

A Brief Introduction to the Basics of Game Theory

Matthew O. Jackson, Stanford University

I provide a (very) brief introduction to game theory. I have developed these notes to provide quick access to some of the basics of game theory; mainly as an aid for students in courses in which I assumed familiarity with game theory but did not require it as a prerequisite. Of course, the material discussed here is only the proverbial tip of the iceberg, and there are many sources that offer much more complete treatments of the subject.¹ Here, I only cover a few of the most fundamental concepts, and provide just enough discussion to get the ideas across without discussing many issues associated with the concepts and approaches. Fuller coverage is available through a free on-line course that can be found via my website: <http://www.stanford.edu/~jacksonm/>

The basic elements of performing a noncooperative ² game-theoretic analysis are (1) framing the situation in terms of the actions available to players and their payoffs as a function of actions, and (2) using various equilibrium notions to make either descriptive or

¹ For graduate-level treatments, see Roger Myerson's (1991) *Game Theory: Analysis of Conflict*, Cambridge, Mass.: Harvard University Press; Ken Binmore's (1992) *Fun and Games*, Lexington, Mass.: D.C. Heath; Drew Fudenberg and Jean Tirole's (1993) *Game Theory*, Cambridge, Mass.: MIT Press; and Martin Osborne and Ariel Rubinstein's (1994) *A Course in Game Theory*, Cambridge, Mass.: MIT Press. There are also abbreviated texts offering a quick tour of game theory, such as Kevin Leyton-Brown and Yoav Shoham's (2008) *Essentials of Game Theory*, Morgan and Claypool Publishers. For broader readings and undergraduate level texts, see R. Duncan Luce and Howard Raiffa (1959) *Games and Decisions: Introduction and Critical Survey*; Robert Gibbons (1992) *Game Theory for Applied Economists*; Colin F. Camerer (2003) *Behavioral Game Theory: Experiments in Strategic Interaction*; Martin J. Osborne (2003) *An Introduction to Game Theory*; Joel Watson (2007) *Strategy: An Introduction to Game Theory*; Avinash K. Dixit and Barry J. Nalebuff (2010) *The Art of Strategy: A Game Theorist's Guide to Success in Business and Life*; Joseph E. Harrington, Jr. (2010) *Games, Strategies, and Decision Making*, Worth Publishing.

²“Noncooperative game theory” refers to models in which each players are assumed to behave selfishly and their behaviors are directly modeled. “Cooperative game theory,” which I do not cover here, generally refers to more abstract and axiomatic analyses of bargains or behaviors that players might reach, without explicitly modeling the processes. The name “cooperative” derives in part from the fact that the analyses often (but not always) incorporate coalitional considerations, with important early analyses appearing in John von Neumann and Oskar Morgenstern's 1944 foundational book “*Theory of Games and Economic Behavior*.”

prescriptive predictions. In framing the analysis, a number of questions become important. First, who are the players? They may be people, firms, organizations, governments, ethnic groups, and so on. Second, what actions are available to them? All actions that the players might take that could affect any player's payoffs should be listed. Third, what is the timing of the interactions? Are actions taken simultaneously or sequentially? Are interactions repeated? The order of play is also important. Moving after another player may give player i an advantage of knowing what the other player has done; it may also put player i at a disadvantage in terms of lost time or the ability to take some action. What information do different players have when they take actions? Fourth, what are the payoffs to the various players as a result of the interaction? Ascertaining payoffs involves estimating the costs and benefits of each potential set of choices by all players. In many situations it may be easier to estimate payoffs for some players (such as yourself) than others, and it may be unclear whether other players are also thinking strategically. This consideration suggests that careful attention be paid to a sensitivity analysis.

Once we have framed the situation, we can look from different players' perspectives to analyze which actions are optimal for them. There are various criteria we can use.

