Paul Klemperer ©1995-2026

MPhil 1st Year Micro Notes

1st version 1995; last revision 2026

My lecture will cover Sections 2 and 5 only.

Sections 0 and 1 are introductory material that you should read on your own (and are examinable). Please at least skim these sections prior to my lecture. (They are elementary, with the possible exception of Section 1.3 which will be illustrated in the problem set.)

If you don't feel very confident with technical material, please also look at Section 2.1 in advance.

Following up references in the footnotes is strictly optional and non-examinable.

Where I refer to "my book", this is *Auctions: Theory and Practice*, downloadable FREE from the "ONLINE BOOK" on my website at www.paulklemperer.org.

ALL the Oxford MPhil in Economics auctions exam questions set in the years 1995-2024 are on my website, together with their answers.

The following questions can be answered well on the basis of this year's course, and would be good practice for all 1st year students:1, 4, 5, 6 (parts i, ii, iii), 11, 12, 13, A8, A11, A16, A17, A20, A25, A28, A30, and A31.

(Questions 9, 10, 15, A3, A13, and A32 use the non-examinable material in Section 4 of the notes; questions 2, 3, 13, 16, 20, A2, and A6 are more challenging than I expect most 1st years will want to attempt, and go a little beyond the material that I will cover, but can be attempted by any 1st years who want to stretch themselves–they all come from past 2nd year MPhil exams; the remaining questions go well beyond the material I will cover this year.)

Questions A16, A22, A27, and A28 were "short" questions; the remainder were all "long" questions, roughly of the length of the questions being set in 2026. However the definitions of "short" and "long" have varied over time.

0 Introduction

0.1 Why study auctions?

Auction theory has been the basis of fundamental *theoretical work* in industrial economics, public economics, labour economics and finance, and has helped our understanding of price formation in markets.[1]

Starting around the 1990s there was an explosion of interest in its *practical applications*, especially in organising the sale of government assets (for example, treasury bonds, radio spectrum licenses, and firms to be privatised) and in developing new markets for electricity, transport, etc.[2] Auctions moved even further centre stage in the financial crisis, when auction theorists were called in to design new markets to replace those that failed.[3]

Furthermore, because auctions are such simple and well-defined economic environments, they provide a valuable testing ground for economic theory that has been increasingly exploited in *experimental and empirical work*.

Finally, because auctions are also simple and well-defined *games*, they provide an excellent revision and application of the ideas that you studied in your game theory course. I'll particularly discuss connections to basic oligopoly theory.

0.2 The four major types of auctions

Four basic types of auctions are widely used: the English auction (also called the open, oral, or ascending-bid auction), the descending-bid auction (used in the sale of flowers in the Netherlands and called the Dutch auction by economists), the first-price sealed-bid auction, and the second-price sealed-bid auction (called

[1] Auction theory was recognised most recently by the 2020 Nobel prize to Paul Milgrom and Bob Wilson. The pioneer was William Vickrey (who received a Nobel prize in 1996 jointly with Jim Mirrlees [Oxford professor 1969-95]); the 2007 prize shared by Leonid Hurwicz, Eric Maskin and Roger Myerson was also in large part for contributions to auction theory.

[2] For example, in 2000 I was the principal designer of the UK government's 3G mobile phone auction that raised £22.5 billion (or 2.5% of GDP).

[3] I designed a new auction (the "Product-Mix Auction", see www.nuff.ox.ac.uk/users/klemperer/productmix.pdf) for the Bank of England to get liquidity to banks, building societies, etc when its auctions failed after the 2007 Northern Rock bank run. (See 5-minute film produced by the *Guardian* newspaper https://www.nuff.ox.ac.uk/users/klemperer/auctions/Auctions_Final.m4v .)

the Vickrey auction by economists and, confusingly, called the Dutch auction by investment bankers).

In the English auction, the price is successively raised until only one bidder remains. This can be done by having the seller announce prices, or by having the bidders call out prices themselves, or by having bids submitted electronically with the best current bid posted. Antiques and artwork are two examples of goods commonly sold using the English auction, and houses are sometimes sold this way, too. One version of the English auction is called the "Japanese auction", and this is the type most commonly studied by auction theorists. In this version, the price rises continuously while bidders gradually quit the auction. Once someone quits, she is not let back in, and all bidders can see who is still present in the auction. Also, there is no possibility for one bidder to pre-empt the process by making a large "jump bid".

The descending-bid auction works in just the opposite way from the English auction. For example, in the Dutch flowers auctions, the potential buyers all sit in a room at desks with buzzers connected to an electronic clock at the front of the room. The interior of the clock has information about what is being sold and the price at which the auction starts. Once the auction begins, a series of lights around the edge of the clock indicate to what percentage of the original asking price the good has fallen. As soon as one bidder buzzes in, she gets the flowers at the price indicated on the clock.[4] Fish are sold this way in Israel, as is tobacco in Canada.[5]

In the first-price sealed-bid auction bidders each submit a single bid, without being able to observe what their competitors are bidding. The highest bidder wins, and pays the amount of her bid. First-price sealed-bid auctions are used in auctioning mineral rights in government-owned land; they are also sometimes

[4] Actually, if there are several lots of the same flowers from the same seller available that morning, the buyer can choose to buy only some of the available lots, and the rest will be re-auctioned.

[5] At least, these things used to be the case–I haven't checked recently.

used in the sales of artwork and real estate. This method is also often used in procurement (that is, competing contractors submit prices and the *lowest* bidder wins and receives her price for fulfilling the contract).

In the second-price sealed-bid auction, bidders submit sealed bids having been told that the best bid(s) will win, but the winner(s) will pay the best losing bid. The multi-unit version is often used by governments when selling foreign exchange and by companies when buying back shares. In most European countries Treasury securities are sold through the multi-unit equivalent of the first-price auction (every winner pays her own bid), and U.S. Treasury auctions used to be run this way too, but the U.S. Treasury now uses the multi-unit version of the second-price sealed-bid auction, and the U.K. has recently been using both methods.

