Lecture 10: Repeated Incomplete Information Games without Discounting

Harry PEI
Department of Economics, Northwestern University

Spring Quarter, 2021

Lectures 1-9: Models with Commitment Types

All the models we have seen so far:

- There exists at least one type whose strategy is exogenously assumed.

Fudenberg and Levine (89,92), Schmidt (93), Evans and Thomas (95), Cripps, Dekel and Pesendorfer (05), Atakan and Ekmekci (12), Pei (20).

- Compute the rational type's payoff from imitating a commitment type.
- This leads to a lower bound on the rational type's equilibrium payoff.

This week: Models *without* Commitment Types

All types' behaviors are endogenous.

- It is hard to compute type a's payoff from imitating type b since type b's behavior is also unknown.
- Deliver sharp predictions on players' payoffs seems unrealistic except for specific classes of games (e.g., zero-sum games).

Results in this literature: Characterize the set of equilibrium payoffs.

Two classes of models:

- Models *without* discounting.
- Models *with* discounting.

Start from models w/o discounting (Shalev 1994) and then explain the connections to models with discounting (Cripps and Thomas 2003).

Repeated Incomplete Info Game without Discounting

- Time $t = 1, 2, \ldots$
- Two patient players: 1 and 2. Actions $a_1 \in A_1$ and $a_2 \in A_2$.
- P1 has private info about a persistent state $\theta \in \Theta$.

  P2's prior belief $\pi \in \Delta(\Theta)$.
- Stage-game payoffs $u_1(\theta, a_1, a_2)$ and $u_2(\theta, a_1, a_2)$.
- Both players can perfectly observe all the past actions.

  Public history $h^t \equiv \{a_{1,s}, a_{2,s}\}_{s \leq t-1} \in \mathscr{H}^t$, with $\mathscr{H} \equiv \cup_{t \in \mathbb{N}} \mathscr{H}^t$.
- Player 1's strategy: $\sigma_1 : \mathscr{H} \times \Theta \to \Delta(A_1)$.

  Player 2's strategy: $\sigma_2 : \mathscr{H} \to \Delta(A_2)$.
- **Assumption:** A_1, A_2, and Θ are finite, $|A_1| \geq 2$, $|A_2| \geq 2$, and π has full support.

Repeated Game Payoffs & Nash Equilibrium

Player i maximizes the limit of $\frac{1}{T}\sum_{t=1}^{T} u_i(\theta, a_{1,t}, a_{2,t})$ as $T \to +\infty$.

Notion of Nash equilibrium in undiscounted games.

Definition: Nash Equilibrium

(σ_1, σ_2) *is a Nash equilibrium if for every* σ_2',

$$\liminf_{T\to+\infty} \mathbb{E}^{\pi}_{(\sigma_1,\sigma_2)}\left[\frac{1}{T}\sum_{t=1}^{T} u_2(\theta, a_{1,t}, a_{2,t})\right] \geq \limsup_{T\to+\infty} \mathbb{E}^{\pi}_{(\sigma_1,\sigma_2')}\left[\frac{1}{T}\sum_{t=1}^{T} u_2(\theta, a_{1,t}, a_{2,t})\right]$$

and for every σ_1' *and* $\theta \in \Theta$,

$$\liminf_{T\to+\infty} \mathbb{E}^{\theta}_{(\sigma_1,\sigma_2)}\left[\frac{1}{T}\sum_{t=1}^{T} u_1(\theta, a_{1,t}, a_{2,t})\right] \geq \limsup_{T\to+\infty} \mathbb{E}^{\theta}_{(\sigma_1',\sigma_2)}\left[\frac{1}{T}\sum_{t=1}^{T} u_1(\theta, a_{1,t}, a_{2,t})\right]$$

Uniform Nash equilibrium in Hart (1985): Equivalent in terms of payoffs in games with one-sided private info.

