Static Games Of Incomplete Information

Objective: Understand what a game of incomplete information is.

In the games we have explored to date, every player knew the payoff functions of all players. We refer to these games as games of complete information. What distinguishes games of incomplete information is the assumption that at least one player does not know the payoff function of another. That is, there is some information unavailable to a player that influences an opponent's payoff. For example, a firm may not be aware of the cost of production of a competitor. A salesman may not know how much you value the new car he is trying to sell you. Insurance companies may not know whether a driver is careful or reckless. In each of these cases, there is information that is unknown to one party and likely to influence the behavior of another. As we will see, this makes the analysis of games at least a little more difficult.

Objective: Understand how static games of incomplete information can be formulated so we can apply much of what we have already learned to understand what people will or should do.

To deal with the problems introduced by incomplete information Harsanyi proposed simply adding a new player to the game. The player is referred to as Nature. Nature gets to start the game by randomly choosing the state of the world. Under different states of the world, strategy profiles can correspond to different payoffs for some or all players. Nature then reveals some information regarding the state of the world to each player. An important characteristic of these types of games is that Nature does not reveal the same information about the state of the world to all players. Therefore, some or all players can have private information.

Specifically, Harsanyi suggested characterizing the different states of the world and a player's private information by defining player types. Nature chooses a player's type and then reveals this type to the player, but not his opponents. A player's payoff function then depends on its type.

Recall that we can describe a complete information static game G by

1. Players: $i = 1, 2, \ldots, n;\ i \in N$
2. Strategies: s_i is a pure strategy (action) for player i, $s_i \in S_i$ is set of possible pure strategies for player i. A strategy profile is defined as $s = (s_1, s_2, \ldots, s_n)$ where $s \in S = \times_{i \in N} S_i$ is the set of all possible strategy profiles.
3. Utility (preferences): $U_i(s)$, describes player i's ranking of possible outcomes. $U = \{U_1(s), U_2(s), \ldots, U_n(s)\}$

such that $G = \{N, S, U\}$.

To extend this description to include static games of incomplete information, we also need

4. Types: t_i is a type for player i, $t_i \in T_i$ is the set of all possible types for player i. The state of the world can be defined as a vector of types, $t = (t_1, t_2, \ldots, t_n)$, one for each player where the set of all possible states of the world is $t \in T = \times_{i \in N} T_i$. We will also find it

useful to define the state of the world exclusive of player i's type, $t_{\sim i} = (t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_n)$, and all possible states of the world exclusive of player i's type, $T_{\sim i} = \times_{i \neq j \in N} T_j$.

5. Probabilities: The probability of any given state of the world is $p = p(t)$ where $1 > p(t) > 0$ and $\sum_{t \in T} p(t) = 1$. The probability that i's type is t_i is defined as $p_i(t_i) = \sum_{t_{\sim i} \in T_{\sim i}} p(t)$. The conditional probability of the state of the world being t given i's type is t_i is $p(t \mid t_i)$. Bayes rule implies $p(t \mid t_i) = p(t)/p_i(t_i)$. $P = \{p(t) \mid t \in T\}$ is then just the probability distribution over all states of the world. Note that all this can be generalized to handle continuous type spaces.

Finally, we need to modify our definition of strategies and payoffs:

2'. Strategies: $s_i(t_i)$ is a pure strategy for player i given t_i. $s_i(t_i) \in S_i$ is the set of possible pure strategies for player i given t_i. A strategy profile conditional on the state of the world is $s(t) = (s_1(t_1), s_2(t_2), \ldots, s_n(t_n))$. A pure strategy for player i is a collection of strategies, one for each type: $s_i = \{s(t_i) \mid t_i \in T_i\}$. A strategy profile and the set of all strategy profiles can then be defined just as before.

3'. Utility (preferences): $U_i(s) = \sum_{t \in T} U_i(s(t), t_i)p(t)$ is the expected utility for player i. It is important to note that for this expected utility the strategy profile is conditional on the state of the world and player i's type. With this modified definition of an individual's payoff function, the payoff space can be defined just as before. Note that if we don't like expected utility theory we can always modify our definition of a player's payoff function. Also note that as with mixed strategies, the intensity of preferences matters, not just the ordinal ranking.

With these additions, we can now define a static game of incomplete information as $G = \{N, S, U, T, P\}$.

Objective: Understand that the Bayesian Nash equilibrium is just a Nash equilibrium for a static game of incomplete information.

The Nash equilibrium requires that no individual player can do better by unilaterally changing its strategy. The Bayesian Nash equilibrium requires that no type of an individual player can do better by unilaterally changing its strategy.