1 Independent private values

1.1 The Basic IPV Model

In the basic independent private values (IPV) model, n bidders compete for a single unit. Bidder i values the unit at v_i, which is private information to her, but it is common knowledge that each v_i is independently drawn from the same continuous distribution $F(v)$ on $[\underline{V}, \overline{V}]$ (so $F(\underline{V}) = 0, F(\overline{V}) = 1$) with density $f(v)$. That is, a bidder's "type" is her value v_i. Bidders do not know each others' types, but they do know the distribution $F(\cdot)$ from which those types are drawn. All bidders are risk-neutral.

1.2 The Uniform Distribution

Working out simple examples often helps to develop one's understanding of an economic model. They can both improve your intuition about the way your current model works, and suggest directions in which it might be interesting to develop it further. In my experience the uniform distribution

$(F(v) = \frac{v - \underline{V}}{\overline{V} - \underline{V}}, f(v) = \frac{1}{\overline{V} - \underline{V}})$ is often particularly easy to work with in auction theory.

Note that the uniform distribution corresponds to linear demand, in that the proportion of buyers with valuations above any price p (that is $\frac{\overline{V} - p}{\overline{V} - \underline{V}}$) is linear in p.

The following fact is very helpful: The expected k^{th} highest value among n values independently drawn from the uniform distribution on $[\underline{V}, \overline{V}]$ is $\underline{V} + \left(\frac{n+1-k}{n+1}\right)(\overline{V} - \underline{V})$.

1.3 The 1st Price Sealed Bid Auction

In this auction every player independently chooses a bid without seeing others' bids, and the object is sold to the bidder who makes the highest bid. The winner pays her bid (i.e. the price is the highest or "first" price bid). We look for a symmetric Nash equilibrium in which a bidder with value v chooses the bid $b(v)$, and assume that b is a continuous strictly increasing function of v, that is, bidders with higher values bid more than bidders with lower values and no range of "types" of bidders all choose the same bid.[6]

Consider player i with value v_i. In Nash equilibrium she assumes all other players will bid according to the function $b(v)$, so how should she bid? Imagine she chooses the bid $\widetilde{b}$. It is useful to think about the value $\widetilde{v}$ of the type of bidder she would just tie with, that is, we let $b(\widetilde{v}) = \widetilde{b}$.[7] Thus bidding $\widetilde{b}$ mimics the way type $\widetilde{v}$ would bid, and so beats any other bidder j with the probability that $v_j < \widetilde{v}$, which equals $F(\widetilde{v})$, in equilibrium. So mimicing $\widetilde{v}$ would beat all the other $(n-1)$ bidders with probability $(F(\widetilde{v}))^{n-1}$ and yields expected surplus to player

[6] Justifying this is left as a (strictly optional) exercise for the technically minded. They may also look at examples 6.3, 6.4 and (especially) 6.5 in Fudenberg and Tirole's book "Game Theory". Fudenberg and Tirole also show that the symmetric equilibrium is the unique equilibrium.

[7] Why does there exist a type $\widetilde{v}$, whose bid our bidder will just tie with? Recall we are looking for an equilibrium in which all types' equilibrium bids are in the range $[b(\underline{V}), b(\overline{V})]$. So if player i bids $\widetilde{b} > b(\overline{V})$ she could bid a bit less and still always win. If player i bids $\widetilde{b} < b(\underline{V})$ she will never win and so makes zero profit. And if she bids $\widetilde{b} \in [b(\underline{V}), b(\overline{V})]$ she will tie with some type because we assumed the bidding function is continuous.

i of

$$S(v_i, \widetilde{v}) = (v_i - b(\widetilde{v}))(F(\widetilde{v}))^{n-1}.$$

Note the tradeoff the bidder faces. Choosing a low $\widetilde{v}$ to mimic, that is, choosing a low bid $b(\widetilde{v})$, increases the value of winning, $(v_i - b(\widetilde{v}))$, but reduces the probability of doing so, $(F(\widetilde{v}))^{n-1}$.

Choosing the best bid to make is equivalent to choosing the best $\widetilde{v}$ to mimic, which we can do by looking at the first-order condition

$$\frac{\partial S(v_i, \widetilde{v})}{\partial \widetilde{v}} = -b'(\widetilde{v})(F(\widetilde{v}))^{n-1} + (v_i - b(\widetilde{v}))(n-1)(F(\widetilde{v}))^{n-2} f(\widetilde{v}).$$

For the bidding function $b(v)$ to be a (symmetric) equilibrium, i's best-response to all others bidding according to this function must be to do likewise, i.e. her optimal choice of $\widetilde{b}$ is $b(v_i)$ and of $\widetilde{v}$ is v_i.

So

$$\frac{\partial S}{\partial \widetilde{v}}(v_i, \widetilde{v}) = 0 \text{ at } \widetilde{v} = v_i$$

$$\Rightarrow \qquad b'(v_i) = (v_i - b(v_i))(n-1)\frac{f(v_i)}{F(v_i)}. \tag{1}$$

Using the boundary condition

$$b(\underline{V}) = \underline{V}$$

(it is obvious type $\underline{V}$ will not bid more than $\underline{V}$, and we assume the auctioneer will not accept lower bids than $\underline{V}$), we can solve the differential equation (1) for the equilibrium.[8]

For example, for the uniform distribution, (1) yields

$$b'(v_i) = (v_i - b(v_i))(n-1)\left(\frac{1}{v_i - \underline{V}}\right)$$

which is solved by[9]

$$b(v_i) = \underline{V} + \left(\frac{n-1}{n}\right)(v_i - \underline{V}).$$

[8] We will give the general solution in Section 2.

[9] As I'll emphasise in the lectures, I am teaching economics not mathematics, so I won't require you to know how to solve differential equations (though I would now expect you to recognise that $b(v) = C + (\frac{K}{K+1})(v - C)$ solves $b'(v) = (v - b(v))\frac{K}{(v-C)}$, for any constants C and K!).

For this example, the seller's expected revenue is

$$E\left\{\max_{i=1,..,n} b(v_i)\right\} = E\left\{\underline{V} + \left(\frac{n-1}{n}\right)\left(\max_{i=1,..,n} v_i - \underline{V}\right)\right\}.$$

But using the result from Section 1.2,

$$E\left\{\max_{i=1,..,n} v_i\right\} = \underline{V} + \left(\frac{n}{n+1}\right)(\overline{V} - \underline{V}),$$

so the expected revenue is $\underline{V} + \left(\frac{n-1}{n+1}\right)(\overline{V} - \underline{V})$.

