MBA201A: Game Theory Handout

Ed Egan

October 2010

Introduction

Games are characterized by three things: **Players**, **strategies** (or actions[1]. **Always label players, actions, and payoffs**), and **payoffs**. An **equilibrium** to a game is set of strategies for each player that meet a certain criteria, not a final payoff - the final payoff is the equilibrium outcome. The 'certain criteria' will depend on the equilibrium concept. You will be working with Nash Equilibria, which is defined below.

There are three types of games that you might encounter:

1. **Strategic Form Games -** These are simultaneous move games that have a payoff grid and are solved by dominance or finding best responses.
2. **Extensive Form Games -** These are sequential move games that have a tree, much like a decision tree but with more than one player and no uncertainty. These are solved by backwards induction.
3. **Formless Games -** These are essentially logic puzzles, where you would use the solution techniques to find an answer but you could not draw a game tree or payoff matrix.

In addition we often talk about one-shot and repeated games. To start with we will discuss just one-shot games. There is a repeated game section later. In repeated games we talk about the 'stage game' which is the one-shot game played at each stage.

In **strategic form** games the payoff matrix looks like the figure below. Note the order of the players, the order of the payoffs, and the labelling of the actions.

In **extensive form** games we draw trees. Always label the player who is making the decision at every decision node, as well as their actions. The payoffs should be written on the terminal node in the order of play (Payoff to first mover, Payoff to second mover, etc). Trees can be drawn in any direction you like, but

[1] We will use the terms actions and strategies interchangeably

		Player 2	
		*Action*1	*Action*2
Player 1	*Action*1	Payoff to 1,Payoff to 2	$\cdots,\cdots$
	*Action*2	$\cdots,\cdots$	$\cdots,\cdots$

Figure 1: A Strategic Form Game

generally it is easy to draw them either from left to right or from top to bottom.

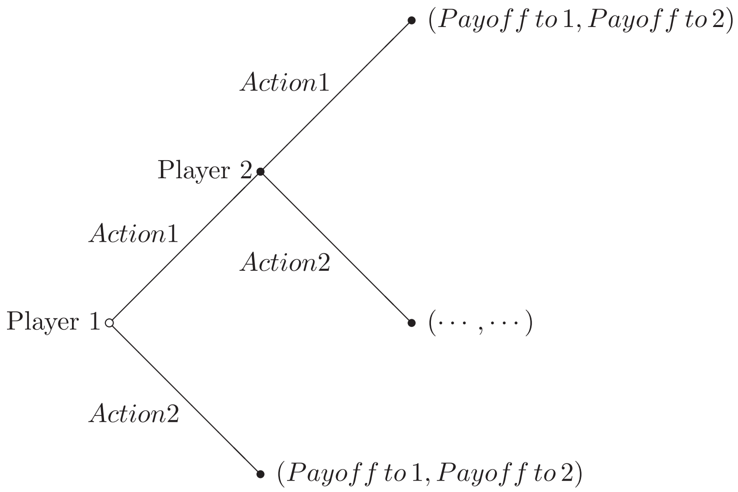

Figure 2: An Extensive Form Game

Dominance

The strategic form (simultaneous) game below gives an example called "the prisoner's dilemma". Ed and Steve have both been arrested and placed in seperate cells and both have been offered the same deal: If they rat on the other and the other stays silent they go free but their partner goes to jail for 10 years; if they both stay silent they are both charged with a minor offense and both go to jail for a single year; and if they both rat then they both go to jail for 5 years. Ed and Steve are both professional economists. They are complete rational[2]. What should they do?

This game can be solved by **interative elimination of dominated strategies**. First, some definitions:

[2] You should take this to mean that they both selfishly maximize their own utilities without any regard for anyone else.

		Steve	
		Silent	*Rat*
Ed	*Silent*	$-1, -1$	$-10, 0$
	Rat	$0, -10$	$-5, -5$

Figure 3: The Prisoner's Dilemma

- A strategy is **dominant** if irrespective of the actions of the other player(s), it always gives a strictly higher utility.
- A strategy is **dominated** if it is in the same action set (at the same time) as a dominant strategy.
- In this context **iterative** means to keep applying a rule until it no longer changes the available strategy set for any player.
- Likewise, in this context **elimination** means to remove a strategy or strategies from a set of possible strategies a player may take. This effectively reduces the game.

