Repeated Games

Mihai Manea

MIT

Repeated Games with Perfect Information

Repeated Games with Perfect Information

- normal-form stage game $G = (N, A, u)$
- players simultaneously play game G at time $t = 0, 1, \ldots$
- at each date t, players observe all past actions: $h^t = (a^0, \ldots, a^{t-1})$
- common discount factor $\delta \in (0, 1)$
- payoffs in the repeated game $RG(\delta)$ for $h = (a^0, a^1, \ldots)$: $U_i(h) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_i(a^t)$
- normalizing factor $1 - \delta$ ensures payoffs in $RG(\delta)$ and G are on same scale
- behavior strategy σ_i for $i \in N$ specifies $\sigma_i(h^t) \in \Delta(A_i)$ for every history h^t

Can check if σ constitutes an SPE using the single-deviation principle.

Minmax

Minmax payoff of player i: lowest payoff his opponents can hold him down to if he anticipates their actions,

$$\underline{v}_i = \min_{\alpha_{-i} \in \prod_{j \neq i} \Delta(A_j)} \left[\max_{a_i \in A_i} u_i(a_i, \alpha_{-i}) \right]$$

- m^i: *minmax profile for i*, an action profile (a_i, α_{-i}) that solves this minimization/maximization problem
- assumes *independent* mixing by i's opponents
- important to consider mixed, not just pure, actions for i's opponents: in the matching pennies game the minmax when only pure actions are allowed for the opponent is 1, while the actual minmax, involving mixed strategies, is 0

Equilibrium Payoff Bounds

In any SPE—in fact, any Nash equilibrium—i's obtains at least his minmax payoff: can myopically best-respond to opponents' actions (known in equilibrium) in each period separately. Not true if players condition actions on correlated private information!

A payoff vector $v \in \mathbb{R}^N$ is *individually rational* if $v_i \geq \underline{v}_i$ for each $i \in N$, and *strictly individually rational* if the inequality is strict for all i.

Feasible Payoffs

Set of *feasible payoffs*: convex hull of $\{u(a) \mid a \in A\}$. For a *common discount factor* δ, normalized payoffs in $RG(\delta)$ belong to the feasible set.

Set of feasible payoffs includes payoffs not obtainable in the stage game using mixed strategies. . . some payoffs require correlation among players' actions (e.g., battle of the sexes).

Public randomization device produces a publicly observed signal $\omega^t \in [0, 1]$, uniformly distributed and independent across periods. Players can condition their actions on the signal (formally, part of history).

Public randomization provides a convenient way to convexify the set of possible (equilibrium) payoff vectors: given strategies generating payoffs v and v', any convex combination can be realized by playing the strategy generating v conditional on some first-period realizations of the device and v' otherwise.

Nash Threat Folk Theorem

Theorem 1 (Friedman 1971)

If e is the payoff vector of some Nash equilibrium of G and v is a feasible payoff vector with $v_i > e_i$ for each i, then for all sufficiently high δ, $RG(\delta)$ has SPE with payoffs v.

Proof.

Specify that players play an action profile that yields payoffs v (using the public randomization device to correlate actions if necessary), and revert to the static Nash equilibrium permanently if anyone has ever deviated. When δ is high enough, the threat of reverting to Nash is severe enough to deter anyone from deviating. $\square$

If there is a Nash equilibrium that gives everyone their minmax payoff (e.g., prisoner's dilemma), then every strictly individually rational and feasible payoff vector is obtainable in SPE.

General Folk Theorem

Minmax strategies often do not constitute static Nash equilibria. To construct SPEs in which i obtains a payoff close to $\underline{v}_i$, need to threaten to punish i for deviations with even lower continuation payoffs. Holding i's payoff down to $\underline{v}_i$ may require other players to suffer while implementing the punishment. Need to provide incentives for the punishers... impossible if punisher and deviator have indetical payoffs.

Theorem 2 (Fudenberg and Maskin 1986)

Suppose the set of feasible payoffs has full dimension $|N|$. Then for any feasible and strictly individually rational payoff vector v, there exists $\underline{\delta}$ such that whenever $\delta > \underline{\delta}$, there exists an SPE of $RG(\delta)$ with payoffs v.

