Lectures 5–6: Repeated Games with Public Monitoring

Alexander Wolitzky

MIT

14.126, Spring 2024

Repeated Games Overview

Repeated games are economists' basic model of long-run relationships. (Dynamic multi-agent moral hazard.)

Baseline model:

- Fixed set of players repeatedly play a fixed game (the **stage game**).
- Players care about payoffs in each period. Typically maximize **expected discounted payoffs**, discount factor $\delta \in (0, 1)$.
  - Player i's repeated-game payoff is $(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i (a_t)$.
- No instrinsic link between periods, but players can condition current play on information about past actions.

Key theme: repeated interaction allows new equilibrium outcomes, because players can use information about past play to "reward" or "punish" opponents.

- Repeated games closely tied to general study of dynamic incentives throughout economics.

Repeated Games Overview (cntd.)

Supporting non-static Nash outcomes requires that:

- Players are patient (high δ), so the "shadow of the future" carries weight.
- There is enough information, so future rewards/punishments depend on actions.
- Rewards/punishments are credible: play is sequentially rational.
  - Typical solution concept: sequential equilibrium or similar.

Analysis often focuses on when these conditions hold (so "cooperation" is possible, or a "folk theorem" holds) or fail to hold (so cooperation is impossible/folk theorem fails).

- Analysis usually characterizes the (large) set of equilibria. Sometimes implicitly assume players coordinate on a cooperative equilibrium when one exists.

Repeated Games Overview (cntd.)

When cooperation is possible, it is also important to understand what **strategies** can support it.

- Sometimes hard, because different strategies can work (**multiplicity**).

Sometimes focus on case where players are very patient ($\delta \to 1$), because usually hard to characterize the equilibrium payoff set for a particular fixed δ (unless exact efficiency is attainable, which it usually isn't with noisy monitoring).

- The **folk theorem** holds if every "feasible and individually rational" payoff vector can be approximated by an equilibrium as $\delta \to 1$.
- Trying to understand the fixed $\delta < 1$ case is also important. This is sometimes possible analytically, more often numerically.
- There are also "intermediate" approaches: e.g., increasing patience and monitoring noise at the same time ("frequent actions"/"continuous time"); approximation/rate of convergence as $\delta \to 1$.

Applications of Repeated Games

Repeated games are widely applied in economics and related fields.

- IO, especially to study collusion among firms.
- Political economy, e.g. to study long-run incentives of politicians, or institutions.
- Organizational economics, especially to study "relational contracts" within organizations, or between fims and their suppliers or customers.
- Macroeconomics, especially to study government credibility.
- Development, e.g. to study informal insurance, risk-sharing, or public goods provision.
- Evolutionary biology, especially to study evolution of self-interested "reciprocal altruism."
- Economics and computer science, e.g. to study behavior on trading platforms like eBay, AirBnB, Uber/Lyft.

Extensions of the Baseline Repeated Game Model

- **Stochastic games:** allow intrinsic payoff link between per-period payoff functions, via an endogenous state variable.
- **Continuous-time games** and other variations on the usual timing of actions, signals, and payoffs.
- **Random matching** or **community enforcement**: people play repeatedly, but with different partners each period (exogenously or endogenously).
- **Repeated games with (persistent) incomplete information**, where dynamically revealing/concealing info matters.

We'll say something about each of these topics, some more than others.

Plan for Lectures 5–6

1. Example of RG w/ public monitoring.
2. Baseline RG model.
3. Recursive structure of RG's w/ public monitoring (APS90).
4. Linear programming characterization of eqm payoffs when $\delta \to 1$ (FL94).
5. Folk theorem w/ public monitoring (FLM94).
6. Patience vs. information arrival (AMP91, FL07, SS07, SW23)

Repeated Games with Perfect Monitoring

I assume familiarity with the basics of repeated games with perfect monitoring: FT Ch. 5.1–5.4 (especially 5.1), or for more detail MS Ch. 2–6 (optional).

Example: A Partnership Game

(Based on repeated prisoner's dilemma example in MS Ch. 7-8.)

- 2 players
- 2 actions each period: *Effort* and *Shirk* (or *Cooperate* and *Defect*).
- Payoff matrix

	E	S
E	$2, 2$	$-1, 3$
S	$3, -1$	$0, 0$

- At end of each period, the players observe a binary public signal $y \in \{\underline{y}, \bar{y}\}$.
- Each player's effort makes $\bar{y}$ more likely:

$$\Pr(\bar{y}|a) = \begin{cases} p & \text{if } a = EE \\ q & \text{if } a = ES \text{ or } SE \\ r & \text{if } a = SS \end{cases}$$

where $0 < r < q < p < 1$.

Aside: Do Players Observe their Own Payoffs?

Do players observe their own payoffs, in addition to the public signal y?

- In most applications, natural to assume that players observe their own payoffs.
- But then apparently payoffs should count as an "extra signal" that players can condition their actions on.
- Resolution: payoff matrix represents **expected** payoffs; each player's **realized** payoff depends only on y and her own action.
- Such a specification of realized payoffs always exists under a full rank condition on the signal distribution. Often left implicit, so model is specified only in terms of the ex ante payoff matrix and the conditional signal distribution (as in the previous slide).
  - Assumption that $u_i(a)$ can be written as expectation of realized payoffs that depend only on y and a_i is called **observed own payoffs**.

Do Players Observe their Own Payoffs? (cntd.)

Caveats:

1. If signals don't have full rank, be careful about whether can interpret players as observing own payoffs (and if not, whether that's reasonable).
2. When do comparative statics on the signal distribution, usually hold ex ante payoffs fixed. This implicitly changes the realized payoffs along with the signal distribution.

Back to the Partnership Game

How can players support cooperation in the partnership game?

One idea: **grim trigger strategies**.

- Each player $i \in \{1, 2\}$ takes $a_i = E$ if all previous signals are $\bar{y}$, takes $a_i = S$ if any previous signal is $\underline{y}$.

Grim trigger is an equilibrium if δ and $p - q$ (difference in $\Pr(\bar{y})$ when i takes E vs. S) are both high enough.

- Effort cost of 1 today outweighed by increased prob of opponent taking E tomorrow (worth 3).

However, since $p < 1$, expected discounted payoffs under grim trigger go to 0 as $\delta \to 1$.

- Recall: payoffs given by $(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(a_t)$.
- Under grim trigger, permanently transition from EE to SS w/ indep prob $1 - p$ each period.
- Given this fixed transition probability, $\lim_{\delta \to 1} \mathbb{E}\left[(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(a_t)\right] = u_i(SS) = 0$.

Tolerant Grim Trigger

A problem with grim trigger: punishment of permanently transitioning to SS is too harsh when δ is high.

- And it happens on-path, since monitoring is noisy.

Can do better with **tolerant grim trigger**.

- Start in EE. If $y = \bar{y}$, stay in EE. If $y = \underline{y}$, stay in EE w/ prob ϕ, permanently transition to SS w/ prob $1 - \phi$.

Tolerant Grim Trigger: Remarks

- Strategy transitions are coordinated. Possible if players observe a public randomizing device at the start of each period. We assume public randomization is available.
- However, instead of a permanent transition to SS w/ prob $1-\phi$, could transition to SS for the next "$1-\phi$ fraction of (discounted) time" for sure. Also works.
- Tolerant grim trigger is *optimal* among **strongly symmetric equilibria (SSE)**: players take same action *at every history*.
- In a strongly symmetric eqm, continuation payoffs stay on the 45° line. The lowest eqm payoff on the 45° line is $(0,0)$, the highest is some (v,v). Rather than moving a little down the 45° line from (v,v), can mix between staying at (v,v) and jumping to $(0,0)$. This is what tolerant grim trigger does.

