Static games of incomplete information

Mattias Polborn, Vanderbilt

YES summer school

July 12, 2024

Introduction

Incomplete Information: At least one player does not know the complete description of the game, i.e. he does not know either one or several of the following:

- The payoff function of the other player(s)
- The available strategies of the other player(s)
- The information available for the other player(s)

Example 1: Test-optional admissions

- You want to apply to Woke University, your dream school.
- There are mandatory and optional parts of your application.
- From the mandatory parts alone, a reasonable admissions officer at Woke U would estimate your quality at 50, but it could be anything between 0 and 100.
- You have optional information that you can include, but don't have to include. The admissions officer knows that you have this information.
- Unfortunately for you, after seeing the optional information, an admissions officer would overall estimate you at a 30.
- Assuming you want to be estimated as high as possible, *and that the admissions officers at Woke U are rational*, should you reveal the optional information?

Example 1: Test-optional admissions

- You want to apply to Woke University, your dream school.
- There are mandatory and optional parts of your application.
- From the mandatory parts alone, a reasonable admissions officer at Woke U would estimate your quality at 50, but it could be anything between 0 and 100.
- You have optional information that you can include, but don't have to include. The admissions officer knows that you have this information.
- Unfortunately for you, after seeing the optional information, an admissions officer would overall estimate you at a 30.
- Assuming you want to be estimated as high as possible, *and that the admissions officers at Woke U are rational*, should you reveal the optional information?

Example 2: Accident at night

- You had an accident on an unlit road outside the city. Your cell phone is dead, so you cannot get help yourself.
- Other people in your country are friendly, but not excessively so. Everyone gets a payoff of 0 if you don't get help, 1 if you do get help, and have a cost c if they are the one to call for help. c is iid drawn from a uniform distribution on $[0, 2]$, and everyone knows their own c_i draw, but not the cost of others.
- Scenario 1: One car is on the horizon, and it is well known to all that this is most likely the only driver that will be in a position to help.
- Scenario 2: Two cars arrive simultaneously at the site (and it is known that these two are likely the only ones in a position to help). Drivers simultaneously decide whether to help. If both do, they each have to pay their cost.
- In which scenario is you chance of receiving help higher?

Example 2: Accident at night

- You had an accident on an unlit road outside the city. Your cell phone is dead, so you cannot get help yourself.
- Other people in your country are friendly, but not excessively so. Everyone gets a payoff of 0 if you don't get help, 1 if you do get help, and have a cost c if they are the one to call for help. c is iid drawn from a uniform distribution on $[0, 2]$, and everyone knows their own c_i draw, but not the cost of others.
- Scenario 1: One car is on the horizon, and it is well known to all that this is most likely the only driver that will be in a position to help.
- Scenario 2: Two cars arrive simultaneously at the site (and it is known that these two are likely the only ones in a position to help). Drivers simultaneously decide whether to help. If both do, they each have to pay their cost.
- In which scenario is you chance of receiving help higher?

Example 3: Garbage dump

- City C needs to have a new garbage dump. There are three potential sites in villages V1, V2, V3.
- Having a garbage dump in the village is unpleasant (d_i) and C cannot force any village to accept it without compensation.
- C's value of having a new garbage dump is 3, and from C's (and everyone else's) point of view, the villages' damages are iid draws from $[0, 1]$
- C wants to set up an optimal mechanism. Should they:
  - Offer a compensation of 0.5 and build a new dump iff at least one village accepts this offer?
  - Starting from 0, slowly increase the compensation they offer until the first village accepts?
  - Ask all villages for their demands and promise (ex-ante) to build in the village that asks for the smallest compensation, but pay them the (higher) demand of the village that made the second-lowest demand?

