A Folk Theorem for Repeated Games with Imperfect Monitoring

R. McLean

I. Obara

A Postlewaite

This work deals with with communication in discounted repeated games with private monitoring. Important antecedents are

- Aoyagi, M., *Economic Theory, 2005*
- Ben-Porath, E. and M. Kahneman, *Journal of Economic Theory,* 1996.
- Compte, O., *Econometrica, 1998.*
- Fudenberg, D. and D. Levine, *Journal of Economic Theory*, 2007.
- Kandori, M. and H. Matsushima, *Econometrica*, 1998.
- Obara, I., *Journal of Economic Theory,* 2009.
- Tomala, T. ,*Games and Economics Behavior*, 2009.

The approach in this paper is presents a communication extension of a repeated game with imperfect private monitoring that is different from those in the previously cited literature and relies on the seminal results for games with imperfect public monitoring of

- Fudenberg, Levine and Maskin, *Econometrica*, 1994

- Abreu, Pearce and Stacchetti, *Econometrica* , 1990

and the mechanism design machinery developed in

- McLean and Postlewaite, *Econometrica*, 2002

Repeated games with imperfect public monitoring: Basics

The set of players: $N = \{1, ..., n\}$.

Player i chooses an action from a finite set A_i. An action profile is denoted by $a = (a_1, ..., a_n) \in \Pi_i A_i := A$.

Actions are not publicly observable, but the players observe a public signal from a finite set Y.

The probability that $y \in Y$ is realized given $a \in A$ is denoted $\pi(y|a)$.

Player i's stage game payoff is $u_i(a_i, y)$

Player i's expected stage game payoff is

$$g_i(a) = \sum_y u_i(a_i, y)\, \pi(y|a).$$

This stage game is denoted by (G, π) where $G = (N, A, g)$.

We normalize payoffs so that each player's *pure strategy* minmax payoff is 0.

The set of feasible payoff profiles is

$$V(G) = co\{g(a) | a \in A\}$$

and

$$V^*(G) = \{v \in V | v \gg \mathbf{0}\}$$

is the set of feasible, strictly individually rational payoff profiles.

Histories and strategies: imperfect public monitoring

Private history for player i at stage t : $h_i^t = \left(a_i^0, ..., a_i^{t-1}\right) \in H_i^t = A_i^t$

Public history at stage t: $h^t = \left(y^0, ..., y^{t-1}\right) \in H^t = Y^t$ with $H_i^0 = H^0 := \{\emptyset\}$.

A pure strategy for player i : $\alpha_i = \left\{\alpha_i^t\right\}_{t=0}^{\infty}$, with $\alpha_i^t : H_i^t \times H^t \to A_i$.

Payoffs and equilibria: Imperfect public monitoring

A pure strategy profile induces a probability measure on A^{∞}. Player i's discounted expected payoff given α and $\delta \in (0,1)$ is

$$w_i^{\alpha,\delta} = (1-\delta)\sum_{t=0}^{\infty}\delta^t E\left[g_i\left(\tilde{a}^t\right)|\alpha\right].$$

We denote this repeated game associated with (G,π) by $G_{\pi}^{\infty}(\delta)$.

A strategy is *public* if it only depends on H^t.

A profile of public strategies is a *perfect public equilibrium* (PPE) if, after every public history, the continuation (public) strategy profile constitutes a Nash equilibrium (Fudenberg, Levine, and Maskin).

Given α, δ and history $h^{t+1} = \left(h^t, y\right) \in H^{t+1} = H^t \times Y$, let $w_i^\alpha\left(h^t, y\right)$ denote player i's continuation payoff from period $t+1$.

Definition: A pure strategy profile α is a perfect public equilibrium (PPE) for $G_\pi^\infty(\delta)$ if for all $h^t \in H^t, t \geq 0$, $a_i' \neq \alpha_i^t\left(h^t\right)$, and $i \in N$.

$$\begin{aligned} & (1-\delta)\, g_i\left(\alpha^t\left(h^t\right)\right) + \delta \sum_{y \in Y} w_i^\alpha\left(h^t, y\right) \pi\left(y | \alpha^t\left(h^t\right)\right) \\ \geq \; & (1-\delta)\, g_i\left(a_i', \alpha_{-i}^t\left(h^t\right)\right) + \delta \sum_{y \in Y} w_i^\alpha\left(h^t, y\right) \pi\left(y | a_i', \alpha_{-i}^t\left(h^t\right)\right) \end{aligned}$$

Remark: PPE reduces to PE in games with perfect monitoring.

