

Static games with incomplete information*

1 Introduction

The next step in games is to consider a case where one (or both) of the players has *incomplete* information. Recall that *imperfect* information deals with not knowing which node of a game a player is at. With incomplete information, there is some uncertainty over the other player's payoff function. A classic example is an auction. In a fairly standard auction setting each bidder knows his own private value but not the value of the other bidders. All bidders have the same payoff function, which is their value minus the price they pay for the object (this is simply a calculation of consumer surplus), but since bidders only know their own value they have incomplete information about the other bidder's actual payoff. Contrast this with the case of Cournot competition we recently considered. All firms knew the demand conditions as well as the cost conditions for all firms. Thus, for any strategy choices by the firms both firms knew exactly what the payoff would be to the other firm. In an incomplete information setting this knowledge of the payoff function is removed.

It is easiest in this setting to talk about the players having a possibility of having different *types*. A type simply represents the possible state of the world that a particular player could be. In an auction setting, there might be 100 different values that a bidder *could* have for an item. Thus, there are 100 possible types that the bidder could have. In a Cournot competition setting there might be 2 different cost functions a firm *could* have, which yields 2 types. These types will be assigned by some probability distribution, which we will treat as known to the players.¹ In an auction setting a bidder's value might be drawn from the uniform distribution between 0 and 100, while in the Cournot competition game a firm may have a probability of θ of having a high cost and a probability of $1 - \theta$ of having a low cost. Since each player has multiple types that he could be we will need a new solution concept.

2 Bayes-Nash equilibrium

In our simultaneous game of complete information we specified a normal form game as $\Gamma_N = [I, \{S_i\}, \{u_i\}]$, where the game consisted of I players, a strategy space S_i for each player, and a payoff function u_i for each player where the payoff is contingent upon the strategies chosen by all of the players (the strategy profile). Now we have two additional features to incorporate, one of which is a player's type space, T_i , as well as a player's belief space, which we denote by p_i . In the games we will initially consider, a specific player's type will yield no information about the other player's type – thus, players' types are independent of one another. Thus, the belief space for these initial game is simply the exogenously given probability distribution of player types. Also, because these games are simultaneous there are no observed actions that a player can use to update his belief about another player's type.² A static game of incomplete information is then represented by $\Gamma_{BN} = [I, \{S_i\}, \{T_i\}, \{p_i\}, \{u_i\}]$, where $\{T_i\}$ represents the type space of all I players and $\{p_i\}$ represents the belief space of all players.

The question now becomes what constitutes an equilibrium in games of incomplete information? Like the other games, an equilibrium will be a set of strategies such that neither player can unilaterally change

*Based on Chapter 3 of Gibbons (1992).

¹It is possible that there are multiple probability distributions from which a player's type could arise and that the other player is uncertain about which of these probability distributions is the correct one. This simply adds an additional layer of complexity to the problem.

²We cover dynamic games of incomplete information in the next set of notes.

his strategy and make himself better off. However, now a strategy consists of actions that must be specified for every *possible* type a player could have. One might ask why a player has to specify an action for every possible type since the player knows his own type. The answer is that while a player A knows his own type the other player B does not, so player A has to consider that player B only knows the probability that player A has a certain type, and not the specific type. This concept is similar to players having to specify their actions in an extensive form game for nodes that are not going to be reached when the game is actually played. The players still need to specify which actions they will take at those nodes because that can impact the game. The same is true for types that are “never reached”. The actions taken by a player for a type that he is not could affect the action choices of the other player. Thus, we now define a strategy for a player i in game Γ_N in terms of the possible types player i could be.

Definition 1 A strategy in game Γ_{BN} is a function $s_i(t_i)$, where for each type t_i in T_i , $s_i(t_i)$ specifies the action from the feasible set of actions A_i that type t_i would choose if drawn by nature.

Now that a strategy is completely characterized we can define a Bayes-Nash equilibrium for static games of incomplete information. The solution is called “Bayes-Nash” because of the method that the players are to use when updating their beliefs after observing information. As I mentioned earlier, we are currently assuming that when a player is told his type this does *not* alter the exogenously given probability distribution about the other player’s types. However, it could, and if it did the player would use Bayes’ rule to update his beliefs about the type of the other player. For instance, it is plausible that in an art auction a bidder’s own value provides some information about the other bidders’ values. If one bidder’s value is high, then it is likely that all bidders’ values are high. Thus, a player’s observation of his own value would cause his belief about the types of the other bidders to change. Again, because we assume that players update using Bayes’ rule we call this a Bayes-Nash equilibrium. We will discuss Bayes’ rule in more detail in the next section.