How Efficient is Tolerant Grim Trigger?

What are the highest payoffs that can be supported with tolerant grim trigger (and hence in any SSE)?

In particular, as $\delta \to 1$, can tolerant grim trigger support payoffs close to the efficient point $(2, 2)$?

- The higher is ϕ, the longer we stay in EE, so the higher are the payoffs.
- But if ϕ is too high, deviating to S is profitable.
- So the best tolerant grim trigger eqm is the one where ϕ just satisfies incentive constraints.

How Efficient is Tolerant Grim Trigger? (cntd.)

Let's compute this. Note that

$$v = (1-\delta)(2) + \delta\left[pv + (1-p)(\phi v + (1-\phi)(0))\right].$$

Solving for v gives

$$v = \frac{2(1-\delta)}{1-\delta(p+(1-p)\phi)}.$$

The incentive constraint is

$$v \geq (1-\delta)(3) + \delta\left[qv + (1-q)(\phi v + (1-\phi)(0))\right].$$

Subbing in v and setting $\geq$ to $=$ gives

$$\phi = \frac{\delta(3p-2q)-1}{\delta(3p-2q-1)}.$$

Note: optimal value of ϕ is in $[0,1]$ iff $\delta(3p-2q) \geq 1$. This is precisely the condition for regular grim trigger to be an equilibrium. As we predicted, it holds if δ and $p-q$ are both high enough.

How Efficient is Tolerant Grim Trigger? (cntd.)

Sub the optimal value of ϕ back into our equation for v to get

$$v = 2 - \frac{1-p}{p-q}.$$

This is the highest payoff a player can get in any SSE.

Note that v does not go to 2 as $\delta \to 1$. So the folk theorem fails *for SSE*.

In fact, v doesn't depend on δ at all.
(Once δ is high enough that a grim trigger eqm exists.)

- $\delta \uparrow \implies$ can reduce transition prob while satisfying incentives, but transition is worse when it happens.

Remarks

- v is determined by "signal informativeness": the likelihood ratio difference

$$\frac{\Pr\left(\underline{y}|ES\right) - \Pr\left(\underline{y}|EE\right)}{\Pr\left(\underline{y}|EE\right)} = \frac{p-q}{1-p}.$$

  - Reminiscient of moral hazard (Holmström 79).
  - If $p \to 1$ then "false positives" $\to 0$, and hence inefficiency $\to 0$.
- In SSE, incentives can be provided only through **continuation value destruction**.
- Expected value destruction required for incentives is determined by signal informativess, not δ.
- To achieve effficiency with noisy signals, must provide incentives through **continuation value transfers** rather than destruction.
- Requires going beyond SSE.

Beyond Strong Symmetry

In our example, the folk theorem actually fails or a much broader class of eqm: **perfect public equilibrium (PPE)**.

- To support payoffs near $(2, 2)$, must incentivize EE.
- Incentivizing EE by threatening value destruction when $y = \underline{y}$ ends up destroying significant value on average.
- Suppose instead incentivize EE by saying that if $y = \underline{y}$ then continuation value is transfered from player 1 to player 2. (E.g., in future player 1 must take E while player 2 takes S.)
- This helps with player 1's incentives, but it makes player 2's incentives even worse. So it fails to incentivize EE.
- Problem: $y = \underline{y}$ is equally "bad news" about **both** players' actions, so not a useful basis for **transfers**.

Formally, EE is not **orthogonally enforceable** (in direction $(1, 1)$)**:** there is no specification of continuation payoffs that incentivizes EE without destroying social value with Pareto weights $(1, 1)$.

More Signals

Consider the same stage game but a different information structure: a 2-dimensional public signal

$$y = (y_1, y_2) \in \{\underline{y}, \bar{y}\} \times \{\underline{y}, \bar{y}\},$$

$$\text{where } \Pr(\bar{y}_i | a) = \begin{cases} p & \text{if } a_i = E \\ q & \text{if } a_i = S \end{cases}$$

independently across dimensions.

- A monitoring structure with conditionally independent signals of each player's action is called **product structure**.
- Product structure monitoring plus observed own payoffs implies the folk theorem for PPE.
  - We'll see a weaker sufficient condition: **pairwise full rank**.

Also, suppose the players can send each other money at the end of each period, and these transfers are observable.

- Not necessary when $\delta \to 1$, because continuation v ue transfers can substitute for monetary transfers.

al

More Signals (cntd.)

Then EE can be incentivized with no efficiency loss, for high enough $\delta < 1$.

- Players take EE each period.
- If $y_i = \underline{y}_i$, player i must pay $\$1/\left(p-q\right)$ to player j. (Or pick any larger number here.)
- If player i doesn't pay, go to SS forever.

Proof.

- Eqm payoffs equal $(2,2)$, as transfers cancel out on average.
- Taking E is optimal, as effort cost is 1 and reduces prob of having to pay $\$1/\left(p-q\right)$ from $1-q$ to $1-p$.
- For high enough δ, paying $\$1/\left(p-q\right)$ when required is optimal, because $\left(1-\delta\right)/\left(p-q\right) < 2$.

Partnership Game: Summary

With 2 signals, folk theorem fails for SSE (we proved it), and also for PPE (gave intuition).

- Can't statistically distinguish deviations by players 1 and 2, so can only provide incentives by value destruction.
- Value destruction + noisy signals = inefficiency, even as $\delta \to 1$.
- Inefficiency determined by signal informativeness.

Partnership Game: Summary (cntd.)

With 2 signals per player + product structure, folk theorem holds with monetary transfers (we proved it), and also without them (gave intuition).

- Can statistically distinguish deviations by players 1 and 2, so can provide incentives by value transfers.
  - Pairwise full rank suffices.
- With monetary transfers, Pareto frontier of eqm payoff set is **linear**, so can transfer value w/ **zero** efficiency cost.
- Without monetary transfers, for any payoff set where boundary is **smooth**, can transfer value w/ **vanishing** efficiency cost as $\delta \to 1$.
  - Intuition for public monitoring folk theorem: identification + approx tangent continuation payoff value movement = vanishing efficiency cost.
- Without monetary transfers, when $\delta < 1$ amount of inefficiency again determined by signal informativeness (and hard to characterize precisely).

Baseline Repeated Game Model

Now turn to formal model and analysis.

In a (discounted) repeated game, a static *stage game* is played repeatedly in periods $t = 1, 2, \ldots$, players get some *signals* before taking actions each period, and players maximize *expected discounted payoffs*.

A **stage game** $G = (I, A, u)$ consists of

- A set of players $I = \{1, \ldots, N\}$.
- A set of actions A_i for each player $i \in I$.
- A payoff function $u_i : A \rightarrow \mathbb{R}$ for each player $i \in I$ (where $A = \times_{i \in I} A_i$ and $u = \times_{i \in I} u_i$).
  - Assume A is finite.

Feasible and Individually Rational Payoffs

A payoff vector $v \in \mathbb{R}^N$ is **feasible** if there exists a distribution $\mu \in \Delta(A)$ (possibly correlated) s.t. $v_i = \sum_{a \in A} \mu(a) u_i(a)$.

- Denote set of feasible payoff vectors by $V \subset \mathbb{R}^N$.
- In other words, $V = \mathrm{conv}(u(A))$.