A pure strategy Bayesian Nash equilibrium is just a Nash equilibrium for the game $G = \{N, S, U, T, P\}$: $s^* \in S$ is a **pure strategy Nash equilibrium** for $G = \{N, S, U, T, P\}$ if for all players $i \in N$, $U_i(s_i^*, s_{\sim i}^*) \geq U_i(s_i, s_{\sim i}^*)$ for all $s_i \in S_i$.

So, why make a distinction?

The reason for the distinction is that the game $G = \{N, S, U, T, P\}$ imposes an important restriction that makes it a subset of a more general class of static games of incomplete information. This subset is referred to as Bayesian games.

The restriction that is important to note is that types are defined so that a player's payoff function is separable in its types:

$$U_i(s_i, s_{\sim i}) = \sum_{t \in T} U_i(s(t), t_i)\, p(t) = \sum_{t_i \in T_i} \left\{ \left[\sum_{t_{\sim i} \in T_{\sim i}} U_i(s(t), t_i) \frac{p(t)}{p_i(t_i)} \right] p_i(t_i) \right\}.$$

Therefore, to optimize $U_i(s_i, s_{\sim i})$, we can optimize $U_i(s_i(t_i), s_{\sim i}) = \sum_{t_{\sim i} \in T_{\sim i}} U_i(s(t), t_i) \frac{p(t)}{p_i(t_i)}$ individually for each type. Typically, the definition for the Bayesian Nash equilibrium distinguishes itself from the definition of a Nash equilibrium by making this restriction more transparent:

Definition: $s^* \in S$ is a **pure strategy Bayesian Nash equilibrium** for $G = \{N, S, U, T, P\}$ if for all players $i \in N$, $s_i(t_i)$ solves $\max_{s_i(t_i) \in S_i} U_i(s_i(t_i), s_{\sim i}^*)$ for all $t_i \in T_i$.

Example:

Consider the game where we have two players A and B. Furthermore, suppose player B has an equal probability of being one of two types: Type 1 (Cooperative) and Type 2 (Uncooperative). Both types have two strategies either *L* or *R*. The payoffs for each type are summarized in the tables below.

Player B, Type 1

		L	*R*
Player A	*U*	25 30	0 0
	D	0 0	25 20

Player B, Type 2

		L	*R*
Player A	*U*	0 30	25 0
	D	25 0	0 20

Question: What is the pure strategy Bayesian Nash equilibrium for this game?

To answer this question, we want to start by defining the strategy space for each player: $S_A = \{U, D\}$ and $S_B = \{(L, L), (L, R), (R, L), (R, R)\}$. The space of strategy profiles is then $S = S_A \times S_B$.

Next we need to define the payoffs over strategy profiles. This can be accomplished in a table:

		Player B (Type 1, Type 2)			
		(L, L)	(L, R)	(R, L)	(R, R)
Player A	U	(25,0) 30/2 + 30/2 = 30	(25,25) 30/2 + 0/2 =15	(0,0) 0/2 + 30/2 = 15	(0,25) 0/2 + 0/2 = 0
	D	(0,25) 0/2 + 0/2 = 0	(0,0) 0/2 + 20/2 = 10	(25,25) 20/2 + 0/2 = 10	(25,0) 20/2 + 20/2 = 20

We can now look for which strategies are best responses for each player, given the strategy played by all other players. For player A, *U* is a best response when $s_B = (L, L)$, (L, R), or (R, L) and *D* is a best response when $s_B = (R, R)$. Player B's best response depends on its type. For type 1, *L* is a best response when player A chooses *U* and *R* is a best response when player A chooses *D*. For type 2, *L* is a best response when player A chooses *D* and *R* is a best response when player A chooses *U*.

We can now denote each player's best responses by putting asterisks next to the corresponding payoffs and then look for a strategy profile where there is an asterisk next to the payoff for player A and the payoff to both types of player B:

		Player B (Type 1, Type 2)			
		(L, L)	(L, R)	(R, L)	(R, R)
Player A	U	(25*,0) 30*	(25*,25*) 15*	(0,0) 15*	(0,25*) 0
	D	(0,25*) 0	(0,0) 10	(25*,25*) 10	(25*,0) 20*

The only strategy profile that is a mutual best response for all players is $s = \{U, (L, R)\}$, which is the Bayesian Nash equilibrium strategy.

