17.810/17.811 – Game Theory

Lecture 6: Static Games of Incomplete Information

Asya Magazinnik

MIT

Where We Are/Where We Are Headed

- Recall that we have been focused thus far on games of complete information: where players know everything about the game, including all players, actions available to them, and associated payoffs
    - In games of imperfect information, we allowed for players not to always know how all players have acted in the past
- But what if we relaxed other forms of knowledge?
    - Other players' payoffs
    - Allowing players to be differentially informed
    - Allowing players to learn about the game as they play

Reading

These slides will focus on the following readings:

- Static Games
  - Gibbons, Chapter 3

Bayesian Games

A Bayesian Game in Normal Form

Bayes' Rule

Examples

Example 1: Jury Voting

Example 2: Public Goods and Incomplete Information

Example 3: Electoral Competition Under Uncertainty

Example 4: Rationalist Explanations for War (Fearon 1995)

Remember the Normal Form

In games of complete information, we had:

- set of N players
- a set of actions for each player: A_i
- a game form that maps action profiles into outcomes
- utilities over outcomes for each player

All of these components — players, actions available to all players, all associated payoffs, and game form — were known to all players.

- The only thing that might be unknown in games of complete information is how players have acted in the past (imperfect information).

Motivation for Bayesian Games

What if we want to think about situations where actors don't know everything or where some actors know more or different things than others? This motivates Bayesian Games!

- In war, leaders may not know their opponents' political costs for fighting or their military strength
- Potential challengers in elections don't know the quality of incumbents
- Leaders may not know the willingness of the public to pay for a new public library and therefore don't know if they should provide it
- Participants in auctions may not know other people's valuations of the item
- Lobbies and interest groups might know more about their policy areas than politicians or the public

Bayesian Games

A Bayesian Game in Normal Form

Bayes' Rule

Examples

Example 1: Jury Voting

Example 2: Public Goods and Incomplete Information

Example 3: Electoral Competition Under Uncertainty

Example 4: Rationalist Explanations for War (Fearon 1995)

A Bayesian Game in Normal Form (The Harsanyi Model)

A Bayesian game in normal form is:

- a set of players: $\{1, \ldots, n\}$
- a set of action spaces for each player: $A_1, \ldots, A_n$
- a set of type spaces for each player: $T_1, \ldots, T_n$
  - a player's type $t_i \in T_i$ is known privately to her but not necessarily to other players
- payoffs that are determined by the players' types and, as before, their actions: $u_i(a_1, \ldots, a_n; t_1, \ldots, t_n)$
- for each player, beliefs about the other players' types: $p_i(t_{-i}|t_i)$

Together, these components constitute a Bayesian game:
$G = \{A_1, \ldots, A_n; T_1, \ldots, T_n; p_1, \ldots, p_n; u_1, \ldots, u_n\}$

Intuition for a Bayesian Game

Think of the process of playing this game as:

- Players lack specific knowledge about the world; instead, they have perceptions or preconceived notions that we call beliefs
- Each player gets a signal about something (about another player's type, about the state of the world)
- Players update their beliefs with this new information using Bayes' Rule
- Then they choose the action that maximizes their expected utility given their best guess about the game they're playing

Intuition for a Bayesian Game

Harsanyi (1967) suggests a useful way of understanding a game of incomplete information as a game of imperfect information with a fictitious player who moves first called Nature:

1. Nature draws a type vector $(t_1, \ldots, t_n)$ for all players
2. For each i, Nature reveals t_i to some players (usually including i) and sends a signal to others
3. Players choose actions a_i simultaneously
4. Payoffs are realized: $u_i(a_1, \ldots, a_n; t_i)$ for all i

Not all players are able to observe Nature's first move, making this a game of imperfect information whereby players do not know the complete history of play.

Strategies in Bayesian Games

- A strategy is a contingent plan that tells the player what to do for every possible type Nature might draw.
- Thus, in a Bayesian Nash equilibrium, each player-type (not player!) chooses a strategy that maximizes her expected utility given the strategies of all the other player-types and the probability distribution over the types.
- Why do we need to consider strategies for a type that will never be realized? That is, once Nature has drawn a type and revealed it to some player *i*, why does it matter what he would have done had he been some other type?
    - Because the others don't know his type, but what they do depends on it, and what player *i* does in turn depends on that

Bayesian Nash Equilibrium

Because strategies are now functions of types, evaluation of best responses is somewhat complicated.

