MS&E 246: Lecture 12
Static games of incomplete information

Ramesh Johari

Incomplete information

- *Complete information* means the entire structure of the game is common knowledge
- *Incomplete information* means "anything else"

Cournot revisited

Consider the Cournot game again:

- Two firms ($N = 2$)
- Cost of producing s_n : $c_n s_n$
- *Demand curve*:

  Price = $P(s_1 + s_2) = a - b\,(s_1 + s_2)$
- Payoffs:

  Profit = $\Pi_n(s_1, s_2) = P(s_1 + s_2)\,s_n - c_n s_n$

Cournot revisited

Suppose c_i is known only to firm i.

Incomplete information:

Payoffs of firms are not common knowledge.

How should firm i *reason* about what it expects the competitor to produce?

Beliefs

A first approach:

Describe firm 1's *beliefs* about firm 2.

Then need firm 2's beliefs about firm 1's beliefs...

And firm 1's beliefs about firm 2's beliefs about firm 1's beliefs...

Rapid growth of complexity!

Harsanyi's approach

Harsanyi made a key breakthrough in the analysis of incomplete information games:

He represented a *static game of incomplete information* as a *dynamic game of complete but imperfect information.*

Harsanyi's approach

- *Nature* chooses (c_1, c_2) according to some probability distribution.
- Firm i now knows c_i, but opponent only knows the *distribution* of c_i.
- Firms maximize *expected* payoff.

Harsanyi's approach

All firms share the *same* beliefs about incomplete information, determined by the *common prior*.

Intuitively:

Nature determines "who we are" at time 0, according to a distribution that is common knowledge.

Bayesian games

In a Bayesian game, player i's preferences are determined by his *type* $\theta_i \in T_i$; T_i is the *type space* for player i.

When player i has type θ_i, his payoff function is: $\Pi_i(\mathbf{s} ; \theta_i)$

Bayesian games

A *Bayesian game* with N players is a dynamic game of $N + 1$ stages.

Stage 0: *Nature* moves first, and chooses a type θ_i for each player i, according to some joint distribution $P(\theta_1, \ldots, \theta_N)$ on $T_1 \times \cdots \times T_N$.

Bayesian games

Player i learns his own type θ_i, but not the types of other players.

Stage i: Player i chooses an action from S_i.

After stage N: payoffs are realized.

Bayesian games

An alternate, equivalent interpretation:

1) Nature chooses players' types.
2) All players *simultaneously* choose actions.
3) Payoffs are realized.

Bayesian games

- How many subgames are there?
- How many information sets does player i have?
- What is a strategy for player i?

Bayesian games

- How many subgames are there?

  ONE – the entire game.

- How many information sets does player i have?

  ONE per type θ_i.

- What is a strategy for player i?

  A function from $T_i \to S_i$.

Bayesian games

One subgame $\Rightarrow$
SPNE and NE are identical concepts.

A *Bayesian equilibrium*
(or *Bayes-Nash equilibrium*)
is a NE of this dynamic game.

Expected payoffs

How do players reason about uncertainty regarding *other players*' *types*?

They use *expected payoffs*.

Given $\mathbf{s}_{-i}(\cdot)$, player i chooses $a_i = s_i(\theta_i)$ to maximize

$$E[\ \Pi_i(\ a_i\ ,\ \mathbf{s}_{-i}(\theta_{-i})\ ;\ \theta_i\ )\ |\ \theta_i\ ]$$

Bayesian equilibrium

Thus $s_1(\cdot)$, ..., $s_N(\cdot)$
is a Bayesian equilibrium if and only if:

$s_i(\theta_i) \in \arg\max_{a_i \in S_i} E[\ \Pi_i(\ a_i\ , \mathbf{s}_{-i}(\theta_{-i})\ ;\ \theta_i)\ |\ \theta_i\ ]$

for all θ_i, and for all players i.