Looking at the prisoner's dilemma it should be clear that if Ed stays *Silent* he gets -1 or -10 if Steve stays *Silent* or *Rat*s respectively, and if Ed *Rat*s then he gets 0 or -5 if Steve stays *Silent* or *Rat*s respectively. *Rat* is therefore a **dominant strategy** for Ed. If Steve stays *Silent* then Ed gets 0 from *Rat* versus -1 from *Silent*, and if Steve *Rat*s then Ed gets -5 from *Rat* and -10 from *Silent*. Irrespective of what Steve does, Ed prefers to *Rat*.

Therefore *Silent* is a **dominated strategy and can be eliminated** for Ed. If we now turn to Steve and iterate this process we find that Ed has only one stragegy remaining, *Rat*, and Steve faces a payoff of -5 from *Rat* and -10 from *Silent* and so *Rat* dominates *Silent* and *Silent* can be eliminated for Steve. There is only one cell left, so the equilibrium must be (Rat, Rat).

Iterative elimination of dominated strategies will not always solve the entire game, leaving only a single cell and so a single pair of strategies, but it is a very fast method of reducing a game down to something smaller. Dominated strategies are never equilibria, and so can always be removed.

Note, however, that to remove a strategy we must have **strict dominance**. Weak dominance, where the pay off to two or more actions is the same for some action of another player, will not suffice. In the game below you can not remove any payoffs using iterative elimination of dominated strategies.

		Player 2	
		Left	*Right*
Player 1	*Up*	$0, 0$	$2, 1$
	Down	$1, 2$	$2, 2$

Figure 4: You can not eliminate weakly dominated strategies

Nash Equilibria

A best response is the best action available - the action that maximizes utility. We talk about best responses given beliefs, but you can also think of best responses as being the best action to respond to some opponents action with.

In class you were given the following definition. A profile of strategies is a **Nash Equilibrium** if:

1. Every player is playing a best response to her beliefs
2. The beliefs of all the players are correct

Let's apply this to the prisoner's dilemma. Suppose that Ed believed that Steve would stay *Silent*. Then Ed's best reponse is to *Rat*. Likewise if Ed believed that Steve would *Rat* then Ed's best response is to *Rat*. The same logic can be said of Steve. In the case of (Rat, Rat), Ed believed that Steve would *Rat* and so chose to *Rat*, and Steve believed Ed would *Rat* and so chose *Rat*. Both players were playing their best responses given their beliefs and their beliefs were correct!

The game below is the "Battle of the Sexes"[3]. Sam wants to go to the *Ballet*, and Alex wants to go to the *Opera*, and both get more enjoyment if they go with the other.

		Alex	
		Ballet	*Opera*
Sam	*Ballet*	$2, 2$	$1, 0$
	Opera	$0, 1$	$2, 2$

		Alex	
		Ballet	*Opera*
Sam	*Ballet*	$\underline{2}, \underline{2}$	$1, 0$
	Opera	$0, 1$	$\underline{2}, \underline{2}$

Figure 5: Battle of the Sexes

To find the Nash equilibria (there are two), you should hold one player's action constant and underline the other player's best response (actually you underline the payoff the best response gets). In the redrawing of the "Battle of the Sexes" this has been done for both players, first imagining that the opponent played *Ballet* and then that the opponent played *Opera*. Any cell that has both of its payoffs underlined is a Nash equilibrium. Remember, once again, that we write equilibriums in terms of strategies, so the set of equilibria in the "Battle of the Sexes" is $\{(Ballet, Ballet), (Opera, Opera)\}$.

The best responses for Ed and for both Steve and Ed for the Prisoner's dilemma are also shown below.

