Formal Analysis
Game Theory: A Primer on Simple Games

Ben Ansell

Hilary Term 2017

Contents

1 Introduction

This lecture moves away from the ‘choice theoretic’ approach associated with microeconomics to a ‘strategic’ approach associated with game theory. In our analysis of public goods and externalities we were already noting interdependencies between agents’ choices and in principal agent problems we were directly dealing with situations where one agent must choose their behavior predicated upon the expected reaction of another. Thus, in many ways game theory as an approach is simply a new way of conceptualizing these other situations, and indeed public goods and principal-agent problems can be written out as games. So.. in this lecture what we are doing is moving to a consistent language for describing a whole series of multi-agent interactions, rather than doing anything completely new. The key insights will be in understanding the two ways in which games can be described: extensive form and normal form, and in understanding new equilibrium concepts: Nash equilibrium, dominated strategies, etc. Accordingly, I begin by going through the drudgery of describing the 'language' of game theory before moving on to examining equilibrium concepts and the solutions to games themselves. Throughout I rely on several sources: McCarty and Meirowitz (2007); Gintis (2009); Bowles (2004); and Ordeshook (1986).

2 The Language of Game Theory

There are two canonical ways of formally defining 'games' among actors: the *normal* and *extensive* forms. The two types of description can theoretically be applied to any game. However, as will become apparent, the normal form is best employed in the case of one-shot simultaneous games, and the extensive form in the case of dynamic games where actors move in a specified order. We will come back to extensive forms again in a later lecture.

2.1 Normal Form Games

A normal form game can be summarized as a collection of sets that defines all the key information embedded in a game. We will write it, following McCarty and Meirowitz as $\langle N, \{S_i, u(.)\}_{i \in n} \rangle$. What on earth does this mean? This defines a game as having N actors, each with their own individual set of strategies they can play S_i, which in combination with the strategies played by other players will produce payoffs that give utility $u_i(.)$. OK, now let's go through each of these in turn.

The first thing we need to know is the number of actors N. In most examples we will simply examine the strategic behavior of two actors, as in, for example, the classic Prisoner's Dilemma or Battle of the Sexes. Many games can be 'easily' extended to the N actor case: for example, the public goods problem resembles an n-actor Prisoner's Dilemma.

The second element is more complex: the set of strategies open to each actor indexed by i: S_i, where an actual choice of strategy is written in the lower case as s_i. What are strategies? These are the choices of actions available to each player. Note the importance of two words here: choices and actions. An action is simply a the behavior an agent engages in during a particular round of play. The choice aspect suggests that agents can condition their current and future actions on the behavior of other players. For example, they could choose to act in one way in the first round of interaction but a different way in future rounds if other actors do not reciprocate (this is the famous 'tit for tat' strategy). Thus a strategy is a set of actions over the set of rounds of interaction, possibly made conditional on the history of previous interactions. In a one round game, like the ones we will mostly discuss in this lecture, a strategy is the same as an action. Can you see why?

The final element of the normal form is the payoff structure of the game, as represented by the utilities $u(.)$ agents get from different strategies, *conditional* on the strategies played by other players. I noted above that utilities depend not only on an actor's own strategy but that of the other players. So if there are just two actors then we write $u_1(s_1, s_2)$ and $u_2(s_1, s_2)$ for the utilities of each player. More generally for n actors we use the expression $\sim i$ to mean 'all actors not including i', so that we can write utility payoffs for actor i as $u_i(s_i, s_{\sim i})$. This may seem rather complex but you should dwell on this point since this will be key for thinking about the concept of Nash Equilibrium.

As an example of a simple normal form we can write out the Prisoner's Dilemma as $N = 2; S_i = S_j = \{\text{Cooperate, Defect}\}; u_i\{s_i = C, s_j = C\} = 3; u_i\{s_i = C, s_j = D\} = 1; u_i\{s_i = D, s_j = C\} = 4; u_i\{s_i = D, s_j = D\} = 2$, where u_i is symmetric for u_j (i.e. both players face the same payoff structure). As with many normal games, we can write this out as a matrix, which is the classic way of representing simple normal form games.

		Player Two	
		C	D
Player One	C	$3, 3$	$1, 4$
	D	$4, 1$	$2, 2$

Prisoners' Dilemma

2.2 Extensive Form Games

Extensive form games are the classic 'game tree' examples you will see in the assigned readings. Since creating these in LaTeXis an absolute bind I will simply describe these and then put examples up on the board in class. The advantage of extensive form games is that they make much more apparent than do normal form games the dynamic structure of play - which actors get to move when, what their choices are, and what information they have.

