Reputation and the Chain Store Paradox

Game Theory · Module 5 — Repeated Games

Node 5.5

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Repeated Games **Node:** 5.5

The chain-store paradox is one of the most important puzzles in game theory, and its resolution through reputation theory is one of the most elegant. The paradox goes like this: a monopolist faces a sequence of T potential entrants, one per period. Against each, the incumbent can "fight" (costly, deters entry) or "accommodate" (cheap, invites entry). The game has a unique Nash equilibrium found by backward induction: accommodate every entrant, because in the last period the threat to fight is not credible. But in practice, real monopolists fight — and it works, because entrants are deterred. The paradox is the gap between the backward-induction prediction and observed behavior. Kreps, Wilson, Milgrom, and Roberts (1982) resolved the paradox by showing that even a **small probability** that the incumbent is a "tough" type — an irrational commitment type who always fights — is sufficient to sustain fighting as equilibrium behavior for the rational type, through a **reputation mechanism**. The rational incumbent mimics the tough type to build a reputation for toughness, which deters later entrants. This is the first formal demonstration that asymmetric information about types (Module 4) can dramatically change the structure of finite repeated games (Module 5.1), and it connects directly to the signaling games of 4.5: fighting is a costly signal of type.

Formal Definition

Chain-store game. A single incumbent (player I) plays a sequential game against T entrants $(E_1, E_2, \ldots, E_T)$, each arriving in one period. In each period t:

1. Entrant E_t observes the history of the incumbent's prior play and chooses: **Enter** (e) or **Stay out** (o).
2. If E_t enters, the incumbent chooses: **Fight** (f) or **Accommodate** (a).
3. If E_t stays out, no action is needed.

Stage payoffs (per period): - Stay out: $(M, 0)$ where $M > 0$ is the monopoly profit for I, and 0 for the entrant. - Enter + Accommodate: (D, d) where $D < M$ and $d > 0$ (duopoly profits). - Enter + Fight: $(F, -c)$ where $F < D < M$ and $c > 0$ (both lose in a fight, but incumbent loses less).

The ordering $F < D < M$ is crucial: accommodation is strictly preferred to fighting for the rational incumbent in a single interaction.

Backward-induction prediction (complete information). In the last period, if entry occurs, the incumbent accommodates (since $D > F$). Knowing this, E_T enters (since $d > 0$). In period $T-1$, the incumbent's fighting does not deter E_T (entry in T is unconditional), so the incumbent accommodates; E_{T-1} enters. By backward induction: every entrant enters, the incumbent always accommodates. Total payoff to incumbent: $T \cdot D$.

The paradox: In reality, monopolists fight early entrants and this deters later ones. The backward-induction solution says this should never happen.

Reputation resolution (Kreps–Wilson–Milgrom–Roberts 1982). Introduce a **small perturbation** of the information structure. With probability $p > 0$ (possibly very small), the incumbent is a "tough type" I_T

who always fights entry (this is their dominant strategy, perhaps due to irrational commitment, managerial incentives, or a different payoff function where $F > D$). With probability $1 - p$, the incumbent is the rational type I_R with payoffs as above.

Equilibrium behavior. In this perturbed game, the rational type I_R fights early entrants to build reputation — to maintain the entrants' posterior belief that the incumbent might be the tough type. As long as the posterior is high enough, later entrants stay out. Near the end of the game, the reputation eventually unravels (the rational type eventually accommodates), but the number of periods lost to accommodation is bounded regardless of T, while the number of periods of monopoly (entrants staying out due to reputation) grows with T. As $T \to \infty$, the fraction of periods with entry goes to zero.

Formal theorem. For any $\epsilon > 0$ and any initial probability $p > 0$ of the tough type, there exists T^* such that for all $T > T^*$, in every sequential equilibrium the rational incumbent's average payoff is within ϵ of the monopoly payoff M.

Intuition

The resolution has a beautiful structure: reputation is **signaling across time**. Each period the incumbent fights is a "signal" that they might be the tough type. Each fight increases (or maintains) the entrants' posterior probability μ_t that they face a tough type. The posterior evolves via Bayes' rule:

$$\mu_{t+1} = \frac{\mu_t}{\mu_t + (1 - \mu_t) \cdot \sigma_t}$$

where σ_t is the probability that the rational type fights in period t. If the rational type fights for sure ($\sigma_t = 1$), the posterior does not change: $\mu_{t+1} = \mu_t$. If the rational type sometimes accommodates ($\sigma_t < 1$), the posterior increases after a fight (Bayesian updating on the less-likely event conditional on rationality). If the rational type accommodates, the posterior drops to 0 — the type is revealed.