Test-optional admissions

- Reasonable type of strategy: Cutoff k such that all types report iff $Q \geq k$
- If all types with $Q \geq k$ report, what should a rational admissions officer infer when there is no report?
- Expected value $= \frac{k+0}{2}$
- There can be no types in $\left[\frac{k}{2}, k\right]$ because these types would make a mistake by not reporting
- Only $k = 0$ is consistent with retionality

Test-optional admissions

- Reasonable type of strategy: Cutoff k such that all types report iff $Q \geq k$
- If all types with $Q \geq k$ report, what should a rational admissions officer infer when there is no report?
- Expected value $= \frac{k+0}{2}$
- There can be no types in $\left[\frac{k}{2}, k\right]$ because these types would make a mistake by not reporting
- Only $k = 0$ is consistent with retionality

Test-optional admissions

- Reasonable type of strategy: Cutoff k such that all types report iff $Q \geq k$
- If all types with $Q \geq k$ report, what should a rational admissions officer infer when there is no report?
- Expected value $= \frac{k+0}{2}$
- There can be no types in $\left[\frac{k}{2}, k\right]$ because these types would make a mistake by not reporting
- Only $k = 0$ is consistent with retionality

Public good provision

- Two players; if at least one contributes, both players get to enjoy a payoff of 1
- Whoever contributes has a cost c_i
- Both players draw their own c_i from cdf $F()$.

	Contribute	Don't cont.
C	$1 - c_1, 1 - c_2$	$1 - c_1, 1$
D	$1, 1 - c_2$	$0, 0$

Public good provision

- Cutoff equilibrium: Contribute iff $c_i \leq c_i^*$

$$1 - c_1^* = F(c_2^*) \text{ and } 1 - c_2^* = F(c_1^*)$$

- Suppose $F(x) = x/2$ for $x \in [0, 2]$ (i.e., uniform on $[0, 2]$)
$\Rightarrow c_1^* = c_2^* = 2/3$
- With two players, each helps with probability $1/3$. The probability that no one helps is $(2/3)^2 = 4/9$.
- With only one player, the probability that no one helps is $1/2$.
- $\Rightarrow$ the chance of help is better with two players, but only slightly so.
- This ranking can also be reversed. Example: Suppose $c \sim U[0, 1]$

Public good provision

- Cutoff equilibrium: Contribute iff $c_i \leq c_i^*$

$$1 - c_1^* = F(c_2^*) \text{ and } 1 - c_2^* = F(c_1^*)$$

- Suppose $F(x) = x/2$ for $x \in [0, 2]$ (i.e., uniform on $[0, 2]$)
  $\Rightarrow c_1^* = c_2^* = 2/3$
- With two players, each helps with probability $1/3$. The probability that no one helps is $(2/3)^2 = 4/9$.
- With only one player, the probability that no one helps is $1/2$.
- $\Rightarrow$ the chance of help is better with two players, but only slightly so.
- This ranking can also be reversed. Example: Suppose $c \sim U[0, 1]$

Auctions

Consider a seller who has a good to sell. The good is non-standard in that these goods are sold only rarely and so no "market" with a well defined market price exists. There are (at least) 2 potential buyers, A and B, and their willingness to pay is denoted v_A and v_B (this is their "type").

Applications:

- Rembrandt paintings
- (all sorts of stuff on Ebay)
- financial instruments (treasury bills, shares in IPO)
- procurement contracts (bidding by sellers!)

Auction formats

Second price auction (SPA):
Players simultaneously hand in bids b_i; the good goes to the bidder with the highest bid, and the price to be paid by the buyer is the second highest bid. If two bidders bid the same highest bid, the buyer is picked at random.

A bit odd: Why should the seller not charge the highest bid?
In fact, a perfectly reasonable way to organize an auction.

Ebay implements second price auctions. (?)

SPA is strategically equivalent to the *English auction* (an open, ascending bids auction)

Auction formats

First price auction (FPA):
The good goes to the highest bidder, and the price equals the winner's bid.

Many procurement auctions are organized as first price auctions (e.g. State wants to build some road; construction firms simultaneously submit bids, i.e., "We are willing to build the road for b_i"; lowest b_i wins.)

Equivalent to Dutch auction

Optimal auction design

What is the optimal way to bid in such a SPA?
Highest bid of the other bidders defines a price at which a bidder can buy the good.
Do you want to buy at this price?
Yes, if $v >$price.
No, if $v <$price
Reporting one's true WTP to the mechanism implements this optimal policy.
$\Rightarrow$ The SPA is a direct revelation mechanism
Difference to FPA:
Clearly, bidding one's true valuation is not optimal in the FPA, but rather (weakly) dominated by reducing ones bid below one's WTP.

