Module 8 — Network Games

Game Theory

hbar.university

Graph Games and Strategic Topology

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Network Games **Node:** 8.1

Module 8 is the point at which game theory stops being a theory of "n players in a room" and becomes a theory of **interaction topologies**. In Modules 1–7, when we wrote $u_i(s_i, s_{-i})$, the index $-i$ meant "everyone else," and the combinatorial object holding the game together was the complete Cartesian product $S_1 \times \cdots \times S_n$. Every player's payoff could, in principle, depend on every other player's action. That is the *complete-graph* assumption: the interaction structure is K_n, and it is typically left implicit. Once we take interaction topology seriously — once we recognize that Alice only interacts with her neighbors, that influence propagates along edges rather than through a room — we are in the world of **graph games**. This node establishes the framework: a game on a graph, players on nodes, payoffs depending only on the actions of graph-neighbors, and the structural consequences of this locality. Every subsequent node in Module 8 — coordination (8.2), diffusion (8.3), public goods (8.4), formation (8.5), complementarities (8.6) — is a specialization of the general graph-game framework introduced here.

Formal Definition

A **graph game** (or **network game**) is a tuple

$$\Gamma = (N, G, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$$

where:

- $N = \{1, \dots, n\}$ is the set of **players**, each identified with a node.
- $G = (N, E)$ is a graph (directed or undirected) on N, called the **interaction network**. For $i \in N$, write $\mathcal{N}_i = \{j : (i, j) \in E\}$ for the **neighborhood** of i, and $d_i = |\mathcal{N}_i|$ for the **degree** of i.
- S_i is player i's strategy set.
- $u_i : S_i \times S_{\mathcal{N}_i} \to \mathbb{R}$ is player i's payoff function, **depending only on i's own strategy and the strategies of i's neighbors**.

The **locality axiom** is the defining constraint: u_i factors through the neighborhood, $u_i(s) = \tilde{u}_i(s_i, s_{\mathcal{N}_i})$. A game where every u_i depends on all of s_{-i} is the special case $G = K_n$ (complete graph), recovering the standard normal-form game.

A **symmetric** graph game is one where u_i has the same functional form for every i; payoff differences come entirely from differences in $s_{\mathcal{N}_i}$ (which depend on the graph position). A **linear-aggregative** graph game is one where $u_i(s_i, s_{\mathcal{N}_i}) = f(s_i, \sum_{j \in \mathcal{N}_i} s_j)$ — payoff depends on own action and the *sum* of neighbor actions. Linear-aggregative games are the workhorse of the network-games literature because they admit closed-form equilibria via the graph Laplacian.

A **Nash equilibrium** of a graph game is a profile s^* such that for each i,

$$s_i^* \in \arg\max_{s_i \in S_i} \tilde{u}_i(s_i, s^*_{\mathcal{N}_i}).$$

The locality axiom means the best-response computation is local: player i only needs to know the actions of neighbors, not of distant players. This local-best-response structure is what makes network games computationally and conceptually tractable.

Intuition

The move from "complete interaction" to "graph interaction" changes the game in three ways. First, **payoffs are local**: each player's welfare depends only on direct connections. Second, **information can be local**: best responses can be computed by looking only at neighbors. Third, and most important, **equilibrium structure inherits graph structure**: the spectrum of the interaction graph (eigenvalues, Bonacich centrality, spectral radius) appears in closed-form equilibrium expressions. A dense interaction graph and a sparse interaction graph produce qualitatively different strategic landscapes, and the graph itself becomes a parameter of the theory.

Think of the difference between "everyone in the room votes" and "everyone talks to their neighbors, and behavior propagates." The first is a normal-form game on K_n; the second is a graph game on a sparse network. In the first, no topology matters — the outcome depends only on payoffs and equilibrium selection. In the second, topology is everything: who is adjacent to whom determines who influences whom, which determines which equilibria are stable, which cascades can propagate, and whose individual position in the graph gives them strategic leverage.

The strategic role of **position**. In a graph game, two otherwise identical players can have radically different equilibrium payoffs because they sit at different positions in the graph. A player at the center of a star has payoffs averaged over many neighbors; a player at the tip of a spoke has a payoff determined by the center alone. Degree, centrality, spectral position, and cluster membership all become strategic assets. This is the game-theoretic foundation of "network power" — the idea that in a connected system, *where* you are matters as much as *what* you do.

A useful framing: a graph game is a factor-graph specification of a strategic situation. The graph tells you *which pieces of the joint strategy profile matter for which player's payoff.* Anything outside the neighborhood is irrelevant to that player's best response, which means best-response dynamics, learning dynamics, and equilibrium computation all decompose along the graph. The locality axiom is what makes network games fit into Module 8 — it is simultaneously the defining restriction and the source of all the tractability gains.

Structural Role

8.1 is the foundational node of Module 8. Every subsequent node in the module is a specialization of the graph-game framework to a particular class of games:

- **8.2 Coordination on networks** specializes to games where $u_i(s_i, s_{\mathcal{N}_i})$ rewards matching with neighbors — producing Morris-style contagion, thresholds, and best-response dynamics on graphs.
- **8.3 Diffusion and cascades** specializes to games where adoption decisions depend on the fraction of adopted neighbors — producing the Kempe–Kleinberg–Tardos framework of influence maximization.
- **8.4 Public goods on graphs** specializes to games where effort is a local public good — producing the Bramoullé–Kranton specialization theory.
- **8.5 Network formation games** lets the graph G itself be a strategic variable, endogenizing the topology.
- **8.6 Strategic complementarities** restricts attention to games where higher actions by neighbors make higher own actions more attractive — producing Topkis's supermodularity theory and the global-games refinement.

Beyond Module 8, the graph-game framework appears in Module 9 (learning on networks), Module 10 (mechanism design on networks), and throughout the strategic systems integration modules. It is also the point of contact with network science, statistical physics (Ising models on graphs are special cases of coordination graph games), and distributed optimization.

Mathematical Development

Linear-quadratic network games

The canonical tractable class is the **linear-quadratic graph game**: each player i chooses an effort level $x_i \in \mathbb{R}_{\geq 0}$ with payoff

$$u_i(x_i, x_{\mathcal{N}_i}) = \alpha_i x_i - \tfrac{1}{2}x_i^2 + \phi \sum_{j \in \mathcal{N}_i} x_i x_j.$$

Here α_i is an idiosyncratic marginal benefit, the $-\frac{1}{2}x_i^2$ term is a private cost, and ϕ is the **strategic interaction parameter** — positive for complementarities (neighbors' efforts make my effort more attractive), negative for substitutes (neighbors' efforts make my effort less attractive).

Best response. Setting $\partial u_i / \partial x_i = 0$:

$$x_i^* = \alpha_i + \phi \sum_{j \in \mathcal{N}_i} x_j^*.$$

In matrix form, let A be the adjacency matrix of G:

$$x^* = \alpha + \phi A x^* \iff (I - \phi A)x^* = \alpha.$$

If $|\phi| < 1/\rho(A)$, where $\rho(A)$ is the spectral radius of A, then $(I - \phi A)$ is invertible and

$$x^* = (I - \phi A)^{-1}\alpha = \sum_{k=0}^{\infty} \phi^k A^k \alpha.$$

This is the **Bonacich centrality** of α under the interaction graph G with decay parameter ϕ. The walk-counting interpretation: $(A^k)_{ij}$ counts walks of length k from i to j, so x_i^* aggregates influences from nodes at all graph distances, weighted by ϕ^k.

Theorem (Ballester–Calvó-Armengol–Zenou, 2006). The unique Nash equilibrium of the linear-quadratic network game is proportional to the Bonacich centrality vector of the interaction graph. Equilibrium effort at node i is monotone in i's Bonacich centrality.

This is the most celebrated result in network games: a player's equilibrium effort is determined by their **graph-theoretic centrality**, and the centrality measure that matters is not degree or betweenness but the walk-counting Bonacich measure indexed by ϕ.

Best-response dynamics on graphs

Starting from any profile, run asynchronous best-response updates: at each step, a player reads their neighbors' current actions and updates their own to a best response. In a linear-quadratic game with $|\phi| < 1/\rho(A)$, this dynamic is a contraction and converges globally to the unique equilibrium. In games with strategic complementarities (8.6), best-response dynamics converge monotonically to equilibria because the joint best-response map is monotone in the Topkis sense.

The **speed of convergence** depends on graph structure — specifically on the spectral gap of A. Well-connected graphs (high spectral gap) converge fast; graphs with bottlenecks or communities converge slowly. This is the same spectral-gap phenomenon that governs mixing times of random walks, and the analogy is not accidental: best-response dynamics on a linear-quadratic game *are* a linear dynamical system driven by A.

Local vs global equilibria

Because payoffs depend only on neighbors, a graph game can have equilibria that are locally consistent (every player is best-responding to their neighbors) but globally inefficient. The gap between the Nash equilibrium and the social optimum depends on graph structure — cut-vertices, clustering, and the presence of long paths all matter. This gap is the origin of the "specialization" result of Bramoullé–Kranton (8.4) and the "stability-efficiency tradeoff" of Jackson–Wolinsky (8.5).

Worked Example

Linear-quadratic game on a star

Consider a star graph $K_{1,n}$ with one central node c connected to n peripheral nodes $p_1, \dots, p_n$. Each player has the same idiosyncratic marginal benefit $\alpha_i = \alpha$. Strategic parameter $\phi > 0$, small enough for uniqueness.

Best responses:

$$x_c^* = \alpha + \phi \sum_{j=1}^{n} x_{p_j}^*, \qquad x_{p_j}^* = \alpha + \phi x_c^*.$$

By symmetry, all peripherals choose the same value x_p^*. Substituting:

$$x_c^* = \alpha + n\phi x_p^*, \qquad x_p^* = \alpha + \phi x_c^*.$$

Solving:

$$x_c^* = \frac{\alpha(1 + n\phi)}{1 - n\phi^2}, \qquad x_p^* = \frac{\alpha(1 + \phi)}{1 - n\phi^2}.$$

For $\phi > 0$, $x_c^* > x_p^*$: the central node exerts higher effort. The ratio $x_c^*/x_p^* = (1 + n\phi)/(1 + \phi)$ grows with n. A player at a more central position chooses a larger action because *more* neighbors' actions feed back to it, amplifying the private return on effort.

Structural reading. The star exhibits a pure centrality premium. Two players with identical preferences choose different actions because they sit at different graph positions. The center's effort is a weighted average of idiosyncratic benefit α and the feedback from n peripherals; the peripherals' effort feels only the center. This is the kernel of the Bonacich-centrality result: in any graph, equilibrium effort tracks walk-weighted centrality.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Nodes N	Players in a community, coaches, content creators, study groups
Interaction graph G	Who shares hands with whom, whose Discord you are in, who you talk strategy with
Neighborhood $\mathcal{N}_i$	Your immediate circle of exchange partners
Strategic action x_i	Time invested in solver study, range construction, review
Complementarity $\phi > 0$	Study is more valuable when peers study (shared language, peer review)
Bonacich centrality	"Plugged-in-ness": reach into the aggregate study network

The poker ecosystem is **not** a complete graph. It is a sparse network of regulars, study groups, Discords, coaches, staked stables, and content creators. Your information and strategy circulate only through the edges you actually have. Two equally talented players who sit at different positions in this graph will end up with different strategies and different win rates, not because of talent but because of topology.

Concrete situation

Consider a mid-stakes NLHE regular, Alex, who is part of a 5-person study group plus two Discord servers. Alex's "interaction graph" includes the other 4 study-group members (dense edges — daily hand discussion), roughly 20 active Discord contacts (sparse edges — occasional spot review), and a larger tail of 200+ lurkers who read but do not contribute. Alex's neighborhood in the graph game is the dense core: 4 study-group members plus maybe 5 high-touch Discord contacts.

Alex's effort level (hours on solver work) is a best response to the study network's effort. If the group is running deep on a particular topic (turn probe-bet frequencies in single-raised pots), Alex's marginal benefit from doing the same is high — they can share findings, compare trees, catch each other's errors. If the group is checked out, Alex's marginal benefit collapses: why solve in isolation when the work will not circulate? This is a **positive-ϕ** linear-quadratic game on Alex's study graph, and Alex's equilibrium effort is proportional to the Bonacich centrality of Alex within that graph.

Now contrast with Blake, a similarly skilled regular who is in no study group and two low-activity Discords. Blake's neighborhood is effectively empty — isolated node. Blake's best-response effort is the base level α with no multiplier. Alex ends up studying much more than Blake even if they start with identical α. And critically, Alex's study time is determined by Alex's position in the graph, not by innate traits.

What this understanding buys you

Your poker "work ethic" is not fully yours. A significant fraction of your study and game effort is a best response to your local network's effort. If you are in a strong study group, you will naturally put in more work. If you are isolated, you will naturally put in less. This is not a moral failing — it is a structural equilibrium of a graph game with strategic complementarities. The right intervention is not to "try harder" but to **change your graph position**: join a study group, cultivate a small set of high-activity contacts, become central in a live network of regulars who care. Changing your neighborhood changes your ϕA multiplier, which changes your equilibrium effort.