The rational type's strategic problem is: how long to mimic the tough type before revealing? The answer depends on the entry-deterrence effect of the reputation. There is a critical posterior μ^* above which entrants stay out and below which they enter. The rational type fights until the posterior is close to μ^*, then mixes (fighting with some probability to maintain the posterior near the threshold), and eventually reveals near the end of the game.

The deep lesson is that **asymmetric information rescues backward-induction games from their paradoxical predictions**. The backward-induction collapse of 5.1 requires common knowledge of rationality. Reputation theory shows that even a tiny departure from this common knowledge — a probability p of irrationality — can change the equilibrium outcome dramatically. This is a structural fragility result for backward induction.

Why does the result require only a "small" probability of the tough type? Because the signaling game is extremely powerful for the informed party. Each time the rational incumbent fights, the entrant's belief moves. The cost of fighting is bounded (it is at most $D - F$ per period), but the gain from deterrence is also bounded per period but accrues over many periods. For T large, the total deterrence gain dominates the total fighting cost, and the rational type is willing to mimic the tough type for the majority of the game. The initial probability p matters because it determines the "starting stock of reputation" — but even a tiny p is sufficient for large T.

Structural Role

5.5 resolves the tension between 5.1 (backward-induction collapse in finite repeated games) and 5.2 (rich equilibrium set in infinite repeated games). It shows that the collapse of 5.1 is fragile: a small informational perturbation (incomplete information about types) replaces the collapse with reputation-building behavior that looks cooperative. This is an alternative route to cooperation — not through infinite patience (5.2) but through asymmetric information.

5.5 connects directly to:

- **4.5 Signaling games** — reputation is multi-stage signaling.
- **4.6 Sequential equilibrium** — the solution concept for the perturbed chain-store game.
- **5.1 Finitely repeated games** — the paradox that reputation theory resolves.
- **9.5 Convergence to equilibrium** — reputation models as a limit of learning dynamics.
- **10.4 Commitment and credibility** — reputation as a substitute for formal commitment.

The reputation mechanism is also the foundation of modern theories of deterrence, brand building, credible threats, and organizational reputation — all of which appear in Module 10.

Mathematical Development

Posterior dynamics and the Bayesian learning structure

Let μ_t be the common prior/posterior at the start of period t that the incumbent is the tough type. At $t = 1$, $\mu_1 = p$.

If the entrant enters and the incumbent fights in period t, the posterior updates:

$$\mu_{t+1} = \frac{\mu_t \cdot 1}{\mu_t + (1-\mu_t) \cdot \sigma_t} = \frac{\mu_t}{\mu_t + (1-\mu_t)\sigma_t}$$

where $\sigma_t \in [0,1]$ is the rational type's mixing probability over fighting.

If the entrant enters and the incumbent accommodates, the posterior drops to 0: the incumbent is revealed as rational. This is why accommodation is "informationally costly" — it destroys reputation.

If the entrant stays out, the posterior is unchanged: $\mu_{t+1} = \mu_t$.

The entrant's entry decision

Entrant E_t enters if the expected payoff from entry is positive:

$$\mu_t \cdot (-c) + (1-\mu_t)[\sigma_t \cdot (-c) + (1-\sigma_t) \cdot d] \geq 0.$$

The left-hand side is: probability of tough type (always fights, gives $-c$) plus probability of rational type (fights with probability σ_t, accommodates otherwise). Entry is profitable when the posterior is low enough (small chance of being fought). There is a critical posterior:

$$\mu^* = \frac{d}{c+d},$$

above which entry gives negative expected payoff (assuming the rational type fights for sure, the "optimistic" bound for the incumbent). Below μ^*, entry is profitable under any rational-type mixing.

Equilibrium characterization

Phase 1: reputation building (early periods, $\mu_t > \mu^*$). Entry is deterred. Entrants stay out. Posterior unchanged. The incumbent enjoys monopoly profit M.

