Repeated Games and Applications- Imperfect Public Monitoring

Wouter Vergote (FUSL and CORE)

FUSL and CORE

December 2009

A motivating example (FLM 1994): prisoners' dilemma with moral hazard

- two players, two un-observable actions: work or shirk
- utility = half total output minus cost of effort equal to 4
  - expected output is 12 if both work, 6 if one works and 0 if no one works
- Payoff table

		player 2	
		work	shirk
player 1	w	2,2	-1,3
	s	3,-1	0,0

- two different probability distributions:
  1. $y = (18, 0)$, with $\rho(18 \mid ww) = 2/3$, $\rho(18 \mid ws) = \rho(18 \mid sw) = 1/3$ and $\rho(18 \mid ss) = 0$
  2. $y = (18, 12, 0)$, with
     - $\rho(18 \mid ww) = 1/3$, $\rho(18 \mid ws) = 1/3$, $\rho(18 \mid sw) = 0$ and $\rho(18 \mid ss) = 0$
     - $\rho(12 \mid ww) = 1/2$, $\rho(12 \mid ws) = 0$, $\rho(12 \mid sw) = 1/2$ and $\rho(12 \mid ss) = 0$

A motivating example (FLM 1994): inefficieny

- Case 1
- Effcient payoffs in $\mathcal{F}'$: $v_1 + v_2 \geq 8/3$
- In any equilibrium of the repeayed game with discount factor less than 1, the discounted average payoffs are no larger than 2.
- Let $\gamma = \max \{v_1 + v_2 \mid (v_1, v_2) \in \mathcal{E}(\delta)\}$
- Choose $v \in \mathcal{E}(\delta)$ such that $v_1 + v_2 = \gamma$. Then we can decompose v so that:

$$v_i = (1-\delta)u_i(\alpha) + \delta(\rho(18 \mid \alpha)w_i(18) + \rho(0 \mid \alpha)w_i(0))$$

- But then

$$w_1(18) - \frac{3(1-\delta)}{\delta} \geq w_1(0)$$

A motivating example (FLM 1994): inefficieny

- Hence

$$\begin{aligned}
v_1 &\leq (1-\delta)u_1(\alpha)+\delta\left[\begin{array}{c}\rho(18\mid\alpha)w_1(18)+\\ \rho(0\mid\alpha)\left(w_1(18)-\frac{3(1-\delta)}{\delta}\right)\end{array}\right]\\
v_1 &\leq (1-\delta)((3\mu_2-\mu_1)-(3-\mu_2-\mu_1))+\delta w_1(18)\\
v_1 &\leq (1-\delta)(4\mu_2-3)+\delta w_1(18)\\
v_2 &\leq (1-\delta)(4\mu_2-3)+\delta w_2(18)\\
\gamma &\leq (1-\delta)2+\delta\left[w_1(18)+w_2(18)\right]\\
\text{Since } \gamma &> 2 \implies w_1(18)+w_2(18)>\gamma \text{ a contradiction}
\end{aligned}$$

A motivating example (FLM 1994): prisoners' dilemma with moral hazard

- At the heart of inefficiency: In order to induce (work, work) we need to punish both players if $y = 0$.
- example 2 solves this issue
    - three outcomes
    - players' actions induce DIFFERENT distributions of public outcomes

Another motivating example: collusion in auctions

- Two bidders. Private values for bidder i :

$$v_i \sim U[0,1] \text{ identically and independently distributed (i.i.d.)}$$
$$\rightarrow F(v) = v, \frac{dF(v)}{dv} = f(v) = 1$$

- Nash Equilibrium? Choosing bid b, expected profits are: $\pi_i(b) = prob(b > b_j)(v_i - b)$

We can solve to obtain: $b^*(v) = \frac{v}{2}$, $u_i^{NE} = \frac{1}{6}$

- Best Collusive outcome: highest bidder obtains good and bids zero, which yields expected payoffs $u_i^C = \frac{1}{3}$.Can this be obtained when $\delta \rightarrow 1$ in the repeated auction?

