

ECO 4400 Supplemental Handout: All About Auctions

Tim Salmon
Florida State University

March 2004

1 Introduction

The basic auction problem that this set of notes attempts to solve involves one agent, called the seller, who has a single indivisible item that he wants to try to sell. He possesses some value for the object $v_0 \geq 0$ and is happy selling it to someone else so long as the purchase price, or revenue, r , is greater than v_0 or $r \geq v_0$. He should be able to do better than just v_0 though and since he is greedy, he wants to get as much money as he can. How can he do that? We will be going through the steps to evaluate a number of different alternative mechanisms he might consider and discuss the good and bad points to each.

One possibility is that we go to our principles of micro story of monopoly pricing. If seller knew the real demand curve he faced, then he can solve his problem quite easily. He finds the price that makes Supply=Demand. One unit, so he sets the price equal to the value of the top person on the demand curve, done. Problem is, most sellers won't know the real demand curve. If a seller wants to try this, they must create a sort of expected demand curve based on his beliefs about the distribution of values among the pool of potential buyers. He could set a monopolistic price based on this demand, but it may not net him much money unless his information is very good.

Foreshadowing of a number of results we are going to see, one way of seeing an auction is as a way of discovering that demand curve. In general, auction institutions that are better at revealing the underlying demand curve are going to tend to perform better so this is a really nice image to set in your mind now.

Before moving on to evaluating different mechanisms, there are a few things that must be explained. First, what do we mean by "function better"? A couple of possibilities. When we are thinking of the seller's problem, he cares about one thing: maximizing $r - v_0$, or just maximizing r since v_0 is fixed and can not be changed. If the seller is a government organization, he might care about maximizing efficiency. Efficiency of an auction is computed by calculating the value generated by the auction (in terms of utility to the holders of the objects at the end of the auction), V , and dividing it by the maximum value the auction could have generated, V^* . or V/V^* . An auction that is 100% efficient is one that achieves the maximum possible value. For the single unit auctions that we will be talking about most of the time, a perfectly efficient outcome is one in which the buyer who values the item most receives it at the end. quick example: $v_1 = 10$, $v_2 = 9$. What is efficiency if 2 gets it? 90% as the maximum attainable value is 10, but only a value of 9 was achieved.

Second, what exactly is an auction? It is a set of rules that define how bids are submitted, how those bids determine the allocation (i.e. who gets the item) and how much each participant has to pay based on their bids and the allocation. There are four classic forms.

1. Sealed Bid First Price: All bidders submit bids in sealed envelopes so that other bidders do not know each others' bids at the time of submission, winner is determined by finding the highest bid and payment is decided by having only the winner pay their bid.
2. English Auction (Ascending Clock Auction): Bids are submitted sequentially in an ascending manner until bidding stops, the last person to submit a bid wins the item, he pays what he bid.
3. Dutch Auction (Descending Clock/Price Auction): Price starts high and descends until one bidder declares he will buy, item is assigned to him, he pays what he bid.
4. Second Price Sealed Bid Auction (Vickrey): Again, bidders submit sealed bids, the item is allocated to the one who submitted the highest bid, the winner pays the price of the second highest bid.

2 Optimal Bidding Strategies

The first step in evaluating which of these four auction formats our seller might want to use, involves understanding how bidders will behave in each. That means we need to solve for the optimal or equilibrium bidding strategies for each auction format. In each case we are going to be finding a Bayes-Nash equilibrium of the game. We will take a simple environment and then investigate the optimal bidding strategy under each of the institutions. The environment we will begin with is referred to as the private values case or more formally, the symmetric independent private values (SIPV) case. The characteristics of this environment are

- single indivisible unit up for auction
- each bidder knows his own value for the item and only his
- all bidders are indistinguishable
- valuations are independently and identically distributed (i.i.d.) and are continuous random variables
- all agents are risk neutral

N – number of bidders participating in auction (fixed exogenously)

v_i - is the value of bidder i . For convenience we will simply label bidders such that $v_1 > v_2 > \dots > v_N$

Each v_i is drawn from a uniform distribution on the range $[0, 1]$. This means that bidders have an equal chance of drawing a value anywhere between 0 and 1. Further the probability of drawing a value less than, say .7, is .7. The probability of drawing a value greater than .7 is 1-.7 and that works for any $x \in [0, 1]$.

b_i – The bid submitted by bidder i .

All of these auctions are going to be Bayesian games. So what we need to do to solve them is derive bid functions that map from the type or value space to bids or actions. This means deriving what bid a bidder would place for any possible value they might have.

2.1 English Auction

Let us do the simple one first. What is the equilibrium bidding strategy?