Phase 2: reputation maintenance (periods where μ_t is near μ^*). The rational incumbent mixes: fights with probability σ_t chosen so that the entrant is indifferent between entering and not. The posterior is pinned near μ^*. Entry occurs sometimes but not always.

Phase 3: endgame unraveling (last few periods). The rational type eventually accommodates (or the game is about to end), revealing the type with positive probability. The last period is always accommodation (by the backward-induction argument of the last period). The unraveling propagates backward from the terminal period but only for a bounded number of periods — the critical number depends on the payoffs and p, not on T.

Key result. The number of periods in which entry occurs and the incumbent accommodates is bounded by a constant $K(p)$ that depends on p but not on T. For T large, the fraction of periods with accommodation goes to zero, and the incumbent's average payoff approaches M.

Fudenberg–Levine (1989, 1992): general reputation results

Fudenberg and Levine extended the chain-store analysis to general repeated games with incomplete information. Their main result:

Theorem (Fudenberg–Levine 1989). Consider a long-run player (the incumbent) facing a sequence of short-run opponents (entrants). Suppose there is probability $p > 0$ that the long-run player is a commitment type who plays a fixed strategy σ^c. Then in any Nash equilibrium, the long-run player's average payoff is at least $BR_{\min}(\sigma^c)$ — the minimum best-response payoff of any short-run opponent against σ^c — minus an $O(1/T)$ error term.

In the chain-store case, the commitment type always fights, and the short-run opponent's best response to "always fight" is "stay out," which gives the incumbent M. So the general result implies the chain-store conclusion.

The Fudenberg–Levine bound is sharp and has wide applications: firm reputation, central-bank credibility, political reputation, and military deterrence.

Reputation fragility

Reputation can be fragile under several perturbations:

- **Two-sided incomplete information.** If both players have private types, reputation building becomes a war of attrition, and the results are qualitatively different (Kreps–Wilson 1982).
- **Imperfect monitoring.** If the entrant cannot perfectly observe the incumbent's action (only a noisy signal), reputation builds more slowly (Fudenberg–Levine 1992). With sufficient noise, reputation effects can vanish.
- **Multiple commitment types.** If there are many possible commitment types, the reputation result weakens — the incumbent cannot build a reputation for a specific strategy if entrants are uncertain about which commitment type they face.

Worked Example

Chain-store game with $T = 5$, $p = 0.2$

Stage payoffs: $M = 5$ (monopoly), $D = 3$ (duopoly), $F = 1$ (fight), entry payoff $d = 2$, fight loss $c = 3$.

Critical posterior: $\mu^* = d/(c + d) = 2/5 = 0.4$.

Period 1: $\mu_1 = 0.2 < \mu^* = 0.4$. The posterior is below threshold! Entry is profitable even if the rational type fights for sure: expected payoff from entry $= 0.2 \cdot (-3) + 0.8 \cdot (-3) = -3$ if fighting for sure. Wait — actually, if the rational type fights for sure ($\sigma_1 = 1$): expected entry payoff $= \mu_1 \cdot (-c) + (1 - \mu_1) \cdot \sigma_1 \cdot (-c) + (1 - \mu_1)(1 - \sigma_1) \cdot d = -c$. With $\sigma_1 = 1$, entry payoff is -3, so entrant stays out. With σ_1 lower, entry payoff increases. At $\sigma_1 = d/(c + d) = 2/5$ (the mix that makes the entrant indifferent between entering and not, at this posterior), entry gives 0.

Let me recompute. Entry expected payoff given posterior μ and rational type mixing σ:

$$E[\text{entry}] = [\mu + (1-\mu)\sigma] \cdot (-c) + (1-\mu)(1-\sigma) \cdot d.$$

Setting this to zero: $[\mu + (1-\mu)\sigma] \cdot c = (1-\mu)(1-\sigma) \cdot d$.

At $\mu = 0.2$: $[0.2 + 0.8\sigma] \cdot 3 = 0.8(1-\sigma) \cdot 2$. $0.6 + 2.4\sigma = 1.6 - 1.6\sigma$. $4\sigma = 1$. $\sigma = 0.25$.

So the rational type mixes: fights with probability 0.25, accommodates with probability 0.75. The entrant is indifferent.