Player i's (mixed action) **minmax payoff** is

$$\underline{u}_i = \min_{\alpha_{-i} \in \times_{j \neq i} \Delta(A_j)} \max_{a_i \in A_i} u_i(a_i, \alpha_{-i}).$$

A payoff vector $v \in \mathbb{R}^N$ is **individually rational (IR)** if $v_i \geq \underline{u}_i$ for all i.

- Denote set of feasible + IR payoff vectors by $V^* \subset \mathbb{R}^N$.

Other minmax notions: pure action minmax (relevant for pure strategy eqm), correlated minmax (take min over $\alpha_{-i} \in \Delta(A_{-i}) \supset \times_{j \neq i} \Delta(A_j)$), relevant when opponents can correlate without i's knowledge).

Baseline Repeated Game Model (cntd.)

A **repeated game** $\Gamma = (G, \delta, Y, p)$ consists of

- A stage game G.
- A discount factor $\delta \in (0, 1)$.
- A **monitoring structure** (Y, p), where in each period if action profile $a \in A$ is played, a **signal** $y \in Y$ (an N-dimensional vector) is drawn according to $p(\cdot|a) \in \Delta(Y)$, and each player i observes the i^{th} component of y.
  - Assume Y is finite.

A **history** $h_i^t = (a_i^t, y_i^t)$ for player i at the beginning of period t (where $a_i^t = (a_{i,1}, \dots, a_{i,t-1})$ and $y_i^t = (y_{i,1}, \dots, y_{i,t-1})$) summarizes player i's information before choosing $a_{i,t}$.

A (behavior) **strategy** σ_i for player i maps histories h_i^t to (mixed) actions $\alpha_{i,t} \in \Delta(A_i)$.

Players choose strategies to max expected discounted utility, $(1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} u_i(a_t)$.

Types of Monitoring Structures

Monitoring structures in repeated games are classified as **perfect**, **public**, or **private**.

Perfect monitoring means all actions are perfectly observed:
$p(y|a) = \mathbf{1}\{y_i = a \; \forall i \in I\}$.

Public monitoring means all players see the same signal:
$p(y|a) > 0 \implies y_i = y_j \; \forall i, j$.

- E.g., N firms, each chooses how much output to produce, then observes the common market-clearing price, which depends on everyone's output and a random market demand shock. (Green Porter 84)

Private monitoring refers to the general case.

- E.g., N firms, each chooses its price, then observes its own sales, which depends on everyone's prices and a random market demand shock. (Stigler 64)

(Non-)Recursive Structure of Repeated Game Equilibria

One reason why it makes a big difference whether monitoring is perfect, public, or private is that this determines whether equilibria have a **recursive structure**.

With perfect monitoring, in any subgame perfect equilibrium, continuation play starting at any history is itself a subgame perfect equilibrium.

(Non-)Recursive Structure (cntd.)

With imperfect (public or private) monitoring, it is **not** necessarily true that continuation play in a sequential equilibrium is itself a sequential equilibrium.

- Players can have different information at the beginning of period t and hence different beliefs, so their continuation strategies need not be mutual best responses.
- With private monitoring, players inevitably have different information, so there is no way to recover a recursive structure (without dramatically restricting strategies).
  - This is the key challenge in repeated games with private monitoring, which we cover next week.
- With public monitoring, players see the same signals y, but they can still have different information because each player knows her own past actions.
  - However, **if** all players use strategies that condition only on the public signal and not on their own actions, this private information is irrelevant, and we recover a recursive structure.

Public Strategies and PPE

A strategy σ_i for player i is **public** if it depends only on the public history y^t: that is, $\sigma_i\left(a_i^t, y^t\right) = \sigma_i\left(\tilde{a}_i^t, y^t\right)$ for all $t, a_i^t, \tilde{a}_i^t, y^t$.

A **perfect public equilibrium (PPE)** is a strategy profile σ where

1. Each σ_i is a public strategy.
2. For each t and y^t, the continuation strategies form a NE starting from public history y^t.

- This is well-defined: there are no proper subgames because actions are private information, but since no one conditions on actions we can check that continuation strategies are mutual best responses.

Properties of Public Strategies and PPE

Whenever players $-i$ use public strategies, player i has a public strategy as a best response.

- Since opponents' play depends only on y^t, and i observes y^t, i has no reason to condition on own past realized actions.

Every **pure strategy** sequential eqm is outcome-equivalent to a pure strategy PPE.

- Suppose i's on-path period 1 action is $a^*_{i,1}$ and her period 2 strategy is $\sigma_i(a_{i,1}, y_1)$. Replace this with $\sigma_i\left(a^*_{i,1}, y_1\right)$ for all $a_{i,1}$. This yields the same outcome distribution for periods 1 and 2 as the original strategy, for any opposing strategies. Recursively, do this for each period, and then for each player.

Properties of PPE (cntd.)

In some games, the set of mixed strategy sequential eqm outcomes is **strictly larger** than the set of mixed strategy PPE outcomes.

- If i and j are both mixing, then i's past action realizations are informative about j's past action realizations, conditional on y^t, and similarly for j. So both may strictly prefer to condition on past realizations. (See FT Exercise 5.10.)
- Considering PPE vs. arbitrary SE makes a big difference in some games.
  - E.g., Players must get at least their mixed-action minmax in every PPE, but not in every SE (although must get at least correlated minmax). *(Why?)*

Recursive Structure of PPE

PPE are not without loss, but they are a large class of SE with a recursive structure.

- In any PPE, the continuation payoff vector starting at any public history y^t is itself a PPE payoff vector.
  - Can't do this for arbitrary SE, as continuation payoff vector conditional on y^t isn't well-defined, and players have different beliefs about continuation payoffs at the full history (a^t, y^t).
- The set E of all PPE payoff vectors can be characterized as the largest set W of payoff vectors that can be "generated" with continuation payoffs drawn from W.
- Any such set W of **self-generating** payoffs is contained in E.
- To characterize E, instead of constructing an eqm to attain each payoff vector, ask what **sets** of payoff vectors W are self-generating.
  - Much more tractable.
  - Key contribution of APS. Extends value recursion from dynamic programming to games.

Enforceability

Definition

For any $W \subset \mathbb{R}^N$, a mixed action profile $\alpha \in \times_i \Delta(A_i)$ is **enforceable on** W if there exists a function $w : Y \to W$ s.t. α is a NE in the stage game augmented with continuation payoffs $w(y)$: that is, for all i,

$$\operatorname{supp} \alpha_i \subset \underset{a_i \in A_i}{\operatorname{argmax}} (1-\delta) u_i(a_i, \alpha_{-i}) + \delta \sum_{y \in Y} p(y | a_i, \alpha_{-i}) w_i(y).$$

Decomposability

Definition

A payoff vector $v \in \mathbb{R}^N$ is **decomposable on** W if there exist $\alpha \in \times_i \Delta\left(A_i\right)$ and $w : Y \to W$ s.t.

1. $v = (1-\delta)\, u(\alpha) + \delta \sum_{y \in Y} p(y|\alpha)\, w(y)$.
   ("Promise Keeping")
2. w enforces α. ("Incentive Compatibility")

- v can be "generated" w/ continuation payoffs in W.
- Note: enforceability/decomposability depend on δ.

Self-Generation

For any $W \subset \mathbb{R}^N$, define

$$B(\delta, W) = \left\{ v \in \mathbb{R}^N : v \text{ is decomposable on } W \text{ w/ discount factor } \delta \right\}.$$

Definition

A set of payoffs $W \subset \mathbb{R}^N$ is **self-generating** if $W \subset B(\delta, W)$.