1.4 The Dutch Auction

In this auction, the auctioneer starts at a very high price, say $\overline{V}$, and then lowers the price continuously. The first bidder who calls out that she will accept the current price wins the object at the current price.

Note that although I have described this as a dynamic problem, each bidder's problem is essentially static. Each bidder must choose a price at which he will call out, conditional on no other bidder having yet called out; and the bidder who chooses the highest price wins the object at the price she calls out. Thus this game is *strategically equivalent* to the 1st Price Sealed Bid Auction, and players' bidding functions and the expected payoff of the seller are therefore exactly those we derived in section 1.3.

Therefore this auction form is sometimes referred to as an Open 1st Price Auction.

1.5 The 2nd Price Sealed Bid Auction

In this auction every player independently chooses a bid, without seeing the other players' bids, and the object is sold to the bidder who makes the highest bid. However, the price she pays is the *second*-highest bidder's bid, or "second price". (This auction is sometimes called a Vickrey auction after William Vickrey who wrote the seminal paper on auctions in 1961 in the *Journal of Finance*.)

In this auction it is optimal for a player i to bid her true value, v_i, whatever other players do. In other words "truth telling" is a dominant strategy equilib-

rium (and so also a Nash equilibrium). We show this by showing that if you (player i) bid less than v_i, you may be worse off but will never be better off. Similarly, if you bid more than v_i, you may be worse off but never better off.

1. (a) Consider bidding $v_i - x$ when your true value is v_i. If the highest bid other than yours is w, then if $v_i - x > w$ you win the auction and pay w, just as if you bid v_i. If $w > v_i$ you lose the auction and get nothing, just as if you bid v_i. But if $v_i > w > v_i - x$, bidding $v_i - x$ causes you to lose the auction and get nothing, whereas if you had bid v_i, you would have won the auction and paid w for a net surplus of $v_i - w$. So you never gain and might lose if you bid $v_i - x$.

   (b) Consider bidding $v_i + x$. If the highest bid other than yours is w, then if $v_i > w$ you win and pay w, just as if you bid v_i. If $w > v_i + x$ you lose and pay nothing, just as if you bid v_i. But if $v_i + x > w > v_i$ having bid $v_i + x$ causes you to "win" an auction you otherwise would have lost, and you have to pay $w > v_i$ so you get negative surplus. So bidding $v_i + x$ may hurt you compared with bidding v_i, but it never helps you.

Since everyone is bidding her true value, the seller's expected revenue is the expected second highest value of the n values. Using the result from section 1.2, for the uniform distribution this is $\underline{V} + \left(\frac{n-1}{n+1}\right)(\overline{V} - \underline{V})$.

1.6 The English and Japanese Auctions

In English auctions the price rises until only one bidder remains. These auctions are run in a variety of ways (see section 0), but are usually modelled as "Japanese auctions" in which the price rises continuously. Clearly it is a dominant strategy to stay in the bidding until the price gets up to your value. The next-to-last person will drop out when her value is reached, so the person with the highest value will win at a price equal to the value of the second-highest bidder.

Thus this auction form is equivalent to that of the previous subsection, and is sometimes referred to as an Open 2nd Price Auction.

However, if there are more than two bidders, the Open and Sealed 2nd Price auctions are *not* strategically equivalent. That is, the sets of strategies available to players are then different in the two auctions. In a Japanese auction, players see when other bidders quit the auction and could, if they wish, condition their behaviour on this information (except with just two bidders, in which case the auction ends as soon as either bidder quits).

In the IPV model bidders do not wish to condition their behaviour on when others quit, because each bidder knows exactly how much the item is worth to her, but in a model in which you were unsure of your value and other players had private information about your value, your quitting price would depend on the prices at which other players quit. Thus in a common-value model (see section 4) the Japanese and sealed 2nd price auctions (and indeed other versions of the English auction which reveal different amounts of information about which bidders are still active) are no longer equivalent when there are more than two bidders.

1.7 Comparison of 1st and 2nd Price Auctions

Notice that while the equilibrium in the second-price auctions was a *dominant strategies* equilibrium - each bidder bidding his true value regardless of what others bid - the equilibrium in the first-price auctions is only a *Bayesian-Nash* equilibrium. For example, if other people were not behaving optimally in the first-price auction, so that none of them bid over £10.00 no matter what their true value, your optimal strategy would never involve bidding more than £10.01. Similarly, getting secret information about an opponent's value is worthless in a second-price auction (since it won't affect your bid) but is potentially valuable in a first-price auction.

Assuming players do bid equilibrium strategies, however, we notice that the

seller's expected revenues are the same in our uniform example for all the auction forms we have looked at. The next section shows that this is actually a very general result for the basic IPV model.

2 The Main Result

2.1 The Revenue Equivalence Theorem

This section discusses my favourite theorem in economics.

Consider *any* mechanism (any single stage or multi-stage game) for allocating the unit among the n bidders. For this mechanism, and for a given bidder i let $S_i(v_i)$ be the expected surplus that bidder i will obtain in equilibrium from participating in the mechanism, as a function of her type, v_i. Let $P_i(v_i)$ be her probability of receiving the object in the equilibrium.

The following equation is absolutely fundamental:

$$S_i(v) \geq S_i(\widetilde{v}) + (v - \widetilde{v})P_i(\widetilde{v}). \tag{2}$$

The right-hand side is the surplus that player i would obtain if she had type v but deviated from equilibrium behaviour, and instead followed the strategy that type $\widetilde{v}$ of player i is supposed to follow in the equilibrium of the game induced by the mechanism. That is, if type v exactly mimics what type $\widetilde{v}$ would do, then v makes the same payments and wins the object as often as $\widetilde{v}$ would. So v gets the same utility that $\widetilde{v}$ would get ($S_i(\widetilde{v})$), except that in states in which $\widetilde{v}$ would win the object (which happens with probability $P_i(\widetilde{v})$) type v values the object at $(v - \widetilde{v})$ more than $\widetilde{v}$ does, and so v obtains an extra $(v - \widetilde{v})\ P_i(\widetilde{v})$ more surplus in all. In an equilibrium, v must prefer not to deviate from equilibrium behaviour, so the left-hand side must (weakly) exceed the right-hand side.