If the rational type fights, posterior updates to:

$$\mu_2 = \frac{0.2}{0.2 + 0.8 \cdot 0.25} = \frac{0.2}{0.4} = 0.5.$$

Now $\mu_2 = 0.5 > \mu^* = 0.4$. So if the incumbent fights in period 1, the posterior jumps above threshold and E_2 stays out. From period 2 onward (assuming the posterior stays above 0.4), entrants stay out and the incumbent earns monopoly profits.

If the rational type accommodates in period 1, the posterior drops to 0 (type revealed), and all subsequent entrants enter and are accommodated: payoff $D = 3$ per period.

Rational type's expected payoffs. Fighting in period 1: pay $F = 1$ this period, gain monopoly for periods 2-5 (4 periods at $M = 5$). Accommodating: pay $D = 3$ this period, earn $D = 3$ in all remaining periods (type revealed). With discount factor $\delta = 0.9$:

Fight: $(1 - 0.9)(1) + 0.9[(1 - 0.9)(5) \cdot 4/(1 - 0.9)]$ — let me just compute the undiscounted sum for simplicity. Fight: $1 + 5 + 5 + 5 + 5 = 21$. Accommodate: $3 + 3 + 3 + 3 + 3 = 15$. Even undiscounted, fighting dominates. The rational type strictly prefers to fight in period 1.

But the equilibrium has the rational type fighting only with probability 0.25 (mixing). Why? Because if the rational type fights for sure, the entrant's best response is to stay out (guaranteed fight), and then the rational type's "best response" is to not fight (nobody enters anyway). The equilibrium requires mixing to pin the entrant at indifference. The key insight: the mixing itself reveals information probabilistically, and the Bayesian dynamics are what sustain the equilibrium.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Incumbent type	Your playing style / identity: are you a maniac (tough type) or a balanced regular (rational type)?
Entrant's belief μ_t	Opponent's belief about whether you are a maniac or a balanced player
Fighting	Making aggressive plays: overbets, 5-bet jams with air, re-raises with blockers
Accommodation	Playing standard balanced poker
Reputation building	Making aggressive plays early to build the image that you might be a maniac, even though you're balanced
Critical posterior μ^*	The belief threshold above which your opponent is deterred from calling your bets
Unraveling	Late in the session, when the "game" is about to end, your reputation incentives weaken

Concrete situation: table image in the first orbit

You sit down at a new 6-max cash game. Nobody knows you. In the first 20 hands, you 3-bet three pots, take down two uncontested, and show a bluff in the third. Your opponents' posterior on "maniac" has just spiked. For the next 100 hands, you play balanced, but you receive far fewer calls on your value bets — opponents give you credit for aggression based on the first orbit. The early aggression was costly (you risked money on marginal spots) but the reputation benefit was enormous: the deterrence effect reduces the number of tough decisions you face and increases your fold equity across the session.

This is the chain-store reputation mechanism. The cost of fighting (the 3-bet bluffs that might get snapped off) is the short-run investment. The payoff is that opponents believe you might be a maniac and adjust accordingly, staying out of pots (folding more) or playing passively. As the session nears its end (last few hands), you might adjust back to playing balanced because your table image will not be useful after the session ends — the endgame unraveling.

Concrete situation: reputation on a poker tour

On the high-stakes live circuit, regulars play each other repeatedly across multiple events. Reputation builds over months and years. Phil Ivey's reputation for being capable of massive bluffs in huge pots means that opponents give him credit in situations where they would not give a random player credit. This reputation is "commitment type" reputation: even if Ivey plays balanced, the prior that he might be a maniac is always positive, and this prior gives his actions credibility.

The reputation depreciates if Ivey consistently plays conservatively for a long stretch (posterior updates downward toward "balanced" type). To maintain the reputation, occasional aggressive plays are needed — these are the "fight" moves that keep the posterior elevated. The cost of these plays (risking chips on thin bluffs) is the maintenance cost of reputation.

What this understanding buys you

Table image is a formal reputation mechanism. When poker players discuss "establishing a table image early," they are describing the chain-store reputation-building strategy: costly aggressive plays early in the session to raise the posterior that you are a maniac, which deters opponents from playing back at you. The formal theory tells you exactly how this works: the posterior evolves by Bayes' rule, the critical threshold depends on the payoff structure, and the cost of reputation building is bounded while the benefit scales with the length of the session.

