

6.207/14.15: Networks

Lectures 19-21: Incomplete Information: Bayesian Nash Equilibria, Auctions and Introduction to Social Learning

Daron Acemoglu and Asu Ozdaglar
MIT

November 23, 25 and 30, 2009

Outline

- Incomplete information.
 - Bayes rule and Bayesian inference.
 - Bayesian Nash Equilibria.
 - Auctions.
 - Extensive form games of incomplete information.
 - Perfect Bayesian (Nash) Equilibria.
 - Introduction to social learning and herding.
-
- **Reading:**
 - Osborne, Chapter 9.
 - EK, Chapter 16.

Incomplete Information

- In many game theoretic situations, one agent is unsure about the preferences or intentions of others.
- Incomplete information introduces additional strategic interactions and also raises questions related to “learning”.
- Examples:
 - *Bargaining* (how much the other party is willing to pay is generally unknown to you)
 - *Auctions* (how much should you bid for an object that you want, knowing that others will also compete against you?)
 - *Market competition* (firms generally do not know the exact cost of their competitors)
 - *Signaling games* (how should you infer the information of others from the signals they send)
 - *Social learning* (how can you leverage the decisions of others in order to make better decisions)

Example: Incomplete Information Battle of the Sexes

- Recall the battle of the sexes game, which was a complete information “coordination” game.
- Both parties want to meet, but they have different preferences on “Ballet” and “Football”.

	B	F
B	(2, 1)	(0, 0)
F	(0, 0)	(1, 2)

- In this game there are two pure strategy equilibria (one of them better for player 1 and the other one better for player 2), and a mixed strategy equilibrium.
- Now imagine that player 1 does not know whether player 2 wishes to meet or wishes to avoid player 1. Therefore, this is a situation of **incomplete information**—also sometimes called *asymmetric information*.

Example (continued)

- We represent this by thinking of player 2 having two different **types**, one type that wishes to meet player 1 and the other wishes to avoid him.
- More explicitly, suppose that these two types have probability $1/2$ each. Then the game takes the form one of the following two with probability $1/2$.

	B	F
B	(2, 1)	(0, 0)
F	(0, 0)	(1, 2)

	B	F
B	(2, 0)	(0, 2)
F	(0, 1)	(1, 0)

- Crucially, player 2 knows which game it is (she knows the **state of the world**), but player 1 does not.
- What are strategies in this game?

Example (continued)

- Most importantly, from player 1's point of view, player 2 has two possible types (or equivalently, the world has two possible states each with $1/2$ probability and only player 2 knows the actual state).
- How do we reason about equilibria here?
- **Idea:** Use Nash Equilibrium concept in an expanded game, where each different type of player 2 has a different strategy
 - Or equivalently, form conjectures about other player's actions in each state and act optimally given these conjectures.

Example (continued)

- Let us consider the following strategy profile $(B, (B, F))$, which means that player 1 will play B , and while in state 1, player 2 will also play B (when she wants to meet player 1) and in state 2, player 2 will play F (when she wants to avoid player 1).
- Clearly, given the play of B by player 1, the strategy of player 2 is a best response.
- Let us now check that player 2 is also playing a best response.
- Since both states are equally likely, the expected payoff of player 2 is

$$\mathbb{E}[B, (B, F)] = \frac{1}{2} \times 2 + \frac{1}{2} \times 0 = 1.$$

- If, instead, he deviates and plays F , his expected payoff is

$$\mathbb{E}[F, (B, F)] = \frac{1}{2} \times 0 + \frac{1}{2} \times 1 = \frac{1}{2}.$$

- Therefore, the strategy profile $(B, (B, F))$ is a (Bayesian) Nash equilibrium.

Example (continued)

- Interestingly, meeting at Football, which is the preferable outcome for player 2 is no longer a Nash equilibrium. Why not?
- Suppose that the two players will meet at Football when they want to meet. Then the relevant strategy profile is $(F, (F, B))$ and

$$\mathbb{E}[F, (F, B)] = \frac{1}{2} \times 1 + \frac{1}{2} \times 0 = \frac{1}{2}.$$

- If, instead, player 1 deviates and plays B , his expected payoff is

$$\mathbb{E}[B, (F, B)] = \frac{1}{2} \times 0 + \frac{1}{2} \times 2 = 1.$$

- Therefore, the strategy profile $(F, (F, B))$ is not a (Bayesian) Nash equilibrium.

Bayesian Games

- More formally, we can define **Bayesian games**, or “incomplete information games” as follows.

Definition

A *Bayesian game* consists of

- A set of players \mathcal{I} ;
 - A set of actions (pure strategies) for each player i : S_i ;
 - A set of types for each player i : $\theta_i \in \Theta_i$;
 - A payoff function for each player i : $u_i(s_1, \dots, s_I, \theta_1, \dots, \theta_I)$;
 - A (joint) probability distribution $p(\theta_1, \dots, \theta_I)$ over types (or $P(\theta_1, \dots, \theta_I)$ when types are not finite).
-
- More generally, one could also allow for a signal for each player, so that the signal is correlated with the underlying type vector.

Bayesian Games (continued)

- Importantly, throughout in Bayesian games, the strategy spaces, the payoff functions, possible types, and the prior probability distribution are assumed to be **common knowledge**.
 - Very strong assumption.
 - But very convenient, because any private information is included in the description of the type and others can form beliefs about this type and each player understands others' beliefs about his or her own type, and so on, and so on.

Definition

A (pure) strategy for player i is a map $s_i : \Theta_i \rightarrow S_i$ prescribing an action for each possible type of player i .

Bayesian Games (continued)

- Recall that player types are drawn from some prior probability distribution $p(\theta_1, \dots, \theta_I)$.
- Given $p(\theta_1, \dots, \theta_I)$ we can compute the conditional distribution $p(\theta_{-i} | \theta_i)$ using **Bayes rule**.
 - Hence the label “Bayesian games”.
 - Equivalently, when types are not finite, we can compute the conditional distribution $P(\theta_{-i} | \theta_i)$ given $P(\theta_1, \dots, \theta_I)$.
- Player i knows her own type and evaluates her expected payoffs according to the **conditional distribution** $p(\theta_{-i} | \theta_i)$, where $\theta_{-i} = (\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_I)$.

