A Brief Introduction to the Basics of Game Theory

Rory Smead
Northeastern University

September 18, 2013

Game theory is a set of mathematical tools that are used to represent and study social interactions. This is a very brief (and *not* comprehensive) introduction to some of the basic ideas of game theory. In many ways, game theory is an extension of decision theory which is the study of rational choice when we are faced with uncertainty. For this reason, some of the basic concepts from decision theory (e.g. probability and utility) are important for understanding the basics of game theory. The first section covers some central ideas from decision theory and the discussion of game theory begins in section two.

1 Decision Theory

What is the best decision to make when we are faced with uncertainty about the outcomes? This is the central question of decision theory. A decision problem typically consists of (i) a set of possible states of the world, (ii) the probability that each state will occur, (iii) a set of available actions to take and (iv) a value assigned to each act in each state of the world.

1.1 Probability

A probability of a state is just a number between 0 and 1 that represents the chance that state will occur. If S is a possible event, $Pr(S)$ is the probability that S will occur. In decision theory, probabilities are usually interpreted as *subjective degrees of confidence* that an event will occur. Probability is governed by a small set of rules (called the Kolmogorov axioms):

1. All probabilities are in the $[0, 1]$ interval: $0 \leq Pr(S) \leq 1$

2. The probability of a necessary event N is equal to 1: $Pr(N) = 1$

3. The probability that one of two mutually exclusive events S and R occurs is equal to the sum of their probabilities: $Pr(S \text{ } or \text{ } R) = Pr(S) + Pr(R)$

These three axioms plus a definition for conditional probability — $Pr(S \text{ } given \text{ } R) = Pr(S \text{ and } R)/Pr(R)$ — are sufficient to define all of elementary probability theory. Here are a few facts of probability that will be helpful for our purposes:

- S and R are *independent* of one another if and only if $Pr(S) = Pr(S \text{ given } R)$.
- If S and R are independent, $Pr(S \text{ and } R) = Pr(S) \times Pr(R)$, otherwise $Pr(S \text{ and } R) = Pr(S) \times Pr(R \text{ given } S)$.
- If S and R are not mutually exclusive, then $Pr(S \text{ or } R) = Pr(S) + Pr(R) - Pr(S \text{ and } R)$.

1.2 Utility

Taking a particular action in a particular state usually has some consequences. The desirability of those consequences is represented by another number called "utility." The utility of an outcome can be any real number (positive or negative) and the higher number represents more desirability. Someone's utility is just a numerical representation of their preferences, whatever those may be. In other words, utility represents what someone cares about and the degree to which they care about it. $u(X)$ represents the utility of outcome X.

The utility of a particular decision across all possible states can be represented in a *decision matrix*. Suppose there are two possible actions one might take (A and B) and three possible relevant states of the world that might occur (S, R, and T). It is typical to assume that these states are mutually exclusive and exhaustive, meaning one and only one will occur. We can represent the decision problem as follows:

States of the World

Actions	S	R	T
A	1	2	3
B	5	0	0

This means that the most desirable outcome is to choose act B in state S followed by act A in states T, R, and S respectively. However, we can only choose our actions and cannot choose which state of the world will occur. So, if you want to know what act is the best to do, you need to know the probability that each state will occur.

1.3 Expected Utility

We can combine the probability that each state will occur with the utilities of choosing an action for each state. This gives us the *expected utility* for each action. Lets take the example above and suppose that the probabilities for each state are as follows: $Pr(S) = 0.4$, $Pr(R) = 0.4$ and $Pr(S) = 0.2$. Given the for each act-state combination provided above we can calculate the expected utility of each action by combining the probabilities with the utilities:

- $EU(A) = Pr(S) \times u(A \text{ in } S) + Pr(R) \times u(A \text{ in } R) + Pr(T) \times u(A \text{ in } T)$
  $= 0.4 \times 1 + 0.4 \times 2 + 0.2 \times 3 = 1.8$
- $EU(B) = Pr(S) \times u(B \text{ in } S) + Pr(R) \times u(B \text{ in } R) + Pr(T) \times u(B \text{ in } T)$
  $= 0.4 \times 5 + 0.4 \times 0 + 0.2 \times 0 = 2.0$

Since utility represents what one cares about (and the degree that they care about it), one commonly accepted principle of rational decision making is to choose whichever action generates the highest expected utility. In this case, action B has higher expected utility than action A.

