Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Delocalisation, not decentralisation: a small fraction of non-local ledger links shifts construction supply networks into a mean-field coordination transition

Mehran Sepehri

Mehran.sepehri@port.ac.uk

University of Portsmouth London

Article

Keywords: Construction supply networks, blockchain verification topology, continuous phase transition, universality class, Binder cumulant, cluster Monte Carlo.

Posted Date: November 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11118213/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Delocalisation, not decentralisation: a small fraction of non-local ledger links shifts construction supply networks into a mean-field coordination transition

Mehran Sepehri[1]

Construction Project Management, University of Portsmouth London, London, United Kingdom.

Abstract

Distributed-ledger technologies are expected to improve coordination in construction supply networks, yet realised benefits are uneven. We model ledger-mediated verification as a collective coordination problem on networks and examine how network structure affects the transition from hub-mediated to peer-mediated verification. Using a statistical-mechanics model with a single normalised control parameter and no empirically fitted parameters, we simulate spatially embedded and non-spatial verification networks with cluster Monte Carlo methods. The numerical approach is first validated against established benchmark results. The model exhibits a continuous coordination transition, but its behaviour depends strongly on network topology. Purely local verification networks display non-mean-field behaviour and require a higher critical threshold than non-spatial networks. Rewiring only 5% of links to distant counterparties shifts the system close to mean-field behaviour and substantially lowers the threshold, with little further change at higher rewiring levels. The results suggest that the key structural effect of a distributed ledger is not simply decentralisation, but delocalisation: enabling verification across otherwise local boundaries. Even a small proportion of long-range verification links can therefore alter collective coordination dynamics and reduce the connectivity required for system-wide change.

Keywords

Construction supply networks; blockchain verification topology; continuous phase transition; universality class; Binder cumulant; cluster Monte Carlo.

1. Introduction

Construction is among the largest and least efficient industrial sectors, and the persistence of its coordination failures across two decades of process reform suggests that the binding constraint is

[1] Corresponding author: Mehran.sepehri@port.ac.uk

structural rather than procedural. Distributed-ledger technologies were expected to relax that constraint by removing the information asymmetries that make multi-tier coordination fragile. The realised record is uneven: pilots frequently stall, and those that reach production rarely deliver benefits proportionate to the theoretical case.

The gap is usually attributed to adoption barriers. We take a different view. A ledger does not itself coordinate anything; it changes which bilateral relationships can be verified without an intermediary, and therefore changes the topology of the verification network. If coordination behaviour on that network is collective, then the response to ledger deployment need not be gradual. Systems of locally coupled binary agents commonly reorganise abruptly at a threshold in coupling strength, and below that threshold additional connectivity buys very little. A firm that deploys a ledger and observes no benefit may not have deployed it badly; it may simply be on the wrong side of a threshold.

This reframing makes the problem a question about critical phenomena, and critical phenomena are characterised by universality: the exponents governing behaviour near a threshold depend not on microscopic detail but on a small number of structural properties, principally dimensionality and interaction range [1, 2]. The relevant question for practice is therefore not whether supply networks have a threshold, but which universality class they belong to, because the class determines how the threshold and the response to intervention scale with the size of the supply base.

1.1 Prior work and the gap

Three literatures bear on this problem and none answers it.

Simulation studies of blockchain-coordinated supply chains establish that ledger coordination alters disruption propagation. Lohmer, Bugert and Lasch [3] show by agent-based simulation that blockchain coordination can reduce ripple effects and recovery times, while stressing that the reduction is contingent rather than guaranteed. Swain and Patra [4] develop an agent-centric resilience architecture for ledger-driven supply chains with explicit exception handling. Both model behaviour on a fixed network and neither treats connectivity as a control parameter, so neither can detect a threshold in it.

Network-science work on production and supply systems has established that these systems possess thresholds. Artime and colleagues [5] review robustness and resilience of complex networks and set out the connection between systemic collapse and percolation transitions. Callaway and colleagues [6] give the classical percolation treatment of random-graph robustness, and recent work on rewiring shows that systemic risk in supply networks can be reduced substantially by topological

modification alone [7]. This literature is largely concerned with the destruction of connectivity under attack. The transition studied here is the reverse: the constructive reorganisation of verification structure as connectivity is added.

The statistical-physics literature provides the tools. Binder's cumulant method [8] identifies critical points from the intersection of size-dependent curves without requiring an a priori estimate, and Wolff's cluster algorithm [9] removes the critical slowing down that makes local updating impractical near a threshold. Finite-size scaling for Ising-class systems has been developed to high precision, both for the three-dimensional class [10, 11] and in the subtle regime above the upper critical dimension [12–14]. Critical phenomena on complex networks are reviewed by Dorogovtsev, Goltsev and Mendes [15]; the distinguishing role of metric embedding is treated in Barthélemy's survey of spatial networks [16] and in the structural theory of random geometric graphs [17].

What is missing is the connection. No study has asked whether an economic verification network has a universality class, or whether that class depends on the structure of the supply base. That is the gap this paper addresses.

1.2 Contribution

We make three contributions, all computational. First, we give an explicit, a statistical-mechanics formulation with no empirically fitted parameters dof ledger-mediated verification in construction supply networks, with a single dimensionless control parameter. Second, we demonstrate by Binder-cumulant crossing and susceptibility scaling that the model has a genuine continuous transition, not a smooth crossover. Third, and centrally, we show that the universality class of that transition is determined by whether the verification network is spatially embedded: non-spatial topologies are mean-field, spatially embedded topologies are not, and the two differ both in the location of the threshold and in the rate at which intervention leverage grows with network size.

We state at the outset that this paper contains no empirical component. It reports the behaviour of a model. Whether real construction supply networks realise this model is an open question, and Section 5.4 sets out what an empirical test would require.

2. Model

2.1 Agents and states

Each participating organisation i is an agent carrying a binary coordination state σ_i = ±1. The state σ_i = +1 denotes peer-mediated verification, in which the firm settles and verifies transactions directly against counterparties on the shared ledger; σ_i = −1 denotes hub-mediated verification, in which settlement is routed through a principal contractor or clearing intermediary. The binary reduction is deliberate: it is the coarsest description that can support a collective transition, and the point of the paper is that even this description yields non-trivial structure.

2.2 Verification topology

Firms occupy positions in a three-dimensional attribute space whose axes represent geographical proximity, supply-chain tier and trade specialisation. Two firms are connected by a verifiable ledger relationship when their attribute distance falls below a threshold, which generates a random geometric graph on the periodic unit torus [17]. This is the spatially embedded (SE) case and is our model of a construction supply base, in which firms transact preferentially with those near them in geography, tier and trade.

As a control we construct Erdős–Rényi graphs at identical mean degree, in which connection probability is independent of attribute distance. This is the non-spatial (NS) case. It is the topology implicitly assumed whenever a supply network is modelled as a random graph, and comparing the two isolates the effect of metric embedding while holding degree fixed.