Bayesian Games (continued)

- Since the payoff functions, possible types, and the prior probability distribution are common knowledge, we can compute expected payoffs of player i of type θ_i as

$$U(s'_i, s_{-i}(\cdot), \theta_i) = \sum_{\theta_{-i}} p(\theta_{-i} | \theta_i) u_i(s'_i, s_{-i}(\theta_{-i}), \theta_i, \theta_{-i})$$

when types are finite

$$= \int u_i(s'_i, s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) P(d\theta_{-i} | \theta_i)$$

when types are not finite.

Bayes Rule

- Quick recap on **Bayes rule**.
- Let $\Pr(A)$ and $\Pr(B)$ denote, respectively, the probabilities of events A and B ; $\Pr(B | A)$ and $\Pr(A \cap B)$, conditional probabilities (one event conditional on the other one), and $\Pr(A \cap B)$ be the probability that both events happen (are true) simultaneously.
- Then Bayes rule states that

$$\Pr(A | B) = \frac{\Pr(A \cap B)}{\Pr(B)}. \quad (\text{Bayes I})$$

- Intuitively, this is the probability that A is true given that B is true.
- When the two events are **independent**, then $\Pr(B \cap A) = \Pr(A) \times \Pr(B)$, and in this case, $\Pr(A | B) = \Pr(A)$.

Bayes Rule (continued)

- Bayes rule also enables us to express conditional probabilities in terms of each other. Recalling that the probability that A is not true is $1 - \Pr(A)$, and denoting the event that A is not true by A^c (for A “complement”), so that $\Pr(A^c) = 1 - \Pr(A)$, we also have

$$\Pr(A | B) = \frac{\Pr(A) \times \Pr(B | A)}{\Pr(A) \times \Pr(B | A) + \Pr(A^c) \times \Pr(B | A^c)}. \quad (\text{Bayes II})$$

- This equation directly follows from (Bayes I) by noting that

$$\Pr(B) = \Pr(A) \times \Pr(B | A) + \Pr(A^c) \times \Pr(B | A^c),$$

and again from (Bayes I)

$$\Pr(A \cap B) = \Pr(A) \times \Pr(B | A).$$

Bayes Rule

- More generally, for a finite or countable partition $\{A_j\}_{j=1}^n$ of the event space, for each j

$$\Pr(A_j | B) = \frac{\Pr(A_j) \times \Pr(B | A_j)}{\sum_{i=1}^n \Pr(A_i) \times \Pr(B | A_i)}.$$

- For continuous probability distributions, the same equation is true with densities

$$f(A' | B) = \frac{f(A') \times f(B | A')}{\int f(B | A) \times f(A) dA}.$$

Bayesian Nash Equilibria

Definition

(Bayesian Nash Equilibrium) The strategy profile $s(\cdot)$ is a (pure strategy) Bayesian Nash equilibrium if for all $i \in \mathcal{I}$ and for all $\theta_i \in \Theta_i$, we have that

$$s_i(\theta_i) \in \arg \max_{s'_i \in S_i} \sum_{\theta_{-i}} p(\theta_{-i} | \theta_i) u_i(s'_i, s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}),$$

or in the non-finite case,

$$s_i(\theta_i) \in \arg \max_{s'_i \in S_i} \int u_i(s'_i, s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) P(d\theta_{-i} | \theta_i) .$$

- Hence a Bayesian Nash equilibrium is a Nash equilibrium of the “expanded game” in which each player i 's space of pure strategies is the set of maps from Θ_i to S_i .

Existence of Bayesian Nash Equilibria

Theorem

Consider a finite incomplete information (Bayesian) game. Then a mixed strategy Bayesian Nash equilibrium exists.

Theorem

Consider a Bayesian game with continuous strategy spaces and continuous types. If strategy sets and type sets are compact, payoff functions are continuous and concave in own strategies, then a pure strategy Bayesian Nash equilibrium exists.

- The ideas underlying these theorems and proofs are identical to those for the existence of equilibria in (complete information) strategic form games.

Example: Incomplete Information Cournot

- Suppose that two firms both produce at constant marginal cost.
- Demand is given by $P(Q)$ as in the usual Cournot game.
- Firm 1 has marginal cost equal to C (and this is common knowledge).
- Firm 2's marginal cost is private information. It is equal to C_L with probability θ and to C_H with probability $(1 - \theta)$, where $C_L < C_H$.
- This game has 2 players, 2 states (L and H) and the possible actions of each player are $q_i \in [0, \infty)$, but firm 2 has two possible types.
- The payoff functions of the players, after quantity choices are made, are given by

$$\begin{aligned}
 u_1((q_1, q_2), t) &= q_1(P(q_1 + q_2) - C) \\
 u_2((q_1, q_2), t) &= q_2(P(q_1 + q_2) - C_t),
 \end{aligned}$$

where $t \in \{L, H\}$ is the type of player 2.

Example (continued)

- A strategy profile can be represented as (q_1^*, q_L^*, q_H^*) [or equivalently as $(q_1^*, q_2^*(\theta_2))$], where q_L^* and q_H^* denote the actions of player 2 as a function of its possible types.
- We now characterize the Bayesian Nash equilibria of this game by computing the best response functions (correspondences) and finding their intersection.
- There are now three best response functions and they are given by

$$B_1(q_L, q_H) = \arg \max_{q_1 \geq 0} \{ \theta(P(q_1 + q_L) - C)q_1 \\ + (1 - \theta)(P(q_1 + q_H) - C)q_1 \}$$

$$B_L(q_1) = \arg \max_{q_L \geq 0} \{ (P(q_1 + q_L) - C_L)q_L \}$$

$$B_H(q_1) = \arg \max_{q_H \geq 0} \{ (P(q_1 + q_H) - C_H)q_H \}.$$

Example (continued)

