1 Introduction

Several auction formats we analyze share the same direct revelation game, at least for certain information sets. Appealing to the revelation principle discussed in the previous chapter, this implies they have the same solution payoffs to each bidder and the auctioneer. We discuss this form of equivalence in Section 5, before introducing revenue equivalence, a weaker form, in Section 6. When bidders are risk neutral, and their valuations are distributed as independently and identically across bidders, then a wide range of auction mechanisms yield the same expected payoffs to the auctioneer and each bidder. This is called the revenue equivalence theorem. Under these conditions, both the auctioneer and the bidders are indifferent between what auction mechanism is used. Indeed it is straightforward to calculate the solution bid functions for any auction satisfying the conditions of the revenue equivalence theorem from the solution to the second price auction (which itself satisfies the conditions).

There are four main reasons why the revenue equivalence theorem might not apply, and the first three of these are discussed in Section 2. When bidders are not risk neutral, say risk averse, or if the private valuations are not identically distributed, then the equivalence may break down. A special kind of asymmetry arises when bidders collude, and this is complicated by the fact that different auctions create different incentives amongst the bidders to collude. Finally the theorem does not apply to common value auctions, where bidders receive different pieces of information about an item they would value the same way if they shared all their information with each other.

Common value auctions are discussed in Sections 3 and 4. We begin with a description of the winner's curse, a pervasive empirical phenomenon that has no rational basis. Then we derive the optimal bidding rules for three different types of common value auctions, and compare the expected revenue generated by each of them. The last topic we investigate, in Section 4, are auctions where the bidders are differentially informed about a common value. In this case uninformed bidders follow a randomized strategy of bidding according to a probability distribution that yields an expected value of zero and also limits the expected information rents of the informed bidder.

2 When does Revenue Equivalence Fail?

There are three main reasons why different types of auctions might not yield the same expected return to the seller: Bidders might be risk averse, not risk neutral. The private valuations of bidders might be drawn from the probability distribution that are not identical. The signals bidders receive about a common value might not be independent. (The revenue equivalence theorem does apply to auctions in which the value that bidders attach to the auctioned item is unknown to them and related, but only if the signals they receive about the value are independent, and they come from the same distribution.)The theorem does not apply when bidders receive signals about the value of the object to them that are correlated with each other.

We investigate these three possibilities in turn. Then we analyze collusion and entry in auctions. We explain why some types of auctions are more susceptible to collusive behavior than others. Finally we ask under what conditions the seller should encourage more bidders by providing information about the value of the auctioned item.

2.1 Attitude towards risk

The revenue equivalence theorem implies that in private value auctions, the bidders and the auctioneer are indifferent between a wide range of auctions if they are risk neutral. For example we showed that the all pay auction, the first price sealed bid auction (which is strategically equivalent to the descending price auction in this case) and the second price sealed bid auction (strategically equivalent to an ascending auction) are all revenue equivalent. But what if bidders are risk averse?

In the case of second price sealed bid auctions with private values, the arguments we used before apply to risk averse bidders in this case. It remains a weakly dominant strategy for each player to bid his or her valuation. Therefore the optimal bidding strategy for the second price sealed bid auction (and also the Japanese and English auctions) is independent of a bidder's attitude towards risk and uncertainty when private values are drawn from a common probability distribution.

The same claims cannot be made for first price sealed bid auctions with private values. In this case, note first that a strategy of bidding your valuation guarantees exactly zero surplus. If you place a lower bid than your valuation your expected surplus initially increases until it reaches the maximum for a risk neutral bidder, and then falls, but the variance of the surplus increases as well. Since a risk averse gambler is willing to trade a lower expected value to reduce the amount of uncertainty, he accordingly bids higher than a risk neutral bidder.

So when comparing first and second price sealed bid auctions, the revenue generated by a second price auction is independent of the bidders' preferences over uncertainty, since bidding is unaffected. Yet the revenue generated by the first price auction is the same as the revenue generated by a second price auction when bidders are risk neutral. Therefore risk averse bidders generate more revenue in a first price auction than they would in a second price auction, and they generate more revenue in a first price auction than do risk neutral bidders.

