Adaptive Resolution Modelling

A Framework for Computationally Efficient Simulation of Complex Systems

Equilibria Research

February 27, 2026

Abstract

Agent-based models (ABMs) are powerful tools for understanding emergent social and biological phenomena, yet they face a fundamental computational barrier: systems where "no model less complex than the system itself can accurately predict in detail" how behavior will unfold [12]. Meanwhile, coarse-grained approximations lose the nonlinear cascade dynamics that matter most at phase transitions. We propose *Adaptive Resolution Modelling* (ARM), a framework that dynamically allocates computational resources based on proximity to criticality. Using percolation measures as early warning indicators and minimum-cut algorithms to identify natural aggregation boundaries, ARM operates on coarse representations during stable periods and descends to fine resolution only as the system approaches phase transitions. This respects the actual information structure of complex dynamics: most of the time, detailed agent states are redundant; near criticality, they become consequential. The framework addresses a key gap in ABM methodology—principled, automatic switching between resolution levels—while preserving the "bottom-up" explanatory power that makes agent-based approaches valuable.

1 Introduction

Agent-based modeling has emerged as a central methodology for understanding complex systems across disciplines—from epidemiology and economics to ecology and urban planning [13, 14]. The appeal is explanatory: rather than positing aggregate dynamics directly, ABMs "grow" collective behavior from the bottom up, revealing how macroscopic phenomena emerge from microscopic interactions [14]. As Epstein argues, this shifts the question from "Can you explain it?" to "Can you grow it?"

Yet ABMs face persistent challenges that limit their scientific and policy utility. Three stand out.

First, **computational complexity scales poorly**. The complexity of agent interactions grows combinatorially with population size; for many systems, "the behavior of one million agents is significantly different from that of 100" [15]. Simulation-based optimization becomes impractical, and even basic sensitivity analysis requires prohibitive computational resources [16].

Second, **emergence detection remains informal**. As Macy observed in the foundational PNAS colloquium on ABM, demonstrations of emergence typically rely on "human observers to declare emergence to have occurred based on graphical computer outputs" [17]. Without computable measures of when collective dynamics shift qualitatively, modelers cannot systematically distinguish genuine phase transitions from noise.

Third, **multi-scale coupling is ad hoc**. Real systems operate across scales—individual decisions aggregate into market dynamics, which feed back into individual choices. Attempts to model everything at once face the "complexity curse": reduced transparency from too many interacting mechanisms [18]. Yet principled methods for moving between scales remain underdeveloped.

These challenges share a common structure: they all concern the relationship between microscopic detail and macroscopic behavior. When does detail matter? When can it be safely ignored? When must we track every agent, and when can we work with aggregates?

We propose that *punctuated equilibria*—the pattern of long stability interrupted by brief criticality, first identified in paleontology [1, 2] and now recognized across complex systems—provides the key. During stable periods, agent-level details average out; coarse representations suffice. Near phase transitions, microscale configurations become consequential; full resolution is required. The question becomes: how do we detect approaching criticality?

Percolation theory, developed in statistical physics, provides exactly this [7]. Near percolation thresholds, correlation length diverges, susceptibility spikes, and cluster distributions become scale-free. These are computable signatures that can be monitored during simulation.

This paper develops *Adaptive Resolution Modelling* (ARM), a framework that combines punctuated equilibria as a conceptual foundation, percolation measures as early warning indicators, and graph-based coarse-graining as the mechanism for moving between scales. The result is a principled approach to the computational complexity problem: simulate at high resolution when it matters, low resolution when it doesn't, and let the system's own dynamics determine which regime applies.

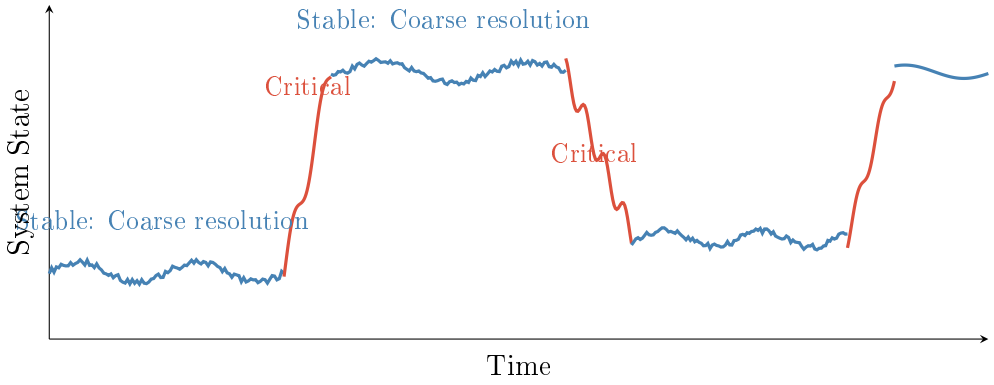

Figure 1: Punctuated equilibria in complex systems. Long periods of stability (blue) are interrupted by brief critical windows (red) where rapid, nonlinear change occurs. Adaptive Resolution Modelling allocates computational resources proportionally to this structure.

2 Punctuated Equilibria as a Universal Pattern

When paleontologists Niles Eldredge and Stephen Jay Gould proposed their theory of punctuated equilibria in 1972 [1], they made a claim that initially seemed confined to evolutionary biology: species remain morphologically stable for vast stretches of time, then change rapidly during brief speciation events. What has become clear in the five decades since is that this pattern is not peculiar to biological evolution. It appears across domains so different that the recurrence demands explanation.

2.1 The Biological Evidence: Stasis is the Rule, Not the Exception

The original evidence came from Eldredge's meticulous study of trilobites in the genus *Phacops*, extinct arthropods abundant in the Middle Devonian rocks of eastern North America [19]. What Eldredge found defied the Darwinian expectation of gradual, continuous change. Across 6 to 8 million years of geological time, these trilobites remained essentially unchanged. The number of lens columns in their compound eyes—a precisely countable morphological feature—stayed constant within populations for millions of years. Then, in geographically restricted populations

at the margins of the species range, rapid change occurred. Speciation took between 5,000 and 50,000 years—a geological eyeblink representing less than 1% of the species' total duration [20].

This is not an isolated case. A meta-analysis of 58 published studies on speciation patterns found that 71% of species exhibited stasis, and 63% showed punctuated patterns of evolutionary change [21]. The fern *Osmunda claytoniana* has remained unchanged at the level of fossilized cell nuclei for 180 million years. Bryozoans of the genus *Metrarabdotos*—which Gould called "the most brilliantly persuasive" example—show the same pattern: long stasis punctuated by rapid cladogenesis.