1 Games in Normal Form

Let us begin with a standard representation of a game, which is known as a *normal form* game, or a game in *strategic form*:

- The set of players is $N = \{1, \dots, n\}$.
- Player i has a set of actions, a_i , available. These are generally referred to as *pure strategies*.³ This set might be finite or infinite.
- Let $a = a_1 \times \dots \times a_n$ be the set of all *profiles* of pure strategies or actions, with a generic element denoted by $a = (a_1, \dots, a_n)$.

³The term "pure" indicates that a single action is chosen, in contrast with "mixed" strategies that I discuss below, in which there is a randomization over actions.

- Player i 's payoff as a function of the vector of actions taken is described by a function $u_i : A \rightarrow \mathbb{R}$, where $u_i(a)$ is i 's payoff if the a is the profile of actions chosen in the society.

Normal form games are often represented by a table. Perhaps the most famous such game is the *prisoners' dilemma*, which is represented in Table 1. In this game there are two players who each have two pure strategies, where $a_i = \{C, D\}$, and C stands for “cooperate” and D stands for “defect.” The first entry indicates the payoff to the row player (or player 1) as a function of the pair of actions, while the second entry is the payoff to the column player (or player 2).

Table 1: A Prisoners' Dilemma Game

		Player 2	
		C	D
Player 1	C	-1, -1	-3, 0
	D	0, -3	-2, -2

The usual story behind the payoffs in the prisoners' dilemma is as follows. The two players have committed a crime and are now in separate rooms in a police station. The prosecutor has come to each of them and told them each: “If you confess and agree to testify against the other player, and the other player does not confess, then I will let you go. If you both confess, then I will send you both to prison for 2 years. If you do not confess and the other player does, then you will be convicted and I will seek the maximum prison sentence of 3 years. If nobody confesses, then I will charge you with a lighter crime for which we have enough evidence to convict you and you will each go to prison for 1 year.” So the payoffs in the matrix represent time lost in terms of years in prison. The term *cooperate* refers to cooperating with the other player. The term *defect* refers to confessing and agreeing to testify, and so breaking the (implicit) agreement with the other player.

Note that we could also multiply each payoff by a scalar and add a constant, which is an equivalent representation (as long as all of a given player's payoffs are rescaled in the same

way). For instance, in Table 2 I have doubled each entry and added 6. This transformation leaves the strategic aspect of the game unchanged.

Table 2: A Rescaling of the Prisoners' Dilemma

		Player 2	
		C	D
Player 1	C	4, 4	0, 6
	D	6, 0	2, 2

There are many games that might have different descriptions motivating them but have a similar normal form in terms of the strategic aspects of the game. Another example of the same game as the prisoners' dilemma is what is known as a *Cournot duopoly*. The story is as follows. Two firms produce identical goods. They each have two production levels, high or low. If they produce at high production, they will have a lot of the goods to sell, while at low production they have less to sell. If they cooperate, then they agree to each produce at low production. In this case, the product is rare and fetches a very high price on the market, and they each make a profit of 4. If they each produce at high production (or defect), then they will depress the price, and even though they sell more of the goods, the price drops sufficiently to lower their overall profits to 2 each. If one defects and the other cooperates, then the price is in a middle range. The firm with the higher production sells more goods and earns a higher profit of 6, while the firm with the lower production just covers its costs and earns a profit of 0.

1.1 Dominant Strategies

Given a game in normal form, we then can make predictions about which actions will be chosen. Predictions are particularly easy when there are “dominant strategies.” A dominant strategy for a player is one that produces the highest payoff of any strategy available *for every possible action by the other players*.

That is, a strategy $a_i \in a_i$ is a *dominant* (or weakly dominant) strategy for player i if

$u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i})$ for all a'_i and all $a_{-i} \in a_{-i}$. A strategy is a *strictly dominant strategy* if the above inequality holds strictly for all $a'_i \neq a_i$ and all $a_{-i} \in a_{-i}$.

Dominant strategies are powerful from both an analytical point of view and a player's perspective. An individual does not have to make any predictions about what other players might do, and still has a well-defined best strategy.

