Econ 121b: Intermediate Microeconomics

Dirk Bergemann, Spring 2012

Week of 4/8 - 4/14

1 Lecture 16: Game Theory Continued

So far we have considered only pure strategies: strategies where the players do not randomize over which action they take. In other words, a pure strategy is a deterministic choice. The following simple example demonstrates that a pure strategy Nash Equilibrium may not always exist.

Example 1. (Matching Pennies) Consider the following payoff matrix:

		Bob H	Bob T
Ann	H	1, −1	−1, 1
	T	−1, 1	1, −1

Here Ann wins if both players play the same strategy, and Bob wins if they play different ones. Clearly there cannot be pure strategy equilibrium, since Bob would have an incentive to deviate whenever they play the same strategy and Ann would have an incentive to deviate if they play the differently. Intuitively, the only equilibrium is to randomize between H and T with probability $\frac{1}{2}$ each.

While the idea of a matching pennies game may seem contrived, it is merely the simplest example of a general class of zero-sum games, where the total payoff of the players is constant regardless of the outcome. Consequently gains for one player can only come from losses of the other. For this reason, zero-sum games will rarely have a pure strategy Nash equilibrium. Examples would be chess, or more relevantly, competition between two candidates or political parties. Cold War power politics between the US and USSR was famously (although probably not accurately) modelled as a zero-sum game. Most economic situations are not zero-sum since resources can be used inefficiently.

Example 2. A slight variation is the game of Rock-Paper-Scissors.

		Bob		
		R	P	S
	R	0, 0	−1, 1	1, −1
Ann	P	1, −1	0, 0	−1, 1
	S	−1, 1	1, −1	0, 0

Definition 1. A *mixed strategy* by player i is a probability distribution $\sigma_i = \left(\sigma_i\left(S_i^1\right), \ldots, \sigma_i\left(S_i^K\right)\right)$ such that

$$\sigma_i\left(s_i^k\right) \geq 0$$
$$\sum_{k=1}^{K} \sigma_i\left(s_i^k\right) = 1.$$

Here we refer to s_i as an action and to σ_i as a strategy, which in this case is a probability distribution over actions. The action space is $S_i = \left\{s_i^1, \ldots, s_i^K\right\}$.

Expected utility from playing action s_i when the other player plays strategy σ_j is

$$u_i\left(s_i, \sigma_j\right) = \sum_{k=1}^{K} \sigma_j\left(s_j^k\right) u_i\left(s_i, s_j^k\right).$$

Example 3. Consider a coordination game (also knows as "battle of the sexes" similar to the one in Example **??** but with different payoffs

			Bob	
			σ_B	$1 - \sigma_B$
			O	C
Ann	σ_A	O	1, 2	0, 0
	$1 - \sigma_A$	C	0, 0	2, 1

Hence Bob prefers to go to the opera (O) and prefers to go to a cricket match (C), but both players would rather go to an event together than alone. There are two pure strategy Nash Equilibria: (O, O) and (C, C). We cannot make a prediction, which equilibrium the players will pick. Moreover, it could be the case that there is a third Nash Equilibrium, in which the players randomize.

Suppose that Ann plays O with probability σ_A and C with probability $1 - \sigma_A$. Then Bob's expected payoff from playing O is

$$2\sigma_A + 0(1 - \sigma_A) \quad (1)$$

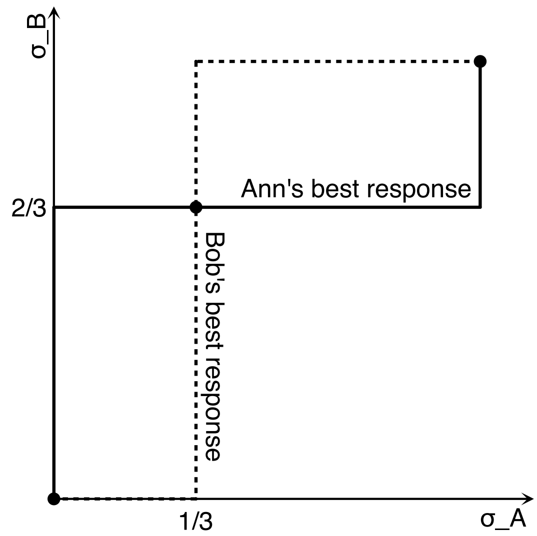

Figure 1: Three Nash Equilibria in battle of the sexes game

and his expected payoff from playing C is

$$0\sigma_A + 1(1 - \sigma_A). \tag{2}$$

Bob is only willing to randomize between his two pure strategies if he gets the same expected payoff from both. Otherwise he would play the pure strategy that yields the highest expected payoff for sure. Equating (1) and (2) we get that