Definition

In the static Bayesian game $G = \{A_1, \ldots, A_n; T_1, \ldots, T_n; p_1, \ldots, p_n; u_1, \ldots, u_n\}$, the strategies $s^* = (s_1^*, \ldots, s_n^*)$ are a pure-strategy Bayesian Nash equilibrium if for each player i and for each of i's types $t_i \in T_i$, $s_i^*(t_i)$ solves:

$$\max_{a_i \in A_i} \sum_{t_{-i} \in T_{-i}} u_i(s_1^*(t_1), \ldots, s_{i-1}^*(t_{i-1}), a_i, s_{i+1}^*(t_{i+1}), \ldots, s_n^*(t_n); t) p_i(t_{-i} | t_i)$$

Bayesian Games

A Bayesian Game in Normal Form

Bayes' Rule

Examples

Example 1: Jury Voting

Example 2: Public Goods and Incomplete Information

Example 3: Electoral Competition Under Uncertainty

Example 4: Rationalist Explanations for War (Fearon 1995)

Bayes' Rule

What does it mean to be rational when you don't have complete information about the world?

- Making the best guess you can based on (1) what you observe and (2) your preexisting beliefs.

Bayes' Rule formalizes rational updating.

This is nothing more than an extension of the basic rules of probability theory, which we will now review.

Conditional Probability

What is the probability that two die sum to 3?

- Before the die are cast, the event (call it A) occurs with probability $\frac{2}{36}$ (1,2; 2,1).

Now suppose the first die comes up a 1, call this event B. What is the conditional probability that the two die sum to 3, given the first die came up 1?

- Consider the reduced sample space:

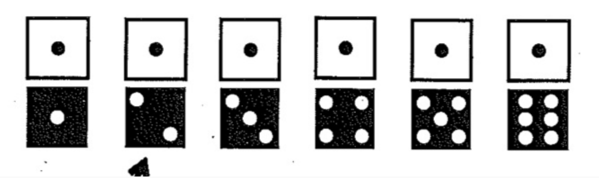

- Only one of these outcomes gives us 3, so $Pr(A|B) = 1/6$.

Conditional Probability

The conditional probability that some event A will happen, given that we know some event B happened, is defined as:

$$Pr(A|B) = \frac{Pr(A \cap B)}{Pr(B)}$$

Similarly,

$$Pr(B|A) = \frac{Pr(B \cap A)}{Pr(A)}$$

where $Pr(B \cap A) = Pr(A \cap B)$.

In the die example,

$$Pr(A|B) = \frac{1/36}{1/6} = \frac{1}{6}$$

Total Probability

- A and A' are mutually exclusive if when one happens the other cannot, e.g. head and tails on a coin flip
- We also say that some events, like heads and tails of a coin flip, are exhaustive if either one or the other must happen and no other event is possible
- A partition of all possible events is a collection of events that is both mutually exclusive and exhaustive

Law of Total Probability

Let $A_1, A_2, \cdots, A_n$ be disjoint events that form a partition of the sample space and assume $P(A_i) > 0$ for all i. Then for any event B we have:

$$\begin{aligned} Pr(B) &= Pr(A_1 \cap B) + \cdots + Pr(A_n \cap B) \\ &= Pr(A_1)Pr(B|A_1) + \cdots + Pr(A_n)Pr(B|A_n) \end{aligned}$$

where the second line comes from the definition of conditional probability.

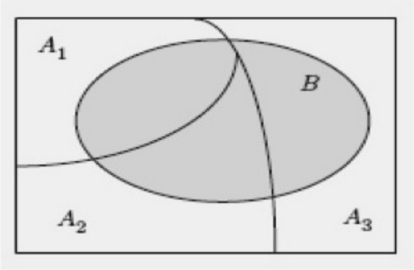

Inverting Conditional Probability: Bayes' Rule

Bayes' rule tells us how to go between $Pr(A|B)$ and $Pr(B|A)$. It states:

$$Pr(A|B) = \frac{Pr(A)Pr(B|A)}{Pr(B)}$$

We can derive this quite easily starting from the definition of conditional probability:

$$Pr(A|B) = \frac{Pr(A \cap B)}{Pr(B)} = \frac{Pr(B \cap A)}{Pr(B)} = \frac{Pr(A)Pr(B|A)}{Pr(B)}$$

We can expand the denominator using the law of total probability:

$$Pr(A|B) = \frac{Pr(A)Pr(B|A)}{Pr(A_1)Pr(B|A_1) + \cdots + Pr(A_n)Pr(B|A_n)}$$

Example: Nuclear Alarm

Example: Nuclear Alarm

- Suppose the alarm system at the nuclear plant is not completely reliable.
- If something is wrong, the probability the alarm goes off is .99.
- On the other hand, the alarm goes off with probability .01 on days when nothing is wrong.
- The historical record shows that the nuclear reactor has something wrong with it, for real, 1 out of 100 days.
- If the alarm just went off, what is the probability that something is actually wrong?