- Every $v \in W$ can be generated w/ continuation payoffs in W.

The Set of PPE Payoffs is Self-Generating

Let $E(\delta)$ =set of all PPE payoff vectors w/ discount factor δ.

Theorem
$E(\delta) \subset B(\delta, E(\delta))$.

- Fix $v \in E(\delta)$. Pick a PPE σ s.t. $U(\sigma) = v$.
- Let α denote the period-1 mixed action profile.
- Note that v is decomposed by α and w given by $w(y_1)$ =continuation payoffs conditional on y_1.
- Since σ is a PPE, continuation strategies conditional on each y_1 form a PPE. Hence, continuation payoffs conditional on each y_1 lie in $E(\delta)$.
- This shows that $v \in B(\delta, E(\delta))$.

Self-Generating Payoffs are PPE Payoffs

Theorem

For any bounded, self-generating set of payoff vectors $W \subset \mathbb{R}^N$, we have $W \subset B(\delta, W) \subset E(\delta)$.

- Fix $v \in B(\delta, W)$. Pick $\alpha^1, w^1 : Y \rightarrow W$ that decompose it.
- For each $w^1(y_1) \in W$, we have $w^1(y_1) \in B(\delta, W)$. *(Why?)* Pick $\alpha^2, w^2 : Y \rightarrow W$ that decompose $w^1(y_1)$. Keep going. This defines a strategy profile.
- The strategy profile yields payoff v. This follows from boundedness: present value is similar to payoff if we truncated at period T and got 0 thereafter ("continuity at infinity"); then take limit on the truncation point.
  - Rules out generating v today by promising v/δ tomorrow, generating that by promising v/δ^2 the day after, etc.
- The strategy profile is a PPE. This follows from continuity at infinity and the 1-shot deviation principle: if strategies were not Nash from some y^t on, there would be a profitable 1-shot deviation, but decomposability implies there isn't one.

Remark: Self-Generation and the Folk Theorem

Fact that any bounded, self-generating set W is contained in $E(\delta)$ is key to proving the folk theorem with imperfect public monitoring.

- We'll show that under some conditions, any smooth set W in the interior of the set of feasible + individually rational payoffs is self-generating when δ is high enough.
- This implies that any F+IR payoff vector can be approximated in some PPE as $\delta \to 1$, i.e. the folk theorem holds.

E(\delta) is the Largest Bounded, Self-Generating Set

Corollary

$E(\delta) = B(\delta, E(\delta))$.

- We showed that $E(\delta)$ is self-generating.
- $E(\delta)$ is contained in the set of feasible payoffs, so it's bounded.
- We showed that for any bounded, self-generating set W we have $W \subset B(\delta, W) \subset E(\delta)$.
- So we have $E(\delta) \subset B(\delta, E(\delta)) \subset E(\delta)$.

Properties of the Operator B

B is a monotone operator: if $W \subset W'$ then $B(\delta, W) \subset B(\delta, W')$.

- v decomposable w/ continuation values in $W \implies$ decomposable w/ continuation values in W'.

B is a compact operator: if W is compact then $B(\delta, W)$ is compact.

- Follows because set of $\alpha \in \times_i \Delta(A_i)$ and $w : Y \to W$ is compact, and enforceability is defined by weak inequalities.

$V \supset B(V) \supset B(B(V)) \supset \ldots \supset B^\infty(V) = E(\delta)$.

- Gives an algorithm for computing $E(\delta)$. This is a key contribution of APS. See Judd Yeltekin Conklin 03, Abreu Sannikov 14 for implementations.
- Chain of inclusions, including $B^\infty(V) \supset E(\delta)$ follows from monotonicity. Can show $B^\infty(V) = E(\delta)$ by compactness arguments. See APS or MS Ch. 7.3.
- Corollary: $E(\delta)$ is compact.

Monotonicity in \delta

Not generally true that $B(\delta, W)$ is monotone in δ.

- Set of feasible payoffs not monotone in δ due to discreteness.

But it is true w/ enough convexity: if W is convex and $\delta < \delta'$, then $B(\delta, W) \cap W \subset B(\delta', W)$.

- If $v \in B(\delta, W) \cap W$ and δ increases, can average between staying at v for all y and using the same continuation values that decomposed v with discount factor δ.
- "Ignoring y" with some probability offsets the increase in δ.
- Can choose v itself as continuation value because $v \in W$.

Corollary

If $E(\delta)$ is convex and $\delta < \delta'$, then $E(\delta) \subset E(\delta')$.

- We have $E(\delta) = B(\delta, E(\delta)) \cap E(\delta) \subset B(\delta', E(\delta))$, so $E(\delta)$ is self-generating with discount factor δ'.
- Implies that $E(\delta) \subset E(\delta')$.

Remark: Monotonicity in \delta and Public Randomization

When public randomization is available, $E(\delta)$ is always convex, and hence is always monotone in δ.

$E(\delta)$ is often convex when δ is high even without public randomization, as deterministic action sequences can mimic public randomization.

We'll see that the set of limit PPE payoffs $\lim_{\delta \to 1} E(\delta)$ is convex, and is the same with and without public randomization.

However, without public randomization there are examples where $E(\delta)$ is non-convex and non-monotone for arbitrarily large $\delta < 1$ (Yamamoto 10).

Monotonicity in Signal Precision

Intuitively, monotonicity in δ is related to monotonicity in signal precision: future rewards/punishments get more weight vs. vary more with actions.

Let P denote the $|A| \times |Y|$ matrix with entries $p(y|a)$.

Definition

A public monitoring structure (Y', p') is a (Blackwell) **garbling** of (Y, p) if there exists a $|Y| \times |Y'|$ row-stochastic matrix Q such that

$$P' = PQ.$$

- Can replicate draw of y' from (Y', p') by first drawing y from (Y, p) and then drawing y' according to $q(y'|y)$.

Monotonicity in Signal Precision (cntd.)

For monitoring structures (Y, p) and (Y', p'), let B and B' denote the corresponding generation operators.

Theorem (Kandori 92)

If (Y', p') is a garbling of (Y, p) and W is convex, then $B'(\delta, W) \subset B(\delta, W)$. In particular, if $E'(\delta)$ is convex, then $E'(\delta) \subset E(\delta)$.

- Given $w' : Y' \to W$, define $w : Y \to \mathbb{R}^N$ as expected value of $w'(y')$ when y' is drawn according to $q(y'|y)$.
- By defn of garbling, expected continuation values conditional on any a are the same under $(Y, p; w)$ and $(Y', p'; w')$, so they decompose the same v's.
- Since W is convex, the image of w lies in W.

Kandori also gives conditions for strict inclusion: $E'(\delta) \subsetneq E(\delta)$.

- Strictly positive Q + "Slater condition."

Bang-Bang Structure of PPE

One more result on general structure of PPE: it's typically without loss to restrict continuation values to *extreme points* of $E(\delta)$.

The **extreme points** of a convex set W are

$$\mathrm{ext}W = \left\{ v \in W : \nexists v', v'' \in W, \beta \in (0,1) \text{ s.t. } v = \beta v + (1-\beta) v'' \right\}.$$

With public randomization, this is immediate: replacing an interior continuation value $w(y)$ with a distribution over $\mathrm{ext}W$ with the same expectation doesn't affect payoffs or incentives.

Without public randomization:

- If signals are discrete, cannot restrict to extreme points.
- If signals (and payoffs) are continuous, can restrict to extreme points. Intuitively, use fine variation in the signal to mimic public randomization. See APS, MS Ch. 7.5.