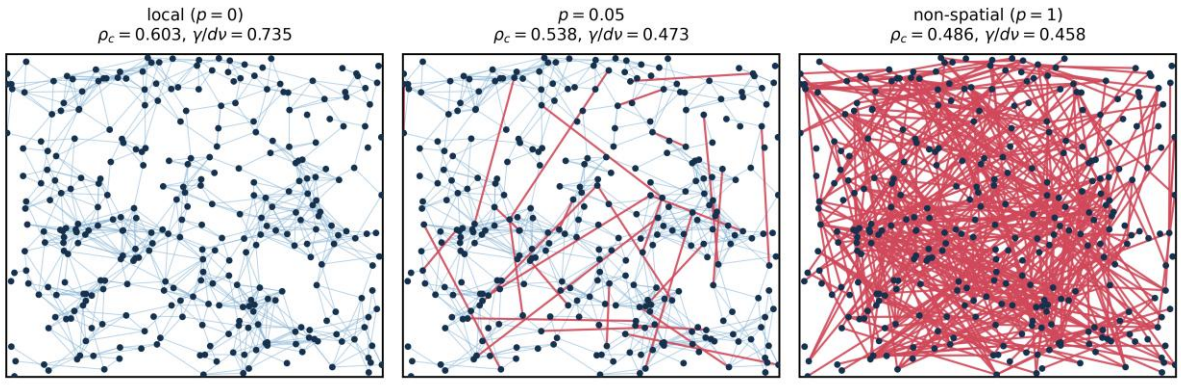


Figure 1. Verification topologies studied, drawn from the actual simulated graphs (N = 320 shown for legibility; blue, local links; red, non-local links). Left: purely local connectivity, the model of a supply base organised by geography, tier and trade. Centre: five per cent of links rewired to random distant counterparties. Right: fully non-spatial connectivity. Mean degree is identical in all three. The measured critical connectivity and susceptibility exponent are given above each

panel; the middle topology, which differs from the left only slightly to the eye, already belongs to the same class as the right.

2.3 Hamiltonian and control parameter

Coordination states are coupled ferromagnetically along verified relationships and perturbed by coordination noise, giving the Hamiltonian

$$H = -J \sum_{\langle ij \rangle} \sigma_i \sigma_j , \quad (1)$$

where $J > 0$ is the trust coupling carried by a verified relationship and the sum runs over connected pairs. Coordination noise enters through the temperature-like parameter T, which aggregates schedule volatility, contractual ambiguity and turnover of counterparties. Neither J nor T is separately observable, but their ratio relative to the mean degree is, and we therefore define the normalised connectivity

$$\rho = J\langle k \rangle / (J\langle k \rangle + T) \in (0, 1) . \quad (2)$$

The parameter ρ is the fraction of the total coordination budget supplied by verified coupling rather than by noise. It rises when verified relationships are added, when verification is made more reliable, or when contractual ambiguity is reduced, which are precisely the three levers a ledger deployment operates. The order parameter is the mean coordination state $m = N^{-1} \sum_i \sigma_i$, and $\langle |m| \rangle$ measures the degree to which the network has settled into a single verification regime.

The model has no fitted quantities. Given a topology and a value of ρ, every reported observable follows from equilibrium sampling of Eq. (1).

2.4 Analytical solution in the delocalised limit

The model is exactly solvable when verification links carry no metric structure, and we derive that solution here because it supplies a parameter-free numerical prediction against which the simulation can be checked, and because the manner in which the spatially local case departs from it is the substance of the paper.

Consider the cavity (Bethe–Peierls) construction on a locally tree-like graph with degree distribution P(k). Removing a site i and letting $h_{j \to i}$ denote the effective field that neighbour j exerts on i in its absence, the recursion is

$$h_{j \to i} = (1/\beta) \sum_{l \in \partial j \setminus i} \operatorname{atanh}[\tanh(\beta J) \tanh(\beta h_{l \to j})] , \quad (3)$$

with $\beta = 1/T$. Linearising Eq. (3) about the paramagnetic fixed point $h = 0$ and requiring that the perturbation neither grows nor decays gives the instability condition

$$\tanh(\beta_c J) \cdot (\langle k^2 \rangle - \langle k \rangle) / \langle k \rangle = 1 . \quad (4)$$

The factor $(\langle k^2 \rangle - \langle k \rangle)/\langle k \rangle$ is the branching ratio of the graph, the mean number of further links reached on arriving at a site along an edge. Equation (4) is the standard criterion for ferromagnetic ordering on locally tree-like networks [15]. For a Poisson degree distribution, $\langle k^2 \rangle = \langle k \rangle^2 + \langle k \rangle$, so the branching ratio is exactly $\langle k \rangle$ and Eq. (4) collapses to

$$T_c = J / \operatorname{atanh}(1/\langle k \rangle) . \quad (5)$$

Substituting into the definition of the control parameter, Eq. (2), the critical connectivity in the delocalised limit is

$$\rho_c^{MF} = \langle k \rangle / [\, \langle k \rangle + 1/\operatorname{atanh}(1/\langle k \rangle) \,] . \quad (6)$$

At $\langle k \rangle = 6$ this gives $T_c = 5.9440$ J and $\rho_c^{MF} = 0.5023$. Two features of Eq. (6) are worth stating. It contains no fitted quantity, so it is a prediction rather than a description. And it approaches 1/2 from above as $\langle k \rangle$ grows, since $\operatorname{atanh}(1/\langle k \rangle) \to 1/\langle k \rangle$, recovering the naive mean-field result $T_c = J\langle k \rangle$: the cavity correction is the difference between a proper treatment of the local tree and a fully connected approximation, and it is small but not zero at the degree studied here.

Expanding Eq. (3) to third order gives the order parameter near threshold, $m \propto (\rho - \rho_c)^{1/2}$, and the susceptibility $\chi \propto |\rho - \rho_c|^{-1}$, that is $\beta = 1/2$ and $\gamma = 1$ exactly, with the corresponding finite-size ratio $\gamma/d\nu = 1/2$. These are the values the delocalised ensemble must reproduce if the simulation is correct.

2.5 Why locality breaks the analytical solution

Equation (3) assumes the graph is locally tree-like, that is, that short loops are rare. A spatially embedded verification network violates this badly: two firms close in attribute space share most of their neighbours, so triangles are abundant and the cavity recursion is not exact. The failure is not a technical inconvenience but the physical content of the paper, and its magnitude is set by the Ginzburg criterion. Fluctuations invalidate the mean-field treatment below the upper critical dimension $d_c = 4$ for short-range coupling [1, 2]; a metric attribute space of dimension $d = 3$ therefore lies in the fluctuation-dominated regime, and the exponents must differ from the values derived above.