Centrality has a dual: exposure. A very central position in the poker information network also means that your strategy is visible to many people. If you are the center of a 20-person Discord, your hand histories circulate faster than a peripheral player's, and anything you do gets seen, studied, counter-exploited. The same centrality that raises your effort also raises your visibility. The optimal graph position trades off effort-amplification against information-leakage, and 8.1 is the framework in which to reason about that tradeoff.

Cross-Links

- **1.1 Players, strategies, payoffs** — graph games are the locality-constrained specialization.
- **1.4 Best response** — extends to the local best response in graph games.
- **2.1 Nash equilibrium** — specializes to equilibria of linear-quadratic graph games via Bonacich centrality.
- **8.2 Coordination on networks** — the next specialization, to coordination payoffs.
- **8.4 Public goods on graphs** — Bramoullé–Kranton substitutes game.
- **8.5 Network formation games** — endogenizing the graph itself.
- **8.6 Strategic complementarities** — the monotone-comparative-statics specialization.
- **9.3 Learning on graphs** — best-response dynamics and convergence via spectral gap.
- **Network science** — the graph-theoretic layer on which everything in Module 8 is built.

Structural Compression

Invariant: A graph game is a strategic situation in which each player's payoff depends only on their own action and the actions of graph-neighbors, and equilibrium structure inherits the spectral and topological properties of the interaction graph.

Mechanism: Locality axiom $u_i = \tilde{u}_i(s_i, s_{\mathcal{N}_i})$; best-response updates depending only on neighbors; equilibrium of linear-quadratic games given by $x^* = (I - \phi A)^{-1}\alpha$, i.e. Bonacich centrality; spectral radius governs existence and uniqueness.

Cross-System Echo: Graph games are the game-theoretic analogue of **local-interaction models in statistical physics** (Ising models, spin glasses) and **factor graphs in probabilistic inference** (belief propagation, Markov random fields). In all three, a global object (joint distribution, joint strategy profile, joint energy) factorizes along a sparse graph, and global structure is controlled by spectral properties of that graph. The convergence of best-response dynamics is structurally identical to the convergence of belief propagation and Glauber dynamics — same spectral-gap bounds, same local-to-global logic.

Exercises

1. For the linear-quadratic game on a cycle C_n with uniform $\alpha_i = \alpha$ and parameter ϕ, compute the symmetric Nash equilibrium effort level as a function of ϕ and n. Compare to the complete-graph and star-graph cases.
2. Show that in a linear-quadratic graph game with $\phi > 0$, the equilibrium effort x_i^* is monotone non-decreasing in d_i holding other neighbors fixed. Is the same true in terms of Bonacich centrality across graphs of different structure?
3. A star $K_{1,n}$ has spectral radius $\rho = \sqrt{n}$. Derive the uniqueness condition $|\phi| < 1/\rho$ for the linear-quadratic game on the star and interpret the failure mode when ϕ exceeds this bound.
4. Construct a graph game on 4 nodes (two disjoint edges) where best-response dynamics converges to one of two asymmetric equilibria depending on initial condition. Identify the structural source of multiplicity.
5. Write down the linear-quadratic game on a graph with $\phi < 0$ (substitutes). Show that effort is *anti-monotone* in neighbors' effort and that equilibria can be non-symmetric even on symmetric graphs. Relate to Bramoullé–Kranton specialization (preview of 8.4).
6. A graph game has the structure: nodes are players, edges indicate "share strategy files," and the strategic action is "depth of solver study (hours)" with complementarities on study partners. Formalize this as a linear-quadratic game, specify parameters, and predict which players will study the most given degree-based graph statistics.
7. **Argue, structurally,** why the graph-game framework is strictly richer than the complete-graph normal-form game, and identify what strategic phenomena become visible in graph games that are invisible in the complete-graph case. (Hint: think about centrality, specialization, local multiplicity, and the role of graph position.)

Coordination on Networks

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Network Games **Node:** 8.2

A **coordination game** rewards matching: if my neighbor plays A, I want to play A; if they play B, I want to play B. On a two-person complete graph this is just the Stag Hunt or its cousins, with two pure equilibria (A, A) and (B, B) and the standard equilibrium-selection question. When we put the same coordination game onto a sparse graph, something qualitatively new happens: players in different regions of the graph can be **locally** coordinated on different conventions, and the question becomes whether one convention can **spread** to the other. The classical analysis is due to Stephen Morris (2000), "Contagion": on an infinite graph with a threshold coordination game, when can a single deviation trigger contagion to the whole population? The answer depends on the graph's **cohesion structure** and on the threshold at which players switch. This node develops the Morris framework, the critical threshold for contagion, best-response dynamics on coordination graphs, and the conceptual bridge to the diffusion models of 8.3.

Formal Definition

A **coordination graph game** is a graph game $(N, G, \{S_i\}, \{u_i\})$ where every player has the same binary strategy set $S_i = \{A, B\}$ and a payoff of the form

$$u_i(s_i, s_{\mathcal{N}_i}) = \sum_{j \in \mathcal{N}_i} \pi(s_i, s_j),$$

where π is a symmetric 2-by-2 coordination payoff matrix with

$$\pi(A, A) > \pi(B, A), \qquad \pi(B, B) > \pi(A, B).$$

Both diagonals dominate the off-diagonals, so coordination with each neighbor is locally optimal. The prototypical normalization:

	A	B
A	(a, a)	$(0, 0)$
B	$(0, 0)$	(b, b)

with $a, b > 0$. Player i's payoff from playing A when a fraction q_i of neighbors play A is $a \cdot d_i q_i$; from playing B is $b \cdot d_i(1 - q_i)$. Best-response condition:

$$\text{play } A \iff aq_i \geq b(1 - q_i) \iff q_i \geq q^* := \frac{b}{a + b}.$$

The threshold $q^* \in (0, 1)$ is the **critical fraction** of A-playing neighbors at which player i is indifferent.

Contagion threshold (Morris). An action A is said to be **contagious** on graph G with threshold q^* if starting from a finite set of A-playing seeds with the rest playing B, the best-response dynamics eventually spread A to the entire graph. The **contagion threshold** $q^*(G)$ is the largest q^* for which A is contagious. Morris shows that on suitably regular infinite graphs, $q^*(G)$ is a well-defined number in $(0, 1/2]$, and it is determined by a **cohesion** property of the graph.

Cohesion. A subset $S \subseteq N$ is **p-cohesive** if every node in S has at least a fraction p of its neighbors inside S. Morris's key theorem: A is contagious from a finite seed iff the complement of the seed contains no infinite $(1 - q^*)$-cohesive set. Equivalently: contagion happens iff there is no "cohesive region" large enough to resist conversion.

Intuition

The two-player story of a coordination game is equilibrium selection: both (A, A) and (B, B) are Nash, and we argue on theoretical grounds (risk dominance, payoff dominance, cheap talk, focal points) about which will be played. The network story is something else. Imagine a large population playing the coordination game, arranged on a sparse graph, each playing a best response to their neighbors. Equilibria include the all-A and all-B profiles, but also **mixed equilibria** where some regions of the graph play A and other regions play B, with the boundary between them kept stable because boundary players have enough same-action neighbors to justify their choice.

Morris asks: starting from the all-B equilibrium, if a small group of seed players switches to A, will the switch **spread**? Not if A is unattractive (high threshold). Not if the rest of the graph is tightly knit around B. The question becomes: is there any graph region that is tightly enough coordinated on B to be immune to the spreading A?

The concept of **cohesion** captures this immunity. A region is cohesive if its nodes have most of their neighbors inside the region. A cohesive B-region is self-sustaining: its border nodes see only a small fraction of outside A-neighbors and so fall below the threshold for switching. The tightness of this cohesion is exactly what q^* measures: a lower q^* (A is more attractive) requires less cohesion to resist; a higher q^* requires more.

The critical-threshold picture ties it together. Every graph has some maximum q^* below which a finite seed can propagate A to the whole graph and above which it cannot. This number is the **contagion threshold** of the graph. On the infinite line $\mathbb{Z}$, the contagion threshold is $1/2$: any action with $q^* < 1/2$ can spread from a single seed. On a 4-regular lattice, it is lower — locally cohesive regions can resist. On random graphs with high clustering, it is lower still.

The deepest insight is that **coordination on a network is a game of cohesion vs reach**. A tightly clustered community can sustain local conventions indefinitely; a well-mixed network is vulnerable to whichever convention is more attractive. This is the game-theoretic foundation of much of social-norms research, language adoption, technology standards, and political polarization.

Structural Role

8.2 sits between 8.1 (general graph games) and 8.3 (diffusion and cascades). It specializes the graph-game framework to coordination payoffs and develops the analytic tools — thresholds, cohesion, contagion bounds — that the diffusion node will generalize.

- **From 8.1:** inherits the graph-game framework and the local-best-response structure.
- **To 8.3:** the Morris contagion model is a special case of the threshold-diffusion model of Kempe–Kleinberg–Tardos, which allows heterogeneous thresholds and richer payoff structures.
- **To 8.6:** coordination games are supermodular (strategic complementarities), so the results here are concrete instances of the general Topkis theory.
- **To Module 9:** best-response dynamics on coordination graphs is a particular case of learning in games, and convergence results here feed back into learning theory.

Mathematical Development

Best-response dynamics on a coordination graph

Consider the discrete-time asynchronous dynamic: at each step, a single player wakes up, reads the current neighbor actions, and switches to the best response. Equivalently, in the synchronous version, all players simultaneously update.

Define $n_A^t(i)$ as the number of A-playing neighbors of i at time t, and $q_i^t = n_A^t(i)/d_i$. Player i's best response is

$$s_i^{t+1} = \begin{cases} A & \text{if } q_i^t \geq q^* \\ B & \text{if } q_i^t < q^* \end{cases}$$

(ties broken arbitrarily). This is a **monotone** dynamic in the following sense: if A is a **complement** (q_i going up makes A more attractive), then the update rule is monotone in the partial order where "more A's" is higher. Consequently, starting from any profile, best-response dynamics converges to an equilibrium (no cycles), and starting from the all-B profile with a seed set S, the dynamics converges to the smallest equilibrium above "A on S, B elsewhere."

Morris's contagion theorem

Theorem (Morris, 2000). Let G be an infinite, bounded-degree graph and let $q^* \in (0, 1)$. A is contagious from some finite seed on G iff there is no infinite subset $T \subseteq N$ that is $(1 - q^*)$-cohesive (every node in T has at least a $1 - q^*$ fraction of its neighbors inside T).

The contagion threshold of G is

$$q^*(G) = \sup\{q^* : \text{no infinite } (1 - q^*)\text{-cohesive set exists}\}.$$

For any finite q^* strictly below $q^*(G)$, a finite seed can convert the whole graph to A. For any q^* strictly above $q^*(G)$, some cohesive B-region always survives.

Theorem. $q^*(G) \leq 1/2$ for every graph G.

Proof sketch: at the boundary between two regions (one all-A, one all-B), the boundary nodes have roughly half their neighbors on each side if the graph is locally homogeneous. So $q^* \leq 1/2$ is necessary for contagion in almost any structured graph.

The contagion threshold on specific graphs

Infinite line $\mathbb{Z}$: $q^*(\mathbb{Z}) = 1/2$. Every node has 2 neighbors; a seed of one A-node flips a neighbor if $1/2 \geq q^*$. Since strict inequality is needed for a single seed to start a chain, the dynamic actually flips when $q^* < 1/2$.

Infinite 2D lattice $\mathbb{Z}^2$: $q^*(\mathbb{Z}^2) = 1/2$ (Morris). A half-plane of A has boundary nodes with $q = 1/2$, so strict threshold $q^* < 1/2$ spreads.

Regular tree of degree d: $q^*(T_d) = 1/d$. Each node has d neighbors; a single A-neighbor is a fraction $1/d$. So any $q^* < 1/d$ spreads.

Random graphs with clustering: contagion thresholds are generally lower than for non-clustered graphs of the same mean degree, because clusters form cohesive B-regions that resist conversion.

Risk dominance vs contagion

In a symmetric 2×2 coordination game, the **risk-dominant** equilibrium is the one with higher q^* compared to its alternative. If $\pi(A, A) = a, \pi(B, B) = b$ with $a > b$, then (A, A) is **payoff-dominant** but its threshold for switching to it is $q^* = b/(a + b) < 1/2$, while (B, B) has threshold $1 - q^* > 1/2$ — (B, B) is harder to break out of, which makes B the **risk-dominant** strategy. Morris's contagion theorem says that **the risk-dominant equilibrium is the one that spreads**: the action that spreads from a finite seed is the one with lower threshold, i.e. the risk-dominant action. This is a beautiful alignment between risk dominance (a solution concept from two-player game theory) and contagion (a property of network dynamics).

Worked Example

Contagion on an infinite line

Let $G = \mathbb{Z}$ (infinite line). Coordination game with $a = 3, b = 2$, so $q^* = 2/5 = 0.4$.

Seed set $\{0\}$ starts playing A; everyone else plays B.

