Lecture 21: Auctions and Incomplete Information

Alexander Wolitzky

MIT

6.207/14.15: Networks, Spring 2022

Last Unit: Social Learning and Information Aggregation in Networks

The last part of this course studies situations where different agents in a group or network have different information.

- In an auction (e.g. the sponsored search auction responsible for Google's revenue), each bidder has private information about her valuation of the good(s).
- In a financial market, prediction market, or election, each participant has private information about the likelihood of different events and/or the quality of different options.
- In a social learning setting (e.g. learning by sharing information on a social network), each individual has access to different pieces of information, which she might or might not share with others.

Games with Incomplete Information

To model strategic interactions in these settings, we need to extend the game theory we've learned so far to cover **games with incomplete information** (also called **asymmetric** or **private** information).

- Today's lecture: game theory and application to auctions
- Next week: information aggregation (elections, prediction markets), social learning
- Last lecture: guest lecture by James Siderius on social media.

Games with Incomplete Information

Basic idea: before playing the game, each agent observes the realization of some (different) random variable.

- This is her called her **private information** or **type**.

Then, when it's time to play the game, each player calculates the posterior distribution of the variables that she does not observe.

- Assume players do this correctly, using **Bayes' rule**.
- So games with incomplete information are also called **Bayesian games.**

The formal, mathematical model of incomplete information games is somewhat complicated, but the key ideas are fairly simple.

- We start with an example, then go to the model.

Example: A Public Good Game

Each of two players has to decide whether to contribute to a public good that benefits both of them.

- E.g. each of two roommates has to decide whether to clean the bathroom.

As long as someone contributes, both players get a benefit of 1, but each player i who contributes also incurs a cost c_i.

Payoff matrix:

$$\begin{array}{ccc} & Contribute & Don't \\ Contribute & 1-c_1, 1-c_2 & 1-c_1, 1 \\ Don't & 1, 1-c_2 & 0, 0 \end{array}$$

So far, this is a standard (complete information) game.

- What are the pure-strategy Nash equilibria?

Example (cntd.)

Twist: Now assume each player knows **her own** cost of contributing c_i, but **not** the other player's cost.

- For example, each player might believe the other player's cost is distributed uniformly on some interval $\left[\underline{c}, \bar{c}\right]$.
- Or player 1 might know c_2, while player 2 believes c_1 is distributed uniformly on some interval $\left[\underline{c}, \bar{c}\right]$.
- These two possibilities correspond to different incomplete information games.

Let's work through the case where each cost c_i is distributed $U\left[0, 2\right]$, independently across players.

- We first ask what the equilibrium "should" be, then formally define the solution concept.

Example (cntd.)

$$\begin{array}{ccc} & Contribute & Don't \\ Contribute & 1-c_1, 1-c_2 & 1-c_1, 1 \\ Don't & 1, 1-c_2 & 0, 0 \end{array}$$

$c_1, c_2 \sim U[0,2]$, independent.

As in dynamic games, a strategy s_i for player i is a complete contingent plan for how she should play the game.

In this case, a (pure) strategy s_i specifies, for each possible realization of player i's cost $c_i \in [0,2]$, should player i contribute or not?

Example (cntd.)

$$\begin{array}{ccc} & Contribute & Don't \\ Contribute & 1-c_1, 1-c_2 & 1-c_1, 1 \\ Don't & 1, 1-c_2 & 0, 0 \end{array}$$

$c_1, c_2 \sim U[0,2]$, independent.

I claim that the following symmetric strategy profile is an equilibrium: for each $i = 1, 2$,

$$s_i(c_i) = \begin{cases} Contribute & \text{if } c_i \leq \frac{2}{3} \\ Don't & \text{if } c_i > \frac{2}{3} \end{cases}$$

Why is this an equilibrium?

Example (cntd.)

If player j follows strategy s_j, he ends up contributing with probability $\Pr\left(c_j \leq \frac{2}{3}\right) = \frac{1}{3}$.

If player i plays *Contribute*, her payoff is $1 - c_i$ regardless of what player j does.

If player i plays $Don't$, her expected payoff is

$$\underbrace{\frac{1}{3}}_{\Pr(j\text{ contributes})}(1) + \underbrace{\frac{2}{3}}_{\Pr(j\text{ doesn't contribute})}(0) = \frac{1}{3}.$$

Therefore, player i should contribute iff $c_i \leq \frac{2}{3}$.

