

2018 Delhi Winter School

Repeated Games: Private Monitoring

George J. Mailath

University of Pennsylvania
and
Australian National University

December 2018



Games with Private Monitoring

- Intertemporal incentives arise when public histories coordinate continuation play.
- Can intertemporal incentives be provided when the monitoring is **private**?
- Stigler (1964) suggested that that answer is often NO, and so collusion is not likely to be a problem when monitoring problems are severe.



The Problem

- Fix a strategy profile σ . Player i 's strategy is **sequentially rational** if, after all **private** histories, the continuation strategy is a best reply to the other players' continuation strategies (which depend on their private histories).
- That is, player i is best responding to the other players' behavior, given his beliefs over the private histories of the other players.
- While player i knows his/her beliefs, the modeler typically does not.
- Most researchers thought this problem was intractable,



The Problem

- Fix a strategy profile σ . Player i 's strategy is **sequentially rational** if, after all **private** histories, the continuation strategy is a best reply to the other players' continuation strategies (which depend on their private histories).
- That is, player i is best responding to the other players' behavior, given his beliefs over the private histories of the other players.
- While player i knows his/her beliefs, the modeler typically does not.
- Most researchers thought this problem was intractable, until Sekiguchi, in 1997, showed:


There exists an almost efficient eq for the PD with conditionally-independent almost-perfect private monitoring.



Prisoners' Dilemma

Conditionally Independent Private Monitoring

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



- Rather than observing the other player's action for sure, player i observes a noisy signal: $\pi_i(y_i = a_j) = 1 - \varepsilon$.
- Grim trigger is **not** an equilibrium: at the end of the first period, it is not optimal for player i to play S after observing $y_i = s_j$ (since in eq, player j played E and so with high prob, observed $y_j = e_i$).
- Sekiguchi (1997) avoided this by having players randomize (we will see how later).

Almost Public Monitoring

- How robust are PPE in the game with public monitoring to the introduction of a **little** private monitoring?
- Perturb the public signal, so that player i observes the conditionally (on y) independent signal $y_i \in \{\underline{y}, \bar{y}\}$, with probabilities given by

$$\pi(y_1, y_2 \mid y) = \pi_1(y_1 \mid y)\pi_2(y_2 \mid y),$$

and


$$\pi_i(y_i \mid y) = \begin{cases} 1 - \varepsilon, & \text{if } y_i = y, \\ \varepsilon, & \text{if } y_i \neq y. \end{cases}$$

- Ex post payoffs are now $u_i^*(a_i, y_i)$.



Prisoners' Dilemma with Noisy Monitoring


Bounded Recall-public monitoring



- Suppose $(3p - 2q - r)^{-1} < \delta < (p + 2q - 3r)^{-1}$, so profile is strict PPE in game with public monitoring.
- $V_i(w)$ is i 's value from being in **public** state w .

Prisoners' Dilemma with Noisy Monitoring


Bounded Recall-private (almost-public) monitoring



- In period t , player i 's continuation strategy after **private** history $h_i^t = (a_i^0, a_i^1, \dots, a_i^{t-1})$ is completely determined by i 's **private** state $w_i^t \in \mathcal{W}$.
- In period t , j sees **private** history h_j^t , and forms belief $\beta_j(h_j^t) \in \mathcal{W}$ over the period t state of player i .

Prisoners' Dilemma with Noisy Monitoring

Bounded Recall-Best Replies



- For all y , $\Pr(y_i \neq y_j \mid y) = 2\varepsilon(1 - \varepsilon)$, and so


$$\Pr(w_j^t \neq w_i^t(h_i^t) \mid h_i^{t'}) = 2\varepsilon(1 - \varepsilon) \quad \forall t' \leq t.$$

- For ε sufficiently small, incentives from public monitoring carry over to game with almost public monitoring, and profile is an equilibrium.