2.2 Asymmetric valuations

In many auctions where there are private valuations, the bidders have different uses for the auctioned object, and this may be common knowledge to all the bidders. If a particular bidder knows the probability distributions that each of the other valuations are drawn from, he will typically use that information when making his own bid. This in turn affects the revenue equivalence theorem, and also the auctioneer's preferences towards different types of auctions.

We consider an example of asymmetry that illustrates an important consideration. Instead of assuming that all bidders appear the same to the seller and to each other, suppose that bidders fall into two recognizably different classes. Instead of there being a single distribution $F(v)$ from which the bidders draw their valuations, there are two cumulative distributions, $F_1(v)$ and $F_2(v)$. Bidders of type $i \in \{1,2\}$ draw their valuations independently from the distribution $F_i(v)$.

Let fi(x) denote the probability density function of Fi(x).

Consider a first price auction where there are only two bidders.

The private valuation of the first bidder is drawn from a probability distribution $F_1(v)$ that stochastically dominates the probability distribution for the other probability distribution $F_2(v)$. In fact we make a stronger assumption, that for all v

$$\frac{F_1'(v)}{F_1(v)} > \frac{F_2'(v)}{F_2(v)}$$

Then $b_1(v) < b_2(v)$. The intuition is to bid aggressively from weakness and vice versa. Suppose each bidder sees his valuation, but does not immediately learn whether he comes from the high or low probability distribution. At that point the bidding strategy cannot depend on which probability distribution the valuation comes from. Then each bidder is told which probability distribution his bid is drawn from. How should he revise his bid? The second (first) bidder learns that the first (second) bidder is more likely to draw a higher (lower) valuation than himself, realizes the probability of winning falls (rises), so adjusts his upwards (downwards).

Since $b_1(v) < b_2(v)$, the possibility arises that the second bidder will win the auction with a lower valuation than the first. In other words the first price sealed bid auction is inefficient, because aggressive bidding by the second bidder, and low balling by the first bidder, sometimes lead to the second bidder winning the auction with a lower valuation than the first. We may contrast this outcome with a second price auction where it remains optimal to bid your valuation, implying that the auction is always won by the bidder with the highest valuation.

2.3 Collusion

Collusion is a special form of asymmetry. Imagine we start out with N independently and identically distributed valuations, we suppose that C of them form a cartel and submit only one bid, or more commonly, agree on who should submit the highest bid. We compare agreements with second form of cartelization, entry deterrence. Whether this is voluntary or enforced is also a question of importance. (In this case the number of bidders falls. A slightly different analysis, since they might be keeping out medium versus low bidders. this is less effective.) Note that a cartel action only has effect if a member would have won anyway, and also continues to win despite the fact that a lower bid has been tendered. With reference to Figure 16.1, cartels are less likely to be effective if there are a large number of bidders. And in the case of a small number of bidders, there are rents to be had to winning bidders regardless of whether there is a cartel or not.

Up until now we have taken as a given that the bidders will follow the rules that characterize the auction. This is not necessarily a reasonable assumption, and some the rules of some auctions may be easier to enforce than others. How easy is it to enforce a collusive practice? In a second price auction or English auction the onus is on the lower valuation bidders to essentially withdraw form the auction. To break any cartel arrangement they must bid above the highest valuation person. his implies that would In a first price auction, the bidder only has to bid above the low value that the highest valuation person submits by submitting a slightly higher value.

3 Common Value

This section of this chapter begin our investigation of common value auctions with a discussion of the winner's curse, a pervasive phenomenon in empirical studies that has no theoretical basis in optimization and equilibrium solutions. We describe how the winner's curse arises, its magnitude show the losses it generates, and thus provide some heuristics for avoiding it. Then we derive the optimal bidding rules in three different types of common value auctions. Specifically we concentrate on the first price sealed bid (and descending)auctions, the second price sealed bid (and limited information ascending) auctions, and the Japanese (or full information) ascending auction. this naturally leads us to an examination of which auction yields the most revenue. We show that the auctions listed in order are ranked from the lowest to the highest. These results are also sensitive to the structure for the same reasons that revenue equivalence breaks down in auctions with private value case, asymmetrically distributed signals (including collusion) and risk aversion