- The Bayesian Nash equilibria of this game are vectors (q_1^*, q_L^*, q_H^*) such that

$$B_1(q_L^*, q_H^*) = q_1^*, \quad B_L(q_1^*) = q_L^*, \quad B_H(q_1^*) = q_H^*.$$

- To simplify the algebra, let us assume that $P(Q) = \alpha - Q$, $Q \leq \alpha$. Then we can compute:

$$q_1^* = \frac{1}{3}(\alpha - 2C + \theta C_L + (1 - \theta)C_H)$$

$$q_L^* = \frac{1}{3}(\alpha - 2C_L + C) - \frac{1}{6}(1 - \theta)(C_H - C_L)$$

$$q_H^* = \frac{1}{3}(\alpha - 2C_H + C) + \frac{1}{6}\theta(C_H - C_L).$$

Example (continued)

- Note that $q_L^* > q_H^*$. This reflects the fact that with lower marginal cost, the firm will produce more.
- However, incomplete information also affects firm 2's output choice.
- Recall that, given this demand function, if both firms knew each other's marginal cost, then the unique Nash equilibrium involves output of firm i given by

$$\frac{1}{3}(\alpha - 2C_i + C_j).$$

- With incomplete information, firm 2's output is less if its cost is C_H and more if its cost is C_L . If firm 1 knew firm 2's cost is high, then it would produce more. However, its lack of information about the cost of firm 2 leads firm 1 to produce a relatively moderate level of output, which then allows firm 2 to be more "aggressive".
- Hence, in this case, firm 2 benefits from the lack of information of firm 1 and it produces more than if 1 knew his actual cost.

Auctions

- A major application of Bayesian games is to **auctions**, which are historically and currently common method of allocating scarce goods across individuals with different valuations for these goods.
- This corresponds to a situation of incomplete information because the valuations of different potential buyers are unknown.
- For example, if we were to announce a **mechanism**, which involves giving a particular good (for example a seat in a sports game) for free to the individual with the highest valuation, this would create an incentive for all individuals to overstate their valuations.
- In general, auctions are designed by **profit-maximizing** entities, which would like to sell the goods to raise the highest possible revenue.

Auctions (continued)

- Different types of auctions and terminology:
 - *English auctions*: ascending sequential bids.
 - *First price sealed bid auctions*: similar to English auctions, but in the form of a strategic form game; all players make a single simultaneous bid and the highest one obtains the object and pays its bid.
 - *Second price sealed bid auctions*: similar to first price auctions, except that the winner pays the second highest bid.
 - *Dutch auctions*: *descending sequential auctions*; the auction stops when an individual announces that she wants to buy at that price. Otherwise the price is reduced sequentially until somebody stops the auction.
 - *Combinatorial auctions*: when more than one item is auctioned, and agents value combinations of items.
 - *Private value auctions*: valuation of each agent is independent of others' valuations;
 - *Common value auctions*: the object has a potentially common value, and each individual's signal is imperfectly correlated with this common value.

Modeling Auctions

- Model of auction:
 - a valuation structure for the bidders (i.e., private values for the case of private-value auctions),
 - a probability distribution over the valuations available to the bidders.
- Let us focus on first and second price sealed bid auctions, where bids are submitted simultaneously.
- Each of these two auction formats defines a static game of incomplete information (Bayesian game) among the bidders.
- We determine Bayesian Nash equilibria in these games and compare the equilibrium bidding behavior.

Modeling Auctions (continued)

- More explicitly, suppose that there is a single object for sale and N potential buyers are bidding for the object.
- Bidder i assigns a value v_i to the object, i.e., a utility

$$v_i - b_i,$$

when he pays b_i for the object. He knows v_i . This implies that we have a **private value** auction (v_i is his “private information” and “private value”).

- Suppose also that each v_i is *independently and identically distributed* on the interval $[0, \bar{v}]$ with cumulative distribution function F , with continuous density f and full support on $[0, \bar{v}]$.
- Bidder i knows the realization of its value v_i (or realization v_i of the random variable V_i , though we will not use this latter notation) and that other bidders' values are independently distributed according to F , i.e., all components of the model except the realized values are “common knowledge”.

Modeling Auctions (continued)

- Bidders are *risk neutral*, i.e., they are interested in maximizing their expected profits.
- This model defines a Bayesian game of incomplete information, where the types of the players (bidders) are their valuations, and a pure strategy for a bidder is a map

$$\beta_i : [0, \bar{v}] \rightarrow \mathbb{R}_+.$$

- We will characterize the symmetric equilibrium strategies in the first and second price auctions.
- Once we characterize these equilibria, then we can also investigate which auction format yields a higher expected revenue to the seller at the symmetric equilibrium.

Second Price Auctions


- Second price auctions will have the structure very similar to a complete information auction discussed earlier in the lectures.
- There we saw that each player had a weakly dominant strategy. This will be true in the incomplete information version of the game and will greatly simplify the analysis.
- In the auction, each bidder submits a sealed bid of b_i , and given the vector of bids $b = (b_i, b_{-i})$ and evaluation v_i of player i , its payoff is

$$U_i((b_i, b_{-i}), v_i) = \begin{cases} v_i - \max_{j \neq i} b_j & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{if } b_i \leq \max_{j \neq i} b_j. \end{cases}$$

- Let us also assume that if there is a tie, i.e., $b_i = \max_{j \neq i} b_j$, the object goes to each winning bidder with equal probability.
- With the reasoning similar to its counterpart with complete information, in a second-price auction, it is a weakly dominant strategy to bid *truthfully*, i.e., according to $\beta^i(v) = v$.

Second Price Auctions (continued)

- This can be established with the same graphical argument as the one we had for the complete information case.
- The first graph shows the payoff for bidding one's valuation, the second graph the payoff from bidding a lower amount, and the third the payoff from bidding higher amount.
- In all cases B^* denotes the highest bid excluding this player.



Second Price Auctions (continued)

- Moreover, now there are no other optimal strategies and thus the (Bayesian) equilibrium will be unique, since the valuation of other players are not known.
- Therefore, we have established:

Proposition

In the second price auction, there exists a unique Bayesian Nash equilibrium which involves

$$\beta^H(v) = v.$$

- It is straightforward to generalize the exact same result to the situation in which bidders' values are not independently and identically distributed. As in the complete information case, bidding truthfully remains a weakly dominant strategy.
- The assumption of private values is important (i.e., the valuations are known at the time of the bidding).