Abreu, Dutta, and Smith (1994) relax the full-dimensionality condition: only need that no two players have the same payoff function (equivalent under affine transformation).

Proof Elements

- Assume first that i's minmax action profile m^i is pure.
- Consider an action profile a for which $u(a) = v$ (or a distribution over actions that achieves v using public randomization).
- By full-dimensionality, there exists v' in the feasible individually rational set with $\underline{v}_i < v'_i < v_i$ for each i.
- Let w^i be v' with ε added to each player's payoff except for i; for small ε, w^i is a feasible payoff.

Equilibrium Regimes

- Phase I: play a as long as there are no deviations. If i deviates, switch to II_i.
- Phase II_i: play m^i for T periods. If player j deviates, switch to II_j. If there are no deviations, play switches to III_i after T periods.
  - If several players deviate simultaneously, arbitrarily choose a j among them.
  - If m^i is a pure strategy profile, it is clear what it means for j to deviate. If it requires mixing. . . discuss at end of the proof.
  - T independent of δ (to be determined).
- Phase III_i: play the action profile leading to payoffs w^i forever. If j deviates, go to II_j.

SPE? Use the single-shot deviation principle: calculate player i's payoff from complying with prescribed strategies and check for profitable deviations at every stage of each phase.

Deviations from I and II

Player i's incentives

- Phase I: deviating yields at most $(1-\delta)M + \delta(1-\delta^T)\underline{v}_i + \delta^{T+1} v_i'$, where M is an upper bound on i's feasible payoffs, and complying yields v_i. For fixed T, if δ is sufficiently close to 1, complying produces a higher payoff than deviating, since $v_i' < v_i$.
- Phase II_i: suppose there are $T' \leq T$ remaining periods in this phase. Then complying gives i a payoff of $(1-\delta^{T'})\underline{v}_i + \delta^{T'} v_i'$, whereas deviating can't help in the current period since i is being minmaxed and leads to T more periods of punishment, for a total payoff of at most $(1-\delta^{T+1})\underline{v}_i + \delta^{T+1} v_i'$. Thus deviating is worse than complying.
- Phase II_j: with T' remaining periods, i gets $(1-\delta^{T'})u_i(m^j) + \delta^{T'}(v_i' + \varepsilon)$ from complying and at most $(1-\delta)M + (\delta-\delta^{T+1})\underline{v}_i + \delta^{T+1} v_i'$ from deviating. For high δ, complying is preferred.

Deviations from *III*

Player *i*'s incentives

- Phase III_i: determines choice of T. By following the prescribed strategies, i receives v_i' in every period. A (one-shot) deviation leaves i with at most $(1-\delta)M + \delta(1-\delta^T)\underline{v}_i + \delta^{T+1}v_i'$. Rearranging, i compares between $(\delta + \delta^2 + \ldots + \delta^T)(v_i' - \underline{v}_i)$ and $M - v_i'$. For any $\underline{\delta} \in (0,1)$, $\exists T$ s.t. former term is grater than latter for $\delta > \underline{\delta}$.
- Phase III_j: Player i obtains $v_i' + \varepsilon$ forever if he complies with the prescribed strategies. A deviation by i triggers phase II_i, which yields at most $(1-\delta)M + \delta(1-\delta^T)\underline{v}_i + \delta^{T+1}v_i'$ for i. Again, for sufficiently large δ, complying is preferred.

Mixed Minmax

What if minmax strategies are mixed? Punishers may not be indifferent between the actions in the support... need to provide incentives for mixing in phase *II*.

Change phase *III* strategies so that during phase II_j player i is indifferent among all possible sequences of T realizations of his prescribed mixed action under m^j. Make the reward ε_i of phase III_j dependent on the history of phase II_j play.

Dispensing with Public Randomization

Sorin (1986) shows that for high δ we can obtain any convex combination of stage game payoffs as a normalized discounted value of a deterministic path $(u(a^t))$... "time averaging"

Fudenberg and Maskin (1991): can dispense of the public randomization device for high δ, while *preserving incentives*, by appropriate choice of which periods to play each pure action profile involved in any given convex combination. Idea is to stay within ε^2 of target payoffs at all stages.