Proposition 1 *In an Symmetric Independent Private Values (SIPV) environment, $b^*(v_i) = v_i$ for all i is a NE of the ascending auction.*

Proof. Think of $b^*(v_i)$ as the point that bidder is going to be willing to drop out of the auction. What this means is that the bidder with the highest value, v_1 , will end up paying a price equal to the point at which the bidder with the second highest value drops out, $b^*(v_2)$. All we need to do is verify that no bidder will do better by altering their bid. Can anyone do better by choosing $\tilde{b}_i(v_i) > v_i$. No. If b_i^* would not have won, but \tilde{b}_i does, then we know the bidder lost money as there was another bid $b_j > b_i^*$ and this is what i will pay. If \tilde{b}_i does not win either, then no help. What about $\tilde{b}_i(v_i) < v_i$? If b_i^* would not have won, \tilde{b}_i certainly won't so no effect. If b_i^* would have won, will \tilde{b}_i help? The only way possible is by lowering amount paid. Will it? No. It will, however, lower the probability of winning. ■

So in the simple SIPV case, the strategy is to just stay in the auction until price hits your value and drop out. NOTE: This is a dominant strategy eq as it is optimal regardless of what your opponents do.

2.2 Vickrey Auction

This auction is defined as follows. Each bidder submits a bid, b_i . The assignment rule is:

$$i^* = \text{Arg max}_i(b_i) \quad (1)$$

The payment rule is

$$p = \max_{j \neq i^*}(b_j) \quad (2)$$

Proposition 2 *In an SIPV environment, $b_i^*(v_i) = v_i$ for all i is a NE of a Vickrey auction.*

It is clear that the same proof as before works here. Also note that this too is a dominant strategy equilibrium. It should be clear to see that the 2nd price and English are what is called strategically equivalent since the same strategy is a NE for both. It should also be clear that the revenue expected by the seller between these two institutions is identical.

How do we find that expected revenue? First, since the bidding strategy is the same for both cases, it should be obvious that the expected revenue will be the same for both. To find what that is, that is done by finding the expectation of the second highest value drawn from the distribution. To do that we use something called an order statistic. In general, finding the order statistic of a distribution is a tricky thing. We will ignore that trickiness, part of why we are using the uniform distribution.

In this case, letting k refer to the order statistic we want or rather to find the expected k 'th highest draw among n values,

$$E[V_{(k)}] = \frac{n - k + 1}{n + 1} \quad (3)$$

The second order statistic is then:

$$E[V_{(2)}] = \frac{n - 1}{n + 1} \quad (4)$$

If we assume that the values of bidders in the English or second price auction are distributed uniformly on $[0, 1]$ then this gives us the expected revenue to the seller. Why does it only depend on n ? If you only have one person, you will get 0 revenue as he does not have to outbid anyone. If you have two, how likely is it that both have a high value, say .9 or larger. The probability of one of them have a value above .9 is .1 and the value of both is $.1^2 = 0.01$. If I have three people, what is the probability that 2 out of the 3 have values greater than .9? This one is more tricky to work out. Think of the people as persons A, B and C. There are three different combinations for which two of these guys have values greater than .9. It could be A and B, B and C or A and C. The probability of two of them having draws above .9 is still .01, but we get three different ways that could happen. That means the probability of at least 2 having values above .9 is $3 * .01 = 0.03$.¹

The idea being that the more people you have, the more likely you are to have two people draw high values and result in a high price. Consequently if we only have two bidders, at least one of them are still likely to have only mediocre a value so the expected revenue is $1/3$. If, however we have 9 bidders, the expected revenue is .8 as there is a much better chance of two bidder having drawn high values.

This fills in a few bits of information in our quest to figure out which mechanism our seller should pick. We know what the optimal bidding strategy is under the English and second price auction. We know that they are strategically equivalent and we know that our expected revenue is $(n - 1)/(n + 1)$.

2.3 Sealed Bid First Price

How do we solve this game? Well first we need to set up the problem.:

$$u_i(B; V) = \begin{cases} v_i - b_i & \text{if } b_i > b(v_j) \text{ for all } j \neq i \\ 0 & \text{else} \end{cases} \quad (5)$$

ignore the possibility of ties and look for a symmetric solution (i.e. all bidders have the same bid function). So what the bidder wants to do is maximize his expected utility:

$$\max_{b_i} (v_i - b_i) Prob(win) \quad (6)$$

I'm not going to go through the general solution to this as it is very complex, but I will explain the intuition which is not so complex and go over a simplified derivation of the solution. In setting up the intuition, we need to look at what the bidder is really trying to do. His first interest is that he wants to win, which means he should bid as high as possible. Problem is that in the event that he does win, he does not want to have bid very much. So there are two components to his expected utility, probability of winning and value of winning and both go in opposite directions based upon an increase or decrease in b_i . If you get too greedy in terms of your value of a win, you lower your probability of winning, but if you raise your bid too high, you make it rather joyless to win. The problem comes down to figuring out how to balance these two things appropriately.