Prisoners' Dilemma with Noisy Monitoring

Grim Trigger

- Suppose $\frac{1}{2p-q} < \delta < 1$, so grim trigger is a strict PPE.
- Strategy in game with private monitoring is



- If $1 > p > q > r > 0$, profile is not a Nash eq (for any $\varepsilon > 0$).
- If $1 > p > r > q > 0$, profile is a Nash eq (but not sequentially rational).

Prisoners' Dilemma with Noisy Monitoring

Grim Trigger, $1 > p > q > r > 0$

- Consider private history $h_1^t = (E\underline{y}_1, S\bar{y}_1, S\bar{y}_1, \dots, S\bar{y}_1)$.
- Associated beliefs of 1 about w_2^t :

$$\Pr(w_2^0 = w_E) = 1,$$

$$\Pr(w_2^1 = w_S \mid E\underline{y}_1) = \Pr(y_2^1 = \underline{y}_2 \mid E\underline{y}_1, w_2^0 = w_E) \approx 1 - \varepsilon < 1,$$

$$\text{but } \Pr(w_2^t = w_S \mid h_1^t)$$

$$\begin{aligned} &= \underbrace{\Pr(w_2^t = w_S \mid w_2^{t-1} = w_S)}_{=1} \Pr(w_2^{t-1} = w_S \mid h_1^t) \\ &\quad + \underbrace{\Pr(y_2^t = \underline{y} \mid w_2^{t-1} = w_E, h_1^t)}_{\approx 0} \Pr(w_2^{t-1} = w_E \mid h_1^t), \end{aligned}$$

and $\Pr(w_2^{t-1} = w_S \mid h_1^t) < \Pr(w_2^{t-1} = w_S \mid h_1^{t-1})$, and so
 $\Pr(w_2^t = w_S \mid h_1^t) \rightarrow \approx 0$, as $t \rightarrow \infty$. **Not** Nash.



Prisoners' Dilemma with Noisy Monitoring

Grim Trigger, $1 > p > r > q > 0$

- Consider private history $h_1^t = (E\underline{y}_1, S\bar{y}_1, S\bar{y}_1, \dots, S\bar{y}_1)$.
- Associated beliefs of 1 about w_2^t :

$$\Pr(w_2^0 = w_E) = 1,$$

$$\Pr(w_2^1 = w_S \mid E\underline{y}_1) = \Pr(y_2^1 = \underline{y}_2 \mid E\underline{y}_1, w_2^0 = w_E) \approx 1 - \varepsilon < 1,$$

$$\begin{aligned} \text{but } \Pr(w_2^t = w_S \mid h_1^t) &= \underbrace{\Pr(w_2^t = w_S \mid w_2^{t-1} = w_S)}_{=1} \Pr(w_2^{t-1} = w_S \mid h_1^t) \\ &\quad + \underbrace{\Pr(y_2^t = \underline{y} \mid w_2^{t-1} = w_E, h_1^t)}_{\approx 0} \Pr(w_2^{t-1} = w_E \mid h_1^t), \end{aligned}$$

and $\Pr(w_2^{t-1} = w_S \mid h_1^t) > \Pr(w_2^{t-1} = w_S \mid h_1^{t-1})$, and so
 $\Pr(w_2^t = w_S \mid h_1^t) \approx 1$ for all t . **Nash.**



Prisoners' Dilemma with Noisy Monitoring

Grim Trigger, $1 > p > r > q > 0$

- Consider private history $h_1^t = (E\underline{y}_1, E\bar{y}_1, E\bar{y}_1, \dots, E\bar{y}_1)$.
- Associated beliefs of 1 about w_2^t :

$$\Pr(w_2^0 = w_E) = 1,$$

$$\Pr(w_2^1 = w_S \mid E\underline{y}_1) = \Pr(y_2^1 = \underline{y}_2 \mid E\underline{y}_1, w_2^0 = w_E) \approx 1 - \varepsilon < 1,$$

