Chapter 9

Repeated Games and the Emergence of Cooperation

The Prisoner's Dilemma predicts defection—yet cooperation among self-interested parties is everywhere: cartels hold, merchants extend credit, countries honor treaties, strangers leave tips at restaurants they'll revisit. The unifying explanation: *the interaction is repeated*, so today's betrayal can be punished tomorrow.

9.1 Finitely repeated games

Let G be a *stage game*, repeated T times with observed history and payoffs summed (or discounted).

Proposition 9.1: Unraveling

If the stage game G has a *unique* Nash equilibrium, then the finitely repeated game has a unique subgame perfect equilibrium: the stage equilibrium is played in every period regardless of history.

Proof. In the last period T, no future exists, so play must be the stage Nash equilibrium in every subgame. Then period $T-1$ actions cannot influence period-T outcomes, so period $T-1$ must also be stage-Nash; induct backwards. □

So the finitely repeated Prisoner's Dilemma yields defection throughout—backward induction's cold logic again (compare the Centipede, Example 6.2). Experiments show cooperation until near the end, collapsing in the final rounds. When the stage game has *multiple* equilibria, however, even finite repetition can support non-equilibrium stage outcomes early on (rewarding good behavior by selecting the better continuation equilibrium).

9.2 Infinitely repeated games

Let G be played in periods $t = 0, 1, 2, \ldots$ with common discount factor $\delta \in (0, 1)$; player i's repeated-game payoff is the normalized sum $(1-\delta)\sum_t \delta^t u_i(a^t)$. Interpret δ either as time preference or as the probability the relationship continues another period—so "infinite" really means *indefinite*: no commonly known final date from which unraveling can start.

Definition 9.1: Trigger strategies

In the repeated Prisoner's Dilemma, *grim trigger* plays C initially and after any history of mutual cooperation, and plays D forever after any defection. *Tit-for-tat* plays C first and thereafter copies the opponent's previous move.

Proposition 9.2: Cooperation via grim trigger

In the Prisoner's Dilemma of Example 3.1 (payoffs $3,3$ / $0,4$ / $4,0$ / $1,1$), grim trigger by both players is a subgame perfect equilibrium iff

$$\underbrace{(1-\delta)\cdot 4+\delta\cdot 1}_{\text{defect today, punished forever}} \;\le\; \underbrace{3}_{\text{cooperate forever}} \quad\Longleftrightarrow\quad \delta \ge \tfrac{1}{3}.$$

Proof. On-path one-period deviation: gain $4-3$ today, lose $3-1$ every period after, discounted—the displayed inequality. Off path (after a defection), both play D forever, a stage Nash equilibrium repeated, hence immune to deviations. By the *one-shot deviation principle* (it suffices to check single-period deviations in discounted repeated games), the profile is SPE. □

Key Idea: The shadow of the future

Cooperation is sustainable when players are patient enough that the long-run value of the relationship outweighs the short-run gain from betrayal. The condition $\delta \ge \delta^{\min}$ formalizes folk wisdom: relationships sustain honesty where one-shot encounters do not; cartels are more stable in frequent-interaction markets; online platforms manufacture a "future" via reputation scores precisely to enable trade among strangers.

9.3 The Folk Theorem

Which payoffs can patient players sustain? Essentially all reasonable ones.

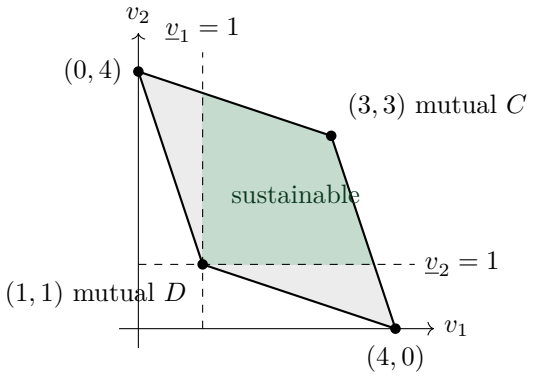
Definition 9.2: Minmax and feasible payoffs

Player i's *minmax payoff* is $\underline{v}_i = \min_{\alpha_{-i}} \max_{a_i} u_i(a_i, \alpha_{-i})$: the worst the others can inflict on a best-responding i. A payoff vector is *feasible* if it is a convex combination of stage-game payoff vectors, and *individually rational* if it gives each player strictly more than $\underline{v}_i$.