Application: The Cournot Duopoly Model Revisited

Suppose there are two firms denoted as 1 and 2 that produce a homogeneous product. Demand is denoted by $p = a - q_1 - q_2$ where *p* is the price of the product and q_i is the output of the *i*th firm. Suppose each firm can have either a high or low marginal cost denoted by c_H and c_L. Each firm knows its cost but not its competitor's cost. Furthermore, suppose that the probability that a firm has a high marginal cost is α such that the probability a firm has low costs is $1 - \alpha$ and that each firm's cost is drawn independently.

Question: How much should each firm produce?

To answer this question we must first ask what is each player's type space: $T_1 = T_2 = (c_H, c_L)$. The states of nature are then $T = \{(c_H, c_H), (c_L, c_H), (c_H, c_L), (c_L, c_L)\}$. The probability distribution for this set of states of nature is $P = \{\alpha^2, \alpha(1- \alpha), (1- \alpha)\alpha, (1- \alpha)^2\}$. A strategy for firm *i* is going to be an output level for each possible type nature can define: $s_i = \{q_i(c_H), q_i(c_L)\}$.

Now we need to know each firm's payoff. Since we know payoffs are separable by type, let c_i be i's realized cost. The expected profit for i assuming risk neutrality is

$$\begin{aligned}\pi_i &= a[(a - q_i - q_j(c_H))q_i - c_i q_i] + (1 - a)[(a - q_i - q_j(c_L))q_i - c_i q_i] \\ &= (a - c_i - q_i - q_j(c_L))q_i + (q_j(c_L) - q_j(c_H))aq_i\end{aligned}$$

The first order condition for an interior solution is

$$a - c_i - 2q_i - q_j(c_L) + (q_j(c_L) - q_j(c_H))a = 0$$

such that

$$q_i(c_i)^* = (a - c_i - aq_j(c_H) - (1 - a)\, q_j(c_L))/2$$

is player i's best response given its type.

If $c_i = c_H$,

$$q_i(c_H)^* = (a - c_H - aq_j(c_H) - (1 - a)\, q_j(c_L))/2.$$

If $c_i = c_L$,

$$q_i(c_L)^* = (a - c_L - aq_j(c_H) - (1 - a)\, q_j(c_L))/2.$$

Therefore, we have four equations,

$$\begin{aligned}q_i(c_H)^* &= (a - c_H - aq_j(c_H)^* - (1 - a)\, q_j(c_L)^*)/2, \\ q_i(c_L)^* &= (a - c_L - aq_j(c_H)^* - (1 - a)\, q_j(c_L)^*)/2, \\ q_j(c_H)^* &= (a - c_H - aq_i(c_H)^* - (1 - a)\, q_i(c_L)^*)/2, \text{ and} \\ q_j(c_L)^* &= (a - c_L - aq_i(c_H)^* - (1 - a)\, q_i(c_L)^*)/2,\end{aligned}$$

and four unknowns,

$q_i(c_H)^*$, $q_i(c_L)^*$, $q_j(c_H)^*$, and $q_j(c_L)^*$.

A lot of tedious algebra (hopefully, no mistakes) yields:

$$q_i(c_i)^* = \frac{2a - 3c_i + ac_H + (1-a)c_L}{6} = \frac{a - (ac_H + (1-a)c_L)}{3} + \frac{ac_H + (1-a)c_L - c_i}{2}$$

or

$$q_i(c_i)^* = \frac{a - 2c_i + c_j}{3} + \frac{E(c_j) - c_j}{3} + \frac{c_i - E(c_i)}{6}$$

where $E(c_i) = E(c_j) = ac_H + (1 - a)c_L$.

This second solution is nice because it shows how private information on a firms' marginal cost influences equilibrium output. The first expression on the right-hand-side is the Nash equilibrium output for the complete information game with costs c_i and c_j. The second and third expressions show how equilibrium output is different when firms have private information.

The second term is positive (negative) when *i*'s competitor's marginal cost is lower (higher) than the expected marginal cost. If *i* expects its competitors marginal cost is lower (higher) than it is, it produces less (more) than if it knew for certain because it expects its competitor to tend to produce more (less).

The third term is positive (negative) when *i*'s marginal cost is higher (lower) than the expected marginal cost. If *i* marginal cost is higher (lower) than its competitor is expecting, its competitor tends to produce more (less) than it would if it knew for certain for the reasons stated above. Therefore, *i* should tend to produce less (more) than it would if its competitor knew its cost for certain.

If we look at total equilibrium output, we get

$$q_i(c_i)^* + q_j(c_j)^* = \frac{2a - c_i - c_j}{3} + \frac{E(c_j) - c_j}{6} + \frac{E(c_i) - c_i}{6},$$

which implies incomplete information leads to higher equilibrium output when expected costs are higher than actual cost for both firms and lower equilibrium output when expected costs are lower than actual costs for both firms. The results is ambiguous when one firm's actual cost is higher than expected and the others is lower than expected.