Imperfect public monitoring: the stage game

- Players $i : 1, ..., N$
- Finite (compact) action sets $A_i \subset \mathbb{R}^k$ for some k,
  - $a_i \in A_i$, $m_i = |A_i|$
  - $A = \prod_i A_i$
  - α_i is mixed action for i
  - set of mixed actions of player $i : \Delta(A_i)$;
  - $\Delta(A) = \prod_i \Delta(A_i)$ is the set of mixed action profiles

The stage game: public signal

- Public signal or outcome y, after players have chosen actions, is drawn from a finite signal space Y, where $m = |Y|$
- Denote $\rho(y \mid a)$ the probability that signal y realizes given action profile a.
  - $\rho : Y \times A \to [0, 1]$ is continuous
  - $\rho(y \mid \alpha)$ is extension to mixed action profiles:

$$\rho(y \mid \alpha) = \sum_{a \in A} \rho(y \mid a)\alpha(a)$$

$$\text{where } \alpha(a) = \prod_{i=1}^{n} \alpha_i(a_i)$$

The stage game: public signal and expected payoffs

- each player $i's$ realized payoff depends on its *own* action and the public signal: $u_i^*(y, a_i)$
- Expected payoffs for player i given profile a :

$$u_i(a) = \sum_{y \in Y} \rho(y \mid a) u_i^*(y, a_i)$$

- Mixed actions:

$$u_i^*(y, \alpha_i) = \sum_{a_i \in A_i} \alpha_i(a_i) u_i^*(y, a_i)$$

$$u_i(\alpha) = \sum_{y \in Y} \sum_{a \in A} \rho(y \mid a) \alpha(a) u_i^*(y, a_i)$$

Stage Game: Assumptions

1. A_i is finite or a continuum action space: a compact and convex subset of $\mathbb{R}^k$ for some k.
2. Y is finite, and if A_i is a continuum action space, then $\rho : Y \times A \rightarrow [0,1]$ is continuous
3. if A_i is a continuum action space, then $u_i^* : Y \times A_i \rightarrow \mathbb{R}$ is continuous and u_i is quasiconcave in a_i.

We will focus, as in FLM (1994) on case where A_i is finite

The Repeated Game

- period t history of public signals, h^t is $\{y^0, y^1, ..., y^{t-1}\}$:
- the set of public histories $\mathcal{H} \equiv \bigcup_{t=0}^{\infty} Y^t$ with $Y^0 \equiv \{\varnothing\}$, typical element: $h^t \in Y^t$
- Set of all possible histories for player i: $\mathcal{H}_i \equiv \bigcup_{t=0}^{\infty} (A_i \times \mathcal{H})^t$, typical element: $h_i^t \in (A_i \times Y)^t$
- A pure strategy for player $i : \sigma_i : \mathcal{H}_i \rightarrow \ A_i$
- A mixed strategy for player $i : \sigma_i : \mathcal{H}_i \rightarrow \ \Delta(A_i)$
- When the signal is random, a pure strategy profile does not induce a unique outcome path.

The Repeated Game: Payoffs

- Players aim to maximize the discounted sum of stage game payoffs $=$ discounted average sum of payoffs. If $\{u_i^t\}$ is a sequence of stage game payoffs, one maximizes:

$$(1-\delta)\sum_{t=0}^{\infty}\delta^t u_i^t$$

- Given a pure strategy profile σ we obtain, given that $a^0 = \sigma(\emptyset)$,

$$U_i(\sigma) = (1-\delta)u_i(a^0) + (1-\delta)\delta \sum_{y\in Y} \rho(y \mid a^0) u_i^*(\sigma_1(h_1^1), \ldots, \sigma_n(h_n^1)) + \ldots$$

Scope of public monitoring

- perfect monitoring
- moral hazard
- adverse selection (private information)
  - example: collusion in auctions
- imperfect observability
- ...

Recursive Structure: Public strategies and Perfect Public Equilibrium

Definition

A (behavior) strategy $\sigma_i = \{\sigma_i^t\}$ is public if, for all t, σ_i^t only depends on h^{t-1} and NOT on h_i^{t-1}.

If the strategy is not public, it is private.

Lemma

(7.1.1) If all players other than i play a public strategy, then player i has a public strategy as a best reply.