Repeated Game Payoffs & Nash Equilibrium

Definition: Nash Equilibrium

(σ_1, σ_2) is a Nash equilibrium if for every σ_2',

$$\liminf_{T\to+\infty} \mathbb{E}^{\pi}_{(\sigma_1,\sigma_2)}\left[\frac{1}{T}\sum_{t=1}^{T} u_2(\theta, a_{1,t}, a_{2,t})\right] \geq \limsup_{T\to+\infty} \mathbb{E}^{\pi}_{(\sigma_1,\sigma_2')}\left[\frac{1}{T}\sum_{t=1}^{T} u_2(\theta, a_{1,t}, a_{2,t})\right]$$

and for every σ_1' and $\theta \in \Theta$,

$$\liminf_{T\to+\infty} \mathbb{E}^{\theta}_{(\sigma_1,\sigma_2)}\left[\frac{1}{T}\sum_{t=1}^{T} u_1(\theta, a_{1,t}, a_{2,t})\right] \geq \limsup_{T\to+\infty} \mathbb{E}^{\theta}_{(\sigma_1',\sigma_2)}\left[\frac{1}{T}\sum_{t=1}^{T} u_1(\theta, a_{1,t}, a_{2,t})\right]$$

If (σ_1, σ_2) is a Nash equilibrium, then

$$\lim_{T\to+\infty} \mathbb{E}^{\theta}_{(\sigma_1,\sigma_2)}\left[\frac{1}{T}\sum_{t=1}^{T} u_1(\theta, a_{1,t}, a_{2,t})\right] \text{ and } \lim_{T\to+\infty} \mathbb{E}^{\pi}_{(\sigma_1,\sigma_2)}\left[\frac{1}{T}\sum_{t=1}^{T} u_2(\theta, a_{1,t}, a_{2,t})\right]$$

exist and are well-defined.

Equilibrium Payoff & Equilibrium Payoff Set

An equilibrium payoff is a $|\Theta|+1$ dimensional real vector $\left((u_\theta)_{\theta\in\Theta}, v\right)$.

Objective: Given the primitives (u_1, u_2, π), characterize the equilibrium payoff set $\mathscr{U} \subset \mathbb{R}^{|\Theta|+1}$, s.t.

1. For every $\left((u_\theta)_{\theta\in\Theta}, v\right) \in \mathscr{U}$, there exists a sequence of equilibria $\{(\sigma_1^n, \sigma_2^n)\}_{n\in\mathbb{N}}$, s.t. type θ's payoff converges to u_θ and P2's payoff converges to v as $n \to +\infty$.
2. For every $\left((u_\theta)_{\theta\in\Theta}, v\right) \notin \mathscr{U}$, players' payoff in every equilibrium is bounded away from $\left((u_\theta)_{\theta\in\Theta}, v\right)$.

Overview of the Literature

Aumann and Maschler (1966): Characterize $\mathscr{U}$ for zero-sum games.

Hart (1985): Characterize $\mathscr{U}$ for every game with one-sided private info.

- Upside: Very general, can handle games with interdependent values.
- Downside: Based on bi-martingales. Hard to compute.

Shalev (1994): Characterize $\mathscr{U}$ for every game with one-sided private info and private values.

- **Assumption**: The value of $u_2(\theta, a_1, a_2)$ does not depend on θ.
- Upside: The characterization is based on linear programming, which is easy to understand and to compute.

Cripps and Thomas (2003): Establish connections between $\mathscr{U}$ and the equilibrium payoff set in games with equal discounting as $\delta \to 1$.

Today: Private Value Games

We make a private value assumption on players' payoffs:

$$u_2(\theta, a_1, a_2) = u_2(a_1, a_2).$$

The French and the Israelis call this case *known own payoffs* since P2 knows his cardinal preferences over (a_1, a_2) despite not knowing θ.

Minmax Payoff

First step: What are players' minmax payoffs?

- Suppose P1 deviates, to what extent can P2 punish P1?
- Challenge: P1 has private info about his payoff.

  Can P2 come up with a strategy that can punish *all* types of P1?

Formally, suppose P1's equilibrium payoff is $\boldsymbol{u} \equiv (u_\theta)_{\theta\in\Theta}$.