Throughout the previous two chapters we maintained the assumption that each bidder knows how to value the object. We relax this assumption in this third chapter on auctions, and explore how that affects the conclusions we have drawn thus far. We proceed in stages. First we relax the assumption that relaxes the assumption that valuations are drawn independently, and asks the same set of agenda questions we laid out in Chapter 18. Much of our discussion will revolve around another extreme case. Instead of assuming values are independently distributed, we shall assume they are exactly the same, and denote the common value by v. Naturally this is not at all an interesting avenue to explore if all the bidders know v. We shall accordingly assume that each bidder has some private information about the common value, but not all of them know it exactly. of Our discussion is accordingly divide into three parts. under what circumstances the revenue equivalaffiliated signals common value with independent and identically distributed signals.

We will assume that

$$u_n(s_1, \ldots, s_N) = u(s_n, s_1, \ldots, s_N)$$

Notice that there are two assumptions embodied in the definition of symmetric valuations. The fact that the arguments can be interchanged means that each bidder does not care about the permutation of the other signals. Noting that the mapping on the right side is not subscripted by $n,$ all bidders have the same utility function

Common value but independent signals

There are many situations in which bidders all know one characteristic of the object, and each bidder has some private information about the other characteristics. For example, suppose the common value of the object is the sum of independent signals received by bidders:

$$v = \sum_{k=1}^{N} s_k$$

Each signal might represent part of the value of the object that each bidder is privy to. In this case the revenue equivalence theorem applies.

Suppose the value of the object to each bidder is the same, but this value is unknown to each bidder. We assume the n^{th} bidder receives a signal s_n which is distribute about the common value v, and write

$$s_n = v + \varepsilon_n$$

where

$$\varepsilon_n \equiv E[v|I_n] - v$$

In general one might expect that ε_n might not be independent of all $\varepsilon_m.$ For example if v is stochastically evolving and bidders have the opportunity to sequentially observed its quality, one might suppose that those bidders who have more recently inspected the object have more precise knowledge.

A slight generalization of this is a model in which each bidder places more significance on their own draw, but does attach some value to the assessments of others too.

$$v_n = s_n + \frac{\alpha}{N} \sum_{k=1}^{N} s_k$$

3.1 The winner’s curse

When other bidders have information that you lack about the value of the object for sale, winning the auction may cause you to decrease your conditionally expected value of the object. Failure to take into account the bad news about others’ signals that comes with any victory is called the winner’s curse. The winner’s curse describes the fact that winning an auction may convey new and unfavorable information about the item. Because all other bids are less than the winning bid, the expected value of the item to the winning bidder might fall when the outcome of the auction is announced.

To fully explain the winner’s curse, it is useful to draw upon the notation we have already developed. As above, consider an object worth v_n to the n^{th} bidder, but this value is unknown to the bidder until after the auction is over. Before the auction begins, suppose the bidder receives a signal s_n about the value v_n. Without loss of generality, one can order the set of possible signals so that higher valued objects tend to generate more favorable signals, or more precisely that $E[s_n|v_n]$ is an increasing function in v_n. The expected value of v_n given the signal s_n is expressed as $E[v_n|s_n]$. We also make the standard assumption that higher bids are induced by better signals, or the bidding function, denoted $b(s_n)$, is increasing in s_n, which is what one would expect in equilibrium.

Under that assumption, if the n^{th} bidder wins the auction, he will realize his signal exceeded the signals of everybody else, that is $s_n \equiv \max\{s_1, \ldots, s_N\}$, so he will condition the expected value of the item on this new information. His expected value is now the expected value of v_n conditional upon observing the maximum signal:

$$E[v_n|\max\{s_1, \ldots, s_N\}]$$

This is the value that the bidder should use in the auction, not $E[v_n|s_n]$, because he should recognize that unless his signal is the maximum he will receive a payoff of zero. The winner's curse is defined as:

$$W(s_n) \equiv E[v_n|s_n] - E[v_n|\max\{s_1, \ldots, s_N\}]$$

Since the $\max$ operator is a convex increasing function of its arguments

$$\begin{aligned} E[v_n|\max\{s_1, \ldots, s_N\}] &\geq \max\{E[v_n|s_1], \ldots, E[v_n|s_N]\} \\ &= \max\{E[v_n|s_1], \ldots, E[v_n|s_N]\} \end{aligned}$$

it follows that $W(s_n)$ is a positive function. Although bidders should take the winner's curse into account, there is widespread evidence that novice bidders do not take this extra information into account when placing a bid.