If we know the distributions of values and assume that everyone else is a good game theorist then what we do is take the derivative of this function, set that equal to 0 and solve for b . One detail to that is what the value is for $Prob(win)$? You win if your bid is higher than everyone else's, so this is the probability that your bid is higher than everyone else's. The general version of this is

¹For those who know something about combinatorics, this isn't quite right but it gets the basic idea across.

tough, so let's take a shortcut (that we could solve for if we wanted). Let's assume that everyone is going to bid some constant fraction of their value or $b_i = kv_i$. In that case, if I want

$$\text{prob}(b_1 > b_2) = \text{prob}(b_1 > kv_2) = \text{prob}\left(\frac{b_1}{k} > v_2\right) = \frac{b_1}{k}$$

If we want to look at this generally, then if there are two other bidders, the probability that my bid is greater than both is $\left(\frac{b_1}{k}\right)^2$ and for $n - 1$ other bidders it is $\left(\frac{b_1}{k}\right)^{n-1}$. So, if we assume a total of n bidders the problem becomes

$$\max_{b_i} (v_i - b_i) \left(\frac{b_i}{k}\right)^{n-1} \quad (7)$$

FOC

$$-\left(\frac{b_i}{k}\right)^{n-1} + (v_i - b_i) \left(\frac{b_i}{k}\right)^{n-1} \frac{n-1}{b_i} = 0$$

Solution is:

$$b_i^* = \frac{n-1}{n} v_i = v_i - \frac{1}{n} v_i \quad (8)$$

So what this is saying is that in a first price auction, you do not want to bid exactly your value, because if you do there is no point in winning. What you want to do is to place a bid in which you shade your value or you bid some fraction of your value. Further, the fraction by which you shade your value is dependent only on the number of other bidders. First, let's note what additional bidders do to your bid. For a given value do you bid more or less with additional bidders? 2 bidders $\implies 1/2$, 3 bidders your bid $2/3$ etc. . .so with additional bidders you raise your bid. Why? Because with more bidders around you think there is a better chance of some drawing a high value. If there is only you and some other guy, there is a pretty decent chance that he drew a low value and will be bidding low. If there are 10 other guys, chances are pretty good that someone else drew a higher value. Consequently, you don't want to run the risk of losing that comes from bidding low.

2.4 Dutch Auction

For the dutch auction, the price starts high and descends until the first person jumps in. To to determine a strategy, a bidder must decide at what price he will jump in? How will he decide this? Well, should he jump in at his value? No. He should let the price drop a little more so that he can make some surplus. So he has to trade off surplus against probability of winning and since he wants to maximize his EU we have

$$\max_{b_i} (v_i - b_i) \text{Prob}(\text{win}) \quad (9)$$

looks pretty similar to sealed bid first price auction. These two are strategically equivalent as it is solved in the exact same manner as above yielding the exact same bidding strategy, $b_i^* = \frac{n-1}{n} v_i$.

What we now want to do is be able to compare the revenue between the first two and second two auctions. We therefore need to find the expected revenue for the first price and Dutch. In this case, the revenue is determined not based on the bid submitted by the second highest bidder but based on the expected value of the highest bidder less any shading. Need to compute the expectation of the highest bid. Since $b_i^* = \frac{n-1}{n} v_i$, we need

$$E[r] = \frac{n-1}{n} E[V_{(1)}] \quad (10)$$

taking note of the trick shown earlier that $E[V_{(k)}] = \frac{n-k+1}{n+1}$, we have

$$E[r] = \frac{n-1}{n} \left(\frac{n}{n+1} \right) \quad (11)$$

$$E[r] = \frac{n-1}{n+1} \quad (12)$$

Note that if you flip back a few pages in your notes, what is the expected revenue from a second price or English auction? Same thing. This is no coincidence and is one of the central results in auction theory and one of the most powerful and useful results we will develop this semester known as revenue equivalence.

3 Revenue Equivalence

Theorem 3 *Assume each of N Risk-Neutral bidders has a privately known value that is independently drawn from a common distribution $F(v)$ that is strictly increasing and atomless on $[\underline{v}, \bar{v}]$. Suppose that no buyer wants more than one of the k available identical indivisible objects. Then any auction mechanism in which*

1. *the objects always go to the k buyers with the highest values (or the probability that a bidder with value v wins an item is the same) and*
2. *any bidder with value \underline{v} expects 0 surplus*

Theorem 4 *yields the same expected revenue and results in a buyer with value v making the same expected payment.*