Second Price Auctions (continued)

- Let us next determine the expected payment by a bidder with value v , and for this, let us focus on the case in which valuations are independent and identically distributed
- Fix bidder 1 and define the random variable y_1 as the highest value among the remaining $N - 1$ bidders, i.e.,

$$y_1 = \max\{v_2, \dots, v_N\}.$$

- Let G denote the cumulative distribution function of y_1 .
- Clearly,

$$G(y) = F(v)^{N-1} \text{ for any } v \in [0, \bar{v}]$$

Second Price Auctions (continued)

- In a second price auction, the expected payment by a bidder with value v is given by

$$\begin{aligned}m''(v) &= \Pr(v \text{ wins}) \times \mathbb{E}[\text{second highest bid} \mid v \text{ is the highest bid}] \\ &= \Pr(y_1 \leq v) \times \mathbb{E}[y_1 \mid y_1 \leq v] \\ &= G(v) \times \mathbb{E}[y_1 \mid y_1 \leq v].\end{aligned}$$

Payment II

- Note that here and in what follows, we can use strict or weak inequalities given that the relevant random variables have continuous distributions. In other words, we have

$$\Pr(y_1 \leq v) = \Pr(y_1 < v).$$

Example: Uniform Distributions

- Suppose that there are two bidders with valuations, v_1 and v_2 , distributed uniformly over $[0, 1]$.
- Then $G(v_1) = v_1$, and

$$\mathbb{E}[y_1 \mid y_1 \leq v_1] = \mathbb{E}[v_2 \mid v_2 \leq v_1] = \frac{v_1}{2}.$$

- Thus

$$m''(v) = \frac{v^2}{2}.$$

- If, instead, there are N bidders with valuations distributed over $[0, 1]$,

$$G(v_1) = (v_1)^{N-1}$$

$$\mathbb{E}[y_1 \mid y_1 \leq v_1] = \frac{N-1}{N} v_1,$$

and thus

$$m''(v) = \frac{N-1}{N} v^N.$$

First Price Auctions

- In a first price auction, each bidder submits a sealed bid of b_i , and given these bids, the payoffs are given by

$$U_i((b_i, b_{-i}), v_i) = \begin{cases} v_i - b_i & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{if } b_i > \max_{j \neq i} b_j. \end{cases}$$

- Tie-breaking is similar to before.
- In a first price auction, the equilibrium behavior is more complicated than in a second-price auction.
- Clearly, bidding truthfully is not optimal (why not?).
 - Trade-off between higher bids and lower bids.
- So we have to work out more complicated strategies.

First Price Auctions (continued)

- Approach: look for a **symmetric (continuous and differentiable) equilibrium**.
- Suppose that bidders $j \neq 1$ follow the symmetric increasing and differentiable equilibrium strategy $\beta^j = \beta$, where

$$\beta_j : [0, \bar{v}] \rightarrow \mathbb{R}_+.$$

- We also assume, without loss of any generality, that β is increasing.
- We will then allow player 1 to use strategy β_1 and then characterize β such that when all other players play β , β_1 is a best response for player 1. Since player 1 was arbitrary, this will complete the characterization of equilibrium.
- Suppose that bidder 1 value is v_1 and he bids the amount b (i.e., $\beta(v_1) = b$).

First Price Auctions (continued)

- First, note that a bidder with value 0 would never submit a positive bid, so

$$\beta(0) = 0.$$

- Next, note that bidder 1 wins the auction whenever $\max_{i \neq 1} \beta(v_i) < b$.
- Since $\beta(\cdot)$ is increasing, we have

$$\max_{i \neq 1} \beta(v_i) = \beta(\max_{i \neq 1} v_i) = \beta(y_1),$$

where recall that

$$y_1 = \max\{v_2, \dots, v_N\}.$$

- This implies that bidder 1 wins whenever $y_1 < \beta^{-1}(b)$.

First Price Auctions (continued)

- Consequently, we can find an optimal bid of bidder 1, with valuation $v_1 = v$, as the solution to the maximization problem

$$\max_{b \geq 0} G(\beta^{-1}(b))(v - b).$$

- The first-order (necessary) conditions imply

$$\frac{g(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))}(v - b) - G(\beta^{-1}(b)) = 0, \quad (*)$$

where $g = G'$ is the probability density function of the random variable y_1 . [Recall that the derivative of $\beta^{-1}(b)$ is $1/\beta'(\beta^{-1}(b))$].

- This is a first-order differential equation, which we can in general solve.

First Price Auctions (continued)

- More explicitly, a symmetric equilibrium, we have $\beta(v) = b$, and therefore (*) yields

$$G(v)\beta'(v) + g(v)\beta(v) = vg(v).$$

- Equivalently, the first-order differential equation is

$$\frac{d}{dv} \left(G(v)\beta(v) \right) = vg(v),$$

with boundary condition $\beta(0) = 0$.

- We can rewrite this as the following optimal bidding strategy

$$\beta(v) = \frac{1}{G(v)} \int_0^v yg(y)dy = \mathbb{E}[y_1 \mid y_1 < v].$$

- Note, however, that we skipped one additional step in the argument: the first-order conditions are only necessary, so one needs to show sufficiency to complete the proof that the strategy $\beta(v) = \mathbb{E}[y_1 \mid y_1 < v]$ is optimal.

First Price Auctions (continued)

- This detail notwithstanding, we have:

Proposition

In the first price auction, there exists a unique symmetric equilibrium given by

$$\beta^I(v) = \mathbb{E}[y_1 \mid y_1 < v].$$

First Price Auctions: Payments and Revenues

- In general, expected payment of a bidder with value v in a first price auction is given by

$$\begin{aligned} m^I(v) &= \Pr(v \text{ wins}) \times \beta(v) \\ &= G(v) \times \mathbb{E}[y_1 \mid y_1 < v]. \end{aligned} \quad \text{Payment I}$$

- This can be directly compared to (Payment II), which was the payment in the second price auction ($m^{II}(v) = G(v) \times \mathbb{E}[y_1 \mid y_1 \leq v]$).
- This establishes the somewhat surprising results that $m^I(v) = m^{II}(v)$, i.e., both auction formats yield the same expected revenue to the seller.