Notice how the role of incentives and plays a role with the nature of inference in this game. If we changed the auction rules and award the object to the person who is closest to the true value of the object, then each bidder would minimize (?)

$$E\left[(b_n - v)^2|s_n\right]$$

and bid the solution to that problem, which is $E[v_n|s_n]$.

To gauge the importance of the winner's curse, it is useful to worth calculating analytically for some probability distributions where this is feasible.

3.2 Symmetric Solutions

We will look at the symmetric solutions.

Japanese auction

In a Japanese (ascending) auction, bidders remaining in the auction observe who has dropped out and at what point. For this reason they are able to infer the valuation of the bidders as they withdraw from the auction.

Consider the first bidder to withdraw from the auction. At that point the only way he can win is for everybody to withdraw simultaneously, which would only occur in a symmetric equilibrium if all bidders received the same signal. In that case the winner of the auction would receive a utility of $u(s^{(N)},\dots,s^{(N)})$. WE consider the properties of the strategy profile defined by the N reservation price functions:

$$b^{(N)}(s^{(N)}) = u(s^{(N)},\dots,s^{(N)})$$

Now set:

$$b^{(N-1)}(s^{(N-1)},s^{(N)}) = u(s^{(N-1)},\dots,s^{(N-1)},s^{(N)})$$

and more generally

$$b^{(k)}(s^{(k)},s^{(k+1)},\dots,s^{(N)}) = u(s^{(k)},\dots,s^{(k)},s^{(k+1)},\dots,s^{(N)})$$

We note first that the strategy reveals the valuations of bidders who withdraw from the auction, allowing the remaining bidders to condition on the value of the object with this new information.

This strategy is the unique symmetric Nash equilibrium solution to the auction game. We first show that if players have followed the strategy so far then the proposed strategy weakly dominates setting a higher reservation price.

Instead of bidding $b^{(k)}(s^{(k)},s^{(k+1)},\dots,s^{(N)})$, suppose the bidder with the k^{th} highest valuation bids $b^{(1)}(s^{(k)}+\Delta,s^{(2)},\dots,s^{(N)})$. The only way this deviation might affect her payoff is if she consequently wins the auction. In that case

$$s^{(k)}+\Delta \geq s^{(1)} \geq \dots > s^{(k)}$$

and the bidder with the k^{th} highest valuation would pay the reservation price of the highest valuation bidder, namely

$$b^{(2)}(s^{(1)},\dots,s^{(k-1)},s^{(k+1)},\dots,s^{(N)})$$

to obtain an item which yields a utility of

$$u(s^{(k)},s^{(1)},\dots,s^{(k-1)},s^{(k+1)},\dots,s^{(N)})$$

Since $u(s_1,s_2,\dots,s_N)$ is increasing in its first argument

$$\begin{aligned} u(s^{(k)},s^{(1)},s^{(2)},\dots,s^{(k-1)},s^{(k+1)},\dots,s^{(N)}) &< u(s^{(1)},s^{(1)},s^{(2)},\dots,s^{(k-1)},s^{(k+1)},\dots,s^{(N)}) \\ &= b^{(2)}(s^{(1)},s^{(2)},\dots,s^{(k-1)},s^{(k+1)},\dots,s^{(N)}) \end{aligned}$$

This establishes that it is not profitable by deviating from the strategy profile by bidding more.

Proving claim about equilibrium is completed by showing setting a reservation price lower than the proposed limit is not a best response to the other bidders playing this strategy profile. If the highest valuation bidder wins she pays $b^{(2)}(s^{(2)},\dots,s^{(N)})$, which does not depend on her own reservation price. Accordingly suppose the highest valuation bidder reduces his reservation price below $b^{(N)}(s^{(1)},\dots,s^{(N)})$, and consequently loses the auction, by dropping out as before the bidder with the k^{th} highest valuation, meaning

$$s^{(1)} > s^{(2)} > \dots > s^{(k)} > s^{(1)} - \Delta$$

Then the utility foregone is greater than the bid because

$$u(s^{(1)}, s^{(2)}, \ldots, s^{(N)}) > u(s^{(2)}, s^{(2)}, \ldots, s^{(N)})$$
$$= b^{(2)}(s^{(2)}, \ldots, s^{(N)})$$

Therefore the only effect from setting a lower reservation price than the one proposed by the formula is to risk losing the auction when it is profitable to win.