Proof. Consider any mechanism for allocating an item among n bidders. In equilibrium it must be the case the the surplus to a bidder with value v for following the equilibrium strategy is greater than for doing anything else (definition of equilibrium) so

$$S(v) \geq S(\tilde{v}) + (v - \tilde{v})P(\tilde{v}) \quad (13)$$

where $S(v)$ is the surplus obtained by following the strategy for value v . The RHS is the surplus that the bidder would get if he had value v , but followed the strategy for \tilde{v} . $P(\tilde{v})$ is the probability of winning if they bid as if they had \tilde{v} . That is, if his value is \$10 = v , and $S(10) = 5$, we want to know what his expected surplus would be if he chose to bid as if his value were \$40 = \tilde{v} , eq bid of 20? He would get the same surplus a \tilde{v} bidder would get if he won, \$20, + the difference between the two values ($v - \tilde{v} = -30$) times the probability of having to pay up. Another way of writing this is by letting $\tilde{v} = v + dv$

$$S(v) \geq S(v + dv) + (-dv)P(v + dv) \quad (14)$$

If there is someone else who actually has the value $\tilde{v} = v + dv$, it must also be the case that

$$S(v + dv) \geq S(v) + (dv)P(v) \quad (15)$$

or if he is of type $v+dv$ he must prefer choosing $S(v + dv)$ as opposed to what a v bidder would do, $S(v)$. Rewriting and doing some fancy algebra implies

$$P(v + dv) \geq \frac{S(v + dv) - S(v)}{dv} \geq P(v) \quad (16)$$

The first inequality comes from solving first eq for $P(v + dv)$ and second from solving second for $P(v)$.

Take the limit of this as $dv \rightarrow 0$ and we have simply


$$\frac{\partial S(v)}{\partial v} = P(v) \tag{17}$$

The left and right sides approach each other, so middle must be equal. Middle is simply the formula for the derivative of $S(v)$. We can then integrate both sides between \underline{v} and v :

$$S(v) = S(\underline{v}) + \int_{\underline{v}}^v P(x)dx \tag{18}$$

This simply states that the expected surplus of any given value, v , is equal to the surplus of the very base value, \bar{v} , plus some cumulative function based only on the probability of winning. So if we know probability of winning, $P(v)$ then we know the slope of the function and if we know $S(\underline{v})$, we know where it starts. Thus we know every point along the function and it will look something like the graph below. Any two mechanisms that have the same $S(\underline{v})$ and the same $P(v)$, have the same $S(v)$ function. This means that any bidder has the same surplus in any mechanism and the seller expects the same revenue in either mechanism. Note that all of the designs we have discussed have the bidder with the highest value winning with probability one in equilibrium and all will imply that the guy with the lowest value receives a surplus of zero. Thus ■

The proof is really not all that difficult, but let us just look at a graphical proof of it. What this is saying is quite elegant and simple. Lets graph value against expected surplus. What this says is then obvious: If two institutions give the same $S(\underline{v})$ and the probability of winning is constant across both mechanisms given the same v (that is what second condition really states), then what has to be the case is that $S(v)$ goes up in the exact same manner. So all you need to know is $S(\underline{v})$ and $P(v)$ and the surplus as a function of value can be drawn as below. This will then hold for every single mechanism that has the same $S(\underline{v})$ and $P(v)$. If you look across the mechanisms we just talked about: Anyone who drew a value of 0 was losing (and a boundary condition you really need to solve the first price is that $b(0) = 0$) and the guy with the highest draw always wins. Thus revenue will be equivalent across the 4 institutions.



So which auction mechanism should the seller choose? Doesn't matter a bit. Any one will yield the same revenue. This is an incredibly powerful result and useful result. A bit later I will show how it can be used to give us insight into things that would normally be very difficult to analyze and that most people probably think auction theory couldn't tell us anything about.

In practice, this result seems kind of hard to believe. How might it go wrong in certain situations? First thing to note is that if an environment in the real world matches this one, it will hold. If it does not hold in reality (as it does not in many situations), it must be the case that there are differences between our hypothesized environment and real one. Let's propose some modifications to the environment and see if they matter.

4 How to Break RE

In the following examples we will only consider the second price vs first price auctions for simplicity

4.1 Risk Aversion

4.1.1 Risk Averse buyers

Theorem 5 *In a first price sealed bid auction, for any two bidders, if bidder 2 is risk averse and bidder 1 is risk neutral then $b_2(x) \geq b_1(x) \forall x$*

I won't go through the derivation but I will talk through the intuition. This theorem states that a risk averse bidder will bid higher than a risk neutral bidder. To figure that out, let us look first about the first price auction. Notice that what he cares about is $u(v-b)P(\text{winning})$. Consider that our standard RA function looks like $u(x) = x^{.5}$ as opposed to RN of $u(x) = x$. What is happening is that you are now caring less about adding more surplus, so adding marginal surplus doesn't help you. Adding more probability of winning however does help you. So you raise your bid, lower your surplus if you do win, but raise your probability of winning. What risk aversion means in this context is that you are averse to the risk of losing the auction. Another way of stating this is that a bidder being risk averse means that he is willing to accept a lower surplus so long as he has a higher probability of getting it. This means explains exactly why a risk averse bidder would be willing to bid more than a risk neutral bidder as he is perfectly happy making the implied trade-off between declining surplus yet increasing probability of winning.