Repeated Games with Fixed Discount Factor

- Folk theorem concerned with limit $\delta \to 1$
  - many payoffs possible in SPE
  - elaborate hierarchies of punishments needed
- Equilibrium outcomes for fixed $\delta < 1$?
- Abreu (1988): equilibrium strategies can be enforced by using worst punishment for every deviator
- Is there a worst possible punishment?

Theorem 3 (Abreu 1988)

Assume that

1. *action sets in stage game are compact subsets of Euclidean spaces;*
2. *payoffs in stage game are continuous in actions;*
3. *there exists a pure-strategy SPE.*

Then among all pure-strategy SPEs, there is one that is worst for player i.

Proof

- $a = (a_t)_{t \geq 0}$: play path
- $U_i(a)$: i's payoff in repeated game for path a
- equilibrium play path: play path of some pure-strategy SPE
- $y(i) = \inf\{U_i(a) | a \text{ equilibrium play path}\}$
- $y(i)$ well-defined, set of equilibrium play paths is nonempty
- $\exists (a^{i,k})_{k \geq 0}$: sequence of equilibrium play paths s.t. $U_i(a^{i,k}) \to y(i)$
- $\prod_{t \geq 0} A$ sequentially compact, convergent subsequence $a^{i,k} \to a^{i,\infty}$ (product topology)
- $U_i(a^{i,\infty}) = y(i)$ by continuity of U_i
- Prove that $a^{i,\infty}$ is an equilibrium play path. Candidate SPE
  - in regime i, players follow strategies $a^{i,\infty}$
  - play starts in regime i
  - deviation by player j from current regime leads to regime j

Proof

One-shot deviation principle

- Consider deviation by player j from stage τ of regime i to action $\widehat{a}_j$
- Show that j's one-shot deviation is not profitable,

$$(1-\delta)\sum_{t=0}^{\infty}\delta^t u_j(a^{i,\infty}(\tau+t)) \geq (1-\delta)u_j(\widehat{a}_j, a^{i,\infty}_{-j}(\tau)) + \delta y(j).$$

- For each k, there is some SPE whose play path is $a^{i,k}$
  - j does not have incentives to deviate to $\widehat{a}_j$ at τ
  - j's continuation payoff is at least $y(j)$ following any deviation (definition of $y(j)$)
- For all k,

$$(1-\delta)\sum_{t=0}^{\infty}\delta^t u_j(a^{i,k}(\tau+t)) \geq (1-\delta)u_j(\widehat{a}_j, a^{i,k}_{-j}(\tau)) + \delta y(j).$$

- Taking limit $k \to \infty$, we obtain desired inequality.

Optimal Penal Codes

- Optimal penal code for i: SPE giving i lowest possible payoff
- Showed existence of optimal penal codes for pure-strategy SPEs
- If stage game is finite, pure-strategy SPEs may not exist.
- Optimal penal codes for mixed-strategy SPEs?

Theorem 4 (Fudenberg and Levine 1983)

Assume that the stage game is finite. Then the set of SPE strategies and payoffs are nonempty and compact. In particular, among all SPEs, there is one that is worst for player i.

Conclusion extends to multistage games with observable actions that have a finite set of actions at every stage and are continuous at infinity. Fudenberg and Levine used ε-perfect equilibria of finite-horizon truncations.

Proof

- $\prod_{h_t} \prod_i \Delta(A_i)$: set of mixed strategies (with product topology)
- $\prod_{h_t} \prod_i \Delta(A_i)$ is compact (Tychonoff's theorem)
- Payoffs in repeated game are continuous in strategies
  - take a sequence of strategies (σ^n) converging to σ as $n \to \infty$
  - for any date t, the distribution over histories/actions induced by σ_n at t also converges to the one induced by σ
  - expected payoffs at t under σ_n converge to those under σ
- Since payoffs are continuous, the set of SPEs is closed.
- Since set of strategies is compact, set of SPEs is compact (closed subsets of compact sets are compact).
- Since payoffs are continuous in strategies and set of SPEs is compact, the set of SPE payoffs is also compact.