English (second price sealed bid) auction

Here we define $\lambda(s, m_n)$ as the expected value of $u(s_1, \ldots, s_N)$ conditional on the bidder's own signal s, and conditional on the highest signal from the $N-1$ other bidders being m_n.

$$\lambda(s, m_n) = E\left[u(s_1, \ldots, s_N) \middle| \begin{array}{l} s_n = s, \\ m_n = \max\{s_1, \ldots, s_{n-1}, s_{n+1}, \ldots, s_N\} \end{array} \right]$$

Note first if there are only two bidders, meaning $N = 2$, then the arguments developed in our discussion of the Japanese auction imply that each bidder sets a reservation price of $u(s_n, s_n) = \lambda(s, m_n)$. More generally, the solution to a second price sealed bid auction with N bidders is to set a reservation price of $\lambda(s_n, s_n)$. We show that this is a symmetric Nash equilibrium. Suppose everyone else is following this strategy. Then the value of setting a reservation price at b is found by integrating the expected value of the object upon winning the auction when your value is v, and the second highest value is μ, after subtracting payment $\lambda(\mu, \mu)$, and integrating over the probability distribution of the second highest valuation in the region $[\underline{v}, b]$

$$\int_{\underline{v}}^{b} [\lambda(v, \mu) - \lambda(\mu, \mu)] G'(\mu|v) d\mu$$

The optimal value of b is deduced from the fact that $\lambda(v, \mu)$ is increasing in μ, and therefore $\lambda(v, \mu) > \lambda(\mu, \mu)$ if and only if $v > \mu$. Consequently the solution to the optimization problem is determined by its unique stationary point. Differentiating with respect to b we directly obtain the first order condition $\lambda(v, b) = \lambda(b, b)$, which, upon solving for b, establishes the claim.

It now follows that the winner of the auction pays the reservation price, or sealed bid, of the bidder with the second highest valuation, namely

$$\lambda(v^{(2)}, v^{(2)}) = E[u(v^{(2)}, v^{(2)}, V^{(3)}, \ldots, V^{(N)}) | v^{(2)} > \max\{V^{(3)}, \ldots, V^{(N)}\}]$$

Exercise *Conduct some English auctions*

1. *Regress the winning bids on the two highest signals*

Dutch (first price sealed bid) auction

The third type of auction we study with a symmetric information structure underlying valuations is a Dutch or a first price sealed bid auction. As before we suppose that there exists a symmetric equilibrium with strategies $b(v_n)$ that are differentiable with respect to v_n, the n^{th} bidder's valuation for each $n \in \{1, \ldots, N\}$.

To derive this equilibrium strategy, suppose the n^{th} bidder's valuation is v_n but she bids $b(v^*)$ instead of $b(v_n)$. Then she wins the auction if $v^* > \mu$, which is the event $1\{v^* > \mu\}$. If she wins the auction, and the valuation of the second highest bidder is μ, her expected utility is $[\lambda(v_n, \mu) - b(v^*)]$. Denote by The

$$\int_{\underline{v}}^{\bar{v}} [\lambda(v_n, \mu) - b(v^*)] 1\{v^* > \mu\} G'(\mu|v_n) d\mu$$
$$= \int_{\underline{v}}^{v^*} [\lambda(v_n, \mu) - b(v^*)] G'(\mu|v_n) d\mu$$