Reputation is more valuable in long sessions. The chain-store theory says the gain from reputation scales with T (the session length). In a 3-hand online match, reputation is worthless. In a 500-hand cash session, it is very valuable. Adjust your early-session aggression accordingly: invest more in image building when the session will be long.

Reputation unravels near the end. The backward-induction logic kicks in: in the last few hands, your reputation has no future value, so the rational calculus changes. Opponents who are paying attention should expect you to stop investing in image near the end of a session. Pro players sometimes exploit this by making aggressive plays in the *very last hand* of a session — when nobody expects it — to leave a residual impression that persists into the next session.

Cross-Links

- **4.5 Signaling games** — reputation is multi-period signaling.
- **4.6 Sequential equilibrium** — the solution concept for reputation games.
- **5.1 Finitely repeated games** — the paradox that reputation resolves.
- **5.3 Trigger strategies** — an alternative cooperation mechanism (explicit threats vs implicit reputation).
- **5.6 Cooperation in repeated Prisoner's Dilemma** — reputation interacts with cooperation.

- **9.5 Convergence to equilibrium** — reputation dynamics as a learning process.
- **10.4 Commitment and credibility** — reputation as a substitute for formal commitment devices.

Structural Compression

Invariant: In a finitely repeated game with even a small probability of a commitment type, the rational player can build a reputation by mimicking the commitment type, sustaining near-optimal payoffs for the majority of the game horizon; reputation unravels only in a bounded-length endgame, independent of the total horizon length.

Mechanism: Bayesian posterior updating on type, driven by observed play; the rational type fights to maintain the posterior above the entry-deterrence threshold; accommodation reveals the type; the number of revealing periods is bounded by the logarithm of the initial prior (how many Bayes updates are needed to move the posterior from its initial value to the critical threshold), giving the result its striking insensitivity to T.

Cross-System Echo: Reputation theory is the game-theoretic formalization of the broader phenomenon of **signal consistency** in biology, economics, and social life. The handicap principle in evolutionary biology (Zahavi 1975) says costly displays signal genetic quality; corporate branding theory says costly advertising signals product quality (Nelson 1974, Milgrom–Roberts 1986); diplomatic deterrence theory says costly military posturing signals willingness to fight. In each case, the mechanism is: costly action + Bayesian updating = reputation building. The chain-store game is the minimal formal model that captures this mechanism, and the KMWR theorem shows it works even when the prior on the "tough" type is arbitrarily small.

Exercises

1. In the chain-store game with $T = 10$, $p = 0.1$, $M = 5, D = 3, F = 1, d = 2, c = 3$, compute the critical posterior μ^* and trace the posterior dynamics if the rational type fights in the first two periods.
2. Show that the number of periods in which the rational type accommodates (revealing their type) in the sequential equilibrium of the chain-store game is bounded by a constant independent of T. (Hint: each accommodation reduces the posterior to 0; the number of such periods is bounded by the "reputation stock.")
3. In the Fudenberg–Levine framework, suppose the commitment type plays tit-for-tat against short-run opponents. What is the long-run player's reputation payoff guarantee?
4. Construct a reputation game with **two-sided** incomplete information (both players have unknown types). Characterize how the equilibrium differs from the one-sided case.
5. Show that the chain-store reputation mechanism fails when monitoring is very noisy: if the entrant cannot reliably observe whether the incumbent fought or accommodated, the posterior does not update sharply, and reputation builds slowly or not at all.
6. In a poker context, estimate the "critical posterior" μ^* for a specific river bet-call scenario. How high does the opponent's belief that you are a maniac need to be for them to fold a bluff-catcher?
7. **Argue, structurally,** why reputation theory resolves the chain-store paradox in a way that is both rigorous and realistic. What makes the resolution robust (works even for tiny p), and what makes it fragile (what perturbations can destroy it)?

Game Theory · Module 5 — Repeated Games · Node 5.5
Concepts: bayesian-updating, incomplete-information, repeated-interaction, bayesian-nash-equilibrium

Prerequisites: 4.5, 5.1, 5.2
Enables: 5.6, 9.5

hbar.university — own.your.cognition