Application: Rent Seeking With One-Sided Asymmetric Information

The first two games we considered focused on discrete types. All of this works equally well with continuous types. Let us see how by extending the classic rent seeking game. Let us also use this game to talk about how information asymmetries influence rent dissipation.

In the standard rent seeking game, we have two players. Let us denote them as *I* for the informed player and *U* for the uninformed player. The ith player can expend costly effort equal to x_i to influence the outcome of a contest that earns it V_i. To keep things simple, we will assume the probability player *i* wins the contest is $P_i(x_i, x_j) = x_i / (x_i + x_j)$. Now for the twist: let us assume that both players know there own value of winning, player *I* knows the value of winning to player *U*, but player *U* does not know the value of winning to player *I*. Let *v* be a random variable with distribution function $F(v)$ and support $v_H \geq v \geq v_L$. Both players know nature uses $F(v)$ to randomly draw player *I*'s value.

Player *I*'s payoff can be written as $p_I = V_I \dfrac{x_I}{x_I + x_U} - x_I$, which leads to the best response function

$$x_I(V_I) = \begin{cases} \sqrt{V_I x_U} - x_U & \text{for } 0.0 < x_U \le V_I \\ 0, & \text{otherwise} \end{cases}.$$

Now what about player *U*? Player *U*'s best payoff is going to depend on player *I*'s effort, which will ultimately depend on player *I*'s type. Therefore, we can write player *U*'s expected payoff as $p_U = \int_{v_L}^{v_U} \left(V_U \dfrac{x_U}{x_I(v) + x_U} - x_U \right) F'(v) dv$. The first order condition for an interior solution is $\int_{v_L}^{v_U} \left(V_U \dfrac{x_I(v)}{(x_I(v) + x_U)^2} - 1 \right) F'(v) dv = 0$, which cannot be solved explicitly for *U*'s best response. However, we can use *I*'s best response and this first order condition to still find the Bayesian Nash equilibrium. To avoid tedium, let us assume *I*'s best response is interior. We can then write *U*'s first order condition as $\int_{v_L}^{v_U} \left(V_U \dfrac{\sqrt{v x_U} - x_U}{v x_U} \right) F'(v) dv = 1$. At this point it is convenient to define $r = \sqrt{\dfrac{V_U}{v}}$ and $G(r)$ as the transformed distribution of $F(v)$, so we can write the first order condition as $\int_{r_L}^{r_U} \left(\sqrt{\dfrac{V_U}{x_U}} r - r^2 \right) G'(r) dr = 1$. Now it is relatively easy to solve for $x_U = V_U \left(\dfrac{m}{1 + s^2 + m^2} \right)^2$ where $m = \int_{r_L}^{r_U} r G'(r) dr$ and $s^2 = \int_{r_L}^{r_U} r^2 G'(r) dr - m^2$. Using *I*'s best response function yields $x_I(v) = vr \left(\dfrac{m}{1 + s^2 + m^2} \right) \left(1 - \dfrac{m}{1 + s^2 + m^2} \right)$. Finally, rent dissipation is $x_U + x_I(v) = \sqrt{v V_U} \dfrac{m}{1 + s^2 + m^2}$.

Now if we set $s^2 = 0$ we get all of our standard complete information rent seeking results. Therefore, for $s^2 > 0$, we see that the uninformed player expends less effort. Why? There is also less rent dissipation. The informed player will increase effort if $\dfrac{m}{1 + s^2 + m^2} > \dfrac{1}{2}$ and decrease effort if $\dfrac{1}{2} > \dfrac{m}{1 + s^2 + m^2}$.

Application: Sealed Bid Auction

The previous application showed how we could apply the Bayesian Nash equilibrium to a case with continuous types, but it made one important simplifying assumption that made life much easier. The assumption was that only one player had incomplete information. Think for a moment about how much more complicated the problem would have been if both players had incomplete information. Now let us take a look at another application where both players have incomplete information with continuous types.

Consider two risk neutral buyers bidding for a good. The value of the good to each buyer is private information. Each buyer submits its bid in a sealed envelope. The buyer who bids the highest gets the good for the price it bid. The low bidder gets nothing. If there is a tie, each bidder gets half the difference between its value and bid. Suppose each bidders value is drawn randomly and independently from the continuous probability distribution $F(v)$ with support v_L and $v_{H.}$

Question: What is the Bayesian Nash equilibrium bid strategy?