Definition 2 (Bayes-Nash equilibrium) In the game Γ_N , a set of strategies $s^* = (s_1^*, \dots, s_I^*)$ are a Bayes Nash equilibrium if for each player i and for each of i ’s types t_i in T_i , $s_i^*(t_i)$ solves:

$$\max_{a_i \in A_i} \sum_{t_{-i} \in T_{-i}} u_i(s_1^*(t_1), \dots, s_{i-1}^*(t_{i-1}), a_i, s_{i+1}^*(t_{i+1}), \dots, s_n^*(t_n); t) p_i(t_{-i}|t_i).$$

2.1 Normal form game example

Consider the following game of incomplete information between Wyatt Earp and a stranger in town. With probability 0.75, Earp believes the stranger is a gunslinger (i.e. a fast draw). With probability 0.25, Earp believes the stranger is a cowpoke (i.e. a slow draw). Earp only knows the probability of each type before taking an action, and does not observe the stranger’s type. This means that Earp believes he is in the matrix on the left 75% of the time, and in the one on the right 25% of the time. The stranger knows his own type, and also knows exactly who Wyatt Earp is. The payoffs to this simultaneous game of incomplete information are as follows:

		Stranger (gunslinger)				Stranger (cowpoke)	
		Draw	Wait			Draw	Wait
Earp	Draw	2, 3	3, 1	Earp	Draw	5, 2	4, 1
	Wait	1, 4	8, 2		Wait	6, 3	8, 4

1. How many types does each player have?
2. Does any player (or any player type) have a dominant strategy?
3. Find all pure strategy Bayes-Nash equilibria to this game.

To answer the first question, there is only one Wyatt Earp and no uncertainty over who he is. Thus, Wyatt Earp has 1 type. However, there is uncertainty over who the stranger is – the stranger could be a (1) gunslinger or a (2) cowpoke. Thus, the stranger has two types, gunslinger and cowpoke.

In this game, notice that regardless of what Wyatt Earp does the gunslinger has a strictly dominant strategy of choosing "Draw". Regardless of what strategy Wyatt Earp chooses (Draw or Wait), the gunslinger type will always choose Draw. So any equilibrium to this game should have the gunslinger type choosing "Draw". Note that no other player has a strictly dominant strategy – for Wyatt Earp, choosing "Wait" is almost better than every other strategy, except that when the gunslinger type chooses "Draw" Earp would also choose "Draw".

Now, finding equilibria in this game. Suppose that Wyatt Earp chooses Draw. The gunslinger type's best response is to choose Draw (since the gunslinger has a strictly dominant strategy). Also, the cowpoke type's best response is to choose Draw, since $2 > 1$. So we have just found the stranger's best response when Earp chooses Draw. Now we need to check: Is Earp's strategy choice of Draw a best response to the stranger's strategy of choosing: Draw if gunslinger, Draw if cowpoke? To do this we need to determine the expected value of Earp choosing Draw and the expected value of Earp choosing Wait. If Earp chooses Draw he receives:

$$\begin{aligned} E[\text{Draw} | \text{Draw if GS, Draw if CP}] &= 2 * \frac{3}{4} + 5 * \frac{1}{4} \\ E[\text{Draw} | \text{Draw if GS, Draw if CP}] &= \frac{11}{4} \end{aligned}$$

The 2 is from Earp's payoff in the upper left corner of the gunslinger matrix, the $\frac{3}{4}$ is the probability Earp believes he is facing a gunslinger, the 5 is from the upper left corner of the cowpoke matrix, and the $\frac{1}{4}$ is the probability that Earp believes he is facing a cowpoke. So, if both types Draw, Earp expects to receive a payoff of $\frac{11}{4}$ from choosing Draw. Now, what if Earp choose to Wait?

$$\begin{aligned} E[\text{Wait} | \text{Draw if GS, Draw if CP}] &= 1 * \frac{3}{4} + 6 * \frac{1}{4} \\ E[\text{Wait} | \text{Draw if GS, Draw if CP}] &= \frac{9}{4} \end{aligned}$$

So, Earp's strategy of choosing Draw is a best response to (Draw if gunslinger, Draw if cowpoke). Thus, one Bayes Nash equilibrium is: Earp Draw, Stranger Draw if gunslinger, Stranger Draw if cowpoke.