At time 1, node 0 is A, all others B. Node 1 has one A-neighbor and one B-neighbor, so $q_1 = 1/2 \geq 0.4$. Best response: A. Similarly node -1 switches to A.

At time 2, nodes $\{-1, 0, 1\}$ play A. Node 2 has one A-neighbor (node 1) and one B-neighbor (node 3), $q_2 = 1/2 \geq 0.4$. Switch to A. Similarly for node -2.

The wave of A spreads one node per time step on each side. Eventually the entire line converts.

Now raise the threshold: $a = 2, b = 3$, so $q^* = 3/5 = 0.6$. Node 1 has $q_1 = 1/2 < 0.6$, so does not switch. The seed $\{0\}$ is isolated at A and everyone else stays at B. The wave does not start.

Critical threshold on $\mathbb{Z}$: $q^*(\mathbb{Z}) = 1/2$. At exactly $q^* = 1/2$, the wave may or may not spread depending on tie-breaking; strict inequality in either direction gives a clean answer.

Contagion on a cohesive cluster

Consider a graph on 20 nodes: a tight cluster of 10 nodes forming a complete subgraph K_{10}, plus 10 additional nodes forming a line attached to one node of K_{10}.

Start with A seeded on the far tip of the line. With $q^* = 0.4$, A spreads along the line (each node has 2 neighbors, converting when one is A). It eventually reaches the attachment point to K_{10}.

Now ask whether A spreads into K_{10}. Each node in K_{10} has 9 neighbors inside the cluster and (for the attachment node) 1 outside. When the outside neighbor is A, the attachment node has $q = 1/10 = 0.1 < 0.4$. It does not switch. K_{10} is a (9/10)-cohesive set, which exceeds $1 - q^* = 0.6$, so it is immune.

Lesson: a tightly cohesive cluster is a firewall against contagion even for low thresholds. A well-connected community can preserve its convention indefinitely against invasions from a sparse periphery. This is the Morris theorem in action: the cluster's cohesion 9/10 exceeds $1 - q^* = 0.6$, so contagion is blocked.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Strategies A, B	Two competing "sizing trees" — e.g., 32% / 75% / overbet vs 50% / 100% / overbet
Payoff of coordinating with a neighbor	Easier to study and discuss hands when you and your study partners use the same size grid
Neighbor set	Your study group, coaching clients, content creators you follow
Threshold q^*	How many of your peers need to use the new sizing tree before you switch
Cohesive cluster	A study group that has converged internally on a convention
Contagion event	A new solver insight spreading through the poker ecosystem
Risk dominance	The sizing tree that is easier to switch to

Concrete situation: spread of a sizing convention

Two years ago, most high-stakes NLHE regulars used a sizing tree with 75% pot as the default c-bet. A small group of solver-heavy regulars started using 33% pot (on certain board textures) because their solvers showed it was marginally higher EV in spots with low equity swings. This was a new "convention" — call it action A. Old convention is action B.

The payoff structure is a coordination game. If your study group uses 33%, you want to use 33% too — otherwise your hand discussions are apples-to-oranges, you cannot compare lines, and your solver runs do not map to the shared vocabulary. Within a tight study group, there are strong complementarities. Across the ecosystem, the graph is very sparse: you only coordinate with people you actually talk to.

Morris asks: does the 33% convention spread from the initial seeds to the whole ecosystem? The answer depends on two things: the **threshold** (how attractive is 33%?) and the **cohesion** of the existing 75% community. In the actual history, 33% did spread through the high-stakes online ecosystem (well-connected, low clustering, so low cohesion of any single subcommunity), but did not fully reach live cash games (tightly clustered local communities, high cohesion, resistance to change).

Concrete situation: Discord server language

A poker Discord server develops local conventions: certain abbreviations (SRP, 3BP, OOP-COOP), certain default sizes to reference in discussion, certain tags for hand history posts. These are pure coordination conventions — no inherent reason to prefer one over another, but if everyone uses them, communication is easier. New members face a coordination decision: do I use the server's conventions or my personal ones?

The game is locally cohesive: your "neighborhood" is the active posters, and the fraction of them using the server convention is near 1. The threshold for switching is low (small gain from matching), but the cohesion is overwhelming — you match. A new convention can only spread through the server if a cohesive subgroup adopts it first (the power-posters, the admins, the solver-study sub-channel) and then propagates outward.

What this understanding buys you

Convention change is a contagion problem, not a correctness problem. When you want to change how a group of players talks about or plays something, the question is not "is my idea correct?" but "what is the cohesion of the group's current convention, and does my idea beat the threshold?" You can be completely right about a new sizing tree or a new concept, but if the existing convention is cohesive enough, your idea will not spread — it will be trapped in a seed set and isolated. The intervention is to attack the cohesion: find and convert the **most-connected node** (a high-Bonacich-centrality regular or content creator), because their conversion breaks the cohesion of their local cluster and lowers q^* for their neighbors.

Your study group's language is a local convention. If you have been coordinating with a small group for a year, your shared vocabulary is a highly cohesive q^*-equilibrium. Joining a bigger community with a different convention feels like a cost — that is the coordination cost of switching, not a sign that the new community is wrong. And conversely, if you want your group to update to new ideas, you need to import a critical mass of converted nodes, not just make the argument once.

Online and live are different graphs. The online regular community is dense and low-clustering — contagion of new ideas is fast. The live cash game community is sparse and highly clustered — conventions persist for decades. This is not about intelligence or modernity; it is about the topology of the interaction graph.

Cross-Links

- **2.1 Nash equilibrium** — coordination games have multiple Nash; networks resolve selection.
- **8.1 Graph games** — coordination is the coordination-payoff specialization.
- **8.3 Diffusion and cascades** — generalizes Morris's contagion to richer threshold models.
- **8.6 Strategic complementarities** — coordination games are the canonical supermodular example.
- **4.6 Global games** — equilibrium-selection refinement that often picks the risk-dominant action, matching Morris's contagion result.

- **7.4 Evolutionary stability** — alternative selection concept giving the same risk-dominant answer in many cases.
- **Network science — epidemic models** — mathematical parallels to threshold contagion.

Structural Compression

Invariant: On a network coordination game with threshold q^*, an action is contagious from a finite seed iff no subset of the remaining nodes is sufficiently cohesive to resist conversion; the risk-dominant action is the one that spreads.

Mechanism: Binary best-response dynamics with threshold rule $s_i = A \iff q_i \geq q^*$; monotone dynamics converging to equilibrium; contagion blocked by $(1 - q^*)$-cohesive subsets.

Cross-System Echo: Morris contagion is the game-theoretic analogue of **bootstrap percolation** in statistical physics — a cellular automaton where cells activate when a fraction of neighbors are active. The same phase transitions, the same cohesion-based obstructions, the same lattice-vs-tree asymmetries appear in both. It is also parallel to the **voter model** of population genetics and to the **SIR threshold models** of epidemiology. The common structural feature: a local update rule depending on a fraction of neighbors in a particular state, propagating a wave whose reach is controlled by the graph's cohesion structure.

Exercises

1. Derive the threshold $q^* = b/(a + b)$ for the symmetric 2×2 coordination game with on-diagonal payoffs a and b and off-diagonal 0. Show that the risk-dominant action is the one with the smaller threshold.

2. Compute the contagion threshold of the infinite 2D lattice $\mathbb{Z}^2$ under 4-neighbor connectivity. Show that $q^*(\mathbb{Z}^2) = 1/2$ and identify what happens at exactly $q^* = 1/2$.

3. Show that on a regular tree T_d with $d \geq 3$, the contagion threshold is $1/d$. Compute the threshold for $d = 3, 4, 5$ and interpret why higher-degree trees are harder to convert.

4. Construct a finite graph with two cohesive communities of sizes 10 and 10 connected by a single edge. For a coordination game with $q^* = 0.3$, determine whether a single seed in one community can propagate to the other. What about $q^* = 0.1$?

5. Show that in a best-response dynamic on a coordination graph game, no cycles can occur (monotone convergence). Use this to bound the convergence time on a graph of size n.

6. Formalize the following as a network coordination game and compute its contagion threshold: a population of 100 poker regulars are each in one of two Discord servers (50 and 50), with complete graphs within each server and 5 bridging edges between. A new sizing convention is introduced in server 1. What is the minimum q^* for which it spreads to server 2?

7. **Argue, structurally,** why contagion on a network identifies the risk-dominant action of a coordination game rather than the payoff-dominant action. Interpret this as a game-theoretic "mechanism design" observation: networks select on risk, not on payoff.

Diffusion and Cascades

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Network Games **Node:** 8.3

Node 8.2 studied coordination contagion: a single binary action spreading on a graph through best-response dynamics. This node generalizes that framework in two directions. First, the **Kempe–Kleinberg–Tardos (KKT) model** (2003, 2005) treats diffusion as an algorithmic problem: given a graph, a threshold model, and a budget of k seed nodes, which seeds maximize the eventual spread? This is the **influence maximization** problem, and it turns out to be NP-hard but admits a greedy $(1-1/e)$-approximation because the influence function is submodular. Second, the **Bass diffusion model** (1969), originally from marketing, describes the aggregate dynamics of adoption in a population where agents adopt based on both external advertising (innovation effect) and peer influence (imitation effect). When we give the Bass model a graph structure and strategic agents, it becomes a network diffusion game. This node develops the threshold model, the KKT submodularity result, the connection to the Bass model, and the strategic interpretation of all three.

Formal Definition

A **threshold diffusion model** on a graph $G = (N, E)$ is specified by:

- A threshold $\theta_i \in [0,1]$ for each node $i \in N$, representing the fraction of i's neighbors that must have adopted before i adopts.
- An initial **seed set** $S_0 \subseteq N$ of nodes that adopt at time 0.
- A deterministic update rule: at each time t, node i adopts if the fraction of its adopted neighbors is at least θ_i:
$$\sigma_i^{t+1} = \begin{cases} 1 & \text{if } \frac{|\{j \in \mathcal{N}_i : \sigma_j^t = 1\}|}{d_i} \geq \theta_i \text{ or } \sigma_i^t = 1 \\ 0 & \text{otherwise.} \end{cases}$$
  Adoption is **irreversible**: once a node adopts, it stays adopted.

The **influence function** $\sigma : 2^N \to \mathbb{Z}_{\geq 0}$ maps a seed set S_0 to the total number of nodes eventually adopted. Formally,
$$\sigma(S_0) = \lim_{t \to \infty} |\{i \in N : \sigma_i^t = 1\}|.$$
The limit is well-defined because the set of adopters is monotonically non-decreasing.

The **influence maximization problem** is: given G, thresholds $\{\theta_i\}$, and a budget k, find
$$S^* = \arg \max_{S_0 \subseteq N, |S_0| = k} \sigma(S_0).$$

Independent cascade model. An alternative stochastic model: when node i adopts, it independently "activates" each non-adopted neighbor j with probability p_{ij}. The influence function $\sigma(S_0) = \mathbb{E}[|\text{eventual adopters}|]$ is now an expected value.

Bass model (aggregate). In a well-mixed population of size N, let $F(t)$ be the fraction adopted by time t. The Bass equation:
$$\frac{dF}{dt} = (1 - F(t))(p + qF(t)),$$
where p is the **innovation coefficient** (external influence) and q is the **imitation coefficient** (peer influence). The solution is an S-curve:
$$F(t) = \frac{1 - e^{-(p+q)t}}{1 + (q/p)e^{-(p+q)t}}.$$
The Bass model on a graph replaces $F(t)$ with the local adoption fraction in each node's neighborhood, coupling the differential equation to the graph structure.

Intuition

The core question of diffusion is: **how far does an initial seed spread, and what determines the reach?** Morris (8.2) answered this for homogeneous-threshold coordination games: contagion happens iff no cohesive region blocks it. The KKT framework generalizes by allowing heterogeneous thresholds, irreversible adoption, and the algorithmic question of *choosing* the seeds.

Irreversibility is the key new assumption. In a coordination game, a player who switched to A can switch back to B if their neighbors switch back. In a diffusion model, adoption is one-way: once you buy the product, join the platform, or learn the concept, you do not un-adopt. This makes the dynamic monotone and the influence function well-defined, but it also means the process is fundamentally different from equilibrium coordination — it is a **one-shot cascade**, not a repeated best-response.

The KKT result that influence is submodular formalizes a powerful intuition: **the marginal value of adding a seed decreases as the seed set grows**. A seed in an empty graph spreads to all reachable nodes; a seed added to a large existing seed set only reaches nodes not already in the cascade. This diminishing return is the structural reason why greedy seed selection works: always pick the seed with the largest marginal influence, and you get a provably near-optimal solution.

The Bass model adds a second dimension: the distinction between **external influence** (advertising, discovery, innovation) and **peer influence** (word of mouth, imitation, social learning). In the pure threshold model, all influence is from peers. In the Bass model, there is an exogenous "push" (parameter p) that causes adoption even without peer pressure. The interplay between p and q determines the shape of the adoption curve: high p gives exponential early adoption (technology-push); high q gives S-curve adoption (word-of-mouth). On a network, high-degree nodes amplify the q-term because they see more peers — which is why hubs adopt early and amplify cascades.