The idea in this example can be extended and expressed very generally. If some action A has many possible consequences $C_1, C_2, ..., C_n$, then to figure out the expected utility for that action, you simply consider the utility of each possible consequence multiplied by the probability of that consequence. The expected utility is the sum of those products:

$$EU(A) = \sum_i u(C_i)Pr(C_i).$$

1.4 Dominance

In some special cases, one action may be better than other actions no matter what state of the world we wind up in. When an one act is better than another no matter what happens, this action is said to be *dominant* over the other. There are two forms of dominance: *weak* and *strong*. An act *strongly dominates* another if and only if it leads to higher utility no matter what happens. An act *weakly dominates* another if and only if it does at least as good as the other no matter what happens and does better in at least one instance. Consider the following example:

States of the World

Actions		S	R	T
	A	5	1	3
	B	2	1	3
	C	4	0	2

In this case, we can see that act A does better than act C no matter if we end up in state S, R or T. So, act A *strictly dominates* act C. Also, act A *weakly dominates* act B since it does better when S occurs and equally well when R or T occurs. Recognizing that one action dominates another can be helpful if we don't know the probabilities of the states. If an act dominates another, it doesn't matter what the probabilities of the states are, we will do at least as well or better by choosing the dominate action. In the example here, it would be irrational to choose any action other than A since it dominates both other actions.

2 Game Theory

A game is essentially a multi-person decision problem where the outcome depends, not on some external state of the world, but on the actions of all the players. A game consists of a set of players, a set of strategies (actions) for each player and a payoff function for each player which specifies the utility that player i gets for every possible combination of strategy choices by the players. The most well known game is probably the Prisoner's Dilemma (shown below).

Player 2

Player 1		Cooperate	Defect
	Cooperate	2,2	0,3
	Defect	3,0	1,1

Prisoner's Dilemma

In this game there are two players {Player 1, Player 2}. Player 1's strategy set is {Cooperate, Defect} and Player 2's strategy set is {Cooperate, Defect}. The payoffs for every combination of strategies is given in the game matrix above. The left number in each cell is the payoff to Player 1, the right number in each cell is the payoff to Player 2. Here are some other important terms and ideas for studying games:

- When a game is expressed in a matrix-form as above it is called then *'Normal Form'* or the *'Strategic Form'* of a game.
- An outcome of a game can be specified by a list of everyones' strategy choices: (Player 1's choice, Player 2's choice). Such a list is called a *'Strategy Profile'*.
- A *best response* to other players' strategies is a strategy that generates the highest payoff given what the other players are doing.

Most games we will consider will have just two players, like the Prisoner's Dilemma. And, in many cases, the players players will have the same strategies available to them and have symmetric payoff functions. These games are called two-player symmetric games. It is possible to have more than two players and for different players to have different strategies available to them. We will consider these more complicated games later on, but first we will focus on two-player symmetric games to introduce the central *solution concept* for game theory: the Nash equilibrium.

3 Nash Equilibrium

If we examine the Prisoner's Dilemma game, we will notice that Player 1 does better by playing 'Defect' no matter what Player 2 does (i.e. if Player 2 cooperates, Player 1 gets 3 by playing 'Defect' and 2 by playing 'Cooperate' and if Player 2 defects, Player 1 gets 1 by playing 'Defect' and 0 by playing 'Cooperate'). And likewise, Player 2 does better by playing 'Defect' no matter what Player 1 does. In other words, the strategy 'Defect' *strictly dominates* the strategy 'Cooperate' for both players.

Since choosing dominant actions over non-dominant actions is considered the *rational* choice, the rational solution to the Prisoner's Dilemma is for both players to choose 'Defect.' In other words the strategy profile (Defect, Defect) is the unique rational solution to the Prisoner's Dilemma.

Player 2

Player 1

	Right	Left
Right	1,1	0,0
Left	0,0	1,1

The Driving Game

In many games, however, reasoning from dominance does not produce a solution. Consider the Driving game (above). In this game each player has two strategies: Right and Left. But, neither strategy dominates the other since the best thing to do depends on what the other player does. If the Player 2 drives on the Right, Player one wants to drive on the Right. And, if Player 2 drives on the Left, Player 1 wants to drive on the Left. So, neither strategy dominates the other.