Why does this happen? The answer lies in the basic mechanics of natural selection operating on populations of different sizes in different environmental conditions.

During stable environmental periods, species occupy ecological niches to which they are already well-adapted. The existing phenotype works: individuals with that body plan can find food, avoid predators, and reproduce successfully. Variants that deviate significantly from this optimum tend to die young or fail to reproduce—they cannot find the specific food sources, cannot escape the particular predators, or cannot attract mates in the ways the population has evolved to do. This is *stabilizing selection*: selection that eliminates variation rather than favoring directional change. The phenotype stays the same not because nothing is happening, but because selection is actively maintaining it.

In large, widespread populations, there is another force preserving stasis: gene flow. Even if a beneficial mutation arises somewhere in the species' range, it gets diluted by interbreeding with the vast majority of individuals who lack it. New variants cannot reach fixation because the population is too large and too interconnected. The genetic signal gets swamped before it can spread.

But environments change. Climate shifts. New competitors arrive. Food sources disappear. When this happens, the old phenotype stops working. Individuals with the ancestral body plan now *cannot* find food, *cannot* survive the new temperatures, *cannot* compete with the newcomers. They die, or they fail to reproduce. This is the brutal simplicity of natural selection: what worked before now kills you.

The crucial insight of Eldredge and Gould, building on Ernst Mayr's earlier work on geographic speciation, is that rapid evolutionary change happens in small, geographically isolated populations at the margins of a species' range. In these peripheral isolates, two conditions combine to enable rapid change. First, gene flow from the main population is cut off, so new variants are not diluted. Second, these marginal populations often face the most extreme environmental stress—they are at the edge of what the species can tolerate, so when conditions shift, the selection pressure is intense. Individuals with the old phenotype die. The few variants that happen to cope better with the new conditions survive and reproduce. Within thousands of years—fast by geological standards but still hundreds of generations—a new species emerges, adapted to the new conditions.

Once this new species is established, it expands its range, and stasis resumes. The new phenotype now works in the new environment. Stabilizing selection kicks in again, eliminating variants. The cycle repeats: millions of years of stasis, then a brief burst of change when environmental shifts make the old form unviable.

2.2 Technology and Economics: Schumpeter's Creative Destruction

The economist Joseph Schumpeter observed the same pattern in technological and economic systems, though he used different language. His concept of "creative destruction" [22] describes capitalism as fundamentally characterized by long periods of relative stability punctuated by waves of radical innovation that sweep away established practices and create new industrial structures.

Economic historian Joel Mokyr made the connection explicit in his 1990 paper "Punctuated Equilibria and Technological Progress" [23], arguing that technological change displays non-

linear dynamics: prolonged periods of incremental improvement succeeded by rapid bursts of radical novelty. The economic literature now recognizes these patterns as "Kondratiev waves" or "techno-economic paradigms" [24]—cycles of approximately 40-60 years in which a general-purpose technology (steam power, electricity, internal combustion, digital computing) transforms the economic structure before giving way to a new paradigm.

Consider the pattern: the railroad industry remained dominant from roughly 1840 to 1890, with incremental improvements in track, locomotives, and scheduling but no fundamental reorganization. Then came a 20-year transition period as electrical systems and internal combustion engines emerged, fundamentally restructuring transportation, manufacturing, and urban geography. Similar patterns appear in the transitions from mainframe to personal computing, from analog to digital communications, and—currently—from human-only to human-AI hybrid economic processes.

The mechanism parallels the biological case. Established technologies and business practices constitute an "organizational phenotype" adapted to a particular techno-economic niche. Within that niche, incremental improvements occur, but radical reorganization is suppressed by sunk costs, institutional inertia, and the difficulty of coordinating change across interdependent systems. When a sufficiently powerful new technology emerges—one that changes the fundamental constraints of the niche—rapid reorganization becomes possible and often inevitable.

2.3 Policy and Institutions: The Baumgartner-Jones Framework

Political scientists Frank Baumgartner and Bryan Jones extended punctuated equilibrium thinking to policy analysis in their influential 1993 work *Agendas and Instability in American Politics* [25]. Their empirical finding was striking: most policy domains remain stable for decades, with only incremental adjustments. But occasionally, issues break through to the "macropolitical" agenda, public attention spikes, and rapid, large-scale policy change occurs.

Baumgartner and Jones documented this pattern across US government budgets from 1800 to 2000. The distribution of annual budget changes is not normal (as incrementalist theory would predict) but leptokurtic—a high central peak (most years, budgets change very little) with fat tails (occasional years see dramatic increases or decreases). The same pattern appears in comparative data from Belgium, Denmark, and other democracies [26].

The mechanism involves "policy monopolies"—stable configurations of interest groups, legislative committees, and regulatory agencies that control how issues are framed and processed. Within a monopoly, only incremental change is possible because alternative framings cannot gain traction. But when external shocks, scandal, or shifting public attention disrupt the monopoly, issues enter competitive, high-attention venues where dramatic change becomes possible. Once a new equilibrium forms—a new policy monopoly with a new dominant framing—stability resumes.

2.4 Why the Same Pattern Appears Across Domains

The mechanisms differ across these domains, but the pattern is the same. In biology, stabilizing selection maintains phenotypes until environmental change makes them lethal, and gene flow in large populations prevents change until geographic isolation removes it. In economics, sunk costs and coordination problems lock in technologies until a sufficiently superior alternative makes the old approach uncompetitive. In policy, institutional inertia and entrenched interests preserve the status quo until attention shifts or crises force reconsideration.

What these mechanisms share is a common structure: forces that resist change dominate during normal periods, but become overwhelmed when conditions shift beyond certain thresholds. The result is a characteristic temporal signature—long stability punctuated by rapid transitions—that emerges from the interaction between stabilizing forces and environmental

change, regardless of whether those forces operate through natural selection, market competition, or political institutions.

Principle 1 (Punctuated Dynamics)**.** *Systems subject to stabilizing forces that resist change, operating in environments that occasionally shift beyond tolerance thresholds, generically exhibit punctuated dynamics: long periods of quasi-stability separated by brief windows of rapid reorganization.*

For agent-based modelers, this pattern has profound methodological implications. The temporal signature of punctuated equilibria—extended stasis interrupted by rapid transitions—corresponds directly to the resolution requirements of simulation. During stable periods, coarse-grained models suffice. During transitions, fine-grained dynamics determine outcomes. The challenge is detecting the approach of a transition before it occurs.