$$\sigma_A^* = \frac{1}{3}.$$

In other words, Ann has to play O with probability $\frac{1}{3}$ to induce Bob to play a mixed strategy as well. We can calculate Bob's mixed strategy similarly to get

$$\sigma_B^* = \frac{2}{3}.$$

Graphically, we can depict Ann's and Bob's best response function in Figure 1. The three Nash Equilibria of this game are the three intersections of the best response functions.

2 Lecture 16: Asymmetric Information: Adverse Selection and Moral Hazard

Asymmetric information simply refers to situations where some of the players have relevant information that other players do not. We consider two types of asymmetric information: adverse selection, also known as hidden information, and moral hazard or hidden action.

A leading example for adverse selection occurs in life or health insurance. If an insurance company offers actuarially fair insurance it attracts insurees with above average risk whereas those with below average risk decline the insurance. (This assumes that individuals have private information about their risk.) In other words, individuals select themselves into insurance based on their private information. Since only the higher risks are in the risk pool the insurance company will make a loss. In consequence of this adverse selection the insurance market breaks down. Solutions to this problems include denying or mandating insurance and offering a menu of contracts to let insurees self-select thereby revealing their risk type.

Moral hazard is also present in insurance markets when insurees' actions depend on having insurance. For example, they might exercise less care when being covered by fire or automobile insurance. This undermines the goal of such insurance, which is to provide risk sharing in the case of property loss. With moral hazard, property loss becomes more likely because insurees do not install smoke detectors, for example. Possible solutions to this problem are copayments and punishment for negligence.

2.1 Adverse Selection

The following model goes back to George Akerlof's 1970 paper on "The market for lemons." The used car market is a good example for adverse selection because there is variation in product quality and this variation is observed by sellers, but not by buyers.

Suppose there is a potential buyer and a potential seller for a car. Suppose that the quality of the car is denoted by $\theta \in [0, 1]$. Buyers and sellers have different valuations/willingness to pay v_b and v_s, so that the value of the car is $v_b\theta$ to the buyer and $v_s\theta$ to the seller. Assume that $v_b > v_s$ so that the buyer always values the car more highly then the seller. So we know that trade is always efficient. Suppose that both the buyer and seller know θ, then we have seen in the bilateral trading section that trade can occur at any price $p \in [v_s\theta, v_b\theta]$ and at that price the efficient allocation (buyer gets the car) is realized (the buyer has a net payoff of $v_b\theta - p$ and the seller gets $p - v_s\theta$, and the total surplus is $v_b\theta - v_s\theta$).

The assumption that the buyer knows the quality of the car may be reasonable in some situations (new car), but in many situations the seller will be much better informed about the car's quality. The buyer of a used car can observe the age, mileage, etc. of a car and so have a rough idea as to quality, but the seller has presumably been driving the car and will know more about it. In such a situation we could consider the quality θ as a random variable, where the buyer knows only the distribution but the seller knows the realization. We could consider a situation where the buyer knows the car is of a high quality with some probability, and low quality otherwise, whereas the seller knows whether the car is high quality. Obviously the car could have a more complicated range of potential qualities. If the seller values a high quality car more, then their decision to participate in the market potentially reveals negative information about the quality, hence the term adverse selection. This is because if the car had higher quality the seller would be less willing to sell it at any given price. How does this type of asymmetric information change the outcome?