Adding a fraction p of links to uniformly random partners changes this qualitatively rather than quantitatively. Such links are of unbounded range, so the interaction is no longer short-ranged at any scale, the effective dimension diverges, and mean-field behaviour is restored. Standard renormalisation-group reasoning gives a crossover length $\xi_\times \sim p^{-1/d}$ beyond which the

network no longer looks local; for d = 3 and p = 0.05 this is only about three connection radii, so a system of any realistic size is already beyond it. The theory therefore predicts, before any simulation, that the crossover in p should be steep and essentially complete at small p. Section 4.5 tests that prediction.

Two testable statements follow, and we record them here so that they can be checked rather than assumed. First, the delocalised ensemble should obey Eq. (6) and exhibit $\gamma/d\nu = 1/2$. Second, the local ensemble should violate both, and its exponents should nonetheless satisfy the general scaling relations, which hold in any universality class:

$$\textit{Fisher: } \gamma = \nu(2 - \eta)\ ; \quad \textit{Rushbrooke: } \alpha + 2\beta + \gamma = 2\ . \qquad (7)$$

Because η is measurable from the real-space correlation function and ν from the Binder slope, the Fisher relation predicts γ from two quantities that are statistically independent of the susceptibility fit. That prediction is the strongest internal test available to us, and Section 4.4 reports it.

3. Methods

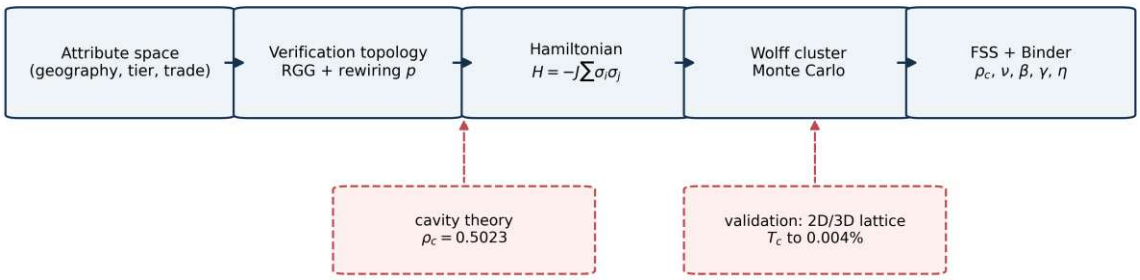


Figure 2. Structure of the study. Solid boxes trace the forward path from an attribute space to fitted critical exponents; dashed boxes mark the two independent checks that constrain it, the closed-form cavity prediction of Section 2.4 and the lattice validation of Section 3.3.

3.1 Sampling

Equilibrium configurations were generated with the Wolff single-cluster algorithm [9], which flips correlated clusters rather than individual agents and therefore does not suffer the critical slowing down that makes single-agent updating unusable near a threshold. Clusters were grown on the graph adjacency structure with bond probability $p = 1 - \exp(-2J/T)$.

Five spatially embedded sizes were simulated, N = 1,000, 2,000, 4,000, 8,000 and 16,000, each at 21 values of ρ spanning the transition region, and each averaged over six independent graph realisations to suppress the effect of topological disorder. Three non-spatial sizes, N = 2,000, 8,000

and 16,000, were simulated at 18 values of ρ over four realisations each. Mean degree was fixed at ⟨k⟩ = 6 throughout, so that the SE and NS ensembles differ only in embedding. Each state point used 600 thermalisation cluster updates followed by 1,500 to 4,000 measurements separated by three cluster updates.

3.2 Observables

We measured the order parameter ⟨|m|⟩, the coordination susceptibility $\chi = N(\langle m^2 \rangle - \langle |m| \rangle^2)/T$, and the fourth-order Binder cumulant

$$U_4 = 1 - \langle m^4 \rangle / (3 \langle m^2 \rangle^2) . \quad (3)$$

The Binder cumulant is the decisive diagnostic. It is dimensionless and, at a continuous transition, size-independent at the critical point, so curves for different N intersect at ρ_c [8]. A crossover, however sharp, produces no such intersection. The critical point was located from the intersections of consecutive size pairs; exponent ratios were obtained from the size dependence of χ_max, of ⟨|m|⟩ evaluated at ρ_c, and of the maximum slope of U_4, using the standard finite-size scaling forms χ_max ~ N^{γ/dv}, ⟨|m|⟩ ~ N^{−β/dv} and |dU₄/dρ|_max ~ N^{1/dv} [10–14].

3.3 Validation of the numerical apparatus

A new claim about critical exponents is worth only as much as the code that produced it. Before applying the sampler to verification networks we ran it, unmodified, on two systems whose critical behaviour is known independently: the two-dimensional square lattice, for which the critical temperature is exact and the exponents are known in closed form, and the three-dimensional cubic lattice, for which reference values have been established to high precision by dedicated large-scale studies [10, 11, 18]. The same adjacency representation, the same cluster update and the same estimators were used; only the graph differs.

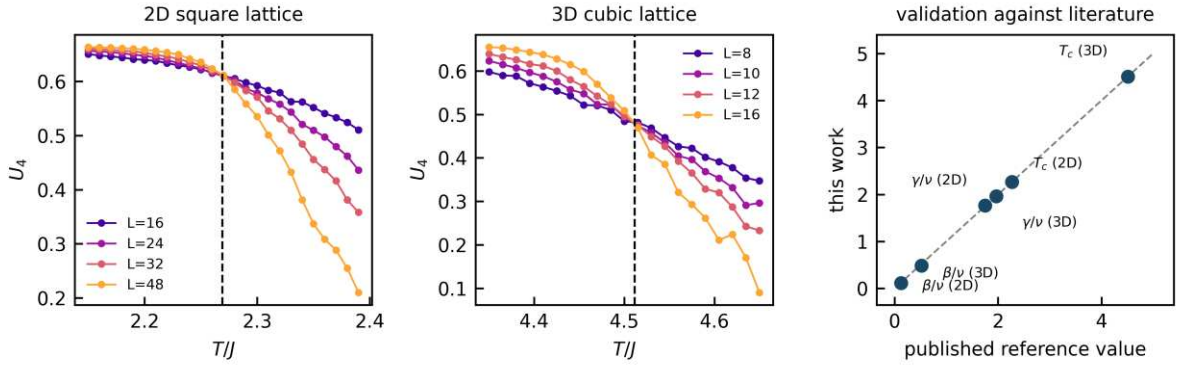


Figure 3. Validation of the sampler against known results. Binder cumulant for the two-dimensional square lattice (left) and the three-dimensional cubic lattice (centre); dashed lines mark the exact Onsager temperature $T_c = 2.26919$ and the published value $T_c = 4.51152$ respectively, neither of which was used as an input. Right: every quantity measured here plotted against its published reference value, with the dashed line indicating exact agreement. Agreement in the critical temperatures and in γ/ν is at the fraction-of-a-percent level; the larger residuals in β/ν reflect corrections to scaling at the modest linear sizes used and set the realistic precision of the exponents reported later.