Example

So for our problem we have:

- Let A be the event that something is wrong with the reactor.
- Let B be the event that the alarm goes off. We are told:

$$Pr(B|A) = .99$$
$$Pr(B|\neg A) = .01$$
$$Pr(A) = .01$$
$$Pr(A|B) = ?$$

What to do?

PANIC

Calculate the Posterior Probability

With A being the event that something is wrong with the reactor and B the event that the alarm goes off, we apply Bayes' rule:

$$Pr(A|B) = \frac{Pr(A)Pr(B|A)}{Pr(A)Pr(B|A) + Pr(\neg A)Pr(B|\neg A)}$$

$$\begin{aligned} Pr(A) &= .01 \\ Pr(B|A) &= .99 \\ Pr(\neg A) = 1 - Pr(A) &= .99 \\ Pr(B|\neg A) &= .01 \end{aligned}$$

Thus we have:

$$Pr(A|B) = \frac{(.01)(.99)}{(.01)(.99) + (.99)(.01)} = \frac{1}{2}$$

Some More Intuition

- Consider 10000 reactor days. On how many days will there be trouble? $.01 \times 10000 = 100$
- On how many of those troubled days will the alarm actually go off? $.99 \times 100 = 99$
- On the other hand, there are 9900 days with no trouble. On how many of those days will there be a false alarm? $.01 \times 9900 = 99$
- So how likely is it that an alarm day is a trouble day? $50/50$

Bayes' Rule in Medicine

Do doctors understand test results?

By William Kremer
BBC World Service

7 July 2014

https://www.bbc.com/news/magazine-28166019

Example: Politician Types

- Suppose there is an incumbent politician who can be of two types (good or bad)
- Good incumbents generate low inflation with probability $2/3$
- Bad incumbents generate low inflation with probability $1/5$
- Ex ante the probability of a good incumbent is $3/4$

What is the likelihood the incumbent is good if we observe low inflation?

Example: Politician Types

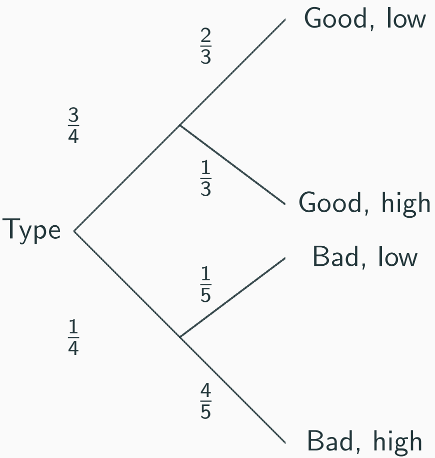

Example: Politician Types

Let A_1 be the event that the incumbent is good, A_2 be the event the incumbent is bad, and L be low inflation. Then Bayes' rule tells us:

$$\Pr(good|L) = \frac{\Pr(L|good)\Pr(good)}{\Pr(L|good)\Pr(good) + \Pr(L|bad)\Pr(bad)}$$

- $\Pr(good) = 3/4$
- $\Pr(bad) = 1/4$
- $\Pr(L|good) = 2/3$
- $\Pr(L|bad) = 1/5$

So we have:

$$\Pr(good|L) = \frac{2/3 \times 3/4}{2/3 \times 3/4 + 1/5 \times 1/4} = \frac{10}{11}$$

Bayesian Games

A Bayesian Game in Normal Form

Bayes' Rule

Examples

Example 1: Jury Voting

Example 2: Public Goods and Incomplete Information

Example 3: Electoral Competition Under Uncertainty

Example 4: Rationalist Explanations for War (Fearon 1995)

Bayesian Games

A Bayesian Game in Normal Form

Bayes' Rule

Examples

Example 1: Jury Voting

Example 2: Public Goods and Incomplete Information

Example 3: Electoral Competition Under Uncertainty

Example 4: Rationalist Explanations for War (Fearon 1995)

Application: Jury Voting

- Suppose that three jurors $N = \{1, 2, 3\}$ collectively choose an outcome $x \in \{c, a\}$.
- The jurors simultaneously cast ballots $v_i \in S_i = \{c, a\}$, and the outcome is chosen by majority rule.
- Each juror is uncertain whether or not the defendant is guilty, G, or innocent, I. So the set of state variables is $\Omega = \{G, I\}$.
- Each juror assigns prior probability $\pi > \frac{1}{2}$ to state G.