Now suppose that Wyatt Earp chooses Wait. The gunslinger's best response is still to choose Draw. But the cowpoke's best response if Earp chooses Wait is also to choose Wait. Now we need to determine if Earp choosing Wait is a best response to the stranger's strategy, which is Draw if gunslinger, Wait if cowpoke. If Earp chooses Wait he receives:

$$\begin{aligned} E[\text{Wait} | \text{Draw if GS, Wait if CP}] &= 1 * \frac{3}{4} + 8 * \frac{1}{4} \\ E[\text{Wait} | \text{Draw if GS, Wait if CP}] &= \frac{11}{4} \end{aligned}$$

The 1 is from the lower left corner of the first matrix (Earp Wait, gunslinger Draw) while the 8 is from the lower right corner of the second matrix (Earp Wait, cowpoke Wait). If Earp chooses to Draw, he receives:

$$\begin{aligned} E[\text{Draw} | \text{Draw if GS, Wait if CP}] &= 2 * \frac{3}{4} + 4 * \frac{1}{4} \\ E[\text{Draw} | \text{Draw if GS, Wait if CP}] &= \frac{10}{4} \end{aligned}$$

Since $\frac{11}{4} > \frac{10}{4}$, Earp will choose to Wait rather than Draw. Thus, a second Bayes Nash equilibrium is: Earp Wait, Stranger Draw if gunslinger, Stranger Wait if cowpoke.

3 Cournot competition with multiple types

Consider now a modified version of Cournot competition. In this game, Firm 1 has a known cost of c_1 . Thus, we could say that Firm 2 has a belief that Firm 1's cost is c_1 with probability 1. However, Firm 2 may have either a high cost, c_H , or a low cost, c_L . Firm 2 has cost of c_H with probability θ and cost of c_L

with probability $1 - \theta$. Before beginning to find the Bayes-Nash equilibrium, it is important to note that the Bayes-Nash equilibrium will consist of 1 set of actions for Firm 1 (since Firm 1 has 1 type) and 2 sets of actions for Firm 2 (since Firm 2 has 2 types). Using the structure from the previous Cournot games, Firm 2 will solve:

$$\max_{q_2 \geq 0} [a - q_1^* - q_2 - c_H] q_2$$

if its cost is c_H and:

$$\max_{q_2 \geq 0} [a - q_1^* - q_2 - c_L] q_2$$

if its cost is c_L . Firm 1 will solve the following problem:

$$\max_{q_1 \geq 0} \theta [a - q_1 - q_2^*(c_H) - c] q_1 + (1 - \theta) [a - q_1 - q_2^*(c_L) - c] q_1$$

Solving these problems yields the following first-order conditions:

$$\begin{aligned} q_2^*(c_H) &= \frac{a - q_1^* - c_H}{2} \\ q_2^*(c_L) &= \frac{a - q_1^* - c_L}{2} \\ q_1^* &= \frac{\theta [a - q_2^*(c_H) - c] + (1 - \theta) [a - q_2^*(c_L) - c]}{2} \end{aligned}$$

Again, these are almost the best response functions for the firms, but recall that if either firm is to produce so much as to make the other firm desire to produce a negative quantity then the firm will choose to produce 0 since it cannot produce a negative quantity. There are 3 equations and 3 unknowns so solving (assuming the restrictions on c_H , c_L , and c are such that all firms produce a positive output) we have:

$$\begin{aligned} q_2^*(c_H) &= \frac{a - 2c_H + c}{3} + \frac{1 - \theta}{6} (c_H - c_L) \\ q_2^*(c_L) &= \frac{a - 2c_L + c}{3} - \frac{\theta}{6} (c_H - c_L) \\ q_1^* &= \frac{a - 2c + \theta c_H + (1 - \theta) c_L}{3} \end{aligned}$$

Recall that in the case where both firms had identical marginal costs of c that $q_1^* = q_2^* = \frac{a-c}{3}$. Note that if $c_H = c_L = c$, then both firms would produce $\frac{a-c}{3}$. So imposing our prior assumptions brings us right back to our prior solution, which is a good thing. Also note that if $\theta = 0$ then this means that Firm 2 has the low cost with probability 1, so that $q_2^*(c_L) = \frac{a-2c_L+c}{3}$ and $q_1^* = \frac{a-2c+c_L}{3}$. These are just the best responses when the two firms have different, but known to all, costs.

4 Auctions and mechanisms

4.1 General Environment

Before discussing the auction formats and the equilibrium strategies we need to set up the general environment. This suggests that if the environment (or pieces of it) change, the NE bidding strategies will change.