First Price Auctions: Uniform Distribution

- As an illustration, assume that values are uniformly distributed over $[0, 1]$.
- Then, we can verify that

$$\beta^I(v) = \frac{N-1}{N}v.$$

- Moreover, since $G(v_1) = (v_1)^{N-1}$, we again have

$$m^I(v) = \frac{N-1}{N}v^N.$$

Revenue Equivalence

- In fact, the previous result is a simple case of a more general theorem.
- Consider any **standard** auction, in which buyers submit bids and the object is given to the bidder with the highest bid.
- Suppose that values are independent and identically distributed and that all bidders are risk neutral. Then, we have the following theorem:

Theorem

Any symmetric and increasing equilibria of any standard auction (such that the expected payment of a bidder with value 0 is 0) yields the same expected revenue to the seller.

Sketch Proof

- Consider a standard auction A and a symmetric equilibrium β of A .
- Let $m^A(v)$ denote the equilibrium expected payment in auction A by a bidder with value v .
- Assume that β is such that $\beta(0) = 0$.
- Consider a particular bidder, say bidder 1, and suppose that other bidders are following the equilibrium strategy β .
- Consider the expected payoff of bidder 1 with value v when he bids $b = \beta(z)$ instead of $\beta(v)$,

$$U^A(z, v) = G(z)v - m^A(z).$$

- Maximizing the preceding with respect to z yields

$$\frac{\partial}{\partial z} U^A(z, v) = g(z)v - \frac{d}{dz} m^A(z) = 0.$$

Sketch Proof (continued)

- An equilibrium will involve $z = v$ (Why?) Hence,

$$\frac{d}{dy}m^A(y) = g(y)y \quad \text{for all } y,$$

implying that

$$m^A(v) = \int_0^v yg(y)dy = G(v) \times \mathbb{E}[y_1 \mid y_1 < v], \quad (\text{General Payment})$$

establishing that the expected revenue of the seller is the same independent of the particular auction format.

- (General Payment), not surprisingly, has the same form as (Payment I) and (Payment II).

Common Value Auctions

- Common value auctions are more complicated, because each player has to infer the valuation of the other player (which is relevant for his own valuation) from the bid of the other player (or more generally from the fact that he has one).
- This generally leads to a phenomenon called **winner's curse**—conditional on winning each individual has a lower valuation than unconditionally.
- The analysis of common value auctions is typically more complicated. So we will just communicate the main ideas using an example.

Common Value Auctions: The Difficulty

- To illustrate the difficulties with common value auctions, suppose a situation in which two bidders are competing for an object that has either high or low quality, or equivalently, value $v \in \{0, \bar{v}\}$ to both of them (for example, they will both sell the object to some third party). The two outcomes are equally likely.
- They both receive a **signal** $s_i \in \{l, h\}$. Conditional on a low signal (for either player), $v = 0$ with probability 1. Conditional on a high signal, $v = \bar{v}$ with probability $p > 1/2$. [Why are we referring to this as a “signal”?]
- This game has no symmetric pure strategy equilibrium.
- Suppose that there was such an equilibrium, in which $b(0) = b_l$ and $b(\bar{v}) = b_h \geq b_l$. If player 1 has type \bar{v} and indeed bids b_h , he will obtain the object with probability 1 when player 2 bids b_l and with probability $1/2$ when player 2 bids b_h .
- Clearly we must have $b_l = 0$ (Why?).

Common Value Auctions: The Difficulty

- In the first case, it means that the other player has received a low signal, but if so, $v = 0$. In this case, player 1 would not like to pay anything positive for the good—this is the **winner's curse**.
- In the second case, it means that both players have received a high signal, so the object is high-quality with probability $1 - (1 - p)^2$. So in this case, it would be better to bid ε more and obtain the object with probability 1, unless $b_h = \left[1 - (1 - p)^2\right] \bar{v}$. But a bid of $b_h = \left[1 - (1 - p)^2\right] \bar{v}$ will, on average, lose money, since it will also win the object when the other bidder has received the low signal.

Common Value Auctions: A Simple Example

- If, instead, we introduce some degree of private values, then common value auctions become more tractable.
- Consider the following example. There are two players, each receiving a signal s_j . The value of the good to both of them is

$$v_i = \alpha s_i + \beta s_{-i},$$

where $\alpha \geq \beta \geq 0$. Private values are the special case where $\alpha = 1$ and $\beta = 0$.

- Suppose that both s_1 and s_2 are distributed uniformly over $[0, 1]$.

Second Price Auctions with Common Values

- Now consider a second price auction.
- Instead of truthful bidding, now the symmetric equilibrium is each player bidding

$$b_i(s_i) = (\alpha + \beta) s_i.$$

- Why?
- Given that the other player is using the same strategy, the probability that player i will win when he bids b is

$$\begin{aligned}\Pr(b_{-i} < b) &= \Pr((\alpha + \beta) s_{-i} < b) \\ &= \frac{b}{\alpha + \beta}.\end{aligned}$$

- The price he will pay is simply $b_{-i} = (\alpha + \beta) s_{-i}$ (since this is a second price auction).

