1 Bayes-Nash equilibria

The problems that have been discussed up to this point in the course have dealt with situations where players know the exact payoff function of the other player. The players may not know exactly what strategy the other player has chosen (as in a simultaneous game) but they know precisely what each player's payoff will be if they take a particular action and the other player takes a particular action. We will discuss two types of games in which there is uncertainty over what the other person's precise payoff is. First we will look at a normal form game with uncertainty. Then we will consider auctions, in which bidders know their own value but only that the other bidder(s) value is drawn from some probability distribution. Thus, given a particular action for both bidders, if the other bidder wins then you will not know precisely what his payoff is because you do not know his precise value.

Finally we will discuss an oligopoly market with incomplete information. So far we have assumed that all firms know all other firms' costs. However, this is unrealistic, as firms do not necessarily know precisely what the other firm's cost is but may only have some idea as to whether it is high or low.

When solving games such as these where players can have different "types" it will be necessary to specify an action for EACH type a player COULD be. Thus, although in an auction a player only draws one value, it is necessary to specify what the player will do for any possible value draw. In the oligopoly model, it will be necessary to specify one action for each cost type each player could have. Of course, when specifying a strategy (the entire collection of actions) it will be necessary to make sure that the players' strategies are best responses to one another. This type of equilibrium is called a Bayes-Nash equilibrium.

2 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)	
		Draw	Wait
Earp	Draw	$2,3$	$3,1$
	Wait	$1,4$	$8,2$

		Stranger (cowpoke)	
		Draw	Wait
Earp	Draw	$5,2$	$4,1$
	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:

$$E\left[Draw|Draw\ if\ GS, Draw\ if\ CP\right] = 2 * \frac{3}{4} + 5 * \frac{1}{4}$$
$$E\left[Draw|Draw\ if\ GS, Draw\ if\ CP\right] = \frac{11}{4}$$

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?

$$E\left[Wait|Draw\ if\ GS, Draw\ if\ CP\right] = 1 * \frac{3}{4} + 6 * \frac{1}{4}$$
$$E\left[Wait|Draw\ if\ GS, Draw\ if\ CP\right] = \frac{9}{4}$$

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:

$$E\left[Wait|Draw\ if\ GS, Wait\ if\ CP\right] = 1 * \frac{3}{4} + 8 * \frac{1}{4}$$
$$E\left[Wait|Draw\ if\ GS, Wait\ if\ CP\right] = \frac{11}{4}$$

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:

$$E\left[Draw|Draw\ if\ GS, Wait\ if\ CP\right] = 2 * \frac{3}{4} + 4 * \frac{1}{4}$$
$$E\left[Draw|Draw\ if\ GS, Wait\ if\ CP\right] = \frac{10}{4}$$

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.

Finally, note that this is really a 2 player game between Earp and the Stranger. As such, we can put the entire game into a single matrix if we can identify the Stranger's strategies. The Stranger has 4 strategies here: (1) Draw if gunslinger, Draw if cowpoke, (2) Draw if gunslinger, Wait if cowpoke, (3) Wait if gunslinger, Draw if cowpoke, and (4) Wait if gunslinger, Wait if cowpoke. Identifying strategies is the easy part – determining the payoffs is slightly more difficult. The matrix, with the payoffs we have already found and letters representing the other payoffs, is:

		Earp	
		Draw	Wait
	D if GS, D if CP	$a, \frac{11}{4}$	$b, \frac{9}{4}$
Stranger	D if GS, W if CP	$c, \frac{10}{4}$	$d, \frac{11}{4}$
	W if GS, D if CP	e, f	g, h
	W if GS, W if CP	i, j	k, l

To find f we need:

$$E_{Earp}\left[Draw|Wait\ if\ GS, Draw\ if\ CP\right] = 3 * \frac{3}{4} + 5 * \frac{1}{4}$$
$$E_{Earp}\left[Draw|Wait\ if\ GS, Draw\ if\ CP\right] = \frac{14}{4}$$

To find h we need:

$$\begin{aligned} E_{Earp}\left[Wait|Wait\ if\ GS, Draw\ if\ CP\right] &= 8*\frac{3}{4}+6*\frac{1}{4} \\ E_{Earp}\left[Wait|Wait\ if\ GS, Draw\ if\ CP\right] &= \frac{30}{4} \end{aligned}$$