Now let us consider the second price auction. What happens to our bidder here? Previously his bidding strategy was $b^*(v) = v$. Does that change here? No. If you go back through the previous proof of the optimal bidding strategy there was never any risk or probabilities involved in evaluating strategies so risk aversion does not enter into consideration.

Result is that we know that in the sealed bid case, a risk averse bidder bids higher than a RA bidder. Therefore revenue will be higher with risk averse bidders. We also know that the second price auction will generate the same amount of revenue with RA or RN bidders. Taken together this implies that with RA bidders, the revenue generated by the first price or dutch auction will be higher than that of second price or English. So if you are a seller and think that you potential bidders are RA then you might want to hold a sealed bid auction as you expect to make more money.

4.1.2 Risk Averse Sellers

What if our bidders are RN but our seller is RA? What should our seller prefer?

We won't do a proof but it can be shown that in this case sellers will also prefer the sealed bid first price auction. To show this, you need to examine the variances of the expected selling prices under the different formats. The base issue is that in a second price or English auction, prices are very steady. In the first or dutch, prices vary much more even though on average they come out the same. The key though is that a high price will emerge with higher probability under the dutch or sealed bid first price. For example, consider the case in which someone draws a value of .8 and bids .4 expecting that to be the second highest value. It may be that the other player drew .05. In this case, the ascending or second price auction would have only raise .05 while the first price, .4. In expectation this all works out to be equal, but the tails are what matter for this case.

4.2 Common Values

Another possibility is that the values between bidders may not be independent but related. In the extreme, let us assume that the value of the object is the same for all bidders. With this assumption though, it would be a pretty boring auction if everyone actually knew what that common value was so let us assume further that instead of each bidder knowing the value of the object each bidder receives a signal t_i of the value of the item. These signals are drawn from some common distribution F and there is some v that is the true value of the object and common to all bidders.

Consider a sample auction in which there are 4 people who have drawn estimates of $t_1 = 15$, $t_2 = 11$, $t_3 = 8$, $t_4 = 5$ and the true value is $v = 9$. Who is likely to win the auction? Probably not #4, but it is much more likely to be bidder #1. This will be the bidder who has drawn the highest signal or who has most severely overestimated the value of the item. If the bidder bids in such a way as to not account for this overestimation, he can end up overpaying for the object. This phenomenon is known as the "Winner's Curse". What happens is that the bidder who just found out he won also just found out that he necessarily overestimated the value of the item and probably just lost money if he didn't take these effects into account. Experimental results show that some people can learn to deal with this effect, but it isn't easy.

So how do you bid to get over it? Well, it isn't easy. I'm going to avoid doing the full solution as it is rather messy. Instead, I will explain the bidding strategy in the English auction and what this tells us about bidding strategy in other auctions.

4.3 English Auction

Remember in an English auction, the price starts low and increases until only one person is left. So what every bidder needs to do is figure out, where to drop out. Should a bidder drop out at their signal? Well, if they wait that high, then they may win and have overestimated the value which means they lose money. Perhaps they should go higher than their signal. Why? Well, if they estimated low, then they would be dropping out well below the value of the item. Look at player 4 above. Sounds plausible, but there seems to be a problem with it.

To determine this, we first need to clarify the environment. We will assume that the actual value of the object is

$$v = \frac{t_i}{N} + \sum_{j \neq i} \frac{t_j}{N} \quad (19)$$

We will further assume that t_i is uniformly distributed on $[0, \bar{t}]$. We will also denote the j^{th} highest signal by $t_{(j)}$. This then means that the actual value of the object is just the average of all of the signals received. Yes this is arbitrary and cooked, but it makes the math easier and changes nothing about the substance of the approach. What I am going to do is go through and explain the full equilibrium bidding strategy. What it involves is each of the bidders taking the best information they have available to estimate the true value of the object and being willing to bid up to that point. As the auction goes on, their information gets better and better.

So let us look at bidder N who has signal of $t_{(N)}$, or the lowest signal. How does he decide a dropout point. He estimates the value of the object based upon all of the information he has available which is his own signal. He knows that the other bidders have signals as well so he needs to come up with some belief about what those are. The one he is going to use is that he is going to believe that every one else has seen the same signal he has. If he does then this means his estimate of the value of the object is:

$$\left(\frac{1}{N} + (N-1)\frac{1}{N}\right)t_{(N)} \quad (20)$$

$$\frac{N}{N}t_{(N)} \quad (21)$$

There are two issues with this. First, when auction starts no one knows that they are bidder N , so no one knows to use this. The way this will work is that everyone assumes they are bidder N until proven otherwise by seeing another person drop out earlier. Another way of viewing this gets at the second point which is, why would the person use their own signal as the estimate of the signals of the other players? As the price starts low and goes up the bidder knows who else is in and can assume that the signals of the other bidders are at least as high as the current price. If their signals are higher than his, he is going to lose the auction so he doesn't care about what happens in that case so he assumes that the signals of other people are below his and above the current price. If he doesn't see anyone else drop out below his signal, then his effective estimate of their signals at that point is that they must all be equal to his.