In the prisoners' dilemma, it is easy to check that each player has a strictly dominant strategy to defect—that is, to confess to the police and agree to testify. So, if we use dominant strategies to predict play, then the unique prediction is that each player will defect, and both players fare worse than for the alternative strategies in which neither defects. A basic lesson from the prisoners' dilemma is that individual incentives and overall welfare need not coincide. The players both end up going to jail for 2 years, even though they would have gone to jail for only 1 year if neither had defected. The problem is that they cannot trust each other to cooperate: no matter what the other player does, a player is best off defecting.

Note that this analysis presumes that all relevant payoff information is included in the payoff function. If, for instance, a player fears retribution for confessing and testifying, then that should be included in the payoffs and can change the incentives in the game. If the player cares about how many years the other player spends in jail, then that can be written into the payoff function as well.

When dominant strategies exist, they make the game-theoretic analysis relatively easy. However, such strategies do not always exist, and then we can turn to notions of equilibrium.

1.2 Nash Equilibrium

A pure strategy Nash equilibrium⁴ is a profile of strategies such that each player's strategy is a best response (results in the highest available payoff) against the equilibrium strategies of the other players.

⁴The concept is named after John Nash, who provided the first existence proof in finite games: Nash, J.F. (1951) Non-Cooperative Games, *Annals of Mathematics* 54:286-295. On occasion it is also referred to as *Cournot-Nash equilibrium*, with reference to Antoine Augustin Cournot, who in the 1830's first developed such an equilibrium concept in the analysis of oligopoly (a set of firms in competition with one another) : Cournot (1838) *Recherches sur les principes mathématiques de la théorie des richesses*, translated as: *Researches into the Mathematical Principles of the Theory of Wealth*, New York: Macmillan (1897).

A strategy a_i is a *best reply*, also known as a *best response*, of player i to a profile of strategies $a_{-i} \in A_{-i}$ for the other players if

$$u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i})$$

for all a'_i . A best response of player i to a profile of strategies of the other players is said to be a *strict best response* if it is the unique best response.

A profile of strategies $a \in A$ is a *pure strategy Nash equilibrium* if a_i is a best reply to a_{-i} for each i . That is, a is a Nash equilibrium if

$$u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i})$$

for all i and a'_i . This definition might seem somewhat similar to that of dominant strategy, but there is a critical difference. A pure strategy Nash equilibrium only requires that the action taken by each agent be best against the actual equilibrium actions taken by the other players, and not necessarily against all possible actions of the other players.

A Nash equilibrium has the nice property that it is stable: if each player expects a to be the profile of actions played, then no player has any incentive to change his or her action. In other words, no player regrets having played the action that he or she played in a Nash equilibrium.

In some cases, the best response of a player to the actions of others is unique. A Nash equilibrium such that all players are playing actions that are unique best responses is called a *strict Nash equilibrium*. A profile of dominant strategies is a Nash equilibrium but not vice versa.

To see another illustration of Nash equilibrium, consider the following game between two firms that are deciding whether to advertise. Total available profits are 28, to be split between the two firms. Advertising costs a firm 8. Firm 1 currently has a larger market share than firm 2, so it is seeing 16 in profits while firm 2 is seeing 12 in profits. If they both advertise, then they will split the market evenly and get 14 in base profits each, but then must also pay the costs of advertising, so they receive net profits of 6 each. If one advertises while the other does not, then the advertiser captures three-quarters of the market (but also pays for advertising) and the non-advertiser gets one-quarter of the market. (There

are obvious simplifications here: just considering two levels of advertising and assuming that advertising only affects the split and not the total profitability.) The net payoffs are given in the Table 3.