3 Percolation Theory: Detecting Transitions Before They Occur

The previous section established that complex systems generically exhibit punctuated dynamics—long periods of stability interrupted by rapid transitions. For agent-based modelers, this creates a practical problem: how do you know when a transition is approaching? If you could detect the approach of a transition before it happens, you could adaptively increase your model's resolution in time to capture the critical dynamics. Percolation theory, originally developed to study fluid flow through porous media [27], provides exactly this capability.

The connection between percolation and phase transitions was recognized early. Broadbent and Hammersley's founding paper noted that the emergence of a connected path through a random medium—what they called "percolation"—exhibits the mathematical structure of a critical phenomenon [27]. Over the following decades, this connection was made precise. Stauffer and Aharony's comprehensive treatment [7] established percolation as a canonical example of a continuous phase transition, complete with critical exponents, scaling laws, and universality classes. The theory was extended to complex networks by Newman, Strogatz, and Watts [28, 29], who showed that percolation on small-world and scale-free networks exhibits similar critical behavior but with threshold values that depend on network topology.

The relevance to social and economic systems became clear through work on cascade dynamics. Watts's 2002 paper "A Simple Model of Global Cascades on Random Networks" [30] demonstrated that social contagion—the spread of behaviors, opinions, or information through a population—maps directly onto percolation. A cascade occurs if and only if the network contains a "percolating vulnerable cluster": a connected component of agents who will adopt a behavior once a threshold fraction of their neighbors have done so. This insight connects the abstract mathematics of percolation to concrete questions about when rumors go viral, when markets crash, and when social movements take off.

The key insight for agent-based modeling is this: percolation theory provides *computable early warning signals* that indicate when a system is approaching a critical transition. Scheffer and colleagues [31] surveyed evidence from ecology, climate science, and finance showing that systems approaching tipping points exhibit characteristic signatures—critical slowing down, increased variance, and growing spatial correlations. These signatures are manifestations of the same underlying phenomenon that percolation theory describes mathematically.

3.1 The Basic Framework: When Local Becomes Global

Consider a social network $G = (V, E)$ where nodes represent agents and edges represent potential pathways for influence, information, or contagion. Not all connections are equally active at all times: a friendship might transmit a rumor today but not tomorrow; a business relationship might propagate a financial shock during a crisis but lie dormant during normal times. We model

this by saying each edge is "active" with probability p, where p reflects the current intensity of interaction in the system.

The central question of percolation theory is: at what threshold probability p_c does a "giant component"—a connected cluster containing a macroscopic fraction of all nodes—first emerge in the network of active edges? Below this threshold, the network fragments into small isolated clusters. Above it, a system-spanning cluster exists through which perturbations can propagate globally.

For random graphs with arbitrary degree distributions, the percolation threshold has a remarkably simple form. Cohen, Erez, ben-Avraham, and Havlin [34] showed that the giant component emerges when:

$$p > p_c = \frac{\langle k \rangle}{\langle k^2 \rangle - \langle k \rangle} \tag{1}$$

where $\langle k \rangle$ is the mean degree and $\langle k^2 \rangle$ is the second moment of the degree distribution. This formula has profound implications. In homogeneous networks where all nodes have similar degree, $\langle k^2 \rangle \approx \langle k \rangle^2$, giving $p_c \approx 1/\langle k \rangle$—you need edges to be active with probability roughly $1/\langle k \rangle$ before global connectivity emerges. But in heterogeneous networks with high-degree hubs, $\langle k^2 \rangle$ can be much larger than $\langle k \rangle^2$, pushing p_c toward zero. In the limit of scale-free networks where the degree distribution follows a power law $P(k) \sim k^{-\gamma}$ with $\gamma \leq 3$, the second moment diverges and $p_c \to 0$: even vanishingly weak connectivity can produce a giant component.

Newman and Watts [28] extended this analysis to small-world networks, showing that even a small density of random "shortcuts"—long-range connections bridging otherwise distant communities—dramatically lowers the percolation threshold. Their key result: the correlation length (the typical distance over which perturbations propagate) scales as $\xi \sim \phi^{-1/2}$ where ϕ is the density of shortcuts. A network with just 1% random long-range connections has a correlation length 10 times larger than one with purely local connections. This explains why social networks, with their characteristic combination of local clustering and occasional long-range ties, are so susceptible to cascades.

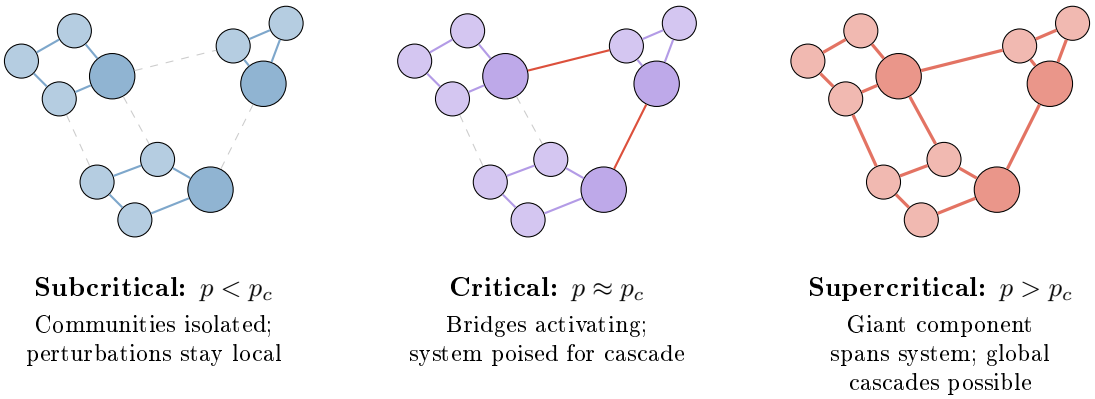

Figure 2: Network percolation in a system with community structure. **Left**: Below threshold, inter-community bridges (dashed) are inactive; each community evolves independently and can be modeled separately. **Center**: At threshold, some bridges activate (red); the system becomes sensitive to which specific bridges are open. **Right**: Above threshold, all communities merge into a giant component; a perturbation anywhere can propagate everywhere. The larger nodes represent high-degree "hubs" whose connections disproportionately influence whether the threshold is crossed.

The structure of the network matters enormously for determining p_c. Figure 2 illustrates the key intuition for networks with community structure—the kind of networks that actually arise in social and economic systems. During normal times (left panel), inter-community connections are sparse or inactive; each community can be modeled independently. As interaction intensity

increases (center panel), bridges between communities begin to activate; the system enters a critical regime where the specific pattern of active bridges determines whether cascades can spread. Above threshold (right panel), a giant component emerges spanning all communities; the system has transitioned to a qualitatively different regime where local events can have global consequences.