You do not need to solve for 'mixed strategy' equilibria - these are where you specify that that some pure strategy (e.g. *Ballet*) be played with some probability p_1 and some other pure strategy (e.g. *Opera*) be played with some other probability p_2, and so forth. However, you do need to be aware of the concept.

[3]This is the Berkeley version, courtesy of John Bates Clark Medal winner Matthew Rabin.

		Steve	
		Silent	*Rat*
Ed	*Silent*	$-1, -1$	$-10, 0$
	Rat	$\underline{0}, -10$	$\underline{-5}, -5$

		Steve	
		Silent	*Rat*
Ed	*Silent*	$-1, -1$	$-10, \underline{0}$
	Rat	$\underline{0}, -10$	$\underline{-5}, \underline{-5}$

Figure 6: Ed's BRs and both Ed & Steve's BRs, for the Prisoner's Dilemma

Crucially you should know that there are **always an odd number of equilibria**. That is every game must have at least one equilibria, and if you have 2 pure strategy equilibria then there will be at least a third mixed strategy equilibria. In the "Batttle of the Sexes" the third, mixed strategy, equilibrium is to mixed the two pure strategy equilibria 50-50.

Sequential Games

Sequential games are even easier to solve than simultaneous games: Just use dominance and backwards induction. That is work from the last decision in the game, and choose the best response (one strategy will always dominate another for you), and then mark that as the action taken and move the payoff backwards, until the entire game is marked.

A strategy in a sequential game is a complete contingent plan. That is, not matter what the other player does, you must have a full account of what you will do. Imagine that you have to hand in your strategy to an agent who will play the game for you - the agent must always know what you want her to do, without consulting anything other than your strategy.

In the simple game below called "Enter and Accomodate", you can refer to the equilibrium as (*Enter*, *Accomodate*). However, formally, the equilibrium strategy plan is *Enter* for player 1, and "if *Enter*, *Accomodate*" and "if *Don't Enter*, *Accomodate*".

First you solve the last stage, and then you move the payoff back to the first stage to solve the game:

Auctions

You should be familar with four auction formats:

1. **English Auction** - The prices are called out, and ascend until only one bidder remains. The dominant strategy is to keep bidding as long as the price is less than your valuation.