The expression above shows gives an expression for the expected utility from bidding $b(v^*)$ instead of $b(v_n)$. But if $b(v_n)$ is a Nash equilibrium strategy Accordingly we differentiate the expression to obtain the first order condition

$$\begin{aligned}[\lambda(v_n, v^*) - b(v^*)] G'(v^*|v_n) &= b'(v^*) \int_{\underline{v}}^{v^*} G'(\mu|v_n) d\mu \\ &= b'(v^*)[G(v^*|v_n) - G(\underline{v}|v_n)] \\ &= b'(v^*) G(v^*|v_n)\end{aligned}$$

because $G(\underline{v}|v_n) = 0$. But if $b(v_n)$ is a Nash equilibrium strategy then

$$[\lambda(v_n, v_n) - b(v_n)] G'(v_n|v_n) = b'(v_n) G(v_n|v_n)$$

Supposing that the minimum acceptable bid is $\underline{v}$, it follows that $b(\underline{v}) = \underline{v}$, meaning that a bidder receiving the lowest valuation bids the minimal acceptable bid, then the differential equation can be solved to obtain

$$b(v_n) = \int_{\underline{v}}^{v^*} \lambda(\mu, \mu) dL(\mu|v_n)$$

where $L(\mu|v_n)$ is defined as

$$L(\mu|v_n) = \exp\left[-\int_{\mu}^{v_n} \frac{G'(t|t)}{G(t|t)} dt \right]$$

The easiest way to check this is a solution is to differentiate $L(\mu|v_n)$ with respect to v_n and substitute it back into the first order equation above.

3.3 Revenue comparisons

In private value auctions we proved a revenue equivalence theorem that showed, amongst other things that the expected revenue to the auctioneer generated by the first price sealed bid, the second price sealed bid and the Japanese auction is the same. This result does not extend to symmetric auctions where the bidders valuations are dependent.

First price versus second price

The second price sealed bid auction yields more revenue than its first price counterpart.

Second price versus ascending

The expected revenue from a Japanese auction is

$$E[b^{(2)}(v^{(2)},\ldots,v^{(N)})] = E[u(v^{(2)},v^{(2)},v^{(3)},\ldots,v^{(N)})]$$

Whether the second price sealed bid or Japanese auctions yields more revenue or not turns on the sign of

$$E[u(v^{(2)},v^{(2)},V^{(3)},\ldots,V^{(N)})|v^{(1)} = v^{(2)} > \max\{V^{(3)},\ldots,V^{(N)}\}]$$
$$-E[u(v^{(2)},v^{(2)},V^{(3)},\ldots,V^{(N)})|V^{(1)} > v^{(2)} > \max\{V^{(3)},\ldots,V^{(N)}\}]$$

In words, do the bottom $N-2$ valuations $(V^{(3)},\ldots,V^{(N)})$ tend to rise with the top valuation $V^{(1)}$, thus raising expected utility when we fix the second highest valuation $V^{(2)}$ at $v^{(2)}$? If so, we can show that the expected revenue from the Japanese auction is higher than the expected revenue from the second price sealed bid auction. Noting the expected revenue from the latter is $E[\lambda(v^{(2)},v^{(2)})]$, the inequality implies Denote the expected revenue from the second price sealed bid auction by $r_2(v)$ and the expected revenue from the Japanese auction by $r_3(v)$. Then

$$\begin{aligned}
&E\{E[\lambda(v^{(2)},v^{(2)})|V^{(2)} = v^{(2)}]\}\\
&= E\{E[u(v^{(2)},v^{(2)},V^{(3)},\ldots,V^{(N)})|v^{(1)} = v^{(2)} > \max\{V^{(3)},\ldots,V^{(N)}\}]|V^{(2)} = v^{(2)}\}\\
&\leq E\{E[u(v^{(2)},v^{(2)},V^{(3)},\ldots,V^{(N)})|V^{(1)} > v^{(2)} > \max\{V^{(3)},\ldots,V^{(N)}\}]|V^{(2)} = v^{(2)}\}\\
&= E[u(V^{(2)},V^{(2)},V^{(3)},\ldots,V^{(N)})]\\
&= E[b^{(2)}(V^{(2)},V^{(3)},\ldots,V^{(N)})]
\end{aligned}$$

and the result is proved.

4 Differential Information

Symmetry, an assumption we have maintained thorough much of this chapter, can be violated in many ways. We have already discussed how the revenue equivalence theorem breaks down when private values are drawn from different probability distributions. We now analyze the effects of differential information in a first price sealed bid common value auction. We consider the extreme case, where one bidder receives a perfect signal that reveals the common value, and where the other bidders do not receive any signal at all.

