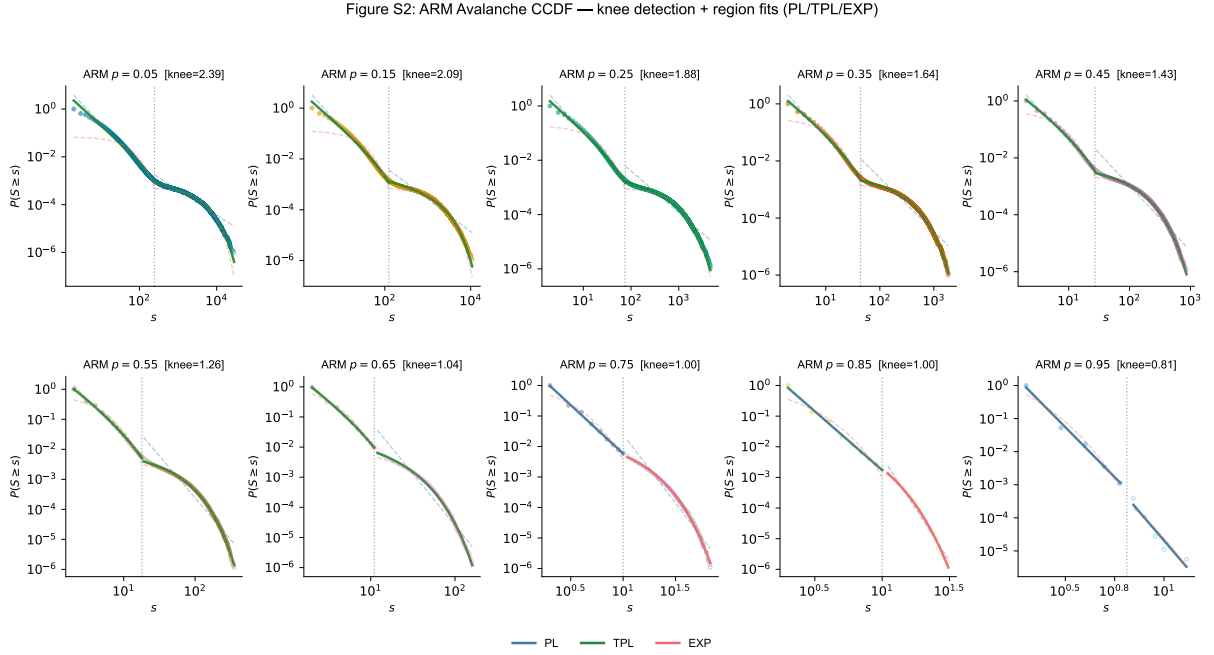


Figure S2: **ARM avalanche statistics exhibit a robust two-scale structure with a universal large-scale exponent across the plateau $p \in [0.05, 0.65]$.** The AIC-best model on the large-scale region ($s > \hat{s}_k$) is a truncated power-law (TPL) with $\bar{\tau}_{\text{large}} = 1.376 \pm 0.081$ (Table S2)—a signature consistent with self-organised criticality—throughout this range; outside the plateau ($p \geq 0.75$), the large-scale region loses scale-free structure and is best fit by a pure power law or an exponential. *Detail:* Filled circles indicate the small-scale region ($s \leq \hat{s}_k$); open circles, the large-scale region ($s > \hat{s}_k$). The vertical dotted line marks the automatically detected knee \hat{s}_k (Methods, “Knee Detection (ARM)”). Fit colour coding as in Figure S1.