- That is, she should follow strategy s_i.

So this strategy profile is an equilibrium.

- In fact, it is the unique equilibrium.

Recap

How did we analyze this example?

- We specified a **prior probability distribution** (in this case, uniform $[0, 2]$, independent) over each player's **type** (c_i).
- We noted that a **strategy** for each player is a mapping from her type to her action ($Contribute, Don't$).
- An **equilibrium** is a strategy profile where each player takes an optimal action, for every type she might have, taking her opponent's strategy as given. . .
- . . . where "optimal" means "optimal in expectation, given the player's uncertainty about her opponent's type (and hence her opponent's action)."

Let's formalize this reasoning.

Incomplete Information Games

A **game of incomplete information** consists of

- A finite set $N = \{1, \ldots, n\}$ of **players.**
- A set $\Theta = \Theta_1 \times \ldots \times \Theta_n$ of **types**, where Θ_i is the **type space** for player i.
- A set $A = A_1 \times \ldots \times A_n$ of **actions**, where A_i is the **action space** for player i.
- **Payoff functions** $u_i : A \times \Theta \rightarrow \mathbb{R}$ for each player i. (Note: payoffs can depend on actions **and** types)
- A **prior joint probability distribution** p on Θ.

Example

In the example:

- $N = \{1, 2\}$ (two players)
- $\Theta = [0, 2] \times [0, 2]$ (player i's type is $c_i \in [0, 2]$)
- $A = \{Contribute, Don't\} \times \{Contribute, Don't\}$
- Payoff functions are given by the payoff matrix.
  (Note: in the example, player i's payoff depends on the actions and **her own type** c_i, but not the other player's type c_j. In general, could depend on all types.)
- p is the uniform distribution on the square $[0, 2] \times [0, 2]$ (=product of two uniform distributions on $[0, 2]$).

Strategies

A **pure strategy** for player i is a function $s_i : \Theta_i \to A_i$.

- For each possible type the player could have, what does she do?

A **mixed strategy** for player i is a function $\sigma_i : \Theta_i \to \Delta(A_i)$, the set of probability distributions on A_i.

Interpretation: player i's type θ_i is "everything she knows".

- Player i's action can depend on θ_i, but it cannot depend on anything else.
- In our example, if player 2 follows the equilibrium strategy "*Contribute* iff $c_2 \leq \frac{2}{3}$", player 1 would love to play "*Contribute* iff $c_1 < 1$ **and** $c_2 > \frac{2}{3}$".
  But this is not a valid strategy for her, because it depends on player 2's type, which player 1 doesn't know.

Bayesian Nash Equilibrium

The main equilibrium concept for incomplete information games is called Bayesian Nash equilibrium.

Intuitively, a Bayesian Nash equilibrium is a strategy profile where each player maximizes her expected payoff given her opponents' strategies.

Formally, a strategy profile σ is a **Bayesian Nash equilibrium (BNE)** if, for all $i \in N$ and all σ_i',

$$\begin{aligned} & E_\theta \left[u_i\left(\sigma_i\left(\theta_i\right), \sigma_{-i}\left(\theta_{-i}\right); \theta_i, \theta_{-i}\right)\right] \\ \geq \quad & E_\theta \left[u_i\left(\sigma_i'\left(\theta_i\right), \sigma_{-i}\left(\theta_{-i}\right); \theta_i, \theta_{-i}\right)\right], \end{aligned}$$

where $E_\theta[\cdot]$ denotes expectation over θ.

Finite Types

When the type space Θ is finite, we have

$$\begin{aligned} & E_\theta \left[u_i \left(\sigma_i \left(\theta_i \right), \sigma_{-i} \left(\theta_{-i} \right); \theta_i, \theta_{-i} \right) \right] \\ = & \sum_{\theta_i, \theta_{-i}} u_i \left(\sigma_i \left(\theta_i \right), \sigma_{-i} \left(\theta_{-i} \right); \theta_i, \theta_{-i} \right) p \left(\theta_i, \theta_{-i} \right). \end{aligned}$$