The role of network hubs is particularly important. High-degree nodes contribute disproportionately to $\langle k^2 \rangle$, and therefore to lowering the percolation threshold. Pastor-Satorras and Vespignani [**?**] showed that in scale-free networks, epidemics can spread even when the transmission probability per contact is arbitrarily small—the hubs create shortcuts that circumvent any attempt at containment. For ABM, this means that models which treat all agents as equivalent (replacing the heterogeneous degree distribution with its mean) will systematically underestimate cascade risk.

3.2 Critical Exponents: What Percolation Theory Predicts

The power of percolation theory lies in its quantitative predictions. Near the percolation threshold, physical quantities exhibit power-law scaling characterized by *critical exponents*. These exponents are universal—they depend only on the dimensionality of the system and certain symmetry properties, not on microscopic details [7]. This universality means that the same mathematical framework applies whether you are modeling disease spread, information diffusion, or financial contagion.

The fundamental critical exponents for percolation are:

The **correlation length exponent** ν governs how the typical cluster diameter ξ diverges as the threshold is approached:

$$\xi \sim |p - p_c|^{-\nu} \tag{2}$$

For two-dimensional lattices, $\nu = 4/3$ exactly [35]. For three dimensions, high-precision numerical studies give $\nu \approx 0.88$ [36]. For random graphs (mean-field limit), $\nu = 1/2$. The correlation length tells us the distance over which perturbations propagate: when ξ is small, shocks stay local; when ξ spans the system, local events have global consequences.

The **susceptibility exponent** γ governs how the mean cluster size χ (excluding the infinite cluster) diverges:

$$\chi \sim |p - p_c|^{-\gamma} \tag{3}$$

For two dimensions, $\gamma = 43/18 \approx 2.39$; for three dimensions, $\gamma \approx 1.80$; for mean-field, $\gamma = 1$. The susceptibility quantifies system sensitivity: a large χ means that perturbing one agent sets off a chain reaction affecting many others.

The **order parameter exponent** β governs how the fraction P_∞ of nodes in the giant component grows above threshold:

$$P_\infty \sim (p - p_c)^\beta \quad \text{for } p > p_c \tag{4}$$

For two dimensions, $\beta = 5/36 \approx 0.14$; for three dimensions, $\beta \approx 0.41$; for mean-field, $\beta = 1$. The small value of β in low dimensions means the giant component emerges gradually—just above threshold, only a tiny fraction of nodes belong to it.

The **cluster size distribution exponent** τ governs the power-law distribution of cluster sizes at criticality:

$$n_s \sim s^{-\tau} \quad \text{at } p = p_c \tag{5}$$

where n_s is the number of clusters containing s nodes. For two dimensions, $\tau = 187/91 \approx 2.05$; for three dimensions, $\tau \approx 2.19$; for mean-field, $\tau = 5/2$. The appearance of a power-law cluster size distribution—lacking any characteristic scale—is a signature of criticality.

These exponents are not independent. They satisfy *scaling relations* that follow from the underlying mathematical structure [7]:

$$\gamma = (3 - \tau)/(\tau - 1) \tag{6}$$

$$\beta = (\tau - 2)/(\tau - 1) \tag{7}$$

$$2 - \alpha = \nu d \quad \text{(hyperscaling)} \tag{8}$$

where α is the specific heat exponent and d is the spatial dimension. These relations provide internal consistency checks: if your measured exponents violate the scaling relations, something is wrong with your analysis.

Exponent	**2D Lattice**	**3D Lattice**	**Mean-Field**
ν (correlation length)	4/3	0.88	1/2
γ (susceptibility)	43/18	1.80	1
β (order parameter)	5/36	0.41	1
τ (cluster distribution)	187/91	2.19	5/2

Table 1: Critical exponents for percolation in different dimensions. The exact values for 2D follow from conformal field theory [35]; 3D values are from high-precision Monte Carlo simulations [36]; mean-field values apply to random graphs and high-dimensional lattices.

The practical implication is that these exponents tell us *how fast* the warning signs grow as we approach a transition. The larger the exponent, the more dramatic the divergence. In two dimensions, where $\gamma \approx 2.4$, the susceptibility grows very rapidly as $p \to p_c$—you get strong early warning. In mean-field (random graphs), $\gamma = 1$, so the warning is less dramatic but still present.

Observation 1 (Percolation Measures as Resolution Triggers)**.** *When correlation length ξ is small relative to system size and susceptibility χ is of order unity, mean-field approximations and coarse-grained models are safe—local details don't propagate to global outcomes. When ξ and χ begin to grow according to the power laws above, approximations fail—you need full resolution to capture the dynamics correctly.*

3.3 Practical Computation

For an agent-based model with interaction graph G, the percolation measures can be computed efficiently during simulation:

The **correlation length** ξ can be estimated from the spatial decay of state correlations. Sample pairs of agents at various graph distances r. Compute the correlation $C(r)$ between their states (opinions, behaviors, resources). Near criticality, $C(r) \sim r^{-(d-2+\eta)} e^{-r/\xi}$, where η is the anomalous dimension exponent. Fitting this form yields ξ.

The **susceptibility** χ equals the mean cluster size and can be computed exactly using union-find algorithms. Identify connected components in the active interaction graph (edges with activation probability above some threshold). Compute $\chi = \langle s^2 \rangle / \langle s \rangle - S_{\max}$, where the sums run over cluster sizes excluding the largest cluster $S_{\max}$. Union-find runs in nearly linear time with path compression.

The **cluster size distribution** is obtained by histogramming component sizes. To test whether the distribution follows a power law, use the Clauset-Shalizi-Newman procedure [37], which provides both a maximum-likelihood estimate of the exponent τ and a p-value for the power-law hypothesis. At criticality, you expect τ to match the theoretical predictions in Table 1.

The computational cost of these measures is negligible compared to the cost of simulating agent dynamics. Union-find on a graph with n nodes and m edges runs in $O((n + m)\alpha(n))$

time, where α is the inverse Ackermann function—effectively constant. Correlation functions require sampling $O(n)$ pairs. Power-law fitting is $O(n \log n)$. You can monitor these measures continuously and use them to trigger resolution changes dynamically.

For agent-based modelling, this suggests a concrete strategy: continuously compute percolation-theoretic measures on the interaction graph, and use these as triggers for resolution adjustment.