The idea of a *Nash equilibrium* will allow us to find the solutions to games where the idea of dominance doesn't help. Players are at a Nash equilibrium of a game each Player's strategy is a best response to the strategies of all the other players. This idea can be captured as follows:

- ***Nash Equilibrium:*** *A strategy profile is a 'Nash equilibrium' if and only if no player could do better by choosing a different strategy given what the other players are doing.*

- In a two player game, a strategy profile (s, r) is a *pure strategy* Nash equilibrium iff $u_1(s, r) \geq u_i(s', r)$ for all s' in Player 1's strategy set AND $u_2(s, r) \geq u_2(s, r')$ for all r' in Player 2's strategy set.

In the Prisoner's Dilemma, there is only one Nash equilibrium: (Defect, Defect). In the Driving Game, there are at least two Nash equilibria: (Right, Right) and (Left, Left) – in either of those states, neither player wants to change given what the other player is doing. There is a third Nash equilibrium in the Driving Game as well, but this third equilibrium is not a *pure strategy* equilibrium and involves what are called *mixed strategies* (see the next section).

The idea of a Nash equilibrium also works in asymmetric games, when players might have different strategies and different payoff functions. Consider the game below where the two players have different strategies and the payoff functions are also not symmetric.

Player 2

Player 1		a	b	c	d
	r	2,1	0,2	2,5	1,0
	s	4,0	2,1	1,1	1,1
	t	3,3	2,0	0,0	1,3

Example asymmetric game

To find all the pure strategy Nash equilibria of this game, we simply have to go through every possible combination of strategies and ask whether either player would want to switch given what the other is doing. Doing an exhaustive search on the example game above shows us that there are four pure strategy Nash equilibria in this game: (r, c), (s, b), (s, d) and (t, d).

Notice that three of these equilibria actually involve (weakly) dominated strategies. For Player 1, s weakly dominates t since it does better against a and c and equally well against b and d. Likewise, for Player 2, c weakly dominates b. Despite this, however, these strategies can be part of Nash equilibria of the game. Nash equilibria can involve strategies that are weakly dominated, but not strategies that are strongly dominated.

So a Nash equilibrium captures a "solution" to the game. But what does that mean? It turns out that there is not a clear consensus on what the "solution" is supposed to tell us. Does it tell us how we *ought* to play the game? Does it make a prediction about how people *will* play the game? Or does it just tell us what idealized (and fictitious) "rational" beings would do if they played a game—and if that is all it tells us, why is that useful to know? It turns out that the concept of a Nash equilibrium is useful no matter what you think the "solution" really means. It tells us what we ought to do if we want to get the most we can out of a game if everyone else is trying to do the same (and we all know that we are all trying to do this, and we all know that we all know, etc). But, the Nash equilibria of a game are also the points that will govern the way that individuals will learn when playing a game. They also tell us something about how mindless agents (like bacteria) might evolve when playing a game. These are topics that will be discussed later on.

One interesting fact abut all finite games—games with a finite number of strategies and a finite number of players—is that all games have at least one Nash equilibrium. This means that *all* games have at least one rational solution. If the games are large and complicated enough, it may be too difficult to find the solution even though we know there is one. And in some games, the only Nash equilibria cannot be expressed in terms of a specific set of strategies and involve players *mixing* between strategies of the game.

4 Mixed Strategies

Consider the well known game Rock-Paper-Scissors (expressed in normal form below). We can work through every possible combination of strategies to prove what we all already know: no matter combination of strategies you choose, someone wants to switch to a different strategy. This means that there is no *pure strategy Nash equilibrium* in this game. Instead, the best thing to do in this game is to be unpredictable.

We might try to choose a strategy at random. Simply playing r in Rock-Paper-Scissors is called a "pure strategy" but when we try to act unpredicatbly—

Player 2

	r	p	s
Player 1 r	0,0	-1,1	1,-1
p	1,-1	0,0	-1,1
s	-1,1	1,-1	0,0

Rock-Paper-Scissors

sometimes using one strategy, sometimes another—we are using a *mixed strategy.* A mixed strategy is defined by giving a probability to each of the pure strategies in the game. We use the symbol σ to represent a mixed strategy and s to represent pure strategies.