Hence, a strategy profile σ is a BNE if, for all $i \in N$ and all σ_i',

$$\begin{aligned} & \sum_{\theta_i, \theta_{-i}} u_i \left(\sigma_i \left(\theta_i \right), \sigma_{-i} \left(\theta_{-i} \right); \theta_i, \theta_{-i} \right) p \left(\theta_i, \theta_{-i} \right) \\ \geq & \sum_{\theta_i, \theta_{-i}} u_i \left(\sigma_i' \left(\theta_i \right), \sigma_{-i} \left(\theta_{-i} \right); \theta_i, \theta_{-i} \right) p \left(\theta_i, \theta_{-i} \right). \end{aligned}$$

Infinite Types

When the type space Θ is infinite (as in the example, where $c_i \in [0,2]$), we have

$$\begin{aligned} & E_\theta \left[u_i \left(\sigma_i \left(\theta_i \right), \sigma_{-i} \left(\theta_{-i} \right); \theta_i, \theta_{-i} \right) \right] \\ = & \int_{\Theta_i \times \Theta_{-i}} u_i \left(\sigma_i \left(\theta_i \right), \sigma_{-i} \left(\theta_{-i} \right); \theta_i, \theta_{-i} \right) dp \left(\theta_i, \theta_{-i} \right). \end{aligned}$$

Hence, a strategy profile σ is a BNE if, for all $i \in N$ and all σ_i',

$$\begin{aligned} & \int_{\Theta_i \times \Theta_{-i}} u_i \left(\sigma_i \left(\theta_i \right), \sigma_{-i} \left(\theta_{-i} \right); \theta_i, \theta_{-i} \right) dp \left(\theta_i, \theta_{-i} \right) \\ \geq & \int_{\Theta_i \times \Theta_{-i}} u_i \left(\sigma_i \left(\theta_i \right), \sigma_{-i} \left(\theta_{-i} \right); \theta_i, \theta_{-i} \right) dp \left(\theta_i, \theta_{-i} \right). \end{aligned}$$

BNE: Alternative Definition

Equivalent definition (often easier to check): rather that requiring that each **player** maximizes her expected payoff (taking the expectation over all players types, including her own type), require that each **type** of each player maximizes her expected payoff (taking the expectation over the other players' types only).

With this definition, if Θ is finite, a strategy profile σ is a **Bayesian Nash equilibrium (BNE)** if, for all $i \in N$, **all** $\theta_i \in \Theta_i$**,** and all $a_i \in A_i$,

$$\begin{aligned} & \sum_{\theta_{-i}} u_i\left(\sigma_i\left(\theta_i\right), \sigma_{-i}\left(\theta_{-i}\right); \theta_i, \theta_{-i}\right) p\left(\theta_{-i} | \theta_i\right) \\ \geq \quad & \sum_{\theta_{-i}} u_i\left(a_i, \sigma_{-i}\left(\theta_{-i}\right); \theta_i, \theta_{-i}\right) p\left(\theta_{-i} | \theta_i\right), \end{aligned}$$

where $p\left(\theta_{-i} | \theta_i\right) = p\left(\theta_i, \theta_{-i}\right) / p\left(\theta_i\right)$, by Bayes' rule.

If Θ is infinite and types are not independent, must define conditional expectation more carefully, but the idea is the same.

Example

Let's show that "*Contribute* iff $c_i \leq \frac{2}{3}$" is the **unique** BNE in the example (up to indifference at $c_i = \frac{2}{3}$).

Suppose $s_i(c_i) = Contribute$ for some type c_i.

This implies

$$E_{c_j}\left[u_i\left(Contribute, s_j(c_j); c_i\right)|c_i\right] \geq E_{c_j}\left[u_i\left(Don't, s_j(c_j); c_i\right)|c_i\right],$$

or equivalently

$$1 - c_i \geq \Pr\left(s_j(c_j) = Contribute\right).$$

This implies that the same inequality holds for every $c_i' < c_i$, so all lower types must also *Contribute*.

Example (cntd.)

This implies that any BNE must take a **cutoff form**: for some cutoffs c_1^*, c_2^*, we have

$$s_1(c_1) = \begin{cases} Contribute & \text{if } c_1 \leq c_1^* \\ Don't & \text{if } c_1 > c_1^* \end{cases},$$
$$s_2(c_2) = \begin{cases} Contribute & \text{if } c_2 \leq c_2^* \\ Don't & \text{if } c_2 > c_2^* \end{cases}.$$

Just need to determine what pairs of cutoffs (c_1^*, c_2^*) form a BNE.

For strategy profile (s_1, s_2) to be a BNE, the cutoff types must be **indifferent** (otherwise, types just on either side of the cutoff would have to be taking the same action).