- What are the conditions on $\boldsymbol{u}$ s.t. there exists a strategy for P2 σ_2^* under which P1's payoff is no more than $\boldsymbol{u}$ regardless of P1's strategy σ_1?

Blackwell's Approachability Theorem

For every $p \in \Delta(\Theta)$, consider a one-shot complete info zero-sum game with P1's payoff being $\widetilde{u}_{1,p}(a_1,a_2) \equiv \sum_{\theta\in\Theta} p(\theta) u_1(\theta,a_1,a_2)$. Let

$$\underline{u}(p) \equiv \min_{\alpha_2 \in \Delta(A_2)} \max_{\alpha_1 \in \Delta(A_1)} \widetilde{u}_{1,p}(a_1,a_2).$$

Theorem: Blackwell Approachability (Blackwell 1956)

There exists a strategy for P2 σ_2^* under which

$$\lim_{T\to+\infty} \sup \mathbb{E}^{\theta}_{(\sigma_1,\sigma_2^*)}\Big[\frac{1}{T}\sum_{t=1}^{T} u_1(\theta,a_{1,t},a_{2,t})\Big] \leq u_\theta \textit{ for every } \theta \in \Theta \textit{ and } \sigma_1 \in \Sigma_1$$

if and only if $\boldsymbol{u} \equiv (u_\theta)_{\theta\in\Theta}$ *is such that*

$$\sum_{\theta\in\Theta} p(\theta) u_\theta \geq \underline{u}(p) \quad \textit{for every} \quad p \in \Delta(\Theta). \tag{1}$$

If P1's payoff $\boldsymbol{u}$ satisfies (1), then there exists a strategy for P2 under which no type of P1 has an incentive to make any observable deviation.

Player 2's Minmax Payoff

Consider a one-shot complete info zero-sum game with P2's payoff being $u_2(a_1, a_2)$. Let

$$\underline{v} \equiv \min_{\alpha_1 \in \Delta(A_1)} \max_{\alpha_2 \in \Delta(A_2)} u_2(a_1, a_2).$$

By definition, there exists a strategy for P1 under which P2's payoff is no more than $\underline{v}$ regardless of P2's strategy.

If P2's payoff is at least $\underline{v}$, then there exists a strategy for P1 under which P2 has no incentive to make any observable deviation.

Shalev: Games with Private Values

Theorem (Proposition 4 in Shalev 1994)

$((u_\theta)_{\theta\in\Theta}, v)$ *is an equilibrium payoff if and only if there exists* $\{\alpha^\theta\}_{\theta\in\Theta}$ *with* $\alpha^\theta \in \Delta(A_1 \times A_2)$ *such that:*

1. ***Feasibility:*** $u_\theta = u_1(\theta, \alpha^\theta)$ *for every* $\theta \in \Theta$. $v = \sum_{\theta\in\Theta} \pi(\theta) u_2(\alpha^\theta)$.
2. ***IR-1:*** $\sum_{\theta\in\Theta} p(\theta) u_\theta \geq \underline{u}(p)$ *for every* $p \in \Delta(\Theta)$.

   IR-2: $u_2(\alpha^\theta) \geq \underline{v}$ *for every* $\theta \in \Theta$.
3. ***IC:*** *For every* $\theta, \theta' \in \Theta$, *type* θ *weakly prefers* α^θ *to* $\alpha^{\theta'}$.

Intuition: Think about a mechanism design problem s.t. P1 reports his type before the game starts and receives an allocation α^θ if he reports θ.