Second Price Auctions with Common Values (continued)

- Conditional on the fact that $b_{-i} \leq b$, s_{-i} is still distributed uniformly (with the top truncated). In other words, it is distributed uniformly between 0 and b . Then the expected price is

$$\mathbb{E} \left[(\alpha + \beta) s_{-i} \mid s_{-i} < \frac{b}{\alpha + \beta} \right] = \frac{b}{2}.$$

- Next, let us compute the expected value of player $-i$'s signal conditional on player i winning. With the same reasoning, this is

$$\mathbb{E} \left[s_{-i} \mid s_{-i} < \frac{b}{\alpha + \beta} \right] = \frac{b}{2(\alpha + \beta)}.$$

Second Price Auctions with Common Values (continued)

- Therefore, the expected utility of bidding b_i for player i with signal s_i is:

$$\begin{aligned}U_i(b_i, s_i) &= \Pr[b_i \text{ wins}] \times \left(\alpha s_i + \beta \mathbb{E}[s_{-i} \mid b_i \text{ wins}] - \frac{b_i}{2} \right) \\ &= \frac{b_i}{\alpha + \beta} \left[\alpha s_i + \frac{\beta}{2} \frac{b_i}{\alpha + \beta} - \frac{b_i}{2} \right].\end{aligned}$$

- Maximizing this with respect to b_i (for given s_i) implies

$$b_i(s_i) = (\alpha + \beta) s_i,$$

establishing that this is the unique symmetric Bayesian Nash equilibrium of this common values auction.

First Price Auctions with Common Values

- We can also analyze the same game under an auction format corresponding to first price sealed bid auctions.
- In this case, with an analysis similar to that of the first price auctions with private values, we can establish that the unique symmetric Bayesian Nash equilibrium is for each player to bid

$$b_i^j(s_i) = \frac{1}{2} = (\alpha + \beta) s_i.$$

- It can be verified that expected revenues are again the same. This illustrates the general result that revenue equivalence principle continues to hold for, value auctions.

Incomplete Information in Extensive Form Games

- Many situations of incomplete information cannot be represented as static or strategic form games.
- Instead, we need to consider extensive form games with an explicit order of moves—or dynamic games.
- In this case, as mentioned earlier in the lectures, we use **information sets** to represent what each player knows at each stage of the game.
- Since these are dynamic games, we will also need to strengthen our Bayesian Nash equilibria to include the notion of **perfection**—as in subgame perfection.
- The relevant notion of equilibrium will be [Perfect Bayesian Equilibria](#), or Perfect Bayesian Nash Equilibria.

Example




Figure: Selten's Horse

Dynamic Games of Incomplete Information

Definition

A dynamic game of incomplete information consists of

- *A set of players \mathcal{I} ;*
- *A sequence of histories H^t at the t th stage of the game, each history assigned to one of the players (or to Nature);*
- *An information partition, which determines which of the histories assigned to a player are in the same information set.*
- *A set of (pure) strategies for each player i , S_i , which includes an action at each information set assigned to the player.*
- *A set of types for each player i : $\theta_i \in \Theta_i$;*
- *A payoff function for each player i : $u_i(s_1, \dots, s_I, \theta_1, \dots, \theta_I)$;*
- *A (joint) probability distribution $p(\theta_1, \dots, \theta_I)$ over types (or $P(\theta_1, \dots, \theta_I)$ when types are not finite).*

Strategies, Beliefs and Bayes Rule

- The most economical way of approaching these games is to first define a **belief system**, which determines a posterior for each agent over to set of nodes in an information set. Beliefs systems are often denoted by μ .
- In Selten's horse player 3 needs to have beliefs about whether when his information set is reached, he is at the left or the right node.
- A **strategy** can then be expressed as a mapping that determines the actions of the player is a function of his or her beliefs at the relevant information set.
- We say that a strategy is **sequentially rational** if, given beliefs, no player can improve his or her payoffs at any stage of the game.
- We say that a belief system is **consistent** if it is derived from equilibrium strategies using Bayes rule.

Strategies, Beliefs and Bayes Rule (continued)

- In Selten's horse, if the strategy of player 1 is D , then Bayes rule implies that $\mu_3(\text{left}) = 1$, since conditional on her information set being reached, player 3's assessment must be that this was because player 1 played D .
- Similarly, if the strategy of player 1 is D with probability p and the strategy of player 2 is d with probability q , then Bayes rule implies that

$$\mu_3(\text{left}) = \frac{p}{p + (1 - p)q}.$$

- What happens if $p = q = 0$? In this case, $\mu_3(\text{left})$ is given by $0/0$, and is thus *undefined*. Under the consistency requirement here, it can take any value. This implies, in particular, that information sets that are not reached along the equilibrium path will have **unrestricted** beliefs.

Perfect Bayesian Equilibria

Definition

A Perfect Bayesian Equilibrium in a dynamic game of incomplete information is a strategy profile s and belief system μ such that:

- *The strategy profile s is sequentially rational given μ .*
 - *The belief system μ is consistent given s .*
-
- Perfect Bayesian Equilibrium is a relatively weak equilibrium concept for dynamic games of incomplete information. It is often strengthened by restricting beliefs information sets that are not reached along the equilibrium path. We will return to this issue below.


Existence of Perfect Bayesian Equilibria

Theorem

Consider a finite dynamic game of incomplete information. Then a (possibly mixed) Perfect Bayesian Equilibrium exists.

- Once again, the idea of the proof is the same as those we have seen before.
- Recall the general proof of existence for dynamic games of imperfect information. Backward induction starting from the information sets at the end ensures perfection, and one can construct a belief system supporting these strategies, so the result is a Perfect Bayesian Equilibrium.

Perfect Bayesian Equilibria in Selten's Horse



- It can be verified that there are two pure strategy Nash equilibria.
 - (C, c, R) and (D, c, L) .

Perfect Bayesian Equilibria in Selten's Horse (continued)

- However, if we look at *sequential rationality*, the second of these equilibria will be ruled out.
- Suppose we have (D, c, L) .
- The belief of player 3 will be $\mu_3(\text{left}) = 1$.
- Player 2, if he gets a chance to play, will then never play c , since d has a payoff of 4, while c would give him 1. If he were to play d , then player 1 would prefer C , but (C, d, L) is not an equilibrium, because then we would have $\mu_3(\text{left}) = 0$ and player 3 would prefer R .
- Therefore, there is a unique pure strategy Perfect Bayesian Equilibrium outcome (C, c, R) . The belief system that supports this could be any $\mu_3(\text{left}) \in [0, 1/3]$.