A mixed strategy σ specifies the chance that a player chooses each pure strategy. For example, the mixed strategy where you choose r, p and s with equal chance is: $\sigma = (1/3, 1/3, 1/3)$. And the mixed strategy where we play r half the time and p and s one quarter of the time each is: $\sigma = (1/2, 1/4, 1/4)$. We can also have mixed strategies that ignore some of the pure strategies. For example, the mixed strategy that only uses r and s with equal chance is: $\sigma = (1/2, 0, 1/2)$.

When you know opponent is using a mixed strategy, you can consider the expected utility for each action you might take. Suppose Player 1 believes Player 2 will use a mixed strategy $\sigma_2 = (1/2, 2/6, 1/6)$ in Rock-Paper-Scissors. If $Pr_2(s)$ is the probability that Player 2 uses strategy s, then Player 1 can calculate how well she expects each of her strategies to do against Player 2 using the idea of expected utility. For example, here is Player 1's expected utility from playing r:

$$EU(r, \sigma_2) = Pr_2(r)u_1(r, r) + Pr_2(p)u_1(r, p) + Pr_2(s)u_1(r, s)$$

$$EU(r, \sigma_2) = (1/2)(0) + (2/6)(-1) + (1/6)(1) = -1/6$$

Using the same method to calculate his expected utility for her other two strategies shows that playing p would give the best expected result:

$$EU(p, \sigma_2) = (1/2)(1) + (2/6)(0) + (1/6)(-1) = 2/6$$

$$EU(s, \sigma_2) = (1/2)(-1) + (2/6)(1) + (1/6)(0) = -1/6$$

So, if Player 1 thinks that Player 2 is going to use this strategy, she should just play p and she will win more often than she loses.

However, if Player 2 can anticipate that Player 1 might figure out how to exploit any tendencies toward one strategy over another, Player 2 might decide to just play all strategies with equal chance. And, if Player 2 uses the mixed strategy $\sigma = (1/3, 1/3, 1/3)$, Player 1 will not be able to take advantage of any tendency. In other words, if Player 2 uses the $\sigma = (1/3, 1/3, 1/3)$ strategy, Player 1 will be indifferent to all her strategies.

Furthermore, if Player 1 reasons the same way and is worried about Player 2 potentially taking advantage of any tendencies, Player 1 will also use the $\sigma = (1/3, 1/3, 1/3)$ strategy. When both players use this strategy in Rock-Paper-Scissors, neither player wants to do anything different, and it is a Nash equilibrium of the game. So, Rock-Paper-Scissors has a Nash equilibrium, but only when both players are using just the right *mixed* strategies.

To determine the payoff that two players get when they are both using mixed strategies, we use the *expected utility* as before, but must now combine both players strategies. If Player 1 has n pure strategies and plays a mixed strategy $\sigma_1 = (x_1, x_2, ..., x_n)$ and Player 2 has m strategies and uses a mixed strategy $\sigma_2 = (y_1, y_2, ..., y_m)$, then Player 1's payoff against Player 2 can be calculated as follows:

$$u(\sigma_1, \sigma_2) = \sum_{i=1}^{n} \sum_{j=1}^{m} u(s_i, r_j) x_i y_j.$$

s_i represents Player 1's ith strategy and r_j represents Player 2's jth strategy. In other words, we consider every possible between the two strategies and multiply the payoff of that outcome with the chance that outcome happens. Then, we add up all those possibilities.

4.1 Finding a mixed strategy equilibrium in 2×2 games

Finding mixed strategy equilibria in games can be very hard, especially if the games have many players or many strategies. But, for small games there are some tricks we can use to find mixed strategy Nash equilibria. This section uses an example to demonstrate how to find a mixed strategy Nash equilibrium in games that have 2-players with 2-strategies each (sometimes called '2×2 games').

Notice that when we found the equilibrium for Rock-Paper-Scissors, it was when both players were using strategies that made the other player not care what which strategy they played. This turns out to be a general fact about the mixed strategies that make up Nash equilibria. So, if we can figure out which combination of strategies make players not care which pure strategy they play, that will tell us where the equilibrium is. This is easiest to see with an example, consider the game below.