4 The Adaptive Resolution Framework

We now present the Adaptive Resolution Modelling framework. The core idea is to maintain a hierarchy of representations at different coarse-graining levels, and dynamically switch between them based on percolation measures.

4.1 Hierarchical Graph Representation

Consider a system of N agents with interaction graph $G_0 = (V_0, E_0)$. We construct a hierarchy of coarse-grained representations:

$$G_0 \to G_1 \to G_2 \to \cdots \to G_L \tag{9}$$

where each $G_{\ell+1}$ is obtained from G_ℓ by aggregating clusters of nodes into "supernodes." At level ℓ, we have $|V_\ell|$ effective agents, with $|V_0| = N$ and $|V_L| \ll N$.

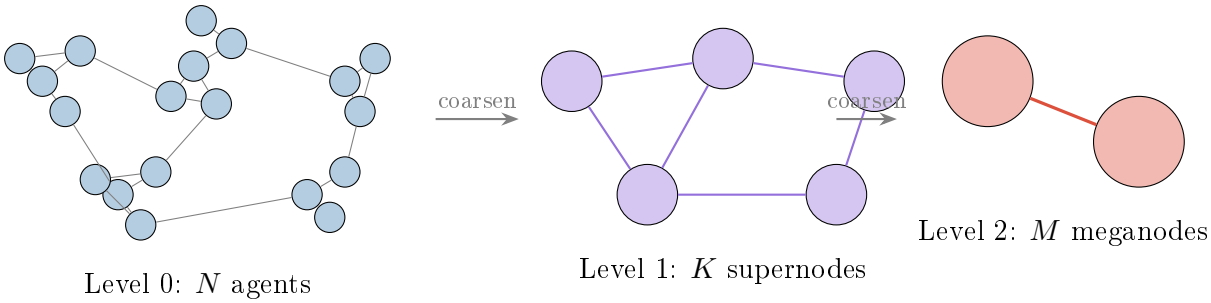

Figure 3: Hierarchical coarse-graining of an agent network. Level 0 contains all N agents. Level 1 aggregates clusters into supernodes, preserving inter-cluster connectivity. Level 2 further aggregates into meganodes. Each level is a valid simulation target with different resolution-computation tradeoffs.

The key insight is that each supernode at level $\ell+1$ inherits all external connections from its constituent nodes at level ℓ. If nodes $\{v_1, \ldots, v_k\} \subset V_\ell$ are aggregated into supernode $u \in V_{\ell+1}$, then u connects to supernode $w \in V_{\ell+1}$ if and only if some v_i connected to some node in the cluster that formed w. The Markov blanket of the supernode is composed from the blankets of its constituents.

4.2 Finding Natural Aggregation Boundaries: Where Should We Draw the Lines?

Not all aggregations are equally good. If you lump the wrong agents together, you destroy important structure; if you keep too much detail, you waste computation. The question is: where are the *natural* boundaries in the system—the places where aggregation loses the least information?

This question has a precise answer rooted in information theory. The key insight is to think about *mutual information flow* through the network.

4.2.1 Mutual Information as Edge Weight

Consider two agents X and Y that interact. The mutual information $I(X;Y)$ measures how much knowing about X tells you about Y [10]:

$$I(X;Y) = H(X) + H(Y) - H(X,Y) \tag{10}$$

where $H(\cdot)$ denotes entropy. If X and Y are independent, $I(X;Y) = 0$—knowing one tells you nothing about the other. If they're perfectly correlated, $I(X;Y) = H(X) = H(Y)$—knowing one completely determines the other.

Now construct a weighted graph where each edge (X,Y) has weight $w(X,Y) = I(X;Y)$. This graph encodes the complete statistical dependency structure of the system. High-weight edges connect agents that "know about" each other; low-weight edges connect agents that are nearly independent.

4.2.2 The Markov Blanket as Minimum Cut

A *Markov blanket* is the minimal set of variables that renders a subsystem statistically independent from everything else [5, 3]. Once you know the Markov blanket, the internal states of the subsystem give you no additional information about the external world, and vice versa.

A *minimum cut* is the minimal set of edges that disconnects a graph into two pieces.

Here's the key observation: *on the mutual information graph, these are the same thing.* The minimum cut finds the partition where the total information flow across the boundary is minimized. This is precisely the Markov blanket—the minimal statistical boundary separating inside from outside.

Why does this matter for agent-based modeling? Because it tells us where we can *safely aggregate*. If a cluster of agents is separated from the rest of the system by a low-weight cut, then:

1. The cluster's internal dynamics have minimal effect on the external world (information flows out slowly)

2. The external world has minimal effect on the cluster's internal dynamics (information flows in slowly)

3. We can therefore model the cluster as a single effective agent with aggregate properties—a *mean-field approximation*—with minimal loss of accuracy

4.2.3 Example: Social Network Segmentation in Marketing

Consider a company modeling consumer behavior on a social network of 100,000 users. Running a full agent-based simulation is expensive. Can they safely use simpler models?

They compute mutual information between users based on behavioral correlations: shared purchases, similar browsing patterns, correlated engagement times. This produces a weighted graph. Running a minimum cut algorithm reveals the natural community structure:

High-cut regions: The fitness enthusiasts form a tight cluster with high internal mutual information (they influence each other's purchase decisions constantly) but low mutual information with the vintage car collectors. These groups can be modeled as separate mean-field populations—the internal details of one group don't affect the other.

Low-cut regions: But certain "bridge" individuals have high mutual information with multiple communities. The fitness enthusiast who also restores classic cars creates an information pathway between otherwise separated groups. Near such bridges, the mean-field approximation breaks down.

Marketing professionals already do something like this intuitively when they create "customer personas"—archetypical profiles representing market segments. What they're implicitly doing is drawing Markov blankets around clusters and treating each cluster as a homogeneous population. Our framework makes this rigorous: the validity of the persona-based approach depends on whether the inter-segment mutual information is actually low.

During normal times, it probably is—each community evolves somewhat independently, and treating them as separate populations is fine. But during viral moments or cultural shifts, information cascades *across* community boundaries. The cut capacity increases. The mean-field approximation fails. This is precisely when you need to switch to higher resolution.