Equilibrium Outcomes and Optimal Penal Codes

Theorem 5 (Abreu 1988)

Any SPE distribution over play paths can be generated by an SPE enforced by optimal penal codes off path, i.e., when i is the first deviator, continuation play follows the optimal penal code for i.

Fix SPE $\hat{\sigma}$. Modify $\hat{\sigma}$ by replacing play off-path by the optimal penal code for i when i is the first deviator. Modified strategy profile is an SPE. Sufficient to check that on-path one-shot deviations are not profitable.

Generalized Cournot Oligopoly

Abreu (1986): application of Abreu (1988) to symmetric settings

- set of actions: $[0, \infty)$
- $\exists M$ s.t. actions above M are never rational
- payoffs continuous and bounded from above
- $C1$: $u_i(a, \ldots, a)$ is quasi-concave and decreases unboundedly in a
- $C2$: $\max_{a_i \in A_i} u_i(a_i, a \ldots, a)$ is decreasing in a

Strongly symmetric equilibria: equilibria in which all players behave identically at every history, including asymmetric histories.

There is a strongly symmetric equilibrium that is worst for all players.

There is also a strongly symmetric equilibrium that is best for everyone.

Extremal Equilibria

Theorem 6

Let e^* and e_* denote the highest and lowest payoff per player in a pure-strategy strongly symmetric equilibrium.

- *The payoff e_* can be attained in an equilibrium with strongly symmetric strategies of the following form: "Begin in phase A, where players choose an action a_* that satisfies*

$$(1-\delta)u(a_*,\ldots,a_*)+\delta e^* = e_*.$$

  If there are no deviations, switch to an equilibrium with payoff e^* (phase B). Otherwise, continue in phase A."

- *Phase B: the payoff e^* can be attained with strategies that play a constant action a^* as long as there are no deviations and switch to the worst strongly symmetric equilibrium (phase A) if there are any deviations.*

Proof of First Part

Existence of a_*

- $\hat{\sigma}$: strongly symmetric equilibrium with payoff e_* and period 1 action a
- continuation payoffs under $\hat{\sigma}$ cannot be more than e^*, so

$$u_i(a, \ldots, a) \geq (-\delta e^* + e_*)/(1 - \delta).$$

- By $C1$, there is $a_* \geq a$ s.t. $u(a_*, \ldots, a_*) = (-\delta e^* + e_*)/(1 - \delta)$.

Let σ_* denote the strategies constructed for phase A. No profitable deviation in first period of $\hat{\sigma}$ $\Rightarrow$ no profitable deviation in phase A of σ_*

- by $C2$ and $a_* \geq a$, short-run gain from deviating in phase A is not more than that in the first period of $\hat{\sigma}$
- punishment for deviating in phase A is worst possible

Phase B: σ_* subgame perfect by definition

Discussion

- $C1$: punishments can be made arbitrarily bad.
- Good equilibrium can be sustained by one-shot punishments.
- Extremal equilibria described by two numbers a_* and a^*.
- Stick and carrot: first-period deviation leads to one period of punishment with $(a_*, \ldots, a_*)$ and playing $(a^*, \ldots, a^*)$ thereafter.
- (a_*, a^*): highest and lowest action s.t.

$$\begin{aligned}
\max_{a_i \in A_i} u_i(a_i, a_{*-i}) - u_i(a_*, \ldots, a_*) &= \delta(u_i(a^*, \ldots, a^*) - u_i(a_*, \ldots, a_*)) \\
\max_{a_i \in A_i} u_i(a_i, a^*_{-i}) - u_i(a^*, \ldots, a^*) &= \delta(u_i(a^*, \ldots, a^*) - u_i(a_*, \ldots, a_*))
\end{aligned}$$

- Typically best outcome is better (and worst is worse) than static Nash.

Repeated Games with Imperfect Public Monitoring

Imperfect Public Monitoring

Players only observe noisy signal of others' past actions

- y: public signal
- $\pi_y(a)$: distribution of y conditional on action profile a
- $r_i(a_i, y)$: i's payoff
- $u_i(a) = \sum_{y \in Y} r_i(a_i, y)\pi_y(a)$: i's expected payoff

Green and Porter's (1984) collusion model

- Players: firms in a cartel setting production quantities
- Public signal: market price, stochastic function of quantities (unobserved demand shock each period)
- Payoff: product of price and own quantity
- Firms are trying to keep quantities low and prices high.
- But low prices may come from deviations or demand shocks.
- How to detect and punish deviations?

