Diffusion and Cascades

Game Theory · Module 8 — Network Games

Node 8.3

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
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."

Game Theory · Module 8 — Network Games · Node 8.3
Concepts: graph-theory, dynamical-systems, computational-complexity
Prerequisites: 8.1, 8.2
Enables: 8.5, 8.6