[Notation: $\mathbf{s}_{-i}(\theta_{-i}) = (s_1(\theta_1), \ldots, s_{i-1}(\theta_{i-1}), s_{i+1}(\theta_{i+1}), \ldots, s_N(\theta_N))$]

Bayesian equilibrium

The conditional distribution $P(\theta_{-i} \mid \theta_i)$ is called player i's *belief.*

Thus in a Bayesian equilibrium, players maximize expected payoffs given their beliefs.

[*Note:* Beliefs are found using *Bayes' rule*:

$$P(\theta_{-i} \mid \theta_i) = P(\theta_1, \ldots, \theta_N)/P(\theta_i)$$]

Bayesian equilibrium

Since Bayesian equilibrium is a NE
of a certain dynamic game
of imperfect information,
it is guaranteed to exist
(as long as type spaces T_i are finite
and action spaces S_i are finite).

Cournot revisited

Back to the Cournot example:

Let's assume c_1 is common knowledge, but firm 1 does not know c_2.

Nature chooses c_2 according to:

c_2 = c_H, w/prob. p

= c_L, w/prob. 1 - p

(where $c_H > c_L$)

Cournot revisited

These are also informally called games of *asymmetric information*:

Firm 2 has information that firm 1 does not have.

Cournot revisited

- The *type* of each firm is their marginal cost of production, c_i.
- Note that c_1, c_2 are *independent*.
- Firm 1's *belief*:

  P(c_2 = c_H | c_1) = 1 - P(c_2 = c_L | c_1) = p
- Firm 2's *belief*: Knows c_1 exactly

Cournot revisited

- The *strategy* of firm 1 is the quantity s_1.
- The *strategy* of firm 2 is a function:

  $s_2(c_H)$: quantity produced if $c_2 = c_H$

  $s_2(c_L)$: quantity produced if $c_2 = c_L$

Cournot revisited

We want a NE of the dynamic game.

Given s_1 *and* c_2, Firm 2 plays a best response.

Thus given s_1 and c_2, Firm 2 produces:

$$R_2(s_1) = \left[\frac{a - c_2}{2b} - \frac{s_1}{2}\right]^+$$

Cournot revisited

Given $s_2(\cdot)$ and c_1,
Firm 1 maximizes *expected payoff*.

Thus Firm 1 maximizes:

$$E[\, \Pi_1 (s_1,\ s_2(c_2)\ ;\ c_1)\ |\ c_1\,] =$$

$$[p\ P(s_1 + s_2(c_H)) + (1 - p)P(s_1 + s_2(c_L))]\ s_1 - c_1 s_1$$

Cournot revisited

- Recall demand is *linear:* $P(Q) = a - b\,Q$
- So expected payoff to firm 1 is:

  $[\, a - b(s_1 + p\, s_2(c_H) + (1 - p)s_2(c_L))\,]s_1 - c_1\, s_1$
- Thus firm 1 plays best response to *expected production of firm 2.*

$$R_1(s_2) = \left[\frac{a - c_1}{2b} - \frac{ps_2(c_H) + (1 - p)s_2(c_L)}{2}\right]^+$$

Cournot revisited: equilibrium

- A Bayesian equilibrium has 3 unknowns:

  s_1, $s_2(c_H)$, $s_2(c_L)$
- There are 3 equations:

  Best response of firm 1 given $s_2(c_H)$, $s_2(c_L)$

  Best response of firm 2 given s_1,
  when type is c_H

  Best response of firm 2 given s_1,
  when type is c_L

Cournot revisited: equilibrium

- Assume all quantities are positive at BNE (can show this must be the case)
- Solution:

s_1 = [a - 2c_1 + pc_H + (1 - p)c_L]/3

$s_2(c_H)$ = [a - 2c_H + c_1]/3 + (1 - p)(c_H - c_L)/6

$s_2(c_L)$ = [a - 2c_L + c_1]/3 - p(c_H - c_L)/6

Cournot revisited: equilibrium

When p = 0, *complete information*:

Both firms know $c_2 = c_L \Rightarrow$ NE $(s_1^{(0)},\ s_2^{(0)})$

When p = 1, *complete information*:

Both firms know $c_2 = c_H \Rightarrow$ NE $(s_1^{(1)},\ s_2^{(1)})$

For 0 < p < 1, note that:

$s_2(c_L) < s_2^{(0)}, \quad s_2(c_H) > s_2^{(1)}, \quad s_1^{(0)} < s_1 < s_1^{(1)}$

Why?