$$\text{but } \Pr(w_2^t = w_S \mid h_1^t)$$

$$\begin{aligned} &= \underbrace{\Pr(w_2^t = w_S \mid w_2^{t-1} = w_S)}_{=1} \Pr(w_2^{t-1} = w_S \mid h_1^t) \\ &\quad + \underbrace{\Pr(y_2^t = \underline{y} \mid w_2^{t-1} = w_E, h_1^t)}_{\approx 0} \Pr(w_2^{t-1} = w_E \mid h_1^t), \end{aligned}$$

and $\Pr(w_2^{t-1} = w_S \mid h_1^t) < \Pr(w_2^{t-1} = w_S \mid h_1^{t-1})$, and so
 $\Pr(w_2^t = w_S \mid h_1^t) \rightarrow \approx 0$, as $t \rightarrow \infty$. **Failure** seq rationality.



Automaton Representation of Strategies

An **automaton** is the tuple $(\mathcal{W}_i, w_i^0, f_i, \tau_i)$, where

- \mathcal{W}_i is set of states,
- w_i^0 is initial state,
- $f_i : \mathcal{W} \rightarrow A_i$ is output function (decision rule), and
- $\tau_i : \mathcal{W}_i \times A_i \times Y_i \rightarrow \mathcal{W}_i$ is transition function.

Any automaton $(\mathcal{W}_i, w_i^0, f_i, \tau_i)$ induces a strategy for i . Define

$$\tau_i(w_i, h_i^t) := \tau_i(\tau_i(w_i, h_i^{t-1}), a_i^{t-1}, y_i^{t-1}).$$

The induced strategy s_i is given by $s_i(\emptyset) = f_i(w_i^0)$ and

$$s_i(h_i^t) = f_i(\tau_i(w_i^0, h_i^t)), \quad \forall h_i^t.$$

Every strategy can be represented by an automaton.



Almost Public Monitoring Games

- Fix a game with imperfect **full support** public monitoring, so that for all $y \in Y$ and $a \in A$, $\rho(y \mid a) > 0$.
- Rather than observing the public signal directly, each player i observes a private signal $y_i \in Y$.
- The game with private monitoring is **ε -close** to the game with public monitoring if the joint distribution π on the private signal profile (y_1, \dots, y_n) satisfies

$$|\pi((y, y, \dots, y) \mid a) - \rho(y \mid a)| < \varepsilon.$$

Such a game has **almost public monitoring**.

- Any automaton in the game with public monitoring describes a strategy profile in all **ε -close** almost public monitoring games.



Almost Public Monitoring

Rich Private Monitoring

- Fix a game with imperfect **full support** public monitoring, so that for all $y \in Y$ and $a \in A$, $\rho(y | a) > 0$.
- Each player i observes a private signal $z_i \in Z_i$, with (z_1, \dots, z_n) distributed according to the joint ds π .
- The game with **rich private monitoring** is **ε -close** to the game with public monitoring if there are mappings $\xi_i : Z_i \rightarrow Y$ such that

$$\left| \sum_{\xi_1(z_1)=y, \dots, \xi_n(z_n)=y} \pi((z_1, \dots, z_n) | a) - \rho(y | a) \right| < \varepsilon.$$

Such a game has **almost public monitoring**.

- Any automaton in the game with public monitoring describes a strategy profile in all **ε -close** almost public monitoring games with rich private monitoring.



Behavioral Robustness I

Definition

An eq of a game with public monitoring is **behaviorally robust** if the **same** automaton is an eq in all ε -close games to the game with public monitoring for ε sufficiently small.



Behavioral Robustness I

Definition

An eq of a game with public monitoring is **behaviorally robust** if the **same** automaton is an eq in all ε -close games to the game with public monitoring for ε sufficiently small.

Theorem

Suppose the public profile $(\mathcal{W}, w^0, f, \tau)$ is a strict equilibrium of the game with public monitoring for some δ and $|\mathcal{W}| < \infty$. For all $\kappa > 0$, there exists η and ε such that if the posterior beliefs induced by the private profile satisfy $\beta_i(\tau(h_i^t) \mathbf{1} | h_i^t) > 1 - \eta$ for all h_i^t , and if π is ε -close to ρ , then the private profile is a sequential equilibrium of the game with private monitoring for the same δ , and the expected payoff in that equilibrium is within κ of the public equilibrium payoff.