Buyer *i*'s strategy is a bid that is dependent on its realized value: $b_i(v_i)$.

Buyer *i*'s expected payoff is

$$\begin{aligned} U_i\left(b_i, b_j\left(v_j\right)\right) &= \left(v_i - b_i\right)\Pr\left(b_i > b_j\left(v_j\right)\right) + \frac{\left(v_i - b_i\right)}{2}\Pr\left(b_i = b_j\left(v_j\right)\right) + 0\Pr\left(b_i < b_j\left(v_j\right)\right) \\ &= \left(v_i - b_i\right)\Pr\left(b_i > b_j\left(v_j\right)\right) + \frac{\left(v_i - b_i\right)}{2}\Pr\left(b_i = b_j\left(v_j\right)\right) \end{aligned} .$$

Since our players are identical, lets assume they use the same bid strategy: $b_i = b_i(v_i) = b(v_i)$. Furthermore, let us assume this bid strategy is continuous, differentiable, and monotonic, which allows us to invert it: $v_i = b^{-1}(b_i)$. Since $\Pr(b^{-1}(b_i) = v_j) = 0$ for a continuous distribution, the payoff function becomes:

$$U_i\left(b_i, b_j\left(v_j\right)\right) = \left(v_i - b_i\right)\Pr\left(b^{-1}\left(b_i\right) > v_j\right) = \left(v_i - b_i\right)F\left(b^{-1}\left(b_i\right)\right).$$

Player ***i*** with value v_i maximizes its expected payoff by solving:

$$-F\left(b^{-1}\left(b_i\right)\right) + \left(v_i - b_i\right)F'\left(b^{-1}\left(b_i\right)\right)\frac{\partial b^{-1}\left(b_i\right)}{\partial b_i} = 0 .$$

It is not too difficult to show that $\frac{\partial b^{-1}\left(b_i\right)}{\partial b_i} = \frac{1}{b'\left(v_i\right)}$:

Let $v_i = b^{-1}(b_i)$ and $b_i = b(v_i)$. Then $dv_i/db_i = \frac{\partial b^{-1}\left(b_i\right)}{\partial b_i}$ and $db_i/dv_i = b'\left(v_i\right)$, so $\frac{1}{b'\left(v_i\right)} = dv_i/db_i =$ $\frac{\partial b^{-1}\left(b_i\right)}{\partial b_i}$.

So we can rewrite the first order condition as

$$b'(v_i)F(v_i)+b(v_i)F'(v_i)=v_iF'(v_i)$$

or

$$\frac{\partial b(v_i)F(v_i)}{\partial v_i}=v_iF'(v_i).$$

Now if we integrate both sides up to v, we get

$$\int_{v_L}^{v}\frac{\partial b(v_i)F(v_i)}{\partial v_i}dv_i=\int_{v_L}^{v}v_iF'(v_i)dv_i$$

$$b(v)F(v)-b(v_L)F(v_L)=\int_{v_L}^{v}v_iF'(v_i)dv_i$$

$$b(v)=\int_{v_L}^{v}v_i\frac{F'(v_i)}{F(v)}dv_i$$

Note that $F'(v_i)$ is the density function evaluated at v_i. $F(v)$ is the probability v_i is less than v. Therefore, $F'(v_i)/F(v)$ is the probability of v_i given $v_i < v$: $\Pr(v_i \mid v_i < v)$. $\int_{v_L}^{v}v_i\frac{F'(v_i)}{F(v)}dv_i$ can then be interpreted as the expected value of v_i given v_i is less than v. So a player's bid should equal the average of its opponent's value given its opponent's value is lower. An important implication of this result is that buyers have the incentive to bid less than they value the item.

For the experiment in class, $F(v)$ was a uniform distribution with $v_L = 25$ and $v_H = 50$ such that $F'(v_i) = 1/25$ and $F(v) = v/25 - 1$. Therefore, the Bayesian Nash (symmetric, continuous, and monotonic) equilibrium bid strategy is $b(v) = (v^2 - 25^2)/2(v - 25)$.

