G5212: Game Theory

Mark Dean

Spring 2017

Bargaining

- We will now apply the concept of SPNE to bargaining
- A bit of background
    - Bargaining is hugely interesting but complicated to model
    - It turns out that the outcome depends a lot on the details of the game
        - i.e. the bargaining protocol you assume
    - Has led some to thing that non-cooperative game theory is not the way to go
        - See Nash 1950

Bargaining

- However, ‘Sequential Bargaining’ is still a classic framework
  - One side makes an offer
  - The other side can either accept or reject
  - If they reject then either the game ends or they get to make a counter offer
- We will analyze a finite (two stage) bargaining game
- Then an infinite game
  - Rubinstein bargaining

Two Stage Bargaining

Example

Two Stage Bargaining

- There are two players bargaining over a pie of size 1.
- There are two rounds:
  - Player 1 first propose a way to split the pie
  - Player 2 either accepts or rejects it. If an offer is accepted, the pie is split accordingly. Otherwise, the pie shrinks by 10% and the game proceed to the next round.
  - Player 2 proposes in the second round, and player 1 either accepts it or rejects it. If the proposal is rejected, the pie disappears. Otherwise, the pie is split accordingly.

Infinite Period Bargaining

Example

Rubinstein Bargaining. Now suppose there are $2n$ **periods**:
*In $t = 2n$, P2 proposes $(0, 1)$;
*in $t = 2n-1$, P1 proposes $(1-\delta, \delta)$;
*in $t = 2n-2$, $((1-\delta)\delta, 1-(1-\delta)\delta) = \left(\delta - \delta^2, 1 - \delta + \delta^2\right)$;
*in $t = 2n-3$, $\left(1 - \delta\left(1 - \delta + \delta^2\right), \delta\left(1 - \delta + \delta^2\right)\right) =$
$\left(1 - \delta + \delta^2 - \delta^3, \delta - \delta^2 + \delta^3\right)$.
*Inductively, we find, in period 1,
$\left(1 - \delta + \delta^2 - \delta^3 + ... + \delta^{2n-2} - \delta^{2n-1}, \delta - \delta^2 + \delta^3..\right) =$
$\left(\sum_{k=0}^{2n-1} (-\delta)^k, 1 - \sum_{k=0}^{2n-1} (-\delta)^k\right) = \left(\frac{1-\delta^{2n}}{1+\delta}, \frac{\delta+\delta^{2n}}{1+\delta}\right)$
*If $n \to \infty$, the division goes to $\left(\frac{1}{1+\delta}, \frac{\delta}{1+\delta}\right)$

Infinite Period Bargaining

Infinite Horizon Bargaining (Rubinstein 1982 Ecma).
There are two players bargaining over a pie of size 1. Player 1 proposes at $t = 1, 3, ...$ and player 2 proposes at $t = 2, 4, ...$ They discounts future payoff with a discount factor $\delta \in (0, 1)$.
Theorem. There is a unique SPNE: in any period, proposer offers $\frac{\delta}{1+\delta}$ to the other and keeps $\frac{1}{1+\delta}$ for himself; responder accepts an offer iff it is at least $\frac{\delta}{1+\delta}$.

Infinite Period Bargaining

Proof of Uniqueness

Uniqueness: Let m be the smallest amount a proposer receives over all SPNE in all subgames. Let M be the highest amount.

Claim 1. $m \geq 1 - \delta M$

The reason is that the responder today will be the proposer tomorrow – this player will get at most M tomorrow. Hence, today's proposer, can ensure an acceptance by offering δM to the receiver. Hence, his payoff must be no smaller than $1 - \delta M$.

Claim 2. $M \leq 1 - \delta m$.

The responder today has a payoff of at least δm (because he gets at least m tomorrow as a proposer). So the responder will not accept anything less than δm today. If δm is accepted today, then the proposer receives a payoff of $1 - \delta m$; if it is rejected, the responder today (i.e., the proposer tomorrow) will not offer more than δM tomorrow (which is only $\delta^2 M$ discounted back today). Therefore, $M \leq \max\left\{1 - \delta m, \delta^2 M\right\}$.

Infinite Period Bargaining

Proof of Uniqueness (conti.)

Putting together the two claims, we have

$$\begin{aligned} m &\geq 1-\delta M \geq 1-\delta\left(1-\delta m\right) \Longrightarrow m \geq \frac{1}{1+\delta} \\ M &\leq 1-\delta m \leq 1-\delta\left(1-\delta M\right) \Longrightarrow M \leq \frac{1}{1+\delta} \end{aligned}$$

Therefore,

$$m = M = \frac{1}{1+\delta}.$$

Infinite Period Bargaining

Comments on Rubinstein Bargaining

- Immediate agreement.
- First mover advantage
- Discounting $1 - \delta$ is a "friction" – basically the pie shrinks over time.
- If the friction is larger, i.e., δ smaller, then the proposer gets a larger share.
- If the friction disappears, i.e., $\delta \to 1$, then shares $\to \frac{1}{2}$ (this is the Nash bargaining solution).