Green and Porter (1984)

Green and Porter focus on threshold equilibria

- Play collusive action profile $\hat{a}$ for a while.
- If $y > \hat{y}$, revert to static Nash for T periods, then return to collusion.

Equilibrium values (after normalizing static Nash payoffs to 0)

$$\hat{v}_i = (1-\delta)u_i(\hat{a}) + \delta\lambda(\hat{a})\hat{v}_i + \delta(1-\lambda(\hat{a}))\delta^T\hat{v}_i,$$

where $\lambda(a) = P(y > \hat{y}|a)$. Rearranging,

$$\hat{v}_i = \frac{(1-\delta)u_i(\hat{a})}{1-\delta\lambda(\hat{a})-\delta^{T+1}(1-\lambda(\hat{a}))}.$$

No player wants to deviate in the collusive phase: for all a_i'

$$\begin{aligned} u_i(a_i', \hat{a}_{-i}) - u_i(\hat{a}) &\leq \frac{\delta(1-\delta^T)(\lambda(\hat{a})-\lambda(a_i',\hat{a}_{-i}))\hat{v}_i}{1-\delta} \\ &= \frac{\delta(1-\delta^T)(\lambda(\hat{a})-\lambda(a_i',\hat{a}_{-i}))u_i(\hat{a})}{1-\delta\lambda(\hat{a})-\delta^{T+1}(1-\lambda(\hat{a}))}. \end{aligned}$$

Discussion

Incentive constraints compare

- short-term incentives to deviate
- relative probability of triggering punishment by deviating
- severity of punishment

Possible to sustain payoffs above static Nash for high δ: incentive constraints hold for $T = \infty$ and $\hat{a}$ just below static Nash, with low $\hat{y}$ and some bounds on the derivative of λ.

Green and Porter did not identify the best possible equilibria. There can be worse punishments than static Nash.

Need a more general theory to find better equilibria.

Abreu, Pearce, and Stacchetti (APS) (1990)

- A_i and Y finite
- $\pi_y(\alpha)$: distribution of y given mixed action profile α
- $h_i^t = (y^0, \ldots, y^{t-1}; a_i^0, \ldots, a_i^{t-1})$: private history
- $h^t = (y^0, \ldots, y^{t-1})$: public history

APS assume continuum of signals and restrict attention to pure strategies.

A strategy for player i is a public strategy if depends only public history.

Lemma 1

Every pure strategy σ_i is equivalent to a public strategy σ_i'.

Proof.

Define σ_i' on lenght-t histories by induction: for each $s < t$

$$\sigma_i'(y^0, \ldots, y^{t-1}) = \sigma_i(y^0, \ldots, y^{t-1}; a_i^0, \ldots, a_i^{t-1}),$$

where $a_i^s = \sigma_i'(y^0, \ldots, y^{s-1})$. σ_i' is equivalent to σ_i: they differ only at "off-path" histories reachable only by deviations of player i. □

Public Perfect Equilibirum (PPE)

Lemma shows that public strategies are without loss if attention is restricted to pure strategies, which is a nontrivial restriction.

Fudenberg, Levine, and Maskin (1994): restrict attention to public (but potentially mixed) strategies.

Lemma 2

If opponents use public strategies, then i has a best response in public strategies.

Proof.

At every date, i always what the other players will play, since their actions depend only on the public history; hence i can just play a best response to their anticipated future play, which does not depend on i's private history of past actions. □

Discussion of PPE

Perfect public equilibrium: public strategies σ s.t., at every public history h^t, the strategies $\sigma_i|_{h^t}$ form a Nash equilibrium of the continuation game.

Allows for deviations to non-public strategies, but such deviations are irrelevant by Lemma 2.

PPE adapts SPE to this setting (in general, no subgames).

Set of PPE is stationary: same at every history

Set of SE is Not Stationary

A player may want to condition his play in one period on the realization of his mixing in a previous period. Correlation across periods can be self-sustaining.