The general name for the environment is the Symmetric Independent Private Values environment (SIPV) with Risk-neutral bidders. We will also assume that we are auctioning off a single, indivisible unit of the good.

1. There needs to be a probability distribution for player values, denoted v_i . We will assume that all player values are drawn from the uniform distribution on the unit interval. This means that all values are drawn from the interval $[0, 1]$ with equal probability. More importantly, if you draw a value of 0.7, then the probability that someone else drew a value less than you is also 0.7. Since probabilities

must add up to 1, and since the other player's value draw must either be greater than your value or less than your value. We will not allow for the fact that someone else could draw the exact same value (theoretically, ties cannot occur with positive probability in a continuous probability distribution). This means that the probability that the other player has a value greater than yours is $1 - 0.7 = 0.3$.

2. The setting is symmetric in the sense that all players know that the other player's value(s) is drawn from the same probability distribution.
3. The setting is independent in the sense that your value draw has NO impact on the value draw of the other player(s).
4. The setting is private in the sense that only you know your value – thus, it is private information.
5. We add the fact that our bidders are risk-neutral,³ as risk aversion will alter some results. Thus, our utility function will be:

$$u(x) = \begin{cases} x & \text{if win the auction} \\ 0 & \text{if don't win} \end{cases}$$

The term x in the utility function can typically that of $v_i - p$, where v_i is the player i 's value and p is the price paid by player i . Note that a player's expected utility in these auctions can be noted as:

$$u_i = \Pr(\text{win}) * (v_i - p) + \Pr(\text{lose}) * 0$$

where $\Pr(\text{win})$ is the probability that bidder i wins the auction and $\Pr(\text{lose})$ is the probability that bidder i loses the auction. If the bidder wins he receives his value minus the price paid, or $(v_i - p)$ and if he loses he receives 0. Thus, for many auctions, the expected utility of a bidder is:

$$u_i = \Pr(\text{win}) * (v_i - p)$$

Note that the difficulty in deriving the theoretical results lies in establishing the probability of winning, $\Pr(\text{win})$ and, in some cases, the price paid, p , particularly when the price paid depends upon another bidder's bid.

4.2 Auction formats

In this section I will describe the four basic auction formats that we will discuss. The description will include the process by which bids are submitted and the assignment rule for the winner. For now, consider only the cases where we have a single, indivisible unit for sale.

4.2.1 1st-price sealed bid auction

Process All bidders submit a bid on a piece of paper to the auctioneer.

Assignment rule The highest bidder is awarded the object. The price that the high bidder pays is equal to his bid.

Examples Many procurement auctions are 1st-price sealed bid. Procurement auctions are typically run by the government to auction off a construction job (such as paving a stretch of highway).

4.2.2 Dutch Auction

Process There is a countdown clock that starts at the top of the value distribution and counts backwards. Thus, the price comes down as seconds tick off the clock. When a bidder wishes to stop the auction he or she yells, "stop".

³Recall that a risk-neutral individual is indifferent between receiving \$5 with certainty and a gamble that pays \$5 on average (like one that has a 50% chance at \$0 and a 50% chance at \$10). A risk averse individual would prefer the certain \$5 over the expected \$5 and a risk loving individual would prefer the expected \$5 over the certain \$5.

Assignment rule The bidder who called out stop wins the auction, and the bidder pays the last price announced by the auctioneer.

Examples The Aalsmeer flower auction, in the Netherlands, is an example of this type of auction. Hmm, wonder where the phrase “Dutch” auction comes from ...

By the way, the eBay dutch auctions are NOT Dutch auctions as we have described them. They are multi-unit ascending $k + 1$ price auctions.

4.2.3 2^{nd} -price sealed bid auction

Process Bidders submit their bids on a piece of paper to the auctioneer.

Assignment rule The highest bidder wins, but the price that the highest bidder pays is equal to the 2^{nd} highest bid. Hence the term 2^{nd} -price auction.

Examples eBay is kind of a warped 2^{nd} -price auction. If you think about the very last seconds of an eBay auction (or if you consider that every person only submits one bid), think about what happens. You are sending in a bid. If you have the highest bid you will win. You will pay an amount equal to the 2^{nd} highest bid plus some small increment. Thus if you submit a bid of \$10 and the second highest bid is \$4, you pay \$4 plus whatever the minimum is (I think it’s a quarter). So you would pay \$4.25.

There are other reasons to think that eBay is not actually a 2^{nd} -price auction but those can be discussed later.