$Player2$

$Player1$		c	d
	a	$0,0$	$1,3$
	b	$2,1$	$0,0$

Example game #2

This game has three equilibria. Two are pure strategy equilibria: (a, d) and (b, c). But, it also has a third where if both players mix between their strategies just right, neither will do better by doing anything differently. To find this mixed equilibrium, we need to find the strategies that make each player indifferent. First consider the game from the perspective of P_1. If P_2 plays c with probability y and d with probability $(1-y)$, we ask: what value of y makes P_1 indifferent between a and b? The expected utility for each of Player 1's strategies can be determined as follows:

$$EU_1(a,y) = Pr(c)u_1(a,c) + Pr(d)u_1(a,d) = y(0) + (1-y)(1)$$
$$EU_1(b,y) = Pr(c)u_1(b,c) + Pr(d)u_1(b,d) = y(2) + (1-y)(0)$$

Player 1 is indifferent between his two strategies a and b whenever the expected utility of his two strategies is the same:

$$EU_1(b,y) = EU_1(a,y) \text{ or...}$$
$$y(2) + (1-y)(0) = y(0) + (1-y)(1).$$

Solving for y gives us $y = 1/3$. This means that Player 1 will be indifferent between her strategies when Player 2 plays c with probability 1/3 and d with probability 2/3.

Now we do the same thing, but from the perspective of Player 2. Suppose Player 1 plays a with probability x and plays b with probability $(1-x)$. We now ask: what value of x makes Player 2 indifferent between her two strategies?

$$EU_2(c,x) = Pr(a)u_2(a,c) + Pr(b)u_2(b,c) = x(0) + (1-x)(1)$$
$$EU_2(d,x) = Pr(a)u_2(a,d) + Pr(b)u_2(b,d) = x(3) + (1-x)(0)$$

Player 2 is indifferent when $EU_2(d,x) = EV(c,x)$ or whenever $x(3)+(1-x)(0) = x(0)+(1-x)(1)$. Solving for x gives us $x = 1/4$, which means that Player 2 will be indifferent between her two strategies whenever Player 1 plays a with probability 1/4 and b with probability 3/4.

By combining these two calculations, we can get the mixed strategy equilibrium. If Player 1 uses the mixed strategy $\sigma_1 = (1/4, 3/4)$ and Player 2 uses the mixed strategy $\sigma_2 = (1/3, 2/3)$, then both players will be indifferent between their two strategies. This means that neither player can do better by switching their strategy given what the other is doing. So, the strategy profile (σ_1, σ_2) is a Nash equilibrium.

5 The Equilibrium Selection Problem

The game in the previous section (Example game #2) has three Nash equilibria: the two pure strategy equilibria (b, c) and (a, d) along with the mixed strategy equilibrium described in the previous section. Many games, and most of the games that are really interesting, have multiple Nash equilibria. What does game theory predict (or recommend, depending on how we interpret a "solution" to the game) in these cases?

If we want to use game theory make a specific prediction, or to recommend a particular strategy when playing a game, we need to choose which of the "solutions" is the real solution. How we choose among the equilibria will, of course, depend on what we are trying to use the solutions to represent. But, no matter what we are trying to capture, we will need to select which of the equilibria are the relevant ones. This is called the *equilibrium selection problem* in game theory.

Consider the Stag Hunt game (below). In this game there are three equilibria: (c,c), (d,d) and a mixed equilibrium where both players choose c 75% of the time and d the rest of the time. So what would game theory predict people will do when playing this game when every action that the players could take form part of some equilibrium or other? Without a way to approach the equilibrium selection problem, game theory tells us nothing about what to predict (or what to do) in this game.

$Player 2$

$Player 1$		c	d
	c	$4, 4$	$0, 3$
	d	$3, 0$	$3, 3$

The Stag Hunt

There are many different aspects of the game and its various Nash equilibria that we might consider when facing this problem. For example, the (c,c) equilibrium gives the best possible payoff to everyone. This equilibrium is called *'pareto optimal'* and is the *payoff dominant* equilibrium. But an argument can be made that the other equilibrium, (d,d) is the more natural "solution" because it would result from both players avoiding risk. Suppose that both are really not sure what the other will do and think there is an equal chance that the other player will pay any of her strategies. Then, choosing d guarantees the players a decent playoff. The (d,d) equilibrium is the *risk dominant* equilibrium. It is these kind of considerations that can help approach the equilibrium selection problem. However, there is not a unique solution to this problem and the best way to approach it will depend on what we are trying to represent using game theory.