You can think of the extensive form structure in the language that Gintis uses in the excerpts you have from this book. An extensive form game is mathematically a 'graph', which means a series of nodes and connecting paths (something you also see, for example, in social network theory). Each node defines a player who gets to make a choice at that node among a number of paths that connect to later nodes, further 'down' the game tree. At so-called 'terminal nodes' the players stop playing the game and receive utility payoffs that are perfectly determined by the series of paths and nodes that have been chosen by all the players through the game. Strategies, then, can be considered to be choices of actions (paths) at each node. Note, that even though in the resulting equilibrium some paths may never be taken and nodes unreached, a full strategy includes the choices that players would make at these unreached nodes. This restriction on behavior 'off the equilibrium path' is related closely to the concept of subgame perfection and credible threats - an issue we will leave for the next game theory lecture.

So... to connect back to the normal form - the number of players is defined by the number of players who have at least one node at which they choose among actions; strategies are defined as a list of actions taken at each possible node an actor can reach; and utility payoffs are associated with the terminal nodes of the game tree.

One other issue is worth noting with the extensive form: the concept of an 'information set'. These are nodes at which an actor does not know which one they are actually at. This occurs, for example, when the previous move by a different player is unobserved - therefore two or more different paths could have been taken meaning that the next actor to play is unaware of which path was taken and hence which node they are now at. They still face the same range of actions they can play themselves but these might produce different terminal payoffs depending on which path the previous actor chose. This is a form of information asymmetry which is easy to follow with extensive form games. With normal form games, as you'll see with McCarty and Meirowitz, the notation is rather complex - essentially, actors are experiencing certain games with different probabilities and are choosing over lotteries.

3 Equilibrium Concepts

How do actors choose among strategies? We begin with simplest cases of simultaneous games with symmetric and complete information. In these cases the concept of the **Nash Equilibrium** is used to describe optimal behavior. The Nash Equilibrium can be in either 'pure' strategies, where actors choose only one strategy, and 'mixed' strategies where actors choose among different strategies probabilistically.

3.1 Strict Domination

The basic way we get to Nash Equilibrium is through the concept of 'strict domination of pure strategies'. What does this mean? Each actor in a simple one period game has a range of strategies - and since they are in a one period game, these can also be called actions - that they can choose from. For each strategy/action, we can examine what the actor's utility payoff will be for each strategy/action chosen by the other player(s). As an example, if an actor chooses Cooperate in Prisoner's Dilemma (PD), they receive 3 if the other actor chooses Cooperate and 1 if the other actor chooses Defect. If, instead, the first actor chooses defect, they receive 4 if the other actor chooses Cooperate, and 2 if the other actor chooses Defect. Now notice something about this choice - regardless of what the other actor chooses, the first actor is *always* better off choosing Defect, since they earn more doing that than Cooperating in both cases (4 versus 3, and 2 versus 1). In this case the strategy of Cooperate is *strictly dominated* by Defect because *utility payoffs are always higher regardless of the strategy choice of the other player*.

Thus if we were to choose strategies for player one we would eliminate Cooperate. But... since this game is symmetric that must also be true for player two. Thus we end up with the only remaining outcome being {Defect, Defect}. This is, in fact, the Nash Equilibrium of the game. But we will return to that shortly. Formally we can write strict dominance as follows: s_i is strictly dominated by s_i' if $u(s_i, s_{\sim i}) < u(s_i', s_{\sim i}) \forall s_{\sim i} \in S_{\sim i}$. In words, a strategy is strictly dominated if for all possible strategy choices by other players, the utility payoff is lower than for another strategy the player could have chosen. Check above to make sure you see why Cooperate is dominated by Defect in the PD game.

Typically we actually have to iterate this procedure player by player since sometimes strategies for player one are not strictly dominated until player two eliminates their own strictly dominated strategies. Thus, the process of *iterated elimination of strictly dominated strategies* is typically necessary. This requires players to think through the process of elimination that the other players must be undertaking - hence all players have to assume one another's rationality *and* that the other player knows that the first player will act similarly. The process of *weak domination* can also be used, where we eliminate any strategy that is no better than another strategy regardless of the other players' strategic choices and is strictly worse for *at least one* strategy played by the other players. As above this process can also be iterated.

3.2 Nash Equilibrium

The concept of a Nash Equilibrium bears much resemblance to the discussion above (formally, we shall see they are identical). It involves the important concept of a *best response* function. This is the choice of a strategy s_i^* that is the best response possible to the strategies played by all other players $s_{\sim i}^*$, which are themselves best responses to all other players (and so on). The best response function (actually mathematically it is a 'correspondence') is defined as $b_i(s_{\sim i}) = \{s_i \in S_i : u_i(s_i, s_{\sim i}) \geq u_i(s_i', s_{\sim i}) \forall s_i' \in S_i$ and $\forall s_{\sim i} \in S_{\sim i}\}$. In English, the best response to other strategies is defined as the individual strategy that gives a higher utility than any other individual strategy for all possible strategy combinations of the other players.