S3.3 Figure S3: Topological Robustness across Network Models

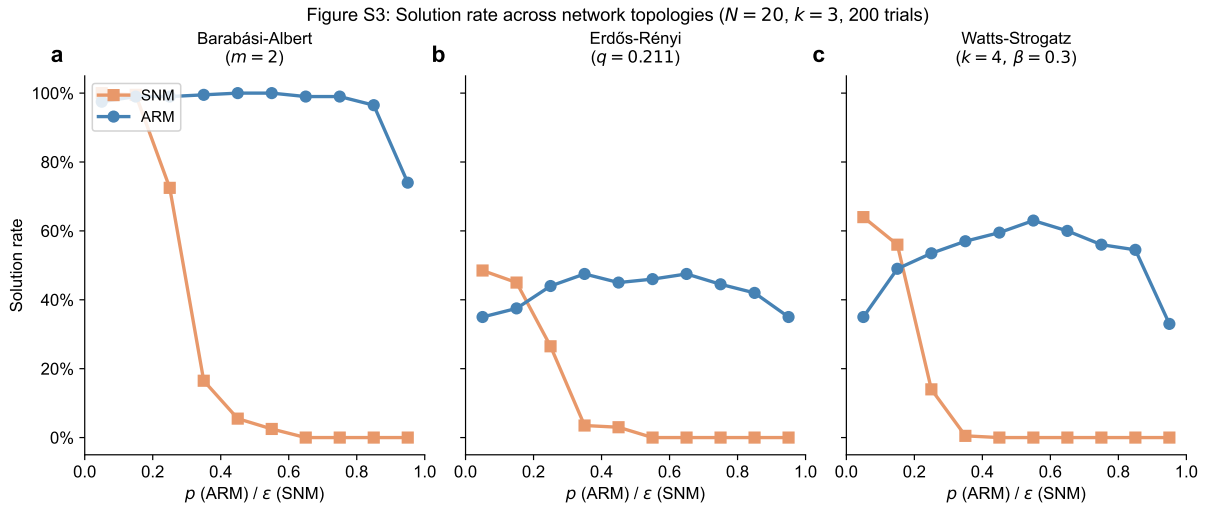


Figure S3: **Topology and update rule contribute orthogonally to coordination performance: topology sets the absolute level (BA > ER, WS), while the qualitative SNM/ARM contrast (narrow ε -peak vs broad p -plateau) is preserved on all three graph classes.** *Detail:* Each panel shows the solution rate within 1,000 steps ($N = 20, k = 3, 200$ trials per point) for the SNM (orange squares) and ARM (blue circles), plotted against the respective parameter ε (SNM) or p (ARM): **(a)** Barabási–Albert ($m = 2$), **(b)** Erdős–Rényi ($q \approx 0.21$, mean degree matched to the BA panel), **(c)** Watts–Strogatz ($k_{\text{ws}} = 4, \beta = 0.3$). The BA network achieves the highest absolute solution rate—its heterogeneous hub structure facilitates cascading conflict resolution. ER and WS networks, with more homogeneous degree distributions, yield lower absolute levels but a qualitatively identical SNM/ARM separation, isolating the update-rule effect from the topology effect.