The strategic layer: when agents can choose their threshold (e.g., how much evidence to require before adopting), diffusion becomes a game. Each agent's threshold is a strategic variable that balances the risk of early adoption (if the innovation is bad) against the cost of late adoption (if the innovation is good and you miss the network effects). The Nash equilibria of this game determine the cascade structure, and the social optimum may involve subsidizing early adopters to reduce their effective threshold.

Structural Role

8.3 generalizes 8.2's coordination-contagion to a richer class of diffusion models and connects to the algorithmic and marketing literatures.

- **From 8.2:** the homogeneous-threshold coordination model is the special case where $\theta_i = q^*$ for all i and adoption is reversible. The irreversible version is the limit of infinite patience (once you switch, you never switch back).
- **To 8.5:** influence maximization is a design problem on networks, parallel to the mechanism-design perspective of network formation.
- **To Module 9:** learning dynamics on networks can be viewed as stochastic diffusion processes, and the convergence results of Module 9 use the same submodularity and monotonicity tools.
- **To Module 6 (mechanism design):** the influence-maximization problem is a mechanism-design problem — choosing seed sets to maximize a social objective under strategic constraints.

Mathematical Development

Submodularity of the influence function

Theorem (Kempe, Kleinberg, Tardos, 2003). Under both the threshold model (with random thresholds drawn uniformly from $[0, 1]$) and the independent cascade model, the influence function $\sigma(S)$ is **monotone**

and **submodular**:

$$\sigma(S \cup \{v\}) - \sigma(S) \geq \sigma(T \cup \{v\}) - \sigma(T) \quad \text{for all } S \subseteq T \subseteq N, v \notin T.$$

Proof sketch (threshold model). Fix a realization of thresholds $\theta_i \sim U[0, 1]$. For a given realization, the set of eventually-adopted nodes is a deterministic function of S_0. Adding a node v to a smaller seed set S can start cascades that have not yet been triggered; adding v to a larger set $T \supseteq S$ can only trigger cascades that were not already triggered by $T \setminus S$, which is a subset. So the marginal gain is non-increasing. Taking expectations over threshold realizations preserves submodularity because non-negative linear combinations of submodular functions are submodular.

Corollary. The greedy algorithm — repeatedly adding the node with largest marginal influence — achieves a $(1 - 1/e)$-approximation to the optimal seed set:

$$\sigma(S_{\text{greedy}}) \geq (1 - 1/e) \cdot \sigma(S^*).$$

This is a direct application of the classic result that greedy maximization of a monotone submodular function subject to a cardinality constraint achieves a $(1 - 1/e)$ factor.

NP-hardness of influence maximization

Theorem (KKT, 2003). Computing the influence-maximizing seed set of size k is NP-hard under both the threshold and independent cascade models, even on simple graph families.

The reduction is from set cover. The practical consequence is that exact optimization is intractable for large networks, but the greedy $(1 - 1/e)$ bound is tight up to lower-order terms under standard complexity assumptions.

Cascade size distribution

In the independent cascade model on an Erdos–Renyi random graph $G(n, p)$, the cascade size exhibits a **phase transition** at edge probability $p = 1/n$:

- For $p < 1/n$: cascades are small (logarithmic in n) with high probability.
- For $p > 1/n$: a single seed can trigger a cascade of size $\Theta(n)$ (giant component).
- At $p = 1/n$: critical regime, cascade sizes follow a power-law distribution $\Pr[|\text{cascade}| = s] \sim s^{-3/2}$.

This phase transition mirrors the percolation threshold in random graphs and the epidemic threshold in SIR models. The structural message: **whether a cascade goes global or stays local depends on the graph's connectivity relative to a critical threshold**.

Bass model on networks

On a graph G, the Bass model becomes a system of coupled ODEs. Let $x_i(t) \in \{0, 1\}$ be node i's adoption status. In continuous time with stochastic adoption:

$$\Pr[\text{node } i \text{ adopts at time } t \mid x_i(t) = 0] = p + q \cdot \frac{\sum_{j \in \mathcal{N}_i} x_j(t)}{d_i}.$$

In the mean-field limit (complete graph), this reduces to the standard Bass ODE. On a sparse graph, the coupling to the adjacency matrix produces heterogeneous adoption times: high-degree nodes adopt earlier (they see more adopted neighbors), and the cascade front propagates along high-connectivity paths.

Strategic Bass model. If agents can choose their threshold (equivalently, their q parameter) at a cost, the game has the structure: each agent chooses a "sensitivity to peers" q_i, and adoption spreads as a function of the joint $(q_1, \ldots, q_n)$. In equilibrium, agents with high degree choose lower q_i (they get enough peer signal) while peripheral agents choose higher q_i (they need to be more responsive to their few connections). This endogenous heterogeneity in thresholds is a prediction of the strategic model that the exogenous-threshold model misses.

Worked Example

Influence maximization on a small network

Consider 8 nodes arranged as two triangles $\{1, 2, 3\}$ and $\{6, 7, 8\}$, connected by a path $3-4-5-6$. Thresholds: $\theta_i = 1/2$ for all nodes. Budget: $k = 2$ seeds.

Option A: seed $\{1, 6\}$ (one in each triangle).

Time 0: nodes 1, 6 adopt. In triangle 1, node 2 has neighbors $\{1, 3\}$, one adopted, fraction $1/2 \geq 1/2$. Node 2 adopts. Similarly node 3 has $\{1, 2\}$, one adopted, fraction $1/2$. Adopts. Triangle 1 fully converts by time 1. Similarly triangle 2. Node 4 has neighbors $\{3, 5\}$, one adopted (node 3), fraction $1/2$. Adopts at time 2. Node 5 has neighbors $\{4, 6\}$, both adopted. Adopts at time 2. Total: 8 nodes.

Option B: seed $\{1, 2\}$ (both in same triangle).

Time 0: nodes 1, 2 adopt. Node 3 has $\{1, 2\}$, both adopted, fraction 1. Adopts. Node 4 has $\{3, 5\}$, one adopted, fraction $1/2$. Adopts at time 2. Node 5 has $\{4, 6\}$, one adopted. Adopts at time 3. Node 6 has $\{5, 7, 8\}$, one adopted, fraction $1/3 < 1/2$. Does **not** adopt. Cascade stalls at node 5. Total: 6 nodes.

Analysis. Option A (diverse seeds) reaches all 8 nodes; option B (clustered seeds) reaches only 6. The greedy algorithm would compute marginal influence of each node as a first seed, then choose the second seed conditional on the first. The first seed in the triangle maximizes $\sigma(\{i\}) = 3$ (triggers whole triangle). The second seed's marginal is maximized by seeding the other triangle (marginal gain 5, reaching the bridge nodes) rather than the same triangle (marginal gain 0).

This illustrates the general principle: **greedy seeding diversifies across clusters** because marginal influence is highest where the cascade has not yet reached. Submodularity formalizes this: the redundancy of placing two seeds in the same cluster is exactly the diminishing-returns property.

Bass model calibration

For the poker ecosystem, consider $p = 0.01$ (1% chance per month of discovering a new concept independently) and $q = 0.3$ (30% peer-influence coefficient). The Bass S-curve gives:

$$F(t) = \frac{1 - e^{-0.31t}}{1 + 30e^{-0.31t}}.$$

Adoption reaches 50% at $t^* = \ln(q/p)/(p + q) = \ln(30)/0.31 \approx 10.97$ months. The peak adoption rate occurs near t^* and equals $(p + q)^2/(4q) = 0.31^2/1.2 \approx 0.08$ per month.

Interpretation. A new solver insight with these parameters takes about 11 months to reach half the player population. The imitation coefficient q dominates: most adoption is driven by peer influence, not independent discovery. Increasing q (denser network, better communication tools, more coaching content) accelerates adoption; increasing p (better discovery tools, more solver access) shifts the curve left but does not change the slope.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Adoption	A player learns and implements a new strategic concept (e.g., probe-bet strategy)
Threshold θ_i	How many of your study partners need to validate a concept before you implement it

Formal object	Poker instantiation
Seed set	The first players to discover and use a new idea (often solver-grinders, theory coaches)
Influence function $\sigma(S)$	Total number of players who eventually learn the concept, given initial seeds S
Innovation coefficient p	Rate of independent discovery (reading papers, running solvers)
Imitation coefficient q	Rate of peer adoption (coaching, Discord discussions, hand reviews)
Cascade	A new concept spreading from a small group of early adopters through the community

Concrete situation: the probe-bet cascade

Around 2018, a small set of high-stakes online players began using "probe bets" — small out-of-position bets on the turn when the in-position player checked back the flop. Solvers had shown this was high-EV in many configurations, but the concept was not widely known. The initial "seed set" was perhaps 10–15 players in two or three interconnected study groups.

The cascade unfolded as follows. Seed players began using probe bets in their games. Their regular opponents noticed the line appearing in hand histories and database filters. Those opponents discussed the line with their own study groups. Content creators picked it up and published videos. Training sites incorporated it into curricula. Within 18 months, probe-betting was standard at mid-stakes and above.

The diffusion followed a threshold model: a player adopted the concept when enough of their information neighbors (study partners, coaching clients, content they consumed) validated it. The cascade was amplified by **high-degree nodes** — content creators with thousands of followers — who acted as super-spreaders. The KKT insight applies: if you wanted to maximize the spread of a new concept, you would seed the highest-influence nodes (the content creators), not random nodes.

Concrete situation: false cascade

Not every cascade is beneficial. In 2019, a widely shared video claimed that "donk-betting is always bad." This was an oversimplification of a nuanced solver result, but it spread rapidly because the seed node (a popular content creator) had high degree and the threshold for adopting a simple heuristic is low. The "donk-betting is bad" meme cascaded through the community and persisted long after more careful analyses showed that donk-betting is correct in many configurations.

This is a **false cascade** — a strategically suboptimal convention that spreads because the diffusion dynamics do not check correctness, only threshold adoption. The KKT framework describes the mechanics of spread; it does not guarantee that what spreads is correct. In poker, false cascades are costly: players who adopt a wrong heuristic lose EV until the cascade is corrected by a counter-cascade from more careful analysis.

What this understanding buys you

You are a node in a diffusion graph. Every concept you have adopted was seeded somewhere and reached you through a cascade path. Understanding diffusion dynamics lets you evaluate the *source quality* of the cascade: did this concept reach you through a high-quality path (solver-validated, peer-reviewed by strong players) or through a low-quality path (viral content, oversimplified heuristics)? The threshold model says you adopted when enough neighbors adopted — but not all neighbors are equal. Being selective about your information neighborhood (your graph edges) is the single most important thing you can do to improve the quality of the cascades that reach you.

Seed selection is a design problem. If you are a coach or content creator, you are a high-degree node with outsized influence on cascades. The greedy algorithm says: target your content at the nodes with highest marginal influence — the mid-level coaches and study-group leaders who will amplify your message to their own communities. Seeding random low-degree nodes is inefficient; seeding the already-converted hub nodes is redundant. The optimal strategy is to find the "bridge" nodes between communities and convert them.

The Bass p/q ratio tells you how ideas spread in your ecosystem. If $q \gg p$, most learning is through peers — invest in your study group. If p is high (you have direct solver access, you read papers), you can be a seed node yourself. Knowing your own p/q ratio tells you whether you are an innovator or an imitator, and the rational response differs: innovators should invest in discovery; imitators should invest in network position.

Cross-Links

- **8.2 Coordination on networks** — Morris contagion is the reversible, homogeneous-threshold special case.
- **8.1 Graph games** — diffusion as a graph game with irreversible dynamics.
- **8.5 Network formation games** — seed selection as a design problem parallels network design.
- **6.1 Mechanism design** — influence maximization as an optimization problem with strategic agents.
- **9.1 Fictitious play** — learning dynamics that produce cascade-like adoption patterns.
- **7.2 Replicator dynamics** — evolutionary spread of strategies in populations, parallel to diffusion on graphs.
- **Marketing and operations research** — Bass model, new-product diffusion, viral marketing.

Structural Compression

Invariant: In a threshold diffusion model on a graph, the influence function is monotone and submodular, cascade reach is determined by network connectivity relative to adoption thresholds, and the greedy algorithm for seed selection achieves a $(1 - 1/e)$-approximation to optimal influence maximization.

Mechanism: Irreversible adoption via local threshold rule; monotone cascade front propagating along edges; marginal influence of additional seeds diminishes (submodularity); phase transition in cascade size at the graph's connectivity threshold.

Cross-System Echo: Diffusion on networks is the game-theoretic version of **epidemic spreading** in mathematical biology (SIR, SIS models), **rumor spreading** in distributed computing, and **percolation** in statistical physics. In all four, a local activation rule propagates through a graph, the reach depends on a threshold-vs-connectivity comparison, and there is a phase transition between local and global cascades. The submodularity of influence mirrors the submodularity of **sensor placement** in monitoring networks and **feature selection** in machine learning — all are instances of diminishing returns under set-function maximization.