4.2.4 Ascending clock auction

Process A clock starts at the bottom of the value distribution. As the clock ticks upward, the price of the item rises with the clock. This is truly supposed to be a continuous process, but it is very difficult to count continuously, so we will focus on one tick of the clock moving the price up one unit. The idea is that this is the smallest amount that anyone could possibly bid – that is how the ticks on the clock move the price up. All bidders are considered in the auction (either they are all standing or they all have their hands on a button – some mechanism to show that they are in). When the price reaches a level at which the bidder no longer wishes to purchase the object, the bidder drops out of the auction (sits down or releases the button). Bidders cannot reenter the auction. Eventually only two bidders will remain. When the next to last bidder drops out, the last bidder wins.

Assignment rule The winning bidder is the last bidder left in the auction. The bidder pays a price equal to the last price on the clock.

Examples The typical example given is Japanese fish markets, though those may be an urban legend. Thus, the English clock auction may only be a theoretical construct.

4.3 Bidding strategies

Bidding strategies for the four auctions formats are described below. They are ordered in terms of deriving the simplest equilibrium bidding strategy to the most difficult.

4.3.1 Ascending clock auction – bidding strategy

Consider the following example. Assume $v_i = 10$. The clock begins at 0 and ticks upward: 0, 1, 2, 3, ..., 9, 10, 11, 12, 13, ... The question is, when should you sit down (or drop out of the auction)? Consider three possible cases:

1. The clock reaches 11:

In this case you should drop out. While you increase your chances of winning the item by staying in, note that you will end up paying more than the item is worth to you. You can do better than this by dropping out of the auction and receiving a surplus of zero. So, as soon as the price on the clock exceeds your value you should drop out.

2. The clock is at some price less than 10:

In this case you should remain in the auction. If you drop out you will receive 0 surplus. However, if you remain in the auction then you could win a positive surplus. If you drop out before your value is reached you are essentially giving up the chance to earn a positive surplus. Since this positive surplus is greater than the 0 surplus you would receive if you dropped out, you should stay in the auction.

3. The clock is at 10:

What happens when the price on the clock reaches your value? Well, if you win the auction you get 0 surplus and if you drop out you get 0 surplus, so regardless of what you do you get 0 surplus. We will say that you stay in at 10, and drop out at 11. For one thing, it makes the NE bidding strategy simple – stay in until your value is reached, then drop out. Another way to motivate this is to consider that peoples values are drawn from the range of numbers $[0.01, 1.01, 2.01, 3.01, \dots]$ instead of $[0, 1, 2, 3, \dots]$. However, assume the prices increase as $[0, 1, 2, 3, \dots]$. It is clear that if you have a value of 3.01 you should be in at 3, while if you have a value of 3.01 you should be out at 4. This is the “add a small amount to your value” approach that I mentioned in class.

So what is the NE strategy? Stay in until your value is reached and drop out as soon as it is passed by the clock. Or, if we let $b_i(v_i)$ represent player i 's bid as a function of his value, we have $b_i(v_i) = v_i$.

4.3.2 2^{nd} -price sealed bid auction – bidding strategy

There is a good for which players have a value, v_i , drawn from a probability distribution $F(\cdot)$. Thus, there are multiple types, each represented by one of the v_i value draws from $F(\cdot)$. All player values are drawn from the same distribution $F(\cdot)$ and the actual value draws are only known to the individual players. The draws are independent so that a draw of a value for one player provides no information as to the value of the other players. All players know these assumptions.

First, note that an action needs to be specified for every type. Thus, it will be easiest to have some function $b_i(v_i)$ represent the player's bid based upon his observed value. Now, consider the following candidate for a best response function:

$$b_i(v_i) = v_i.$$

This function states that the each bidder should simply bid his value. Is this a Bayes-Nash equilibrium for this game? Suppose that a player bids his value. The player either wins the auction or loses the auction. If the player wins the auction does he want to change his strategy? No, because he won. If he submits a bid higher than his value then this does not alter his payoff since his payment is tied to the second highest bid (which he has no control over). If he lowers his bid then he runs the risk of losing the auction and receiving 0 and lowering his bid also does not alter his payoff. So if the player wins this strategy is at least as good as some other strategy. If he loses by submitting a bid equal to his value then he cannot increase his payoff by lowering his bid because he still loses. He could raise his bid in an attempt to win the good, but this would lead to a negative payoff due to the fact that the winning bid must have been greater than his value, so in order to win he would need to submit a winning bid greater than his value AND he would pay an amount greater than his value. Note that this yields a negative payoff. The following example may make this a little clearer.