Let $E(\delta)$ denote the set of PPE payoff profiles.

FLM prove a PPE folk theorem when π satisfies certain "distiguishability" conditions.

Definition (informal): π satisfies distinguishabilty if, given any pair of players i and j, (i) a deviation by either player is statistically detectable and (ii) a deviation by one player can be distinguished from a deviation by the other. FLM refer to these as individual full rank and pairwise full rank.

Theorem (FLM): Suppose that (G, π) satisfies distinguishability. If the feasible set $V^*(G)$ has non-empty interior, then for every compact, convex smooth set $W \subseteq intV^*(G)$, there exists $\underline{\delta} \in (0,1)$ such that, $W \subseteq E(\delta)$ for each $\delta \in (\underline{\delta}, 1)$.

Repeated Games with Imperfect Private Monitoring:Basics

The set of players: $N = \{1, ..., n\}$.

Player i chooses an action from a finite set A_i. An action profile is denoted by $a = (a_1, ..., a_n) \in \Pi_i A_i := A$.

Actions are not publicly observable, but the players observe a private signal from a finite set S_i. A private signal profile is denoted $s = (s_1, .., s_n) \in \Pi_i S_i := S$.

The probability that $s \in S$ is realized given $a \in A$ is denoted $p(s|a)$.

Let

$$p(s_{-i}|a, s_i) := \frac{p(s_i, s_{-i}|a)}{p(s_i|a)}$$

denote the conditional probability of $s_{-i} \in S_{-i}$ given (a, s_i).

Player i's stage game payoff: $v_i(a_i, s_i)$

Player i's expected stage game payoff is

$$g_i'(a) = \sum_s v_i(a_i, s_i)\, p(s|a)$$

We denote this private monitoring stage game by (G', p), where $G' = (N, A, g')$.

Let $V(G')$ and $V^*(G')$ be the feasible payoff set and the set of individually rational and feasible payoffs for G'. Discounted average payoffs are defined as in the public monitoring case. Let $G_p'^{\infty}(\delta)$ be the corresponding repeated game with private monitoring given $\delta \in (0, 1)$.

Histories and strategies: Imperfect private monitoring

private history for player i at stage t:

$$h_t^i \in H_i^t = (a_i^0, ..., a_i^{t-1} s_i^0, ..., s_i^{t-1}) \in A_i^t \times S_i^t$$

A pure strategy for player i : $\alpha_i = \left\{\alpha_i^t\right\}_{t=0}^{\infty}$, with $\alpha_i^t : H_i^t \longrightarrow A_i$.

Set of pure strategies for player i : $\mathbf{\Sigma}_i$.

Strategy profile: $\alpha = \{\alpha_i\}_{i \in N} \in \mathbf{\Sigma} := \times_i \mathbf{\Sigma}_i$.

Public Communication extension

We consider a communication extension of the game (G', p).

A *public coordination device* is a function $\phi : S \rightarrow \Delta(Y)$ where Y is a finite set of public signals.

A *public communication device* for (G', p) is a collection

$$\Phi = \{\phi_{h^t} : h^t \in Y^t, t \geq 0\}$$

where each $\phi_{h^t} : S \rightarrow \Delta(Y)$ is a *public coordination device*.

Note: The stage t output of the public communication device only depends on the history $h^t \in Y^t$ of public signals up to stage t. Consequently, Φ is special type of autonomous ccommunication device in the sense of Forges.

How does play proceed in the public communication extension?

$t = 0$:

• At the start of period 0, players choose an action profile $a^0 \in A$

• Players then receive a private signal profile $s^0 \in S$ with probability $p(\cdot|a^0) \in \Delta(S)$

• Contingent on (a_i^0, s_i^0), i makes a public report r_i^0

• Given the report profile r^0, the public communication device chooses a public signal y^0 according to probability $\phi(\cdot|r^0)$.