Table 3: An Advertising Game

		Firm 2	
		Not	Adv
Firm 1	Not	16, 12	7, 13
	Adv	13, 7	6, 6

To find the equilibrium, we have to look for a pair of actions such that neither firm wants to change its action given what the other firm has chosen. The search is made easier in this case, since firm 1 has a strictly dominant strategy of not advertising. Firm 2 does not have a dominant strategy; which strategy is optimal for it depends on what firm 1 does. But given the prediction that firm 1 will not advertise, firm 2 is best off advertising. This forms a Nash equilibrium, since neither firm wishes to change strategies. You can easily check that no other pairs of strategies form an equilibrium.

While each of the previous games provides a unique prediction, there are games in which there are multiple equilibria. Here are three examples.

EXAMPLE 1 *A Stag Hunt Game* The first is an example of a coordination game, as depicted in Table 4. This game might be thought of as selecting between two technologies, or coordinating on a meeting location. Players earn higher payoffs when they choose the same action than when they choose different actions. There are two (pure strategy) Nash equilibria: (S, S) and (H, H) .

This game is also a variation on Rousseau’s “stag hunt” game.⁵ The story is that two hunters are out, and they can either hunt for a stag (strategy S) or look for hares (strategy H). Succeeding in getting a stag takes the effort of both hunters, and the hunters are separated

⁵To be completely consistent with Rousseau’s story, (H, H) should result in payoffs of $(3, 3)$, as the payoff to hunting for hare is independent of the actions of the other player in Rousseau’s story.

Table 4: A Coordination Game

		Player 2	
		S	H
Player 1	S	5, 5	0, 3
	H	3, 0	4, 4

in the forest and cannot be sure of each other's behavior. If both hunters are convinced that the other will hunt for stag, then hunting stag is a strict or unique best reply for each player. However, if one turns out to be mistaken and the other hunter hunts for hare, then one will go hungry. Both hunting for hare is also an equilibrium and hunting for hare is a strict best reply if the other player is hunting for hare. This example hints at the subtleties of making predictions in games with multiple equilibria. On the one hand, (S, S) (hunting stag by both) is a more attractive equilibrium and results in high payoffs for both players. Indeed, if the players can communicate and be sure that the other player will follow through with an action, then playing (S, S) is a stable and reasonable prediction. However, (H, H) (hunting hare by both) has properties that make it a useful prediction as well. It does not offer as high a payoff, but it has less risk associated with it. Here playing H guarantees a minimum payoff of 3, while the minimum payoff to S is 0. There is an extensive literature on this subject, and more generally on how to make predictions when there are multiple equilibria.⁶

EXAMPLE 2 *A “Battle of the Sexes” Game The next example is another form of coordination game, but with some asymmetries in it. It is generally referred to as a “battle of the sexes” game, as depicted in Table 5.*

The players have an incentive to choose the same action, but they each have a different favorite action. There are again two (pure strategy) Nash equilibria: (X, X) and (Y, Y) . Here, however, player 1 would prefer that they play equilibrium (X, X) and player 2 would prefer (Y, Y) . The battle of the sexes title refers to a couple trying to coordinate on where to meet for a night out. They prefer to be together, but also have different preferred outings.

⁶See, for example, the texts cited in Footnote 1.

Table 5: A “battle of the sexes” game

		Player 2	
		X	Y
Player 1	X	3, 1	0, 0
	Y	0, 0	1, 3

EXAMPLE 3 *Hawk-Dove and Chicken Games* There are also what are known as anti-coordination games, with the prototypical version being what is known as the hawk-dove game or the chicken game, with payoffs as in Table 6.

Table 6: A “hawk-dove” game

		Player 2	
		Hawk	Dove
Player 1	Hawk	0, 0	3, 1
	Dove	1, 3	2, 2

Here there are two pure strategy equilibria, $(Hawk, Dove)$ and $(Dove, Hawk)$. Players are in a potential conflict and can be either aggressive like a hawk or timid like a dove. If they both act like hawks, then the outcome is destructive and costly for both players with payoffs of 0 for both. If they each act like doves, then the outcome is peaceful and each gets a payoff of 2. However, if the other player acts like a dove, then a player would prefer to act like a hawk and take advantage of the other player, receiving a payoff of 3. If the other player is playing a hawk strategy, then it is best to play a dove strategy and at least survive rather than to be hawkish and end in mutual destruction.