FPA: Optimal bidding

Suppose that the players' valuations are iid draws from a uniform distribution on $[0, 1]$. Each player knows his own valuation, but not his opponent's valuation.

We are looking for a pure strategy BNE, i.e. two bidding strategies that map valuations (types) into bids (actions).
Suppose B's bidding strategy can be written $b_B = B(V_B)$, where $B(\cdot)$ is an increasing and differentiable function

Observation: Any reasonable (un-dominated) bidding strategy for A can be written $B(x(V_A))$:
(A behaves as if he had valuation x and followed B's bidding strategy; the optimal x depends, of course, on the realization of V_A) (why?)

FPA: Optimal bidding

The expected payoff for A is

$$F(x)[V_A - B(x)] + (1 - F(x)) \cdot 0$$

$F()$: cumulative distribution function; for the uniform distribution on $[0, 1]$, $F(x) = x$.
The optimal x satisfies

$$F'(x)[V_A - B(x)] - F(x)B'(x) = 0$$

First term: Benefit of bidding more. Bidding like type $x + dx$ increases the probability of winning by $F'(x)dx$. In this case, A gets a rent of $V_A - B(x)$.

Second term: Cost of bidding more. Type $x + dx$ bids $B'(x)dx$ more than type x, and with probability $F(x)$, A would win anyway, so a higher bid just means more money for the seller.

FPA: Optimal bidding

In a symmetric equilibrium, A behaves just like B, so we can substitute $x = V_A$ in the first order condition:

$$F'(V_A)[V_A - B(V_A)] - F(V_A)B'(V_A) = 0$$

This is a differential equation that we can solve for B:

$$F'(V_A)V_A = F'(V_A)B(V_A) + F(V_A)B'(V_A)$$

$$= \frac{d}{dV_A}(F(V_A)B(V_A))$$

Integrate both sides:

$$\int_0^V F'(V_A)V_A dV_A = \int_0^V \frac{d}{dV_A}(F(V_A)B(V_A))dV_A$$

$$= F(V)B(V) \Rightarrow B(V) = \frac{\int_0^V F'(V_A)V_A dV_A}{F(V)}$$

FPA: Optimal bidding

$$B(V) = \frac{\int_0^V F'(V_A) V_A dV_A}{F(V)}$$

Substitute $F(V) = V$, $F'(V) = 1$ for the uniform distribution on $[0; 1]$:

$$B(V) = V/2$$

This function in fact satisfies what we had to assume (B is increasing and differentiable), so we have found a BNE.

Revenue equivalence

What is the better form to organize an auction, from the point of view of the seller? We could calculate expected revenue for the FPA and the SPA directly:

$$\int_0^1 \int_0^1 max\left(\frac{v_1}{2}, \frac{v_2}{2}\right) dv1dv2 \gtreqless \int_0^1 \int_0^1 min\left(v_1, v_2\right) dv1dv2$$

The following way is both easier and more generalizable:

- Social welfare is equal to the rent of the seller (i.e., the price) and the rent of the buyer (his valuation minus the price).
- Social welfare is equal under both auctions (equal to the highest valuation, why?).

Hence, if we can prove that expected rents in the FPA and the SPA are equal for all types of buyers, then this implies that the expected price is also the same.

Revenue equivalence

Let $S_A(v, \hat{v})$ be the expected utility of type v if he behaves (bids) like type $\hat{v}$ in auction A.
In an equilibrium, all types behave like they should, so $S_A(v, v)$ is the equilibrium utility of type v, $EU_A(v)$.
How does the equilibrium utility change, if type goes $v \to v + dv$?
Envelope theorem argument (why?):

$$EU_A'(v) = \frac{\partial S_A(v, v)}{\partial v}$$

In both first and second price auctions with symmetric bidders,

$$\frac{\partial S_A(v, v)}{\partial v} = F(v)$$

Revenue equivalence

The expected rent of a buyer with valuation v in a SPA is

$$\int_0^v (v-x)f(x)dx = (v-x)F(x)|_{x=0}^{x=v} - \int_0^v F(x)(-1)dx$$

$$\int_0^v F(x)dx = EU_{SPA}(v)$$

In a FPA, the following must hold because of the envelope theorem:

$$\frac{dEU_{FPA}(v)}{dv} = F(v)$$

Integrating (and using $EU_{FPA}(0) = 0$) yields

$$EU_{FPA}(v) = \int_0^v F(x)dx$$

Revenue equivalence

The equality of revenue in FPA and SPA is called *Revenue Equivalence Theorem* and is valid more generally.