Suppose instead that the buyer only knows that $\theta \sim U[0,1]$. That is that the quality is uniformly distributed between 0 and 1. The the seller is willing to trade if

$$p - v_s\theta \geq 0 \tag{3}$$

and the buyer, who does not know θ, but forms its expected value, is willing to trade if

$$E[\theta]v_b - p \geq 0. \tag{4}$$

However, the buyer can infer the car's quality from the price the seller is asking. Using condition (3), the buyer knows that

$$\theta \leq \frac{p}{v_s}$$

so that condition (4) becomes

$$E\left[\theta \middle| \theta \leq \frac{p}{v_s}\right] v_b - p = \frac{p}{2v_s} v_b - p \geq 0, \tag{5}$$

where we use the conditional expectation of a uniform distribution:

$$E[\theta|\theta \leq a] = \frac{a}{2}.$$

Hence, simplifying condition (5), the buyer is only willing to trade if

$$v_b \geq 2v_s.$$

In other words, the buyer's valuation has to exceed twice the seller's valuation for a trade to take place. If

$$2v_s > v_b > v_s$$

trade is efficient, but does not take place if there is asymmetric information.

In order to reduce the amount of private information the seller can offer a warranty of have a third party certify the car's quality.

If we instead assumed that neither the buyer or the seller know the realization of θ then the high quality cars would not be taken out of the market (sellers cannot condition their actions on information they do not have) and so we could have trade. This indicates that it is not the incompleteness of information that causes the problems, but the asymmetry.

2.2 Moral Hazard

Moral hazard is similar to asymmetric information except that instead of considering hidden information, it deals with hidden action. The distinction between the two concepts can be seen in an insurance example. Those who have pre-existing conditions that make them more risky (that are unknown to the insurer) are more likely, all else being equal, to buy insurance. This is adverse selection. An individual who has purchased insurance may become less cautious since the costs of any damage are covered by insurance company. This is moral hazard. There is a large literature in economics on how to structure incentives to mitigate moral hazard. In the insurance example these incentives often take the form of deductibles and partial insurance, or the threat of higher premiums in response to accidents. Similarly an employer may structure a contract to include a bonus/commission rather then a fixed wage to induce an employee to work hard. Below we consider an example of moral hazard, and show that a high price may signal an ability to commit to providing a high quality product.

Suppose a cook can choose between producing a high quality meal ($q = 1$) and a low quality meal ($q = 0$). Assume that the cost of producing a high quality meal is strictly higher than a low quality meal ($c_1 > c_0 > 0$). For a meal of quality q, and price p the benefit to the customer is $q - p$ and to the cook is $p - c_i$. So the total social welfare is

$$q - p + p - c_i = q - c_i$$

and assume that $1 - c_1 > 0 > -c_0$ so that the high quality meal is socially efficient. We assume that the price is set beforehand, and the cook's choice variable is the quality of the meal. Assume that fraction α of the consumers are repeat clients who are informed about the meal's quality, whereas $1 - \alpha$ of the consumers are uniformed (visitors to the city perhaps) and don't know the meal's quality. The informed customers will only go to the restaurant if the meal is good (assume $p \in (0, 1)$). These informed customers allow us to consider a notion of reputation even though the model is static.

Now consider the decision of the cook as to what quality of meal to produce. If they produce a high quality meal then they sell to the entire market so their profits (per customer) are

$$p - c_1$$

Conversely, by producing the low quality meal, and selling to only $1-\alpha$ of the market they earn profit

$$(1-\alpha)(p-c_0)$$

and so the cook will provide the high quality meal if

$$p-c_1 \geq (1-\alpha)(p-c_0)$$

or

$$\alpha p \geq c_1-(1-\alpha)c_0$$

where the LHS is the additional revenue from producing a high quality instead of a low quality meal and the RHS is the associated cost. This corresponds to the case

$$\alpha \geq \frac{c_1-c_0}{p-c_0}.$$

So the cook will provide the high quality meal if the fraction of the informed consumers is high enough. So informed consumers provide a positive externality on the uninformed, since the informed consumers will monitor the quality of the meal, inducing the chef to make a good meal.