1.3 Randomization and Mixed Strategies

In each of the above games, there was at least one pure strategy Nash equilibrium. There are also simple games for which pure strategy equilibrium do not exist. To see this, consider the

following simple variation on a penalty kick in a soccer match. There are two players: the player kicking the ball and the goalie. Suppose, to simplify the exposition, that we restrict the actions to just two for each player (there are still no pure strategy equilibria in the larger game, but this simplified version makes the exposition easier). The kicking player can kick to the left side or to the right side of the goal. The goalie can move to the left side or to the right side of the goal and has to choose before seeing the kick, as otherwise there is too little time to react. To keep things simple, assume that if the player kicks to one side, then she scores for sure if the goalie goes to the other side, while the goalie is certain to save it if the goalie goes to the same side. The basic payoff structure is depicted in Table 7.

Table 7: A Penalty-Kick Game.

		Goalie	
		L	R
Kicker	L	-1, 1	1, -1
	R	1, -1	-1, 1

This is also the game known as “matching pennies.” The goalie would like to choose a strategy that matches that of the kicker, and the kicker wants to choose a strategy that mismatches the goalie’s strategy.⁷

It is easy to check that no pair of pure strategies forms an equilibrium. What is the solution here? It is just what you see in practice: the kicker randomly picks left versus right, in this particular case with equal probability, and the goalie does the same. To formalize this observation we need to define randomized strategies, or what are called *mixed strategies*. For ease of exposition suppose that a_i is finite; the definition extends to infinite strategy spaces with proper definitions of probability measures over pure actions.

⁷For an interesting empirical test of whether goalies and kickers on professional soccer teams randomize properly, see Chiappori, Levitt, and Grosseclouse (2002) Testing Mixed-Strategy Equilibria When Players Are Heterogeneous: The Case of Penalty Kicks in Soccer, *American Economic Review* 92(4):1138 - 1151; and see Walker and Wooders (2001) Minimax Play at Wimbledon, *American Economic Review* 91(5):1521 - 1538. For an analysis of randomization in the location of tennis serves in professional tennis matches.

A *mixed strategy* for a player i is a distribution s_i on a_i , where $s_i(a_i)$ is the probability that a_i is chosen. A profile of mixed strategies (s_1, \dots, s_n) forms a *mixed-strategy Nash equilibrium* if

$$\sum_a \left(\prod_j s_j(a_j) \right) u_i(a) \geq \sum_{a_{-i}} \left(\prod_{j \neq i} s_j(a_j) \right) u_i(a'_i, a_{-i})$$

for all i and a'_i .

So a profile of mixed strategies is an equilibrium if no player has some strategy that would offer a better payoff than his or her mixed strategy in reply to the mixed strategies of the other players. Note that this reasoning implies that a player must be indifferent to each strategy that he or she chooses with a positive probability under his or her mixed strategy. Also, players' randomizations are independent.⁸ A special case of a mixed strategy is a pure strategy, where probability 1 is placed on some action.

It is easy to check that each mixing with probability 1/2 on L and R is the unique mixed strategy of the matching pennies game above. If a player, say the goalie, places weight of more than 1/2 on L, for instance, then the kicker would have a best response of choosing R with probability 1, but then that could not be an equilibrium as the goalie would want to change his or her action, and so forth.

There is a deep and long-standing debate about how to interpret mixed strategies, and the extent to which people really randomize. Note that in the goalie and kicker game, what is important is that each player not know what the other player will do. For instance, it could be that the kicker decided before the game that if there was a penalty kick then she would kick to the left. What is important is that the kicker not be known to always kick to the left.⁹

We can begin to see how the equilibrium changes as we change the payoff structure. For example, suppose that the kicker is more skilled at kicking to the right side than to the left.