Application: Jury Voting

- The jurors each receive one unit of utility when the jury as a whole matches its verdict to the state of the world, and zero otherwise:
  - If the defendant is guilty, the jurors receive one unit of utility from convicting and zero from acquitting.
  - If the defendant is innocent, the jurors receive one unit from acquitting and zero from convicting.
- Before voting, each juror receives an independent, private signal about the defendant's guilt, $\theta_i \in \{0, 1\}$.
- The signal is informative, so that a juror is more likely to receive the signal $\theta_i = 1$ when the defendant is guilty than when the defendant is innocent.
- To keep things simple,
  - $\Pr(\theta_i = 1 \mid \omega = G) = \Pr(\theta_i = 0 \mid \omega = I) = p > \frac{1}{2}$
  - $\Pr(\theta_i = 0 \mid \omega = G) = \Pr(\theta_i = 1 \mid \omega = I) = 1 - p < \frac{1}{2}$

Application: Jury Voting

After receiving her signal, voter i selects her vote $v(\theta_i)$ to maximize the probability of a correct decision.

- Recall that each player-type should choose a strategy.
- Suppose that each voter uses the sincere strategy, $v_i(1) = c$ and $v_i(0) = a$.
- Sincere strategies constitute a Bayesian Nash equilibrium only if voter 1 uses this strategy when she believes that voters 2 and 3 also use it.
- We will now see if there is a Bayesian Nash equilibrium in which all voters play the sincere strategies $v_i(1) = c$ and $v_i(0) = a$.

Application: Jury Voting

- Generically, voter 1's expected utility of convict is:

$$\begin{aligned}
&\Pr(\theta_2 = 1 \text{ and } \theta_3 = 0 \text{ and } \omega = G \mid \theta_1) + \\
&\Pr(\theta_2 = 0 \text{ and } \theta_3 = 1 \text{ and } \omega = G \mid \theta_1) + \\
&\Pr(\theta_2 = 1 \text{ and } \theta_3 = 1 \text{ and } \omega = G \mid \theta_1) + \\
&\Pr(\theta_2 = 0 \text{ and } \theta_3 = 0 \text{ and } \omega = I \mid \theta_1)
\end{aligned}$$

- By similar logic, the expected utility of voting to acquit is:

$$\begin{aligned}
&\Pr(\theta_2 = 1 \text{ and } \theta_3 = 0 \text{ and } \omega = I \mid \theta_1) + \\
&\Pr(\theta_2 = 0 \text{ and } \theta_3 = 1 \text{ and } \omega = I \mid \theta_1) + \\
&\Pr(\theta_2 = 0 \text{ and } \theta_3 = 0 \text{ and } \omega = I \mid \theta_1) + \\
&\Pr(\theta_2 = 1 \text{ and } \theta_3 = 1 \text{ and } \omega = G \mid \theta_1)
\end{aligned}$$

- In all other cases, voter 1 gets 0 utility.

Application: Jury Voting

- The last two terms of each sum are the same and cancel out when comparing the utilities.

- Accordingly, voting to convict is a best response if and only if:

$$\Pr(\theta_2 = 1 \text{ and } \theta_3 = 0 \text{ and } \omega = G \mid \theta_1) + \Pr(\theta_2 = 0 \text{ and } \theta_3 = 1 \text{ and } \omega = G \mid \theta_1) \geq \Pr(\theta_2 = 1 \text{ and } \theta_3 = 0 \text{ and } \omega = I \mid \theta_1) + \Pr(\theta_2 = 0 \text{ and } \theta_3 = 1 \text{ and } \omega = I \mid \theta_1)$$

- We want to check that voting to convict is a best response when $\theta_1 = 1$ and that voting to acquit is a best response when $\theta_1 = 0$.

First let's check that voting to convict is a best response for juror 1 when $\theta_1 = 1$. Let's compute:

$$\Pr(\theta_2 = 1 \text{ and } \theta_3 = 0 \text{ and } \omega = G \mid \theta_1 = 1)$$

which is the same as:

$$\Pr(\theta_2 = 0 \text{ and } \theta_3 = 1 \text{ and } \omega = G \mid \theta_1 = 1)$$

(why?)