Bang-Bang Structure (cntd.)

The bang-bang result is especially useful if we restrict attention to strongly symmetric equilibria.

- Strongly symmetric eqm $\implies$ continuation values on 45° line $\implies$ 2 extreme points.
- $w : Y \to \mathbb{R}^N$ simplifies to $w : Y \to [0,1]$, prob of picking the good continuation value.
- We saw how this works in the partnership example, where we could calculate the optimal SSE.

For general PPE, $E(\delta)$ is typically N-dimensional, so bang-bang result is often not as useful analytically.

- Still can be useful numerically.
- Also useful analytically in principal-agent models ($N = 1$) and some 2-player games where can extract qualitative properties from the APS algorithm. Example: recent papers on repeated delegation by Padro i Miquel Yared, Fong Li, Guo Horner, Li Matouschek Powell, Lipnowski Ramos.

Characterizing PPE Payoffs as \delta–>1

So far, analyzed structure of PPE for arbitrary fixed $\delta < 1$.

- Recursive structure, monotonicity in δ and signal precision, bang-bang structure.
- APS's recursive algorithm is useful both conceptually and numerically. But typically don't have an explicit, analytic characterization of $E(\delta)$.

Next topic: characterizing limit PPE payoffs, $\lim_{\delta\to1} E(\delta)$.

- Taking limit allows explicit characterization, new insights on role of identification, and connection between repeated games and static contracting.
- We'll characterize $\lim_{\delta\to1} E(\delta)$ in terms of solutions to a collection of linear programs/static contracting problems. (FL94)
- Then use this characterization to give conditions for the folk theorem: when does $\lim_{\delta\to1} E(\delta) = V^*$? (FLM94)

Limit PPE Payoffs as Intersection of Half-Spaces

We'll see that $\lim_{\delta \to 1} E(\delta)$ is convex, so it can be written as an intersection of *half-spaces*.

- Given a **direction** (unit vector) $\lambda \in \mathbb{R}^N$ and a **score** (scalar) $k \in \mathbb{R}$, let $H(\lambda, k) = \{v \in \mathbb{R}^n : \lambda \cdot v \leq k\}$.
  - Payoff vectors with "average payoff" $\leq k$ for Pareto weights λ.
- We'll define the *maximum score* $k^*(\lambda)$ that can be attained in direction λ using only transfers or value destruction (no value creation).
  - Variation in continuation play can transfer or destroy value.
  - Can't create more value when we're already at the max.
- Then we'll show that $\lim_{\delta \to 1} E(\delta) = \bigcap_{\lambda} H(\lambda, k^*(\lambda))$.

Maximum Score for a Given Action Profile

Definition

For any $\alpha \in \times_i \Delta(A_i)$, the **maximum score attainable by α in direction λ** is defined as

$$k^*(\alpha, \lambda) = \max_{v \in \mathbb{R}^N, x: Y \to \mathbb{R}^N} \lambda \cdot v \qquad s.t.$$

1. $v_i = u_i(\alpha) + \sum_y p(y|\alpha) x_i(y)$ for all i.
   (Promise Keeping)
2. $v_i \geq u_i(a_i, \alpha_{-i}) + \sum_y p(y|a_i, \alpha_{-i}) x_i(y)$ for all i, a_i.
   (Incentive Compatibility)
3. $\lambda \cdot x(y) \leq 0$ for all $y \in Y$. (Generation by $H(\lambda, \lambda \cdot v)$)

The constraints say that v is decomposable on $H(\lambda, \lambda \cdot v)$, starting with action profile α (where $w(y) = v + \frac{1-\delta}{\delta} x(y)$).

If exist v and x that satisfy the constraints with equality in 3 for all y, we say that α is **orthogonally enforceable in direction λ**.

Stage Game Payoffs Bound the Maximum Score

Lemma

$k^* (\alpha, \lambda) \leq \lambda \cdot u (\alpha)$.

- By definition of the max score, there exists v such that v is a convex combination of $u(\alpha)$ and points in $H(\lambda, k^*(\alpha, \lambda))$, and v is an extremal point of $H(\lambda, k^*(\alpha, \lambda))$.
- So, $u(\alpha)$ cannot be in the interior of $H(\lambda, k^*(\alpha, \lambda))$.

Orthogonal Enforceability Implies Bound Attained

Lemma

If α is orthogonally enforceable in direction λ then $k^*(\alpha, \lambda) = \lambda \cdot u(\alpha)$.

- With orthogonal enforceability, we have

$$k^*(\alpha, \lambda, \delta) = \lambda \cdot v = \lambda \cdot \left(u(\alpha) + \delta \sum_y p(y|\alpha) x(y) \right) = \lambda \cdot u(\alpha).$$

- Intuitively, orthogonal enforcement means no value destruction with Pareto weights λ.

Maximum Score for Any Action Profile

To find the overall max score, optimize over the first period action profile α.

- Let $k^* (\lambda) = \sup_\alpha k^* (\alpha, \lambda)$.
  - Max score in direction λ (with any α).
  - Solution to a static optimization problem.
- Let $H^* (\lambda) = H (\lambda, k^* (\lambda))$.
  - Max half-space in direction λ.
- Let $M = \bigcap_\lambda H^* (\lambda)$.
  - Intersection of max half-spaces.

We wish to characterize $\lim_{\delta \to 1} E (\delta)$ as the intersection of maximal half spaces: $\lim_{\delta \to 1} E (\delta) = M$.

- "When δ is high, continuation value transfers are as good as monetary transfers."

Characterization of Limit PPE Payoffs

Let $E^*(\delta) = \mathrm{conv} E(\delta)$. ($=E(\delta)$ w/ pub rand.)

Theorem

1. *$E^*(\delta) \subset M$ for all δ.*
2. *If $\dim M = N$, then $\lim_{\delta\to1} E(\delta) = \lim_{\delta\to1} E^*(\delta) = M$.*

Thus, M bounds PPE payoffs for any δ, and (when it's full-dimensional) it characterizes PPE payoffs as $\delta \to 1$.

- Reduces problem of characterizing $\lim_{\delta\to1} E(\delta)$ to a set of static optimization problems.
- Implication: the folk theorem holds iff $M = V^*$.
- Once we prove the above theorem, we can then prove the folk theorem by giving conditions under which $M = V^*$.
  - Key: give conditions for orthogonal enforcement, as this implies $k^*(\alpha, \lambda) = \lambda \cdot u(\alpha)$.

Proof of 1: Self-Generation Implies Bounded by M

Part 1 is rather immediate. Intuition: variation in continuation play can't create value when already at the max.

Proof.

- Since $E^*(\delta)$ is convex, if it's not contained in M then it extends farther than M in some direction λ. So $E(\delta)$ also extends farther than M in direction λ.
- Pick $v \in \mathrm{argmax}_{v' \in E(\delta)} \lambda \cdot v'$. Note that $\lambda \cdot v > k^*(\lambda)$.
- Since $E(\delta)$ is self-generating, v is decomposable on $E(\delta)$, and hence it's decomposable on the larger set $H(\lambda, \lambda \cdot v)$.
- But if v is decomposable on $H(\lambda, \lambda \cdot v)$, then by definition $k^*(\lambda) \geq \lambda \cdot v$. This is a contradiction.

Proof of 2: Smoothness Implies Limit Self-Generation

Part 2 follows from the following theorem, which is one of the key insights of FL/FLM (also Matsushima 1989).