Behavioral Robustness II

Definition

A public automaton $(\mathcal{W}, w^0, f, \tau)$ has **bounded recall** if there exists L such that after any history of length at least L , continuation play only depends on the last L periods of the public history (i.e., $\tau(w, h^L) = \tau(w', h^L)$ for all $w, w' \in \mathcal{W}$ reachable in the same period).

Theorem

Given a finite memory public profile, for all $\eta > 0$, there exists $\varepsilon > 0$ such that if π is ε -close to ρ , the posterior beliefs induced by the private profile satisfy $\beta_i(\tau(h_i^t) \mathbf{1} | h_i^t) > 1 - \eta$ for all h_i^t .



Behavioral Robustness III

An eq is **behaviorally robust** if the **same** profile is an eq in near-by games.
A public profile has **bounded recall** if there exists L such that after any history of length at least L , continuation play only depends on the last L periods of the public history.

Theorem (Mailath and Morris, 2002)

A strict PPE with bounded recall is behaviorally robust to private monitoring that is almost public.



Behavioral Robustness III

An eq is **behaviorally robust** if the **same** profile is an eq in near-by games.
A public profile has **bounded recall** if there exists L such that after any history of length at least L , continuation play only depends on the last L periods of the public history.

Theorem (Mailath and Morris, 2002)

A strict PPE with bounded recall is behaviorally robust to private monitoring that is almost public.

“Theorem” (Mailath and Morris, 2006)

If the private monitoring is sufficiently rich, a strict PPE is behaviorally robust to private monitoring that is almost public if and only if it has bounded recall.



Illustration of Mailath and Morris 2006

Grim trigger in PD

- Suppose $\frac{1}{2p-q} < \delta < 1$, so grim trigger is a strict PPE.

- Signal structure: $a_1 a_2$

	\underline{y}_2	\bar{y}_2'	\bar{y}_2''
\underline{y}_1	$(1 - \alpha)(1 - 3\varepsilon)$	ε	ε
\bar{y}_1	ε	$\alpha'(1 - 3\varepsilon)$	$(\alpha - \alpha')(1 - 3\varepsilon)$

$$\alpha = \begin{cases} p, & a = EE, \\ q, & a = SE, ES, \\ r, & a = SS, \end{cases} \quad \alpha' = \begin{cases} p', & a = EE, \\ q', & a = SE, ES, \\ r', & a = SS. \end{cases}$$

- If $1 > p > q > r > 0$, profile is not a Nash eq of private monitoring game.
- Suppose $1 > p > r > q > 0$, and $\alpha' = \alpha/2$, profile is a Nash eq of private monitoring game.



Illustration of Mailath and Morris 2006

Grim trigger in PD

- Suppose $\frac{1}{2p-q} < \delta < 1$, so grim trigger is a strict PPE.

- Signal structure:

$a_1 a_2$	\underline{y}_2	\bar{y}'_2	\bar{y}''_2
\underline{y}_1	$(1 - \alpha)(1 - 3\varepsilon)$	ε	ε
\bar{y}_1	ε	$\alpha'(1 - 3\varepsilon)$	$(\alpha - \alpha')(1 - 3\varepsilon)$

$$\alpha = \begin{cases} p, & a = EE, \\ q, & a = SE, ES, \\ r, & a = SS, \end{cases} \quad \alpha' = \begin{cases} p', & a = EE, \\ q', & a = SE, ES, \\ r', & a = SS. \end{cases}$$

- If $1 > p > q > r > 0$, profile is not a Nash eq of private monitoring game.
- Suppose $1 > p > r > q > 0$, but $1 > p' > q' > r' > 0$, profile is not a Nash eq.

Bounded Recall

It is tempting to think that bounded recall provides an attractive restriction on behavior. But:

Folk Theorem II (Hörner and Olszewski, 2009)

The public monitoring folk theorem holds using bounded recall strategies. The folk theorem also holds using bounded recall strategies for games with almost-public monitoring.