4.2.4 Computing the Minimum Cut

The **Stoer-Wagner algorithm** [6] finds the global minimum cut efficiently without requiring us to specify which partition we're looking for:

Algorithm 1 Stoer-Wagner Minimum Cut [6]

procedure GLOBALMINCUT($G = (V, E, w)$)
  best_cut $\leftarrow \infty$
  while $|V| > 1$ **do**
    $(s, t, \text{cut\_value}) \leftarrow$ MINCUTPHASE(G)
    if cut_value $<$ best_cut **then**
      best_cut $\leftarrow$ cut_value
      Record the partition separating t
    end if
    Merge nodes s and t in G
  end while
  return best_cut and its partition
end procedure

procedure MINCUTPHASE(G)
  Pick arbitrary starting node a
  $A \leftarrow \{a\}$
  while $A \neq V$ **do**
    Add to A the node $v \notin A$ maximizing $\sum_{u \in A} w(u, v)$
  end while
  Let s be second-to-last added, t the last
  return $(s, t, \sum_{v \neq t} w(v, t))$
end procedure

The algorithm works by growing a cluster greedily—always adding the node with highest total edge weight to the current cluster—then recording the cut that separates the last-added node. After each phase, it merges the final two nodes and repeats until the graph collapses to a single node. The smallest cut encountered is the global minimum. The complexity is $O(VE + V^2 \log V)$, tractable for moderately sized systems.

For hierarchical coarse-graining, we apply this recursively: find the minimum cut, aggregate one side into a supernode, then find the next minimum cut, and so on. This produces a hierarchy of natural boundaries at different scales.

4.2.5 The Central Question: When Can You Use Mean-Field Approximations?

We can now state the central practical question precisely: a mean-field approximation is valid when the Markov blanket around a cluster has low capacity—when the total mutual information

crossing the boundary is small. The adaptive resolution framework monitors this continuously:

- When cut capacities are low and correlation length is small: use coarse-grained mean-field models. The internal details of each cluster don't matter for inter-cluster dynamics.
- When cut capacities increase and correlation length grows: switch to finer resolution. Information is flowing across boundaries that previously contained it; the mean-field assumption is breaking down.
- At criticality: no aggregation is safe. Information flows at all scales, and you need full resolution to capture the dynamics.

This is the payoff of the framework: it tells you not just *whether* to aggregate, but *where* the natural boundaries are and *when* those boundaries become porous.

4.3 The Switching Protocol

The complete protocol operates as follows:

1. **Initialize** at some resolution level ℓ.
2. **Simulate** dynamics at level ℓ for time window Δt.
3. **Compute percolation measures**: correlation length ξ, susceptibility χ, largest cluster fraction.
4. **Compare to thresholds**:
   - If measures exceed upper threshold θ_+: descend to finer resolution $\ell - 1$
   - If measures fall below lower threshold θ_-: ascend to coarser resolution $\ell + 1$
   - Otherwise: remain at current level
5. **If switching resolution**:
   - Descending: instantiate fine-grained states consistent with current coarse state
   - Ascending: aggregate fine states into coarse representation
6. **Repeat** from step 2.

4.4 Computational Complexity Analysis

Let N be the number of agents at the finest level, and suppose each coarse-graining step reduces node count by factor r (so level ℓ has N/r^ℓ nodes). Let T be total simulation time and Δt the window between percolation checks.

In the worst case (always at finest resolution), complexity is $O(T/\Delta t \cdot f(N))$ where $f(N)$ is the per-step simulation cost.

With adaptive resolution, if the system spends fraction α_ℓ of time at level ℓ, total cost becomes:

$$\text{Cost} = \frac{T}{\Delta t} \sum_{\ell=0}^{L} \alpha_\ell \cdot f(N/r^\ell) \tag{11}$$

For systems exhibiting punctuated equilibria, α_0 (time at finest resolution) may be small—perhaps 5-10% corresponding to critical windows. The bulk of simulation time occurs at coarse levels with $f(N/r^\ell) \ll f(N)$, yielding substantial computational savings.

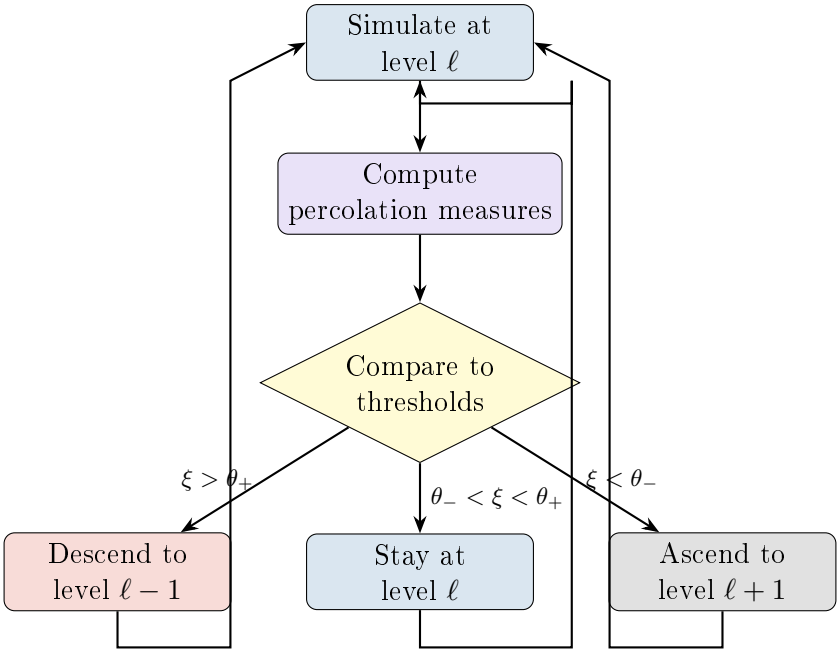

Figure 4: The adaptive resolution switching protocol. Percolation measures are computed after each simulation window; resolution is adjusted based on proximity to critical thresholds.

5 Theoretical Foundations

5.1 Why Percolation Predicts Cascade Susceptibility

The connection between percolation measures and system predictability is not merely empirical. Near percolation thresholds, the system exhibits *critical slowing down*: perturbations decay more slowly, correlations extend further, and fluctuations grow. These are precisely the conditions under which coarse-grained descriptions fail.

Consider a perturbation δ applied to a single node. In the subcritical regime, this perturbation affects on average $O(\chi)$ other nodes, where χ is the susceptibility. When χ is small, effects remain local; coarse descriptions capture the relevant dynamics. When χ diverges (near threshold), the perturbation can reach $O(N)$ nodes; microscale details determine macroscale outcomes.

The correlation length ξ tells us the spatial scale of these effects. When $\xi \ll L$ (system size), the system decomposes into effectively independent subsystems that can be simulated separately. When $\xi \sim L$, the whole system is correlated; no decomposition preserves the relevant dynamics.