Exercises

1. Prove that the influence function $\sigma(S)$ under the threshold model with thresholds drawn uniformly from $[0, 1]$ is equivalent (in distribution) to the influence function under the independent cascade model with appropriate edge probabilities. (This is KKT's equivalence theorem.)
2. Compute the greedy seed set of size 2 for a star graph $K_{1,5}$ with uniform thresholds $\theta = 1/2$. Compare the greedy solution to the optimal solution and verify the $(1 - 1/e)$ bound.
3. For a complete graph K_n with uniform threshold θ, derive the minimum seed set size needed to trigger a full cascade. Express this as a function of n and θ.

4. Consider the Bass model with parameters $p = 0.03, q = 0.4$. Compute the time to 50% adoption, the peak adoption rate, and the time to 90% adoption. Interpret the relative contributions of innovation and imitation.

5. Construct a graph on 12 nodes (two cliques of 5 connected by a bridge of 2) with heterogeneous thresholds such that a single seed in clique 1 triggers a full cascade, but no single seed in clique 2 does. Identify the structural asymmetry.

6. In the independent cascade model, show that the cascade size on $G(n, p)$ undergoes a phase transition at $p = 1/n$. Derive the expected cascade size just above and just below the threshold.

7. **Argue, structurally,** why submodularity of the influence function implies that greedy seed selection diversifies across network communities rather than concentrating seeds in a single cluster. Connect this to the intuition that "redundant seeds waste budget."

Best-Shot and Public Goods on Graphs

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Network Games **Node:** 8.4

In the linear-quadratic graph games of 8.1, the interaction parameter ϕ could be positive (complements) or negative (substitutes). This node specializes to the case of **strategic substitutes** — games where a neighbor's effort *reduces* my incentive to exert effort — and the canonical application: **local public goods**. The foundational model is Bramoullé and Kranton (2007): each player on a graph provides a costly effort that benefits themselves and their neighbors; because neighbors' efforts substitute for own effort, equilibrium produces **specialization** — some players provide high effort while their neighbors free-ride. The structure of the interaction graph determines *who* specializes and *who* free-rides, and the equilibrium set is characterized by the **independent sets** and **maximal independent sets** of the graph. This is the game-theoretic foundation of free-riding on networks, division of labor, and the economics of local public goods.

Formal Definition

A **public goods game on a graph** is a graph game $(N, G, \{S_i\}, \{u_i\})$ where each player i chooses an effort level $x_i \in [0, \bar{x}]$ and has payoff

$$u_i(x_i, x_{\mathcal{N}_i}) = f\left(x_i + \sum_{j \in \mathcal{N}_i} x_j\right) - c(x_i),$$

where f is a concave, increasing benefit function and c is a convex, increasing cost function. The benefit depends on the **total effort in i's closed neighborhood** $\bar{\mathcal{N}}_i = \{i\} \cup \mathcal{N}_i$.

Best-shot game. The extreme case of perfect substitutes: $f(y) = y$, but with $c(x_i)$ convex enough that it is optimal to have exactly one provider per neighborhood. In the pure best-shot model:

$$u_i = \max\left(x_i, \max_{j \in \mathcal{N}_i} x_j\right) - c(x_i).$$

Your benefit is the maximum effort in your neighborhood, not the sum. Only the best contributor matters.

Linear-quadratic substitutes. The tractable benchmark from Bramoullé–Kranton:

$$u_i = \alpha\left(x_i + \sum_{j \in \mathcal{N}_i} x_j\right) - \frac{1}{2}\left(x_i + \sum_{j \in \mathcal{N}_i} x_j\right)^2 - \frac{\gamma}{2} x_i^2,$$

where $\gamma > 0$ is a private cost parameter. With the simpler specification:

$$u_i = x_i + \sum_{j \in \mathcal{N}_i} x_j - \frac{1}{2}\left(x_i + \sum_{j \in \mathcal{N}_i} x_j\right)^2,$$

the first-order condition gives the best response

$$x_i^* = \max\left(0,\ 1 - \sum_{j \in \mathcal{N}_i} x_j^*\right).$$

This says: **provide effort up to the point where the total neighborhood effort (including your own) reaches 1, but never go negative**. If your neighbors already provide enough, you provide zero — you free-ride.

Nash equilibrium characterization (Bramoullé–Kranton). In the binary effort version ($x_i \in \{0, 1\}$), a Nash equilibrium is a profile where every player with $x_i = 1$ has no neighbor with $x_j = 1$, and every player with $x_i = 0$ has at least one neighbor with $x_j = 1$. The set of providers $P = \{i : x_i = 1\}$ is a **maximal independent set** of G: independent (no two providers are adjacent) and maximal (every non-provider is adjacent to a provider). Conversely, every maximal independent set of G corresponds to a Nash equilibrium.

Intuition

The Bramoullé–Kranton model captures a simple but powerful idea: **effort on a network is a local public good, and local public goods produce specialization**. When my neighbor works hard, I do not need to — I benefit from their work for free. This is strategic substitutability: my neighbor's high effort reduces my marginal benefit of effort, so my best response is to lower my own. Conversely, if my neighbor shirks, I must step up. The equilibrium is a division of labor: some nodes are "providers" and their neighbors are "free-riders."

The graph structure determines which equilibria exist. On a star, the center can be the sole provider (all peripherals free-ride) or a peripheral can be a provider (and the center and others free-ride, but this is less stable). On a cycle, alternating nodes provide — the equilibrium is a coloring. On a random graph, the set of providers is an independent set, and the number of equilibria is the number of maximal independent sets, which can be exponential.

The key structural insight: **the equilibrium correspondence between Nash equilibria and maximal independent sets means that the combinatorial structure of independent sets in graphs governs the economics of local public goods**. Graphs with few maximal independent sets have few equilibria (tight predictions); graphs with many have many equilibria (coordination is hard). Trees have polynomial numbers; dense random graphs have exponential numbers.

Free-riding is not pathological in this model — it is *efficient* in the sense that total effort is minimized for a given level of coverage. The inefficiency comes from a different source: the wrong set of players may specialize. A low-cost provider might free-ride while a high-cost neighbor provides, because the graph topology forced the wrong assignment. This misallocation is the welfare cost of decentralized provision on networks.

Structural Role

8.4 is the strategic-substitutes counterpart to 8.2's coordination (strategic complements). Together they span the two fundamental interaction types in network games.

- **From 8.1:** inherits the graph-game framework; specializes to substitutes ($\phi < 0$).
- **From 8.2/8.6:** contrast with coordination and complementarities — complements produce conformity; substitutes produce specialization.
- **To 8.5:** the efficiency of public-goods equilibria motivates the network-formation question: which graphs produce efficient public-goods outcomes?
- **To Module 6:** mechanism design for public goods on networks — how to design transfers that correct the misallocation from decentralized provision.

Mathematical Development

Best-response characterization

In the linear-quadratic model, write $y_i = \sum_{j \in \mathcal{N}_i} x_j$ for the total neighbor effort. The best response is

$$x_i^*(y_i) = \max(0, 1 - y_i).$$

This is a piece-wise linear, decreasing function of y_i. An interior equilibrium satisfies $x_i = 1 - y_i$ for all active providers (those with $x_i > 0$) and $x_i = 0$ with $y_i \geq 1$ for free-riders.

Equilibrium structure on specific graphs

Complete graph K_n. Any single-provider profile is a Nash equilibrium: one player provides $x = 1$, all others provide 0, everyone gets benefit 1. There are n such equilibria. There is also a symmetric interior equilibrium $x_i = 1/(n)$ for all i, where each player provides a small amount. The single-provider equilibria are maximal independent sets of size 1 (every node is adjacent to every other).

Star $K_{1,n}$. Two types of equilibria: (a) center provides $x_c = 1$, all peripherals provide 0. (b) Any single peripheral p_j provides $x_{p_j} = 1$, center provides 0, all other peripherals provide 0. Type (a) is the "center specializes" equilibrium; type (b) is the "peripheral specializes" equilibrium, but the center still gets the good because p_j is the center's neighbor.

Cycle C_n. For even n: alternating providers (every other node provides 1, rest provide 0). Two such equilibria (offset by one position). For odd n: no perfect alternating equilibrium exists; instead, one "gap" must exist where two adjacent nodes both provide 0, and a third node covers both.

General theorem. The number of Nash equilibria equals the number of maximal independent sets of G, which is bounded above by $3^{n/3}$ (Moon–Moser bound) and can be exponential in n for general graphs.

Welfare analysis

Social optimum. A social planner would choose efforts to maximize total welfare $\sum_i u_i$. In the binary model, the planner wants to find the **minimum dominating set** — the smallest set of providers such that every node is either a provider or adjacent to one. This is NP-hard in general but polynomial on trees and some special graph classes.

Efficiency gap. Nash equilibria (maximal independent sets) are generally **larger** than the minimum dominating set — too many players provide, not too few. The inefficiency of decentralized provision is over-provision by poorly positioned players, not under-provision.

Theorem (Bramoullé–Kranton). On any graph, the Nash equilibria of the public-goods game are Pareto-ranked: the equilibrium with the smallest provider set is Pareto-best, and the one with the largest provider set is Pareto-worst. The efficient equilibrium corresponds to the smallest maximal independent set.

Heterogeneous costs

When players have different costs $c_i(x_i) = \gamma_i x_i^2/2$, the best response becomes $x_i^* = \max(0, (1 - y_i)/(1 + \gamma_i))$. High-cost players are more likely to free-ride; low-cost players are more likely to provide. On a bipartite graph, the efficient equilibrium assigns provision to the low-cost side. On a general graph, the assignment of providers to independent-set positions depends on the interplay between cost structure and graph topology.

Worked Example

Public goods on a path graph

Consider a path graph P_5 with 5 nodes: $1 - 2 - 3 - 4 - 5$. Binary effort, threshold model.

Maximal independent sets: $\{1, 3, 5\}$, $\{1, 4\}$, $\{2, 4\}$, $\{2, 5\}$. Each is a Nash equilibrium.

Equilibrium $\{1, 3, 5\}$: three providers, two free-riders. Total effort: 3. Every node is covered (within distance 1 of a provider).

Equilibrium $\{2, 4\}$: two providers, three free-riders. Total effort: 2. Nodes 1, 3, 5 are all adjacent to a provider. This is the minimum dominating set and the Pareto-best equilibrium.

Equilibrium $\{1,4\}$: two providers. Node 1 covers nodes 1 and 2. Node 4 covers 3, 4, 5. Total effort: 2. Also minimum.

Welfare comparison. In the $\{2,4\}$ equilibrium, total cost is $2c$. In the $\{1,3,5\}$ equilibrium, total cost is $3c$. Both achieve full coverage, but the smaller provider set wastes less effort. Decentralized dynamics can get stuck at the larger equilibrium.

Decentralized dynamics. Start with everyone providing ($x_i = 1$ for all i). Under best-response updates: node 2 sees neighbors 1 and 3 providing, so free-rides ($x_2 = 0$). Node 4 sees neighbors 3 and 5 providing, so free-rides. Now $\{1,3,5\}$ provide, $\{2,4\}$ free-ride. This is an equilibrium. The dynamics converged to the Pareto-worst equilibrium because they started from a state with too much provision and shaved off the "unnecessary" providers, but not enough of them.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Effort x_i	Depth of solver work on a particular spot or board texture
Neighbor benefit	Your study partners benefit from your solver output without running it themselves
Free-riding	Relying on study partners' solver work instead of doing your own
Provider (independent set)	The study-group member who actually runs the sims and shares the results
Maximal independent set	The smallest subset of a group such that everyone is adjacent to a provider
Strategic substitutability	"If you already ran this sim, I do not need to"

Concrete situation: solver work in a study group

A 5-person study group works on a shared topic — say, turn play in 3-bet pots out of position. Each person could run their own solver trees, but solver work is expensive (time, compute, mental effort). The output is a local public good: if Alex runs the sim, Alex's study partners (neighbors in the group graph) can read Alex's output and get 80% of the benefit without doing the work.

The equilibrium is specialization: one or two members become the "solver grinders" while the rest free-ride on their output. This is not laziness — it is the Nash equilibrium of a public-goods game with strategic substitutes on the group's interaction graph. The free-riders contribute in other ways (e.g., providing hand histories, running live-play experiments, reviewing results), but the core solver work is provided by a maximal independent set of the group.

The graph structure matters. If the group is fully connected (K_5), a single provider suffices. If the group has subgroups (a 3-person cluster and a 2-person cluster connected by one bridge), the minimum dominating set needs one provider per cluster. The worst outcome is everyone providing — redundant effort. The second-worst is the wrong person providing (a player who is slow at solver work but well-connected).

Concrete situation: coaching ecosystem

In a coaching stable, the head coach produces content (solver analyses, conceptual frameworks, hand reviews). Coaching clients are nodes in a graph where the coach is the hub. The coach's effort is a public good for all clients. If the coach provides, no client needs to independently derive the material — they free-ride. If the coach stops providing, clients must step up or find another provider.

The pairwise stability question from 8.5 connects here: is this coaching graph stable? It is, as long as the coach's effort exceeds the clients' threshold for benefit and the coaching fee is below the clients' willingness to pay for the public good. If another coach opens a competing stable, clients at the boundary face a public-goods decision: which coach's output to free-ride on.