S3.4 Figure S4: Hybrid Parameter Sweeps

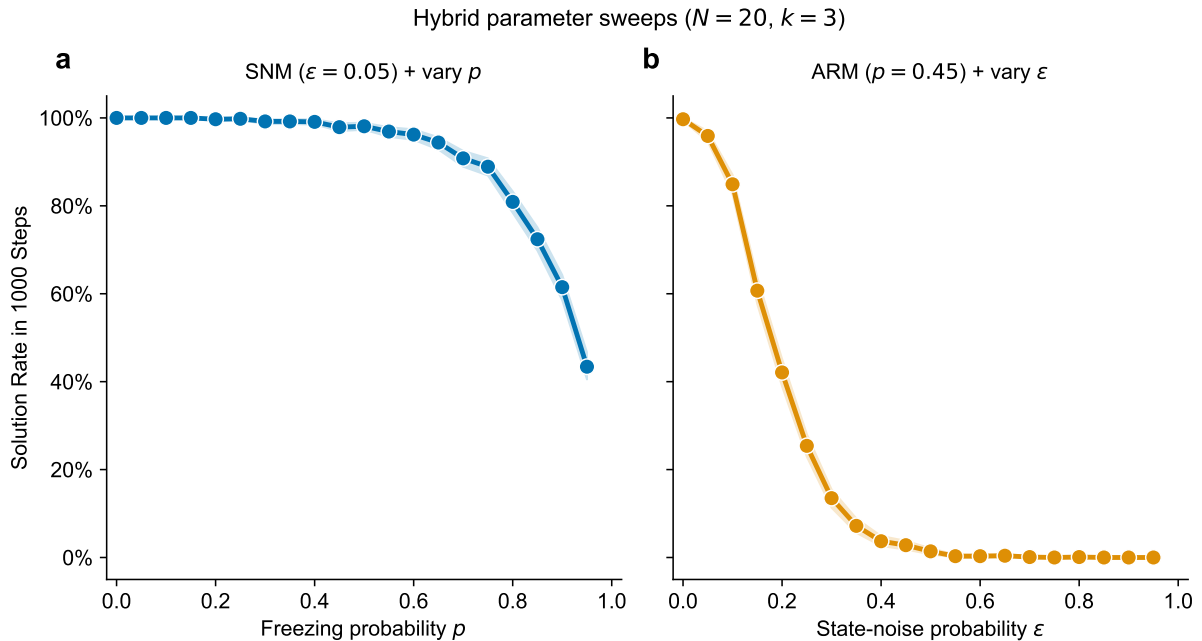


Figure S4: **Solution rate under hybrid parameter combinations** ($N = 20, k = 3$). **(A)** SNM with optimal noise fixed at $\varepsilon = 0.05$, with the asynchronous freezing probability p varied from 0 to 1. The solution rate declines smoothly as p increases, showing that introducing asynchrony into the SNM progressively disrupts the noise-driven coordination mechanism. **(B)** ARM with optimal freezing probability fixed at $p = 0.45$, with state-noise ε varied from 0 to 1. The solution rate similarly declines with increasing ε , indicating that adding synchronous noise degrades the asynchronous cascade mechanism. Together, panels A and B confirm that the two perturbation mechanisms (synchronous noise and asynchronous freezing) are not interchangeable: each model relies on its own distinct dynamical pathway, and cross-contamination degrades performance monotonically.

S3.5 Figure S5: Finite-size scaling of the ARM solution plateau

The main-text analysis identifies two distinct “plateaus” in the ARM parameter space at $N = 20, k = 3$: a *solution-rate plateau* ($0.02 \lesssim p \lesssim 0.92$ at the $\geq 90\%$ threshold; main-text Fig. 3b, see also main-text Fig. 5b,d) over which the network is reliably solved, and a narrower *critical-dynamics window* ($0.05 \leq p \leq 0.65$; main-text Fig. 4e,f, Table S2) over which the avalanche CCDF is best described by a truncated power-law with a universal exponent $\bar{\tau}_{\text{large}} \approx 1.38$. The critical-dynamics window occupies the low- p half of the solution-rate plateau, whereas the high- p half is in a non-critical regime (EXP-preferred in Table S2) yet still solves the colouring. Figure S5 shows how this two-part structure behaves as network size increases.




Figure S5: **Finite-size scaling of the ARM solution plateau ($k = 4$, BA $m = 2$, 200 trials per (N, p) point).** (a) Heatmap of ARM solution rate over freezing probability p (horizontal) and network size N (vertical). Overlaid curves show the plateau lower edge $p_{\min}(N)$ at three thresholds (solution rate $\geq 0.99, 0.95, 0.90$). As N grows, p_{\min} shifts to increasingly higher p , whereas the plateau upper edge ($p \approx 0.95$) is preserved at all tested sizes: the plateau retreats from the low- p side only. (b) The active fraction at the plateau edge, $1 - p_{\min}(N)$, versus N on log–log axes. A least-squares fit on $N \geq 60$ (threshold 0.99) gives $(1 - p_{\min}) = 6.74 N^{-0.49}$, statistically indistinguishable from the power law $N^{-1/2}$ (grey dotted reference). Equivalently, the expected number of simultaneously active agents at the plateau edge scales as $N(1 - p_{\min}) \propto \sqrt{N}$. (c) Scaling collapse: solution rate plotted against the rescaled active count $(1 - p)\sqrt{N}$. Curves for $N \in \{20, \dots, 100\}$ collapse onto a single master curve (vertical dashed line at $(1 - p)\sqrt{N} = 6.74$ marks the estimated breakdown threshold), confirming that $(1 - p)\sqrt{N}$ is the natural scaling variable for the ARM plateau.