Application: Jury Voting

$$\Pr(\underbrace{\theta_2 = 0 \text{ and } \theta_3 = 1 \text{ and } \omega = G}_{A} \mid \underbrace{\theta_1 = 1}_{B})$$

Recall Bayes' Rule:

$$Pr(A|B) = \frac{Pr(B|A)Pr(A)}{Pr(B)}$$

Going piece by piece,

$$\begin{aligned} Pr(B|A) &= Pr(\theta_1 = 1 \mid \theta_2 = 0 \text{ and } \theta_3 = 1 \text{ and } \omega = G) \\ &= Pr(\theta_1 = 1 \mid \omega = G) = p \end{aligned}$$

by the independence of voters' signals.

Application: Jury Voting

$$\Pr(\underbrace{\theta_2 = 0 \text{ and } \theta_3 = 1 \text{ and } \omega = G}_{A} \mid \underbrace{\theta_1 = 1}_{B})$$

Recall Bayes' Rule:

$$Pr(A|B) = \frac{Pr(B|A)Pr(A)}{Pr(B)}$$

Going piece by piece,

$$\begin{aligned} Pr(A) = Pr(\theta_2 = 0 \text{ and } \theta_3 = 1 \text{ and } \omega = G) \\ = Pr(\omega = G) * Pr(\theta_2 = 0 \mid \omega = G) * Pr(\theta_3 = 1 \mid \omega = G) \\ = \pi(1 - p)(p) \end{aligned}$$

by the LTP and again the independence of voters' signals.

Application: Jury Voting

$$\Pr(\underbrace{\theta_2 = 0 \text{ and } \theta_3 = 1 \text{ and } \omega = G}_{A} \mid \underbrace{\theta_1 = 1}_{B})$$

Recall Bayes' Rule:

$$Pr(A|B) = \frac{Pr(B|A)Pr(A)}{Pr(B)}$$

Going piece by piece,

$$\begin{aligned} Pr(B) &= Pr(\theta_1 = 1) \\ &= Pr(\theta_1 = 1 \mid \omega = G)Pr(\omega = G) + Pr(\theta_1 = 1 \mid \omega = I)Pr(\omega = I) \\ &= p\pi + (1-p)(1-\pi) \end{aligned}$$

by the LTP.

Application: Jury Voting

Putting it together,

$$\Pr(\theta_2 = 0 \text{ and } \theta_3 = 1 \text{ and } \omega = G \mid \theta_1 = 1) = \frac{\pi p^2 (1-p)}{\pi p + (1-\pi)(1-p)}$$

Exercise: Verify that, by similar logic:

$$\begin{aligned}
& \Pr(\theta_2 = 1 \text{ and } \theta_3 = 0 \text{ and } \omega = I \mid \theta_1 = 1) \\
= \; & \Pr(\theta_3 = 1 \text{ and } \theta_2 = 0 \text{ and } \omega = I \mid \theta_1 = 1) \\
= \; & \frac{(1-\pi)\, p\, (1-p)^2}{\pi p + (1-\pi)(1-p)}
\end{aligned}$$

Application: Jury Voting

- Thus, $v_i(1) = c$ is optimal for juror 1 if:

$$2\frac{\pi p^2(1-p)}{\pi p + (1-\pi)(1-p)} \geq 2\frac{(1-\pi)p(1-p)^2}{\pi p + (1-\pi)(1-p)}$$

- After simplifying and rearranging, this inequality becomes:

$$\frac{\pi p^2(1-p)}{\pi p^2(1-p) + (1-\pi)p(1-p)^2} \geq \frac{1}{2}$$

- The left-hand side is simply the conditional probability of guilt given two signals of $\theta = 1$ and one signal of $\theta = 0$.
  - Agent 1 wants to vote to convict if she believes that the defendant is more likely to be guilty than innocent after conditioning on her signal and the belief that she is pivotal.

Application: Jury Voting

- Similarly, the requirement for a vote of innocence conditional on a signal of 0 is:

$$\frac{\pi p(1-p)^2}{\pi p(1-p)^2 + (1-\pi)p^2(1-p)} \leq \frac{1}{2}$$

- The upshot is that truthful behavior among jurors requires:

$$\frac{2\pi p}{\pi p + (1-\pi)(1-p)} \geq 1 \geq \frac{2\pi(1-p)}{\pi(1-p) + (1-\pi)p}$$

- We can reduce this messy thing with the help of Wolfram Alpha, which tells us:

$$\pi \leq p \leq \frac{\pi}{2\pi - 1} \text{ for } \frac{1}{2} < \pi < 1$$

- We simply require $\pi \leq p$. The signal has to be good relative to the probability of guilt.