- Assume i and j both mixed at a previous period.
- The signal in that period informs i about the realization of j's mixing.
- Hence it is informative about what j will do in the current period, affecting i's current best response.
- Some third player k may be unable to infer what j will do in the current period, since he does not know what i played in the earlier period.
- Different players can have different beliefs about play at t.
- Stationarity is lost.

Enforceability

$(\alpha, v) \in \prod_i \Delta(A_i) \times \mathbb{R}^n$ enforceable w.r.t. $W \subseteq \mathbb{R}^n$ if there is $w : Y \to W$ s.t.

$$v_i = (1 - \delta)u_i(\alpha) + \delta \sum_{y \in Y} \pi_y(\alpha) w_i(y)$$

and for all i and $a_i' \in A_i$,

$$v_i \geq (1 - \delta)u_i(a_i', \alpha_{-i}) + \delta \sum_{y \in Y} \pi_y(a_i', \alpha_{-i}) w_i(y).$$

Incentive-compatible for each player to play according to α in the present period if continuation payoffs are given by w; expected present payoffs are v.

$B(W)$: set of v enforceable w.r.t. W for some α

Theorem 7

Let E be the set of PPE payoff vectors. Then $E = B(E)$.

Proof

$E \subseteq B(E)$

- Suppose $v \in E$, generated by PPE strategies σ.
- Let $\alpha_i = \sigma_i(\emptyset)$ and $w_i(y)$ be continuation payoff under σ given y.
- $w(y) \in E$ since play in subsequent periods forms a PPE.
- $v \in B(E)$ since (α, v) is enforced by w on E.

$B(E) \subseteq E$

- Suppose $v \in B(E)$ with (α, v) enforced by w on E.
- Define public strategies σ
  - Play α in the first period.
  - After y, follow PPE with payoffs $w(y)$.
- σ is a PPE by the one-shot deviation principle.
- $v \in E$ since v is the payoff vector generated by σ.

Self-Generating Sets

Definition 1

$W \subseteq \mathbb{R}^n$ is self-generating if $W \subseteq B(W)$.

E is self-generating.

Theorem 8

If W is a bounded self-generating set, then $W \subseteq E$.

Proof

Given $v \in W$, we construct a PPE with payoffs v recursively by specifying current play and continuation payoffs in W for every public strategy

- Base case (before game starts): continuation payoffs v
- Assume we have specified
  - Play for periods $0, \ldots, t-1$
  - Continuation payoffs in W for each history $y^0, \ldots, y^{t-1}$
- Fix some history $y^0, \ldots, y^{t-1}$ and continuation payoffs $v' \in W$.
- $W \subseteq B(W)$: pick α and $w : Y \to W$ s.t. (α, v') enforced by w
  - Specify play α at history $y^0, \ldots, y^{t-1}$.
  - Let $w(y)$ be continuation payoffs if signal y is observed next.

Do strategies implement desired payoffs?

$$v - (1-\delta)\sum_{t=0}^{T} \delta^t E[u(\alpha^t)] = \delta^{T+1} E[w(y^T)],$$

Right-hand side goes to zero as $T \to \infty$ (W bounded).

Constructed strategies form a PPE by the one-shot deviation principle.

Monotonicity of PPE Payoffs in δ

$E(\delta)$: set of PPE payoffs when discount factor is δ

Theorem 9

Assume $E(\delta)$ is convex (e.g., public randomization device). If $\delta_1 < \delta_2$, then $E(\delta_1) \subseteq E(\delta_2)$.

Proof.

By previous theorem, enough to show that $E(\delta_1) \subseteq B(E(\delta_1), \delta_2)$. Take $v \in E(\delta_1) = B(E(\delta_1), \delta_1)$. Find α and w that enforce v for discount factor δ_1. By replacing w by a suitable convex combination of w and (the constant function) v, we can enforce (α, v) for δ_2. $\square$

Characterization of PPE Payoffs

Properties of B operator

- If W is compact, then $B(W)$ is compact.
- B is monotone: if $W \subseteq W'$, then $B(W) \subseteq B(W')$.
- If W is nonempty, then $B(W)$ is nonempty.
  - If α is a Nash equilibrium of the stage game, $w : Y \to W$ a constant function, and v the resulting payoffs, then $v \in B(W)$.