Job Market Signaling

- Consider the following simple model to illustrate the issues.
- There are two types of workers, high ability and low ability.
- The fraction of high ability workers in the population is λ .
- Workers know their own ability, but employers do not observe this directly.
- High ability workers always produce y_H , low ability workers produce y_L .

Baseline Signaling Model (continued)

- Workers can invest in education, $e \in \{0, 1\}$.
- The cost of obtaining education is c_H for high ability workers and c_L for low ability workers.
- **Crucial assumption** (“single crossing”)

$$c_L > c_H$$

- That is, education is more costly for low ability workers. This is often referred to as the “single-crossing” assumption, since it makes sure that in the space of education and wages, the indifference curves of high and low types intersect only once. For future reference, I denote the decision to obtain education by $e = 1$.
- To start with, suppose that education does not increase the productivity of either type of worker.
- Once workers obtain their education, there is competition among a large number of risk-neutral firms, so workers will be paid their *expected productivity*.

Baseline Signaling Model (continued)

- Two (extreme) types of equilibria in this game (or more generally in signaling games).
- ① **Separating**, where high and low ability workers choose different levels of schooling.
- ② **Pooling**, where high and low ability workers choose the same level of education.

Separating Equilibrium

- Suppose that we have

$$y_H - c_H > y_L > y_H - c_L. \quad (**)$$

- This is clearly possible since $c_H < c_L$.
- Then the following is an equilibrium: all high ability workers obtain education, and all low ability workers choose no education.
- Wages (conditional on education) are:

$$w(e = 1) = y_H \text{ and } w(e = 0) = y_L$$

- Notice that these wages are conditioned on education, and *not directly on ability*, since ability is not observed by employers.

Separating Equilibrium (continued)

- Let us now check that all parties are playing best responses.
- Given the strategies of workers, a worker with education has productivity y_H while a worker with no education has productivity y_L . So no firm can change its behavior and increase its profits.
- What about workers?
- If a high ability worker deviates to no education, he will obtain $w(e = 0) = y_L$, but

$$w(e = 1) - c_H = y_H - c_H > y_L.$$

Separating Equilibrium (continued)

- If a low ability worker deviates to obtaining education, the market will perceive him as a high ability worker, and pay him the higher wage $w(e = 1) = y_H$. But from (**), we have that

$$y_H - c_L < y_L.$$

- Therefore, we have indeed an equilibrium.
- In this equilibrium, education is valued simply because it is a signal about ability.
- Is “single crossing important”?

Pooling Equilibrium

- The separating equilibrium is not the only one.
- Consider the following allocation: both low and high ability workers do not obtain education, and the wage structure is

$$w(e = 1) = (1 - \lambda) y_L + \lambda y_H \text{ and } w(e = 0) = (1 - \lambda) y_L + \lambda y_H$$

- Again no incentive to deviate by either workers or firms.
- Is this Perfect Bayesian Equilibrium reasonable?

Pooling Equilibrium (continued)

- The answer is no.
- This equilibrium is being supported by the belief that the worker who gets education is no better than a worker who does not.
- But education is more costly for low ability workers, so they should be less likely to deviate to obtaining education.
- This can be ruled out by various different refinements of equilibria.

Pooling Equilibrium (continued)

- Simplest refinement: **The Intuitive Criterion** by Cho and Kreps.
- The underlying idea: if there exists a type who will never benefit from taking a particular deviation, then the uninformed parties (here the firms) should deduce that this deviation is very unlikely to come from this type.
- This falls within the category of “forward induction” where rather than solving the game simply backwards, we think about what type of inferences will others derive from a deviation.

Pooling Equilibrium (continued)

- Take the pooling equilibrium above.
- Let us also strengthen the condition (**) to

$$y_H - c_H > (1 - \lambda) y_L + \lambda y_H \text{ and } y_L > y_H - c_L$$

- Consider a deviation to $e = 1$.
- There is no circumstance under which the low type would benefit from this deviation, since

$$y_L > y_H - c_L,$$

and the low ability worker is now getting

$$(1 - \lambda) y_L + \lambda y_H.$$

- Therefore, firms can deduce that the deviation to $e = 1$ must be coming from the high type, and offer him a wage of y_H .
- Then (**) ensures that this deviation is profitable for the high types, breaking the pooling equilibrium.

Pooling Equilibrium (continued)

- The reason why this refinement is called **The Intuitive Criterion** is that it can be supported by a relatively intuitive “speech” by the deviator along the following lines:

“you have to deduce that I must be the high type deviating to $e = 1$, since low types would never ever consider such a deviation, whereas I would find it profitable if I could convince you that I am indeed the high type).”

- Of course, this is only very loose, since such speeches are not part of the game, but it gives the basic idea.
- **Overall conclusion:** separating equilibria, where education is a valuable signal, may be more likely than pooling equilibria.
- We can formalize this notion and apply it as a refinement of Perfect Bayesian Equilibria. It turns out to be a particularly simple and useful notion for signaling-type games.

Social Learning

- An important application of dynamic games of incomplete information as to situations of **social learning**.
- A group (network) of agents learning about an underlying state (product quality, competence or intentions of the politician, usefulness of the new technology, etc.) can be modeled as a dynamic game of incomplete information.
- The simplest setting would involve **observational learning**, i.e., agents learning from the observations of others in the past.
 - Here a **Bayesian** approach; we will discuss non-Bayesian approaches in the coming lectures.
- In observational learning, one observes past actions and updates his or her beliefs.
- This might be a good way of aggregating dispersed information in a large social group.
- We will see when this might be so and when such aggregation may fail.

The Promise of Information Aggregation: The Condorcet Jury Theorem

- An idea going back to Marquis de Condorcet and Francis Galton is that large groups can aggregate dispersed information.
- Suppose, for example, that there is an underlying state $\theta \in \{0, 1\}$, with both values ex ante equally likely.
- Individuals have **common values**, and would like to make a decision $x = \theta$.
- But nobody knows the underlying state.
- Instead, each individual receives a signal $s \in \{0, 1\}$, such that $s = 1$ has *conditional* probability $p > 1/2$ when $\theta = 1$, and $s = 0$ has conditional probability $p > 1/2$ when $\theta = 0$.