Repeated Games

- Think back to the prisoner's dilemma game
- Is it reasonable to think that people cannot sustain the 'good' outcome of keeping quiet?
- Perhaps the problem is that people play prisoner's dilemma more than once
- Can that solve the problem?

Finitely Repeated Games

Example

	C	D
C	$1, 1$	$-1, 2$
D	$2, -1$	$0, 0$

This game is repeatedly played for T periods,. All previous actions are observable (perfect monitoring). Each player values the sum of payoffs over all periods.

- Can the repetition help us avoid the Prisoner's Dilemma outcome?
- No! Use backward induction

Finitely Repeated Games

Theorem

If the stage game has a unique Nash Equilibrium, then the unique SPNE of the finitely repeated game is for both players to play the Nash Equilibrium in each round

- Note that this is not the same as saying that finite repetition has no power!

Finitely Repeated Games

Example

	L	C	R
T	1, 1	5, 0	0, 0
M	0, 5	4, 4	0, 0
B	0, 0	0, 0	3, 3

- Pure Strategy NE: (T,L), (BR)
- Suppose the game is played twice?
- Finite repetition can help
- Infinite repetition can help more!

Infinitely Repeated Game

- Stage game: $G = \{(A_i, u_i)\}$.
  - A_i is the action space for player i.
  - $a = (a_1, ..., a_n)$ is the action profile.
- The game is played repeatedly at $t = 0, 1, ...$ with discount factor $\delta \in (0, 1)$
- History up to date t: $h^t := \left(a^0, ..., a^{t-1}\right) \in A^t =: H^t$; $H^0 = \{\varnothing\}$.
- Set of possible histories $H := \cup_{t=0}^{\infty} H^t$.
- Pure strategy of player i: $s_i : H \to A_i$
- An outcome path induced by $s = (s_1, ..., s_n)$ is denoted by $a(s) = \left(a^0(s), a^1(s), ...\right)$
- Payoff

$$U_i(s) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_i\left(a^t(s)\right)$$

Infinitely Repeated Prisoner's Dilemma

	C	D
C	$1, 1$	$-1, 2$
D	$2, -1$	$0, 0$

- Consider the following strategies of players 1 and 2: play C no matter what:

$$s_i\left(h^t\right) = C \text{ for any } h^t \in H$$

- Player 1's payoff is

$$U_1\left(s_1, s_2\right) = (1-\delta)\sum_{t=0}^{\infty}\delta^t u_1\left(C, C\right) = (1-\delta)\sum_{t=0}^{\infty}\delta^t = 1.$$

Infinitely Repeated Prisoner's Dilemma

	C	D
C	$1, 1$	$-1, 2$
D	$2, -1$	$0, 0$

- Consider the following strategy of players 1 and 2, "grim trigger":
  - Play C if your opponent has never played D
  - Play D if your opponent has ever played D
- Player 1's payoff is

$$\begin{aligned} U_1(s_1, s_2) &= (1-\delta)\sum_{t=0}^{\infty} \delta^t u_1\left(a_1^t(s), a_2^t(s)\right) \\ &= (1-\delta)\sum_{t=0}^{\infty} \delta^t = 1. \end{aligned}$$

Infinitely Repeated Prisoner's Dilemma

- Is this an equilibrium?
- How can we check?
- Use the one shot deviation principle!

Definition

Player i has a profitable one-shot deviation from his strategy $s_i : H \to A_i$ for a fixed profile of his opponents' strategies s_{-i}, where $S_j : H \to A_j$, if there exists a history h^t in which player i can profitably deviate from $s_i\left(h^t\right)$ to some other action at h^t and reverts to play s_i afterwards.

One Shot Deviation Principle

- We showed that, in finite games we could test for SPNE by testing the one-shot deviation principle
- Can we do the same thing here?
- Not necessarily
- Consider the following (very silly) game
  - One player
  - Pick a or b in every period
  - If always pick a get 1
  - Otherwise get 0
- The strategy of playing b in every period is not SPNE, but survives one shot deviation

One Shot Deviation Principle

Theorem

For $\delta < 1$,a strategy profile is subgame perfect if there are no profitable one-shot deviations.

Infinitely Repeated Prisoner's Dilemma

	C	D
C	$1, 1$	$-1, 2$
D	$2, -1$	$0, 0$

- Is 'play C no matter what' immune to one shot deviations?

$$s_i\left(h^t\right) = C \text{ for any } h^t \in H$$

- Consider the one shot deviation of playing D in period 0. Call this strategy s_1'.
  - Player 1's payoff from (s_1', s_2) is

$$\begin{aligned} U_1\left(s_1', s_2\right) &= (1-\delta)\sum_{t=0}^{\infty} \delta^t u_1\left(a_1^t\left(s_1', s_2\right), a_2^t\left(s_1', s_2\right)\right) \\ &= (1-\delta)\left(2 + 1\cdot\delta + 1\cdot\delta^2 + ...\right) = 2-\delta > 1. \end{aligned}$$