Here is how students did in the Fall of 2003:

Value	Bid	Predicted	Percent Error
27	15	26	-42.3%
36	0	30.5	-100.0%
45	40	35	14.3%
40	35	32.5	7.7%
38	37	31.5	17.5%
27	24	26	-7.7%
34	20	29.5	-32.2%
31	30	28	7.1%

Application: Double Auction

The last application considered two buyers competing in an auction to secure some good. The players in this auction had private information about how they valued the good, but the same beliefs about how their competitors valued the good. With this we argued that the players would behave identically, which facilitated finding a solution to the game. The game we will look at now has only two players. However, we will complicate things by assuming one is selling something that the other wants to buy. Therefore, we will not be able to assume the players will behave identically. As before, we will assume that how each player values the good is private information. In this game, the seller will submit an asking price and the buyer will submit a bid simultaneously. If the buyer bids at least as much as the seller is asking, trade will occur at the average price. If the buyer submits a bid lower than the asking price, no trade will occur.

To formalize this set up, we have two players, a buyer denoted by b and a seller denoted by s. Let v_b be the value to the buyer of the good in question. We will also use this value to denote the buyer's type. Let v_s be the value to the seller of the good in question and the seller's type. Assume v_b and v_s are uniformly distributed on the unit interval. We can specify the buyer's strategy as $P_b(v_b)$ and the seller's strategy as $P_s(v_s)$. If trade occurs, $P_b(v_b) \geq P_s(v_s)$, the price paid by the buyer and received by the seller is $P(v_b, v_s) = \frac{P_b(v_b) + P_s(v_s)}{2}$. Therefore, the buyer's risk neutral payoff is

$$U_b(v_b) = \begin{cases} v_b - P(v_b, v_s), & \mathit{for}\ P_b(v_b) \geq P_s(v_s) \\ 0, & \mathit{otherwise} \end{cases},$$

while the seller's is

$$U_s(v_s) = \begin{cases} P(v_b, v_s) - v_s, & \mathit{for}\ P_b(v_b) \geq P_s(v_s) \\ 0, & \mathit{otherwise} \end{cases}.$$

To find a Bayesian Nash equilibrium for this game we need to find $P_b(v_b)^*$ and $P_s(v_s)^*$ such that

$$\max_{P_b} \left(v_b - E\left(\frac{P_b + P_s(v_s)^*}{2} \mid P_b \geq P_s(v_s)^* \right) \right) \Pr\left(P_b \geq P_s(v_s)^*\right)$$

and

$$\max_{P_s} \left(E\left(\frac{P_b(v_b)^* + P_s}{2} \mid P_b(v_b)^* \geq P_s \right) - v_s \right) \Pr\left(P_b(v_b)^* \geq P_s\right).$$

This is not a trivial problem because there are lots of potential forms $P_b(v_b)^*$ and $P_s(v_s)^*$ could take. To start, let us assume that $P_b(v_b)^*$ and $P_s(v_s)^*$ are of a particularly simple form:

$$P_b(v_b)^* = \begin{cases} x, & for\ v_b \geq x \\ 0, & otherwise \end{cases}$$

and

$$P_s(v_s)^* = \begin{cases} x, & for\ x \geq v_s \\ 1, & otherwise \end{cases}$$

where $x \in [0, 1]$.

Question: For what values of x, if any, will these strategies be a Bayesian Nash equilibrium?

First, let us consider the buyer. Suppose $v_s \geq x$, which implies $P_s(v_s)^* = 1$. This also implies $U_b(v_b) = 0$ for all $x \in [0, 1]$. If the buyer deviates by choosing $x' \neq x$, the only way a trade could take place is if $x' = 1$ and $v_b = 1$, which yields 0. If trade doesn't take place ($x' \neq 1$ or $v_b \neq 1$) the payoff is again 0. Therefore, there is no way the buyer can deviate from x and increase its payoff. But, what if $v_s \leq x$, which means there are prospects for a trade. If $v_b \geq x$, the buyer's payoff is $v_b - x$ in equilibrium. If the buyer deviates by choosing $v_b \geq x' > x$, its payoff is $v_b - (x' + x)/2 < v_b - x$. If the buyer deviates by choosing $x' > v_b$, trade doesn't take place and its payoff is $0 < v_b - x$. If the buyer deviates by choosing $x > x'$, trade again cannot take place and it payoffs is $0 < v_b - x$. Alternatively, if $x > v_b$, then the buyer's payoff is 0 in equilibrium. If the buyer deviates by choosing $x' > x$ or $x > x'$, trade will not take place and its payoff is still 0. Therefore, $P_b(v_b)^* = \begin{cases} x, & for\ v_b \geq x \\ 0, & otherwise \end{cases}$ is a best response for the buyer given $P_s(v_s)^* = \begin{cases} x, & for\ x \geq v_s \\ 1, & otherwise \end{cases}$. Similar arguments can be used to show that $P_s(v_s)^* = \begin{cases} x, & for\ x \geq v_s \\ 1, & otherwise \end{cases}$ is a best response to $P_b(v_b)^* = \begin{cases} x, & for\ v_b \geq x \\ 0, & otherwise \end{cases}$. So, this is a Bayesian Nash equilibrium.