So, since type v mustn't want to mimic type $v + dv$, we have

$$S_i(v) \geq S_i(v + dv) + (-dv)P_i(v + dv) \tag{3}$$

(this is just (2) with $\widetilde{v}$ substituted by $v + dv$), and since $v + dv$ mustn't want to mimic type v we have

$$S_i(v + dv) \geq S_i(v) + (dv)P_i(v). \tag{4}$$

Reorganising (3) and (4) yields

$$P_i(v + dv) \geq \frac{S_i(v + dv) - S_i(v)}{dv} \geq P_i(v)$$

and taking the limit as $dv \to 0$ we obtain

$$\frac{dS_i}{dv} = P_i(v). \tag{5}$$

Integrating up,

$$S_i(v) = S_i(\underline{V}) + \int_{x=\underline{V}}^{v} P_i(x)dx. \tag{6}$$

Equation (6) gives us a picture like Figure 1.

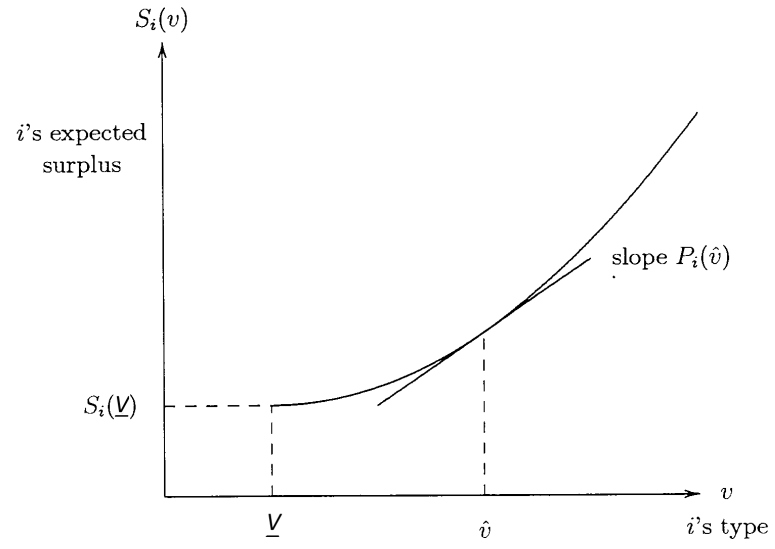


FIGURE 1: Bidder i's expected surplus as a function of her type.

At any type $\widehat{v}$ the slope of the surplus function is $P_i(\widehat{v})$, so if we know where the surplus function starts (i.e. know $S_i(\underline{V})$) we know the entire picture.

Now consider any two mechanisms which have the same $S_i(\underline{V})$ and the same $P_i(v)$ function for all v and for every player i. They have the same $S_i(v)$ function. So any given type of player i makes the same expected payment in each of the

two mechanisms (since $S_i(v) = vP_i(v) - E(\text{payment by type } v \text{ of player } i)$, since the bidder is risk-neutral). So also i's expected payment averaged across her different possible types, v, is the same for both mechanisms. Since this is true for all bidders, i, the mechanisms yield the same expected revenue for the auctioneer. This is the Revenue Equivalence Theorem. There are many different statements of the RET, but they all essentially give the results of this paragraph in a more or less special form.

In particular any mechanism which always gives the object to the highest-value bidder in equilibrium (all the auctions of section 1 do this) has $P_i(v) = (F(v))^{n-1}$ (since a bidders' probability of winning is just the probability that all $(n-1)$ other bidders have lower values then she does), and many mechanisms (including all those of section 1) give a bidder of the lowest feasible type no chance of any surplus, i.e $S_i(\underline{V}) = 0$, so all these mechanisms will yield the same expected payment by each bidder and the same expected revenue for the auctioneer. Since one of these mechanisms is the ascending (Japanese) auction this expected revenue is just the expected value of the second-highest value bidder, and the expected payment of a bidder of type v is just $P_i(v)$ times the expectation of the highest of the remaining $(n-1)$ values conditional on all these values being below v. Since the density of the highest of $(n-1)$ values is $(n-1)f(x)(F(x))^{n-2}$,[10] this last expectation can be written as

$$\frac{\int_{x=\underline{V}}^{v} x(n-1)f(x)(F(x))^{n-2}dx}{\int_{x=\underline{V}}^{v}(n-1)f(x)(F(x))^{n-2}dx}$$

which, after integrating the numerator by parts,[11] yields

$$v - \frac{\int_{x=\underline{V}}^{v}(F(x))^{n-1}dx}{(F(v))^{n-1}}.$$

[10] The easiest way to obtain this formula is to differentiate the cumulative distribution function of the highest of $(n-1)$ values, which is $(F(x))^{n-1}$. (Intuitively, for any x, there are $n-1$ ways of picking one value at x and $n-2$ values below x). You can if you wish find more about order statistics at, e.g., http://probabilityandstats.wordpress.com/2010/02/20/the-distributions-of-the-order-statistics/

[11] The denominator integrates to $(F(v))^{n-1}$, and the numerator $\int_{x=\underline{V}}^{v} x[(n-1)f(x)(F(x))^{n-2}]dx = \int_{x=\underline{V}}^{v} x.d((F(x)^{n-1}) = [x.(F(x))^{n-1}]_{\underline{V}}^{v} - \int_{x=\underline{V}}^{v} 1.(F(x))^{n-1}dx$ which yields the result.

This gives us the following statement of the RET in a form which is quite useful:

Revenue Equivalence Theorem (IPV Case): *Assume each of n risk-neutral potential buyers of an object has a privately known value independently drawn from a common distribution $F(v)$ that is strictly increasing and atomless on $[\underline{V}, \overline{V}]$. Any auction mechanism in which (i) the object always go to the buyer with the highest value, and (ii) any bidder with value $\underline{V}$ expects zero surplus, yields the same expected revenue, and results in a buyer with value v making the same expected payment $(F(v))^{n-1}v - \int_{x=\underline{V}}^{v}(F(x))^{n-1}dx$.*

So for the first-price auction, for example, v bids according to

$$b(v) = v - \frac{\int_{x=\underline{V}}^{v}(F(x))^{n-1}dx}{(F(v))^{n-1}}$$

(since v's expected payments are $b(v)P_i(v)$), and it is easy to confirm that this satisfies the differential equation (1). Using this theorem is often a very much more efficient way to compute bidding strategies than solving for them directly as in section 1.3.[12]

The analysis of this section can be seen as an application of the Revelation Principle that says that we can always restrict attention to direct revelation mechanisms that satisfy incentive compatibility. That is, any mechanism is equivalent to another mechanism in which agents report their types, v, and wish to do so truthfully. Here we have analysed *any* auction by focusing attention on the equivalent truthful direct revelation mechanism. In our problem the incentive compatibility (truth-telling) constraints (2) completely pin down the payments that must be made to each type of agent once $P_i(v)$ and $S_i(\underline{V})$ have been specified.