⁸An alternative definition of correlated equilibrium allows players to use correlated strategies but requires some correlating device that only reveals to each player his or her prescribed strategy and that these are best responses given the conditional distribution over other players' strategies.

⁹The contest between pitchers and batters in baseball is quite similar. Pitchers make choices about the location, velocity, and type of pitch (e.g., whether various types of spin are put on the ball). If a batter knows what pitch to expect in a given circumstance, that can be a significant advantage. Teams scout one another's players and note any tendencies or biases that they might have and then try to respond accordingly.

In particular, keep the payoffs as before, but now suppose that the kicker has an even chance of scoring when kicking right when the goalie goes right. This leads to the payoffs in Table 8.

Table 8: A biased penalty-kick game

		Goalie	
		L	R
Kicker	L	-1, 1	1, -1
	R	1, -1	0, 0

What does the equilibrium look like? To calculate the equilibrium, it is enough to find a strategy for the goalie that makes the kicker indifferent, and a strategy for the kicker that makes the goalie indifferent.¹⁰

Let s_1 be the kicker's mixed strategy and s_2 be the goalie's mixed strategy. It must be that the kicker is indifferent. The kicker's expected payoff from kicking L is $-1 \cdot s_2(L) + 1 \cdot s_2(R)$ ¹¹ and the payoff from R is $1 \cdot s_2(L) + 0 \cdot s_2(R)$, so that indifference requires that

$$-s_2(L) + s_2(R) = s_2(L),$$

which implies that $2s_2(L) = s_2(R)$. Since these must sum to one (as they are probabilities), this implies that $s_2(L) = 1/3$ and $s_2(R) = 2/3$. Similar calculations based on the requirement that the goalie be indifferent lead to

$$s_1(L) - s_1(R) = -s_1(L),$$

¹⁰This reasoning is a bit subtle, as we are not directly choosing actions that maximize the goalie's payoff and maximize the kicker's payoff, but instead are looking for a mixture by one player that makes the other indifferent. This feature of mixed strategies takes a while to grasp, but experienced players seem to understand it well, as discussed below.

¹¹To see where this payoff comes from, note that there is a $s_2(L)$ chance that the goalie also goes L and then the kicker loses and gets a payoff of -1, and a $s_2(R)$ chance that the goalie goes right and then the kicker wins and gets a payoff of 1; thus the expected payoff is $-1 \cdot s_2(L) + 1 \cdot s_2(R)$

and so the kicker's equilibrium strategy must satisfy $2s_1(L) = s_2(R)$, which this implies that $s_1(L) = 1/3$ and $s_1(R) = 2/3$.

Note that as the kicker gets more skilled at kicking to the right, they both adjust to using the right strategy more. The goalie ends up using the R strategy with higher probability than before even though that strategy has gotten worse for the goalie in terms of just looking at each entry of Table 8 compared to Table 7. This reflects the strategic aspect of the game: each player's strategy reacts to the other's strategy, and not just absolute changes in payoffs as one might superficially expect. The kicker using R more means that the goalie is still indifferent with the new payoffs, and the goalie has to adjust to using R more in order to keep the kicker indifferent.¹²

While not all games have pure strategy Nash equilibrium, every game with a finite set of actions has at least one mixed strategy Nash equilibrium (with a special case of a mixed strategy equilibrium being a pure strategy equilibrium), as shown in the seminal paper by John Nash (1951) Non-Cooperative Games, *Annals of Mathematics* 54:286 - 295.

2 Sequentiality, Extensive Form Games, and Backward Induction

Let us now turn to the question of timing. In the above discussion it was implicit that each player was selecting a strategy with beliefs about the other players' strategies but without knowing exactly what they were.

