Private Information and the Problem of Coordinating Punishments

repeated game equilibria have a self-referential nature: players don't do things because they are afraid they will be punished, and they punish because they are afraid if they do not they will be punished for that and so forth

- this requires players to know when they are being punished
- this is difficult with signals that are not common knowledge
- is my price low because you deviated or because you got a signal that you should punish me? In the former case I should punish you, in the latter case if I do I trigger off a war that unravels the equilibrium

Stage Game

two players $i = 1,2$

chooses an *actions* a_i from a finite set A_i

payoff to an action profile $g_i(a)$

$\overline{g} = \max_{i,a} |g_i(a)|$

each player observes a private signal z_i in a finite set Z_i

action profiles induce a probability distribution π_a over *outcomes* z

end of each stage of the game, players make *announcements* $y_i \in Y*$, where $Y*$ is a finite set that is the same for each player

stage game strategy $s_i = (a_i, m_i)$: choice of action a_i and map $m_i : Z_i \rightarrow Y*$ from private signal to announcements

Remark on the Two Player Case

more players is easier: can compare announcement by different players

Repeated Game

each period $t = 1, 2, \ldots$, stage game is played

public randomization device each period uniform $w \in [0,1]$

public history at time t, $h(t)$: announcements and realization of w signals in all previous periods, and also the realization of w in period t, so

$$h(t) = (w(1), y(1), w(2), y(2), \ldots, w(t-1), y(t-1), w(t)).$$

private history for player i at time t is

$$h_i(t) = (a_i(1), z_i(1), a_i(2), z_i(2), \ldots, a_i(t-1), z_i(t-1)).$$

Strategies

strategy for player i is a sequence of maps $\sigma_i(t)$ mapping the public and private histories $h(t), h_i(t)$ to probability distributions over S_i

partial strategy is the strategy conditional on the initial realization of the public randomization device

public strategy is a strategy that depends only on $h(t)$

null private history for player i is $h_i(1)$

initial public history is $h(1)$

for each public history $h(t)$ the public strategy profile σ induces partial strategy profile over the repeated game beginning at t; denote by $[\sigma \mid h(t)]$

Preferences

discount factor by δ, use average present value

given strategy profile σ expected average present value of payoffs generated by partial strategy profiles $[\sigma \mid w(1)]$ denoted by $G_i(\sigma,\delta)$

perfect public equilibrium a public strategy profile σ such that for any public history $h(t)$ and any private partial strategy $\tilde{\sigma}_i$ by any player i we have

$$G_i([\sigma \mid h(t)],\delta) \geq G_i((\tilde{\sigma}_i,[\sigma_{-i} \mid h(t)]),\delta).$$

by standard dynamic programming arguments sufficient to consider deviations to public strategies.

Structure of Information

convenient to think of players “agreeing” if they make same announcement as each other

think of Y^* as being the subset of Y in which $y_1 = y_2$: called *diagonal*

given message profile m the information structure π induces a distribution over the diagonal of announcement profiles

probability of diagonal point $\pi_a^m(y^*) = \sum_{z|m_1(z_1)=m_2(z_2)=y^*} \pi_a(z)$,

probability of joint announcement conditional on diagonal

$$\pi_a^m(y^* \mid Y^*) = \frac{\pi_a^m(y^*)}{\sum_{y \in Y^*} \pi_a^m(y)}$$

probability opponent’s message given positive probability signal

$$\pi_a^m(y_{-i} \mid z_i) = \sum_{z_{-i}|m_{-i}(z_{-i})=y_{-i}} \pi_a(z_{-i} \mid z_i)$$

Almost Public Messaging

Definition 1: A game has (ε,ν) public information with respect to m if for all action profiles a,

(1) $\overline{\pi}_a^m \equiv \sum_{y^* \in Y^*} \pi_a^m(y^*) \geq 1 - \varepsilon$

(2) if $\pi_a(z) > 0$ then for all $y_{-i} \neq m_i(z_i)$,
$\pi_a^m(y_{-i} \mid z_i) \leq \pi_a^m(m_i(z_i) \mid z_i) - \nu$

most of the time, each player fairly confident of the other player's message

limit case of (0,ν)-public information two players' messages are perfectly correlated, so public information

Versus "Close to Public Monitoring"

similar to the Mailath and Morris "ε-close to public monitoring" but weaker in two ways

1. Mailath and Morris suppose each players private signal z_i lie in same set as the signals in the limiting pubic-information game meaning $\#Y^* = \#Z_i$

2. they suppose that in public information limit, every signal has strictly positive probability under every action profile, and that the distribution of each player's private signals is close to this limit

these imply condition (1) a stronger version of condition (2): $\lim_{\varepsilon \to 0} \pi_a^m(m_i(z_i) \mid z_i) = 1$

given $\#Y^* = \#Z_i$ conditions equivalent

many private signals per public message (2) weaker: allows private signals to differ in how informative they are about the message the opposing player will send

Further Discussion

easier to satisfy with coarse message maps

vacuously satisfied if $m_1 \text{ and } m_2$ are equal to the same constant

condition will have force when combined with assumption that messages “reveal enough” about the action profile that generated the underlying signals.

except in the trivial case of perfect information (2) rules out z_1, z_2 independent conditional on a

requires if one player receives a signal unlikely conditional on a, it is likely that the other player receive the corresponding unlikely signal

Information Matrix

consider $\pi_a^m(\bullet|\, Y^*)$ as row vector

construct a matrix $\Pi_a^{m,i}$ by stacking row vectors corresponding to $(\tilde{a}_i, a_{-i})$ as $\tilde{a}_i$ ranges over A_i

stack two matrices corresponding to the two players to get a $(\# A_1 + \# A_2) \times \# Y^*$ matrix Π_a^m

this matrix has two rows (both corresponding to a) that are identical.