Finally notice that price signals quality here: the higher the price the smaller the fraction of informed consumers necessary ensure the high quality meal. If the price is low ($p \approx c_1$) then the cook knows he will lose $p-c_1$ from each informed consumer by producing a low quality meal instead, but gains c_1-c_0 from each uninformed consumer (since the cost is lower). So only if almost every consumer is informed will the cook have an incentive to produce the good meal. As p increases so does $p-c_1$, so the more is lost for each meal not sold to an informed consumer, and hence the lower the fraction of informed consumers necessary to ensure that the good meal will be provided. An uninformed consumer, who also may not know α, could then consider a high price a signal of high quality since it is more likely that the fraction of informed consumers is high enough to support the good meal the higher the price.

2.3 Second Degree Price Discrimination

In Section **??** we considered first and third degree price discrimination where the seller can identify the type of potential buyers. In contrast, second degree price discrimination occurs when the firm cannot observe to consumer's willingness to pay directly. Consequently they elicit these preferences by offering different quantities or qualities at different prices. The consumer's type is revealed through which option they choose. This is known as screening.

Suppose there are two types of consumers. One with high valuation of the good θ_h, and one with low valuation θ_l. θ is also called the buyers' marginal willingness to pay. It tells us how much a buyer what be willing to pay for an additional unit of the good. Each buyer's type is his private information. That means the seller does not know ex ante what type a buyer he is facing is. Let α denote the fraction of consumers who have the high valuation. Suppose that the firm can produce a product of quality q at cost $c(q)$ and assume that $c'(q) > 0$ and $c''(q) > 0$.

First, we consider the efficient or first best solution, i.e., the case where the firm can observe the buyers' types. If the firm knew the type of each consumer they could offer a different quality to each consumer. The condition for a consumer of type $i = h, l$ buying an object of quality q for price p voluntarily is

$$\theta_i q - p(q) \geq 0$$

and for the firm to participate in the trade we need

$$p(q) - c(q) \geq 0.$$

Hence maximizing joint payoff is equivalent to

$$\max_q \theta_i q - p(q) + p(q) - c(q)$$

or

$$\max_q \theta_i q - c(q).$$

The FOC for each quality level is

$$\theta_i - c'(q) = 0,$$

from which we can calculate the optimal level of quality for each type, $q^*(\theta_i)$. Since marginal cost is increasing by assumption we get that

$$q^*(\theta_l) < q^*(\theta_h),$$

i.e., the firm offers a higher quality to buyers who have a higher willingness to pay in the first best case. In the case of complete information we are back to first degree price discrimination and the firm sets the following prices to extract the entire gross utility from both types of buyers:

$$p_h^* = \theta_h q^*(\theta_h) \quad \text{and} \quad p_h^* = \theta_l q^*(\theta_l)$$

so that buyers' net utility is zero. In Figure 2, the buyers' gross utility, which is equal to the price charged, is indicated by the rectangles $\theta_i q_i^*$.

In many situations, the firm will not be able to observe the valuation/willingness to pay of the consumers. That is, the buyers' type is their private information. In

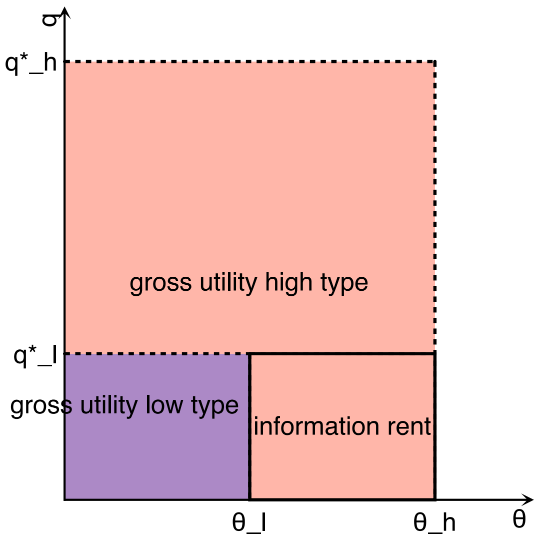

Figure 2: Price discrimination when types are known to the firm

such a situation the firm offers a schedule of price-quality pairs and lets the consumers self-select into contracts. Thereby, the consumers reveal their type. Since there are two types of consumers the firm will offer two different quality levels, one for the high valuation consumers and one for the low valuation consumers. Hence there will be a choice of two contracts (p_h, q_h) and (p_l, q_l) (also called a menu of choices). The firm wants high valuation consumers to buy the first contract and low valuation consumers to buy the second contract. Does buyers' private information matter, i.e., do buyers just buy the first best contract intended for them? High type buyers get zero net utility from buying the high quality contract, but positive net utility of $\theta_h q^*(\theta_l) - p_l > 0$. Hence, high type consumers have an incentive to pose as low quality consumers and buy the contract intended for the low type. This is indicated in Figure 2 as "information rent," i.e., an increase in high type buyers' net utility due to asymmetric information.