- This private monitoring folk theorem is **not** behaviorally robust.

Bounded Recall

It is tempting to think that bounded recall provides an attractive restriction on behavior. But:

Folk Theorem II (Hörner and Olszewski, 2009)

The public monitoring folk theorem holds using bounded recall strategies. The folk theorem also holds using bounded recall strategies for games with almost-public monitoring.

- This private monitoring folk theorem is **not** behaviorally robust.


Folk Theorem III (Mailath and Olszewski, 2011)

The perfect monitoring folk theorem holds using bounded recall strategies with uniformly strict incentives. Moreover, the resulting equilibrium is behaviorally robust to almost-perfect almost-public monitoring.

Prisoners' Dilemma

Conditionally Independent Private Monitoring

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



Player i observes a noisy signal: $\pi_i(y_i = a_j) = 1 - \varepsilon$.

Theorem (Sekiguchi, 1997)

For all $\psi > 0$, there exists $\eta'' > \eta' > 0$ such that for all $\delta \in (1/3 + \eta', 1/3 + \eta'')$, there is a Nash equilibrium in which each player randomizing over the initial state, with the probability on w_E exceeding $1 - \psi$.

Proof and extend to all high δ

Proof of theorem

Optimality of grim trigger after different histories:

- E_S : updating given original randomization $\implies S$ optimal.

Proof and extend to all high δ

Proof of theorem

Optimality of grim trigger after different histories:

- Es : updating given original randomization $\implies S$ optimal.
- Ee, Ee, \dots, Ee : perpetual e reassures i that j is still in w_E .

Proof and extend to all high δ

Proof of theorem

Optimality of grim trigger after different histories:

- Es : updating given original randomization $\implies S$ optimal.
- Ee, Ee, \dots, Ee : perpetual e reassures i that j is still in w_E .
- Ee, Ee, \dots, Ee, Es . Most likely events: either j is still in w_E and s is a mistake, or j received an erroneous signal in the previous period. Odds slightly favor j receiving the erroneous signal, and because δ low, S is optimal.

Proof and extend to all high δ

Proof of theorem

Optimality of grim trigger after different histories:

- Es : updating given original randomization $\implies S$ optimal.
- Ee, Ee, \dots, Ee : perpetual e reassures i that j is still in w_E .
- Ee, Ee, \dots, Ee, Es . Most likely events: either j is still in w_E and s is a mistake, or j received an erroneous signal in the previous period. Odds slightly favor j receiving the erroneous signal, and because δ low, S is optimal.
- $Ee, Ee, \dots, Ee, Es, Se, \dots, Se$. This period's S will trigger j 's switch to w_S , if not there already.

Proof and extend to all high δ

Proof of theorem

Optimality of grim trigger after different histories:

- Es : updating given original randomization $\implies S$ optimal.
- Ee, Ee, \dots, Ee : perpetual e reassures i that j is still in w_E .
- Ee, Ee, \dots, Ee, Es . Most likely events: either j is still in w_E and s is a mistake, or j received an erroneous signal in the previous period. Odds slightly favor j receiving the erroneous signal, and because δ low, S is optimal.
- $Ee, Ee, \dots, Ee, Es, Se, \dots, Se$. This period's S will trigger j 's switch to w_S , if not there already.

To extend to all high δ , lower effective discount factor by dividing games into N interleaved games.

Belief-Free Equilibria

Another approach is to specify behavior in such a way that the beliefs are irrelevant.

Suppose $n = 2$.

Definition

The profile $((\mathcal{W}_1, w_1^0, f_1, \tau_1), (\mathcal{W}_2, w_2^0, f_2, \tau_2))$ is a **belief-free eq** if for all $(w_1, w_2) \in \mathcal{W}_1 \times \mathcal{W}_1$, $(\mathcal{W}_i, w_i, f_i, \tau_i)$ is a best reply to $(\mathcal{W}_j, w_j, f_j, \tau_j)$, all $i \neq j$.