What this understanding buys you

Free-riding in your study group is structurally predicted, not a character flaw. If you find that one person does all the solver work and everyone else reads the results, that is the Nash equilibrium of a substitutes game on your group graph. To change the equilibrium, change the graph: split the group into smaller subgroups where each person must be a provider for at least one topic. Or change the cost structure: make solver work cheaper (shared compute, better tools) so that the marginal cost of providing drops below the marginal benefit of not relying on a single provider.

Your position in the study graph determines your role. If you are the most-connected node (the hub), you will end up as the provider — you have the most neighbors to cover. If you are peripheral, you will free-ride. This is not a choice you make — it is the structural consequence of your graph position. Understanding this lets you predict your role and, if you want to change it, restructure your connections.

Over-provision is the inefficiency, not under-provision. In poker study groups, the typical pathology is not "nobody does the work" — it is "everyone redundantly does the same work." Three people in a group independently run the same solver tree because they do not trust each other's output or do not have efficient sharing mechanisms. This is the $\{1, 3, 5\}$ equilibrium on the path graph — Pareto-dominated by $\{2, 4\}$. The fix is better sharing infrastructure, not more individual effort.

Cross-Links

- **8.1 Graph games** — public goods as the substitutes specialization of linear-quadratic graph games.
- **8.2 Coordination on networks** — contrasting case: complements produce conformity, substitutes produce specialization.
- **8.5 Network formation** — which graphs arise endogenously when agents can form links to access public goods.
- **6.1 Mechanism design** — how to design transfer schemes for efficient public-goods provision on networks.
- **2.1 Nash equilibrium** — maximal-independent-set characterization of equilibria.
- **Graph theory** — independent sets, dominating sets, chromatic number as game-theoretic solution concepts.

Structural Compression

Invariant: In a local public-goods game with strategic substitutes, Nash equilibria correspond to maximal independent sets of the interaction graph; the smallest such set is Pareto-best, and free-riding is the structural consequence of substitutability, not a behavioral pathology.

Mechanism: Effort is a local public good benefiting the closed neighborhood; best response is decreasing in neighbors' effort (substitutes); equilibrium produces specialization where providers form an independent set and non-providers free-ride; efficiency requires the minimum dominating set, which is generically smaller than the equilibrium provider set.

Cross-System Echo: The Bramoullé–Kranton structure is the game-theoretic analogue of **distributed resource allocation** in computer networks (where servers provide a shared resource and clients consume it), **coloring problems** in combinatorial optimization (independent sets are dual to proper colorings), and **division of labor** in organizational economics (team members specialize in complementary tasks). The common structure: a shared benefit, a private cost, and a graph that determines who benefits from whose provision. The equilibrium is always specialization, and the efficiency of specialization is always governed by the graph's independent-set structure.

Exercises

1. Enumerate all maximal independent sets of the cycle C_6 and verify that each corresponds to a Nash equilibrium of the binary public-goods game. Identify the Pareto-best equilibrium.
2. Show that on a tree T, the minimum dominating set can be computed in polynomial time via dynamic programming. Derive the recurrence.
3. For the linear-quadratic public-goods game on K_n, compute the symmetric interior equilibrium and compare total welfare to the single-provider equilibrium. Show that the single-provider equilibrium is Pareto-superior for sufficiently large n.
4. Construct a graph on 8 nodes where the number of maximal independent sets is maximized. How close can you get to the Moon–Moser bound $3^{8/3} \approx 6.16$?
5. In a study group of 6 players on a cycle graph, each player independently decides whether to run a solver sim (cost $c = 0.4$, benefit 1 to self and neighbors). Find all Nash equilibria and compute total welfare for each.
6. Add heterogeneous costs to the public-goods game on a star $K_{1,4}$: the center has cost $\gamma_c = 0.5$, peripherals have cost $\gamma_p = 1.0$. Derive the equilibrium and show that the low-cost center is more likely to be the provider.
7. **Argue, structurally,** why the inefficiency of Nash equilibria in local public-goods games comes from over-provision (too many providers) rather than under-provision, and explain how graph topology determines the severity of this inefficiency.

Network Formation Games

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Network Games **Node:** 8.5

In Nodes 8.1–8.4, the graph G was exogenous: given to us as a fixed interaction structure, with analysis focused on how equilibrium behavior depends on that structure. This node flips the question: **what if the graph itself is a strategic variable?** Players choose which links to form, and the resulting graph determines both the interaction structure and the payoffs. The foundational model is Jackson and Wolinsky (1996), which introduces **pairwise stability** — a bilateral equilibrium concept for network formation — and the central result that **efficient networks need not be stable and stable networks need not be efficient**. The gap between efficiency and stability is the defining tension of network formation theory, and it has structural parallels throughout game theory (analogous to the gap between Nash equilibrium and social optimum in standard games). This node develops the Jackson–Wolinsky framework, the connections model, pairwise stability, efficiency, and the stability–efficiency gap.

Formal Definition

A **network formation game** is specified by:

- A set of players $N = \{1, \dots, n\}$.
- A **network** (undirected graph) $g \subseteq \binom{N}{2}$, where $\binom{N}{2}$ is the set of all unordered pairs. Write $ij \in g$ if the link between i and j is present.
- A **value function** $v : \mathcal{G} \to \mathbb{R}$, where $\mathcal{G} = 2^{\binom{N}{2}}$ is the set of all networks on N. $v(g)$ is the total surplus generated by network g.
- An **allocation rule** $Y : \mathcal{G} \times \mathcal{V} \to \mathbb{R}^n$ specifying how the surplus is distributed: $Y_i(g, v)$ is player i's payoff from network g under value function v, with $\sum_i Y_i(g, v) = v(g)$.

Pairwise stability (Jackson–Wolinsky). A network g is **pairwise stable** if: 1. **No player wants to sever a link:** for every $ij \in g$, $Y_i(g, v) \geq Y_i(g - ij, v)$ and $Y_j(g, v) \geq Y_j(g - ij, v)$. 2. **No pair wants to form a link:** for every $ij \notin g$, if $Y_i(g + ij, v) > Y_i(g, v)$ then $Y_j(g + ij, v) < Y_j(g, v)$.

Link formation requires bilateral consent (both endpoints must benefit or be indifferent). Link deletion is unilateral (either endpoint can sever). This asymmetry — **bilateral formation, unilateral deletion** — is the defining feature of pairwise stability and distinguishes it from Nash equilibrium of a simultaneous link-announcement game.

Efficiency. A network g^* is **strongly efficient** if $v(g^*) \geq v(g)$ for all $g \in \mathcal{G}$. It maximizes total surplus without regard to distribution.

The connections model (Jackson–Wolinsky). Players benefit from both direct and indirect connections, with benefits decaying in distance. For parameters $\delta \in (0, 1)$ (decay) and $c > 0$ (cost per link):

$$Y_i(g) = \sum_{j \neq i} \delta^{d_{ij}(g)} - c \cdot |\{j : ij \in g\}|,$$

where $d_{ij}(g)$ is the shortest-path distance between i and j in g (infinite if disconnected, contributing 0).

Intuition

The move from "game on a fixed graph" to "game that forms the graph" is the game-theoretic analogue of the move from "optimization on a fixed domain" to "optimization over domains." In 8.1–8.4, we asked: given this network, what is the equilibrium? Now we ask: what network will rational agents build?

The Jackson–Wolinsky connections model captures the fundamental tradeoff in network formation: **direct links are costly but indirect links are valuable.** Forming a link costs c to each endpoint, but provides access to all nodes reachable through the network, with decaying benefit. A player at the center of a star pays for many direct links but benefits enormously from being close to everyone. A player at the periphery pays for one link but benefits from the indirect connections it provides.

The stability concept — pairwise stability — is weaker than Nash equilibrium. It only checks one-link deviations (add or delete a single link), not multi-link deviations. This makes it tractable but leaves open the possibility that a pairwise-stable network is vulnerable to coordinated multi-link changes. Various strengthened concepts (pairwise Nash stability, strong stability) exist but are harder to work with.

The central result of Jackson–Wolinsky is the **tension between stability and efficiency**: the efficient network (maximizing total surplus) may not be pairwise stable (individual players want to deviate), and the pairwise-stable network may not be efficient. This tension arises because the value of a link depends on the entire network (through indirect connections), but the cost is borne by the two endpoints. Externalities — the fact that forming a link ij affects the payoffs of third parties through changed path lengths — are the structural source of the gap.

Theorem (Jackson–Wolinsky, 1996). No allocation rule Y that is anonymous and component-balanced simultaneously guarantees that all efficient networks are pairwise stable (for all value functions).

This impossibility result is the network-formation analogue of the first welfare theorem's failure in the presence of externalities. It says: you cannot design a "fair" way to share surplus such that the socially optimal network always emerges from decentralized link-formation decisions.

Structural Role

8.5 endogenizes the graph that was exogenous in 8.1–8.4. It is the mechanism-design layer of Module 8: instead of taking the network as given, we ask what networks arise from strategic interaction and what the welfare consequences are.

- **From 8.1:** the graph-game framework, now with the graph as a strategic variable.
- **From 8.4:** public-goods equilibria depend on the graph; network formation determines which graphs are available.
- **To 8.6:** the stability concept of 8.5 interacts with the complementarity structure of 8.6 — on supermodular games, stable networks tend to be denser.
- **To Module 6:** the stability–efficiency gap parallels the incentive-compatibility constraints in mechanism design.

Mathematical Development

The connections model: efficient and stable networks

Efficient networks. Jackson and Wolinsky characterize efficient networks in the connections model:

- If $c < \delta - \delta^2$: the **complete graph** K_n is efficient. Links are cheap enough that every direct connection is worth forming.
- If $\delta - \delta^2 < c < \delta$: the **star** $K_{1,n-1}$ is efficient. Direct links are moderately costly, so it is efficient to route through a hub.
- If $c > \delta$: the **empty graph** is efficient. Links cost more than the direct benefit.

Pairwise-stable networks. Under the equal-split allocation rule $Y_i(g) = \sum_{j \neq i} \delta^{d_{ij}} - c \cdot d_i$:

- If $c < \delta - \delta^2$: K_n is pairwise stable (efficient and stable coincide).

- If $\delta - \delta^2 < c < \delta$: the star is efficient but **not necessarily pairwise stable**. The center of the star pays $(n-1)c$ in link costs but gets $(n-1)\delta$ in direct benefits. For the center, the star is attractive iff $\delta > c$. But peripherals might want to form links among themselves: peripheral i forming a link to peripheral j gains $\delta - \delta^2$ (replacing an indirect connection of length 2 with a direct one). So if $c < \delta - \delta^2$, peripherals want to add links, and the star is not stable. In the intermediate range $\delta - \delta^2 < c < \delta$, the star is stable.
- If $c > \delta$: the empty graph is pairwise stable.

The gap. The tension appears in the regime $\delta - \delta^2 < c < \delta$ with n large. The star is both efficient and stable. But consider an allocation rule that gives the center a smaller share of the surplus. Then the center may want to sever links, making the star unstable even though it is efficient. Jackson–Wolinsky show that no anonymous, component-balanced allocation rule can guarantee stability of all efficient networks across all value functions.

Pairwise stability vs Nash equilibrium

In the **strategic link-formation game** (Myerson, 1991), each player simultaneously announces a set of links they are willing to form: $L_i \subseteq N \setminus \{i\}$. Link ij forms iff $j \in L_i$ and $i \in L_j$. A Nash equilibrium of this game is a profile $(L_1, \ldots, L_n)$ where no player can improve by changing their announcement.

Every pairwise-stable network can be supported as a Nash equilibrium of the link-formation game (with appropriate strategies). But Nash equilibria can also include networks that are not pairwise stable — for example, the empty graph is always a Nash equilibrium (if no one announces any links, no one can improve by unilaterally announcing, because link formation requires bilateral consent). This multiplicity is why pairwise stability, despite being weaker, is the standard concept.

Network formation with transfers

If players can make side payments when forming links, the efficiency–stability gap can sometimes be closed. A link ij forms with transfer t_{ij} from i to j such that both are at least as well off. With unrestricted transfers and complete information, efficient networks can always be supported as pairwise stable (Bloch and Jackson, 2007). The practical limitation is that transfers require enforceable contracts, which may not exist in informal network settings.

Dynamic network formation

Consider a dynamic model: at each period, a random pair (i, j) is selected and can form or sever a link. The process converges to a stochastically stable network — the network most robust to random perturbations. Jackson and Watts (2002) show that stochastically stable networks are always pairwise stable, but not every pairwise-stable network is stochastically stable. The stochastic-stability refinement selects among pairwise-stable networks.

Worked Example

The connections model on 4 nodes

Four players, $N = \{1, 2, 3, 4\}$. Parameters: $\delta = 0.6, c = 0.2$. Since $c = 0.2 < \delta - \delta^2 = 0.6 - 0.36 = 0.24$, we are in the regime where the complete graph is efficient.

Complete graph K_4. Each player has 3 direct links. Payoff:

$$Y_i = 3(0.6) - 3(0.2) = 1.8 - 0.6 = 1.2.$$

Total: $4 \times 1.2 = 4.8$.