Definition 2: A game has *pure-strategy pairwise full rank* with respect to m if for every pure profile a the rank of Π_a^m is $(\# A_1 + \# A_2) - 1$.

never satisfied in games such as Green and Porter where players have the same sets of feasible actions, and the distribution of signals satisfies symmetry condition that $\pi_{(\alpha,\alpha')} = \pi_{(\alpha',\alpha)}$

is satisfied for set of probability measures π_a of full Lebesgue measure.

Nash Threats Folk Theorem

v^* be static Nash payoff vector normalized so $v^* = 0$

consider a sequence of games indexed by n

Corollary: Fix a message profile m, and suppose that $g^n \to g$, $\pi^n \to \pi$, that game n has (ε^n, ν) public information with respect to m, that $\varepsilon^n \to 0$, that $\pi_a^m(\cdot \mid Y^*)$ has pure-strategy pairwise full rank with respect to m, and that each g^n has a static equilibrium with payoffs converging to 0. Then there is a sequence $\gamma^n \to 0$ such that for any feasible interior vector of payoffs $v > 0$ there exists $\delta^* < 1$ and an N such that for any $n > N$ and all $\delta \geq \delta^*$ there is a perfect public equilibrium in the game n with payoffs v^n satisfying $\|v^n - v\| < \gamma^n$.

Idea of Proof

Find an auxiliary game where there is no disagreement

Prove a uniform version of the folk theorem in that public information game: using the arguments from Fudenberg, Levine, and Maskin [8] and McLean, Obara and Postlewaite [15]

Map back to the original game and punish players for disagreeing

Not so likely to disagree on equilibrium path

Use of Public Information

announcements are public information

why not use the regular folk theorem for that case?

FLM folk theorem limited to the convex hull of the set of profiles that satisfy enforceability plus pairwise identifiability

fix profile, including a strategy for sending messages

a player can randomize announcements independent of private information while preserving the marginal distribution of messages: “faking the marginal”

pairwise identifiability fails, because player one faking his marginal and player two faking hers are observationally equivalent

Information Aggregation

make same announcement for several different private signals.

two effects:

1. increases degree to which each player can forecast the other player's message, reducing role of private information

2. reduces the informativeness of the messages, making it less likely that the assumption of pairwise full-rank is satisfied

Notions of Equilibrium

ε*-sequential*

every player following every of his private histories and public history has consistent beliefs such that conditional on his information he loses no more than ε in average present value measured at that time by deviating

uniform equilibrium with respect to time averaging

1. time average converges on equilibrium path

2. for any $\rho > 0, \tau$ there exists $T > \tau$ such that any deviation loses at least ρ in finite time T average

Approximate Equilibrium and Time Averaging

Theorem: Suppose $T^n, \varepsilon^n > 0, \sigma^n$ such that σ^n is T^n finite horizon ε^n-approximate Nash equilibrium with payoff v^n, and that $\varepsilon^n \to 0$, $v^n \to v$. Then:

A. There exist $\delta^n \to 1$, $\varepsilon^n \to 0$, σ^n such that σ^n is ε^n-sequential for δ^n and the equilibrium average present values converge to v

B. There exists a uniform equilibrium with payoff v

mutual threat point a payoff vector v such that there exists a *mutual punishment action*: mixed action profile α such that $g_i(\alpha_i', \alpha_{-i}) \leq v$

consider *enforceable mutually punishable set* V^*: intersection of closure of the convex hull of the payoff vectors that weakly pareto dominate a mutual punishment point and the closure of the convex hull of the enforceable payoffs

difference with standard folk theorem: can exclude unenforceable actions and the minmax point may not be mutually punishable

Informational Connectedness

relevant only with more than two players

player i is *directly connected* to player $j \neq i$ despite player $k \neq i, j$ if exists mixed profile α and mixed action $\hat{\alpha}_i$ such that

$$\pi_j(\cdot \mid \hat{\alpha}_i, \alpha_k', \alpha_{-i-k}) \neq \rho_j(\alpha) \text{ for all } \alpha_k'.$$

i is *connected* to j if for every $k \neq i, j$ there is a sequence of players $i_1, \ldots, i_n$ with $i_1 = i, i_n = j$ and $i_p \neq k$ for any p such that player i_p is directly connected to player i_{p+1} despite player k

game is *informationally connected* if there are only two players, or if every player is connected to every other player

Theorem 8.1: In an informationally connected game if $v \in V^*$ then there exists a sequence of discount factors $\delta_n \to 1$, non-negative numbers $\varepsilon_n \to 0$ and strategy profiles σ_n such that σ_n is an ε_n-sequential equilibrium for δ_n and equilibrium payoffs converge to v.

- Use communication and punishment phases that are a small fraction of the total time
- Aggregate information over a long time before deciding what to do
- Need the epsilon so you if you've generated really good signals you don't cheat as you approach the assessment phases

Belief Free Equilibrium and Friends

construct equilibria with the property that my best play does not depend on what I believe about your history

- gets around the coordination problem
- a possibly small subset of all equilibria
- but big enough that you can prove some folk theorems this way without communication