Bayesian Games

A Bayesian Game in Normal Form

Bayes' Rule

Examples

Example 1: Jury Voting

Example 2: Public Goods and Incomplete Information

Example 3: Electoral Competition Under Uncertainty

Example 4: Rationalist Explanations for War (Fearon 1995)

Application: Public Goods and Incomplete Information

- Let's consider a public good provision game similar to the crime reporting game that we have studied before
- A version of the Palfrey-Rosenthal contribution game in which potential contributors are uncertain about the contribution costs of other players
- Every agent receives a utility of 1 if at least k agents contribute and 0 otherwise
- Agent i pays a cost c_i to contribute where c_i is distributed uniformly on $[0, 1]$
- Each agent learns her own cost, but remains uncertain about the other players' costs

The case of $k = 1$

- Consider the case where the good is provided for everyone if at least one person contributes ($k = 1$)

- Let's focus on symmetric equilibria where all players play a cutpoint strategy in which agent i contributes if and only if $c_i < \widehat{c}_n$ where $\widehat{c}_n$ is an equilibrium cutpoint for the game with n players

- If everyone else also plays this cutpoint strategy, player i's utility from contributing is $1 - c_i$

- If she does not contribute, she receives 1 if there is at least one contributor and 0 otherwise

The case of $k = 1$

- Because c is distributed uniformly on $[0, 1]$, each player contributes with probability:

$$Pr(c_i \leq \hat{c}_n) = \frac{\hat{c}_n - 0}{1 - 0} = \hat{c}_n$$

where we have simply used the CDF of the uniform distribution.

- Thus the probability that none of the other players contributes is $[1 - \widehat{c}_n]^{n-1}$
- Agent i's utility from not contributing is:

$$\begin{aligned} &Pr(\text{no one contributes}) * 0 + Pr(\text{someone else contributes}) * 1 \\ &= 1 - [1 - \widehat{c}_n]^{n-1} \end{aligned}$$

- Accordingly, agent i contributes so long as:

$$1 - c_i \geq 1 - [1 - \widehat{c}_n]^{n-1} \implies [1 - \widehat{c}_n]^{n-1} \geq c_i$$

The case of $k = 1$

- Because an agent with cost $\widehat{c}_n$ must be indifferent over her choices, a Bayesian Nash equilibrium requires:

$$[1 - \widehat{c}_n]^{n-1} = \widehat{c}_n$$

- This equilibrium condition implies:

$$n - 1 = \frac{\ln(\widehat{c}_n)}{\ln(1 - \widehat{c}_n)}$$

Note that as n grows large, $\hat{c}_n$ must get smaller (closer to 0) to make the equality hold. Thus, familiarly, a player's probability of making a contribution decreases with group size. (This is our classic free-riding result.)

The case of $k = 1$

- However, it is also the case that the probability that no agent contributes, $[1 - \widehat{c}_n]^n$, converges to 0 as n grows large
  - This is because:
    $[1 - \widehat{c}_n]^n = [1 - \widehat{c}_n]^{n-1}(1 - \widehat{c}_n) = \widehat{c}_n(1 - \widehat{c}_n) \quad 0$

- So in this model, although the probability that any particular agent contributes vanishes as n gets large, the probability of public good provision $(1 - [1 - \widehat{c}_n]^n)$ converges to 1

- Compare this result against the result we had with complete information, where we found that the probability no one else calls did not depend on n

Bayesian Games

- A Bayesian Game in Normal Form
- Bayes' Rule

Examples

- Example 1: Jury Voting
- Example 2: Public Goods and Incomplete Information
- Example 3: Electoral Competition Under Uncertainty
- Example 4: Rationalist Explanations for War (Fearon 1995)

Electoral Competition Under Uncertainty

- We now return to the Hotelling-Downs model of candidate competition with positioning on the real line
- Recall that with complete information, this gave us the median voter theorem: candidates position at the median voter's ideal point regardless of their own preference
- We now consider an extension: candidates have ideal points 0 (Democrat) and 1 (Republican) as before, but they don't know the location of the median voter — only that it is distributed uniformly on [0,1]
    - Median voter's location is a random variable ω, with $F(\omega) = \omega$ on $[0, 1]$
- Candidates' preferences are common knowledge

Electoral Competition Under Uncertainty

Assume policy-motivated candidates with quadratic utilities:

$$u_1(x) = -x^2$$
$$u_2(x) = -(x-1)^2$$

Given two platforms $s_1 < s_2$, candidate 1 wins if the median is closer to s_1 than s_2. This is true if the median is less than $\frac{s_1+s_2}{2}$.