Proof Sketch (Necessity of feasibility constraints)

- Given equilibrium (σ_1, σ_2), P1's and P2's payoffs conditional on θ are:

$$u_\theta = \lim_{T \to +\infty} \mathbb{E}^{\theta}_{(\sigma_1,\sigma_2)}\left[\frac{1}{T}\sum_{t=1}^{T} u_1(\theta, a_{1,t}, a_{2,t})\right]$$

$$v_\theta = \lim_{T \to +\infty} \mathbb{E}^{\theta}_{(\sigma_1,\sigma_2)}\left[\frac{1}{T}\sum_{t=1}^{T} u_2(\theta, a_{1,t}, a_{2,t})\right].$$

- Hence, there exists $\alpha^\theta \in \Delta(A_1 \times A_2)$ such that

$$u_\theta = \sum_{(a_1,a_2) \in A_1 \times A_2} \alpha^\theta(a_1,a_2) u_1(\theta, a_1, a_2)$$

$$v_\theta = \sum_{(a_1,a_2) \in A_1 \times A_2} \alpha^\theta(a_1,a_2) u_2(a_1, a_2).$$

- Since prior is $\pi \in \Delta(\Theta)$, P2's equilibrium payoff is $v = \sum_{\theta \in \Theta} \pi(\theta) v_\theta$.

Proof Sketch (Necessity of IR constraints)

(IR-1): $\sum_{\theta\in\Theta} p(\theta)u_\theta \geq \underline{u}(p)$ for every $p \in \Delta(\Theta)$.

Suppose $(u_\theta)_{\theta\in\Theta}$ is P1's equilibrium payoff but violates (IR-1), then

Theorem: Blackwell Approachability (Blackwell 1956)

There exists a strategy for P2 under which P1's time-average payoff is no more than $\boldsymbol{u} \equiv (u_\theta)_{\theta\in\Theta}$ regardless of P1's strategy if and only if

$$\sum_{\theta\in\Theta} p(\theta)u_\theta \geq \underline{u}(p) \quad \textit{for every} \quad p \in \Delta(\Theta).$$

no matter what P2's strategy is, there exists $\theta \in \Theta$ s.t. type θ's payoff is strictly more than u_θ.

- P1 has a profitable deviation, which leads to a contradiction.

Proof Sketch (Necessity of IR constraints)

(IR-2): $\sum_{(a_1,a_2)\in A_1\times A_2} \alpha^\theta(a_1,a_2)u_2(a_1,a_2) \geq \underline{v}$ for every $\theta \in \Theta$.

Recall that by our construction of α^θ, P2's time-average payoff conditional on θ equals his stage-game payoff from α^θ.

Suppose P2's payoff conditional on θ, v_θ, is less than $\underline{v}$.

- Since $\{\pi_t\}_{t\in\mathbb{N}}$ is a bounded martingale, it converges a.s.
- There exists a limiting belief π_∞ s.t. $\theta \in \text{supp}(\pi_\infty)$ and at π_∞, P2's payoff conditional on θ is less than $\underline{v}$.
- Since P2's payoff is no less than $\underline{v}$, there exists $\theta' \in \text{supp}(\pi_\infty)$ s.t. at π_∞, P2's payoff conditional on θ' is more than $\underline{v}$.
- Since P2's payoff does not depend on θ, the continuation play at π_∞ is different under θ and θ', i.e., types θ and θ' behave differently.
- This contradicts the presumption that π_∞ is a limiting belief.

Proof Sketch (Necessity of IC constraint)

(IC): For every $\theta, \theta' \in \Theta$, type θ weakly prefers α^{θ} to $\alpha^{\theta'}$.

If (IC) is violated, then some type of player 1 strictly prefers the equilibrium strategy of another type rather than his own equilibrium strategy.

Proof Sketch (Sufficiency)

- Let n be the smallest integer s.t. $2^n \geq |\Theta|$, and $\{a_1', a_1''\} \subset A_1$.
- From period 1 to n, each type of P1 plays a different sequence in $\{a_1', a_1''\}^n$. P2's actions are arbitrary.

  In equilibrium, P2 learns P1's type in period $n+1$.
- If P1's type is θ, then players' action frequencies in the continuation game is given by α^θ.
- If P1 makes any off-path deviation, then P2 punishes him using the strategy that drives his payoff below $\boldsymbol{u}$.
- If P2 makes any off-path deviation, then P1 punishes him using the strategy that drives his payoff below $\underline{v}$.