Suppose uninformed bidders always makes the same positive bid, denoted b. This is an example of a pure strategy. Is this pure strategy part of a Nash equilibrium? The best response of the informed bidder is to bid a little more than b when the value of the object, denoted v, is worth more than b, and less than b otherwise. Therefore the uninformed bidder makes an expected loss by playing a pure strategy in this auction, because he incurs a loss whenever he submits the highest bid. The pure strategy of bidding below $\underline{v}$ is not in equilibrium either. If that were the case, then the best response of the informed bidder would be to bid $\underline{v}$ at most and win the auction every time, garnering a rent of at least $E[v]-\underline{v}$. But if the informed player bids $\underline{v}$, then the best response of the uninformed bidder is not to bid a price below $\underline{v}$, but some price above the informed bid, but less than the unconditional mean $E[v]$. This proves that there is no equilibrium in which the uniformed player chooses a pure strategy. These arguments can also be extended to auctions where is more than one uninformed agent. To summarize, uninformed players in sealed bid first price auctions must make random bids to avoid being exploited by the informed players.

Notice that for risk neutral bidders, if you do not know your value v for the object, then you should form a subjective probability distribution for this random variable. You could then compute its expected value $E[v]$. Knowing $E[v]$ just as good as knowing v itself if no one else knows anything about v that you do not know, but not if your competitors have some information about v that you do not have.

The argument above shows that the uninformed bidder plays a mixed strategy in this game. One can show that in equilibrium the informed bidder bids according to the strategy of treating the auction like a private value auction when calculating the informed bidder's equilibrium submission, and that the uninformed bidder should bid according to the unconditional distribution of the informed bidder.

In equilibrium, we can prove the informed bidder chooses

$$b(v) = E[\mu|\mu \leq v]$$

and the uninformed bidder chooses a bid at random from the interval $[\underline{v}, E[v]]$ according to the probability distribution H defined by

$$H(b) = \Pr[b(v) \leq b]$$

If the informed bidder bids $b(v)$ then the uninformed bidder follows the prescribed strategy then the probability that the informed bidder wins is

$$\Pr[E[\mu|\mu \leq v] \leq b] = H(b)$$

For consider the problem facing the informed bidder, when the uninformed bidder follows the mixed bidding strategy prescribed for him. If the informed bidder with valuation v submits $b(v^*)$ instead of $b(v)$, his expected net payoff is

$$H(b(v^*))(v - b(v^*)) = F(v^*)(v - b(v^*))$$

Differentiating with respect to v^* we obtain

$$\frac{d[F(v^*)b(v^*)]}{dv^*} = F(v^*)v$$

At the stationary point $v = v^*$. Therefore for all v opitmality requires:

$$\frac{d}{dv}[F(v)b(v)] = F(v)v$$

Integrating both sides of the equation with respect to v from $\underline{v}$ to any $v_n \in [\underline{v}, \bar{v}]$ we obtain:

$$F(v_n)b(v_n) = \int_{\underline{v}}^{v_n} F(v)vdv$$

or

$$b(v_n) = \frac{1}{F(v_n)}\int_{\underline{v}}^{v_n} F(v)vdv = E[\mu|\mu \leq v]$$

as required.

With regards the uninformed bidder, we now show that the payoff from bidding any value in the range $[\underline{v}, E[v]]$ yields the same expected payoff, because bidding any $b_2 \in [\underline{v}, E[v]]$ yields

$$\begin{aligned} E[v|b(v) < b_2] - b_2 &= E[v|v < b^{-1}(b_2)] - b_2 \\ &= b[b^{-1}(b_2)] - b_2 \\ &= b_2 - b_2 \\ &= 0 \end{aligned}$$

Note that bidding more than $E[v]$ would imply that he wins the auction with probability one, for an expected payoff of $E[v]$. Therefore it is not optimal for the uninformed bidder to defect from the proposed strategy by bidding less than $\underline{v}$ or more than $E[v]$.

5 Summary

This chapter explored the reasons why revenue equivalence fails. There are four conditions that are used to prove this result, and we discussed each of them in turn. The fourth, auctions with dependent signals, lead into a more general discussion of equilibrium bidding rules for symmetric auctions, and how their expected revenues are ranked . Finally our discussion or differential information in auctions shows how mixed bidding strategies is likely to be endemic to the bidding process when the symmetry assumption is violated.