Say that a compact set $W \subset \mathbb{R}^N$ is **smooth** if it has non-empty interior and its boundary is C^2.

- Unique unit normal vector at each point suffices.

Theorem

For every smooth convex set $W \subset \mathrm{int} M$, there exists $\bar{\delta}$ such that, for any $\delta > \bar{\delta}$, W is self-generating: hence,

$$W \subset B(\delta, W) \subset E(\delta).$$

Intuition

If the boundary of W were linear, then each boundary point $v \in \mathrm{argmax}_{v' \in W} \lambda \cdot v'$ would be decomposable, because all continuation payoffs in $H(\lambda, \lambda \cdot v)$ would be available and $\lambda \cdot v \leq k^*(\lambda)$.

Since W is smooth, the boundary in any direction is "locally linear."

When δ is high, small continuation payoff movements suffice to provide incentives. So local linearity is good enough.

More precisely: continuation payoff movements of order $1 - \delta$ suffice to provide incentives, but smoothness implies that orthogonal continuation payoff movements of order $\sqrt{1-\delta}$ stay in W. Since $\sqrt{1-\delta} \gg 1 - \delta$ when $\delta \approx 1$, W is self-generating.

Self-Generation of Boundary Points

We show that any boundary point v of a smooth, convex set $W \subset \mathrm{int} M$ is decomposable with continuation payoffs in W, when δ is high enough.

- By compactness, there exists $\bar{\delta}$ such that the entire boundary is decomposable when $\delta > \bar{\delta}$.
- When δ is high, can decompose interior points with static NE, so this proves the theorem.

Decomposing Boundary Points

Fix a boundary point v. Let λ be the unique unit normal.
Let $k = \lambda \cdot v$.

Since $v \in W \subset \mathrm{int} M$, we have $k < k^* (\lambda) \equiv k^*$. So, there exist α and v^* such that $\lambda \cdot v^* = k^* > k$, and v^* is generated by α and continuation payoffs in $H(\lambda, k^*)$.

So v is generated by α and translation by $-(1/\delta)(k^* - k)$ of continuation payoffs used to generate v^*.

The translated continuation payoffs lie in
$H\left(\lambda, k^* - \frac{1}{\delta}(k^* - k)\right) = H\left(\lambda, k - \frac{1-\delta}{\delta}(k^* - k)\right)$.

So v is decomposable with continuation payoffs in
$H\left(\lambda, k - \frac{1-\delta}{\delta}(k^* - k)\right)$.

Decomposing Boundary Points (cntd.)

Now, decompose v in this way, and let $w = \sum_y p(y|\alpha) w(y)$.

Since $v = (1-\delta) u(\alpha) + \delta w$, we have
$v - w = (1-\delta)(u(\alpha) - w)$, and hence $\lambda \cdot (v - w) = O(1-\delta)$.

- w is distance $1-\delta$ from the boundary of W in direction λ.

Since stage game payoffs get $1-\delta$ weight, we can take $w(y)$ s.t.
$|w - w(y)| \leq O(1-\delta)$.

Since W is smooth, distance from w to closest point in
$W^c \cap H(\lambda, k - O(1-\delta))$ is $O\left(\sqrt{1-\delta}\right)$.

- Follows from Pythagorean theorem. Choose coordinates s.t. $(u(\alpha) - v) \perp H(\lambda, k)$. Let r be radius of W, d be desired distance. Then $r^2 = (r - O(1-\delta))^2 + d^2$, so $d = O\left(\sqrt{1-\delta}\right)$. (See MS Ch. 9.1 for details.)

Hence, for high enough δ, every $w(y) \in H(\lambda, k - O(1-\delta))$ s.t.
$|w - w(y)| \leq O(1-\delta)$ lies in W.

Toward the Folk Theorem

We showed that

1. $\lim_{\delta \to 1} E(\delta) = M$ (when M is full-dimensional).
2. If α is orthogonally enforceable in direction λ then $k^*(\alpha, \lambda) = \lambda \cdot u(\alpha)$.

So, if all profiles are orthogonally enforceable in all directions then $M = V^*$ and hence $\lim_{\delta \to 1} E(\delta) = V^*$, so the folk theorem holds.

This is too much to ask for: if $\lambda = (1, 0, 0, \ldots)$ then can only orthogonally enforce action profiles where player 1 takes a static best reply, as all x s.t. $\lambda \cdot x(y) \leq 0$ are punishments for player 1.

So, a different argument is needed for **coordinate directions** (λ has 1 non-zero component) and **regular directions** (λ has > 1 non-zero component).

Pairwise Hyperplanes

Say that a direction λ is **pairwise** if it has exactly 2 non-zero components.

Lemma

If α is orthogonally enforceable in all pairwise directions, then it is orthogonally enforceable in all regular directions.

Intuition: If can enforce α while fixing $\lambda_i w_i + \lambda_j w_j$ (i.e., via transfers between i and j), and can enforce α while fixing $\lambda_j w_j + \lambda_k w_k$ (via transfers between j and k), can also enforce α while fixing $\lambda_i w_i + \lambda_j w_j + \lambda_k w_k$ (via transfers among i, j, k).

Full Rank Conditions

The key statistical identification condition that ensures orthognal enforceability in pairwise direction is *pairwise full rank*.

Start with a weaker condition: *individual full rank*.

For each action profile α and each player i, define the $|A_i| \times |Y|$ matrix $P_i\left(\alpha_{-i}\right)$ with elements

$$\left[P_i\left(\alpha_{-i}\right)\right]_{a_i,y} = p\left(y|a_i,\alpha_{-i}\right).$$

- The $\left(a_i, y\right)$ entry is the probability of signal y when player i deviates to a_i.

If $P_i\left(\alpha_{\ i}\right)$ has full row rank $(=|A_i|)$, player i's actions lead to linearly independent distributions over signals, i.e. they are statistically distinguishable.

- In this case, say that α has **individual full rank** for player i.

Individual Full Rank Implies Enforceability

Individual full rank implies that deviations from profile α by player i are detectable, so correct behavior by player i can be incentivized.

- Consider the equation

$$(1-\delta)\left[u_i\left(a_i,\alpha_{-i}\right)\right]_{a_i} = -\delta P_i\left(\alpha_{-i}\right)\cdot\left[w_i\left(y\right)\right]_y .$$

- When w_i satisfies this equation, player i is indifferent among all her actions.
- Sufficient condition for existence of a solution w_i: $P_i\left(\alpha_{-i}\right)$ has full row rank, i.e. individual full rank holds.
- So if α has individual full rank for every player i, then α is enforceable (on $\mathbb{R}^N$).

Note: If $P_i\left(\alpha_{\;\;i}\right)$ barely has full rank, required $w_i\left(y\right)$'s are very large: if signals not very informative, need large swings in continuation payoffs to provide incentives. This requires δ very close to 1. But OK as far as the folk theorem is concerned.

Pairwise Full Rank

A stronger condition is required to also imply that deviations by player i are distinguishable from deviations by player j, so correct behavior by player i can be incentivized by transfers to/from player j (rather than value destruction).

Definition

Profile α has **pairwise full rank** for players i and j if the $(|A_i| + |A_j|) \times |Y|$ matrix

$$P_{ij}(\alpha) = \begin{bmatrix} P_i(\alpha_{-i}) \\ P_j(\alpha_{-j}) \end{bmatrix}$$

has rank $|A_i| + |A_j| - 1$.

- This is the maximum possible rank, because $\alpha_i P_i(\alpha_{-i}) = [p(y|\alpha)]_{y \in Y} = \alpha_j P_j(\alpha_{-j})$.
- So pairwise full rank is stronger than individual full rank.