Interpretation: two ARM regimes and their N -dependence. The asymmetric shrinkage of the ARM plateau with N admits a natural interpretation in terms of the two fluctuation mechanisms identified in main-text Fig. 4. In the *low- p* portion of the plateau (high update rate), many agents reselect colours concurrently whenever gridlocks arise: conflicts propagate as scale-free cascades, producing the critical-like $\eta_{\text{large}} \approx 1.38$ tail observed for $p \leq 0.65$ at $N = 20$. A simple branching-process argument explains why these cascades become self-interfering at large N . A random reselection at a node of degree d propagates to its neighbours, each of which is active in the same step with probability $(1 - p)$; the effective branching factor is therefore $b \sim (1 - p)d$. In a BA network with $m = 2$ the cascade is dominated by hubs, whose degree scales as the maximum degree $k_{\text{max}} \sim \sqrt{N}$. Cascades remain critical (branching ~ 1) only when $(1 - p)\sqrt{N}$ is bounded, and become over-branching—interfering with their own downstream propagation—once $(1 - p)\sqrt{N}$ exceeds the hub-scale cutoff. This predicts exactly the observed scaling $(1 - p_{\min}) \propto N^{-1/2}$ and aligns the plateau breakdown locus with the $\sqrt{2N}$ critical connectivity already featured in the main-text scalability heatmaps. In the *high- p* portion ($p \gtrsim 0.65$), updates are quasi-sequential: each random reselection propagates locally before the next collision occurs, avalanches are truncated (the EXP-preferred regime in Table S2), and the dynamics are non-critical but size-independent. The ARM’s scalability therefore draws on two qualitatively distinct regimes: critical-like cascade propagation at low p —effective at small N but with an $N^{-1/2}$ -scaled lower edge that retreats in larger networks—and non-critical quasi-sequential coordination at high p , which is slower per cascade but intrinsically size-robust. The solution-rate plateau is the union of these two regimes; the critical-dynamics window is the low- p sub-range where the cascade mechanism is statistically dominant.

S3.6 Figure S6: Solution rate landscape across noise parameter and time budget

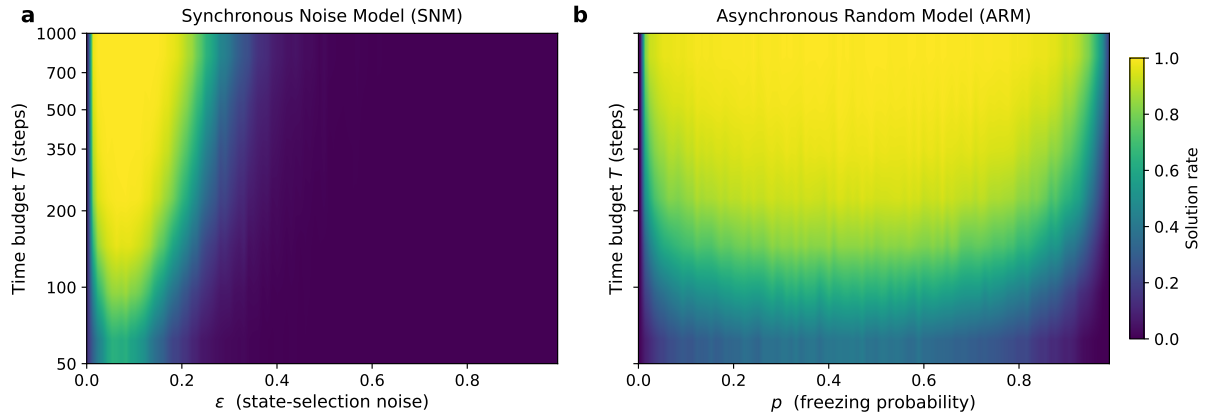


Figure S6: **Solution rate landscape across noise parameter and evaluation horizon T .** (a) SNM solution rate as a function of ε and time budget $T \in \{50, 100, 200, 350, 500, 700, 1000\}$ steps. The narrow ridge near $\varepsilon \approx 0.05$ persists at every T , confirming that the SNM’s parameter fragility is not an artefact of the 1,000-step cutoff used in main-text Fig. 3a. (b) ARM solution rate as a function of p and the same T range. The broad plateau region expands monotonically with T and approaches an asymptotic envelope, demonstrating that the operating range observed in main-text Fig. 3b is robust to (and only widens with) the choice of evaluation horizon. Colour: solution rate (viridis, 0–1). Both panels: $N = 20$, $k = 3$, BA $m = 2$, 1,000 trials per parameter; data computed by re-extracting solve times from the same phase-space sweep used for Fig. 3a, b.

References

- [1] Beggs, J.M. & Plenz, D. Neuronal avalanches in neocortical circuits. *J. Neurosci.* **23**, 11167–11177 (2003).
- [2] Clauset, A., Shalizi, C.R. & Newman, M.E.J. Power-law distributions in empirical data. *SIAM Rev.* **51**, 661–703 (2009).