This implies

$$1 - c_1^* = \Pr(s_2(c_2) = Contribute) = \frac{c_2^*}{2}.$$

Symmetrically,

$$1 - c_2^* = \frac{c_1^*}{2}.$$

Example (cntd.)

We have

$$\begin{aligned} 1 - c_1^* &= \frac{c_2^*}{2}, \\ 1 - c_2^* &= \frac{c_1^*}{2}. \end{aligned}$$

This system of equations has unique solution

$$c_1^* = c_2^* = \frac{2}{3}.$$

Therefore, "*Contribute* iff $c_i \leq \frac{2}{3}$" is the unique BNE.

How to Find BNE in Incomplete Information Games

When A is discrete and Θ is continuous, look for "cutoff strategies" as in the example.

When A and Θ are both discrete, the game is often simple enough to check all possible strategies, as in static games of complete information.

When A and Θ are both continuous (e.g. in an auction where a bidder's value for the good is a continuous variable, and she can bid any amount in the auction), finding equilibria can be trickier and you won't be asked to do this without a lot of guidance.

Summary

- Strategic situations where one or more parties have private information are modeled as **games of incomplete information**.
- A strategy in an incomplete information game is a function from a player's private information (or **type**) to her action.
- In a Bayesian Nash equilibrium, each player chooses her strategy to maximize her expected utility; equivalently, each type of each player chooses her action to maximize her expected utility.

Auctions

Auctions are a leading application of incomplete-information games.

In an auction, one or several goods are up for sales, and multiple buyers (or **bidders**) have private information about how much they want the good(s) (their **valuations**, which are their types in this context).

The bidders then place bids, which determine how the good is allocated and how much is paid, according to the rules of the auction.

Examples of Auctions

- Auction houses (Christie's, Sotheby's, etc.) selling art and other valuables.
- US government selling treasury bills or natural resource rights (timber, oil, spectrum).
- Search engines selling advertising rights for keywords.

Many others: real estate, livestock/produce, electricity, corporate debt, used cars, ...

Auction Formats

There are many different ways of running an auction, called **auction formats**.

There are four classic auction formats used for the sale of a single good (together with many variants):

- **English auction**: price gradually rises until all but one bidder drop out.
- **Dutch auction:** price gradually falls until one bidder claims the object.
- **First-price sealed-bid auction:** all bidders simultaneously submit sealed bids; highest bidder wins and pays bid.
- **Second-price sealed-bid auction:** all bidders simultaneously submit sealed bids; highest bidder wins and pays the **second-highest** bid.

However, there are many other formats.

Generalized Second-Price Auction

One new and important format is the **generalized second-price (GSP) auction**, used (with modifications) by search engines to offer keyword advertising slots.

When you type a search term ("query") into Google, several ads may appear above or below the search results.

Which ad goes in which slot is determined by a keyword-specific auction.

- Each advertiser places a per-click bid b_i.
- The advertiser with the highest bid gets the first slot, the advertiser with the second-highest bid gets the second slot, and so on.
- Each advertiser pays per-click price equal to the bid of the **next-highest** advertiser (hence, "generalized second-price").
- There may also be a minimum price, below which no further slots are allocated.

Common Value Auctions

Another important possibility in auctions is that a bidder may have private information that is relevant not only for her own value, but also for other bidders' values.

- E.g. different oil companies bidding for rights to the same oil tract may have different information about the amount of oil in the tract, which affects all of their values for the tract.

These are called auctions with **interdependent values**: the extreme case where bidders have different information but the same final value for the good is called **common values**.

Note that, in a common values auction, winning the auction is "bad news" about the value of the good: the fact that others did not bid aggressively suggests that they had less favorable information about the good than you did.

Failing to take this effect into account leads to the **winner's curse**, which has been documented in natural resource auctions, corporate IPOs, free agency in professional sports, etc.

How Zillow Lost $881M Buying Houses in 2021

2018: Launched Zillow Offers.

- Tons of data and state-of-the-art algorithm to predict home prices.
- Offered to buy houses a bit below predicted price, then do minor repairs and sell.
- Simulated strategy for years, back-tested, found they would have made a lot of money.
- After launched, bought tens of thousands of houses.
  Lost $881M, 25% of workforce, 25% of market cap.