From here it is easy to get to the two (mathematically identical) definitions of a Nash Equilibrium. A Nash Equilibrium is a strategy profile for all actors s^* satisfying $s_i^* \in b_i(s_{\sim i}^*) \forall i \in N$. Written differently it is a s^* such that $u_i(s_i^*, s_{\sim i}^*) \geq u_i(s_i', s_{\sim i}^*) \forall s_i' \in S_i$ and $\forall i \in N$. Again, in English these two equations can be written as (a) a Nash Equilibrium is a set of strategies for all players such that each strategy choice is in the set of best responses for each player; and (b) a Nash Equilbrium is a set of strategies for all players such that the utility for the strategy choice of each player, given the optimal strategies of all other players, is higher than the utility for any other strategy.

It turns out that this definition of a NE is formally equivalent to the statement that a strategy profile s^* is one where none of its strategies can be eliminated by iterated elimination of strictly dominated strategies. Thus, a simple way to solve for NEs is to apply iterated strict dominance. Hence {Defect, Defect} is a Nash Equilibrium.

		Violetta	
		Gambling	Opera
Alfredo	Gambling	$2, 1$	$0, 0$
	Opera	$0, 0$	$1, 2$

Battle of the Sexes

3.3 Mixed Strategies

In fact, it turns out there are more NEs than those discussed above since we were dealing with only *pure strategies*. If we allow players to choose probabilistically among strategies then we can have Nash Equilibria in *mixed strategies*. We write a mixed strategy as $\sigma_i = p_1 s_{i1} + \cdots + p_k s_{ik}$ where there player i has k available pure strategies $s_{i1}, \ldots s_{ik}$ and $\sum_i p_i = 1$. Thus, a mixed strategy is one where a player chooses strategies with a given probability. The set of all mixed strategies for a player can be written as ΔS_i.Thus we can rewrite a Nash Equilibrium in the general case (involving either pure or mixed strategies) as σ^* such that $u_i(\sigma_i^*, \sigma_{\sim i}^*) \geq u_i(\sigma_i', \sigma_{\sim i}^*) \forall \sigma_i' \in \Delta S_i$ and $\forall i \in N$.

Two important results emerge from considering mixed strategies. First, Nash (1950) proved that if players only have a fixed number of strategies, there is *at least one Nash Equilibrium in (possibly) mixed strategies*. This *existence theory* demonstrates the broad applicability of the equilibrium concept. Second, if it is the case that a mixed strategy is being used where at least two pure strategies are employed with a positive probability (i.e. it is not a pure strategy equilibrium) then each of these strategies must be producing the same utility against the strategies of the other players $\sigma_{\sim i}$. Otherwise, only one of the pure strategies would be used. Generally this means we can solve by considering the (probably) mixed strategy of the other player and equating the utilities of each used pure strategy for the first player. I now provide an example using the famous Battle of the Sexes game below (taken from Gintis, 2009).

Battle of the Sexes describes the situation in which a couple would both prefer to spend time with one another rather than apart but have different preferences about what activity to engage in, with Alfredo preferring gambling and Violetta preferring opera. For whatever reason they cannot coordinate ahead of time and must choose simulataneously. You can probably see two easy pure strategy Nash Equilibria (GG and OO: though note that these cannot be found by iterated elimination of strictly dominated strategies) but there is also a mixed strategy equilibrium. How do we solve for the mixed strategy NE? We assume as above that if Alfredo, for example, chooses both G and O with strictly positive probability, he must in an NE have equal utility for choosing both (otherwise he would simply choose the strategy with higher utility). So it must be that (where α is the probability that Violetta chooses G):

$$U_A^G = \alpha \cdot 2 + (1-\alpha) \cdot 0 = \alpha \cdot 0 + (1-\alpha) \cdot 1 = U_A^O$$

Solving for the α we find $\alpha^* = 1/3$, we find that if Alfredo is choosing a mixed strategy, Violetta must be choosing to gamble one third of the time. So what kind of mixed strategy is Alfredo choosing? For that we simply solve the reverse problem for Violetta, assuming Alfredo chooses G with probability β.

$$U_V^G = \beta \cdot 1 + (1-\beta) \cdot 0 = \beta \cdot 0 + (1-\beta) \cdot 2 = U_V^O$$

Solving this we find that Alfredo chooses to Gamble with $\beta^* = 2/3$. So the mixed strategy is for Alfredo to choose gambling 2/3 of the time, and opera 1/3 of the time, while Violetta chooses gambling 1/3 of the time and opera 2/3 of the time. As Gintis notes, this may be a NE but it is not Pareto optimal unlike the two pure strategy Nash Equilibria. We can see this by noting that Alfredo's utility from this mixed NE is $U_A = \alpha\beta \cdot 2 + (1-\alpha)(1-\beta) \cdot 1 = 2/3$ with Violetta receiving the same utility from a symmetric calculation. Thus each does worse from this outcome than from coordinating on their least preferred activity since they fail to meet up with a positive probability. But once they are playing the mixed strategy they will be 'locked in' since given the behavior of the other, choosing strategies probabilistically remains the ‘best response’.