The Condorcet Jury Theorem (continued)

- Suppose now that a large number, N , of individuals obtain independent signals.
- If they communicate their signals, or if each takes a preliminary action following his or her signal, and N is large, from the (strong) **Law of Large Numbers**, they will identify the underlying state with probability 1.

Game Theoretic Complications

- However, the selfish behavior of individuals in game theoretic situations may prevent this type of efficient aggregation of dispersed information.
- The basic idea is that individuals will do whatever is best for them (given their beliefs) and this might prevent information aggregation because they may not use (and they may not reveal) their signal.
- Specific problem: **herd behavior**, where all individuals follow a pattern of behavior regardless of their signal (“bottling up” their signal).
- This is particularly the case with observational learning and may prevent the optimistic conclusion of the Condorcet Jury Theorem.
- To illustrate this, let us use a model origin only proposed by Bikchandani, Hirshleifer and Welch (1992) “A Theory of Fads, Fashion, Custom and Cultural Change As Informational Cascades,” and Banerjee (1992) “A Simple Model of Herd Behavior.”

Observational Social Learning

- Consider the same setup as above, with two states $\theta \in \{0, 1\}$, with both values ex ante equally likely.
- Once again we assume **common values**, so each agent would like to take a decision $x = \theta$. But now decisions are taken individually.
- Each individual receives a signal $s \in \{0, 1\}$, such that $s = 1$ has *conditional* probability $p > 1/2$ when $\theta = 1$, and $s = 0$ has conditional probability $p > 1/2$ when $\theta = 0$.
- **Main difference:** individuals make decisions **sequentially**.
 - More formally, each agent is indexed by $n \in \mathbb{N}$, and acts at time $t = n$, and observes the actions of all those were the fact that before t , i.e., agent n' acting at time t' observes the sequence $\{x_n\}_{n=1}^{t'-1}$.
- We will look for a Perfect Bayesian Equilibrium of this game.

The “Story”

- Agents arrive in a town sequentially and choose to dine in an Indian or in a Chinese restaurant.
- One restaurant is strictly better, underlying state $\theta \in \{Chinese, Indian\}$. All agents would like to dine at the better restaurant.
- Agents have independent binary private signals s indicating which risk I might be better (correct with probability $p > 1/2$).
- Agents observe prior decisions (who went to which restaurant), but not the signals of others.

Perfect Bayesian Equilibrium with Herding

- Let x^n be the history of actions up to and including agent n , and $\#x^n$ denote the number of times $x_{n'} = 1$ in x^n (for $n' \leq n$).

Proposition

There exists a pure strategy Perfect Bayesian Equilibrium such that $x_1(s_1) = s_1$, $x_2(s_2, x_1) = s_2$, and for $n \geq 3$,

$$x_n(s_n, x^{n-1}) = \begin{cases} 1 & \text{if } \#x^{n-1} > \frac{n-1}{2} + 1 \\ 0 & \text{if } \#x^{n-1} < \frac{n-1}{2} \\ s_n & \text{otherwise.} \end{cases}$$

- We refer to this phenomenon as **herding**, since agents after a certain number “herd” on the behavior of the earlier agents.
- Bikchandani, Hirshleifer and Welch refer to this phenomenon as **informational cascade**.

Proof Idea

- If the first two agents choose 1 (given the strategies in the proposition), then agent 3 is better off choosing 1 even if her signal indicates 0 (since two signals are stronger than one, and the behavior of the first two agents indicates that their signals were pointing to 1).
- This is because

$$\Pr[\theta = 1 \mid s_1 = s_2 = 1 \text{ and } s_3 = 0] > \frac{1}{2}.$$

- Agent 3 does not observe s_1 and s_2 , but according to the strategy profile in the proposition, $x_1(s_1) = s_1$ and $x_2(s_2, x_1) = s_2$. This implies that observing x_1 and x_2 is equivalent to observing s_1 and s_2 .
- This reasoning applies later in the sequence (though agents rationally understand that those herding the not revealing information).

Illustration


- Thinking of the sequential choice between Chinese and Indian restaurants:

Illustration

- Thinking of the sequential choice between Chinese and Indian restaurants:

Illustration

- Thinking of the sequential choice between Chinese and Indian restaurants:



Signal = 'Chinese'
Decision = 'Chinese'

Illustration


- Thinking of the sequential choice between Chinese and Indian restaurants:



Decision = 'Chinese' **Signal = 'Chinese'**
Decision = 'Chinese'

Illustration


- Thinking of the sequential choice between Chinese and Indian restaurants:



Decision = 'Chinese' **Decision = 'Chinese'** **Signal = 'Indian'**
Decision = 'Chinese'

Illustration

- Thinking of the sequential choice between Chinese and Indian restaurants:



Decision = 'Chinese' Decision = 'Chinese' Decision = 'Chinese'

Herding

- This type of behavior occurs often when early success of a product acts as a **tipping point** and induces others to follow it.
- It gives new insight on “diffusion of innovation, ideas and behaviors”.
- It also sharply contrasts with the efficient aggregation of information implied by the Condorcet Jury Theorem.
 - For example, the probability that the first two agents will choose 1 even when $\theta = 0$ is $(1 - p)^2$ which can be close to $1/2$.
 - Moreover, a herd can occur not only following the first two agents, but later on as indicated by the proposition.
- This type of behavior is also inefficient because of an **informational externality**: agents do not take into account the information that they reveal to others by their actions. A social planner would prefer “greater experimentation” earlier on in order to reveal the true state.

Other Perfect Bayesian Equilibria

- There are other, mixed strategy Perfect Bayesian Equilibria. But these also exhibit herding.
- For example, the following is a best response for player 2:

$$x_2(s_2, x_1) = \begin{cases} x_1 & \text{with probability } q_2 \\ s_2 & \text{with probability } 1 - q_2. \end{cases}$$

- However, for any $q_2 < 1$, if $x_1 = x_2 = x$, then $x_3(s_3, x^2)$ must be equal to x . Therefore, herding is more likely.
- In the next lecture, we will look at more general models of social learning in richer group and network structures.