This approach is due to Piccione (2002), with a refinement by Ely and Valimaki (2002). Belief-free eq are characterized by Ely, Hörner, and Olszewski (2005).



Illustration of Belief Free Eq

The product-choice game

	<i>c</i>	<i>s</i>
<i>H</i>	2, 3	0, 2
<i>L</i>	3, 0	1, 1


- Row player is a firm choosing *High* or *Low* quality.
- Column player is a **short-lived** customer choosing the customized or standard product.
- In the game with perfect monitoring, grim trigger (play *Hc* till 1 plays *L*, then revert to perpetual *Ls*) is an eq if $\delta \geq \frac{1}{2}$.

The **belief-free eq** that achieves a payoff of 2 for the row player:

- Row player always plays $\frac{1}{2} \circ H + \frac{1}{2} \circ L$. (Trivial automaton)
- Column player's strategy has one period memory. Play c for sure after H in the previous period, and play

$$\alpha^L := \left(1 - \frac{1}{2\delta}\right) \circ c + \frac{1}{2\delta} \circ s$$

after L in the previous period. Player 2's automaton:



- Let $V_1(w; a_1)$ denote player 1's payoff when 2 is in state w , and 1 plays a_1 . Then (where $\alpha = 1 - 1/(2\delta)$),

$$\begin{aligned} V_1(w_c; H) &= (1 - \delta)2 + \delta V_1(w_c) \\ &= V_1(w_c; L) = (1 - \delta)3 + \delta V_1(w_{\alpha L}), \\ V_1(w_{\alpha L}; a_1 = H) &= (1 - \delta)2\alpha + \delta V_1(w_c) \\ &= V_1(w_{\alpha L}; a_1 = L) = (1 - \delta)(2\alpha + 1) + \delta V_1(w_{\alpha L}). \end{aligned}$$

- Then, $V_1(w_c) - V_1(w_{\alpha L}) = (1 - \delta)/\delta$.
- Which is true when $\alpha = 1 - 1/(2\delta)$.

Belief-Free Eq in the Prisoners' Dilemma

Ely and Valimaki (2002)

- Perfect monitoring PD.

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

- Player i 's automaton, $(\mathcal{W}_i, w_i, f_i, \tau_i)$:

$$\mathcal{W} = \{w_i^E, w_i^S\},$$

$$f_i(w_i^a) = \begin{cases} 1, & a = E, \\ \alpha \circ E + (1 - \alpha) \circ S, & a = S, \alpha := 1 - 1/(3\delta), \end{cases}$$

$$\tau_i(w_i, a_i a_j) = w_i^{a_j}.$$

- Both $(\mathcal{W}_1, w_1^E, f_1, \tau_1)$ and $(\mathcal{W}_1, w_1^S, f_1, \tau_1)$ are best replies to both $(\mathcal{W}_2, w_2^E, f_2, \tau_2)$ and $(\mathcal{W}_2, w_2^S, f_2, \tau_2)$.

Belief-Free in the Prisoners' Dilemma-Proof

- Let $V_1(aa')$ denote player 1's payoff when 1 is in state w_1^a and 2 is in state $w_2^{a'}$. Then

$$V_1(E E) = (1 - \delta)2 + \delta V_1(E E),$$

$$V_1(E S) = (1 - \delta)(3\alpha - 1) \\ + \delta[\alpha V_1(E E) + (1 - \alpha)V_1(S E)],$$

$$V_1(S E; a_1 = E) = (1 - \delta)2 + \delta V_1(E E) \\ = V_1(S E; a_1 = S) = (1 - \delta)3 + \delta V_1(E S),$$

$$V_1(S S : a_1 = E) = (1 - \delta)(-1) \\ + \delta[\alpha V_1(E E) + (1 - \alpha)V_1(S E)] \\ = V_1(S S : a_1 = S) = \delta[\alpha V_1(E S) + (1 - \alpha)V_1(S S)].$$

- Then, $V_1(E E) - V_1(E S) = V_1(S E) - V_1(S S) = (1 - \delta)/\delta$.
- Which is true when $\alpha = 1 - 1/(3\delta)$.