Notice that nothing about this argument (except the computation of the expected payment) relied on there being only a single object. Thus the theorem extends easily to the case of $K > 1$ objects being sold.[13]

[12] For examples of this see Appendix A to chapter 1, and chapter 2, of my book.
[13] But see the final paragraph of section 3.

The theorem also extends to simple common-value models (see section 4), but it breaks down if bidders are not risk-neutral or have budget constraints, or if bidders' valuations are not independent or are not drawn from identical distributions, or if bidders can collude or their number is endogenous to the auction design.[14]

2.2 Optimal Auctions

The RET shows all the main auction types are equally good under the assumptions of the theorem. This does not demonstrate that they are optimal. In fact, they are (all) optimal under these assumptions, under a mild additional condition and under the additional assumption that the seller must always sell the unit.[15] Note that the additional assumption does preclude the use of a reserve price below which the seller refuses to sell.

There are strong connections between the analysis of how to construct an optimal auction in the general case, including the optimal setting of a reserve price, etc., and the theory of a price discriminating monopolist (see Bulow and Roberts (*JPE*, 1989)). It follows that insights from monopoly theory are useful in auction theory and vice versa.

2.3 Some Review

Sections 1 and 2 show two different ways of solving for equilibrium bidding behaviour. To understand their relationship better, write (with some abuse of notation) $S_i(b, v)$ as type-v-of-bidder-i's surplus when she bids b.

(So $S_i(b, v) = vP_i(v)$ *minus* (a function of b that depends on the rules of the auction).)

[14] It also breaks down if $F(\cdot)$ is not strictly increasing and not atomless, for example if v takes only a finite number of values. Moreover, even in the basic model, the second-price auction's revenue is riskier than the first-price auction's even though the means are the same.

For more discussion see solution to additional problem A2 on my website, and also chapters 1 and 3 of my book.

[15] The result is originally due to Myerson (*Maths of O.R.*, 1981). See also Riley and Samuelson (*AER*, 1981).

Then the first approach is to say that in equilibrium

$$\frac{\partial S_i}{\partial b} = 0$$

which yields an equation for $b(v)$. In section 1.3 we did this for the first price auction where $S_i(b, v) = (v - b)P_i(v) = (v - b)(F(v(b)))^{n-1}$, and obtained a differential equation for $b(v)$.

The second approach (section 2.1) is to use the envelope theorem to write

$$\frac{dS_i}{dv} = \frac{\partial S_i}{\partial v}$$

(to be longwinded, $\frac{dS_i}{dv} = \frac{\partial S_i}{\partial v} + \frac{\partial S_i}{\partial b}\frac{db}{dv}$ but $\frac{\partial S_i}{\partial b} = 0$ when b is chosen optimally).

But $\frac{\partial S_i}{\partial v} = P_i(v)$, since if the bidder's report is unchanged, the incremental utility from a value dv higher is $P_i(v)dv$. So we obtain $S_i(v) = S_i(\underline{V}) + \int_{\underline{V}}^{v} P_i(x)dx$ (the RET) as before, and so for the first-price auction, for example, can now back out what the bidding strategies, $b(v)$, must be, in equilibrium, by making substitutions $S_i(v) = (v - b(v))P_i(v)$, $S_i(\underline{V}) = 0$, and $P_i(v) = (F(v))^{n-1}$.

3 *Extremely* Brief Introduction to Multi-unit Auctions

The material in this section is no longer taught in this course (and so it is NOT examinable).

The two most common auction mechanisms used in the real world, when there are multiple identical units of a good for sale, are natural generalisations of the 1st price and 2nd price sealed-bid auctions for a single object.

In both auctions, bidders can make as many bids (each for a single unit) as they wish.

In a Discriminatory auction (also called a "Pay-your-bid" auction), an auctioneer of N units sells to the N highest bids at the prices actually bid.

In a Uniform price auction, an auctioneer of N units sells to the N highest bids, but charges the highest losing bid, that is, the price of the $(N+1)$st highest

bid, or the "market–clearing price".[16] Observe that, by contrast with a 2nd price sealed-bid auction, a bidder in a uniform price auction does *not* simply "tell the truth", that is, does not simply bid her actual values. The bidder will bid her actual value for the first unit she bids for in a uniform price auction with private values. But she will bid less than her actual values on subsequent units, because reducing her bid on a unit will lower the equilibrium price if the bid turns out to be the $(N+1)$st highest bid, and will in this case save money on all the other units she buys. (The logic is exactly the same as for an oligopolist that faces a residual supply that is not perfectly elastic and therefore chooses a price above its marginal cost–or a quantity that results in a price above its marginal cost–in order to make more money on those units she does sell.) However, the "smaller" the bidder is relative to the size of the market, the less important is this effect, so a sufficiently small bidder almost "tells the truth".[17]

One alternative to these sealed-bid auctions is an ascending auction in which the asking price rises, and each bidder states how many units she wishes to buy at each successive asking price; the auction ends when the total number of units demanded equals the number of units available, and all winning bidders pay that final price for each unit they won. Many spectrum auctions for multiple *heterogeneous* objects use a version of this in which the price rises on each object separately, so different objects can have different asking prices.

Since, it will be clear, different multi-unit auctions generally allocate objects differently among bidders (i.e., the "$P_i(.)$" functions of section 2 are generally different for the different bidders), revenue equivalence will generally not apply.[18] More generally, making analogies between single-unit auctions and multi-unit auctions is dangerous. In particular, although the rules of the discriminatory

[16] In the real world, the price charged is usually the price of the lowest winning bid, that is, the price of the Nth highest bid, but this distinction is unimportant in practice for large N.