Definition

A Perfect Public Equilibrium (PPE) is a profile of public strategies σ that, for any $h^t \in \mathcal{H}$,specifies a NE for the repeated game: for all t and all $h^t \in Y^t, \sigma_{|h^t}$ is a NE.

Recursive Structure: Public strategies and Perfect Public Equilibrium

- Deviations to a non-public strategy are irrelevant.
- restricting attention to public strategies, every public history induces a continuation game that is strategically identical.
- for a public strategy σ, history h_t induces a well defined continuation public strategy, $\sigma_i\mid_{h_t}$
- Automaton representation of a public strategy.
  - $(\mathcal{W}, w^0, f, \tau)$ where $\mathcal{W}$ is a set of states, w^0 is the initial state, a decision function $f : \mathcal{W} \to \prod_i \Delta(A_i)$ and a transition function $\tau : \mathcal{W} \times Y \to \mathcal{W}$.
  - extend the domain of τ to $\mathcal{W} \times \mathcal{H} \setminus \{\emptyset\}$ by recursively defining

$$\tau(w, h^t) = \tau(\tau(w, h^{t-1}), y^{t-1}))$$

  - $w' \in \mathcal{W}$ is accessible from another state $w \in \mathcal{W}$ if there exists a sequence of public signals, h^t, such that $w' = \tau(w, h^t)$.

One shot deviation principle

Lemma

(prop 7.1.1.) σ *is a PPE iff there are no profitable one shot deviations: iff for all* $h^t \in Y^t, \sigma(h^t)$ *is a NE of the normal form game with payoffs:*

$$g_i(a) = (1-\delta)u_i(a) + \delta \sum_{y \in Y} \rho(y \mid a) U_i(\sigma \mid_{h^t, y})$$

Denote the set of PPE payoff vectors as $\mathcal{E}(\delta) \subset \mathbb{R}^n$.
$\Longrightarrow$ Can we obtain a folk theorem along the lines of the perfect monitoring case?

Enforceability and Self-generation

Definition

Let δ and $W \subset \mathbb{R}^n$ be given. Profile α is enforceable with respect to δ and W if there exists $v \in \mathbb{R}^n$ and a function $\gamma : Y \to W$ such that for all i :

$$\begin{aligned} v_i &= (1-\delta)u_i(a_i, \alpha_{-i}) + \delta \sum_y \rho(y \mid a_i, \alpha_{-i})\gamma_i(y) \\ &: \text{ for all } a_i \text{ such that } \alpha_i(a_i) > 0 \\ v_i &\geq (1-\delta)u_i(a_i, \alpha_{-i}) + \delta \sum_y \rho(y \mid a_i, \alpha_{-i})\gamma_i(y) \\ &: \text{ for all } a_i \text{ such that } \alpha_i(a_i) = 0 \end{aligned}$$

We say that γ enforces α on W. It describes the expected payoffs from future play (continuation promises).

Enforceability and Self-generation

Definition

We say that $\{\gamma(y)\}_{y \in Y}$ enforces α with respect to δ and W and that v is decomposable with respect to α, δ and W.

Definition

Let $B(W, \delta)$ be the set of all payoff vectors that are decomposable as we vary the profile α with respect to δ and W.

Definition

If $W \subset B(W, \delta)$ then W is self-decomposable (self generating).

Theorem

(prop 7.3.1) APS (1990) have shown that if $W \subset \mathbb{R}^n$ is bounded and $W \subset B(W, \delta)$, then $W \subset B(W, \delta) \subset \mathcal{E}(\delta)$

Factorization

Theorem

(Prop 7.3.2) $\mathcal{E}(\delta) = B(\mathcal{E}(\delta))$

Proof.