Thus the probability that candidate 1 wins is given by:

$$Pr(1 \text{ wins}) = Pr\left(\omega < \frac{s_1+s_2}{2}\right) = \frac{s_1+s_2}{2}$$

Electoral Competition Under Uncertainty

Let's look for an equilibrium in which $s_1 < s_2$ (why does this make sense?)

Writing down their expected utilities:

$$EU_1(s_1, s_2) = \left(\frac{s_1 + s_2}{2}\right)(-s_1^2) + \left(1 - \frac{s_1 + s_2}{2}\right)(-s_2^2)$$

$$EU_2(s_1, s_2) = \left(\frac{s_1 + s_2}{2}\right)(-(s_1 - 1)^2) + \left(1 - \frac{s_1 + s_2}{2}\right)(-(s_2 - 1)^2)$$

So player 1 chooses s_1 to optimize:

$$\max_{s_1}\left\{\left(\frac{s_1 + s_2}{2}\right)(-s_1^2) + \left(1 - \frac{s_1 + s_2}{2}\right)(-s_2^2)\right\}$$

$$= \max_{s_1}\left\{-\frac{s_1^3}{2} - \frac{s_1^2 s_2}{2} - s_2^2 + \frac{s_1 s_2^2}{2} + \frac{s_2^3}{2}\right\}$$

Yielding the FOC: $-\frac{3s_1^2}{2} - s_1 s_2 + \frac{s_2^2}{2} = 0$

Electoral Competition Under Uncertainty

$$-\frac{3s_1^2}{2} - s_1 s_2 + \frac{s_2^2}{2} = 0 \to (-3s_1 + s_2)\left(\frac{s_1}{2} + \frac{s_2}{2}\right) = 0$$
$$\to s_1 = \frac{s_2}{3} \text{ or } s_1 = -s_2$$

Since we are constrained on [0,1], the only solution is:

$$s_1^*(s_2) = \frac{s_2}{3}$$

Repeat the same exercise for player 2 (check this) and you get:

$$s_2^*(s_1) = \frac{2}{3} + \frac{s_1}{3}$$

Solving this system of equations, the unique solution is:

$$s_1^* = \frac{1}{4} \text{ and } s_2^* = \frac{3}{4}$$

Electoral Competition Under Uncertainty

Thus uncertainty about the median voter's location gives us divergence of candidate platforms!

- The candidates take a gamble that the median voter will be on their side and that they'll get policy closer to their ideal point.

Bayesian Games

A Bayesian Game in Normal Form

Bayes' Rule

Examples

Example 1: Jury Voting

Example 2: Public Goods and Incomplete Information

Example 3: Electoral Competition Under Uncertainty

Example 4: Rationalist Explanations for War (Fearon 1995)

Why do states go to war?

- Note that in nearly every model of conflict over a divisible resource that we've studied thus far, there was no fighting in equilibrium
  - War is costly.
  - With complete information, rational actors who know their relative strength and the costs of war can usually bargain into a better outcome
  - Why, then, do wars ever occur?
- We might consider three types of argument:
  1. Leaders are irrational.
  2. Leaders are rational but do not pay the costs of war.
  3. Leaders make rational miscalculations due to incomplete information.
- Let's focus on this third (neorealist) perspective.

A Toy Model

- Consider two states, A and B, with competing claims over a territory
- Let $x \in [0, 1]$ be the share controlled by A and $1 - x$ be the share controlled by B
- The state that wins gets the whole territory but both states pay costs of war c_A and c_B
- Let $p \in [0, 1]$ be the probability that A wins if they go to war
- Consider simple linear utilities: $u_A(x) = x$ and $u_B(1 - x) = 1 - x$

A Toy Model

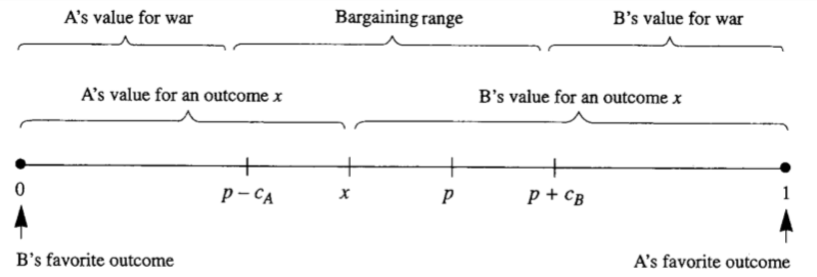

The expected utilities of war are:

$$EU_A(war) = p(1) + (1-p)(0) - c_A = p - c_A$$
$$EU_B(war) = p(0) + (1-p)(1) - c_B = 1 - p - c_B$$

Both states strictly prefer a negotiated settlement in the interval $(p - c_A, p + c_B)$ to fighting.