Crucial requirements for the RET to hold between two auctions

- Social welfare is the same $\Leftarrow$ the good is always (for all bidder type combinations) allocated to the same buyer type. In particular, highest type always gets the good in both auctions is sufficient for this point to be satisfied.
  NOT satisfied if different players have different bidding functions so that sometimes player 1 wins the auction even though player 2 has the higher valuation, or vice versa.
- Each type gets the same expected utility under both auctions. Satisfied if the good always goes to the highest valuation bidder because of the envelope argument.

Revenue equivalence in our example

Revenue in a second price auction:

$$\int_0^1 mg(m)dm$$

where $m = \min(v_1, v_2)$. What is $g(m)$?

$$G(m) = 1 - (1 - F(x))^2 \Rightarrow g(m) = 2(1 - F(m))f(m) = 2(1 - m)$$

$$\int_0^1 m2(1 - m)dm = \left[m^2 - \frac{2}{3}m^3\right]_{m=0}^{m=1} = \frac{1}{3}$$

Revenue equivalence in our example

Revenue in a first price auction:

$$\frac{1}{2}E(\max(v_1, v_2))$$

Let $k = \max(v_1, v_2)$. What is $h(k)$?

$$H(k) = F(k)^2 \Rightarrow h(k) = 2F(k)f(k) = 2k$$

$$\int_0^1 k2kdk = \left[\frac{2}{3}k^3\right]_{k=0}^{k=1} = \frac{2}{3}$$

Thus, revenue is $\frac{1}{2}\frac{2}{3} = \frac{1}{3}$, same as in first price auction.
Note: For a given pair of valuation types, *realized* revenue may be different in the two auction types.

Efficient provision of an indivisible public good

Problem: How could the state find out how much of the public good to supply, if individual demand functions are unknown (for the state; of course, people know their utility).
Indivisible good, costs 1 (if provided)
A: $v_A \in (0,1)$
B: $v_B \in (0,1)$
Efficiency: The good should be provided if and only if $v_A + v_B > 1$. Can the efficient allocation be implemented *even if the state does not know the individuals' WTP in the beginning?*

Clarke–Groves Mechanism

Mechanism

1. Both people announce m_A and m_B as their willingness to pay (they can, of course, lie)
2. If $m_A + m_B > 1$, the good is provided and A pays $(1 - m_B)$, B pays $(1 - m_A)$.
3. If $m_A + m_B < 1$, the good is not provided and no payments are made.

Observation: The report m_A affects A's payoff only if it changes whether the good is provided; the price A has to pay (if the good is provided) is independent of m_A and depends only on B's report!

Truthful revelation

Suppose A knew B's report, and $v_A + m_B > 1$. Then announcing $m_A = v_A$ is optimal for A (why?).

Now suppose $v_A + m_B < 1$. Then announcing $m_A = v_A$ is again optimal for A (why?).

This is true for every value of m_B: Announcing $m_A = v_A$ is a (weakly) dominant strategy for A! In particular, this is completely independent of whether B told the truth.

The same argument holds for B. Under this mechanism, both people announce the truth and the efficient solution can be implemented.

Intuition

$1 - m_B$: Net social cost of the public good for A (if B told the truth).
Do you want to buy the PG if you have to pay these net social cost?
Yes, if $v_A > 1 - m_B$
No, if $v_A < 1 - m_B$
Reporting $m_A = v_A$ to the CG-mechanism implements exactly this policy.

Who pays?

Sum of the payments by A and B: $1 - m_B + 1 - m_A = 2 - (m_A + m_B)$.
Whenever the PG is provided, $(m_A + m_B) > 1$, so payments by A and B are never sufficient to cover the cost of the PG ("no budget balance").
A third party has to put in some money.
However: One could charge from both people an additional lump sum payment (i.e., the same amount, whether or not the good is provided) to offset this.