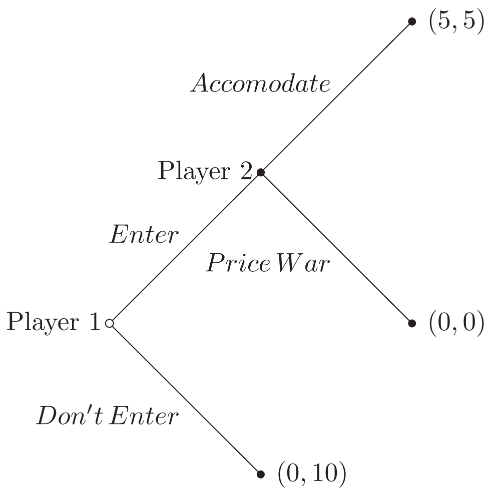

Figure 7: Enter and Accomodate

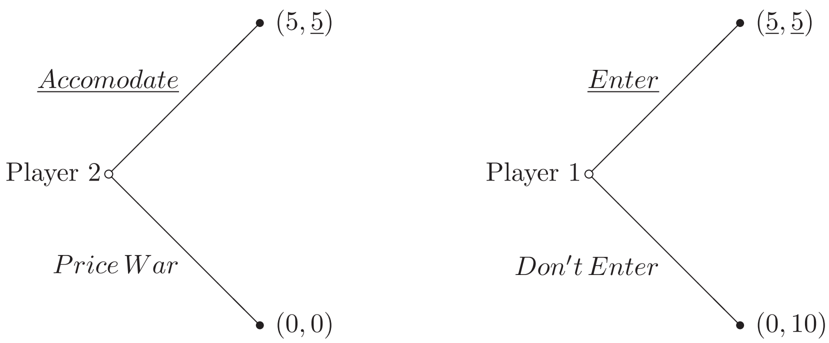

Figure 8: Enter and Accomodate, by backwards induction

2. **Second Price Sealed Bid Auction -** Bids are submitted in sealed envelopes. The highest bid wins and pays the second highest price. The dominant strategy is to bid your valuation.
3. **First Price Sealed Bid Auction -** Bids are submitted in sealed envelopes. The highest bid wins and the winner pays their bid. There is no dominant strategy. You should bid less than your valuation, but how much you should 'shave' your valuation depends on a trade-off. The more you shave the more benefit you get if you win, but the lower your chances of winning.
4. **Dutch Auction -** The price starts high and the drops. The first person to raise thier hand (i.e. bid) wins and pays the current price. Again there is no dominant strategy. You should bid less than your valuation. How much less depends on the same trade off as for the first price sealed bid auction.

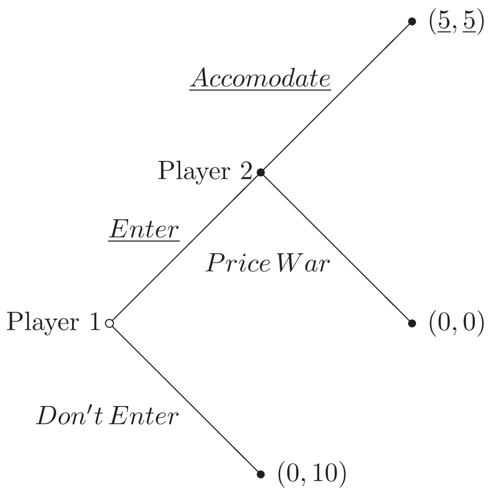

Figure 9: Enter and Accomodate, with the equilibrium

The English auction and the second price sealed bid auction are strategically equivalent, as are the Dutch auction and the first price sealed bid auction. Strategically equivalent auctions result in the same 'outcome'. That is the same person wins and pays the same amount. An auction is considered "perfectly discrimating" if the person who values the good the highest, wins the auction. Thus the English auction and the second price sealed bid auction are perfectly discriminating. Because of the shaving trade-off, both the Dutch auction and the first price sealed bid auction are imperfectly discriminating.

If your valuation is **not known with certainty**, as is the case for most firms valuing a potential future income stream (like an oil well, or spectrum rights), then there is **information in the other participants bids**. Suppose that every bidder receives an unbiased signal of the true value randomly drawn from some distribution. The best unbiased estimator of the true value is the mean (average) of those signals. Further suppose that everyone bids their valuation (or shaves their valuation in the same way), then the person that wins had a valuation above the mean (by definition). This is called the **Winner's Curse**. In a nut shell, the winner of an uncertain value prize has overpaid. In such as case of uncertainty in the value, information about other people's valuations, including their bids if the bids are not sealed, is clearly of great use.

Bertrand Trap

The Bertrand Trap is a famous game that predates game theory. Two (or more) firms produce an identical homogenous good that will be purchased by a large number of consumers who can observe both (or all) prices. There are no frictions in the market - no travel costs, no brand preferences, or anything else.

The market 'rules' are that who ever prices the lowest gets all of the customers, and if firms price the same they split the market evenly.

What happens? Firms competing on price must price at marginal cost. It is the unique Nash Equilibrium. Suppose a firm priced above marginal cost, then the best response of a competitor would be to undercut them. Pricing below marginal cost is clearly a dominated strategy. Pricing at marginal cost is a best response given that you believe other firms are pricing at marginal cost.

There are six conditions that must hold for the Bertrand Trap to be unavoidable, so to avoid the trap you have to break at least one of them:

1. No cost advantage
2. Unlimited capacity
3. No differentiation
4. No switching costs
5. Prices are known to all customers
6. No collusion

Explicit collusion is illegal (at least in the United States, but Cartels like OPEC do exist and prosper elsewhere). Tacit collusion is not illegal, and is facilitated by trade associations, price advertising, and price signalling. The theory of repeated games gives an explanation of how collusion can be sustained.

Repeated Games

There are essentially two ways that a game can be repeated: finitely or infinitely. If at the end of every stage there is some non-zero probability that the game will be played again, then this is equivalent to an infinitely repeated game.