How Zillow Lost $881M Buying Houses in 2021

What went wrong? (probably)

- For sure, **on average** Zillow can guess market price for a house better than its owner.
- But which owners will accept Zillow's offer?
  The ones who know there's a problem with their house!
  (More generally, the ones whose estimate of the value of their house is much lower than Zillow's.)

Another way to look at it:

- Suppose Zillow and homeowner get different signals of market price.
- Even if Zillow's signal is much more accurate than owner's, Zillow's offer largely reveals its signal, so the owner ends up with the informational advantage after seeing Zillow's offer plus her own information.
- Zillow only gets homes whose owners think Zillow bid too much, after taking the bid itself into account.

Plan

We'll study bidding and revenue in standard auction formats.

- Skip GSP auction. It's covered in EK Chapter 15 if you're curious.
- Auctions with interdependent values would be covered in classes on game theory or market design.

Modeling Auctions

Classic model of single-item auction:

- n bidders.
- Each bidder i has a **value** $v_i \in [0, \bar{v}]$ for the good. If bidder i wins the good and pays price p_i, her payoff is

$$v_i - p_i.$$

- Bidder i's value is her **type/private information**. It is distributed according to some cdf F_i on $[0, \bar{v}]$.
- We assume values are independent and identically distributed (can be relaxed).
- A **strategy** for bidder i is a function $\beta_i : [0, \bar{v}] \to \mathbb{R}_+$, indicating how much she bids as a function of her value.

Note: The assumption that only bidder i's own private information affects her value is called **private values**. This is the opposite of **common values**, which arise in the context of the winner's curse.

Modeling Auctions (cntd.)

We have not yet fully defined a game, because we have not specified how the bids determine how the good is allocated and who pays what (that is, the auction format).

We will analyze the four standard auction formats, comparing the bidders' equilibrium strategies and the expected revenue raised by the auction.

Second-Price Auction

All bidders simultaneously submit sealed bids; highest bidder wins and pays the **second-highest** bid.

The second-price auction has the following very special property:

- Recall that a strategy is **weakly dominant** if it is a best response to any opposing strategies.

Theorem

In a second-price auction, it is a weakly dominant strategy for each bidder to bid her true value. That is, $\beta_i(v_i) = v_i$ *for all* v_i *is always an optimal strategy.*

- In a second-price auction it is impossible to increase your payoff by bidding anything other than your true value for the good.

Proof

We show that, for any profile of opposing bids b_{-i}, it is a best response for player i to bid v_i.

- Let $\bar{b}$ be the highest bid among the b_{-i}. Consider 2 cases.

Case 1: $\bar{b} > v_i$.

- If i bids $b_i = v_i$, she loses the auction and gets payoff 0.
- If she bids anything less than $\bar{b}$, she still loses and gets payoff 0.
- If she bids more than $\bar{b}$, she wins the auction but has to pay $\bar{b}$, which gives her payoff $v_i - \bar{b} < 0$.
- Hence, she cannot gain by bidding any $b_i \neq v_i$.
- "If you lose the auction when you bid your value, there's no way you can win the auction without paying too much."

Proof (cntd.)

Case 2: $\bar{b} < v_i$**.**

- If i bids $b_i = v_i$, she wins the auction and has to pay $\bar{b}$, which gives her payoff $v_i - \bar{b} > 0$.
- If she bids anything more than $\bar{b}$, she still wins and has to pay $\bar{b}$, which again gives payoff $v_i - \bar{b}$.
  - This is the key feature of the 2nd-price auction that makes truthful bidding optimal: conditional on the event that you win, the price is independent of your bid, so there is nothing to be gained by bidding "dishonestly".
- If she bids less than $\bar{b}$, she loses and gets payoff 0.
- Hence, she cannot gain by bidding any $b_i \neq v_i$.
- "If you win the auction when you bid your value, there's no way to reduce the price you pay."

Expected Revenue

What is the expected revenue raised by the 2nd-price auction?

It is equal to the expectation of the second-highest value.

- Given a vector $v \in \mathbb{R}^n$, let $v^{(2)}$ denote the value of the second-highest component of v.
- Expected revenue in the 2nd-price auction equals $E\left[v^{(2)}\right]$.
- For example, if each of the n bidders has $v_i \sim U\left[0, \bar{v}\right]$, can show that
$$E\left[v^{(2)}\right] = \frac{n-1}{n+1}\bar{v}.$$
(Fact: if n random variables are distributed iid $U\left[0, x\right]$, the expectation of the k^{th} highest draw is $\frac{n+1-k}{n+1}x$.)