What will happen then, is that everyone uses this to estimate the value of the object and prepares to drop out at their signal of the value. As the price rises, it eventually hits the actual $t_{(N)}$ and the real bidder N drops out.

So the first bidder has dropped. Where does bidder $N-1$ dropout at? All of the bidders now update their estimation of the value of the object. They still know their own signal but now they can also infer what bidder N 's signal was. This gives them 2 out of N data points they need to estimate the value accurately. Since they don't know the others, again they make their best estimate which is that they are equal to their own value. Bidder $N-1$'s estimate and drop out point then is:

$$\frac{1}{N}t_{(N)} + \frac{N-1}{N}t_{(N-1)} < t_{(N-1)} \quad (22)$$

The point to realize is that he is dropping out at a point below his signal. Again, this bidder doesn't know he is bidder $N-1$, but everyone constructs their next possible dropout point like this and as it turns out bidder $N-1$ hits his first. This process continues with the bidder updating their value estimates based on backing out what signal bidder $N-1$ must have had to drop out there and then bidder $N-2$ drops out at some point. The others update their estimate again and this continues on until we get down to the last two bidders. The drop out point for bidder 2 is:

$$\frac{1}{N} \sum_{j=3}^N t_{(j)} + \frac{2}{N}t_{(2)} < v \ll t_{(1)} \quad (23)$$

This is also the final price paid by bidder 1. The important point to see is that this is less than the value of the item and in most cases this will be well below the value for $t_{(1)}$ and $t_{(2)}$ as both are likely to be greater than v .

So the strategy in this case is to use the information conveyed by where other bidders have dropped out to improve your estimation of the value and to drop out where you estimate the value to be. Why, though, should you assume that everyone else who hasn't dropped has the same value as yourself? I tried one attempt at explaining this below, but this ultimately relies on showing that this is an equilibrium strategy by verifying that no one has an incentive to deviate. This is easiest to verify for bidder 2 and it should be obvious that if it works for him, it will work for all others.

The only way 2 can deviate and have it matter to him is if he decides to drop out at a higher price and wins. If bidder 2 drops out at a price under $\frac{1}{N} \sum_{j=3}^N t_{(j)} + (\frac{1}{N} + \frac{1}{N})t_{(1)}$ (where bidder 1 is intending to drop out after the other $N - 2$ bidders have dropped) then he has no change in his outcome as he still loses. If he drops out at a point above this he knows that he is going to pay more than the value of the object, since 1 is following the equilibrium strategy

$$\frac{1}{N} \sum_{j=3}^N t_{(j)} + \frac{2}{N} t_{(1)} > v \tag{24}$$

thus the only way he can win the auction is by bidding above the value of the item. So he won't do it.

The intuition from the private case still holds: you drop out where you think the value of the item is, given your information. The difference is that in the IPV case, you have full info on that matter, in this case you have partial info. The moral of the story is that in this case where you only have partial information, you use what you can to your best ability or you will end up bankrupt.

I won't go through and prove it, but it turns out that revenue equivalence is also going to hold in the common value environment. I won't derive these for you either, but in case you are curious, the bid function in the second price auction is now

$$b = \frac{2 + N}{2N} t_i$$

and for the first price sealed bid the bid function is

$$b(t_i) = \frac{2 + N}{2N} \frac{N - 1}{N} t_i$$

In both of these cases, what you have to do is come up with your best estimate of what the signals are that the other people must have had if yours was the highest. Based on what those signals, in the second price case you place a bid equal to that estimate and in the first price you try to place a bid just above where the person with the second highest signal would be willing to bid. This is essentially the same approach as in the SIPV environment, the only problem is that estimating that point is more difficult.

4.4 Optimal Auction Design

To this point we still have not solved the problem we started with which was what is the actual optimal auction design. The real answer is that the optimal auction design depends on the specifics

of the environment. I will discuss the classic solution for the SIPV environment, but for other environments the answer changes.

In getting to the point of designing an optimal auction we have derived several intermediate steps.

1. We know how to model what bidders are going to do in a number of different institutions.
2. Even better, we know that, at least in certain cases, expected revenue is going to be the same across a broad range of auction forms. <Revenue Equivalence> This means I don't have to check 50 different auction formats for which one makes me the most money.
3. Another fundamental and important result that we haven't talked about but which I will mention is something called the Revelation Principle. What it tells me is that the outcome of any dominant strategy mechanism can be achieved in a direct revelation mechanism for which truth-telling and participation is optimal.