If we wish to be more explicit about timing, then we can consider what are known as games in *extensive form*, which include a complete description of who moves in what order and what they have observed when they move.¹³ There are advantages to working with

¹²Interestingly, there is evidence that professional soccer players are better at playing games that have mixed strategy equilibria than are people with less experience in such games, which is consistent with this observation (see Palacios-Huerta and Volij (2008) "Experientia Docet: Professionals Play Minimax in Laboratory Experiments," *Econometrica*, 76:1, pp 71 - 115.

¹³One can collapse certain types of extensive form games into normal form by simply defining an action to be a complete specification of how an agent would act in all possible contingencies. Agents then choose these actions simultaneously at the beginning of the game. But the normal form becomes more complicated

extensive form games, as they allow more explicit treatments of timing and for equilibrium concepts that require credibility of strategies in response to the strategies of others.

Definitions for a general class of extensive form games are notationally intensive. Here I will just discuss a special class of extensive form games—finite games of perfect information—which allows for a treatment that avoids much of the notation. These are games in which players move sequentially in some pre-specified order (sometimes contingent on which actions have been chosen), each player moves at most a finite number of times, and each player is completely aware of all moves that have been made previously. These games are particularly well behaved and can be represented by simple trees, where a “nonterminal” node is associated with the move of a specified player and an edge corresponds to different actions the player might take, and “terminal” nodes (that have no edges following them) list the payoffs if those nodes are reached, as in Figure 1. I will not provide formal definitions, but simply refer directly to games representable by such finite game trees.




Figure 1: A Game Tree with 3 Players and Two Actions Each.

Each node has a player’s label attached to it. There is an identified *root node* that corresponds to the first player to move (player 1 in Figure 1) and then subsequent nodes

than the two-by-two games discussed above.

that correspond to subsequent players who make choices. In Figure 1, player 1 has a choice of moving either left or right. The branches in the tree correspond to the different actions available to the player at a given node. In this game, if player 1 moves left, then player 2 moves next; while if player 1 moves right, then player 3 moves next. It is also possible to have trees in which player 1 chooses twice in a row, or no matter what choice a given player makes it is a certain player who follows, and so forth. The payoffs are given at the end nodes and are listed for the respective players. The top payoff is for player 1, the second for player 2, and the bottom for player 3. So the payoffs depend on the set of actions taken, which then determines a path through the tree. An equilibrium provides a prediction about how each player will move in each contingency and thus makes a prediction about which path will be taken; we refer to that prediction as the *equilibrium path*.

We can apply the concept of a Nash equilibrium to such games, which here is a specification of what each player would do at each node with the requirement that each player's strategy be a best response to the other players' strategies. Nash equilibrium does not always make sensible predictions when applied to the extensive form. For instance, reconsider the advertising example discussed above in Table 3. Suppose that firm 1 makes its decision of whether to advertise before firm 2 does, and that firm 2 learns firm 1's choice before it chooses. This scenario is represented in the game tree pictured in Figure 2.

To apply the Nash equilibrium concept to this extensive form game, we must specify what each player does at each node. There are two Nash equilibria of this game in pure strategies. The first is where firm 1 advertises, and firm 2 does not (and firm 2's strategy conditional on firm 1 not advertising is to advertise). The other equilibrium corresponds to the one identified in the normal form: firm 1 does not advertise, and firm 2 advertises regardless of what firm 1 does. This is an equilibrium, since neither wants to change its behavior, given the other's strategy. However, it is not really credible in the following sense: it involves firm 2 advertising even after it has seen that firm 1 has advertised, and even though this action is not in firm 2's interest in that contingency.