To further illustrate the point consider the following table when there are two bidders. Suppose that bidder 1 has a value of 12.

	Bidder 1's bid ($v_1 = 12$)		
Other bidder's bid	$b_1 = 10$	$b_1 = 12$	$b_1 = 14$
$b_2 < 10$	$12 - b_2$	$12 - b_2$	$12 - b_2$
$10 < b_2 < 12$	0	$12 - b_2$	$12 - b_2$
$12 < b_2 < 14$	0	0	$(12 - b_2)$
$b_2 > 14$	0	0	0

Note that $(12 - b_2)$ is NEGATIVE. We have now determined that submitting a bid equal to our value is at least as good as submitting a bid greater than or lower than the value in some cases, and strictly better in other cases. Therefore, submitting a bid equal to your value is a weakly dominant strategy. Thus, the Bayes-Nash equilibrium for a 2^{nd} -price auction: Submit a bid equal to your value, or $b_i(v_i) = v_i$. One way to “find” Nash equilibria in games of this type is to propose that a particular strategy is a Bayes-Nash equilibrium and then determine whether or not it actually is.

4.3.3 1^{st} -price sealed bid and descending clock auction – bidding strategies

This first part relaxes some assumptions to show that bidders will bid below value.

A bidder's expected utility in the 1^{st} -price auction consists of two parts: (1) his utility of money income from winning and (2) the probability that he gains that income. Let $(v_i - b_i)$ be bidder i 's money income if his bid is the highest, and 0 be his money income if not. Let $F_i(b_i)$ be bidder i 's subjective probability that he will win with a bid of b_i . Bidder i 's expected utility is then:

$$U_i(b_i) = F_i(b_i) u_i(v_i - b_i). \quad (1)$$

Note that this is simply the probability with which bidder i believes he will win with a bid of b_i multiplied by the utility of his money income if he wins with a bid of b_i . Now, we make 4 assumptions:

1. Amount bid is a continuous variable (for mathematical tractability)
2. Interval $[x_i, \bar{x}_i]$ is the support of the probability distribution
3. U_i is a quasiconcave function
4. There exists a unique positive expected utility maximizing bid b_i^0

Now, b_i^0 satisfies:

$$\frac{\partial U_i(b_i)}{\partial b_i} = F_i'(b_i^0) u_i(v_i - b_i^0) - F_i(b_i^0) u_i'(v_i - b_i) = 0. \quad (2)$$

Assume b_i^0 satisfies the 2^{nd} -order condition for a maximum ($0 > U_i''(b_i^0)$). The first and second order conditions and implicit function theorem imply there exists a differentiable function Ψ_i such that $b_i^0 = \Psi_i(v_i)$. Thus, the function Ψ_i takes any value that bidder i could have and transforms it into the expected utility maximizing bid. The function Ψ_i is called the bid function. From the first-order condition, it is known that:

$$b_i^0 = v_i - u_i^{-1}(u_i'(v_i - b_i^0) F_i(b_i^0) / F_i'(b_i^0)), \quad (3)$$

where u_i^{-1} is the inverse of the utility of money income function. This relationship is important because it shows that the expected utility maximizing bid is less than the bidder's value. Thus, the 1^{st} -price sealed bid auction is not demand revealing. The amount by which v_i exceeds b_i^0 depends on u_i and F_i , which represent bidder i 's risk preferences and subjective probability of winning. Since these factors may differ among bidders, the 1^{st} -price sealed bid auction may not be Pareto efficient, in that the item may not be awarded to the bidder with the highest value in equilibrium. If all bidders have the same risk preferences and subjective probability of winning functions then the 1^{st} -price sealed bid auction will be Pareto efficient.

Imposing some simplifying assumptions We “know” the 2^{nd} -price auction has a weakly dominant strategy where each player bids his own value, and this is a weakly dominant strategy regardless of risk preferences or subjective probabilities of winning (as long as it is in the SIPV framework). We have also seen that the 1st-price auction bids depend on risk preferences and subjective probabilities (that is the previous section). To derive Nash equilibrium bidding functions for these auctions we will impose that risk preferences and subjective probabilities of winning are the same for all individuals, and we will specifically impose that all bidders are risk neutral.