[17] A third kind of auction, the Vickrey auction, removes this "demand-reduction" incentive, so is important theoretically, but it is (for good reasons) rarely used in practice.

[18] See solution to additional problem A2(d) on my website for detailed discussion; the main exception is that revenue equivalence *does* apply if each bidder is restricted to a single unit only.

and uniform price auctions are natural generalisations of 1st price and 2nd price sealed-bid auctions, respectively, the behaviour that these rules induce in the multi-unit context is very different from that in the single-unit context.

4 Common Values

The material in this section is no longer taught in this course (and so it is NOT examinable).

4.1 The General Model

Our analysis so far has assumed "private values", that is, each person knows her own value which is unaffected by anyone else's valuation. However, this may not be a very good assumption. For example, when bidding for a painting you may care about how much other bidders like it, because this affects the resale value and the prestige you will get from buying it. When bidding for a oil lease (or the pot of coins that was auctioned in class) the ex-post value may be the same to everyone, but you may care about people's estimates of this value, because knowing those estimates would affect your own beliefs.

In the general model introduced by Milgrom and Weber (*Econometrica,* 1982) each of bidders $i = 1, ..., n$ has a private signal, z_i, about the object for sale, and i's value for the object is $v_i(z_1, ..., z_n)$. We assume v_i is increasing in z_i and weakly increasing in z_j, $\forall j$, so information that makes one player value the good more does not make other players value it less. Thus with pure private values, v_i is a function only of z_i while with pure common values, $v_i(z_1, .., z_n) = v_j(z_1, ..., z_n) \forall i, j, \forall z_1, ..., z_n$. For example, in the "wallet game" discussed in class in which two bidders compete for the combined contents of their wallets, $v_1 = v_2 = z_1 + z_2$ in which z_i is i's observation of the amount in i's wallet.[19]

[19]The "wallet game" was introduced in my 1997 Discussion Paper (with Bulow), and my 1998 *European Economic Review* paper. See also my 2002 *Rand* paper (with Bulow). Now that (especially since the pandemic) people carry less change in their wallets, I play a telephone version instead.

4.2 Bidding in an Open Ascending Auction

We now informally show how to solve for the price $b_i(z)$ that a player with signal z will quit at in a symmetric open ascending auction (modelled as a Japanese auction) with two bidders.[20]

For example, let

$$\begin{cases} v_1 &= 4z_1 + z_2 \\ v_2 &= z_1 + 4z_2 \end{cases}$$

Assuming the equilibrium is symmetric, type z of player 1 will quit the bidding (at price $b_1(z)$) at the same price as a player 2 with the same type, $z_2 = z$. So if type z finds herself a winner at a price (close to) $b_1(z)$, the type she will have beaten will have type (close to) z. So player 1 must be just indifferent between winning and losing at price $b_1(z)$ on the assumption that $z_2 = z_1 = z$. (If instead type z of player 1 strictly preferred to win at price $b_1(z)$, that type would prefer to stay in the auction a little beyond this price and win a little more often; if type z was unhapppy to find herself a winner at this price she should have quit earlier.) But $z_2 = z_1 = z \Longrightarrow v_1 = 4z + z = 5z$, so it follows that $b_1(z) = 5z$.

That is, each player simply bids up to five times her own signal. For example a player with signal 10 bids up to 50. To check this, note that if the player waited and won at a price $50+\varepsilon$ the equilibrium value of the object to her would be only $40 + \left(\frac{50+\varepsilon}{5}\right)$(which is less than $50 + \varepsilon$), since the opponent's signal is one-fifth of what the opponent bids. But if the player quit when the price reached $50-\varepsilon$, the opponent's signal would then be at least $\left(\frac{50-\varepsilon}{5}\right)$, in equilibrium, yielding a value of at least $40 + \left(\frac{50-\varepsilon}{5}\right) > 50 - \varepsilon$, so quitting early gives up the opportunity of making some money (in the states in which the opponent would have quit close to this price).

Note that this equilibrium does not depend on the distributions of the signals, or on bidders being risk-neutral.

[20] Those interested in the general asymmetric case should refer to note 141 of chapter 1 of my book.

It is easy to check that a player's equilibrium quitting price is the same price that he would bid in a second-price sealed-bid auction.

The last result relies on there being just two bidders: the quitting price of the top two bidders depends on the points at which lower bidders have quit in the Japanese auction (and obviously cannot do so in a sealed-bid auction). For example, in equilibrium in the Japanese auction with three bidders, if $v_1 = 4z_1+z_2+z_3$ (and $v_2 = z_1+4z_2+z_3$ and $v_3 = z_1+z_2+4z_3$) each player initially bids up to six times her own signal (the price at which she would be just indifferent about winning if both her opponents quit simultaneously at that price). If the first dropout is at p_1, the two remaining players infer that the quitter's signal was $p_1/6$ and so each of them now bids up to five times her own signal *plus* $p_1/6$.

Finally, in the pure common value case (but not generally otherwise) there are multiple equilibria.[21]

4.3 Revenue Equivalence in the General Model

It is not hard to show (Bulow and Klemperer, *AER*, 1996) that the RET continues to apply in the general model provided bidders' signals are independent. That is, it is the independence of bidders' signals, not the privateness of their values, that is important for the theorem.[22]

Thus we have:

Revenue Equivalence Theorem (General Case) *Assume each of n risk-neutral potential buyers of an object has a privately known signal independently drawn from a common distribution $F(z)$ that is strictly increasing and atomless on $[\underline{Z}, \overline{Z}]$. Any auction mechanism in which (i) the object always go to the buyer with the highest signal, and (ii) any bidder with signal $\underline{Z}$ expects zero surplus,*

[21] In class we look at the unique *symmetric* equilibrium of the "wallet game" but there are other asymmetric equilibria.

[22] Independence is ofter a rather poor assumption in a common-values setting. For example, in the jar of coins game, bidders' estimates of the value will be correlated (because when the jar is relatively full we are all likely to have high estimates; when the jar is relatively empty we will typically have low estimates). However independence might be a reasonable assumption for the "wallet game".

yields the same expected revenue, and results in a buyer with signal z making the same expected payment.