To find j we need:

$$\begin{aligned} E_{Earp}\left[Draw|Wait\ if\ GS, Wait\ if\ CP\right] &= 3*\frac{3}{4}+4*\frac{1}{4} \\ E_{Earp}\left[Draw|Wait\ if\ GS, Wait\ if\ CP\right] &= \frac{13}{4} \end{aligned}$$

To find l we need:

$$\begin{aligned} E_{Earp}\left[Wait|Wait\ if\ GS, Wait\ if\ CP\right] &= 8*\frac{3}{4}+8*\frac{1}{4} \\ E_{Earp}\left[Wait|Wait\ if\ GS, Wait\ if\ CP\right] &= \frac{32}{4} \end{aligned}$$

That finishes Earp's payoffs. Now, to determine the Stranger's payoffs, if Earp uses Draw and the Stranger uses Draw if GS, Draw if CP, then a is:

$$\begin{aligned} E_{Stranger}\left[Draw\ if\ GS,\ Draw\ if\ CP|Draw\right] &= 3*\frac{3}{4}+2*\frac{1}{4} \\ E_{Stranger}\left[Draw\ if\ GS,\ Draw\ if\ CP|Draw\right] &= \frac{11}{4} \end{aligned}$$

To find b we need:

$$\begin{aligned} E_{Stranger}\left[Draw\ if\ GS,\ Draw\ if\ CP|Wait\right] &= 4*\frac{3}{4}+3*\frac{1}{4} \\ E_{Stranger}\left[Draw\ if\ GS,\ Draw\ if\ CP|Wait\right] &= \frac{15}{4} \end{aligned}$$

To find c we need:

$$\begin{aligned} E_{Stranger}\left[Draw\ if\ GS,\ Wait\ if\ CP|Draw\right] &= 3*\frac{3}{4}+1*\frac{1}{4} \\ E_{Stranger}\left[Draw\ if\ GS,\ Wait\ if\ CP|Draw\right] &= \frac{10}{4} \end{aligned}$$

To find d we need:

$$\begin{aligned} E_{Stranger}\left[Draw\ if\ GS,\ Wait\ if\ CP|Wait\right] &= 4*\frac{3}{4}+4*\frac{1}{4} \\ E_{Stranger}\left[Draw\ if\ GS,\ Wait\ if\ CP|Wait\right] &= \frac{16}{4} \end{aligned}$$

To find e we need:

$$\begin{aligned} E_{Stranger}\left[Wait\ if\ GS,\ Draw\ if\ CP|Draw\right] &= 1*\frac{3}{4}+2*\frac{1}{4} \\ E_{Stranger}\left[Wait\ if\ GS,\ Draw\ if\ CP|Draw\right] &= \frac{5}{4} \end{aligned}$$

To find g we need:

$$\begin{aligned} E_{Stranger}\left[Wait\ if\ GS,\ Draw\ if\ CP|Wait\right] &= 2*\frac{3}{4}+3*\frac{1}{4} \\ E_{Stranger}\left[Wait\ if\ GS,\ Draw\ if\ CP|Wait\right] &= \frac{9}{4} \end{aligned}$$

To find i we need:

$$E_{Stranger}\left[Wait\ if\ GS,\ Wait\ if\ CP|Draw\right] = 1*\frac{3}{4}+1*\frac{1}{4}$$
$$E_{Stranger}\left[Wait\ if\ GS,\ Wait\ if\ CP|Draw\right] = \frac{4}{4}$$

To find k we need:

$$E_{Stranger}\left[Wait\ if\ GS,\ Wait\ if\ CP|Wait\right] = 2*\frac{3}{4}+4*\frac{1}{4}$$
$$E_{Stranger}\left[Wait\ if\ GS,\ Wait\ if\ CP|Wait\right] = \frac{10}{4}$$

The actual matrix, with payoffs, looks like:

		Earp	
		Draw	Wait
	D if GS, D if CP	$\boxed{\frac{11}{4}}, \boxed{\frac{11}{4}}$	$\frac{15}{4}, \frac{9}{4}$
Stranger	D if GS, W if CP	$\frac{10}{4}, \frac{10}{4}$	$\boxed{\frac{16}{4}}, \boxed{\frac{11}{4}}$
	W if GS, D if CP	$\frac{5}{4}, \frac{14}{4}$	$\frac{9}{4}, \boxed{\frac{30}{4}}$
	W if GS, W if CP	$\frac{4}{4}, \frac{13}{4}$	$\frac{10}{4}, \boxed{\frac{32}{4}}$

Note that the two cells with both payoffs circled correspond to the two equilibria we found earlier. There is probably a mixed strategy Bayes Nash equilibrium, where Earp mixes over Draw and Wait while the Stranger mixes over D if GS, D if CP and D if GS, W if CP.