So what does this do? It means that for any dominant strategy equilibrium of any auction game, there is an incentive compatible direct auction that gives us the same outcome (that is an auction where the bid a bidder send in is equal to their true value). So I can set up the game such that I can ask people for bids and get one result. If I can do that then I know that I could also set up a different game and ask them to submit their values, they will, and get the exact same result. How? The basic idea is that you submit those values to an independent third party who then makes optimal choices from the point of view of the bidders and does what they should have done in the original game. For example:

Let us go back to our first price auction in the SIPV RN case. We showed in the case in which values were uniformly distributed on $[0, 1]$ that the optimal strategy was $b_i^* = \frac{n-1}{n}v_i$. So we set up the mechanism as follows. All bidders send in their bids b_i . The bidder with the highest bid will win the item and will pay $p = \frac{n-1}{n}b_i$. It is easy to show that in this case, the mechanism elicits all participants to submit $b_i^* = v_i$.

What points 2 and 3 do for me is make it to where I do not have to spend an eternity checking out every bizarre mechanism someone dreams up to see if it can do better than what I'm about to show you. Since I have 2 and 3, my problem is now very easy (in relative terms).

Deriving this result in full generality is difficult, but again the intuition should be clear. The basic point is that the mechanism ends up looking like the following.

1. Ask bidders to submit a bid (as it will turn out they should submit their true value)
2. Compute what is essentially the marginal revenue, sometimes called priority level, of each bid
or

$$\gamma_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \tag{25}$$

for the uniform distribution case we have been using this is just

$$\gamma_i(v_i) = v_i - (1 - v_i) = 2v_i - 1$$

3. Keep the item if all MR's are below v_0 (value of item to seller).
4. Give the item to the bidder with the highest MR. (works out to be to give the object to the bidder with the highest bid in most cases)

5. Winning bidder pays a price equal to the minimum he could have reported in the place of his true value and still won the auction.

In the case of $F_i(\cdot) = F_j(\cdot)$ for all i and j , then this is just a second price auction with a reserve price, but the reserve price isn't v_0 , it is $\gamma^{-1}(v_0)$. The idea that this is suggesting is that if the true value of the seller is .2, then they should set a reserve price equal to

$$\begin{aligned} .2 &= 2(r) - 1 \\ r &= .6 \end{aligned}$$

and hold a second price auction. If no bids are received above .6, then no sale is made. If only a single bid is above .6, that bidder will pay a price of .6. If two bidders submit bids above .6, then standard second price auction rules hold. This is a curious idea, though, that the seller should actually refuse Pareto improving offers. If a bidder bid .4 and another .3, this would imply a sale price of $.3 > .2$ or the value to the seller. I am arguing he should forego that trade. Why?

This is derived from the seller looking at the expected profit or $(r - v_0) \text{Pr ob}(r)$. If the seller gives up a few transactions between their value and some reserve price, fact is they are not giving up much in expectation. They only lose in those relatively few cases in which there are not two values greater than $\gamma^{-1}(v_0)$ but greater than v_0 . On the other hand for those cases in which there is only a single value above the reservation price, the seller makes more money. In expectation it works out that this is a good trade-off.

5 Legal Systems

I want to take a bit of a diversion from auction theory specifically to show you one of the cooler applications of the RET. What I want to do is look at legal systems and in particular how lawyers are compensated by their clients. In 1991 Vice President Dan Quayle decided that the American legal system needed to be reformed to reduce expenditures on legal costs. Many people have been concerned over how much money is being spent on legal bills and part of that concern is over the fact that it clogs up court systems and makes guaranteeing the right to a “speedy trial” difficult to do. One of his proposals was to change who bears the legal expenses in a lawsuit. He proposed that the losing party pay to the winner an amount equal to the loser’s own legal expenses. The result is that the loser’s legal expenses would be effectively doubled. He based the justification of this argument on a marginal cost basis that if the costs of pressing a lawsuit might cost you double, then less would be spent pursuing lawsuits. What we want to do is determine if his claim was correct.

Analyzing this in full generality gets quite difficult, so we will analyze it in the context of a simple model. Assume that each party in the lawsuit has a privately-known value of winning the lawsuit relative to losing that is independently drawn from a common distribution say from the range $[\underline{v}, \bar{v}]$. What this represents for the plaintiff is what he expects to win, and for the defendant this is more or less the negative of what he expects to lose. Further assume that both parties simultaneously and independently decide how much to spend on legal expenses and then whoever spends the most will win the lawsuit. This is what is known as an all-pay auction. That is bidders place their bids, the highest bidder wins but everyone pays their bid. It is by no means an exact model of the court system, but it really isn’t such a bad one.

What we want to look at is if the regime proposed by Quayle reduces legal expenditures as compared to the standard US system. To do that we need to compare the revenue generated by the

two mechanisms as their bids are essentially what they are paying to lawyers. So if one generates higher revenue, this means it is generating higher legal fees.