Thus, for example, reconsider the example of the previous sub-section.

$$\begin{cases} v_1 &= 4z_1 + z_2 \\ v_2 &= z_1 + 4z_2 \end{cases}$$

How will players bid in a first-price auction?

Assume, for example, that z_1 and z_2 are independently and uniformly distributed on [0,200] and bidders are risk-neutral. (Recall we didn't need these extra assumptions in section 4.2.) A bidder with signal z wins if her opponent has a signal below z, so conditional on the bidder winning, the opponent's signal is uniformly distributed on $[0, z]$. Now we showed in the previous sub-section that in an ascending auction the opponent would bid up to five times her signal. So conditional on the bidder with signal z winning an ascending auction, the opponent would quit at $\frac{5z}{2}$ on average, that is, the bidder would pay $\frac{5z}{2}$, on average (and her unconditional expected payment would be $(\frac{z}{200})(\frac{5z}{2})$). It follows by the RET that in a first-price auction the bidding strategies must be $b_1(z_1) = \frac{5z_1}{2}, b_2(z_2) = \frac{5z_2}{2}$, so that in equilibrium the expected payment of a bidder with signal z is the same as it is in the ascending auction.[23]

4.4 The Winner's Curse

Again consider the example

$$\begin{cases} v_1 &= 4z_1 + z_2 \\ v_2 &= z_1 + 4z_2 \end{cases}$$

with z_1 and z_2 both independently and uniformly distributed on [0,200]. Consider the case in which $z_1 = 10$. The expected value of z_2 is 100, so the expected value

[23] It is easy to check that if there are other bidders whose values are $\alpha v_1 + (1-\alpha)v_2$ for any α, but who have no information, then the equilibrium is that these bidders do not bid, while the bidders with information bid as in the text. This is an instance of the general principle that bidders who have neither private information nor private value earn no rents. (See Milgrom (*Econometrica*, 1981) and Engelbrecht-Wiggans, Milgrom, and Weber (*Jnl of Mathematical Economics*, 1983) for this result in second-price and first-price auctions, respectively.)

of v_1 is 140. Yet according to the equilibrium we computed, bidder 1 must quit at 50, which is far below even the expectation of z_2. So why can't player 1 stay longer in the bidding?

The reason, of course, is that conditional on 1 winning at a price close to 50, z_2 will be close to 10, so indeed 1 should quit at 50. The point is that what is relevant to bidder 1 about z_2 is the expectation of z_2 *conditional* on 1 winning, that is, the expectation of z_2 conditional on z_2 being lower than z_1. That is, winning the auction is bad news about others' signals; this is called the *winner's curse.* Your estimate of the value of the object must take the winner's curse into account, because only when you win the object do you care about its value.[24]

Exactly the same effect arises in other auction types. In the first-price auction in this example, bidder 1 bids just 25 when her signal is 10 (see section 4.3).

Of course, so far we've assumed equilibrium behaviour in which all bidders recognise and optimally adjust for the winner's curse. In reality naive bidders (like most of you when bidding for the jar of coins) fail to account for the winner's curse and therefore bid too much and lose money.

This may be less of a problem for bidders in an ascending auction (where bidders observe their opponents quitting "surprisingly" low and work out what is going on) than in a first-price auction (when you may only understand the problem when its too late).[25]

It is said that companies who bid for oil tracts initially lost money because they didn't understand the winner's curse.[26]

[24] Another way of understanding why you must drop out "so soon" is to recognise that, in equilibrium, staying in influences the other bidder's behaviour; your staying in makes the other bidder believe your signal is higher than it is, so the other bidder will push the auction price above the object's value before quitting.

[25] But, in an equilibrium, bidders may do worse in second-price than first-price auctions. See, if you wish, Milgrom and Weber (*Econometrica,* 1982), but also section 4.2 of my book.

[26] But there are other stories about what happened, in which some companies first deliberately bid aggressively (and lost money) to develop a reputation that would be helpful in the future. (However, economists are generally biased towards stories in which everybody's behaviour is explained as being rational.)

4.5 Auctions with Almost-Common Values

Consider the "wallet game", very slightly modified so that

$$\begin{cases} v_1 &= z_1 + z_2 + \text{ 1 penny} \\ v_2 &= z_1 + z_2. \end{cases}$$

That is, we have a pure common values game except that player 1 has a very small private-value advantage.

What happens in an ascending auction? The answer is that player 1 can follow a strategy of *never* quitting. The intuition is that if there's any price at which it's rational for player 2 to stay in the game, it must be rational for player 1 to stay in also. Likewise player 2 knows he should never win. So player 2 cannot risk bidding beyond z_2.[27,28] Furthermore, if there is even a trivial cost of bidding, then 2 will not bid at all and 1 can win at a price of zero.

One way of thinking about this equilibrium is that it illustrates "strategic substitutes" in oligopoly.[29] With pure private values, your bidding is unaffected by how aggressive your opponent is. But with a common-values element, the more aggressively your opponent bids, the worse is your winner's curse, so the less aggressively you will want to bid. That is, we have strategic substitutes. And this effect is greater, the greater is the common-values element. With almost-common values, the bidders' reaction curves are not only downward sloping, but

[27] Another way to say this is that if at any point the lowest possible types of 1 and 2 are known to be z_1 and z_2, the lowest type of player 2 will bid up to $z_1 + z_2$ while the lowest type of player 1 will bid up to $z_1 + z_2 + 1$. So whatever are the current lowest possible types, there are always more of 2's types who must quit before any of 1's types leave.

[28] The equilibria of the game are for player 1 to stay in until z_1+(maximum possible value of z_2)+ 1 penny or longer, and player 2 to quit anywhere between 0 and z_2+ 1 penny. (More precisely, these are the Nash equilibria in undominated strategies–there are also "unreasonable" Nash equilibria such as player 2 stays in forever and player 1 quits immediately.) A feature of this game is that given 1 is playing an equilibrium strategy, 2 would lose nothing by deviating to bidding up to the maximum possible value of z_2 plus one penny. But if 2 were to do this, then 1 would not stick to her strategy (when z_1 is small), so it cannot be an equilibrium for 2 to stay in beyond $z_2 + 1$ penny. There are more realistic versions of this game that give a unique strict equilibrium (that is, any player who deviated from equilibrium would strictly expect to lose money) with the same general properties as this one, that is, player 1 wins a great majority of the time and at very low prices. However, I think the current game makes the point most simply. For more discussions see my 1998 *European Economic Review* paper.