5.2 Connection to Renormalization Group

Our coarse-graining procedure has deep connections to renormalization group (RG) methods in statistical physics. RG theory systematically integrates out short-wavelength fluctuations to obtain effective theories at larger scales. The key insight is that most microscopic parameters become *irrelevant* under RG flow—they don't affect long-wavelength behavior.

Near fixed points (including critical points), the system becomes *scale-invariant*: it looks statistically similar at all scales. This is precisely when coarse-graining fails to simplify the problem—there is no natural scale at which to truncate.

Our adaptive scheme respects this structure: far from criticality, coarse-graining is justified because irrelevant details have negligible macroscopic effects; near criticality, we maintain resolution because scale invariance means details at all scales matter.

5.3 Connection to Active Inference and Nested Markov Blankets

The coarse-graining procedure we describe has deep connections to the free energy principle and active inference [3, 11]. In this framework, a Markov blanket defines the statistical boundary of an agent—the minimal set of states that separates internal states from external states [5].

Crucially, Kirchhoff et al. [4] show that autonomous systems are hierarchically composed of "Markov blankets of Markov blankets"—all the way down to individual cells, all the way up to organisms and social groups. This nested structure is precisely what our coarse-graining hierarchy captures: each supernode has its own Markov blanket composed from the blankets of its constituents.

When we aggregate a cluster of agents into a supernode, we are identifying a collective Markov blanket—a boundary across which information flow is minimized. The minimum cut on the mutual information graph, computed via Stoer-Wagner, discovers these natural boundaries without assuming the structure in advance.

This connection suggests that our adaptive resolution framework is not merely a computational convenience but reflects the actual multi-scale organization of agency in complex systems. The marketing team's "customer personas" are not arbitrary abstractions—they are collective Markov blankets, natural units of statistical independence in the social network. When we say that a mean-field approximation is valid, we are saying that these collective blankets have low permeability: information stays mostly inside the cluster or mostly outside, crossing the boundary rarely.

Coarse-graining is justified when nested Markov blankets are well-defined and stable; it breaks down near criticality when these statistical boundaries become porous and information flows across all scales. The percolation measures we monitor are, in effect, measuring the integrity of the blanket hierarchy: as correlation length grows, blankets at successively larger scales become leaky, until at criticality the entire hierarchy dissolves into a single connected system where every agent's state is correlated with every other's.

5.4 Information-Theoretic Interpretation

From an information-theoretic perspective, our framework allocates bits of simulation precision proportionally to the information content of different temporal regimes.

During stable periods, the system's trajectory is highly predictable from coarse statistics; detailed microstate information is redundant. The entropy rate is low. Coarse simulation captures essentially all the relevant information.

During critical windows, the system's trajectory becomes sensitive to microscopic fluctuations; knowing the coarse state tells you little about the fine state (or the future coarse state). The entropy rate spikes. Only fine simulation captures the relevant dynamics.

Adaptive resolution thus implements a form of *rate-distortion optimal encoding* of the system trajectory: we allocate representational precision where it matters for predicting outcomes.

6 Discussion

6.1 Relation to Existing Approaches

Adaptive mesh refinement (AMR) in computational physics uses related ideas: refine spatial resolution where gradients are steep, coarsen where fields are smooth. Our contribution adapts this principle to agent-based models with graph structure, using percolation theory to detect where "gradients" in collective behavior demand finer resolution.

Multi-scale modeling approaches in materials science similarly combine atomic-level and continuum descriptions. The key challenge—consistent coupling between scales—is addressed in our framework through the inherited connectivity of supernodes.

6.2 Limitations and Extensions

The framework as presented assumes that percolation measures are computationally cheap relative to simulation. This holds for sparse graphs but may require approximation algorithms for dense networks.

The choice of thresholds θ_+ and θ_- involves tradeoffs: too aggressive switching wastes resources on unnecessary resolution changes; too conservative switching misses critical windows. Adaptive threshold learning based on prediction error is a natural extension.

The instantiation problem when descending resolution—sampling fine states consistent with coarse states—requires care. Maximum entropy methods provide principled approaches, but computational cost and consistency with dynamics remain open challenges.

6.3 Applications

The framework applies naturally to:

Epidemic modeling: Disease spread exhibits threshold effects; near epidemic thresholds, individual transmission events matter; far from threshold, aggregate dynamics suffice.

Financial contagion: Systemic risk involves cascade effects through network connections; near fragility thresholds, detailed counterparty relationships matter.

Opinion dynamics: Social influence models show phase transitions between consensus and fragmentation; near these transitions, local interaction patterns shape global outcomes.

Ecological networks: Ecosystem stability involves threshold effects; near tipping points, species-level dynamics determine whether cascades occur.

7 Conclusion

Complex systems demand simulation approaches that respect their actual dynamics. Punctuated equilibria—long stability interrupted by brief criticality—is the generic pattern. Adaptive Resolution Modelling responds to this pattern by treating computational resources as scarce and allocating them where they yield information: at and near phase transitions.

Percolation theory provides the early warning system, detecting approaching criticality through diverging correlation length and susceptibility. Graph-based coarse-graining provides the multi-resolution hierarchy, aggregating agents into supernodes that inherit connectivity and can be treated as effective agents in their own right.

The result is a framework that promises substantial computational savings for systems exhibiting the punctuated pattern—which is to say, most complex systems of interest. By computing more when it matters and less when it doesn't, we can extend the reach of agent-based modeling to larger systems and longer timescales.

A Percolation Measures: Computational Details

For a graph $G = (V, E)$, key percolation measures can be computed as follows.

Largest component fraction: Run connected components algorithm (breadth-first search from each unvisited node); return size of largest component divided by $|V|$. Complexity: $O(|V| + |E|)$.

Mean cluster size (susceptibility proxy):

$$\chi = \frac{\sum_s s^2 n_s}{\sum_s s \cdot n_s} \tag{12}$$

where n_s is the number of clusters of size s. This weights clusters by their size squared, emphasizing large clusters.

Correlation length (approximation): For computational efficiency, use the mean distance within clusters as a proxy:

$$\xi \approx \left\langle \frac{1}{|C|^2} \sum_{u,v \in C} d(u,v) \right\rangle_C \tag{13}$$

where the average is over clusters C and $d(u,v)$ is graph distance.

B Coarse-Graining Consistency Conditions

For the hierarchy to be consistent, coarse dynamics must be derivable from fine dynamics in an appropriate limit. Specifically, if $\phi_\ell : G_\ell \to G_{\ell+1}$ is the coarse-graining map, we require:

$$\mathbb{E}[\phi_\ell(X_t^\ell) | X_0^{\ell+1}] \approx X_t^{\ell+1} \tag{14}$$

where X_t^ℓ is the state at level ℓ at time t, and the expectation averages over fine-scale fluctuations. This ensures that coarse dynamics are the appropriate average of fine dynamics—not an arbitrary simplification.