First-Price Auction

All bidders simultaneously submit sealed bids; highest bidder wins and pays **her own** bid.

In a 1st-price auction, it is *not* an equilibrium for everyone to bid their true values.

- If you bid $b_i = v_i$, you're *guaranteed* to get payoff 0: either you lose the auction, or you win the auction but pay exactly as much as the good is worth to you.
- If you instead bid $b_i < v_i$ ("shade your bid"), you still might win the auction, and when you do you get a positive payoff.
- Hence, in equilibrium in the first-price auction, bidders will use strategies where $\beta_i(v_i) < v_i$ for all v_i.

The question is *by how much* players shade their bids.

- Note: holding bids fixed, 1st-price auction generates more revenue than 2nd-price.
- But we have seen that eqm bids are lower in 1st-price.
- Which format yields higher expected revenues overall?

Example

Let's focus on the case where all bidder values are independently distributed $U\left[0, \bar{v}\right]$.

We claim that it's a BNE for each bidder to bid $\frac{n-1}{n}$ times her value: that is, each bidder shades down by fraction $\frac{1}{n}$.

- *Deriving* this bidding strategy is somewhat complicated (we skip it), but *checking* that it's an equilibrium is easy.

Example (cntd.)

Let $v^{(1)}$ denote the value of the highest component of v.

If everyone else bids $\frac{n-1}{n}$ times their value, player i wins the auction with bid b iff $v_{-i}^{(1)} \leq \frac{n}{n-1}b$.

This happens with probability $\left(\frac{n}{n-1}\frac{b}{\bar{v}}\right)^{n-1}$.

Hence, player i with value v_i should choose b to maximize

$$\left(\frac{n}{n-1}\frac{b}{\bar{v}}\right)^{n-1}(v_i - b)\text{, or equivalently } b^{n-1}v_i - b^n.$$

FOC:

$$\begin{aligned}(n-1)\,b^{n-2}v_i &= nb^{n-1}, \text{ or equivalently}\\ \frac{n-1}{n}v_i &= b.\end{aligned}$$

Hence, bidding $b = \frac{n-1}{n}v_i$ is indeed optimal.

Revenue Comparison

With n bidders and independent $v_i \sim U[0, \bar{v}]$:

Expected revenue in the 1st-price auction equals

$$E\left[\frac{n-1}{n} v^{(1)}\right] = \frac{n-1}{n} \times \frac{n}{n+1} \bar{v} = \frac{n-1}{n+1} \bar{v}.$$

Expected revenue in the 2nd-price auction equals

$$E\left[v^{(2)}\right] = \frac{n-1}{n+1} \bar{v}.$$

So, the 1st-price and 2nd-price auctions yield *exactly the same* expected revenue!

- Compared to 2nd-price auction, in the 1st-price auction bidders shade their bids by just enough to compensate for the higher payments for fixed bids.

Coincidence?

The fact that the 2nd-price and 1st-price auctions yield the same expected revenue is not a coincidence. It is a special case of an important result called the **revenue equivalence theorem**: when bidder values are independently distributed, **any** two auction formats that always give the good to the highest bidder must generate the same expected revenue.

- Beyond our scope.

Coincidence?

The fact that the 2nd-price and 1st-price auctions yield the same expected revenue is not a coincidence. It is a special case of an important result called the **revenue equivalence theorem**: when bidder values are independently distributed, **any** two auction formats that always give the good to the highest bidder must generate the same expected revenue.

- Beyond our scope.

Note: the revenue equivalence theorem allows the possibility that more revenue can be raised by *not* always giving the good to the highest bidder.

- For example, auctions typically raise more revenue if they use a **reserve price**: a minimum price below which the seller keeps the good.
- Example: with only **one** bidder with $v_i \sim U[0,1]$, the 2nd-price or 1st-price auction without a reserve price raises 0 revenue, but with a reserve price of $\frac{1}{2}$ either "auction" raises expected revenue of $\frac{1}{4}$.

Summary

- Auctions are a leading example of games with incomplete information, and are important in both online and offline markets.
- Many common single-good auction formats yield the same expected revenue, due to the **revenue equivalence theorem**. However, the auctioneer can increase revenue by sometimes withholding the good through the use of a **reserve price**.
- Other important topics in auction theory include auctions for multiple goods such as the GSP auction, and auctions with common values.