Theorem 9.3: Folk Theorem (Fudenberg–Maskin 1986)

Let v be any feasible and strictly individually rational payoff vector of the stage game. Then (under a mild dimensionality condition) for all δ sufficiently close to 1, the infinitely repeated game has a *subgame perfect* equilibrium with payoffs v.

For the Prisoner's Dilemma the theorem's set is the shaded region below:



Remark: Blessing and curse

The Folk Theorem is good news (cooperation is possible) and an embarrassment of riches (so is almost everything else, including unequal and inefficient outcomes). Repetition converts the theory from a prediction machine into a *possibility* theory; which equilibrium prevails is then a matter of history, norms, focality, and renegotiation—active research areas.

9.4 Cartels, price wars, and imperfect monitoring

Example 9.1: Collusion in Cournot competition

Two Cournot firms (Example 4.2) can sustain the monopoly output (each producing $\frac{a-c}{4}$, earning $\frac{(a-c)^2}{8}$ each) with grim trigger reversion to Cournot ($\frac{(a-c)^2}{9}$ each) iff $\delta \geq \delta^{\min} = \frac{9}{17}$. (The optimal single-period cheat against a colluding rival earns $\frac{9(a-c)^2}{64}$; plug into the deviation inequality.) Antitrust implications: collusion is easier with fewer firms, frequent interaction, transparency of rivals' prices—which is why publishing detailed transaction prices can be *anti*-competitive.

When rivals' actions are observed only through noisy signals (market price moves with demand shocks), punishments trigger *on the equilibrium path*: occasional price wars are not collusion breaking down but the very machinery that sustains it (Green–Porter 1984). The general theory of repeated games with imperfect monitoring (Abreu–Pearce–Stacchetti; Fudenberg–Levine–Maskin) is a highlight of modern economic theory.

9.5 Reputation

Example 9.2: Chain store paradox and reputation

An incumbent faces entrants in T markets sequentially (Example 6.1 repeated). Backward induction says: accommodate everywhere, so all enter—Selten's *chain store paradox*. Kreps–Milgrom–Roberts–Wilson (1982) resolution: if entrants assign even tiny probability ε to the incumbent being a "tough" type who enjoys fighting, the rational incumbent *mimics* the tough type by fighting early entrants, deterring entry for most of a long horizon. A drop of incomplete information plus repetition generates the reputation effects that plain backward induction destroys. (Tools for this analysis: Chapter 11.)

Historical Note: Axelrod's tournaments

In Robert Axelrod's famous computer tournaments (1980), strategies for the repeated Prisoner's Dilemma were submitted by game theorists and competed round-robin. The winner—both times—was the simplest entry: Anatol Rapoport's TIT-FOR-TAT. Axelrod distilled its virtues: it is *nice* (never defects first), *retaliatory* (punishes immediately), *forgiving* (returns to cooperation), and *clear*. Caveat: TFT is not subgame perfect, and in noisy environments two TFT players fall into echoing vendettas—more forgiving variants ("generous TFT", "win-stay lose-shift") do better under noise.

Exercises

Exercise 9.1: Threshold discount factors

For the Prisoner's Dilemma with payoffs $C, C \to (2, 2)$; $D, C \to (5, 0)$; $C, D \to (0, 5)$; $D, D \to (1, 1)$: find the minimum δ sustaining mutual cooperation via (a) grim trigger; (b) one-period punishment (defect once, then return to cooperation if the opponent also defected during the punishment). Explain why (b) requires a higher δ.

Exercise 9.2: Tit-for-tat is not subgame perfect

In the Prisoner's Dilemma of Example 3.1, show that mutual tit-for-tat fails the one-shot deviation test after the history in which player 1 defected once: alternating punishment cycles give lower payoffs than an immediate return to cooperation for some δ. Find the range of δ (if any) for which TFT survives the relevant one-shot deviations.

Exercise 9.3: Folk theorem geometry

For the Battle of the Sexes (Example 3.3), draw the feasible set, compute both players' minmax payoffs, and shade the payoffs sustainable by patient players. Can the players sustain the alternating scheme (Ballet in even periods, Football in odd)? Compute its payoff vector for $\delta \to 1$.

Exercise 9.4: Community enforcement

A large population is randomly matched each period to play the Prisoner's Dilemma with a new partner. Personal punishment is impossible (you never meet again), but each player carries a public label "good" or "bad": defecting against a good partner makes you bad. Design a strategy rule under which cooperation with good partners is an equilibrium for large δ, and discuss what happens if labels are sometimes wrong (interpret: gossip, credit scores, platform ratings).