Blackwell Approachability Theorem

Theorem: Blackwell Approachability (Blackwell 1956)

There exists a strategy for P2 σ_2^* under which

$$\lim \sup_{T\to+\infty} \mathbb{E}^{\theta}_{(\sigma_1,\sigma_2^*)}\left[\frac{1}{T}\sum_{t=1}^{T} u_1(\theta, a_{1,t}, a_{2,t})\right] \le u_\theta \textit{ for every } \theta\in\Theta \textit{ and } \sigma_1\in\Sigma_1$$

if and only if $\boldsymbol{u}$ *is such that*

$$\sum_{\theta\in\Theta} p(\theta)u_\theta \ge \underline{u}(p) \quad \textit{for every} \quad p\in\Delta(\Theta).$$

You can view this as a generalization of the minmax theorem to multi-dimensional objectives.

- P2 tries to make the average outcome belonging to $\{\boldsymbol{w}\in\mathbb{R}^{|\Theta|}|\boldsymbol{w}\le\boldsymbol{u}\}$
- P1 tries to make the average outcome escaping this set.

When can P2 successfully do that? (conditions on $\boldsymbol{u}$)

Proof Sketch: Blackwell Approachability

If $\boldsymbol{u}$ is such that

$$\sum_{\theta\in\Theta} p(\theta)u_\theta \geq \underline{u}(p) \quad \text{for every} \quad p\in\Delta(\Theta),$$

let's construct a strategy for P2 s.t. P1's average payoff belongs to $S \equiv \{\boldsymbol{w}\in\mathbb{R}^{|\Theta|} | \boldsymbol{w}\leq \boldsymbol{u}\}$.

Let $\boldsymbol{u}_t \in \mathbb{R}^{|\Theta|}$ be P1's stage-game payoff in period t. Let

$$\overline{\boldsymbol{u}}_t \equiv \frac{1}{t}\sum_{s=1}^{t} \boldsymbol{u}_s.$$

- In period 0 or when $\overline{\boldsymbol{u}}_t \in S$, P2's action is arbitrary.
- If $\overline{\boldsymbol{u}}_t \notin S$, then there exists $\boldsymbol{x}_t \in S$ that minimizes the distance to $\overline{\boldsymbol{u}}_t$ among all elements in S.

  Let $\boldsymbol{p}\in\Delta(\Theta)$ be such that $\boldsymbol{p}\cdot(\overline{\boldsymbol{u}}_t - \boldsymbol{x}_t) = 0$. P2 plays his minmax strategy in the zero-sum game $\sum_{\theta\in\Theta} p(\theta)u_1(\theta,a_1,a_2)$.

Blackwell Approachability Theorem

P2's Blackwell strategy:

- In period 0 or when $\bar{\boldsymbol{u}}_t \in S$, P2's action is arbitrary.
- If $\bar{\boldsymbol{u}}_t \notin S$, then there exists $\boldsymbol{x}_t \in S$ that minimizes the distance to $\bar{\boldsymbol{u}}_t$ among all elements in S.

  Let $\boldsymbol{p} \in \Delta(\Theta)$ be such that $\boldsymbol{p} \cdot (\bar{\boldsymbol{u}}_t - \boldsymbol{x}_t) = 0$. P2 plays his minmax action in the zero-sum game $\sum_{\theta \in \Theta} p(\theta) u_1(\theta, a_1, a_2)$.

Lemma

If P2 uses this strategy and

$$\sum_{\theta \in \Theta} p(\theta) u_\theta \geq \underline{u}(p) \quad \textit{for every} \quad p \in \Delta(\Theta),$$

then for every $\varepsilon > 0$, there exists $T \in \mathbb{N}$ s.t. regardless of P1's strategy,

$$\Pr\Big(d(\bar{\boldsymbol{u}}_t, S) \geq \varepsilon \textit{ for some } t \geq T\Big) \leq \varepsilon.$$

Blackwell Approachability Theorem

Next Lecture

Cripps and Thomas (2003):

- Discounted games with two patient players.
- Shalev's characterization approximately applies.

Pei (2021 TE), Pei (2021 Working Paper).

- Discounted repeated games with short-lived uninformed players.