3 Auctions

3.1 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.

3.1.1 1^{st}-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 1^{st}-price sealed bid. Procurement auctions are typically run by the government to auction off a construction job (such as paving a stretch of highway).

3.1.2 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”.

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. Hmmm, 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.

3.1.3 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.

3.1.4 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.

3.2 Bidding strategies

The previous section is meant to introduce you to the auction formats. In this section we will discuss the NE bidding strategies. We will derive the bid functions for some simple cases. For those truly interesting in the gory details, I suggest the book by Wolfstettar (1999) *Topics in Microeconomics: Industrial Organization, Auctions, and Incentives.*

3.2.1 General Environment

Before discussing the bidding 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,[1] as risk aversion will alter some results. Thus, our utility function will be:

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

The term x in the utility function can typically that of as $v_i - p$, 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:

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

where $\Pr(win)$ is the probability that bidder i wins the auction and $\Pr(lose)$ 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:

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

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

3.2.2 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

[1] Recall 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.

– 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 $[0.01, 1.01, 2.01, 3.01, ...]$ instead of $[0, 1, 2, 3, ...]$. However, assume the prices increase as $[0, 1, 2, 3, ...]$. 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 $b_i(v_i) = v_i$.

3.2.3 2^{nd}-price sealed bid auction – bidding strategy

In this auction you submit a bid and pay a price equal to that of the 2^{nd} highest bid. How should you bid?

One method of finding a NE (or a solution in general) is to propose that a strategy is a NE and then verify it. Naturally, it is a good idea to propose the right strategy the first time. So, consider the strategy: submit your value. Is this a good strategy?

What else could we do? We could submit a bid greater than the value or less than the value. Let's examine each of these.

Bid above your value Suppose we submit a bid above our value. What could this possibly change? Well, if we were to win when submitting our value then absolutely nothing changes – we still pay the same price since the price (if we win) is not tied to our bid. What happens if we submit a bid greater than our value and this causes us to switch from losing the auction to winning the auction? Suppose our value is 12 and the other player's value is 14. The other player submits 14 and we submit 12. We lose and earn 0 surplus. Now suppose we were to bid 15. We win, which is good, but we have to pay 15 for something that is only worth 12 to us. So we earn a surplus of (-3). This is bad. We could have done better by placing a bid of 10 (our value) and earning 0. So placing a bid equal to our value is better than placing a bid above the value in this case.

Bid below your value Suppose we submit a bid below our value. What could this possibly change? Well, if we were going to lose by submitting our value, then we still lose when submitting a bid below the value. So this changes nothing (at least not for us – it would help the highest bidder if we were the 2^{nd} highest bid!) as we still receive 0 surplus. Suppose we lower our bid and still win – again nothing changes because the 2^{nd} highest bidder has still submitted the same bid. It is possible though that we lower our bid and lose – here's where the problem occurs. Suppose our value is 12 and the other value is 11. We submit a bid of 12, we win, and we get a surplus of $(12 - 11) = 1$. Now suppose we submit a bid of 8 – we go from getting a surplus of 1 to getting a surplus of 0. It would be much better to submit a bid equal to your value and get a surplus of 1.

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, so again we have $b_i(v_i) = v_i$.

You should note that the 2^{nd}-price sealed bid auction and the ascending clock auction are strategically equivalent. This means that all players have the same bidding strategies in either auction, even though the mechanism that produces the winner of the auction is slightly different.

3.2.4 1^{st}-price sealed bid auction – bidding strategy

In this auction you pay an amount equal to your bid if you win. The first question is, should you submit a bid equal to your value?

Bid equal to your value If you submit a bid equal to your value then you will expect to earn 0 surplus. If you win, then you will have to pay an amount equal to your value and if you lose you receive nothing. It stands to reason that you may be able to do better than this by submitting a bid below your value. The question is how far below your value?

Bid equal to the lowest possible value If you submit a bid equal to the lowest possible value that could be drawn then you will also receive 0 surplus. The reason is that you will never win because your bid was so low. Taken together with the fact that you will bid below your value, this means your actual bid should fall between the lowest possible value and your value draw.