Pairwise Full Rank (cntd.)

- Pairwise full rank implies that there are at least $|A_i| + |A_j| - 1$ signals.
- E.g. Partnership game w/ 2 signals satisfied individual full rank but failed pairwise full rank. Partnership game w/ 4 signals satisfied pairwise full rank.

Pairwise Full Rank Implies Orthogonal Enforceability

Lemma

If α has pairwise full rank for all pairs of players, it is orthogonally enforceable in all regular directions.

- Since $\operatorname{rk}(P_{ij}(\alpha)) = |A_i| + |A_j| - 1$, there exist $\left[(w_i(y), w_j(y))\right]_{y \in Y}$ that make both players indifferent among all actions and satisfy an additional linear relationship between their continuation values.

Coordinate Directions

Lemma

If every pure action profile has individual full rank, then for any positive coordinate direction e_i, we have $k^*(e_i) = \max_{a \in A} u_i(a)$.

- Decomposability on $H(e_i, e_i \cdot v)$ reduces to $w_i(y) \leq v_i$ for all y. (Arbitrary rewards/punishments for players $-i$, only punishments for player i.)
- Fix any $a \in \operatorname{argmax}_{a'} u_i(a')$.
- By individual full rank, can enforce a_j for each $j \neq i$, as no restrictions on $w_{-i}(y)$.
- Can enforce a_i by setting $w_i(y) = v_i$ for all y, as i is taking a static best reply at a.

Nash-Threat Folk Theorem

Theorem

Assume that V^* has non-empty interior and that all Pareto efficient action profiles have pairwise full rank for all pairs of players.
For any $v \in \operatorname{int}(V^*)$ that Pareto-dominates a static NE, there exists $\bar{\delta} < 1$ such that, for all $\delta > \bar{\delta}$, we have $v \in E(\delta)$.

- Fix any α^{NE}, and let W be any smooth, convex set of points in $\operatorname{int}(V^*)$ that Pareto-dominate it.
- By the previous two lemmas, $k^*(\lambda) = \max_{\alpha \in \times_i \Delta(A_i)} u(\alpha)$ for all regular λ and all positive coordinate λ. So any $v \in W$ satisfies $\lambda \cdot v \leq k^*(\lambda)$ for all such λ.
- For any negative coordinate direction $-e_i$, $k^*(-e_i) \geq -u_i\left(\alpha^{NE}\right)$. So any $v \in W$ satisfies $\lambda \cdot v = -v_i \leq -u_i\left(\alpha^{NE}\right) \leq k^*(\lambda)$ for these λ as well.
- Hence, W is a smooth, convex set in $\operatorname{int}(M)$, so $W \subset E(\delta)$ for high enough δ.

Extensions

- If for each i there exists a minmax profile (a_i, α_{-i}) with individual full rank for each $j \neq i$, get the **minmax-threat** folk theorem: for all $v \in \mathrm{int}(V^*)$, there exists $\bar{\delta} < 1$ such that, for all $\delta > \bar{\delta}$, we have $v \in E(\delta)$. (In this case, $M = V^*$.)
- If every pure action profile has pairwise full rank for all pairs of players and $p(y|a) > 0$ for all y, a, and if for every i there exists a minmax profile in pure strategies, then the minmax-threat folk theorem holds in **strict** PPE.
- With known own payoffs, can relax pairwise full rank to **pairwise identifiability**:
  $\mathrm{rk}(P_{ij}(\alpha)) = \mathrm{rk}(P_i(\alpha_{-i})) + \mathrm{rk}(P_j(\alpha_{-j})) - 1.$
  This is important for games where $|Y| < |A_i|$.
- With a **product structure** and known own payoffs, then the Nash-threat folk theorem holds without any full rank assumptions.
- Results extend to games with long-run and **short-run players**, if restrict attention to α s.t. all short-run players take static best responses.

Low Discounting vs. Low Information

So far, we covered general results on the structure of PPE payoffs for arbitrary fixed $\delta < 1$, and characterized PPE payoffs in the $\delta \to 1$ limit.

In $\delta \to 1$ limit, only identification properties of monitoring matters, not informativeness (e.g., likelihood ratios, matrix inverse norm).

There are important classes of games where $\delta \to 1$ and at the same time informativeness vanishes. The standard $\delta \to 1$ analysis does not apply in these situations.

Low Discounting vs. Low Information: Examples

Frequent actions.

- Suppose actions affect an underlying continuous-time signal process, and the period length Δ measures how frequently players observe the process and potentially change actions.
- E.g., sales are essentially continuous, but management observes them and adjusts strategy every Δ days.
- Then $\Delta \rightarrow 0$ implies little discounting between updates, but also little information.

Large populations.

- Suppose players are patient ($\delta \rightarrow 1$) but there are many of them ($(1-\delta)N \nrightarrow 0$), and they are monitored through some "aggregate signal."
- What properties of the monitoring structure determine the prospects for cooperation in this case?

Poisson Signals: Bad News Case

The first paper on information vs. timing in repeated games was Abreu Milgrom Pearce 91.

- Focused on Poisson signals.
- We present an example that gives some key ideas.

Recall the partnership game. Suppose the discount factor and monitoring structure are parameterized by $\Delta > 0$ as follows:

- $\delta = e^{-r\Delta}$. (Act every Δ units of time, discount rate $r > 0$.)
- $\exists 0 < \beta < \mu$ s.t.

$$\Pr(\bar{y}|a) = \begin{cases} e^{-\beta\Delta} & \text{if } a = EE \\ e^{-\mu\Delta} & \text{otherwise} \end{cases}$$

("Bad" signal $\underline{y}$ arrives according to a Poisson process at rate β if $a = EE$, arrives at rate $\mu > \beta$ otherwise. "Good" signal $\bar{y}$ means no arrival. This situation where effort reduces a Poisson arrival rate is called the **Poisson bad news case**.)

Note: $\Delta \to 0 \implies \delta \to 1$, but also $\Pr(\bar{y}|a) \to 1 \ \forall a$.

Bad News Case (cntd.)

We showed that effort can be supported in an SSE iff $\delta(3p - 2q) \geq 1$ (where $p = \Pr(\bar{y}|EE)$, $q = \Pr(\bar{y}|ES \text{ or } SE)$), i.e.

$$e^{-r\Delta}\left(3e^{-\beta\Delta} - 2e^{-\mu\Delta}\right) \geq 1.$$

When $\Delta = 0$ the LHS equals 1 and its derivative equals $2\mu - 3\beta - r$. So the inequality holds for small Δ if

$$r < 2\mu - 3\beta.$$

When this holds, we saw that the greatest SSE payoff is

$$2 - \frac{1-p}{p-q} = 2 - \frac{1 - e^{-\beta\Delta}}{e^{-\beta\Delta} - e^{-\mu\Delta}}.$$

As $\Delta \to 0$, this converges to

$$2 - \frac{\beta}{\mu - \beta}.$$

- As in the $\delta \to 1$ limit, the greatest SSE payoff is in between

Good News Case

Now suppose that $\exists 0 < \mu < \beta$ s.t.

$$\Pr(\bar{y}|a) = \begin{cases} 1 - e^{-\beta\Delta} & \text{if } a = EE \\ 1 - e^{-\mu\Delta} & \text{otherwise} \end{cases}$$

("Good" signal $\bar{y}$ arrives according to a Poisson process at rate β if $a = EE$, arrives at rate $\mu < \beta$ otherwise. This situation where effort increases a Poisson arrival rate is called the **Poisson good news case**.)