At the start of period $t \geq 1$, we have;

public signal history: $h^t \in H^t = Y^t$

public reporting history: $h_R^t \in H_R^t = S^t$

private history for player i: $h_t^i \in H_i^t = A_i^t \times S_i^t$

Pure strategy for player i: $\sigma_i = (\alpha_i, \rho_i)$ where $\alpha_i = (\alpha_i^0, \alpha_i^1, ...), \rho_i = (\rho_i^0, \rho_i^1, ...)$ and for each t

$$\alpha_i^t : H_i^t \times H^t \times H_R^t \longrightarrow A_i$$

is i's "action strategy" and

$$\rho_i^t : H_i^t \times H^t \times H_R^t \times A_i \times S_i \longrightarrow S_i$$

$t \geq 1$:

period t begins with i-private histories h_i^t, a public signal history h^t and a public reporting history h_R^t

• At the start of period t, player i chooses an action profile $a^t = \alpha^t \left(h_i^t, h^t, h_R^t\right)$

• Players then receive private signal profile s^t according to the distribution $p(\cdot|a^t) \in \Delta(S)$.

• Player i makes a public announcement $r_i^t = \rho_i^t(h_i^t, h^t, h_R^t, a_i^t, s_i^t)$

• Given the report profile r^t, the public communication device chooses a public signal y^t according to the probability $\phi_{h^t}(\cdot|r^t)$.

Strategies and payoffs: The public communication extension

As in the repeated game without communication, pure strategies induce probability measures on A^{∞}.

Player i' s discounted expected payoff in $G_p'^{\infty}(\delta, \Phi)$ is

$$w_i^{\sigma}(\Phi) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t E\left[g_i'\left(\tilde{a}^t\right) | \sigma, \Phi\right]$$

A strategy $\sigma_i = (\alpha_i, \rho_i)$ for player i is *truthful* if player i reports her private signal truthfully whenever she played according to α_i in the same period, i.e.,

$$\rho_i^t \left(h_i^t, h^t, h_R^t, \alpha_i^t \left(h_i^t, h^t, h_R^t\right), s_i\right) = s_i$$

for every $\left(h_i^t, h^t, h_R^t\right)$ and s_i. Note that we allow players to lie immediately after a deviation in action. That is, we allow $\rho_i^t \left(h_i^t, h^t, h_R^t, a_i, s_i\right) \neq s_i$ if $a_i \neq \alpha_i^t \left(h_i^t, h^t, h_R^t\right)$.

A strategy $\sigma_i = (\alpha_i, \rho_i)$ is *public* if α_i^t only depends on $h^t = \left(y^0, ..., y^{t-1}\right) \in H^t$ and ρ_i^t depends only on $\left(h^t, a_i, s_i\right)$.

A strategy profile σ is $\eta-$*uniformly strict perfect public equilibrium with communication* if two conditions are satisfied:

First, player i would lose at least η in term of discounted average payoff at any stage when she deviates from the *equilibrium action*. In particular,

$$(1-\delta)\, g_i\left(\alpha^t\left(h^t\right)\right) + \delta \sum_{s\in S}\left[\sum_{y\in Y} w_i^\sigma\left(h^t, y\right)\phi_{h^t}\left(y|s_i, s_{-i}\right)\right] p(s|\alpha^t\left(h^t\right)) - \eta \geq$$

$$(1-\delta)\, g_i\left(a_i, \alpha_{-i}^t\left(h^t\right)\right) + \delta \sum_{s}\left[\sum_{y} w_i^\sigma\left(h^t, y\right)\phi_{h^t}\left(y|f_i(s_i), s_{-i}\right)\right] p(s|a_i, \alpha_{-i}^t\left(h^t\right)$$

for all $h^t \in H^t, t \geq 0$, $a_i \neq \alpha_i^t\left(h^t\right)$, $f_i : S_i \to S_i$ and $i \in N$.