Actual problem The actual problem facing someone bidding in a 1^{st}-price sealed bid auction is to maximize their expected utility. We have seen that in general we have:

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

In the case of the 1^{st}-price sealed bid auction, we know that if the bidder wins he will end up paying a price equal to his bid, so $p = b_i$. Thus we have:

$$u_i = \Pr(win) * (v_i - b_i)$$

For the example we have discussed, with two bidders and values uniformly distributed on $[0, 1]$, the $\Pr(win) = b_i$. Thus, if you bid 0.7 then you have a 70% chance of winning the auction; if you bid 0.4 then you have a 40% chance of winning the auction.[2] The problem is now a simple maximization problem:

$$\begin{aligned} \max_{b_i} u_i &= b_i * (v_i - b_i) \\ \frac{du_i}{db_i} &= v_i - 2b_i \\ 0 &= v_i - 2b_i \\ 2b_i &= v_i \\ b_i &= \frac{1}{2}v_i \end{aligned}$$

Thus, given the SIPV-RN environment with values distributed uniformly on $[0, 1]$ and 2 bidders we have that player i's bid function is $b_i(v_i) = \frac{1}{2}v_i$. We can extend this result to the case of N bidders fairly easily. Now you have to consider the fact that your bid has to be higher than $N-1$ other bidders' bids. This alters the expected utility function to:

$$u_i = (b_i)^{N-1} * (v_i - b_i)$$

The reason that the b_i is raised to the $N - 1^{st}$ power is because the bidder now has to have a higher bid than $N - 1$ other bidders. While the problem is slightly more complicated than with 2 bidders it is still a

[2] Recall that we are assuming that the bid function is strictly monotone increasing, meaning that bidders with higher values will submit strictly higher bids.

fairly easy problem to solve:

$$\begin{aligned}
\max_{b_i}\ u_i &= (b_i)^{N-1} * (v_i - b_i) \\
\frac{du_i}{db_i} &= (N-1)(b_i)^{N-2} * (v_i - b_i) + (b_i)^{N-1} * (-1) \\
0 &= (N-1)(b_i)^{N-2} * (v_i - b_i) + (b_i)^{N-1} * (-1) \\
0 &= (N-1)(b_i)^{N-2} * (v_i - b_i) - (b_i)^{N-1} \\
& \text{divide by } (b_i)^{N-2} \\
0 &= (N-1)(v_i - b_i) - b_i \\
0 &= (N-1)v_i - (N-1)b_i - b_i \\
0 &= (N-1)v_i - Nb_i \\
Nb_i &= (N-1)v_i \\
b_i &= \frac{N-1}{N}v_i
\end{aligned}$$

Thus, for the general case of N bidders, the bid function in a 1^{st}-price sealed bid auction is $b_i(v_i) = \frac{N-1}{N}v_i$. Note that when $N=2$ we have that $b_i(v_i) = \frac{1}{2}v_i$, which is what we found above. Thus you are shaving your bid depending on how many other bidders there are. The more bidders, the less you shave your bid.

3.2.5 Dutch auction – bidding strategy

Recall that with a Dutch auction the bidder watches as the clock descends, and then calls out when he sees a price that he wishes to pay. The problem facing the bidder is to maximize their expected utility. Their expected utility can be written as:

$$\begin{aligned}
u_i &= \Pr(win) * (v_i - b_i) + \Pr(lose) * 0 \\
u_i &= \Pr(win) * (v_i - b_i)
\end{aligned}$$

Notice that this is the same problem faced in the first price auction, if we make all the same assumptions we did when deriving the bid function for the 1^{st}-price sealed bid auction. This shows that the 1^{st}-price and the Dutch auctions are strategically equivalent. Thus, the bidding strategy in the Dutch auction is to yell out stop when the clock reaches $\frac{N-1}{N}$ of your value, where N is the number of total bidders in the auction (including yourself).

3.3 Which format is "better"?

Now that we have seen the different formats, the question turns to which one is better. Better can mean many things, but we will focus on two meanings of better. From the standpoint of a benevolent social planner, better could mean more efficient. We will say that an auction is efficient if the item goes to the person with the highest value. Of course, an individual seller does not necessarily care about social goals such as efficiency, but about the revenue that the auction will generate for himself. The relevant question for the individual seller is then which format generates more revenue. We will look at both of these notions of "better".