Finitely repeated games do not add much for us. We solve them by backwards induction, begining with the last stage. In finitely repeated games there are no new 'cooperative' equilibrium.

In the "Bertrand Trap Like Game" there is a crude attempt at producing a strategic form version of the Bertrand Trap. There are two firms, who would like to collude to share monopoly profits of 2 equally for a payoff of $(1, 1)$. If one defects but the other maintains the collusion price then the defector gets the entire market for a payoff of $(2, 0)$ or $(0, 2)$ (depending on who defected). And if they both defect then the Bertrand Trap gets them and they are forced to price at marginal cost and earn no profits for a payoff of $(0, 0)$.

The stage game has three pure strategy Nash equilibrium:

$$\{(Collude, Defect), (Defect, Collude), (Defect, Defect)\}$$

		Firm2	
		Collude	*Defect*
Firm1	*Collude*	$1, 1$	$0, 2$
	Defect	$2, 0$	$0, 0$

Figure 10: Bertrand Trap Like Game

Only the $(Collude, Collude)$ pair of strategies is not an equilibrium. Can we make this pair a new 'cooperative' equilibrium?

In a finitely repeated game the answer is no. In the last period there are no rewards to collusion and no punishment from defecting. Given this in the second last period there is no incentive to collude beyond that in the stage game. There is no new reward awaiting good behaviour and no new punishment awaiting bad behaviour: the game is the same. This logic follows forward until the first period.

However, in infinitely repeated games there is a way to get the new 'cooperative' equilibrium, at least if people value the future highly enough. There are two 'punishment' strategies that you should know:

- **Tit-For-Tat**: If the opponent colludes then collude. If an opponent defects on you, then for one period you punish the opponent before returning to collusion.
- **Grim Trigger**: If the opponent colludes then collude. If an opponent defects on you,then you punish them for all time and never, ever return to colluding.

It should be obvious that the Grim Trigger is a much harsher punishment than Tit-For-Tat, in fact it is the most extreme punishment that can be given in infinitely repeated games. We use the Grim Trigger as an example as if cooperation can't be sustained with the Grim Trigger then cooperation can't be sustained!

When we consider whether collusion is preferred to defection, we are really considering whether the NPV of the infinite stream of payoffs from collusion are greater than the NPV of the infinite stream of payoffs from defection given than players are playing the Grim Trigger. We use a discount rate of delta where $\delta \in [0, 1]$ each period:

The solution[4] says that for any discount rate greater than $\frac{1}{2}$ we can sustain collusion. If a softer punishment, such as Tit-For-Tat were used, we would expect that this discount rate would have to be substantially higher (i.e. people would have to value the future more so the punishment is felt with more force), and for some games may even exceed 1 (making cooperation impossible under normal conditions).

[4]The solution uses the equation for the sum of an infinite geometric series, which is: $\sum_{t=0}^{\infty} \alpha\beta^t = \frac{\alpha}{1-\beta}$.

$$\pi(Collude) > \pi(Defect)$$

$$\underbrace{1 + (\delta \cdot 1) + (\delta^2 \cdot 1) + (\delta^3 \cdot 1) + \ldots}_{\text{Cooperate}} > \underbrace{2}_{Defect} + \underbrace{(\delta \cdot 0) + (\delta^2 \cdot 0) + (\delta^3 \cdot 0) + \ldots}_{\text{Get Punished}}$$

$$\sum_{t=0}^{\infty} \delta^t \cdot 1 > 2$$

$$\frac{1}{1-\delta} > 2$$

$$\therefore \delta > \frac{1}{2}$$

One frequent question is: Is it an equilibrium to play the Grim Trigger, wouldn't you want to defect? The answer is simple, given that I believe that the other person is playing the Grim Trigger then if the discount rate is high enough I will cooperate. If my opponent believe that I am playing the Grim Trigger and sees me cooperating then he too will cooperate. Thus we are playing best responses to best responses. The actual exercising of the Grim Trigger is 'off-the-equilibrium-path' and doesn't happen.