Government 2005: Formal Political Theory I

Lecture 10

Instructor: Tommaso Nannicini
Teaching Fellow: Jeremy Bowles

Harvard University

November 2, 2017

Overview

- Bayesian games
- Bayesian Nash equilibrium
- Simple (textbook) examples
- Cournot duopoly with asymmetric information
- Jury voting

Bayesian games

- We need to extend our definition of the normal-form representation of a game to account for **incomplete information**
- The normal-form representation of **Bayesian games** specifies
  - Players
  - Strategy spaces
  - Type spaces
  - Beliefs
  - Payoff functions
- **Bayesian Nash equilibrium** is like NE but accounts for incomplete information about players' payoffs

Meet the Nature

- Incomplete information raises the necessity to consider players' beliefs about other players' preferences, the second-order beliefs about these (first-order) beliefs, and so on
- We can sidestep this challenge following Harsanyi (1967):
  - Players' preferences are realizations of random variables
  - Nature moves first choosing preference types
  - Probability distributions of types are common knowledge
  - Players only observe subset of realizations (e.g., their own)
  - With this trick, from incomplete to imperfect information

Definition. A **Bayesian game** is made up of $\langle I, S_i, u_i(.), \Theta, F(.)\rangle$, where I are the players; $\theta_i \in \Theta_i$ is player i's type, with $\Theta = \Theta_1 \times ... \times \Theta_I$; player i's payoff function $u_i(s_i, s_{-i}, \theta_i)$ depends on her type; pure strategies are given by the functions $s_i(\theta_i) : \Theta_i \rightarrow S_i$; and $F(\theta_1, ..., \theta_I)$ is the joint probability distribution of players' types.

Bayesian Nash equilibrium (BNE)

- BNE is just the NE of a properly defined Bayesian game

Definition. A **Bayesian Nash equilibrium** for the above Bayesian game is a profile of decision rules $(s_1(.), ..., s_I(.))$ such that $\forall i \in I$:

$$E_\theta \left[u_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i) \right] \geq E_\theta \left[u_i(s_1'(\theta_1), s_{-i}(\theta_{-i}), \theta_i) \right]$$

for all alternative decision rules $s_i'(.) : \Theta_i \rightarrow S_i$ (in pure strategies).

Proposition. A profile of decision rules $(s_1(.), ..., s_I(.))$ is a BNE iff $\forall i \in I$ and $\forall \theta_i' \in \Theta_i$ occurring with positive probability:

$$E_{\theta-i} \left[u_i(s_i(\theta_i'), s_{-i}(\theta_{-i}), \theta_i' | \theta_i') \right] \geq E_{\theta-i} \left[u_i(s_1'(\theta_1'), s_{-i}(\theta_{-i}), \theta_i') | \theta_i') \right]$$

for all alternative decision rules $s_i'(.)$, and where expectation is taken over the other players' types conditional on i's realization of her type.

First (textbook) example: DA's brother

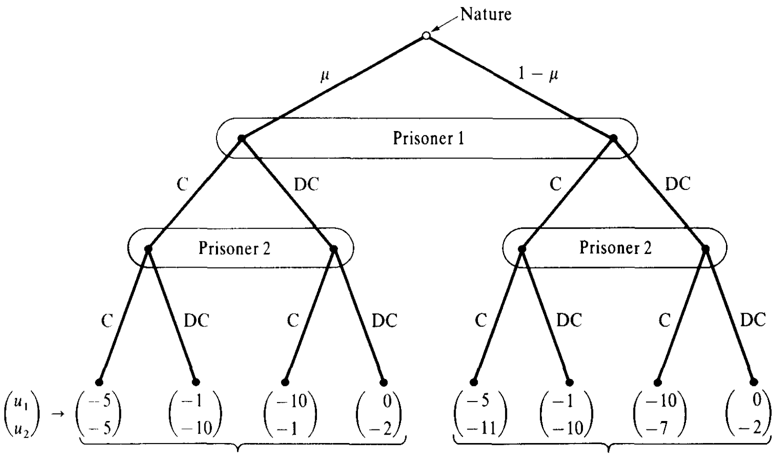

- Same PD but players 2 has two types: Second type, observed with probability $1-\mu$, has additional 6-year cost of confessing

First (textbook) example: DA's brother (contd.)