1. If $\mathcal{E}(\delta) \subset B(\mathcal{E}(\delta))$ boundedness of $\mathcal{E}(\delta) \Rightarrow B(\mathcal{E}(\delta)) \subset \mathcal{E}(\delta)$
2. Suppose $B(\mathcal{E}(\delta)) \subset \mathcal{E}(\delta)$. Take $v \in \mathcal{E}(\delta)$ with value $v = U(\sigma)$ and let $\alpha \equiv \sigma(\emptyset)$ and $\gamma(y) = U(\sigma_{|y})$.Then

$$\begin{aligned} V(\alpha, \gamma) &= (1-\delta)u(\alpha) + \delta \sum_y \rho(y \mid \alpha)\gamma(y) \\ &= (1-\delta)u(\alpha) + \delta \sum_y \rho(y \mid \alpha) U(\sigma_{|y}) \\ &= U(\sigma) = v \end{aligned}$$

Since σ is a PPE, $\sigma_{|y}$ is also a PPE and so $\gamma : Y \rightarrow \mathcal{E}(\delta)$. σ is a PPE, so there are no profitable one shot deviations and hence α is enforced by γ on $\mathcal{E}(\delta)$ and $v \in B(\mathcal{E}(\delta))$. □

Factorization

Lemma

B is a monotone operator: $W' \subset W \Longrightarrow B(W') \subset B(W)$

Lemma

If W is compact, $B(W)$ is closed

- We have that $\mathcal{E}(\delta) \subset B^k(\mathcal{F}') \subset \mathcal{F}'$ for all k
- In fact $\left\{B^k(\mathcal{F}')\right\}_k$ is a decreasing sequence. Define $\mathcal{F}'_\infty = \bigcap_k^\infty B^k(\mathcal{F}')$, which is compact and non-empty.

Theorem

$\mathcal{F}'_\infty$ is self generating and so $\mathcal{F}'_\infty = \mathcal{E}(\delta)$.

- As a consequence $\mathcal{E}(\delta)$ is compact (FLM 1994, Lemma 4.1.)

Monotonicity of PPE payoffs

Lemma

Suppose $0 < \delta_1 < \delta_2 < 1$. $W \subset W'$ *and* $W \subset B(W'; \delta_1, \alpha)$ *for some* α. *Then* $W \subset B(co(W'); \delta_2, \alpha)$.

- Consequence: if $W \subset W'$ and $W \subset B(W'; \delta_1)$, then $W \subset B(co(W'); \delta_2)$

Local Self-Decomposability

Definition

(def 7.3.5) $W \subset \mathbb{R}^n$ is locally self-decomposable (locally self-generating) if for each $v \in W$ there exists a $\delta < 1$ and an open set U_v containing v such that $U \cap W \subseteq B(W, \delta)$.

Lemma

(lemma 4.2. FLM, Corr 7.3.3.)If $W \subset \mathbb{R}^n$ is compact, convex and locally self-decomposable, then $\exists \delta' < 1$ such that $W \subset \mathcal{E}(\delta)$ for all $\delta \in \left(\delta', 1\right)$

Main Idea

MAIN QUESTION: **under what conditions is a smooth interior set of V^* achievable in equilibrium for some $\delta < 1$?**

- If any smooth interior subset W of V^* is locally Self-Decomposable: Approximate Folk Theorem. WHY?

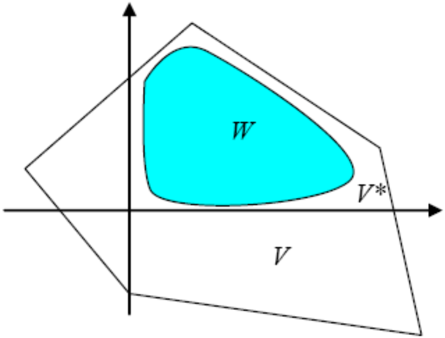

Main Idea

MAIN QUESTION: **under what conditions is a smooth interior set of V^* achievable in equilibrium for some $\delta < 1$?**

- If any smooth interior subset W of V^* is decomposable on (regular) tangent hyperplanes $\Longrightarrow$ locally self-decomposable
- RANK conditions: If a profile α
  - has pairwise full rank for every pair of players
  - is enforceable pure strategy profile that is pairwise-identifiable for every pair of players

  then α is enforceable with respect to all regular hyperplanes
- idea: action profiles need to be 'statistically' identifiable in order to be able to punish some players while rewarding others, a key feature in order to obtain the folk theorem (see example)

Enforceability w.r.t. hyperplanes

Definition

A hyperplane is a set of the form $P = v' + H = \{v' + w \mid w \in H\}$ where v' is a fixed vector and H is a (n-1) dimensional linear subspace of $\mathbb{R}^n$.P can be expressed as:$P = \{v \in R^n \mid \sum \beta_i v_i = c\}$ for some non-zero vector $(\beta_1, ..., \beta_n)$ and a constant c.