V: set of all feasible payoffs; nonempty compact

$(B^k(V))_{k \geq 0}$: $B^0(V) = V$ and $B^k(V) = B(B^{k-1}(V))$ for $k = 1, 2, \ldots$

$B^\infty(V) = \cap_k B^k(V)$: intersection of a decresing sequence of non-empty compact sets, hence nonempty and compact

Theorem 10

$E = B^\infty(V)$. *In particular, E is nonempty and compact.*

Proof

Prove that $B^\infty(V) \subseteq E$ by showing that $B^\infty(V)$ is self-generating

- Suppose $v \in B^\infty(V)$.
- $\exists(\alpha^k, w^k)_{k\geq 1}$: (α_k, v) enforced by $w^k : Y \to B^{k-1}(V)$
- $(\alpha^\infty, w^\infty)$: some limit point of $(\alpha^k, w^k)_{k\geq 0}$ (exists by compactness)
- $w^\infty(y) \in B^\infty(V)$ since for all k, $w^\infty(y)$ limit point of closed $B^k(V)$
- By continuity, (α^∞, v) enforced by $w^\infty : Y \to B^\infty(V)$

$E \subseteq B^\infty(V)$

- $E \subseteq V$ implies $E = B(E) \subseteq B(V)$
- By induction, $E \subseteq B^k(V)$ for all k.

Bang-Bang Property

$ext(W)$: set of extreme points of W

$w : Y \to W$ has the bang-bang property if $w(y) \in ext(W)$ for each y

APS: finite action spaces and continuous signals (common support)

APS: if (α, v) enforceable on compact W, then enforceable on $ext(W)$

Corollary: every vector in E can be achieved as the vector of payoffs from a PPE s.t. the vector of continuation payoffs at every history lies in $ext(E)$

Folk Theorem for Imperfect Public Monitoring

Fudenberg, Levine, and Maskin (1994) approximate set of feasible individually rational payoffs with "smooth" convex sets W that are self-generating for δ high enough.

Key assumptions on informativeness of public signals about actions

- Individual full-rank: given α_{-i}, the different signal distributions generated by varying i's pure actions a_i are linearly independent... If α_i and α_i' generate the same distribution over signals given α_{-i} and $u_i(\alpha_i', \alpha_{-i}) > u_i(\alpha)$, then it is impossible to enforce α in equilibrium.
- Pairwise full rank: deviations by player i are "statistically distinguishable" from deviations by player j
  - For any α, build two matrices whose rows represent signal distributions from (a_i, α_{-i}) and (a_j, α_{-j}) as a_i and a_j varies, resp.
  - The stacked matrix has rank $|A_i| + |A_j| - 1$.
  - Intuitively, players need to know who to punish.

Locally Self-Generating Sets

W locally self-generating: $\forall v \in W$, $\exists$ open neighborhood U and $\underline{\delta} < 1$ s.t. $U \cap W \subseteq B(W)$ when $\delta > \underline{\delta}$.

If W compact and convex, then local self-generation implies self-generation for δ high enough.

We can use full-rank conditions to show that W is locally self-generating.

Focus on boundary points W... interior points can then be achieved with public randomization devices.

“Proof” of Local Self-Generation

- By pairwise full-rank, we can “transfer” continuation payoffs between two players in any desired ratio, depending on the signal. . . pairwise hyperplanes.
- We can “decompose” regular tangent hyperplanes using pairwise hyperplanes.
- Need individual full rank to enforce payoffs on coordinate hyperplanes
- Both individual and pairwise full rank needed only for a subset of action profiles. . . “mixing in” any such profile with small probability generates a dense set of profiles with full rank.
- Then any payoff vector v on the boundary of W is enforceable by continuation payoffs that lie below the tangent hyperplane to W at v.
- Using “discount factor transformation” (Theorem 9), the continuation payoffs that enforce v contract toward v at rate $1 - \delta$ as $\delta \to 1$.
- Since W is smooth, a translate of the tangent hyperplane at distance of order $1 - \delta$ has “diameter” of order $\sqrt{1 - \delta}$, so continuation payoffs lie inside W for high δ.