The firm, not knowing the consumers' type, however, can make the low quality bundle less attractive to high type buyers by decreasing q_l or make the high quality contract more attractive by increasing q_h or decreasing p_h. The firm's profit maximization problem now becomes

$$\max_{p_h, p_l, q_h, q_l} \alpha \left(p_h - c(q_h)\right) + (1 - \alpha)\left(p_l - c(q_l)\right). \tag{6}$$

There are two type of constraints. The consumers have the option of walking away, so the firm cannot demand payment higher than the value of the object. That is, we must have

$$\theta_h q_h - p_h \geq 0 \quad (7)$$
$$\theta_l q_l - p_l \geq 0. \quad (8)$$

These are known as the individual rational (IR) or participation constraints that guarantee that the consumers are willing to participate in the trade. The other type of constraints are the self-selection or incentive compatibility (IC) constraints

$$\theta_h q_h - p_h \geq \theta_h q_l - p_l \quad (9)$$
$$\theta_l q_l - p_l \geq \theta_l q_h - p_h, \quad (10)$$

which state that each consumer type prefers the menu choice intended for him to the other contract. Not all of these four constraints can be binding, because that would determine the optimal solution of prices and quality levels. The IC for low type (10) will not binding because low types have no incentive to pretend to be high types: they would pay a high price for quality they do not value highly. On the other hand high type consumers' IR (7) will not be binding either because we argued above that the firm has to incentivize them to pick the high quality contract. This leaves constraints (8) and (9) as binding and we can solve for the optimal prices

$$p_l = \theta_l q_l$$

using constraint (8) and

$$p_h = \theta_h(q_h - q_l) + \theta_l q_l$$

using constraints (8) and (9). Substituting the prices into the profit function (6) yields

$$\max_{q_h, q_l} \alpha \left[\theta_h(q_h - q_l) + \theta_l q_l - c(q_h)\right] + (1-\alpha)\left(\theta_l q_l - c(q_l)\right).$$

The FOC for q_h is simply

$$\alpha\left(\theta_h - c'(q_h)\right) = 0,$$

which is identical to the FOC in the first best case. Hence, the firm offers the high type buyers their first best quality level $q_R^*(\theta_h) = q^*(\theta_h)$. The FOC for q_l is

$$\alpha(\theta_l - \theta_h) + (1-\alpha)\left(\theta_l - c'(q_l)\right) = 0,$$

which can be rewritten as

$$\theta_l - c'(q_l) - \frac{\alpha}{1-\alpha}(\theta_l - \theta_h) = 0.$$

The third term on the LHS, which is positive, is an additional cost that arises because the firm has to make the low quality contract less attractive for high type buyers. Because of this additional cost we get that $q_R^*(\theta_l) < q^*(\theta_l)$: the the quality level for low types is lower than in the first best situation. This is depicted in Figure . The low type consumers' gross utility and the high type buyers' information rent are decreased, but The optimal level of quality offered to low type buyers is decreasing in the fraction of high type consumer α:

$$\frac{dq_R^*(\theta_l)}{d\alpha} < 0$$

since the more high types there are the more the firm has to make the low quality contract unattractive to them.

This analysis indicates some important results about second degree price discrimination:

1. The low type receives no surplus.
2. The high type receives a positive surplus of $q_l(\theta_h - \theta_l)$. This is known as an information rent, that the consumer can extract because the seller does not know his type.
3. The firm should set the efficient quality for the high valuation type.
4. The firm will degrade the quality for the low type in order to lower the rents the high type consumers can extract.