Infinitely Repeated Prisoner's Dilemma

	C	D
C	$1, 1$	$-1, 2$
D	$2, -1$	$0, 0$

- Is Grim Trigger immune to one shot deviations?
- Consider the following one-shot deviation of player 1:
  - Play D at $t = 0$ and then follow s_1 afterwards. Call this s_1'
  - Player 1's payoff from (s_1', s_2) is

$$\begin{aligned} U_1(s_1', s_2) &= (1-\delta)\sum_{t=0}^{\infty}\delta^t u_1\left(a_1^t(s_1', s_2), a_2^t(s_1', s_2)\right) \\ &= (1-\delta)\left(2 + -\delta 1 + \delta^2 0 ...\right) \\ &= (1-\delta)(2-\delta). \end{aligned}$$

- Note that, after the 1 shot deviation will play C next period and will get punished, hence -1 in period 2

Infinitely Repeated Prisoner's Dilemma

	C	D
C	$1, 1$	$-1, 2$
D	$2, -1$	$0, 0$

- If $\delta > 0.382$ this is not profitable
- BUT, for SPNE have to check each subgame!
- Let's imagine that play in the first period was (C, D)
- According to the grim trigger,
  - Player 1 will play D in period 2
  - Player 2 will play C in period 2
- After which both players will play D forever

Infinitely Repeated Prisoner's Dilemma

	C	D
C	$1, 1$	$-1, 2$
D	$2, -1$	$0, 0$

- Consider only the payoff from period 2 onwards
- Payoff of Player 2 playing C in period 2 is

$$(1-\delta)(-1+0+0+....)$$

- Payoff of the 1 shot deviation to D in period 2 is

$$(1-\delta)(0+0+0+....)$$

Infinitely Repeated Prisoner's Dilemma

	C	D
C	$1, 1$	$-1, 2$
D	$2, -1$	$0, 0$

- Grim Trigger not SPNE as stated
- Need a slightly modified version
    - Play C if no one has ever played D
    - Play D if **anyone** has ever played D
- Can show that this is SPNE as long as $\delta > \frac{1}{2}$

Possible Outcomes in Infinitely Repeated Games

- We have shown that, in the infinitely repeated prisoner's dilemma we can support the outcome $(1, 1)$ as long as people are patient enough
- A natural question is: What is the range of possible outcomes that can be supported in a repeated game?

Example

	C	D
C	$5, 5$	$0, 6$
D	$6, 0$	$1, 1$

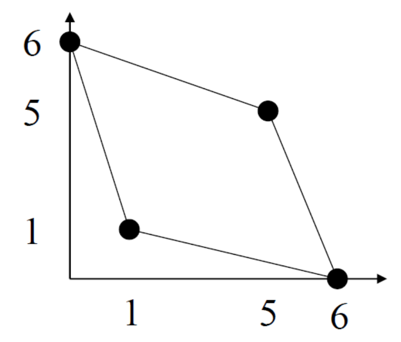

Can there be a SPNE with payoffs $(5.9, 0.1)$?

Possible Outcomes in Infinitely Repeated Games

- It seems that this is unlikely, as by playing D, player 2 can guarantee themselves a payoff of 1
- Why would they every accept 0.1?
- Pure strategy **minmax payoff** for player i:

$$\min_{a_{-i}} \max_{a_i} u_i \left(a_i, a_{-i}\right)$$

- We say a payoff for player i is individually rational if it is above their minmax payoff
- We would not expect to be able to support any payoffs that are not IR
- Are there any other restrictions?

Possible Outcomes in Infinitely Repeated Games

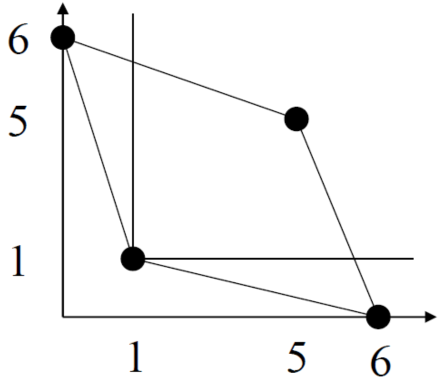

No! every payoff vector that is strictly individually rational for each player can be supported as SPNE if the discount factor is large enough.

Folk Theorem

- More formally:
- Feasible payoff:

$$con\left\{\left(u_1\left(a_1,a_2\right),u_2\left(a_1,a_2\right)\right):a_1\in A_1,a_2\in A_2\right\}.$$

- Folk theorem
  - Every action profile a that gives a payoff that is strictly individually rational can be played on the equilibrium path repeatedly on a SPNE if δ is large enough.
  - Every payoff profile that is feasible and strictly individually rational can be supported as a SPNE payoff if δ is large enough.

Folk Theorem

- Sketch of proof:

1. All players start by playing a and continue to play a if no deviation occurs.
2. If any one player, say player i, deviated, play the strategy profile m which minmaxes i for N periods.
3. If no players deviated from phase 2, all player $j \neq i$ gets rewarded ε above j's min-max forever after, while player i continues receiving his min-max
4. If player j deviated from phase 2, all players restart phase 2 with j as target.

Comment on Folk Theorem:

- Bad result: everything is feasible, and hence no predictive power at all.
- Good result: shows the power of dynamic incentives.