As mentioned above, the equilibrium selection problem arises in many games. This is especially true in games where several moves are made and the order of the moves matters. These games are difficult to represent in the way that we have been doing so far (the normal form), but can be represented in other ways.

6 Extensive Form Games

Not all games are best represented in *normal form*. In some games, one player moves first and the other player can respond. In these games, the order of actions makes a difference and this is not easily represented in the normal form matrices. To capture games where the sequence of actions is important, we use what is called an *'extenstive form'* representation.

Consider the chain-store game. In this game, one company (P1) must choose between entering a new market (e) or not (n). If they choose not to, then everything remains as it is. If they choose to enter the market, another company (P2) must choose whether or not to challenge (c or d). If P2 decides not to challenge, the players split the market. If P2 decides to challenge, they must pay a heavy cost since both players will be forced to sell very cheaply or lose to the other player.

This game is best represented in the *extensive form* (see Figure 1). Which captures the sequence of move in a tree-form. The game can also be expressed in normal form (see below), but this form does not capture the fact that the moves are made in a certain sequence.

$P2$

$P1$		c	d
	e	$-1, -1$	$2, 2$
	n	$0, 4$	$0, 4$

The Chain-Store Game (normal form)

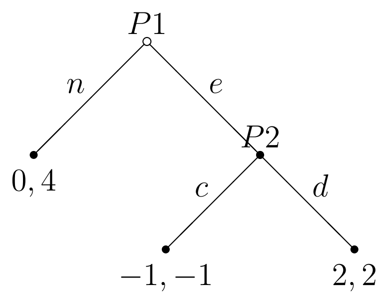

Figure 1: The chain-store game in extensive form

Examining the normal form representation of the chain-store game, we can see that there are two pure-strategy Nash equilibria: (e, d) and (n, c). The first of these is where P1 enters the market and P2 gives up half of it. The second is where P1 does not enter the market and P2 would challenge if P1 were to enter the market. Notice that there is something a little strange about the (n, c) equilibrium—if P1 ever actually chose e, it wouldn't make sense for P2 to chose c, as they would be choosing -1 over 2. It only makes sense for P2 to choose c if P2 never has to make a choice! This means that the (n, c) equilibrium is maintained by a threat that wouldn't be rational to carry out if it came down to it. This equilibrium is not *subgame pefect* (see the next section).

We can also represent uncertainty within extensive form games. And, we can use this uncertainty to represent simultaneous moves using the extensive form. Consider the Driving Game (from above) where r = Right and l = Left. We can suppose that Player 1 does move first, but that Player 2 simply does not know what Player 1 chose. The dotted line between the two choice points indicates that P2 does not know which choice point she is located.

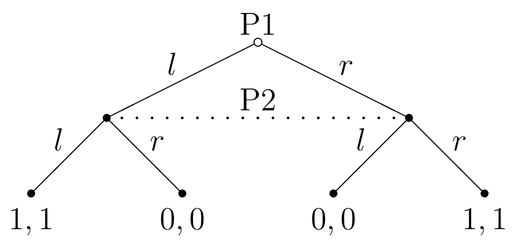

Figure 2: The Driving Game in extensive form

One famous game that is best expressed in extensive form is the Ultimatum Game. Suppose that there are two player that must divide 10 units of some good: the Proposer and the Responder. The Proposer proposes a split to the Responder, who can *accept* or *reject* the split. If the Responder accepts, each gets their share of the proposed split. If the Responder rejects, each gets nothing. To simplify things, we can suppose that the Proposer only has two possible proposals to make: an *unfair* 9-1 split or a *fair* 5-5 split. This simplified game is represented in Figure 3.

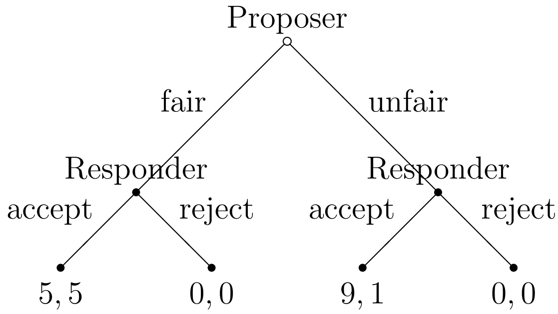

Figure 3: A Simplified Ultimatum Game

A strategy in an extensive form game has to tell the player what to do at every single choice point that they might make. So, the Proposer's strategy set is {fair, unfair}. The Responder's has four possible strategies, which specify what she would say to each possible offer: {accept both, accept fair/reject unfair, reject fair/accept unfair, reject both} or {aa, ar, ra, rr} for short.