Lemma

Suppose P is a hyperplane in $\mathbb{R}^n$ containing the origin. If α is enforceable with respect to some δ and P, then

1. *α is enforceable w.r.t. any translate $P' = v' + P$ and any $\delta' > 0$*
2. *$\exists$ a constant k such that, $\forall\ \delta' > 0$ and $v' \in \mathbb{R}^n$, $\exists \gamma' : Y \rightarrow P'$ such that*
   - *$\{\gamma'(y)\}_{y \in Y}$ enforces α with respect to δ' and v'*
   - *$v' = \sum_y \rho(y \mid \alpha)\gamma'(y)$ and $\|\gamma'(y) - v'\| < k\frac{1-\delta'}{\delta'}$ for all $y \in Y$.*

Enforceability and Self-generation

Definition

(FLM def. 4.4) A smooth subset $W \subset V'$ is decomposable on tangent hyperplanes if $\forall\ w \in \partial W\ \exists \alpha$ s.t.(i) $u(\alpha)$ is separated from W by the (unique) hyperplane P_w that is tangent to W at w and (ii) $\exists \{w(y)\}_{y \in Y}$ in P_v that enforce α.

Theorem

(Theorem 4.1) If a smooth set $W \subset \mathbb{R}^n$ is decomposable on tangent hyperplanes, then it is locally self decomposable. Hence, $\exists \delta' < 1$ such that $W \subset \mathcal{E}(\delta)$ for all $\delta \in (\delta', 1)$.

- If any w on the boundary of W is generated using continuation values on a tangent hyperplane, then W is achievable in equilibrium for δ sufficiently close to 1.

What does it mean for w to be generated with continuation values on a tangent hyperplane?
$\Longrightarrow$ there is a mixed action profile α with payoff $u(\alpha)$ separated from V^* by the tangent hyperplane and transitions $\gamma(y)$ for each signal $y \in Y$ such that:

- promise keeping:

$$w = (1-\delta)u(\alpha) + \delta \sum_y \rho(y \mid \alpha) \left(w + \frac{(1-\delta)\gamma(y)}{\delta} \right)$$

Enforceability w.r.t. hyperplanes: Intuition

- such that α is IC given those continuation values:actions in the support of α maximize:

$$w_i = (1-\delta)u_i(\alpha) + \delta \sum_y \rho(y \mid \alpha) \left(w_i + \frac{(1-\delta)\gamma_i(y)}{\delta} \right)$$

- and scuh that $\gamma(y) \bullet n =$constant for all $y \in Y$, where n is the unit normal vector at w :

$$\begin{bmatrix} \beta_1 \\ \beta_n \end{bmatrix} \begin{bmatrix} w_1 - \gamma_1(y) \\ w_n - \gamma_n(y) \end{bmatrix} = 0$$

Enforceability w.r.t. hyperplanes: Intuition

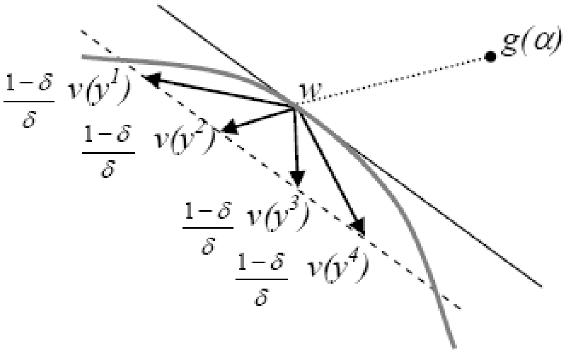

- Note that $w + \frac{(1-\delta)\gamma(y)}{\delta} \in W$ for δ sufficiently close to 1.

Enforceability with respect to hyperplanes and Rank conditions

- link between enforceabilty wrt hyperplanes and the ability to identify one another's actions statistically on the basis of public outcomes
- One can think of $\gamma(. \mid ., \alpha_{-i}) = \Gamma_i(\alpha_{-i})$ as a $m_i \times m$ matrix whose rows correspond to actions a_i and columns to public outcomes y.