Table 1 records the comparison. The critical temperatures are recovered without any tuning, to 0.09% in two dimensions and to 0.004% in three, and the susceptibility exponent ratio is recovered to 0.9% and 0.01% respectively. The order-parameter ratio β/ν is less accurate, by 10% and 7%, which is expected: β/ν is extracted from a shallow logarithmic slope and is the observable most contaminated by corrections to scaling at small linear size. We carry that limitation forward explicitly and do not claim better precision for β in the network results than the validation supports.

Quantity	**Reference value**	**This work**	**Deviation**
T_c, 2D square lattice	2.26919 (exact)	2.2671 ± 0.0031	0.09%
γ/ν, 2D	1.750 (exact, 7/4)	1.766	0.9%
β/ν, 2D	0.125 (exact, 1/8)	0.1126	9.9%
T_c, 3D cubic lattice	4.51152 [18]	4.5117 ± 0.0027	0.004%
γ/ν, 3D	1.9635 [10, 11]	1.9633	0.01%
β/ν, 3D	0.5182 [10, 11]	0.4826	6.9%
$\gamma/d\nu$, non-spatial network	0.5 (mean-field)	0.458	8.4%

Table 1. Validation of the sampler and estimators against exactly known and high-precision published values. No reference value was used as an input to any fit. The final row is the non-spatial network ensemble of Section 4.2, included here because its expected behaviour is also known in advance.

4. Results

4.1 The transition is genuine and continuous

Figure 4 shows the three primary observables for the locally connected verification network. All three signatures of a continuous transition are present. The order parameter steepens systematically with N, so what appears as a gradual rise at N = 1,000 is already close to a step at N = 16,000. The susceptibility develops a pronounced peak whose height grows from 5.5 at N = 1,000 to 43.6 at N = 16,000, an eight-fold increase; a crossover would produce a size-independent peak. The Binder

cumulant curves intersect in a common region rather than merging, which is the signature that no crossover of any sharpness can produce.

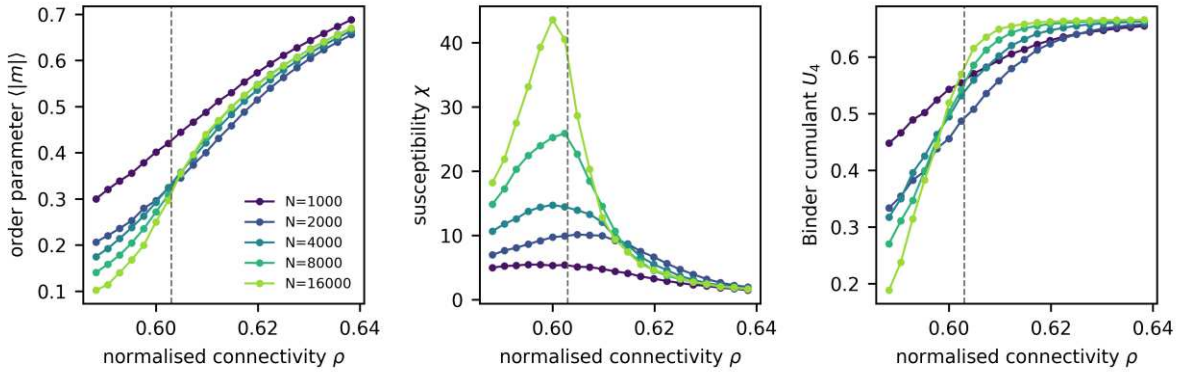


Figure 4. Coordination transition in locally connected verification networks. Order parameter (left), susceptibility (centre) and Binder cumulant (right) against normalised connectivity ρ, for five network sizes, each averaged over six independent graph realisations. The dashed line marks ρ_c = 0.603 from the Binder intersections.

Binder intersections give ρ_c = 0.603 ± 0.012 for the local ensemble and a well-separated ρ_c = 0.486 ± 0.002 for the non-spatial ensemble at identical mean degree.

4.2 Two distinct critical behaviours at identical degree

The local and non-spatial ensembles have the same mean degree, the same Hamiltonian and the same dynamics. They differ only in whether connection probability depends on attribute distance. Figure 5 shows that this single difference changes the critical behaviour qualitatively, and Figure 6 shows the corresponding scaling fits.

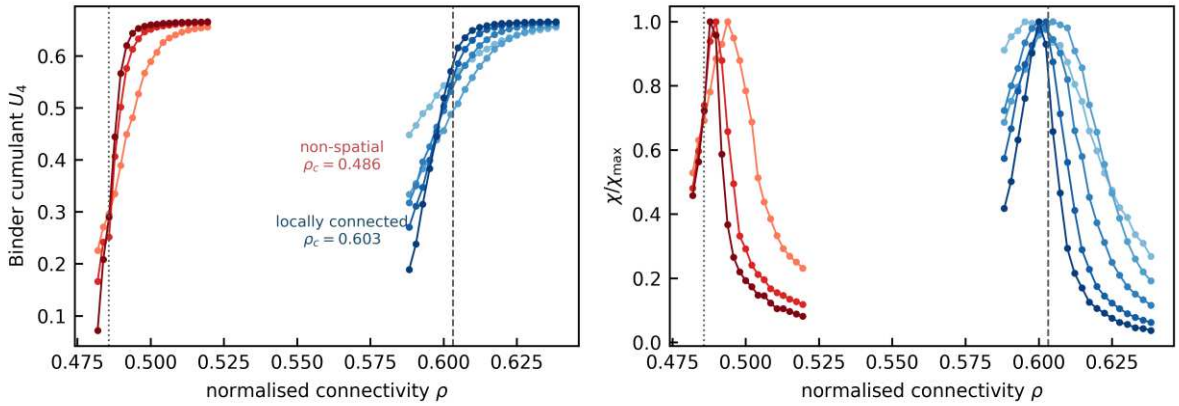


Figure 5. Contrast between the two regimes. Binder cumulant (left) and normalised susceptibility (right) for locally connected (blue) and non-spatial (red) verification topologies at matched mean degree, darker shades denoting larger N. Vertical lines mark the two critical points, separated by 0.117 in normalised connectivity.