Our goal is to find a strategy function $b^*(v)$, such that for every value $v \in [\underline{v}, \bar{v}]$ an individual specifies a bid. This is a Bayes-Nash equilibrium. For simplicity, assume that values are distributed uniformly on the unit interval $[0, 1]$ so that we are looking for a $b^* : [0, 1] \rightarrow \mathbb{R}_+$. Also, there is an inverse bid function, $b^{*-1}(b)$, which specifies the value an individual has if he makes a specific bid. Assume that values are drawn from a continuous distribution and that $b^*(v)$ is strictly monotone increasing, meaning that if $v_1 > v_2$ then $b^*(v_1) > b^*(v_2)$. Those 2 assumptions imply that ties have zero probability.

In order to specify a Nash equilibrium we will look at one bidder’s strategy choice assuming all other bidders are using the same strategy. Think about it this way – if all bidders are identical, and the N^{th} bidder’s best response to $N - 1$ bidders using a particular strategy is the same strategy, then solving this problem for all N bidders becomes redundant as the strategies they are using will be best responses to one another. The first thing we must determine is a bidder’s probability of winning when bidding b against rivals who use $b^*(v)$:

$$\begin{aligned} p(b) &= \Pr \{b^*(v_j) < b, \forall j \neq i\} \\ &= \Pr \{v_j < \sigma(b), \forall j \neq i\}, \text{ where } \sigma := b^{*-1}. \\ &= F(\sigma(b))^{N-1} \end{aligned}$$

Essentially, an individual’s probability of winning is equal to the probability that the individual has a higher value than the other $N - 1$ individuals.

In a Nash equilibrium, an individual must be using his best response, which means he must be maximizing his expected utility (well, we will use expected utility as the objective function). So an individual must maximize:

$$U_i(b, v) = p(b) * (v - b).$$

Again, this individual is risk-neutral, so we can specify $u_i(v - b) = v - b$. Maximizing expected utility we find:

$$\frac{\partial U_i}{\partial b} = p'(b) * (v - b) - p(b) = 0.$$

Now, we must have $v = \sigma(b)$ since this is what $\sigma(b)$ specifies and we also know that someone with a value of 0 must bid 0, so $\sigma(0) = 0$. Using $p(b) = F(\sigma(b))^{N-1}$ and $v = \sigma(b)$ we have:

$$\begin{aligned} \frac{\partial U_i}{\partial b} &= (N - 1) F(\sigma(b))^{N-2} F'(\sigma(b)) * (\sigma(b) - b) - F(\sigma(b))^{N-1} = 0 \\ &\text{or} \\ \frac{\partial U_i}{\partial b} &= (N - 1) F'(\sigma(b)) * (\sigma(b) - b) - F(\sigma(b)) = 0. \end{aligned}$$

Since values are drawn from the uniform distribution on $[0, 1]$, $F(\sigma(b)) = \sigma(b)$ and thus $F'(\sigma(b)) = \sigma'(b)$, so that:

$$\frac{\partial U_i}{\partial b} = (N - 1) \sigma'(b) * (\sigma(b) - b) - \sigma(b) = 0.$$

We now have a differential equation, which has the solution:

$$\sigma(b) = \left(\frac{N}{N - 1} \right) b.$$

Thus, $\sigma(b)$ is a best response by bidder i if all other bidders are using $\sigma(b)$. We should check to see if this

is actually a solution:

$$\begin{aligned}
(N-1) \left(\frac{N}{N-1} \right) * \left(\left(\frac{N}{N-1} \right) b - b \right) - \left(\frac{N}{N-1} \right) b &= 0 \\
N * \left(\frac{Nb - Nb + b}{N-1} \right) - N * \left(\frac{b}{N-1} \right) &= 0 \\
N * \left(\frac{b}{N-1} \right) - N * \left(\frac{b}{N-1} \right) &= 0.
\end{aligned}$$

Solving $\sigma(b)$ for b we find:

$$\begin{aligned}
b^* &= \frac{N-1}{N} \sigma(b) \\
&\text{or} \\
b^* &= \frac{N-1}{N} v.
\end{aligned}$$

A useful extension of this result modifies the range of values from $[0, 1]$ to $[\underline{v}, \bar{v}]$. If values are uniformly distributed along $[\underline{v}, \bar{v}]$, then b^* is:

$$b^*(v_i) = \frac{N-1}{N} (v_i - \underline{v}) + \underline{v}.$$

All we are doing is placing the minimum bid at \underline{v} (if $v = \underline{v}$ then $b = \underline{v}$, just like when values were distributed from $[0, 1]$ if $v = 0$ then $b^* = 0$) and then taking a proportion of the difference between v and \underline{v} .

5 Some useful results

Among the many results for Bayesian games there are two that are used quite a bit. The first result is the Revelation Principle and the second result is the Revenue Equivalence Theorem.