Second, for every public history h^t, after players have chosen the equilibrium action profile $\alpha^t\left(h^t\right)$ in stage t, no player has an incentive to misreport his private signal to the public communication device: for each $s_i \in S_i$,

$$\sum_{s-i}\left[\sum_y w_i^\sigma\left(h^t, y\right)\phi_{h^t}\left(y|s_i, s_{-i}\right)\right]p(s_{-i}|\alpha^t\left(h^t\right), s_i)$$
$$\geq \sum_{s_{-i}}\left[\sum_y w_i^\sigma\left(h^t, y\right)\phi_{h^t}\left(y|s_i', s_{-i}\right)\right]p(s_{-i}|\alpha^t\left(h^t\right), s_i)$$

This incentive compatibility requirement is the main technical hurdle to be overcome.

If IC were not an issue, i.e., if players always submitted honest public reports, then a folk theorem is straightforward by applying the ideas in FLM.

Now we have the ingredients for a proof strategy for a folk theorem for the private monitoring game (G', p).

Step 1: Choose $v \in intV^*(G')$. Suppose that ϕ is a public coordinating device and **assume** that players' public announcements of their signals are truthful at each stage. If p^{ϕ} satisfies the FLM distinguishability conditions, then v is enforceable as a PPE of a game with imperfect public monitoring where $\pi = p^{\phi}$.

Step2: The device ϕ can be perturbed at each history induced by the PPE of step 1 so as to ensure honest announcements at each history. These perturbations of ϕ define a public communication device $\Phi = \{\phi_{h^t} : h^t \in Y^t, t \geq 0\}$ that yields the desired folk theorem.

Step 2 can be interpreted as an implementation problem where

Y = public signals = set of social outcomes

$w_i(y)$ = i's continuation payoff = player i's evaluation of outcome $y \in Y$

ϕ = social choice rule that chooses outcome $y \in Y$ with probability $\phi(y|s)$ when players report the profile s.

The SCR is IC if

$$\sum_{s_{-i}} \left[\sum_y w_i(y)\, \phi(y|s_{-i}, s_i) \right] p(s_{-i}|a, s_i) \geq \sum_{s_{-i}} \left[\sum_y w_i(y)\, \phi(y|s_{-i}, s_i') \right] p(s|a, s_i).$$

How do we implement the SCR ϕ without transfers?

Replace ϕ with a SCR that is "close" to ϕ and then identify conditions under which the perturbed SCR is IC.

With prob $1 - \lambda$, y is chosen with prob $\phi(y|s)$.

With prob $\frac{\lambda}{n}$, one player is chosen for "scrutiny".

Suppose player j is chosen for scrutiny,

Then

with prob $\psi_j(a, s)$, y is chosen with prob $\overline{\mu}_j(y)$ where supp$\overline{\mu}_j = \arg\max_{y'} w_j(y')$

with prob $1-\psi_j(a, s)$, y is chosen with prob $\underline{\mu}_j(y)$ where supp$\underline{\mu}_j = \arg\min_{y'} w_j(y')$

This yields a perturbed SCR ϕ^λ where

$$\begin{aligned}\phi^\lambda(y|s) &= (1-\lambda)\phi(y|s) + \lambda \underbrace{\left[\frac{\sum_{j=1}^n \left[\psi_j(a,s)\overline{\mu}_j(y) + (1-\psi_j(a,s))\underline{\mu}_j(y)\right]}{n}\right]}_{\phi'(y|s)} \\ &= (1-\lambda)\phi(y|s) + \lambda\phi'(y|s)\end{aligned}$$

To ensure IC of the perturbed SCR/coordinating device ϕ^λ, two ideas come into play.

A player's incentive to misreport should diminish as the player's perceived influence on the public coordinating signal ϕ becomes "smaller."

The following index measures the size of this influence for each player.

Definition: Player i's **informational influence** $v_i^\phi\left(s_i, s_i', a\right)$ given a public coordinating device ϕ and $\left(s_i, s_i', a\right) \in S_i \times S_i \times A$ is defined as

$$v_i^\phi\left(s_i, s_i', a\right) = \sum_{s_{-i}} \left\|\phi\left(\cdot|s_i, s_{-i}\right) - \phi\left(\cdot|s_i', s_{-i}\right)\right\| p(s_{-i}|a, s_i).$$

If $v_i^\phi\left(s_i, s_i', a\right)$ is small, then conditional on $\left(s_i, a\right)$, player i 's conditional expected "influence" on the public signal distribution is small.