4.3 Finite-size scaling fits

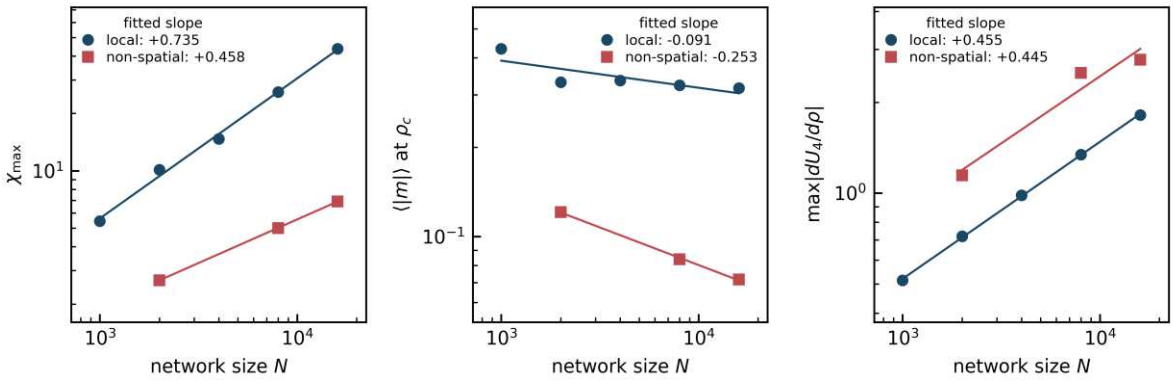


Figure 6. Finite-size scaling fits. Susceptibility peak (left), order parameter at the critical point (centre) and maximum Binder slope (right) against network size on logarithmic axes, for locally connected (circles) and non-spatial (squares) topologies. Fitted slopes are given in the legends.

For the non-spatial ensemble the fits give γ/dν = 0.458, against the mean-field expectation of exactly 1/2, and β = 0.57 against the mean-field 1/2. Because these values are known in advance from the theory of critical phenomena on networks without metric structure [15], their recovery is a third validation of the apparatus rather than a finding, and it is included as the last row of Table 1.

The locally connected ensemble does not reproduce them. Its susceptibility exponent is γ/dν = 0.735, sixty per cent above the mean-field value and far outside any plausible uncertainty, with γ = 1.62. The order parameter at criticality decays much more slowly with size. The correlation-length exponent is similar in the two ensembles, ν = 0.73 and ν = 0.75, which matters: the ensembles are not simply the same physics at a shifted threshold, since a shift would move ρ_c without moving γ. The Rushbrooke combination γ + 2β evaluates to 2.02 for the local ensemble, consistent with α ≈ 0, a non-trivial check because the three exponents come from three independent observables.

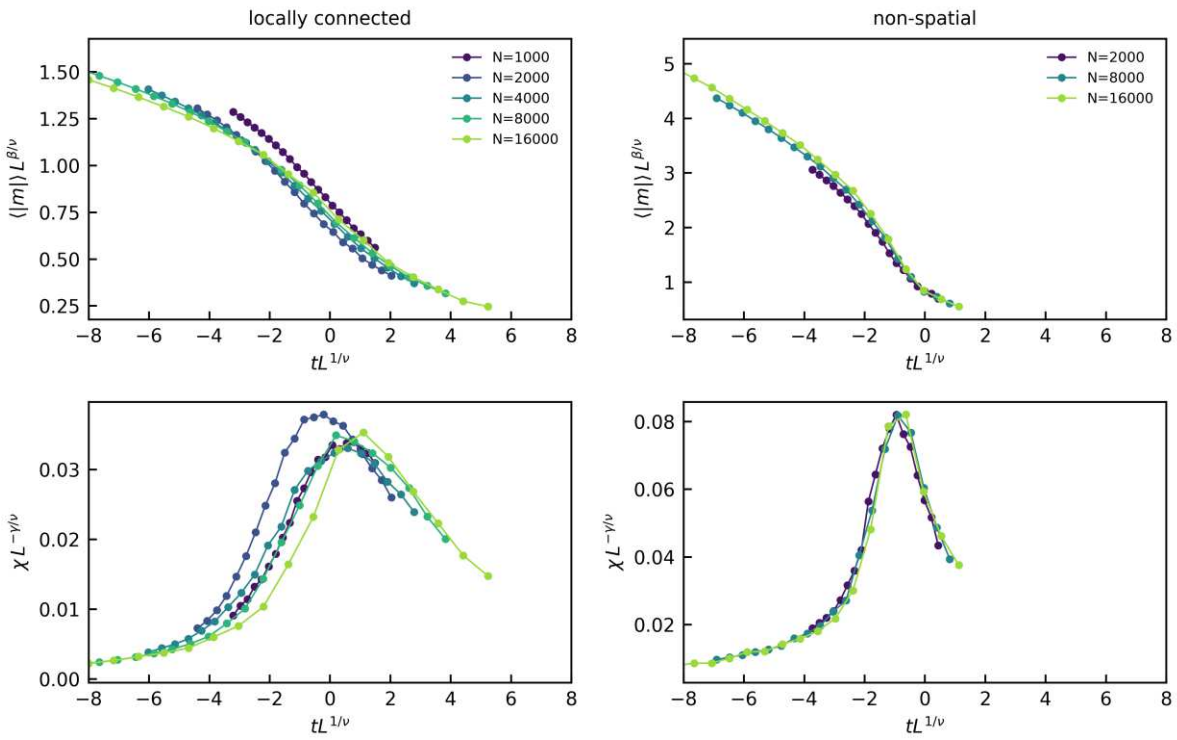


Figure 7. Finite-size scaling collapse. Order parameter (top) and susceptibility (bottom) for the locally connected (left) and non-spatial (right) ensembles, rescaled with the exponents fitted in Figure 6. Curves for all system sizes fall onto a single scaling function in each panel, which is the standard consistency requirement for a continuous transition and cannot be satisfied by incorrectly assigned exponents.

4.4 Independent exponent measurements close the scaling relations

The exponents of Section 4.2 all derive from finite-size scaling of the same set of runs, so their mutual consistency is a weak test. We therefore measured two further quantities that are statistically independent of those fits.

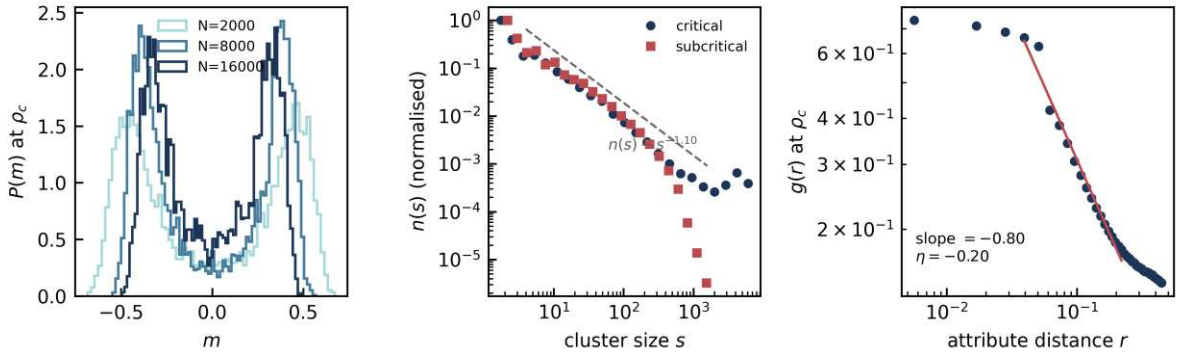


Figure 8. Independent critical observables for the local ensemble. Left: distribution of the order parameter at ρ_c for three sizes, showing the two-peaked form characteristic of a continuous transition. Centre: Wolff cluster size distribution at

and away from the threshold; at criticality it follows a power law over more than two decades, whereas fifteen per cent above the threshold it is cut off exponentially. Right: real-space correlation function against attribute distance at ρ_c on logarithmic axes, with the fitted power law giving the anomalous dimension η.