Changing the Information Structure with Time Period

Suppose time is continuous.

Players can only update actions at dates $t, 2t, \ldots$.

r: discount rate; $\delta = e^{-rt}$: discount factor

$\delta \to 1$ because

- $r \to 0$: players become patient
- $t \to 0$: periods become short

Abreu, Milgrom, and Pearce (1991): limits $r \to 0$ and $t \to 0$ may lead to distinct predictions if quality of public signals deteriorates as $t \to 0$.

Setup

Stage game: partnership game (prisoners' dilemma)

- c: common payoff when both cooperate
- $c + g$: payoff from defecting when other cooperates
- 0: payoff if both defect

AMP: Poisson signal (number of "successes") with intensity λ if both cooperate, μ if only one does.

For small t, prob. of observing more than one success is of order t^2. To simplify analysis, assume binary signal—0 or 1 successes—with prob. $e^{-\theta t}$ and $1 - e^{-\theta t}$, resp., for $\theta \in \{\lambda, \mu\}$. Signals represent "good news," $\lambda > \mu$.

Pure strategy strongly symmetric equilibria with public randomization

- 0: worst equilibrium payoff (minimax, static Nash)
- v^*: best equilibrium payoff
- $\alpha(i)$ $(1 - \alpha(i))$: prob. of reverting to static Nash (playing best equilibrium) if signal $i \in \{0, 1\}$ is observed

PPE Constraints

Best equilibrium should specify cooperation in the first period.

$$\begin{aligned} v^* &= (1-e^{-rt})c + e^{-rt}\left(e^{-\lambda t}(1-\alpha(0)) + (1-e^{-\lambda t})(1-\alpha(1))\right)v^* \\ &\geq (1-e^{-rt})(c+g) + e^{-rt}\left(e^{-\mu t}(1-\alpha(0)) + (1-e^{-\mu t})(1-\alpha(1))\right)v^* \end{aligned}$$

Solve for v^*,

$$v^* = \frac{(1-e^{-rt})c}{1-e^{-rt}\left(1-\alpha(1)-e^{-\lambda t}(\alpha(0)-\alpha(1))\right)}.$$

Incentive constraint

$$(1-e^{-rt})g \leq e^{-rt}(e^{-\mu t}-e^{-\lambda t})(\alpha(0)-\alpha(1))v^*$$

simplifies to

$$g \leq \frac{ce^{-rt}(e^{-\mu t}-e^{-\lambda t})(\alpha(0)-\alpha(1))}{1-e^{-rt}\left(1-\alpha(1)-e^{-\lambda t}(\alpha(0)-\alpha(1))\right)}.$$

v^* is decreasing in $\alpha(1)$ and the incentive constraint is relaxed by decreasing $\alpha(1) \Rightarrow \alpha(1) = 0$.

$t \to 0$ versus $r \to 0$

$\alpha(1) = 0 \Rightarrow$

$$v^* = \frac{(1 - e^{-rt})c}{1 - e^{-rt}(1 - e^{-\lambda t}\alpha(0))} \quad \text{and} \quad g/c \leq \frac{e^{-rt}(e^{-\mu t} - e^{-\lambda t})\alpha(0)}{1 - e^{-rt}(1 - e^{-\lambda t}\alpha(0))}$$

Possible to satisfy the inequality for $\alpha(0) \leq 1$ only if

$$g/c \leq \frac{e^{-rt}(e^{-\mu t} - e^{-\lambda t})}{1 - e^{-rt}(1 - e^{-\lambda t})} \leq e^{(\lambda - \mu)t} - 1.$$

$t \to 0$: we cannot do better than static Nash since $e^{(\lambda-\mu)t} \to 1$

- $e^{(\lambda-\mu)t}$: likelihood ratio for no success
- As $t \to 0$, almost certainly no success is observed.
- Informativeness of public signal is poor.

$r \to 0$: we can do better than static Nash for some values of c, g, t

- Incentive constraint becomes $g/c \leq e^{(\lambda-\mu)t} - 1$.
- "Optimal" $\alpha(0)$: $v^* \to c - \frac{g}{e^{(\lambda-\mu)t} - 1} > 0$