Belief-Free in the Prisoners' Dilemma

Private Monitoring

- Suppose we have conditionally independent private monitoring.
- For ε small, there is a value of α satisfying the analogue of the indifference conditions for perfect monitoring (the system of equations is well-behaved, and so you can apply the implicit function theorem).
- These kinds of strategies can be used to construct equilibria with payoffs in the square $(0, 2) \times (0, 2)$ for sufficiently patient players.



- Histories are **not** being used to coordinate play! There is no common understanding of continuation play.
- This is to be contrasted with strict PPE.
- Rather, lump sum taxes are being imposed after “deviant” behavior is “suggested.”
- This is essentially what we do in the repeated prisoners’ dilemma.
- Folk theorems for games with private monitoring have been proved using belief free constructions.
- These equilibria seem crazy, yet Kandori and Obayashi (2014) report suggestive evidence that in some community unions in Japan, the behavior accords with such an equilibrium.



Imperfect Monitoring

- This works for public and private monitoring.
- No hope for behavioral robustness.

“Theorem” (Hörner and Olszewski, 2006)

The folk theorem holds for games with private almost-perfect monitoring.

- Result uses belief-free ideas in a central way, but the equilibria constructed are not belief free.

Role of Mixing

	H_2	T_2
H_1	1, -1	-1, 1
T_1	-1, 1	1, -2

Mixed strategy eq:

$$(\frac{3}{5} \circ H_1 + \frac{2}{5} \circ T_1, \frac{1}{2} \circ H_2 + \frac{1}{2} \circ T_2)$$

- Standard complaint about mixing:
 - People don't randomize
 - Randomization probabilities for i determined by need to keep j indifferent.

Role of Mixing

	H_2	T_2
H_1	1, -1	-1, 1
T_1	-1, 1	1, -2

Mixed strategy eq:

$$(\frac{3}{5} \circ H_1 + \frac{2}{5} \circ T_1, \frac{1}{2} \circ H_2 + \frac{1}{2} \circ T_2)$$

- Standard complaint about mixing:
 - People don't randomize
 - Randomization probabilities for i determined by need to keep j indifferent.
- Response:
 - players have private information—players don't randomize, but opponents face nontrivial distribution of behavior.

Suppose player i has payoff irrelevant private information $t_i \sim \mathcal{U}([0, 1])$.

$$\sigma_1(t_1) = \begin{cases} H_1, & t_1 \leq \frac{3}{5}, \\ T_1, & t_1 > \frac{3}{5}, \end{cases} \quad \sigma_2(t_2) = \begin{cases} H_2, & t_2 \leq \frac{1}{2}, \\ T_2, & t_2 > \frac{1}{2}. \end{cases}$$

Harsanyi (1973) Purification

Getting the right probabilities

- Make the private information payoff information:

	H_2	T_2
H_1	$1 - \varepsilon t_1, -1 - \varepsilon t_2$	$-1, 1$
T_1	$-1, 1$	$1, -2$

$$t_i \sim \mathcal{U}([0, 1]).$$

$$\sigma_i(t_i) = \begin{cases} H_i, & t_i \leq \bar{t}_i(\varepsilon), \\ T_i, & t_i > \bar{t}_i(\varepsilon), \end{cases}$$

- As $\varepsilon \rightarrow 0$, $\bar{t}_1(\varepsilon) \rightarrow \frac{3}{5}$ and $\bar{t}_2(\varepsilon) \rightarrow \frac{1}{2}$.