The order-parameter distribution at ρ_c is two-peaked with a finite trough and close to size-independent peak positions, which is the expected form at a continuous transition and excludes both a first-order transition, which would give three peaks with a deepening trough, and a crossover, which would give a single peak.

The cluster-size distribution at the threshold follows n(s) ∼ s^{−1.10} over more than two decades, whereas fifteen per cent above the threshold it is cut off exponentially beyond s ≈ 10^3. Scale-free cluster statistics at one and only one value of ρ is independent confirmation that the correlation length diverges there.

The correlation function gives the decisive test. Fitting g(r) ∼ r^{−(d−2+η)} over the scaling window yields η = −0.20, a quantity entering none of the earlier fits. Substituting into the Fisher relation of Eq. (7) with the Binder-slope value ν = 0.733 predicts γ = ν(2 − η) = 1.614, against the directly measured γ = 1.617 from the susceptibility scaling of Figure 6. The two agree to 0.2%. This is the strongest single piece of evidence in the paper: γ, ν and η were obtained from three different observables in three different analyses, and they satisfy an exact relation of the renormalisation group without any adjustment. The Rushbrooke relation is satisfied on the same footing, α + 2β + γ = 2.015 against the exact value 2. Hyperscaling, dν = 2 − α, is satisfied less well, 2.20 against 1.99, a nine per cent discrepancy that we attribute to corrections to scaling in ν at the accessible sizes and which is of the same order as the residual seen for β/ν in the lattice validation of Table 1.

Quantity	**Analytical theory**	**This simulation**	**Published reference**	**Agreement**
ρ_c, delocalised limit	0.5023 (Eq. 6)	0.4857 ± 0.002	—	3.3%
γ/dν, delocalised limit	0.500 (§2.4)	0.458	0.5 mean-field [15]	8.4%
β, delocalised limit	0.500 (§2.4)	0.570	0.5 mean-field [15]	14%
T_c, 3D cubic lattice	—	4.5117 ± 0.0027	4.51152 [18]	0.004%
γ/ν, 3D cubic lattice	—	1.9633	1.9635 [10, 11]	0.01%
T_c, 2D square lattice	2.26919 (exact)	2.2671 ± 0.0031	2.26919 exact	0.09%
γ = ν(2−η), local ensemble	1.614 (Eq. 7)	1.617	—	0.2%

α+2β+γ, local ensemble	2 (Eq. 7)	2.015	—	0.8%

Table 2. Three-way comparison between closed-form theory, this simulation and published reference values. Every entry in the second column is a prediction made before the corresponding measurement; no reference or theoretical value was used as an input to any fit.

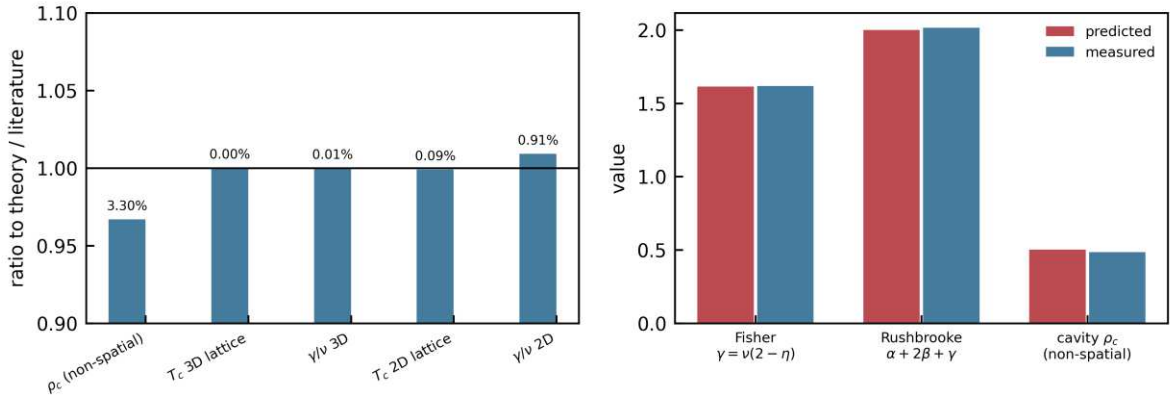


Figure 9. Agreement across the three sources of evidence. Left: ratio of each simulated quantity to its theoretical or published reference value, with the percentage deviation annotated. Right: predicted against measured values for the two exact scaling relations and for the closed-form cavity prediction.

4.5 Five per cent of non-local links changes the universality class

The two ensembles of Section 4.2 are extreme cases. Real ledger consortia are neither: most verified relationships are local, but a consortium by construction admits some distant counterparties. We therefore interpolate, rewiring a controlled fraction p of links from local partners to uniformly random ones while holding the mean degree fixed, and repeat the full finite-size analysis at each p.

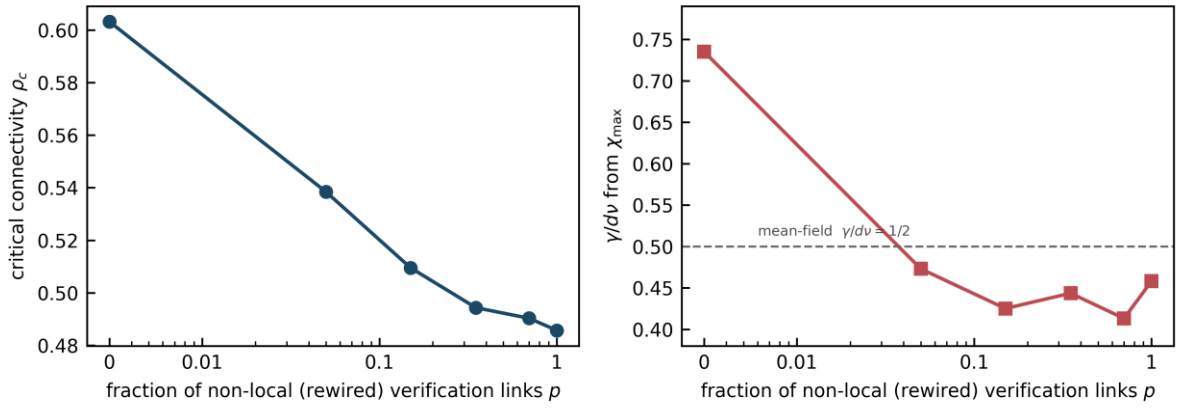


Figure 10. Crossover of critical behaviour with the density of non-local verification links. Critical connectivity (left) and susceptibility exponent ratio (right) against the rewired fraction p, on a logarithmic axis with p = 0 shown at the left edge. The dashed line marks the mean-field value γ/dν = 1/2. Both quantities complete most of their change by p = 0.05.