Definition

The profile α has individual full rank for player i if $\Gamma_i(\alpha_{-i})$ has rank equal to m_i. If this is true for all i, then α is individual full rank.

Lemma

If α is individual full rank, then α is enforceable

- If, given α_{-i}, we can identify player $i's$ actions statistically, then we can choose continuation payoffs to make i indifferent between all his actions.

Rank conditions

Definition

The profile α has pairwise full rank for player i if $\Gamma_{ij}(\alpha) = \binom{\Gamma_i(\alpha_{-i})}{\Gamma_j(\alpha_{-j})}$ has rank equal to $m_i + m_j - 1$.

When satisfied, any deviation of different players can be distinguised statistically and any two deviations of the same player can be statistically distinguished.

Lemma

If the profile α has pairwise full rank for players i and j, then it is strongly enforceable wrt to all $ij-$pairwise hyperplanes.

Proof.

idea behind the proof: the incentive constraints form a system of linear equations (with equality) in which payoffs are unknown, which must have a solution if the rank conditions are satisfied □

Folk theorem

CONDITION I Every pure-action, pareto-efficient profile is pairwise identifiable for all players

Lemma

Any Pareto-efficient profile is enforceable

CONDITION II For all pairs i, j there exists a profile α that has pairwise full rank for that pair.

Lemma

If condition II holds, then there exists an open and dense set of profiles each of which has pairwise full rank for all pairs of players

Theorem

Suppose either condition I or II holds. Consider a NE α^N of the stage game. Let V^0 be the convex hull opf the set consisting of $u^*(\alpha^N)$ and the PE vectors pareto-dominating $u^*(\alpha^N)$.Let W be a smooth set in the interior of V^0.Then there exists $\delta^* < 1$ such that for all $\delta > \delta^*$, $W \subset \mathcal{E}(\delta)$.

Games with a product structure

Definition

A game has a product structure if we can write $y = (y_1, y_2, ..., y_n)$ for all $y \in Y$. where, for all i and $\alpha = (\alpha_i, \alpha_{-i})$, $\rho_i(y_i \mid \alpha)$ satisfies:

$$\rho_i(y_i \mid \alpha) = \rho_i(y_i \mid \alpha_i, \alpha'_{-i}) \text{ for all } \alpha'_{-i}$$

$$\rho(y \mid \alpha) = \prod_{i=1}^{n} \rho_i(y_i \mid \alpha) \text{ for all } y$$

Lemma

If the game has a product structure, then any pure-action profile is pairswise identifiable for all pairs of players.

Theorem

Suppose the game has a product structure and let W be a smooth set in the interior of V^0. Then there exists $\delta^* < 1$ such that for all $\delta > \delta^*$, $W \subset \mathcal{E}(\delta)$.

Adverse selection

- each player i draws a type $z_i \in Z_i$. We suppose that Z_i is finite with distribution $\widehat{\rho}_i$.
- after learning their types, players move simultaneously, and player $i's$ move $y_i \in Y_i$ is publicly observable.
- A (pure) action for player i is a mapping $a_i : Z_i \to Y_i$.
- We suppose that the (realized) payoff of player i depends on her own type and the public outcome only: $u_i^*(z_i, y)$.
- Example: a (first) price auction

Adverse selection

- We can write the ex-ante payoff as follows

$$u_i(a_i, y) = \frac{\sum_{z_i \in a_i^{-1}(y_i)} u_i^*(z_i, y)\widehat{\rho}_i(z_i)}{\sum_{z_i \in a_i^{-1}(y_i)} \widehat{\rho}_i(z_i)}$$

- Distribution of outcomes corresponding to action profile a :

$$\rho(y \mid a) = \prod_{i=1}^{n} \rho_i(y_i \mid a_i) \text{ where}$$

$$\rho_i(y_i \mid a_i) = \sum_{z_i \in a_i^{-1}(y_i)} \widehat{\rho}_i(z_i) \text{ if } a_i^{-1}(y_i) \text{ is non-empty, } 0 \text{ otherwise}$$

- Thus player $i's$ ex ante expected payoff from a :

$$u_i(a) = \sum_{y \in Y} \rho(y \mid a) u_i(a_i, y)$$