War of attrition with unknown valuations

- Two players with valuations θ_1, θ_2; iid drawn from cdf $P/$density p.
- Players choose stopping times $s_1(\theta_1)$, $s_2(\theta_2)$

$$u_1(s_1, s_2, \theta_1) = \begin{cases} -s_1 & \text{if } s_2 > s_1 \\ \theta_1 - s_2 & \text{if } s_1 > s_2 \end{cases}$$

- Look for a symmetric equilibrium with increasing strategies

War of attrition with unknown valuations

Suppose P2 plays according to increasing function $S(\theta_2)$.
P1's choice can be thought of as choosing a marginal type t_0 to beat.
Expected payoff

$$\int_0^{t_0} p(t)[\theta_1 - S(t)]dt - [1 - P(t_0)]S(t_0)$$

A marginally more aggressive choice $(t_0 + dt)$ changes payoff by

$$p(t_0)[\theta_1 - S(t_0)]dt + p(t_0)S(t_0)dt - [1 - P(t_0)]S'(t_0)dt$$

In an optimum, this is zero

$$\Rightarrow p(t_0)\theta_1 = [1 - P(t_0)]S'(t_0)$$

War of attrition with unknown valuations

Suppose P2 plays according to increasing function $S(\theta_2)$.
P1's choice can be thought of as choosing a marginal type t_0 to beat.
Expected payoff

$$\int_0^{t_0} p(t)[\theta_1 - S(t)]dt - [1 - P(t_0)]S(t_0)$$

A marginally more aggressive choice $(t_0 + dt)$ changes payoff by

$$p(t_0)[\theta_1 - S(t_0)]dt + p(t_0)S(t_0)dt - [1 - P(t_0)]S'(t_0)dt$$

In an optimum, this is zero

$$\Rightarrow p(t_0)\theta_1 = [1 - P(t_0)]S'(t_0)$$

War of attrition with unknown valuations

$$p(t_0)\theta_1 = [1 - P(t_0)]S'(t_0)$$

In a symmetric equilibrium, each player plays the same strategy.
$\Rightarrow$ Each type chooses to beat exactly all types of the other player that are below him, i.e., $t_0(\theta) = \theta$

$$S'(\theta) = \frac{p(\theta)}{[1 - P(\theta)]}\theta$$

War of attrition with unknown valuations

$$S'(\theta) = \frac{p(\theta)\theta}{[1-P(\theta)]} \Rightarrow S(\theta) = \int_0^\theta \frac{p(x)x}{[1-P(x)]}dx$$

- If θ is exponentially distributed: $p(\theta) = e^{-\theta}$, $P(\theta) = 1 - e^{-\theta}$

$$S'(\theta) = \frac{e^{-\theta}\theta}{e^{-\theta}} = \theta \Rightarrow S(\theta) = A + \frac{1}{2}\theta^2$$

  Boundary condition: $S(0) = 0 \Rightarrow A = 0 \Rightarrow S(\theta) = \frac{1}{2}\theta^2$

- If θ is uniform on $[0,1]$

$$S'(\theta) = \frac{\theta}{1-\theta} \Rightarrow S(\theta) = -\theta - \ln(1-\theta) = \ln\left(\frac{1}{1-\theta}\right) - \theta$$

War of attrition with unknown valuations

$$S'(\theta) = \frac{p(\theta)\theta}{[1 - P(\theta)]} \Rightarrow S(\theta) = \int_0^\theta \frac{p(x)x}{[1 - P(x)]}dx$$

- If θ is exponentially distributed: $p(\theta) = e^{-\theta}$, $P(\theta) = 1 - e^{-\theta}$

$$S'(\theta) = \frac{e^{-\theta}\theta}{e^{-\theta}} = \theta \Rightarrow S(\theta) = A + \frac{1}{2}\theta^2$$

  Boundary condition: $S(0) = 0 \Rightarrow A = 0 \Rightarrow S(\theta) = \frac{1}{2}\theta^2$