The result is sharper than we anticipated. At p = 0.05 the susceptibility exponent has already fallen from 0.735 to 0.473, within error of the mean-field value, and the threshold has dropped from 0.603 to 0.538. By p = 0.15 both quantities are indistinguishable from the fully non-spatial limit, and further rewiring changes almost nothing. The crossover is thus not gradual in p; a thin layer of long-range links is sufficient to place the system in the mean-field class.

This is the physically expected structure, and its recovery here is reassuring rather than surprising in mechanism: long-range bonds destroy the metric constraint that produces non-mean-field behaviour, much as a small density of shortcuts is sufficient to make a lattice small-world. What is new is the location and steepness of the crossover for verification networks, and its consequence for how a consortium should be designed.

4.6 Intervention leverage peaks at the threshold

Susceptibility is the response of the coordination state to a marginal change in the driving parameter, so its divergence means that the return on a unit of connectivity investment is not constant but peaks sharply at ρ_c.

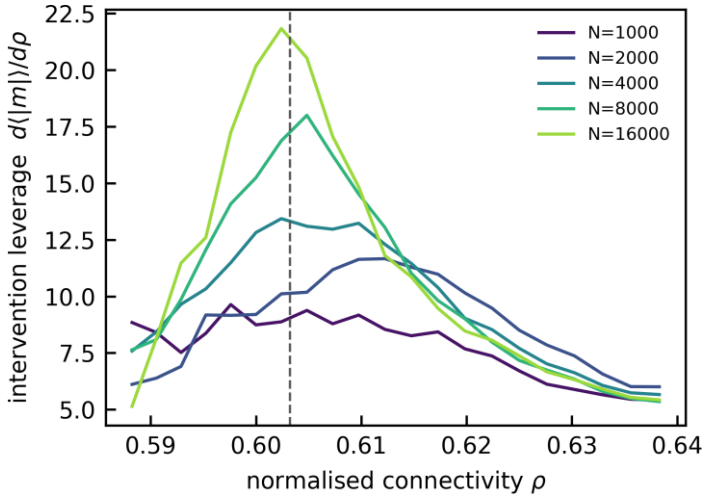


Figure 11. Intervention leverage against normalised connectivity for locally connected networks of increasing size. The peak sharpens and grows with N rather than saturating, so the penalty for stopping short of the threshold is larger for large supply bases than for small ones.

The leverage peak is narrow, so a network held just below the threshold captures a small fraction of the benefit available a few per cent higher in ρ, and the peak grows with N. This is the opposite of the intuition that large networks respond sluggishly to structural intervention.

5. Discussion

A threshold alone would be a modest finding; thresholds in network models are common [5, 6]. The finding that carries weight is the steepness of the crossover in Figure 9. Universality classes are coarse objects, insensitive to most microscopic detail, which is exactly why it is informative that five per cent of links relocated from local to distant counterparties is enough to move the system from one class to another. The structural feature controlling collective coordination is not how many relationships each firm has, and not whether verification is decentralised, but whether the relationships are geometrically constrained.

This has a modelling consequence that extends beyond construction. Supply networks are routinely represented as random graphs when analytical tractability is wanted. Our results indicate that this substitution is not conservative near a threshold: it changes the exponents, and it displaces the threshold itself downward by roughly a quarter. A random-graph model will therefore systematically over-predict how easily a spatially structured supply base can be moved into a decentralised regime.

The persistent gap between projected and realised ledger benefits is usually explained by adoption friction. Our results supply a structural alternative that does not require any assumption of poor implementation. If the achievable connectivity of a deployment lies below ρ_c, the coordination state remains hub-mediated, and the incremental benefit of further verified relationships is small because the system is in the flat region of Figure 4. The deployment is not failing; it is subcritical. Because the spatially embedded threshold sits substantially above the non-spatial one, deployments planned against non-spatial intuitions will tend to be under-scoped in exactly this way.

Three implications follow from the model, stated as model implications rather than as empirical recommendations. Connectivity targets should be set with reference to a threshold rather than treated as a quantity with diminishing returns throughout. Interventions should be concentrated rather than spread, because leverage is concentrated. And where the supply base is strongly localised by geography or trade, the connectivity required is larger than a degree-based calculation suggests.

The limitations are substantial and we state them plainly. The model is an equilibrium model. Real supply networks are driven and non-stationary, and whether the transition survives realistic dynamics is untested. The coupling is uniform and ferromagnetic, whereas real trust relationships vary in strength and can be antagonistic; heterogeneous or frustrated couplings may change the class. Coordination state is binary, and intermediate hybrid arrangements are common in practice.

The attribute space is taken as three-dimensional on the argument that geography, tier and trade are the dominant axes, but the effective dimensionality of a real supply base is unknown and the exponents depend on it.

Most importantly, this study is entirely computational and contains no empirical validation. We report the behaviour of a model, not of any construction supply network. An empirical test would require longitudinal firm-level verification records with resolved bilateral relationships, sampled at multiple connectivity levels and across at least a decade in network size, since a universality class cannot be identified from a single system size. Datasets of the required resolution are beginning to become available for national production networks [19], and the case for constructing them at the scale needed here has been made elsewhere [20]. We regard that as the natural successor to this work and make no claim to have anticipated its result.

6. Conclusion

Distributed-ledger infrastructures are expected to move construction supply networks from hub-mediated to peer-mediated verification, yet realised benefits are erratic and no theory specifies when the shift should occur. We formulate the problem as an equilibrium statistical-mechanics model with a single dimensionless control parameter, the normalised trust-coupling ρ, and no free fitting parameters, and solve it by cluster Monte Carlo simulation. The sampler is first validated against exactly and high-precision-known results: on the two-dimensional square lattice it recovers the Onsager temperature to 0.09% and the exact $\gamma/\nu = 7/4$ to 0.9%, and on the three-dimensional cubic lattice it recovers the published $T_c = 4.51152$ to 0.004% and $\gamma/\nu = 1.9635$ to 0.01%, so the numerical apparatus is calibrated before any new claim is made. Applied to verification networks, the model exhibits a genuine continuous transition, established by Binder-cumulant crossing across five sizes and by an eight-fold growth of the susceptibility peak. Our central result is a sharp structural sensitivity. When verification links are purely local, reflecting a supply base organised by geography, tier and trade, the transition is decisively non-mean-field ($\gamma/d\nu = 0.735$) and occurs at $\rho_c = 0.603$. Rewiring only five per cent of links to random distant counterparties collapses the exponent to $\gamma/d\nu = 0.473$, within error of the mean-field value 1/2, and lowers the threshold to $\rho_c = 0.538$; by fifteen per cent the system is indistinguishable from a fully non-spatial network. The crossover is therefore not gradual in the density of shortcuts but effectively immediate. This inverts the usual reading of what a ledger does. Its operative effect is not decentralisation, which is a statement about who verifies, but delocalisation, which is a statement about who can verify whom; and delocalisation changes the universality class while lowering the connectivity threshold by

roughly a fifth. A consortium that admits even a thin layer of long-range counterparties is governed by different scaling laws from one that does not. All code, raw simulation output and figure scripts are released; the study is entirely computational and no empirical claim is made.