You will note that the only thing we said about x is that it lies between 0 and 1. Therefore, we have actually identified a multitude of Bayesian Nash equilibria, which may be a bit disheartening.

Figure 1 illustrates the implications of this equilibrium and another reason why the results are rather disheartening. In Figure 1, there is a large region where trades do not occur (Efficient With No Trades), but this is fine because $v_s > v_b$, which means any trades would be inefficient anyway. There is also a region where trades occur that are efficient (Efficient With Trades) because $v_b > v_s$. Unfortunately, there are also two large regions where no trades occur even though they would be efficient because $v_b > v_s$ (Inefficient With No Trades). Therefore, this equilibrium can result in inefficient outcomes.

We were able to find lots of different Bayesian Nash equilibria with some very simple assumptions regarding the nature of players' equilibrium strategies. Unfortunately, this is not the end of the story. Suppose buyers and sellers used linear strategies:

$P_b(v_b)^* = a_b + c_b v_b$ and

$P_s(v_s)^* = a_s + c_s v_s$.

Question: Can we find linear strategies that form a Bayesian Nash equilibrium?

The answer to this question is yes. To see why, let us start with the buyer. If v_s is uniformly distributed on the unit interval,

$$\Pr(P_b > P_s(v_s)^*) = \int_0^{\frac{P_b - a_s}{c_s}} dv_s = \frac{P_b - a_s}{c_s} \text{ and}$$

$$E\left(\frac{P_b + P_s(v_s)^*}{2} \mid P_b \geq P_s(v_s)^*\right) = \int_0^{\frac{P_b - a_s}{c_s}} \frac{P_b + a_s + c_s v_s}{2} \frac{1}{\frac{P_b - a_s}{c_s}} dv_s = \frac{3P_b + a_s}{4}.$$

Therefore, our optimization problem becomes

$$\max_{P_b} \left(v_b - \frac{3P_b + a_s}{4}\right)\left(\frac{P_b - a_s}{c_s}\right),$$

which has the first order condition

$$-\frac{3(P_b - a_s)}{4c_s} + \left(v_b - \frac{3P_b + a_s}{4}\right)\frac{1}{c_s} = 0 \text{ or}$$

$$P_b(v_b)^* = \frac{2}{3}v_b + \frac{1}{3}a_s.$$

Now, let us turn to the seller. If v_b is uniformly distributed on the unit interval,

$$\Pr(P_b(v_b)^* > P_s) = \int_{\frac{P_s - a_b}{c_b}}^{1} dv_b = 1 - \frac{P_s - a_b}{c_b} = \frac{c_b - P_s + a_b}{c_b} \text{ and}$$

$$E\left(\frac{P_b(v_b)*+P_s}{2} \mid P_b(v_b)* \geq P_s\right) = \int_{\frac{P_s - a_b}{c_b}}^{1} \frac{a_b + c_b v_b + P_s}{2} \frac{1}{\frac{c_b - P_s + a_b}{c_b}} dv \quad = \left(\frac{a_b + 3P_s + c_b}{4}\right).$$

Therefore, our optimization problem becomes

$$\max_{P_s} \left(\frac{a_b + 3P_s + c_b}{4} - v_s\right)\left(\frac{c_b - P_s + a_b}{c_b}\right),$$

which has the first order condition

$$\frac{3(c_b - P_s + a_b)}{4c_b} - \left(\frac{a_b + 3P_s + c_b}{4} - v_s\right)\frac{1}{c_b} = 0 \text{ or}$$

$$P_s(v_s)* = \frac{c_b + a_b}{3} + \frac{2}{3}v_s .$$

Combining these two results yields $a_b = \frac{1}{3}a_s$, c_b = 2/3, $a_s = \frac{c_b + a_b}{3}$, and c_s = 2/3, which imply $P_b(v_b)* = \frac{1}{12} + \frac{2}{3}v_b$ and $P_s(v_s)* = \frac{1}{4} + \frac{2}{3}v_s$. Note that with these results, trade occurs when $P_b(v_b)* \geq P_s(v_s)*$ or when $v_b \geq \frac{1}{4} + v_s$.

Figure 2 show the implications of this Bayesian Nash equilibrium. As before, there is a significant region where no trade occurs even though trade would be efficient. Indeed, Myerson and Satterthwaite (1983) shows that there will always be a region where trade should occur and it doesn’t in a double auction as long as there is some probability that no trade should occur. They go further and also show that the linear Bayesian Nash strategy is the best of all possible Bayesian Nash strategies in the sense of minimizing the inefficiency.