Coordination game

Consider coordination game with incomplete information:

		Player 2	
		L	*R*
Player 1	*l*	$(2 + t_1, 1)$	(0,0)
	r	(0,0)	$(1, 2 + t_2)$

Coordination game

t_1, t_2 are independent uniform r.v.'s on $[0, x]$.

Player i learns t_i, but not t_{-i}.

Player 2

		L	*R*
Player 1	*l*	$(2 + t_1, 1)$	(0,0)
	r	(0,0)	$(1, 2 + t_2)$

Coordination game

Types: t_1, t_2

Beliefs:

Types are independent, so player i believes t_{-i} is uniform[0, x]

Strategies:

Strategy of player i is a function $s_i(t_i)$

Coordination game

We'll search for a specific form of Bayesian equilibrium:

Assume each player i has a *threshold* c_i, such that the strategy of player 1 is:

$s_1(t_1)$ = *l* if $t_1 > c_1$; = *r* if $t_1 \leq c_1$.

and the strategy of player 2 is:

$s_2(t_2)$ = *R* if $t_2 > c_2$; = *L* if $t_2 \leq c_2$.

Coordination game

Given $s_2(\cdot)$ and t_1, player 1 maximizes expected payoff:

If player 1 plays *l*:

E [Π_1(*l*, $s_2(t_2)$; t_1) | t_1] =

$(2 + t_1) \cdot P(t_2 \leq c_2) + 0 \cdot P(t_2 > c_2)$

If player 1 plays *r*:

E [Π_1(*r*, $s_2(t_2)$; t_1) | t_1] =

$0 \cdot P(t_2 \leq c_2) + 1 \cdot P(t_2 > c_2)$

Coordination game

Given $s_2(\cdot)$ and t_1, player 1 maximizes expected payoff:

If player 1 plays *l*:

E [Π_1(*l*, $s_2(t_2)$; t_1) | t_1] =

$(2 + t_1)(c_2/x)$

If player 1 plays *r*:

E [Π_1(*r*, $s_2(t_2)$; t_1) | t_1] =

$1 - c_2/x$

Coordination game

So player 1 should play *l* if and only if:

$t_1 > x / c_2 - 3$

Similarly:

Player 2 should play *R* if and only if:

$t_2 > x / c_1 - 3$

Coordination game

So player 1 should play *l* if and only if:

$t_1 > x / c_2 - 3$

Similarly:

Player 2 should play *R* if and only if:

$t_2 > x / c_1 - 3$

The right hand sides must be the thresholds!

Coordination game: equilibrium

Solve: $c_1 = x / c_2 - 3$, $c_2 = x / c_1 - 3$

Solution:

$$c_i = \frac{\sqrt{9+4x} - 3}{2}$$

Thus *one* Bayesian equilibrium is to play strategies $s_1(\cdot)$, $s_2(\cdot)$, with thresholds c_1, c_2 (respectively).

Coordination game: equilibrium

What is the *unconditional* probability that player 1 plays *r* ?

= $c_1 / x \to$ 1/3 as $x \to 0$

Thus as $x \to 0$, the unconditional distribution of play matches the *mixed strategy NE of the complete information game*.

Purification

This phenomenon is one example of *purification*:

recovering a mixed strategy NE via Bayesian equilibrium of a perturbed game.

Harsanyi showed mixed strategy NE can "almost always" be purified in this way.

Summary

To find Bayesian equilibria,
provide strategies $s_1(\cdot)$, ..., $s_N(\cdot)$ where:

For each type θ_i, player i chooses $s_i(\theta_i)$ to maximize expected payoff given his belief P(θ_{-i} | θ_i).