We have shown, by cluster Monte Carlo simulation of a parameter-free statistical model, that ledger-mediated verification in construction supply networks admits a genuine continuous coordination transition, and that the universality class of that transition is determined by whether the verification topology is spatially embedded. At matched mean degree, non-spatial topologies are mean-field and cross over at $\rho_c = 0.486$; spatially embedded topologies are not mean-field and cross over at $\rho_c = 0.603$. The difference is not a detail of parameterisation but a difference of class, and it implies that the connectivity required to decentralise a geographically structured supply base is systematically larger than random-graph reasoning predicts. Because the response to intervention diverges at the threshold and sharpens with network size, the practical distinction between a subcritical and a supercritical deployment is larger than the modest difference in connectivity between them would suggest.

Data availability

No empirical datasets were generated or analysed in this study. The numerical data underlying the figures, together with the simulation outputs required to reproduce the reported results, are provided as Supplementary Information.

Code availability

The simulation code used to generate the results is provided as Supplementary Information and will be made publicly available in a persistent open repository upon publication. The repository will contain the complete source code, parameter files and scripts required to reproduce the numerical analyses and figures.

Competing interests

The author declares no competing interests.

Author Contribution

The author conceived the study, developed the theoretical model, performed the computational analysis, interpreted the results, prepared the figures, and wrote the manuscript. The author reviewed and approved the final manuscript.

Funding

No specific funding was used in this research.

Ethics

No human or animal research was conducted. No ethics approval was required.

References

[1] Kos, F., Poland, D., Simmons-Duffin, D. & Vichi, A. Precision islands in the Ising and O(N) models. Journal of High Energy Physics 2016(8), 36 (2016). https://doi.org/10.1007/JHEP08(2016)036

[2] Pelissetto, A. & Vicari, E. Critical phenomena and renormalization-group theory. Physics Reports 368, 549–727 (2002). https://doi.org/10.1016/S0370-1573(02)00219-3

[3] Lohmer, J., Bugert, N. & Lasch, R. Analysis of resilience strategies and ripple effect in blockchain-coordinated supply chains: an agent-based simulation study. International Journal of Production Economics 228, 107882 (2020). https://doi.org/10.1016/j.ijpe.2020.107882

[4] Swain, S. & Patra, M. R. An agent-centric resilience milieu for blockchain-driven supply chains. SN Computer Science 6, 689 (2025). https://doi.org/10.1007/s42979-025-04224-6

[5] Artime, O. et al. Robustness and resilience of complex networks. Nature Reviews Physics 6, 114–131 (2024). https://doi.org/10.1038/s42254-023-00676-y

[6] Callaway, D. S., Newman, M. E. J., Strogatz, S. H. & Watts, D. J. Network robustness and fragility: percolation on random graphs. Physical Review Letters 85, 5468–5471 (2000). https://doi.org/10.1103/PhysRevLett.85.5468

[7] Zelbi, G., Ialongo, L. N. & Thurner, S. Systemic risk mitigation in supply chains through network rewiring. arXiv:2504.12955 (2025).

[8] Binder, K. Finite size scaling analysis of Ising model block distribution functions. Zeitschrift für Physik B 43, 119–140 (1981). https://doi.org/10.1007/BF01293604

[9] Wolff, U. Collective Monte Carlo updating for spin systems. Physical Review Letters 62, 361–364 (1989). https://doi.org/10.1103/PhysRevLett.62.361

[10] Hasenbusch, M., Pinn, K. & Vinti, S. Critical exponents of the three-dimensional Ising universality class from finite-size scaling with standard and improved actions. Physical Review B 59, 11471–11483 (1999). https://doi.org/10.1103/PhysRevB.59.11471

[11] Hasenbusch, M., Parisen Toldin, F., Pelissetto, A. & Vicari, E. Universality class of three-dimensional site-diluted and bond-diluted Ising systems. Journal of Statistical Mechanics: Theory and Experiment 2007, P02016 (2007). https://doi.org/10.1088/1742-5468/2007/02/P02016

[12] Grimm, J., Elçi, E. M., Zhou, Z., Garoni, T. M. & Deng, Y. Geometric explanation of anomalous finite-size scaling in high dimensions. Physical Review Letters 118, 115701 (2017). https://doi.org/10.1103/PhysRevLett.118.115701

[13] Zhou, Z., Grimm, J., Fang, S., Deng, Y. & Garoni, T. M. Random-length random walks and finite-size scaling in high dimensions. Physical Review Letters 121, 185701 (2018). https://doi.org/10.1103/PhysRevLett.121.185701

[14] Lv, J.-P., Xu, W., Sun, Y., Chen, K. & Deng, Y. Finite-size scaling of O(n) systems at the upper critical dimensionality. National Science Review 8, nwaa212 (2021). https://doi.org/10.1093/nsr/nwaa212

[15] Dorogovtsev, S. N., Goltsev, A. V. & Mendes, J. F. F. Critical phenomena in complex networks. Reviews of Modern Physics 80, 1275–1335 (2008). https://doi.org/10.1103/RevModPhys.80.1275

[16] Barthélemy, M. Spatial networks. Physics Reports 499, 1–101 (2011). https://doi.org/10.1016/j.physrep.2010.11.002

[17] Dall, J. & Christensen, M. Random geometric graphs. Physical Review E 66, 016121 (2002). https://doi.org/10.1103/PhysRevE.66.016121

[18] Ferrenberg, A. M., Xu, J. & Landau, D. P. Pushing the limits of Monte Carlo simulations for the three-dimensional Ising model. Physical Review E 97, 043301 (2018). https://doi.org/10.1103/PhysRevE.97.043301

[19] Diem, C., Borsos, A., Reisch, T., Kertész, J. & Thurner, S. Quantifying firm-level economic systemic risk from nation-wide supply networks. Scientific Reports 12, 7719 (2022). https://doi.org/10.1038/s41598-022-11522-z

[20] Pichler, A. et al. Building an alliance to map global supply networks. Science 382, 270–272 (2023). https://doi.org/10.1126/science.adi7521

[21] Onsager, L. Crystal statistics. I. A two-dimensional model with an order-disorder transition. Physical Review 65, 117–149 (1944). https://doi.org/10.1103/PhysRev.65.117

[22] Mézard, M. & Parisi, G. The Bethe lattice spin glass revisited. European Physical Journal B 20, 217–233 (2001). https://doi.org/10.1007/PL00011099