Figure 1: Inefficiencies of a fixed price strategy in the double auction.

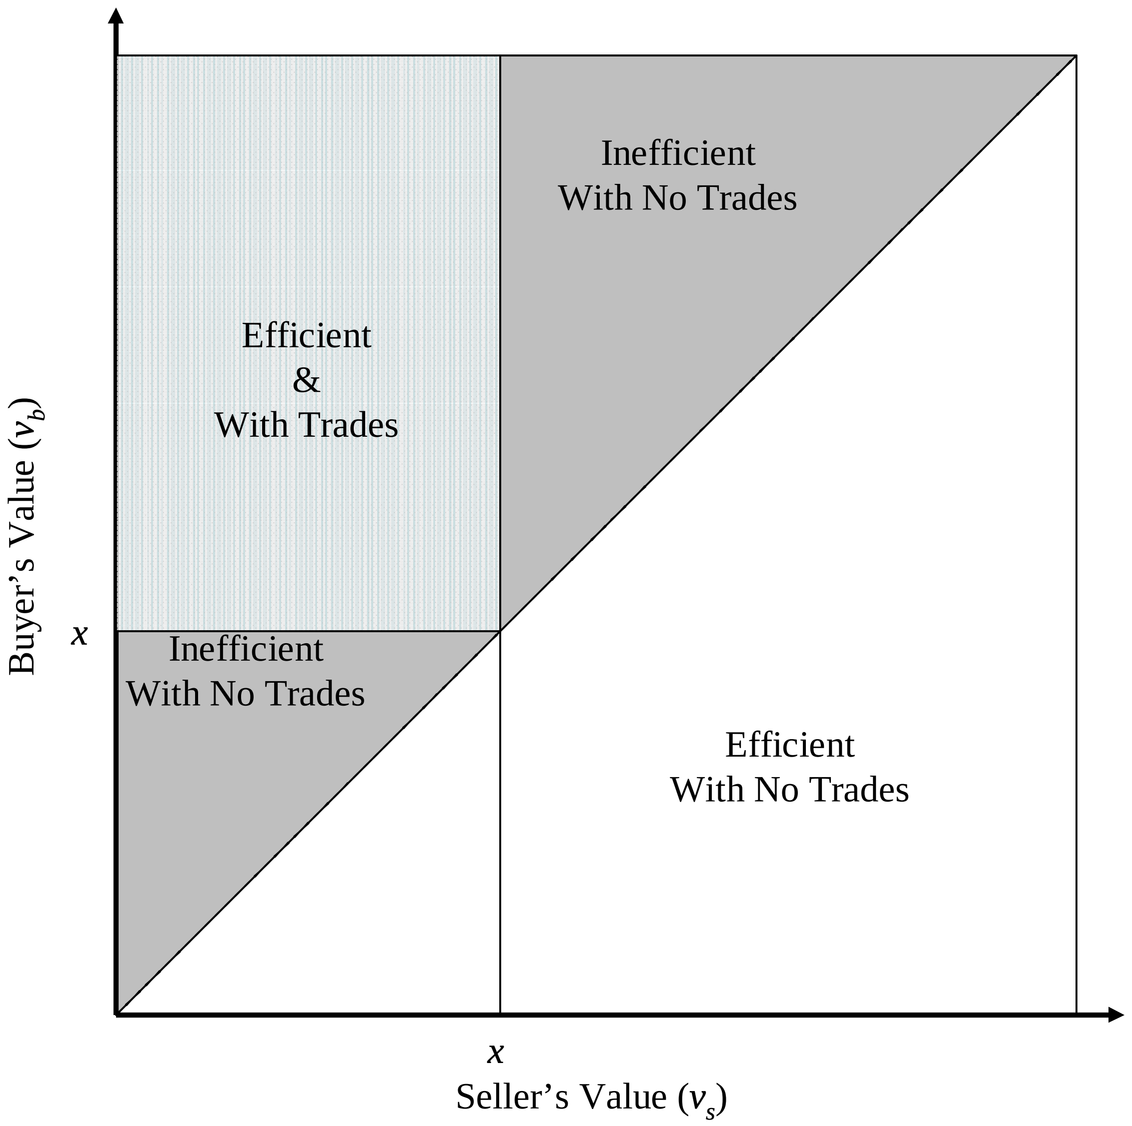

Figure 2: Inefficiencies of a linear price strategy in the double auction.