5.1 Revelation principle

Suppose that a seller wishes to sell an object using some mechanism – the precise mechanism is left unspecified as long as the following conditions are met:

1. The bidders simultaneously make claims about their types. Bidder i can claim to be any type from his feasible set of types.
2. Given the bidders' claims, bidder i pays an amount that is a function of all the reported types and receives the good with some probability based upon the reported types (in an auction the bidder with the highest reported type receives the good with probability 1).

Games that satisfy these criteria are known as direct mechanisms, because the only action is to submit a claim about a type.

Suppose the seller wants to restrict attention to direct mechanisms in which the bidders only report truthful values. It is possible, choosing the appropriate game, to take any direct mechanism in which a Bayes-Nash equilibrium exists which is not a truth-telling equilibrium and transform that game into a direct mechanism in which truth-telling is the equilibrium strategy. Thus:

Any Bayesian Nash equilibrium of any direct mechanism can be represented by a truth-telling direct mechanism.

This result is useful in its own right because it is sometimes difficult to pin down the properties of certain games, but by being able to relate one Bayes-Nash equilibrium to another in a different game we can get a clearer view of the complicated game's properties.

5.2 Revenue equivalence

To compare expected revenue between the 1st and 2nd-price sealed bid auctions we need to compare the 2nd-highest value drawn in a 2nd-price sealed bid auction (this is the amount the bidder with the highest value will pay in equilibrium) with the BID by the bidder with the highest value in the 1st-price auction.

Before beginning the discussion of revenue equivalence we need to determine the amount of revenue a seller of an item expects to receive from a particular auction design. Consider the 2nd-price sealed bid auction. If players follow equilibrium (why shouldn't they?), then the seller can expect to receive an amount equal to the expected value of the 2nd highest value drawn. In a 2nd-price auction of N bidders, the implied probability distribution of revenue is that of the $(N - 1)$ th order statistic.

What is an order statistic? Suppose we draw from some probability distribution $F(X)$ N times (so there are N value draws). We rank the realizations in increasing order, so:

$$x_{v_1} \leq x_{v_2} \leq \dots \leq x_{v_k} \leq \dots \leq x_{v_N}. \quad (4)$$

The k^{th} order statistic is the function $X_{(k)}$ that assigns to each realization for the series (X_1, \dots, X_N) the k^{th} smallest value x_{v_k} . Order statistics are random variables and have means and variances. For the uniform distribution over $[\underline{v}, \bar{v}]$, the expected value of the $(N - 1)$ th order statistic is:

$$\frac{(N - 1)(\bar{v} - \underline{v})}{N + 1} + \underline{v}. \quad (5)$$

Thus, this is the expected revenue in the 2nd-price sealed bid auction, as well as the expected value of the 2nd highest value draw. If $N = 2$, $\bar{v} = 1$, and $\underline{v} = 0$ then the expected value of the second highest value is $\frac{1}{3}$. If $N = 4$, then the expected value of the second highest value is $\frac{3}{5}$. Note that as N increases the expected value of the second highest value increases. Also, we can let N remain a variable and let $\bar{v} = 1$, and $\underline{v} = 0$. Then the expected revenue is $\frac{N-1}{N+1}$ if we assume that values are uniformly distributed on $[0, 1]$.

Now, we need to compare this revenue with that from the 1st-price sealed bid auction. Simply for expositional ease, continue with the assumption that values are distributed uniformly. We know $\frac{N-1}{N+1}$ is the expected value of the $(N - 1)^{th}$ order statistic. In a 1st-price auction we need the expected value of the N^{th} order statistic (the highest draw). This is just $\frac{N}{N+1}$. To find the expected revenue, simply find the bid that would be made by the highest draw in a 1st-price auction. This bid is just $\frac{N-1}{N} * \frac{N}{N+1}$ and we have that expected revenue from the 1st-price auction is $\frac{N-1}{N+1}$. Thus, in the SIPV-RN environment, we have at least one example (uniformly distributed values) where expected revenue is equivalent across auction institutions. This result holds generally for auction institutions in the SIPV-RN environment, as long as any bidder who draws the lowest possible value (\underline{v}) receives an expected surplus of zero and the auction is Pareto efficient (it awards the object(s) to the highest valued user(s) – at least theoretically). Be aware that this result is ONLY for the SIPV-RN environment – if values are not independent, or bidders are risk averse (even if they have the same level of risk aversion) this will not necessarily hold.