- Player 1 has no private info, and then 2 strategies: C, DC
- Player 2 has 4 strategies: (C if 1, C if 2), (C if 1, DC if 2), (DC if 1, C if 2), (DC if 1, DC if 2)
- C dominant for type-I player 2: $s_2(I) = C$
- DC dominant for type-II player 2: $s_2(II) = DC$
- Therefore:
  - $E_\theta[u_1(C)] = -5\mu - (1-\mu)$
  - $E_\theta[u_1(DC)] = -10\mu$
- $s_1 = C$ iff $\mu > 1/6$
- Unique BNE depending on parametric distribution μ

Second (textbook) example: Information may hurt

- Again, Player 2 has two types: $P(\omega_1) = 1/2$, $P(\omega_2) = 1/2$
- Payoffs are given (respectively) by:

		Type-ω_1 player 2		
		L	M	R
Player 1	T	(1,2ϵ)	(1,0)	(1,3ϵ)
	B	(2,2)	(0,0)	(0,3)

		Type-ω_2 player 2		
		L	M	R
Player 1	T	(1,2ϵ)	(1,3ϵ)	(1,0)
	B	(2,2)	(0,3)	(0,0)

- Where: $0 < \epsilon < 1/2$

Second (textbook) example: Information may hurt (contd.)

- Assume there's no private information
  - $s_2 = L$ as $br_2(T) = L$ and $br_2(B) = L$
  - $s_1 = B$ as $br_1(L) = B$
  - Equilibrium outcome is (2,2)
- Assume there's private information (2 knows her type)
  - Type ω_1: $s_2(\omega_1) = R$ as it's dominant
  - Type ω_2: $s_2(\omega_2) = M$ as it's dominant
  - Player 1: $br_1(s_2(.)) = T$
  - Equilibrium outcome is $(1,3\epsilon)$
  - As both 1 and 3ϵ are smaller than 2, everybody is worse off with incomplete information

Cournot duopoly with asymmetric information

- Two firms are engaged in **Cournot competition**, but one firm has private information about its costs
- Firm 1's cost function is $c_1(q_1) = cq_1$
- Firm 2's cost function is
  - $c_2(q_2) = c_H q_2$ with probability θ, and
  - $c_2(q_2) = c_L q_2$ with probability $1-\theta$, where $c_L < c_H$
- Firm 2 knows which cost function it has, but firm 1 does not $\rightarrow$ It only knows the **distribution** of firm 2's costs
- Both firms know the aggregate demand function, which is described by $p(Q) = a - Q$
- How do we find the Bayesian Nash Equilibrium of this game?

Cournot duopoly with asymmetric information (contd.)

- Let q_1^* be firm 1's optimal quantity choice
- Let q_{2H}^* and q_{2L}^* be firm 2's optimal choices when it has high and low costs, respectively
- If firm 2's cost is c_i, it will choose q_{2i}^* by maximizing

$$[(a - q_1^* - q_2) - c_i]q_2 = [a - q_1^* - c_i]q_2 - q_2^2$$

- Firm 1 will choose q_1^* by maximizing

$$\begin{aligned} \theta[(a - q_{2H}^* - q_1) - c]q_1 + (1-\theta)[(a - q_{2L}^* - q_1) - c]q_1 \\ = [a - \theta q_{2H}^* - (1-\theta)q_{2L}^* - c]q_1 - q_1^2 \end{aligned}$$

Cournot duopoly with asymmetric information (contd.)

- The FOCs for these 3 optimization problems are

$$\begin{aligned} q_{2H}^* &= \tfrac{1}{2}[a - q_1^* - c_H] \\ q_{2L}^* &= \tfrac{1}{2}[a - q_1^* - c_L] \\ q_1^* &= \tfrac{1}{2}[a - \theta q_{2H}^* - (1-\theta) q_{2L}^* - c] \end{aligned}$$

- Solving these equations yields

$$\begin{aligned} q_1^* &= \tfrac{1}{2}[a - \tfrac{1}{2}\theta(a - q_1^* - c_H) - \tfrac{1}{2}(1-\theta)(a - q_1^* - c_L) - c] \\ \text{or } q_1^* &= \tfrac{1}{3}[a - 2c + \theta c_H + (1-\theta)c_L] \end{aligned}$$

$$\begin{aligned} \text{and } q_{2H}^* &= \tfrac{1}{3}[a - 2c_H + c] + \tfrac{1}{6}(1-\theta)(c_H - c_L) \\ q_{2L}^* &= \tfrac{1}{3}[a - 2c_L + c] - \tfrac{1}{6}\theta(c_H - c_L) \end{aligned}$$

- The strategies $(q_1^*, q_{2H}^*, q_{2L}^*)$ constitute a Bayesian Nash equilibrium of the game

Cournot duopoly with asymmetric information (contd.)

- So far, we have not allowed for firm 2 to reveal its type to firm 1
- If it could, would firm 2 reveal its type to firm 1?
- To see this, let us contrast the BNE of incomplete-information Cournot to the NE complete-information Cournot where firm 2's costs are c_H

$$(q_1^{**}, q_2^{**}) = (\tfrac{1}{3}[a - 2c + c_H], \tfrac{1}{3}[a - 2c_H + c])$$

- Note that $q_{2H}^* > q_2^{**}$
- When firm 2's costs are high, firm 2 (1) produces *more* (*less*) in the incomplete-information game than it would in the complete-information game

Cournot duopoly with asymmetric information (contd.)