Small informational influence alone is not enough to induce honest reporting. Since players may still have a small but positive incentive to misreport their signals

We need to introduce some scheme to punish dishonest reporting. To that end, define

$$p^{\phi}(y|a) = \sum_{s \in S} \phi(y|s) p(s|a)$$

and

$$p^{\phi}(y|a, s_i) = \sum_{s_{-i} \in S_{-i}} \phi(y|s_{-i}, s_i) p(s|a, s_i)$$

Given a public coordinating device $\phi : S \rightarrow \Delta(Y)$ and $\left(s_i, s_i', a\right) \in S_i \times S_i \times A$

$$\Lambda_i^{\phi}\left(s_i, s_i', a\right) = \min_{s_i' \neq s_i} ||p^{\phi}\left(\cdot|a, s_i\right) - p^{\phi}\left(\cdot|a, s_i'\right)||^2$$

This measures the extent to which player i's conditional beliefs regarding the public coordinating signal are different given s_i and s_i' (assuming honest reporting by others).

We use this variation of player i's beliefs to induce her to report her private signals truthfully.

Given a public coordinating device $\phi : S \rightarrow \Delta(Y)$, the measure p^{ϕ} is $\gamma-$regular for ϕ if

$$v_i^{\phi}\left(s_i, s_i', a\right) \leq \gamma \Lambda_i^{\phi}\left(s_i, s_i', a\right)$$

for all u, $s_i \in S_i$, $s_i' \in S_i$, $a \in A$ and $i \in N$.

Lemma: If (G', p) is a private monitoring game and if $\lambda \in (0, 1)$, then there exists a $\gamma > 0$ such that the following holds: if p^{ϕ} is $\gamma-$regular for some ϕ, then for any $a \in A$ and any payoff function $w : Y \rightarrow \mathbb{R}^n$, there exists a public coordination device $\phi'_{a,w} : S \rightarrow \triangle(Y)$ such that truthful reporting is a Bayesian Nash equilibrium for the one-shot information revelation game with public coordinating device $\phi^{\lambda}_{a,w} = (1 - \lambda)\phi + \lambda\phi'_{a,w}$.

Combining these ideas, we obtain the following result:

Theorem: Fix any private monitoring game (G', p). Suppose that $intV^*(G') \neq \varnothing$ and there exists $\phi : S \rightarrow \Delta(Y)$ such that p^{ϕ} is distinguishable. Then there exists a $\gamma > 0$ such that, if p^{ϕ} is $\gamma-$regular, then the following holds: for every convex,compact, smooth set $W \subseteq intV^*(G')$, there exists an $\eta > 0$ and a $\underline{\delta} \in (0, 1)$ such that, for each $\delta \in (\underline{\delta}, 1)$ and for each $v \in W$, there exists a public communication device Φ and a $(1-\delta)\,\eta-$uniformly strict truthful PPE of $G_p'^{\infty}(\delta, \Phi)$ with payoff v.

Remark: Distiguishability and $\gamma-$regularity are properties of p^{ϕ} determined by ϕ. When such a ϕ exists, it "works" for all convex,compact, smooth set $W \subseteq intV^*(G')$.

To state the theorem precisely, we need to define the distinguishability conditions. Given p and ϕ, let

$$T_i^{p^\phi}(a) = co\left\{p^\phi\left(\cdot|a_i', a_{-i}\right) - p^\phi(\cdot|a) : a_i' \neq a_i\right\}$$

and

$$\widehat{T}_i^{p^\phi}(a) = co\left\{T_i^{p^\phi}(a) \cup \{0\}\right\}$$

We say that p^ϕ satisfies *distinguishability* at $a \in A$ if for each pair of distinct players i and j, the following conditions are satisfied:

$$0 \notin T_i^{p^\phi}(a) \cup T_j^{p^\phi}(a)$$

$$\widehat{T}_i^{p^\phi}(a) \cap \widehat{T}_j^{p^\phi}(a) = \{0\}$$

$$\left(-\widehat{T}_i^{p^\phi}(a)\right) \cap \widehat{T}_j^{p^\phi}(a) = \{0\}$$

The first condition implies that a deviation by i from a_i to a_i' is statistically detectable. The second and third conditions implies that a deviation by player i from a_i to a_i' and a deviation by player j from a_j to a_j' are statistically distinguishable.