[29] The terms "strategic complements" and "strategic substitutes" were invented, and these ideas developed, in Bulow, Geanakoplos and Klemperer (*JPE*, 1985), so I naturally prefer them to Fudenberg and Tirole's alternative, simultaneously-developed, "animals" terminology.

are almost coincident, so a small change in one bidder's reaction curve moves the equilibrium a long way.[30,31]

So a bidder who can increase the direct value to himself of winning the auction, and so credibly commits to bidding more aggressively, gets a strategic benefit (from the effect on his opponent) that is additional to the direct benefit. Increasing your value of winning the auction (or any other strategic move that commits to more aggressive behaviour) can deter entry into an auction, in exactly the same way as strategic entry deterrence works in any "ordinary" economic market.[32]

Many real-life settings are very much like the "wallet game". For example, when the Wellcome drugs company was taken over in 1995, it was probably generally believed that, although the exact value of Wellcome was hard to estimate, it was worth a similar amount to each of Glaxo, Zeneca, and Roche (among other drug companies), but that it was worth a bit more to Glaxo than to any of the others. It is possible that Wellcome cost its shareholders £1 billion or more by running an ascending auction rather than using an alternative form of auction.[33]

What should Wellcome have done? Note that the first-price auction of the "wallet game" above is almost unaffected by the addition of the penny to 1's

[30] A tatonnement logic for the equilibrium is as follows:

if player 2 were to bid as in the symmetric equilibrium of the symmetric game ($v_1 = v_2 = z_1 + z_2$), that is, bid up to twice his signal, then player 1 would get a 1 penny larger prize whenever she wins, so can bid up to 2 x 1 penny = 2 pence more than before for any given signal of hers. [If she wins at price p, she infers $z_2 = p/2$, and so her value is $v_1 = z_1 + p/2 + 1$. She bids until she is indifferent between winning and losing, so bids until $p = v_1 = z_1 + p/2 + 1$, so $p = 2z_1 + 2$, or 2 pence more than she bid originally.]

In this case, player 2 will receive 2 pence ÷ 2 = 1 penny less if he wins, so should quit 2 x 1 penny = 2 pence earlier than before (for any signal of his). [If player 2 wins at p, he infers $z_1 = (p-2)/2$. So his value is now $v_2 = (p-2)/2 + z_2$, that is, 1 penny less than before. Player 2 bids until he is indifferent, so bids 2 pence less than before.]

So player 1 gets another additional 2 pence ÷ 2 = 1 penny whenever she wins, so should bid up to 2 x (1 penny + 1 penny) = 4 pence more than before (for any signal of hers). So player 2 will then actually receive 4 pence ÷ 2 = 2 pence less if he wins, so should quit 4 pence earlier than before (for any signal). So player 1 actually gets an additional 4 pence ÷ 2 = 2 pence whenever she wins, in addition to the 1 penny larger prize, so should actually bid .... so player 2 should quit bidding when the price reaches his signal, z_2.

[31] If you are interested in even more detail, you can look at the closely-related 2011 Advanced Micro exam question and solution on my website.

[32] The analysis exactly parallels Dixit (*Econ Jnl*, 1980)'s classic analysis of entry deterrence in oligopoly (which implicitly assumed strategic substitutes–see Bulow, Geanakoplos and Klemperer, *Econ Jnl*, 1985).

[33] Of course this case is vastly more complicated than the way I've presented it here, and there are arguments in favour of the way the auction was handled.

value. (This should seem intuitive, but it's too hard to prove here.) So the first-price auction is almost revenue-equivalent to the original, symmetric, version of the wallet game. So a first-price auction is often very much more profitable for the seller in this example.

Another possibility is that Wellcome could have tried to "level the playing field" in some way by giving a compensating advantage (perhaps share options) to Glaxo's rivals.[34]

5 Concluding Remarks

5.1 Practical Auction Design

I have emphasised the simplest models, and those that revise the ideas in your game theory course and teach/revise basic oligopoly theory. For issues such as the effects of correlation, risk-aversion, asymmetries, budget constraints, etc., you can, if you wish (strictly optional, of course), see chapter 1 of my book. However what matters most in auction design are the same issues any industry regulator would recognise as key concerns: discouraging collusive and predatory behaviour, and encouraging entry. So good auction design is mostly good elementary economics. It also follows that auctions provide excellent illustrations of simple oligopoly theory and I will give examples in the lecture.[35] Much more discussion is in chapter 3, "What Really Matters in Auction Design", of my book.[36,37]

[34] Another application of the wallet game is that analysis similar to this section's was crucial to the UK government's decision to block BSkyB's proposed takeover of Manchester United football club in 1999. See the Monopoly and Mergers Commission's 1999 report, especially the sections on "toeholds" 2.106-2.117, 2.140, and 4.140-4.152, which were based on my 1999 *JPE* paper (with Bulow and Huang) and my 1998 *EER* paper. (I did not advise any party in this case, and it was decided before I was appointed a Member of the Competition Commission.)

[35] Section 4.5 of these notes gave a good example of entry-deterring behaviour.

[36] For my more recent work on auction design, see www.nuff.ox.ac.uk/users/klemperer/productmix.pdf.

[37] Chapter 4 of the book discusses how to use economic theory, more generally, in practice, and may also be of interest.

5.2 Applying Auction Theory to "Standard" Economics

I will also use my lecture to illustrate the strong connections between auction theory and "standard" economic theory—auction-theoretic tools and intuitions can provide useful arguments and insights in many economic contexts that do not, at first sight, look like auctions. Examples include takeover battles, financial crashes, price discrimination, the value of competition, wars of attrition, monopolistic rationing, comparative legal systems, and oligopoly pricing. For more discussion of this, see chapter 2 of my book.

5.3 Other Issues

Important topics that I have no time to cover include:

–how to design a revenue-maximising auction (see, especially, Myerson (*Maths of O.R.*, 1981), and the interpretation in Bulow and Roberts (*JPE*, 1989));

–multi-unit auctions–these are both important (e.g., in financial markets and energy markets), and also importantly different, from single unit auctions;

–more about practical auction design, and recent advances in auction design.