Now, $\delta(3p - 2q) \geq 1$ iff

$$e^{-r\Delta}\left(1 - 3e^{-\beta\Delta} + 2e^{-\mu\Delta}\right) \geq 1.$$

When $\Delta = 0$, LHS=0. So, for small Δ, there is **no** SSE with strictly positive payoffs!

- Moreover, this holds uniformly over r.

Good News vs. Bad News

Intuition:

- In either case, in each period get "no news" w/ high prob, "news" w/ low prob.
- As $\Delta \to 0$, "no news" becomes uninformative (likelihood ratio difference$\to 0$), so cannot provide incentives through continuation value movement after no news. So continuation value must stay fixed after no news.
- Let v be the best SSE payoff.
- In bad news case, can decompose $v > 0$ into $u(EE) > v$, $w(\text{no news}) = w(\bar{y}) = v$, $w(\text{news}) = w(\underline{y}) < v$.
- In good news case, cannot decompose $v > 0$ into $u(EE) > v$, $w(\text{no news}) = w(\underline{y}) = v$, $w(\text{news}) = w(\bar{y}) < v$, as this violates IC. (And $w(\bar{y}) \geq v$ would violate PK.)

Normal Signals

Now suppose that $y \in \mathbb{R}$ is distributed according to cdf

$$F(y|a) = \begin{cases} \Phi_{1/\Delta}(y-1) & \text{if } a = EE \\ \Phi_{1/\Delta}(y) & \text{otherwise} \end{cases}$$

where $\Phi_{1/\Delta}$ is the $N(0, 1/\Delta)$ cdf. That is, $y = 1 + \varepsilon$ if $a = EE$ and $y = \varepsilon$ otherwise, where $\varepsilon \sim N(0, 1/\Delta)$.

- Interpretation: observe Brownian process every Δ periods, so variance is proportional to $1/\Delta$.
- Again, let $\delta = e^{-r\Delta}$ for fixed $r > 0$.

Theorem

With normal noise with variance inversely proportional to Δ, for any $r > 0$, there exists $\bar{\Delta} > 0$ such that, for every $\Delta < \bar{\Delta}$, there is no SSE with strictly positive payoffs.

Normal Signals (cntd.)

- Let v be the best SSE payoff.
- Best scheme to decompose v is a **tail test:** $\exists y^*$ s.t. $w(y) = v \; \forall y > y^*$, $w(y) = 0 \; \forall y < y^*$.
- Let ϕ and Φ be the standard normal pdf/cdf. Switching to z-scores, IC requires

$$\underbrace{\sqrt{\Delta}\phi(z^*)}_{\text{pivot prob}} \underbrace{\frac{v}{1-\delta}}_{\text{penalty}} \geq \underbrace{1}_{\text{effort cost}} .$$

- PK requires

$$\underbrace{\Phi(z^*)}_{\text{penalty prob}} \underbrace{\frac{v}{1-\delta}}_{\text{penalty}} \leq \underbrace{2}_{\text{max E[penalty]}} .$$

- Divide the first inequality by the second to get

$$\frac{\sqrt{\Delta}\phi(z^*)}{\Phi(z^*)} \geq \frac{1}{2}.$$

Normal Signals (cntd.)

$$\frac{\sqrt{\Delta}\phi\left(z^*\right)}{\Phi\left(z^*\right)} \geq \frac{1}{2}.$$

- $\phi\left(z^*\right)/\Phi\left(z^*\right)$ is the **score** of a normal tail test. It is approximately equal to $-z^*$ for $z^* < 0$.
- Hence, z^* must decrease at least linearly with $1/\sqrt{\Delta}$.
- But $\phi\left(z^*\right)$ decreases exponentially with z^*, and hence with $1/\sqrt{\Delta}$.
- Hence, $\sqrt{\Delta}\phi\left(z^*\right)$ decreases exponentially with Δ. But $1-\delta \approx r\Delta$ decreases only linearly with Δ. So IC is violated.

More generally, if allow $\sigma^2 \to \infty$ and $\delta \to 1$ at arbitrary rates, there is no SSE w/ positive payoffs if there exists $\rho > 0$ s.t. $(1-\delta)\exp\left(\sigma^{2(1-\rho)}\right) \to \infty$.

- To support positive payoffs, need $\delta \to 1$ "almost exponentially faster" than $\sigma^2 \to \infty$.

Application 1: Collusion w/ Flexible Production

Sannikov Skrzypacz 07 consider a duopoly model where firms choose quantities q_i every Δ units of time.

- Price equals $P\left(q_1+q_2\right)+\varepsilon$, where $\varepsilon \sim N\left(0, \sigma^2/\Delta\right)$.
  - Interpretation: Observe average price over Δ seconds, where each second price is hit by iid $N\left(0, \sigma^2\right)$ noise.
- Similar argument implies that collusion is impossible (i.e., profits equal static Cournot Nash) in SSE as $\Delta \to 0$.
- Sannikov Skrzypacz also show that collusion is imposible for any PPE as $\Delta \to 0$.
  - Intuition: pairwise full rank fails, so can't use value transfers to improve on value destruction.
- However, Rahman 14 shows that collusion is possible with more general sequential eqm.

Sannikov 07 and Sannikov Skrzypacz 10 develop general theory of PPE (a la APS, FLM) w/ continuous actions + Brownian noise.

Application 2: Large Populations

To model large populations of patient players, Sugaya Wolitzky 23 consider sequences of games where δ and N vary together.

- Payoffs can also vary, as long as each player's payoffs are uniformly bounded.

Assume product structure monitoring with $q < p < 1$ s.t.

$$\Pr\left(\bar{y}_i | a\right) = \begin{cases} p & \text{if } a_i = E \\ q & \text{if } a_i = S \end{cases}$$

Assume that $u_i\left(E, a_{-i}\right) - u_i\left(S, a_{-i}\right) \geq 1$ for all i, a_{-i}.

Large Populations (cntd.)

Theorem

Fix $q < p < 1$. For any sequence of (δ, N) s.t. there exists $\rho > 0$ satisfying $(1-\delta)\exp\left(N^{1-\rho}\right) \to \infty$, for any pair (δ, N) far enough along the sequence and any SSE, all players take S in every period.

- A tail test is again optimal: $\exists n$ s.t. $w(y) = v$ $\forall y : \#\{i : y_i = \bar{y}_i\} > n$, $w(y) = 0$ $\forall y : \#\{i : y_i = \bar{y}_i\} < n$.
- With fixed $q < p < 1$, $\mathrm{Var}(\#\{i : y_i = \bar{y}_i\})$ is proportional to N.
- To support effort, need $\delta \to 1$ almost exponentially faster than $\mathrm{Var} \to \infty$.

Arbitrary Equilibria and Signals

Above results are essentially for SSE.

However, it's obvious that to support non-Nash outcomes, players must be sufficiently patient relative to the signal informativeness, for any class of equilibria (not just SSE or PPE) and monitoring structures (not just public monitoring).

Several papers work out necessary conditions on δ vs. informativeness.

- Fudenberg Levine Pesendorfer 98, al-Najjar Smorodinsky 00, Awaya Krishna 16.
- Sugaya Wolitzky 23 give tight conditions: folk thm if $\delta \to 1$ faster than info measure$\to 0$ (under $\approx$pairwise full rank), myopic play if $\delta \to 1$ slower than info measure$\to 0$.

Active area of research, with interesting connections between repeated games and probability.