When are these conditions satisfied ?

Suppose that $S_i = Y$ for each i and that

$$p(s|a) = \sum_{y \in Y} \prod_i q_i(s_i|y) \pi(y|a)$$

where $q_i(y|y) \geq \beta$ for any y and i. Let ϕ_M be the "majority rule", which is a public coordination device that chooses y reported by the largest number of players (with some tie-breaking rule). Then for every γ, there exists a β such that

$$p^{\phi_M}(y|a) = \sum_{s \in S} \phi_M(y|s) p(s|a) = \sum_{s \in S} \phi_M(y|s) \sum_{y \in Y} \left[\prod_{i=1}^{n} q_i(s_i|y) \right] \pi(y|a)$$

is γ-regular. and distiguishability is satisfied.

How does the proof work? We adapt the ideas in APS and FLM.

We extend the notions of enforceability and decomposability to our framework.

An action profile $a \in A$ is $\eta-$*enforceable* with respect to $W \subset \mathbb{R}^n$ and $\delta \in (0,1)$ if there exists a public coordinating device $\phi : S \rightarrow \Delta(Y)$ and $w : Y \rightarrow W$ such that for all $i \in N$

$$(1-\delta)\, g_i\,(a) + \delta \sum_{s \in S} \left[\sum_{y \in Y} w_i\,(y)\, \phi\,(y|s_i, s_{-i}) \right] p(s|a) - \eta \geq$$

$$(1-\delta)\, g_i \left(a_i', a_{-i}\right) + \delta \sum_{s} \left[\sum_{y} w_i \left(h^t, y\right) \phi_{h^t}\,(y|f_i(s_i), s_{-i}) \right] p(s|a_i, \alpha_{-i}^t \left(h^t\right))$$

for all $a_i' \neq a_i$, $f_i : S_i \rightarrow S_i$ and for each s,

$$\sum_{s} \left[\sum_{y} w_i\,(y)\, \phi\,(y|s_i, s_{-i}) \right] p(s|a) \geq \sum_{s} \left[\sum_{y} w_i\,(y)\, \phi \left(y|s_i', s_{-i}\right) \right] p(s|a)$$

If $a \in A$ is $\eta-$enforceable with respect to W and δ with some ϕ and w and if $v = (1 - \delta)\, g(a) + \delta E^{\phi}[w_i(\cdot)\,|a]$, then we say that the triple (a, ϕ, w) $\eta-$*enforces* v with respect to W and δ.

We say that v is $\eta-$*decomposable* with respect to W and δ when there exists a triple (a, ϕ, w) that $\eta-$enforces v with respect to W and δ.

Next define the set of $\eta-$decomposable payoffs with respect to W and δ as follows.

$$B(\delta, W, \eta) := \{v \in \mathbb{R}^n | v \text{ is } \eta - \text{decomposable with respect to } W \text{ and } \delta\}.$$

We say that W is $\eta-$*self decomposable* with respect to $\delta \in (0, 1)$ if $W \subset B(\delta, W, \eta)$.

A "uniformly strict" version of Theorem 1 in Abreu, Pearce, and Stacchetti holds here when $\eta > 0$: if W is $\eta-$self decomposable with respect to δ, then every $v \in W$ can be supported by a $\eta-$uniformly strict PPE of $G_p'^{\infty}(\delta, \Phi)$ for some public communication device Φ. Note that each payoff profile may need to be supported by using a different public coordinating device. Hence different public coordinating devices need to be used at different public histories. More formally, we have:

Lemma: If $W \subset \mathbb{R}^n$ is bounded and $\eta-$self decomposable with respect to $\delta \in (0, 1)$, then for any $v \in W$, there exists Φ such that $v \in E(\delta, \Phi, \eta)$.