There are three pure-strategy Nash equlibria of this game: (unfair, accept both), (unfair, reject fair/accept unfair), and (fair, accept fair/reject unfair). This is easiest to see if we represent the game in normal form (below).

Responder

Proposer		aa	ar	ra	rr
	fair	5,5	5,5	0,0	0,0
	unfair	9,1	0,0	9,1	0,0

Simplified Ultimatum Game (normal form)

7 Subgame Perfection

Two of the equilibria in the simplified Ultimatum Game have a rather strange feature if they are to be considered "rational" solutions. Both the (unfair, reject fair/accept unfair), and (fair, accept fair/reject unfair) equilibria involve rejections where the Responder effectively chooses nothing over something. The only reason that these are equilibria is just that the Responder never actually has to make that choice given what the Proposer is offering.

But, if the Proposer were to ever offer anything different, it wouldn't be in the Responder's interest to reject the offer (something is always better than nothing).

So, these equilibria are maintained by the threat of a possible future action that is irrational.

Some games, such as the simplified Ultimatum Game, contain what are called *'subgames'*. A subgame is a game-within-a-game. The simplified Ultimatum Game (in Figure 3) has three subgames: (i) the whole game, (ii) the fork in the tree on the left side beginning at the Responder's choice, and (iii) the fork in the tree on the right side beginning at the Responder's choice. Using this idea, we can define a more "rational" version of the Nash equilibrium: the *subgame perfect equilibrium.*

- ***Subgame Perfect Equilibrium:*** *A Nash equilibrium of a game that also forms a Nash equilibrium of all the subgames.*

In the case of the simplified Ultimatum Game, there is only one subgame perfect equilibrium: (unfair, accept both). Thinking of the two small subgames, it is irrational for the Responder to choose reject in either case. And, given that it is always rational for the Responder to accept, it is rational for the Proposer to offer the unfair split. In the other two Nash equilibria of the game, the Responder has a strategy that does not constitute a Nash equilibrium of one or other of the subgames.

The idea of a *subgame perfect equilibrium* is one way that we might go about solving the equilibrium selection problem in extensive form games. This might be of special interest if we want our solutions to capture the "rational" outcomes of a game.

7.1 Backward Induction

One common method for finding a subgame perfect equilibrium is called *backward induction.* This method begins by looking at the very end of the game and asking: if this decision point were reached, what would a rational player do? Then imagining that the game ends with that rational decision, you move on to the previous decision in the game and ask again: given that the last choice will be rational, what is the rational choice to make at the second-to-last choice points? And so on, to the beginning of the game.

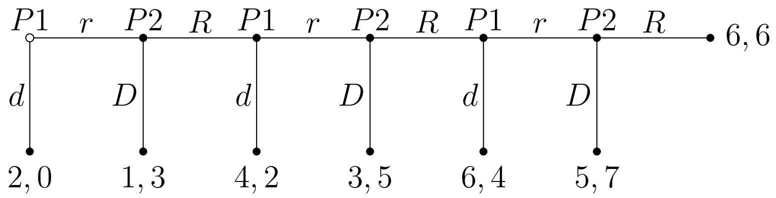

Figure 4: The Centipede Game

Consider the Centipede Game (Figure 4). In this game, players take turns choosing whether or not to take the majority of a pool of resources and end the game or to pass the choice to the other player, thereby increasing the pool of resources. The game continues for a fixed number of rounds, if no one takes the majority, the game ends with an equal split.

According to the method of backward induction, lets begin at the end of the game. At the end of the game, Player 2 should choose to take the majority and end the game (play D). Given that player two should do this, Player 1 should take the majority one round sooner and end the game. Given that Player 1 should do this at the second-to-last point, Player 2 should take the majority at the the third-to-last-point, and so on. The subgame perfect equilibrium is for both players to take the majority at every possible choice point: (ddd, DDD).