A Toy Model

- A concrete example: players bargaining over \$100.
- They can come to an agreement or they can play a costly lottery: 50/50 chance of winning, \$20 to play
- Choosing the lottery has expected value \$30 $(.5 * 100 + .5 * 0 - 20)$
- Thus any agreement between (\$31, \$69) and (\$69, \$31) is mutually beneficial (assuming money is divisible into \$1 increments)
- Risk aversion or discounting will increase this range even further

War Due to Private Information

We consider three rationalist explanations for war relating to players' private information:

1. Disagreements about relative power (p)
2. Miscalculation about opponent's willingness to fight (c_A and c_B)
3. Incentives to misrepresent private information (p, c_A, c_B) during the bargaining process

(1) Disagreements About Relative Power

- Staying with the \$100 example, suppose both players believe that they will win with probability 0.8. Then each will have an expected utility of playing the costly lottery (war) of $0.8 * 100 + 0.2 * 0 - 20 = 60$
    - Then there is no division of the \$100 that both will accept!
- Is this truly a rationalist explanation? How can rational states come to hold conflicting views?
    - Mutual optimism, emotional commitments, nationalism
    - Same information $\rightarrow$ different conclusions
    - Different information $\rightarrow$ different conclusions
- Harsanyi: Two rational agents presented with the same information should reach the same conclusions. So only the third is a fully rationalist account.

(2) Opponent's Willingness to Fight

- Recall that willingness to fight is determined by the bargaining range, $(p - c_A, p + c_B)$
- We have considered states holding different beliefs about p, but beliefs about c_A and c_B also affect this bargaining range
- Suppose A gets to propose an allocation. It will propose $p + c_B$, the largest share to A that makes B indifferent between compromise and war
- But what if A doesn't know exactly what c_B is?
- As in the electoral competition game, A's uncertainty about the location of B's indifference point creates incentives to reach for a bigger power grab
- But if A overreaches $\rightarrow$ war

(2) Opponent's Willingness to Fight

- Historical examples abound:
  - Germany miscalculated Russian/British willingness to fight in 1914
  - Japan miscalculated U.S. willingness to engage in long conflict over South Pacific in 1941
  - U.S. miscalculated China's willingness to defend North Korea
- Again we ask: are such explanations truly rationalist?
  - Going to war when there was a nonempty bargaining range is an inefficient outcome that both parties have incentives to avoid → why not simply communicate private information in that case?
  - Thus a fully rationalist explanation has to explain why leaders chose not to faithfully communicate their private information

(3) Incentives to Misrepresent in Bargaining

- While states have an incentive to reach agreement, they also have an incentive to misrepresent their private information if it helps them get a better bargain
- Consider a simple adaptation where, before bargaining, B gets to make an announcement f, which can be any statement that's informative about its power and/or costs
- It can be shown that if this announcement does not change anyone's payoffs, it does not change equilibrium outcomes
  - B has an incentive to communicate information to increase its share of the territory
  - Knowing this, A ignores B's message completely
- (Note, however, that sometimes cheap talk *can* matter in bargaining.)

(3) Incentives to Misrepresent in Bargaining

Now what if we made the signal f costly to send?

- e.g. building weapons, mobilizing troops, signing alliance treaties, supporting troops in a foreign land, creating domestic political costs that would be paid if the announcement is false ("audience costs")

To be effective, the signal has to be costly such that a state with less resolve or capability may not wish to send it.

- Actions that generate a higher risk of war, such as troop mobilizations or belligerent rhetoric before a domestic audience, satisfy this requirement
- Thus the benefit of communication is offset by some real risk of war

Other considerations: reputational effects, armed conflict as signal (for both strong and weak states)

Other Rationalist Explanations?

Fearon also offers some additional rationalist explanations for war beyond informational ones:

- Commitment problems
- Indivisible resources