- In the complete-information game, firm 1 knows that firm 2 has high costs, and it exploits firm 2's high costs by increasing its own output
- In the incomplete-information game, firm 1 does not know whether firm 2 has high or low costs, so it produces a lower, "intermediate" level of output
- As a result, a firm 2 with high costs exploits its informational advantage
- The reverse happens with the complete-information game where firm 2's costs are c_L
- In the incomplete-information game, high-cost firm 2 has the incentive to keep private the information about its costs
- On the contrary, a low-cost firm 2 has the incentive to disclose information

Jury voting

- Pool of jurors ($i \in \{1, ..., I\}$) must decide whether a defendant is guilty or innocent
- True state of the world (unobserved by jurors) is one of the two: $\omega \in \{G, B\}$ (where B stands for innocent/blameless)
- Common prior about state of the world: $Prob(\omega = G) = \pi$
- But then each juror receives (private) signal about state of the world: $s_i \in \{g, b\}$
  - $Prob(s = g|\omega = G) = p$
  - $Prob(s = b|\omega = G) = 1 - p$
  - $Prob(s = b|\omega = B) = q$
  - $Prob(s = g|\omega = B) = 1 - q$
- In order for private signals to be informative, we must have:
  - $p > 1/2$, $q > 1/2$
  - and hence $p > 1 - q$
- Each signal realization is observed only by the receiving juror, and it thus ends up being her type

Jury voting (contd.)

- Each juror can vote either to convict or to acquit the defendant: $v_i \in \{c, a\}$
- Voting is by unanimity, that is, the defendant is convicted if the I-vector voting profile is $v = (c, ..., c)$, and she's acquitted otherwise
- To close the representation of the Bayesian games, we need to specify the jurors' payoffs:
  - $u_i = 0$ if $v = (c, ..., c)$ & $\omega = G$ or if $v \neq (c, ..., c)$ & $\omega = B$
  - $u_i = -z$ if $v = (c, ..., c)$ & $\omega = B$
  - $u_i = -(1 - z)$ if $v \neq (c, ..., c)$ & $\omega = G$
  - With $0 \leq z \leq 1$
- If $r > z$, the juror with posterior r prefers the defendant to be convicted
- Clearly, $z \to 1$ for Cesare Beccaria-like preferences, and $z \to 0$ for Avengers-like preferences

Jury voting (contd.)

One juror

- We want to check if sincere/informative strategy (i.e., voting according to the received signal) is an equilibrium
- The juror gets either b or g as signal. By Bayes' rule:

$$r = Prob(\omega = G|s = b) = \frac{\pi(1-p)}{\pi(1-p)+(1-\pi)q} \equiv \underline{z}_1$$

$$r = Prob(\omega = G|s = g) = \frac{\pi p}{\pi p+(1-\pi)(1-q)} \equiv \overline{z}_1$$

- If $z \geq \underline{z}_1$: acquittal is at least as good as conviction after receiving b
- If $z \leq \overline{z}_1$: conviction is at least as good as acquittal after receiving g
- Therefore, sincere/informative strategy is optimal iff: $\underline{z}_1 \leq z \leq \overline{z}_1$

Jury voting (contd.)

Two jurors

- We want to check if sincere/informative strategies are BNE
- Postulate that juror 2 votes a if b and c if g
- Consider the problem of type-b juror 1
- If juror 2 receives b, juror 1's vote has no effect (as you need unanimity for conviction)
- Therefore, juror 1 must update her posterior also to infer the probability of being pivotal

$$r = Prob(\omega = G | s_1 = b, s_2 = g) = \frac{\pi p(1-p)}{\pi p(1-p) + (1-\pi)(1-q)q} \equiv \underline{z}_2$$

- If $z \geq \underline{z}_2$: type-b juror 1 votes a

Jury voting (contd.)

Two jurors (contd.)

- Consider the problem of type-g juror 1

$$r = Prob(\omega = G | s_1 = g, s_2 = g) = \frac{\pi p^2}{\pi p^2 + (1-\pi)(1-q)^2} \equiv \overline{z}_2$$

- If $z \leq \overline{z}_2$: type-g juror 1 votes c
- Therefore, sincere/informative strategies are BNE iff: $\underline{z}_2 \leq z \leq \overline{z}_2$
- Note that $\underline{z}_2 > \underline{z}_1$
  - Less likely than with one juror to vote a if b
  - Why? Each juror less worried about convicting an innocent because she may not be pivotal
  - Problem get worse as I increases (free-riding annihilates Cesare Beccaria)
- Note also that $\overline{z}_2 > \overline{z}_1$

Where are we?

- We have (briefly) studied **static games of incomplete information** (or **Bayesian games**)
- References:
  - Lecture slides $\rightarrow$ 10 (final folder)
  - Osborne $\rightarrow$ chapter 9
  - Gibbons $\rightarrow$ chapter 3
- But the most interesting class of games of incomplete information involves some dynamics (and thus some information transmission). That's our next topic