Harsanyi (1973) Purification

Getting the right probabilities

- Make the private information payoff information:

	H_2	T_2
H_1	$1 - \varepsilon t_1, -1 - \varepsilon t_2$	$-1, 1$
T_1	$-1, 1$	$1, -2$

$$t_i \sim \mathcal{U}([0, 1]).$$

$$\sigma_i(t_i) = \begin{cases} H_i, & t_i \leq \bar{t}_i(\varepsilon), \\ T_i, & t_i > \bar{t}_i(\varepsilon), \end{cases}$$

- As $\varepsilon \rightarrow 0$, $\bar{t}_1(\varepsilon) \rightarrow \frac{3}{5}$ and $\bar{t}_2(\varepsilon) \rightarrow \frac{1}{2}$.
- Belief-free equilibria typically have the property that players randomize the same way after different histories (and so with different beliefs over the private states of the other player(s)).

Purification of Belief Free Eq

- Belief-free equilibria typically have the property that players randomize the same way after different histories (and so with different beliefs over the private states of the other player(s)).
- Can we purify belief-free equilibria (Bhaskar, Mailath, and Morris, 2008) by introducing iid payoff shocks over time?
 - The one period memory belief free equilibria of Ely and Valimaki (2002), as exemplified above, is not purifiable using one period memory strategies.
 - They are purifiable using unbounded memory strategies.
 - Open question: can they be purified using bounded memory strategies? (It turns out that for sequential games, only Markov equilibria can be purified using bounded memory strategies, Bhaskar, Mailath, and Morris 2013).



What about noisy monitoring?

Current best result is Sugaya (2013):

“Theorem”

The folk theorem generically holds for the repeated two-player prisoners' dilemma with private monitoring if the support of each player's signal distribution is sufficiently large. Neither cheap talk communication nor public randomization is necessary, and the monitoring can be very noisy.

Ex Post Equilibria

- The belief-free idea is very powerful.
- Suppose there is an unknown state determining payoffs and monitoring.

ω_E	E	S
E	1, 1	-1, 2
S	2, -1	0, 0

ω_S	E	S
E	0, 0	2, -1
S	-1, 2	1, 1

- Let $\Gamma(\delta; \omega)$ denote the complete-information repeated game when state ω is common knowledge. Monitoring may be perfect or imperfect public.

Perfect Public Ex Post Equilibria

$\Gamma(\delta; \omega)$ is complete-information repeated game at ω .

Definition

The public strategy profile σ^* is a **perfect public ex post eq** if $\sigma^*|_{h^t}$ is a Nash eq of $\Gamma(\delta; \omega)$ for all $h^t \in H$, where $\sigma^*|_{h^t}$ is continuation public profile induced by h^t .

- These equilibria can be strict; histories **do** coordinate play.
- But the eq are **belief free**.



Perfect Public Ex Post Equilibria

$\Gamma(\delta; \omega)$ is complete-information repeated game at ω .

Definition

The public strategy profile σ^* is a **perfect public ex post eq** if $\sigma^*|_{h^t}$ is a Nash eq of $\Gamma(\delta; \omega)$ for all $h^t \in H$, where $\sigma^*|_{h^t}$ is continuation public profile induced by h^t .

- These equilibria can be strict; histories **do** coordinate play.
- But the eq are **belief free**.

“Theorem” (Fudenberg and Yamamoto 2010)

Suppose the signals are statistically informative (about actions and states).
The folk theorem holds state-by-state.

These ideas also can be used in some classes of reputation games (Hörner and Lovo, 2009) and in games with private monitoring (Yamamoto, 2014).



Conclusion

- The current theory of repeated games shows that the long relationships can discourage opportunistic behavior, it does not show that long run relationships will discourage opportunistic behavior.
- Incentives can be provided when histories coordinate continuation play.
- Punishments must be credible, and this can limit their scope.
- Some form of monitoring is needed to punish deviators.
- This monitoring can occur through communication networks.
- Intertemporal incentives can also be provided in situations when there is no common understanding of histories, and so of continuation play.



What is left to understand

- Which behaviors in long-run relationships are plausible?
- Why are formal institutions important?
- Why do we need formal institutions to protect property rights, for example?
- Communication is not often modelled explicitly, and it should be. Communication make things significantly easier (see Compte, 1998, and Kandori and Matsushima, 1998).
- Too much focus on patient players (δ close to 1).