- If θ is uniform on $[0, 1]$

$$S'(\theta) = \frac{\theta}{1 - \theta} \Rightarrow S(\theta) = -\theta - \ln(1 - \theta) = \ln\left(\frac{1}{1 - \theta}\right) - \theta$$

War of attrition with unknown valuations – interpretation

- War of attrition with known valuations, symmetric equilibrium
  - Mixed strategies: Need to randomize in order to keep the other player indifferent between different stopping times
  - Players' ex-ante expected surplus: Zero
- War of attrition with unknown valuations, symmetric equilibrium
  - Pure strategies played by each type: No need to randomize because the other player's ignorance of one's type generates enough uncertainty
  - Players' ex-ante expected surplus: Positive for all types $\theta > 0$.

War of attrition with unknown valuations – interpretation

- War of attrition with known valuations, symmetric equilibrium
  - Mixed strategies: Need to randomize in order to keep the other player indifferent between different stopping times
  - Players' ex-ante expected surplus: Zero
- War of attrition with unknown valuations, symmetric equilibrium
  - Pure strategies played by each type: No need to randomize because the other player's ignorance of one's type generates enough uncertainty
  - Players' ex-ante expected surplus: Positive for all types $\theta > 0$.

War of attrition with unknown valuations – ex-ante utility

Is it true in general that all types $\theta > 0$ in the war of attrition with unknown valuation have a positive ex-ante utility?
Ex-ante utility of type θ who behaves likes type θ':

$$V(\theta, \theta') = \int_0^{\theta'} p(t)[\theta - S(t)]dt - [1 - P(\theta')]S(\theta')$$

In equilibrium, the optimal θ' is θ. $U(\theta) = V(\theta, \theta)$
Change in expected utility as true type θ changes:

$$\frac{dU}{d\theta} = \frac{\partial V}{\partial \theta} + \frac{\partial V}{\partial \theta'} \frac{d\theta'}{d\theta}$$

Envelope theorem:

$$\frac{dU}{d\theta} = \frac{\partial V}{\partial \theta}$$

$$\Rightarrow \frac{dU}{d\theta} = \int_0^{\theta} p(t)dt = P(\theta) \Rightarrow U(\theta) = \int_0^{\theta} P(t)dt$$

War of attrition with unknown valuations – ex-ante utility

Is it true in general that all types $\theta > 0$ in the war of attrition with unknown valuation have a positive ex-ante utility?
Ex-ante utility of type θ who behaves likes type θ':

$$V(\theta, \theta') = \int_0^{\theta'} p(t)[\theta - S(t)]dt - [1 - P(\theta')]S(\theta')$$

In equilibrium, the optimal θ' is θ. $U(\theta) = V(\theta, \theta)$
Change in expected utility as true type θ changes:

$$\frac{dU}{d\theta} = \frac{\partial V}{\partial \theta} + \frac{\partial V}{\partial \theta'} \frac{d\theta'}{d\theta}$$

Envelope theorem:

$$\frac{dU}{d\theta} = \frac{\partial V}{\partial \theta}$$

$$\Rightarrow \frac{dU}{d\theta} = \int_0^{\theta} p(t)dt = P(\theta) \Rightarrow U(\theta) = \int_0^{\theta} P(t)dt$$

War of attrition with unknown valuations – ex-ante utility

Is it true in general that all types $\theta > 0$ in the war of attrition with unknown valuation have a positive ex-ante utility?
Ex-ante utility of type θ who behaves likes type θ':

$$V(\theta, \theta') = \int_0^{\theta'} p(t)[\theta - S(t)]dt - [1 - P(\theta')]S(\theta')$$

In equilibrium, the optimal θ' is θ. $U(\theta) = V(\theta, \theta)$
Change in expected utility as true type θ changes:

$$\frac{dU}{d\theta} = \frac{\partial V}{\partial \theta} + \frac{\partial V}{\partial \theta'}\frac{d\theta'}{d\theta}$$

Envelope theorem:

$$\frac{dU}{d\theta} = \frac{\partial V}{\partial \theta}$$

$$\Rightarrow \frac{dU}{d\theta} = \int_0^{\theta} p(t)dt = P(\theta) \Rightarrow U(\theta) = \int_0^{\theta} P(t)dt$$