3.3.1 Efficiency

We will define the level of efficiency in an auction as $\frac{V_w}{V_H}$, where V_w is the value of the winning bidder and V_H is the value of the high bidder. Note that if the winner is the high bidder, then efficiency is 1 or 100%. The question is, in all of our auction formats will the bidder with the highest value bid more than, less than, or an amount equal to bidders with lower values? It is easy to see that in an ascending clock or 2^{nd}-price sealed bid auction that higher values lead to higher bids because bidders simply submit their values as bids. In the Dutch and 1^{st}-price auctions, the bid function is $b_i = \frac{N-1}{N}v_i$. The question is, who will submit the highest bid? It should be fairly easy to see that higher values will submit higher bids. Technically, we can say that the bid function is increasing in the value draw – as the value draw increases, the bid increases. Thus, bidders with higher values will submit higher bids, and the bidder with the highest value will submit the highest bid. These auctions will also be 100% efficient, assuming that all of our conditions hold and bidders use the Nash equilibrium bidding strategies. Thus, theoretically there is no difference between the efficiency of the 1^{st}-price sealed bid auction, the Dutch auction, the 2^{nd}-price sealed bid auction, or the ascending clock auction

3.3.2 Revenue

As for revenue, we know that the 1^{st}-price and Dutch auctions are strategically equivalent and that the ascending clock auction and the 2^{nd}-price are strategically equivalent. Thus we know that the revenue from the 1^{st}-price and Dutch will be equal and the revenue from the ascending clock auction and the 2^{nd}-price will be equal. The question is, does the 1^{st}-price generate more revenue than the 2^{nd}-price?

Let V_1 be the highest value and V_2 be the second highest value. Then the expected revenue of the 1^{st}-price auction is:

$$Revenue\left(1^{st} - price\right) = \frac{N-1}{N} E\left[V_1\right]$$

The expected revenue of the 2^{nd}-price auction is:

$$Revenue\left(2^{nd} - price\right) = E\left[V_2\right]$$

We will assume that there are the same number of bidders in each auction. We now need to know what $E\left[V_1\right]$ and $E\left[V_2\right]$ are in order to answer which of the auctions will generate more revenue. To do this we use the concept of an order statistic – basically, an order statistic tells us what the expected value of the k^{th} highest draw from a distribution will be given that we make N draws from the distribution. In our case, we are using the uniform distribution over the range 0 to 1. We find that the k^{th} highest value will be equal to:

$$\frac{N-k+1}{N+1}$$

Think about what this means. When there are 2 bidders, on average the highest value draw will be $\frac{2}{3}$, and on average the 2^{nd} highest value draw will be $\frac{1}{3}$. When there are 3 bidders, on average the highest value draw will be $\frac{3}{4}$, the 2^{nd} highest value draw will be $\frac{2}{4}$, and the 3^{rd} highest value draw will be $\frac{1}{4}$. Using this formula we have:

$$E\left[V_1\right] = \frac{N}{N+1}$$

$$E\left[V_2\right] = \frac{N-1}{N+1}$$

Now if we plug these variables into our revenue for the 1^{st} and 2^{nd}-price auctions we get:

$$\begin{aligned} Revenue\left(1^{st} - price\right) &= \frac{N-1}{N} * \left(\frac{N}{N+1}\right) = \frac{N-1}{N+1} \\ Revenue\left(2^{nd} - price\right) &= \frac{N-1}{N+1} \end{aligned}$$

This means that the expected revenue from the 1^{st}-price auction is equal to $\frac{N-1}{N+1}$ and the expected revenue from the 2^{nd}-price auction is also equal to $\frac{N-1}{N+1}$. Thus, both auction formats are expected to generate the same revenue.

While that may surprise you, we have a more powerful result called the revenue equivalence theorem. Essentially, if the conditions of the theorem (laid out below) are met, then any mechanism designed will lead to the same expected revenue.

3.3.3 Revenue Equivalence Theorem

Assume our set-up – SIPV with N risk-neutral agents. Values are drawn from some distribution $F(v)$ that is strictly increasing and atomless on $[\underline{v}, \overline{v}]$. Strictly increasing simply means that the probability of drawing a value less than a number X must be less than the probability of drawing a value less than the number Y if $Y > X$. Atomless is similar to continuous, and really means that the probability of drawing a particular value is 0 (because we are focusing on drawing one particular value out of an infinite number of possible values). Suppose there is one object for sale (this can be extending to the multiple object case as long as no buyer wants more than 1 of k identical, indivisible objects for sale).

If the above conditions are met as well as the following two conditions:

1. The object always goes to the buyer with the highest value (in the multiple object case, the objects go to the bidders with the k highest values)
2. any buyer with value $v = \underline{v}$ expects 0 surplus

then any mechanism that satisfies the above assumptions yields the same expected revenue and results in a buyer with value v_i making the same expected payment across all mechanisms.