Star with center 1. Center payoff: $Y_1 = 3(0.6) - 3(0.2) = 1.2$. Peripheral payoff: $Y_2 = 0.6 + 2(0.36) - 0.2 = 0.6 + 0.72 - 0.2 = 1.12$. Total: $1.2 + 3(1.12) = 4.56 < 4.8$.

Pairwise stability of K_4. Delete link 12: player 1 goes from payoff 1.2 to $2(0.6)+0.36-2(0.2) = 1.56-0.4 = 1.16 < 1.2$. Player 1 does not want to delete. Similarly for all links. Add a link? K_4 is already complete. So K_4 is pairwise stable.

Now change to $c = 0.3$. We have $c = 0.3 > \delta - \delta^2 = 0.24$ but $c = 0.3 < \delta = 0.6$. The star is efficient.

Star payoff. Center: $Y_1 = 3(0.6) - 3(0.3) = 0.9$. Peripheral: $Y_2 = 0.6 + 2(0.36) - 0.3 = 1.02$. Total: $0.9 + 3(1.02) = 3.96$.

Complete graph payoff. Each: $Y_i = 3(0.6) - 3(0.3) = 0.9$. Total: $3.6 < 3.96$. Complete graph is not efficient.

Is the star pairwise stable? Peripheral i considers adding link to peripheral j: marginal benefit $\delta - \delta^2 = 0.24$, cost $c = 0.3$. Net: $-0.06 < 0$. No peripheral wants to add a link. Center considers severing a link: drops from payoff 0.9 to $2(0.6) + 0.36 - 2(0.3) = 1.56 - 0.6 = 0.96 > 0.9$. Wait — the center *gains* from severing a link. So the star is **not** pairwise stable at this allocation.

Rechecking: center with 3 links gets 0.9. With 2 links: $2(0.6) + 0.36 - 2(0.3) = 1.2 + 0.36 - 0.6 = 0.96$. With 1 link: $0.6 + 0 + 0 - 0.3 = 0.3$. The center prefers 2 links to 3. But the severed peripheral loses access. The pairwise-stable network in this range is a "partial star" or path-like structure, not the full star. **This is the stability–efficiency gap in action**: the efficient star is not pairwise stable because the center bears too much cost under equal-split allocation.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Network g	The current graph of study partnerships, coaching relationships, Discord memberships
Link ij	A bilateral study or information-sharing relationship between players i and j
Link cost c	Time, effort, and potential information leakage of maintaining a study relationship
Decay δ	Degradation of information value as it passes through intermediaries
Pairwise stability	Neither partner wants to end the study relationship, and no new pair wants to form one
Efficient network	The graph that maximizes total poker knowledge and edge across the community
Stability–efficiency gap	Individually rational study partnerships may not produce the best collective outcome

Concrete situation: study group formation

Five mid-stakes regulars live in the same city. Each could pair up with any other for study sessions. A study partnership costs time (c) but provides direct benefits (shared solver output, hand discussion) plus indirect benefits (access to your partner's other partners' insights, decayed by one hop).

The efficient network might be a star — one central organizer who connects everyone — because it minimizes total links while keeping everyone within 2 hops. But the center bears $4c$ in link costs while each peripheral bears only c. Under equal effort-sharing, the center may refuse to be the hub — the stability–efficiency gap.

In practice, poker study networks tend to be **small dense clusters** rather than stars, because the cost of maintaining a study relationship is shared and the benefit is direct and high. This matches the predictions of the connections model in the low-c regime where the complete graph (within a small group) is both efficient and stable. The groups stay small (5–8 people) because adding more members raises cost (more sessions, more coordination) faster than it raises benefit (diminishing marginal knowledge from additional partners).

Concrete situation: coaching market as network formation

Consider the coaching market. A coach (h) and a client (p) form a link if both benefit: the coach gets paid ($Y_h > 0$), the client gets value ($Y_p > 0$). The network of coaching relationships is pairwise stable if no coach wants to drop a client, no client wants to drop a coach, and no unmatched coach-client pair would both benefit from forming a link.

In practice, top coaches limit their client count (they sever links when the marginal time cost exceeds the marginal revenue). Clients leave coaches when they feel the benefit has plateaued (they sever links when Y_p drops below c). New coach-client links form when a referral creates a match that both parties find profitable. The resulting coaching network is sparse, hierarchical (a few high-degree coaches connected to many low-degree clients), and pairwise stable in the Jackson–Wolinsky sense.

The efficiency question is: **is the coaching market's pairwise-stable network also efficient?** Probably not. The efficient coaching network might involve more peer coaching (dense clusters of similar-skill players coaching each other at zero monetary cost), but peer coaching links are harder to form (bilateral benefit is less clear) and easier to sever (no financial commitment). The stability–efficiency gap here is between the sparse, hierarchical, pairwise-stable coaching market and the dense, egalitarian, efficient peer-learning network.

What this understanding buys you

Your poker network is a strategic choice, not an accident. Every study partnership, coaching relationship, and Discord membership is a link you formed (or maintained) because the benefit exceeded the cost. The network you are in right now is pairwise stable — otherwise you would have already changed it. But pairwise stability does not mean efficiency. You may be in a locally optimal network that is globally suboptimal.

The hub position is costly. If you are the central organizer of a study group, you bear disproportionate link costs (scheduling, moderating, curating). The stability–efficiency gap says that the hub position may not be individually rational even when it is socially efficient. If you find yourself in this position and feeling it is unsustainable, that is the Jackson–Wolinsky gap in action — the efficient network (star) is not stable under the current allocation of costs.

Transfers close the gap. In the theoretical framework, side payments can make efficient networks stable. In the poker context, this means: if the hub is compensated (the group pays for the organizer's time, or the organizer gets first access to shared content), the efficient star can become stable. Many successful study groups have an explicit or implicit compensation structure for the organizer — this is transfers making the efficient network pairwise stable.

Cross-Links

- **8.1 Graph games** — the framework where the graph is fixed; 8.5 endogenizes it.
- **8.4 Public goods on graphs** — efficiency of public-goods provision depends on the graph, which 8.5 determines.
- **6.1 Mechanism design** — the stability–efficiency gap parallels the incentive-compatibility constraints.
- **5.4 Folk theorem** — repeated interaction can sustain cooperative network structures that are not stable in a one-shot formation game.
- **2.1 Nash equilibrium** — pairwise stability is a bilateral refinement of Nash for network games.
- **Matching theory** — two-sided network formation has structural parallels to stable matching.

Structural Compression

Invariant: In a network formation game, pairwise-stable networks need not be efficient and efficient networks need not be pairwise stable; the gap arises from externalities in link formation (forming a link affects third parties through changed path lengths), and no anonymous, component-balanced allocation rule eliminates the gap for all value functions.

Mechanism: Bilateral link formation (consent required) and unilateral link deletion; value function depending on the entire network through path-length effects; externalities from indirect connections creating a wedge between private and social returns to link formation.

Cross-System Echo: The Jackson–Wolinsky stability–efficiency gap is the network analogue of the **first welfare theorem's failure under externalities** in general equilibrium theory. It is also structurally parallel to the **price of anarchy** in algorithmic game theory (the gap between the Nash equilibrium and the social optimum in routing, congestion, and resource-allocation games). In all three settings, decentralized decision-making fails to achieve the social optimum because agents do not internalize the full externalities of their choices — and the structure of the externality (positive/negative, local/global) determines the direction and magnitude of the gap.

Exercises

1. For the connections model with $n = 3$, $\delta = 0.5$, enumerate all 8 possible networks and compute $v(g)$ and $Y_i(g)$ for each. Identify the efficient and pairwise-stable networks as a function of c.
2. Prove that the empty graph is always pairwise stable when $c > \delta$. Show that it is the unique pairwise-stable network in this regime.
3. Show that under the equal-split allocation rule in the connections model, the complete graph is pairwise stable iff $c < \delta - \delta^2$. Interpret the condition: adding a direct link replaces an indirect connection of length 2, and the net benefit must be positive.
4. Construct a 5-node network that is pairwise stable but not efficient (under the connections model with specific δ, c). Compute the efficiency loss.
5. In the strategic link-formation game (simultaneous announcements), show that the empty graph is always a Nash equilibrium. Explain why this does not contradict pairwise stability.
6. A poker study group of 6 players has link cost $c = 0.15$ and decay $\delta = 0.7$. Determine whether the complete graph, the star, and the cycle are pairwise stable. Which is efficient?
7. **Argue, structurally,** why the stability–efficiency gap is an inherent feature of network formation with externalities, not an artifact of a particular model. Connect this to the general principle that decentralized decisions with externalities are generically inefficient.

Strategic Complementarities

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Network Games **Node:** 8.6

Strategic complementarities are the structural force behind coordination, contagion, and tipping points in network games. An action exhibits strategic complementarity when each player's incentive to take that action *increases* with the number (or proportion) of neighbors who take it. This is the formal content of "doing what others do is rational" – the property that generates bandwagon effects, technology adoption cascades, bank runs, and coordination failures. Nodes 8.1-8.5 introduced games on graphs, coordination, diffusion, public goods, and network formation. Node 8.6 synthesizes these with a unified mathematical treatment grounded in the theory of **supermodular games** (Topkis 1979, Milgrom-Roberts 1990, Vives 1990). The key results are: supermodular games always have pure-strategy Nash equilibria (by Tarski's fixed-point theorem), the equilibrium set has a lattice structure with greatest and least elements, and comparative statics are monotone – an increase in any player's incentive to act raises everyone's equilibrium action. These properties make supermodular games the analytically tractable workhorse of applied game theory, especially in network and industrial-organization contexts.

Formal Definition

Supermodular game. A game $G = (N, \{S_i\}, \{u_i\})$ is **supermodular** if each S_i is a complete lattice and the payoff functions satisfy:

1. **Supermodularity in own action.** For each i and each fixed s_{-i}, $u_i(\cdot, s_{-i})$ is supermodular on S_i:

$$u_i(s_i \vee s_i', s_{-i}) + u_i(s_i \wedge s_i', s_{-i}) \geq u_i(s_i, s_{-i}) + u_i(s_i', s_{-i})$$

   for all $s_i, s_i' \in S_i$, where $\vee$ and $\wedge$ are the lattice join and meet.

2. **Increasing differences (strategic complementarity).** For each i, u_i has increasing differences in (s_i, s_{-i}): if $s_{-i}' \geq s_{-i}$ (in the product-lattice order), then

$$u_i(s_i', s_{-i}') - u_i(s_i, s_{-i}') \geq u_i(s_i', s_{-i}) - u_i(s_i, s_{-i})$$

   for all $s_i' \geq s_i$.

Equivalently, for smooth payoffs on $\mathbb{R}$, strategic complementarity is $\partial^2 u_i / \partial s_i \partial s_j \geq 0$ for all $j \neq i$: increasing others' actions raises the marginal return to one's own action.

Strategic complementarity on a network. In a network game on graph $G = (N, E)$, player i's payoff depends on their own action s_i and the actions of their neighbors $\mathcal{N}_i$. Strategic complementarity on the network means

$$\frac{\partial^2 u_i}{\partial s_i \partial s_j} \geq 0 \quad \text{for all } j \in \mathcal{N}_i.$$

The canonical example is the **linear-quadratic network game**:

$$u_i(s_i, s_{-i}) = a_i s_i - \frac{1}{2} s_i^2 + \lambda \sum_{j \in \mathcal{N}_i} g_{ij} s_i s_j,$$

where a_i is a standalone benefit, $\lambda > 0$ is the complementarity parameter, and g_{ij} is the edge weight. The best response is linear:

$$s_i^* = a_i + \lambda \sum_{j \in \mathcal{N}_i} g_{ij} s_j,$$

which in matrix form is $s^* = a + \lambda G s^*$, giving the Nash equilibrium $s^* = (I - \lambda G)^{-1} a$ when λ is small enough for convergence.

Strategic substitutes. The opposite: $\partial^2 u_i/\partial s_i \partial s_j \leq 0$. Each player's incentive to act *decreases* with neighbors' actions. Free-riding and anti-coordination games exhibit strategic substitutes.

Intuition

Strategic complementarity formalizes the idea that "doing more is more attractive when others do more." This positive feedback loop is the engine of coordination games, technology adoption, social conventions, and financial contagion. The formal consequence is that the game has clean analytical properties: equilibria exist, they are ordered, and they respond monotonically to parameter changes. This makes supermodular games the first-choice model for any applied setting where coordination effects are present.

The contrast with strategic substitutes is sharp. Under substitutes, the feedback is negative: doing more is less attractive when others do more. This produces anti-coordination (Hawk-Dove, free-riding on public goods) rather than coordination. The equilibrium structure under substitutes is different: uniqueness is more common, and the comparative statics are ambiguous.

On networks, strategic complementarity interacts with the graph structure to produce the phenomena studied in 8.2 (coordination), 8.3 (cascades), and 8.4 (public goods). High-degree nodes (hubs) have more neighbors, so their complementarity effect is amplified: a hub's action influences more neighbors, and more neighbors influence the hub. This produces the **key-player effect** (Ballester-Calvo-Armengol-Zenou 2006): in a network game with complementarities, the player whose removal most reduces total equilibrium activity is the one with the highest "intercentrality" – a network centrality measure derived from the equilibrium formula $s^* = (I - \lambda G)^{-1} a$.