We have two systems to analyze, one we will call the American system in which what people pay is what they bid regardless of whether they won or lost. The Quayle system would have people pay twice their bid in the event that they lose and the winner would pay their bid but get back the amount of the bid by the loser. In order to compare these two systems we could go through and find the equilibrium strategy for our players and then use that to compute expected revenue. As you might imagine, actually doing that would be very difficult. The beauty of the RET is that it gives us another option. At least for starting the analysis. We can ask the simple question, will they generate the same revenue in terms of payments to lawyers?

At this point, all we need to do is see if the RET holds. That means verifying two conditions:

1. the winner will be the player with the highest value and
2. any bidder with value v expects 0 surplus

In this case, the guy who draws v can guarantee a surplus of 0 by bidding 0 in either mechanism. He knows he has no probability of winning, so why spend anything on lawyers? Further, there is no reason that the bidder with the lower value would ever be able to outbid the guy with the higher value in eq for either mechanism. The one with the higher value will always be willing to bid more. Consequently, the RET holds and in both institutions the lawyers expect to take the same exact amount of money from people with no real difference and the expected surplus to the actual participants is the same as well. The implication is that Dan Quayle's proposal should not be expected to make any difference.

The key insight as to why is that he left out one point in his analysis. He looked at the losing side and said "If court is more expensive you will go less often." The problem is that he ignored what happened when you win. If you do win, you make even more money than you would under the base American system as the other guy pays part of your bill. As it turns out, these two effects will off set each other resulting in no change in revenue or even in the expected surplus to the participants.

That doesn't mean, however that some form of his idea would not have a change on the system. Others who believe legal expenses and lawsuits have gotten out of hand in this country have proposed adopting a European or British style system. In this system, the loser would pay a fraction, perhaps 25%, of the winners legal expenses. What effect should this be expected to have on legal expenditures? Well, lets try to compare the revenue generated in the two systems and try to figure it out. Again it would be really cool if we could use the RET. One thing we also want to try to get an idea about is which system people would be happiest in.

To apply the RET we would need to show that those two conditions hold. In the British system, however, it is no longer true that the person with the lowest value can guarantee themselves a payment of 0 as when they lose they will have to pay part of the winners legal expenses. They have effectively lost control of their legal payments when they lose in the UK system. This means that the RET will not hold and revenue will be different between the two institutions.

We can, however, still get something out of this. We already showed that people are indifferent between the current American system and the Quayle system. In the proof of the RET, we relied on starting from the lowest value bidder's expected utility and then showing that everyone else's depended only on that and their probability of winning. Well the probability of winning should be the same across all three systems, but the expected value to the lowest value person is lower in the

British system. So what happens is that the expected value graph will start at a lower intercept but then go up in the same slope. This demonstrates quite clearly that everyone is worse off in the British system. People would prefer either the American or Quayle systems.

What is accounting for this preference? It would take a while to prove it formally but it should be relatively clear. What is going on is that in the British system, legal expenses are raised, so lawyers get more money. As in moving to the Quayle system, in the British system, winning becomes more valuable and losing does become more costly but these effects no longer exactly offset. The British system causes people to pay more to lawyers in order to avoid losing and because when you win the other guy subsidizes your legal fees. The idea is that if you win you only have to pay 75% of your legal bills so you are willing to pay more. In fact, some proposals for systems of this sort want the loser to pay 100% of the legal bills of the winner. The equilibrium of this game is that both bidders spend an infinite amount on legal bills as no matter what your opponent pays you want to be \$1 higher.

What this tells us is that if a lawsuit is filed, then legal expenses are higher in the British system, but in reality that doesn't necessarily mean that total receipts should be higher. Here we get to some of the limitations of the simple model I have set up. It ignores the fact that once a plaintiff sees the rules of the system and their expected value from pressing a lawsuit, they might decide not to go to court.

If we looked at this issue, again there should be no difference between the US and Quayle systems. Expected value is the same, probability of winning is the same so if you want to press a lawsuit in one, you do in the other. In the UK system, though, plaintiffs with lower values might not think it is worthwhile to push a lawsuit since their expected value is lower. So the prediction on total legal bills is uncertain.

This does give us other empirically testable predictions that we should see more cases in the US but higher legal bills per case in Britain. This prediction can be tested both experimentally and with field data from US and UK trials comparing size of legal bills and number of cases filed (corrected for population and different legal systems). It works out in both cases that the data support the prediction. Fewer cases are filed in the UK but the per case legal expenses are huge.

So should we adopt the British system? Depends. Since this model isn't allowing the dropout decision, we can't say for sure which one will really do what in terms of total revenue. We can say without that the British system unambiguously reduces welfare (unless the fewer cases allowed cases to go to trial more quickly and that made up for the increased cost). It is quite likely therefore that of the three systems, we are fine with the one we have. There may be better systems out there, but the Quayle system does nothing and the UK system actually increases legal fees and reduces welfare (except for the lawyers).

So this is a very nice clean argument and result that, while based on a simple model, actually does capture what is going on in these various legal systems fairly well. Doing this analysis without the RET would have taken us days, but we could do it all very quickly using this theorem.