3.4 Breaking revenue equivalence and efficiency

If all formats are perfectly efficient and generate the same revenue in expectation, why do auctioneers prefer one type or the other? In the sections below we will look at how to "break" the results from above.

3.4.1 Breaking revenue equivalence

Suppose that instead of risk-neutral agents we had risk-averse bidders. They still have the exact same problem as before – they want to maximize their expected surplus. In the 2^{nd}-price and ascending clock auctions, there was no "maximization" problem – bidders simply submitted their bids or dropped out when the clock reached their value. Thus, the strategy should not change in these types of auctions if bidders are risk averse since they can do no better following another strategy. Since the strategy does not change the expected revenue from the 2^{nd}-price auction is still the same.

Consider the 1^{st}-price auction. Bidders wanted to maximize their expected utility, given by:

$$u_i = b_i * (v_i - b_i)$$

However, in the risk averse case with 2 bidders they want to maximize something like:

$$u_i = b_i * \sqrt{(v_i - b_i)}$$

This maximization problem is slightly more difficult, but nonetheless tractable:

$$\begin{aligned}
\max_{b_i} \ u_i &= b_i * (v_i - b_i)^{1/2} \\
\frac{du_i}{db_i} &= (v_i - b_i)^{1/2} + b_i \left(\frac{1}{2}\right) (v_i - b_i)^{-\frac{1}{2}} (-1) \\
0 &= (v_i - b_i)^{1/2} + b_i \left(\frac{1}{2}\right) (v_i - b_i)^{-\frac{1}{2}} (-1) \\
0 &= (v_i - b_i)^{1/2} - b_i \left(\frac{1}{2}\right) (v_i - b_i)^{-\frac{1}{2}} \\
b_i \left(\frac{1}{2}\right) (v_i - b_i)^{-\frac{1}{2}} &= (v_i - b_i)^{1/2} \\
\frac{1}{2} b_i &= v_i - b_i \\
\frac{3}{2} b_i &= v_i \\
b_i &= \frac{2}{3} v_i
\end{aligned}$$

Recall that when we had 2 risk-neutral bidders in the 1^{st}-price sealed bid auction each bidder used $b_i(v_i) = \frac{1}{2}v_i$. With 2 risk averse bidders, each bidder will bid $\frac{2}{3}$ of his value. Thus, we can see that the bidder is going to bid more in the risk averse case. Intuitively, if the bidder were to bid $\frac{1}{2}$ of his value in the risk averse case the marginal benefit from increasing the bid (the increase in the probability of winning) would be greater than the marginal cost (the amount of surplus lost). So we increase the bid until the marginal benefit of increasing the bid equals the marginal cost, just like we do with many other applications in economics. Note that now in the 2 bidder case the expected revenue is:

$$\begin{aligned}
Revenue\left(1^{sst} - price\right) &= \frac{2}{3} * E\left[V_1\right] \\
\text{where } E\left[V_1\right] &= \frac{2}{3} \\
Revenue\left(1^{sst} - price\right) &= \frac{2}{3} * \frac{2}{3} \\
Revenue\left(1^{sst} - price\right) &= \frac{4}{9}
\end{aligned}$$

In the 2 bidder case for the 2^{nd}-price sealed bid auction the expected revenue is $E[V_2]$, which with 2 bidders and values distributed uniformly on $[0, 1]$ would be $\frac{1}{3}$. So with risk averse bidders the 1^{st}-price and Dutch auctions generate more expected revenue than the 2^{nd}-price and ascending clock.

3.4.2 Breaking efficiency

Suppose we want to break efficiency. The true version of the ascending clock auction has the price moving up continuously with the tick of the clock. However, we know that people do not have continuous values, or, even if they do, there is some rational minimum amount by which their values must increase. In the US the smallest value one can have for a good is a penny, so it is not a stretch to think that the smallest unit in which values can be denominated is a penny. If this is a case, then a clock which moves at the rate of 1 penny per second (or 1 penny per hour or 1 penny per half-second – the rate is not important, but the units that it counts are) will still be perfectly efficient in the sense that the highest valued bidder will get the object. However, consider a clock that increases the price at a rate of 1 penny per second. Now consider the following prices and the corresponding amount of time it will take to auction off objects of these values:

\$10 – 16.67 minutes

\$1 *million* – 3.17 years

\$1 *billion* – 3170 years

It doesn't really seem "efficient" to take 3170 years to auction off an item. In fact, it seems quite inefficient. So what auctioneers will typically do is impose a minimum bid increment. This minimum bid increment is the minimum amount by which the clock will increase (or the minimum amount by which bidders must increase the bid if they wish to place a new bid). While this speeds up the process, the introduction of the minimum increment can also destroy the efficiency results of auctions. For instance, suppose 2 players have values of \$14.08 and \$14.92 respectively. If the clock ticks up at \$1 per second, then both bidders will drop out at \$14. In this case, a tie is declared and we must use the tie-breaking mechanism. The tie-breaking mechanism is usually a coin flip or some other equal probability game. Thus, on average, the bidder with a value of \$14.08 will get the item half of the time. As you can see, the minimum increment introduces the possibility of inefficiency into the auction process.

4 Duopoly model

Now consider a simultaneous pricing game with a few changes. Two firms compete by simultaneously choosing prices (a Bertrand game). If Firms 1 and 2 choose prices p_1 and p_2, respectively, the quantity that consumers demand from each firm is:

$$\begin{aligned}
q_1\left(p_1, p_2\right) &= a - p_1 + bp_2 \\
q_2\left(p_1, p_2\right) &= a - p_2 + bp_1
\end{aligned}$$

Firm 2 has constant marginal cost c. With probability θ Firm 1 has constant marginal cost c_H while with probability $(1-\theta)$ Firm 1 has constant marginal cost c_L. Assume that there are no fixed costs. Prices must be nonnegative ($p_1 \geq 0$, $p_2 \geq 0$) and firms wish to maximize profit. **Important note**: Even though it is a Bertrand game, the demand functions are differentiable and continuous for both firms. This means that we can just use calculus to solve.

The first thing to note is how many types each firm has and what those types are. Firm 2 has a known cost of c, so Firm 2 has one type. However, Firm 1 can have a cost of c_L or c_H, so Firm 1 has two types. The number of types tells us how many maximization problems each firm solves as they need to solve one per type. When setting up the problems be careful to note that the firm with the different types (Firm 1 in this case) will need to specify two (possibly different) actions for its strategy. Thus, we need to have two different price variables for Firm 1 because it has two different costs. Let p_H represent Firm 1's price if it has cost c_H and p_L represent Firm 1's price if it has cost c_L. If Firm 1 has cost c_H it solves:

$$\max_{p_H} \Pi_H = p_H * (a - p_H + bp_2) - c_H * (a - p_H + bp_2)$$

The key is to start from the basic maximization problem and then just proceed from there. Taking the derivative with respect to p_H we have:

$$\begin{aligned}
\frac{\partial \Pi_H}{p_H} &= a - 2p_H + bp_2 + c_H \\
0 &= a - 2p_H + bp_2 + c_H \\
2p_H &= a + bp_2 + c_H \\
p_H &= \frac{a + bp_2 + c_H}{2}
\end{aligned}$$

This gives us Firm 1's best response function if it has cost c_H. Now we do the same thing for Firm 1 if it has cost c_L:

$$\begin{aligned}
\max_{p_L} \Pi_L &= p_L * (a - p_L + bp_2) - c_L * (a - p_L + bp_2) \\
\frac{\partial \Pi_L}{p_L} &= a - 2p_L + bp_2 + c_L \\
2p_L &= a + bp_2 + c_L \\
p_L &= \frac{a + bp_2 + c_L}{2}
\end{aligned}$$

And now we have Firm 1's best response function if its cost is c_L. The last problem to solve is for Firm 2. Note that Firm 2's problem is a little different. If Firm 1 has cost c_H then Firm 2 solves the following problem:

$$\max_{p_2} \Pi_2 = p_2 * (a - p_2 + bp_H) - c * (a - p_2 + bp_H)$$

because Firm 1 will be using $p_1 = p_H$. However, if Firm 1 has cost c_L, Firm 2 will solve:

$$\max_{p_2} \Pi_2 = p_2 * (a - p_2 + bp_L) - c * (a - p_2 + bp_L)$$

because Firm 1 will be using $p_1 = p_L$. The question is how often does it solve the first problem and how often it solves the second problem? Well, Firm 1 has cost c_H with probability of θ, so Firm 2 solves the first problem θ percent of the time. Firm 1 has cost c_L with probability $1-\theta$, so Firm 2 solves the second problem $(1-\theta)$ percent of the time. Putting these together we get:

$$\max_{p_2} \Pi_2 = \theta * (p_2 * (a - p_2 + bp_H) - c * (a - p_2 + bp_H)) + (1-\theta) * (p_2 * (a - p_2 + bp_L) - c * (a - p_2 + bp_L))$$

Yes, Firm 2 has a slightly more complicated problem but the derivative is not that difficult, it's just long.