The Ballester et al. result connects network game theory to policy: if you want to reduce total crime in a network (where crime exhibits complementarities), removing the key player (the one with the highest intercentrality, not necessarily the highest degree or the most active criminal) has the largest effect. This is a non-obvious conclusion that requires the full supermodular-game-on-network machinery.

Structural Role

8.6 is the capstone of Module 8, providing the unified mathematical framework that connects 8.1-8.5:

- **8.1 (Graph games)** introduced the network structure; 8.6 adds the strategic complementarity that makes the structure matter analytically.
- **8.2 (Coordination on networks)** is the special case of binary-action supermodular games on graphs.
- **8.3 (Diffusion and cascades)** is the dynamic version: cascades happen because complementarity creates tipping points.
- **8.4 (Public goods on graphs)** can exhibit either complementarity or substitutability depending on the production function.
- **8.5 (Network formation)** asks which networks are formed strategically; 8.6 analyzes play on a given network.

Beyond Module 8, strategic complementarities connect to:

- **9.5 (Convergence to equilibrium)** – supermodular games are the class where learning dynamics are best-behaved (monotone convergence to the equilibrium set).
- **10.1 (Games embedded in institutions)** – institutions often create complementarities by design.
- **Statistics (Ising models)** – the Ising model on a graph is a supermodular game with binary strategies and nearest-neighbor complementarities.

Mathematical Development

Tarski's fixed-point theorem and equilibrium existence

Theorem (Tarski 1955). If L is a complete lattice and $f : L \to L$ is an order-preserving (monotone) function, then the set of fixed points of f is nonempty and itself forms a complete lattice.

Application to supermodular games. Define the joint best-response correspondence $B(s) = (B_1(s_{-1}), \ldots, B_n(s_{-n}))$. Under strategic complementarity, B is monotone: if $s' \geq s$, then $B(s') \geq B(s)$. By Tarski's theorem, the set of Nash equilibria (fixed points of B) is nonempty and forms a complete lattice with a greatest element $\bar{s}$ and a least element $\underline{s}$.

The greatest and least equilibria are reached by iterated best response from the top and bottom of the strategy space respectively. This gives a constructive algorithm for finding them and establishes that best-response dynamics from any initial condition converge to the equilibrium set.

Monotone comparative statics

Theorem (Milgrom-Roberts 1990). In a supermodular game, if a parameter t increases the marginal return to each player's action (formally: u_i has increasing differences in (s_i, t)), then the greatest and least equilibria are weakly increasing in t.

This is the formal version of "if incentives increase, equilibrium actions increase." It is the reason supermodular games are the workhorse of applied theory: you can do comparative statics without solving for the full equilibrium – just check the direction of the parameter change and the complementarity structure.

The Ballester-Calvo-Armengol-Zenou model

In the linear-quadratic network game, the Nash equilibrium is $s^* = (I - \lambda G)^{-1} a$, which can be expanded as a power series:

$$s_i^* = a_i + \lambda \sum_j g_{ij} a_j + \lambda^2 \sum_j \sum_k g_{ij} g_{jk} a_k + \cdots$$

This is a **Katz-Bonacich centrality** measure: player i's equilibrium action is proportional to a weighted count of all walks of all lengths emanating from i in the network, with longer walks discounted by powers of λ. The equilibrium action is higher for more central players (by the Bonacich centrality measure).

Key-player identification. The key player is the one whose removal causes the largest drop in total equilibrium action $\sum_i s_i^*$. Ballester et al. show this is the player with the highest **intercentrality**, defined as

$$c_i^* = \frac{b_i(G, \lambda)^2}{m_{ii}(G, \lambda)},$$

where b_i is the Bonacich centrality and m_{ii} is the (i, i)-entry of $(I - \lambda G)^{-1}$.

Convergence of learning dynamics in supermodular games

Theorem (Milgrom-Roberts 1990). In a finite supermodular game, iterated elimination of dominated strategies yields the same result from the top and bottom of the strategy space: the greatest and least surviving strategy profiles are the greatest and least Nash equilibria. Moreover, any learning dynamics that are "adaptive" (each player's strategy is eventually a best response to some belief in the range of opponents' recent play) converge to the interval $[\underline{s}, \bar{s}]$.

This is one of the strongest convergence results in game theory. It says that in supermodular games, strategic complementarity pins down the learning dynamics: they always converge to the equilibrium set. This is the connection to Module 9: supermodular games are the "well-behaved" class where learning works.

Binary-action supermodular games and thresholds

When actions are binary ($s_i \in \{0, 1\}$), strategic complementarity reduces to: each player i has a **threshold** k_i such that they choose action 1 iff at least k_i of their neighbors choose 1. This is the threshold model of 8.2 and 8.3, now derived as a special case of the supermodular framework. The equilibria are the fixed points of the threshold dynamics; the greatest and least equilibria are the "maximal coordination" and "minimal coordination" outcomes.

Worked Example

Linear-quadratic game on a star network

Network. 5 players: a hub (player 1) connected to 4 spokes (players 2-5). No spoke-to-spoke connections.

Payoffs. $u_i = as_i - s_i^2/2 + \lambda s_i \sum_{j \in \mathcal{N}_i} s_j$, with $a = 1$ for all players, $\lambda = 0.2$.

Best responses. Hub: $s_1^* = 1 + 0.2(s_2 + s_3 + s_4 + s_5)$. Spoke i: $s_i^* = 1 + 0.2s_1$.

Equilibrium. By symmetry, spokes are identical: $s_2 = s_3 = s_4 = s_5 \equiv s_{\text{sp}}$. Hub: $s_1 = 1 + 0.8s_{\text{sp}}$. Spoke: $s_{\text{sp}} = 1 + 0.2s_1 = 1 + 0.2(1 + 0.8s_{\text{sp}}) = 1.2 + 0.16s_{\text{sp}}$. So $0.84s_{\text{sp}} = 1.2$, $s_{\text{sp}} \approx 1.429$. Hub: $s_1 = 1 + 0.8 \times 1.429 \approx 2.143$.

Bonacich centrality. The hub's equilibrium action (2.143) is much larger than each spoke's (1.429), reflecting the hub's higher centrality. If we remove the hub, the spokes become isolated and each plays $s_i^* = 1$. Total action drops from $2.143 + 4 \times 1.429 = 7.86$ to $4 \times 1 = 4$. Removing a spoke instead: remaining network has hub + 3 spokes. New spoke equilibrium: $s'_{\text{sp}} = 1 + 0.2s'_1$, $s'_1 = 1 + 0.6s'_{\text{sp}}$. Solving: $s'_{\text{sp}} \approx 1.363$, $s'_1 \approx 1.818$. Total: $1.818 + 3 \times 1.363 = 5.91$. Drop: $7.86 - 5.91 = 1.95$.

Removing the hub drops total action by 3.86; removing a spoke drops it by 1.95. The hub is the key player – its intercentrality is highest. Policy implication: targeting the hub has roughly double the impact of targeting a spoke.

Comparative statics: increasing the complementarity parameter

Increase λ from 0.2 to 0.3. New spoke equilibrium: $s_{\text{sp}} = 1 + 0.3s_1$, $s_1 = 1 + 1.2s_{\text{sp}}$. $s_{\text{sp}} = 1 + 0.3(1 + 1.2s_{\text{sp}}) = 1.3 + 0.36s_{\text{sp}}$. $0.64s_{\text{sp}} = 1.3$, $s_{\text{sp}} \approx 2.031$. $s_1 \approx 3.438$. Total: $3.438 + 4 \times 2.031 = 11.56$.

As λ increases, everyone's equilibrium action increases – monotone comparative statics, exactly as the theory predicts. The amplification is stronger at the hub because the hub has more neighbors to amplify through.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Strategic complementarity	The more your opponents bluff, the more profitable it is for you to bluff-catch (and vice versa)
Strategic substitutability	The more your opponents value-bet, the less profitable it is for you to value-bet the same spot
Supermodular structure	Games where aggressive play breeds aggressive response – escalation dynamics
Network complementarity	In a table with regulars, one player's aggression raises all others' aggression through strategic response
Key player (hub)	The player whose removal most changes the table dynamic – the meta-setter

Formal object	Poker instantiation
Bonacich centrality	A player's "influence" measured by how many strategic interactions they mediate

Concrete situation: table dynamics driven by one aggressive player

A 6-max table has five balanced regulars and one hyper-aggressive LAG. The LAG's aggression creates complementarity: the regulars' best response to the LAG is to widen their calling ranges, which in turn makes the LAG's bluffs less profitable, which adjusts the LAG's strategy, and so on. The equilibrium of the full table is determined by the complementarity structure: the LAG is the hub, the regulars are the spokes, and the equilibrium aggression levels are higher than if the LAG were replaced by a balanced player.

If the LAG leaves the table, the entire table dynamic shifts: everyone tightens up, average pot size drops, and the aggression level decreases. The LAG is the key player – their removal has the largest effect on total table action. This is the Ballester et al. result applied to the poker table.

What this understanding buys you

Complementarity explains why poker metas are "sticky." Once a table converges to a high-aggression equilibrium (everyone 3-betting light), it stays there because any individual deviation to tightness is punished by the complementarity – your opponents are already playing aggressively, so tightening up costs you pots. The complementarity creates a coordination trap: everyone would prefer lower aggression (more readable play), but no one can unilaterally move there. This is the supermodular game's greatest-equilibrium trap.

Key-player removal is the fastest way to shift a meta. If you want to change a table's dynamic (e.g., from hyper-aggressive to balanced), the most efficient intervention is removing the key player – the one whose style has the most amplified effect through the complementarity structure. In practice, this means: if the LAG leaves, the table calms down much faster than if a nit leaves.

Cross-Links

- **8.1 Graph games** – the network structure on which complementarities operate.
- **8.2 Coordination on networks** – binary-action complementarities as a special case.
- **8.3 Diffusion and cascades** – complementarity-driven tipping points.
- **8.4 Public goods on graphs** – complementarity vs substitutability in provision.
- **9.5 Convergence to equilibrium** – supermodular games are where learning converges cleanly.
- **10.1 Games embedded in institutions** – institutional complementarities.
- **Statistics (Ising models)** – binary supermodular games as lattice spin models.
- **Economics (industrial organization)** – Cournot and Bertrand games with complementary products.

Structural Compression

Invariant: In supermodular games (games with strategic complementarities), the Nash equilibrium set is nonempty and forms a complete lattice with greatest and least elements; comparative statics are monotone (increasing parameters raise equilibrium actions); and adaptive learning dynamics always converge to the equilibrium set. On networks, complementarities amplify through the graph structure, with the key player determined by Bonacich intercentrality.

Mechanism: Increasing differences in the payoff function ensure that the best-response correspondence is monotone; Tarski's fixed-point theorem gives existence and lattice structure; the power-series expansion $(I - \lambda G)^{-1}a$ in linear-quadratic network games links equilibrium actions to walk-based centrality measures; comparative statics follow from monotone selection within the lattice.

Cross-System Echo: Strategic complementarity is the game-theoretic formalization of **positive feedback** in systems theory and **ferromagnetic interaction** in statistical physics. The Ising model on a graph with ferromagnetic couplings is a binary supermodular game; the Gibbs measure is the stochastic-stability-selected equilibrium; the phase transition (spontaneous magnetization) is the cascade threshold of 8.3. The lattice structure of equilibria mirrors the lattice of Gibbs states in statistical mechanics. The connection is not metaphorical – the mathematics is identical, and insights from one field transfer directly to the other.

Exercises

1. Verify that the linear-quadratic network game with $\lambda > 0$ satisfies the strategic complementarity condition (positive cross-partial derivatives). Then verify supermodularity.
2. For the star network with 5 nodes from the worked example, compute the Nash equilibrium for $\lambda = 0.1, 0.2, 0.3$ and verify the monotone comparative statics result.
3. Compute the Bonacich centrality vector for a 4-node cycle graph with uniform edge weights and $\lambda = 0.15$. Identify the key player (by symmetry, all are equivalent – verify this).
4. In a binary-action supermodular game on a complete graph (K_n), derive the threshold model: each player adopts action 1 iff at least k others do, where k depends on the payoff parameters. Characterize all Nash equilibria.
5. Show that in a game with strategic substitutes ($\partial^2 u_i / \partial s_i \partial s_j \leq 0$), the best-response correspondence is decreasing, and Tarski's theorem still applies (the equilibrium set is nonempty) but the greatest and least equilibria have different comparative statics. Verify with a Cournot duopoly example.
6. For the linear-quadratic game on a random Erdos-Renyi graph $G(n, p)$ with $n = 20, p = 0.2, \lambda = 0.1$, compute (numerically) the equilibrium and identify the key player. Verify that the key player has the highest intercentrality, not necessarily the highest degree.
7. **Argue, structurally,** why supermodular games are the "well-behaved" class of games for learning and dynamics. What specific property of the best-response correspondence makes convergence guaranteed, and why does this fail for general games?