References

[1] Eldredge, N. & Gould, S.J. (1972). Punctuated equilibria: an alternative to phyletic gradualism. In T.J.M. Schopf (Ed.), *Models in Paleobiology* (pp. 82–115). Freeman, Cooper & Co., San Francisco.

[2] Gould, S.J. (2002). *The Structure of Evolutionary Theory*. Harvard University Press, Cambridge, MA.

[3] Friston, K. (2013). Life as we know it. *Journal of the Royal Society Interface*, 10(86), 20130475.

[4] Kirchhoff, M., Parr, T., Palacios, E., Friston, K., & Kiverstein, J. (2018). The Markov blankets of life: autonomy, active inference and the free energy principle. *Journal of the Royal Society Interface*, 15(138), 20170792.

[5] Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann.

[6] Stoer, M. & Wagner, F. (1997). A simple min-cut algorithm. *Journal of the ACM*, 44(4), 585–591.

[7] Stauffer, D. & Aharony, A. (1994). *Introduction to Percolation Theory* (2nd ed.). Taylor & Francis, London.

[8] Dorogovtsev, S.N., Goltsev, A.V., & Mendes, J.F.F. (2008). Critical phenomena in complex networks. *Reviews of Modern Physics*, 80, 1275–1335.

[9] Li, M., Liu, R.-R., Lü, L., Hu, M.-B., Xu, S., & Zhang, Y.-C. (2021). Percolation on complex networks: Theory and application. *Physics Reports*, 907, 1–68.

[10] Cover, T.M. & Thomas, J.A. (2006). *Elements of Information Theory* (2nd ed.). Wiley-Interscience.

[11] Parr, T., Pezzulo, G., & Friston, K.J. (2022). *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. MIT Press.

[12] Lempert, R. (2002). Agent-based modeling as organizational and public policy simulators. *Proceedings of the National Academy of Sciences*, 99(suppl 3), 7195–7196.

[13] Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, 99(suppl 3), 7280–7287.

[14] Epstein, J.M. & Axtell, R.L. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. MIT Press, Cambridge, MA.

[15] Parry, H.R. & Bithell, M. (2012). Large scale agent-based modelling: A review and guidelines for model scaling. In A.J. Heppenstall, A.T. Crooks, L.M. See, & M. Batty (Eds.), *Agent-Based Models of Geographical Systems* (pp. 271–308). Springer.

[16] An, G., Fitzpatrick, B.G., Christley, S., Federico, P., Kanarek, A., Miller Neilan, R., Oremland, M., Salerno, R., Laubenbacher, R., & Lenhart, S. (2017). Optimization and control of agent-based models in biology: A perspective. *Bulletin of Mathematical Biology*, 79(1), 63–87.

[17] Macy, M.W. & Willer, R. (2002). From factors to actors: Computational sociology and agent-based modeling. *Annual Review of Sociology*, 28, 143–166.

[18] Filatova, T., Polhill, J.G., & van Ewijk, S. (2025). The power of bridging decision scales: Model coupling for advanced climate policy analysis. *Proceedings of the National Academy of Sciences*, 122(37), e2411592122.

[19] Eldredge, N. (2008). The early "evolution" of "punctuated equilibria". *Evolution: Education and Outreach*, 1, 107–113.

[20] Lieberman, B.S. & Eldredge, N. (2008). Punctuated equilibria. *Scholarpedia*, 3(10), 3806.

[21] Hunt, G., Hopkins, M.J., & Lidgard, S. (2015). Simple versus complex models of trait evolution and stasis as a response to environmental change. *Proceedings of the National Academy of Sciences*, 112(16), 4885–4890.

[22] Schumpeter, J.A. (1942). *Capitalism, Socialism and Democracy*. Harper & Brothers, New York.

[23] Mokyr, J. (1990). Punctuated equilibria and technological progress. *American Economic Review*, 80(2), 350–354.

[24] Perez, C. (2010). Technological revolutions and techno-economic paradigms. *Cambridge Journal of Economics*, 34(1), 185–202.

[25] Baumgartner, F.R. & Jones, B.D. (1993). *Agendas and Instability in American Politics*. University of Chicago Press.

[26] Baumgartner, F.R., Breunig, C., Green-Pedersen, C., Jones, B.D., Mortensen, P.B., Nuytemans, M., & Walgrave, S. (2009). Punctuated equilibrium in comparative perspective. *American Journal of Political Science*, 53(3), 603–620.

[27] Broadbent, S.R. & Hammersley, J.M. (1957). Percolation processes: I. Crystals and mazes. *Mathematical Proceedings of the Cambridge Philosophical Society*, 53(3), 629–641.

[28] Newman, M.E.J. & Watts, D.J. (1999). Scaling and percolation in the small-world network model. *Physical Review E*, 60(6), 7332–7342.

[29] Newman, M.E.J., Strogatz, S.H., & Watts, D.J. (2001). Random graphs with arbitrary degree distributions and their applications. *Physical Review E*, 64(2), 026118.

[30] Watts, D.J. (2002). A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9), 5766–5771.

[31] Scheffer, M., Bascompte, J., Brock, W.A., Brovkin, V., Carpenter, S.R., Dakos, V., Held, H., van Nes, E.H., Rietkerk, M., & Sugihara, G. (2009). Early-warning signals for critical transitions. *Nature*, 461(7260), 53–59.

[32] Kesten, H. (1980). The critical probability of bond percolation on the square lattice equals 1/2. *Communications in Mathematical Physics*, 74(1), 41–59.

[33] Erdős, P. & Rényi, A. (1960). On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5, 17–61.

[34] Cohen, R., Erez, K., ben-Avraham, D., & Havlin, S. (2000). Resilience of the Internet to random breakdowns. *Physical Review Letters*, 85(21), 4626–4628.

[35] Smirnov, S. (2001). Critical percolation in the plane: Conformal invariance, Cardy's formula, scaling limits. *Comptes Rendus de l'Académie des Sciences - Series I - Mathematics*, 333(3), 239–244.

[36] Wang, J., Zhou, Z., Zhang, W., Garoni, T.M., & Deng, Y. (2013). Bond and site percolation in three dimensions. *Physical Review E*, 87(5), 052107.

[37] Clauset, A., Shalizi, C.R., & Newman, M.E.J. (2009). Power-law distributions in empirical data. *SIAM Review*, 51(4), 661–703.