We have:

$$
\begin{aligned}
\frac{\partial \Pi_2}{\partial p_2} &= \theta * (a - 2p_2 + bp_H + c) + (1-\theta) * (a - 2p_2 + bp_L + c) \\
0 &= \theta * (a - 2p_2 + bp_H + c) + (1-\theta) * (a - 2p_2 + bp_L + c) \\
0 &= a\theta - 2p_2\theta + bp_H\theta + c\theta + a - a\theta - 2p_2 + 2p_2\theta + bp_L - bp_L\theta + c - c\theta \\
0 &= bp_H\theta + a - 2p_2 + bp_L - bp_L\theta + c \\
2p_2 &= a + bp_H\theta + (1-\theta)\, bp_L + c \\
p_2 &= \frac{a + bp_H\theta + (1-\theta)\, bp_L + c}{2}
\end{aligned}
$$

This is Firm 2's best response function. Note that it is very similar to Firm 1's best response function, which is derived for the case where it knows Firm 2's cost with certainty, only with Firm 2's best response function we need to take into consideration how often Firm 1 chooses p_H and how often it chooses p_L. Now, there are 3 unknowns in this problem (p_H, p_L, p_2), and we have 3 best response functions so we just need to solve for each one. To make it easier let's use the following parameters: $\theta = \frac{1}{2}$, $a = 108$, $b = \frac{1}{2}$, $c_H = 16$, $c_L = 8$, and $c = 12$. Thus we have the following 3 equations:

$$
\begin{aligned}
p_H &= \frac{108 + \frac{1}{2}p_2 + 16}{2} \\
p_L &= \frac{108 + \frac{1}{2}p_2 + 8}{2} \\
p_2 &= \frac{108 + \frac{1}{2}p_H\frac{1}{2} + \left(1 - \frac{1}{2}\right)\frac{1}{2}p_L + 12}{2}
\end{aligned}
$$

Simplifying further we have:

$$
\begin{aligned}
p_H &= \frac{124 + \frac{1}{2}p_2}{2} \\
p_L &= \frac{116 + \frac{1}{2}p_2}{2} \\
p_2 &= \frac{120 + \frac{1}{4}p_H + \frac{1}{4}p_L}{2}
\end{aligned}
$$

Now substituting p_H and p_L into the equation for p_2 we have:

$$
\begin{aligned}
p_2 &= \frac{120 + \frac{1}{4}\left(\frac{124+\frac{1}{2}p_2}{2}\right) + \frac{1}{4}\left(\frac{116+\frac{1}{2}p_2}{2}\right)}{2} \\
2p_2 &= 120 + \frac{1}{4}\left(\frac{124 + \frac{1}{2}p_2}{2}\right) + \frac{1}{4}\left(\frac{116 + \frac{1}{2}p_2}{2}\right) \\
2p_2 &= 120 + \frac{124 + \frac{1}{2}p_2}{8} + \frac{116 + \frac{1}{2}p_2}{8} \\
16p_2 &= 960 + 124 + \frac{1}{2}p_2 + 116 + \frac{1}{2}p_2 \\
16p_2 &= 1200 + p_2 \\
15p_2 &= 1200 \\
p_2 &= 80
\end{aligned}
$$

Using that $p_2 = 80$ we can easily find p_H and p_L:

$$
\begin{aligned}
p_H &= \frac{124 + \frac{1}{2}p_2}{2} \\
p_H &= \frac{124 + \frac{1}{2} * 80}{2} \\
p_H &= \frac{124 + 40}{2} \\
p_H &= 82 \\
& \quad \text{and} \\
p_L &= \frac{116 + \frac{1}{2}p_2}{2} \\
p_L &= \frac{116 + \frac{1}{2} * 80}{2} \\
p_L &= \frac{116 + 40}{2} \\
p_L &= 78
\end{aligned}
$$

Thus, our Bayes-Nash equilibrium for this game is: Firm 2 charges a price $p_2 = 80$; Firm 1 charges a price $p_L = 78$ if its cost is c_L; and Firm 1 charges a price $p_H = 82$ if its cost is c_H.