To capture the idea that each player's strategy has to be credible, we can solve the game backward. That is, we can look at each decision node that has no successor, and start by making predictions at those nodes. Given those decisions, we can roll the game backward




Figure 2: Advertising Choices of Two Competitors

and decide how player's will act at next-to-last decision nodes, anticipating the actions at the last decision nodes, and then iterate. This is called *backward induction*. Consider the choice of firm 2, given that firm 1 has decided not to advertise. In this case, firm 2 will choose to advertise, since 13 is larger than 12. Next, consider the choice of firm 2, given that firm 1 has decided to advertise. In this case, firm 2 will choose not to advertise, since 7 is larger than 6. Now we can collapse the tree. Firm 1 will predict that if it does not advertise, then firm 2 will advertise, while if firm 1 advertises then firm 2 will not. Thus when making its choice, firm 1 anticipates a payoff of 7 if it chooses not to advertise and 13 if it chooses to advertise. Its optimal choice is to advertise. The backward induction prediction about the actions that will be taken is for firm 1 to advertise and firm 2 not to.

Note that this prediction differs from that in the simultaneous move game we analyzed before. Firm 1 has gained a first-mover advantage in the sequential version. Not advertising is no longer a dominant strategy for firm 1, since firm 2's decision depends on what firm 1 does. By committing to advertising, firm 1 forces firm 2 to choose not to advertise. Firm 1 is better off being able to commit to advertising in advance.

A solution concept that capture found in this game and applies to more general classes of

games is known as *subgame perfect equilibrium* (due to Reinhard Selten (1975) Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games, International Journal of Game Theory 4:25 - 55). A subgame in terms of a finite game tree is simply the subtree that one obtains starting from some given node. Subgame perfection requires that the stated strategies constitute a Nash equilibrium in every subgame (including those with only one move left). So it requires that if we start at any node, then the strategy taken at that node must be optimal in response to the remaining specification of strategies. In the game between the two firms, it requires that firm 2 choose an optimal response in the subgame following a choice by firm 1 to advertise, and so it coincides with the backward induction solution for such a game.

It is worth noting that moving first is not always advantageous. Sometimes it allows one to commit to strategies which would otherwise be untenable, which can be advantageous; but in other cases it may be that the information that the second mover gains from knowing which strategy the first mover has chosen is a more important consideration. For example, suppose that the matching pennies game we discussed above were to be played sequentially so that the kicker had to kick first and the goalie had time to see the kicker's action and then to react and could jump left or right to match the kicker's choice: the advantage would certainly then tip towards the goalie.

This concludes our whirlwind tour of some of the basic tools of game theory. There are many important subjects that I have not touched upon here, including analyses that incorporate incomplete information, repeated games, and behavioral game theory. However, this should provide you with some feeling for a few of the most prominent concepts, and some of the approaches that form the backbone of game theoretic analyses.

3 Some Exercises

EXERCISE 1 *Product Choices.*

Two electronics firms are making product development decisions. Each firm is choosing between the development of two alternative computer chips. One system has higher efficiency, but will require a larger investment and will be more costly to produce. Based on estimates

of development costs, production costs, and demand, the following present value calculations represent the value of the alternatives (high efficiency chips or low efficiency chips) to the firms.

Table 9: A production-choice game

		Firm 2	
		High	Low
Firm 1	High	1, 2	4, 5
	Low	2, 7	5, 3

The first entry in each box is the present value to firm 1 and the second entry is the present value to firm 2. The payoffs in the above table are not symmetric. Firm 2 has a cost advantage in producing the higher efficiency chip, while firm 1 has a cost advantage in producing the lower efficiency chip. Overall profits are largest when the firms choose different chips and do not compete head to head.

- (a) Firm 1 has a dominant strategy. What is it?
- (b) Given your answer to part a), what should firm 2 expect firm 1's choice to be? What is firm 2's optimal choice given what it anticipates firm 1 to do?
- (c) Do firm 1's strategy (answer to (a)) and firm 2's strategy (answer to (b)) form an equilibrium? Explain.
- (d) Compared to (c), firm 1 would make larger profits if the choices were reversed. Why don't those strategies form an equilibrium?
- (e) Suppose that firm 1 can commit to a product before firm 2. Draw the corresponding game tree and describe the backward induction/subgame perfect equilibrium.

EXERCISE 2 *Hotelling's Hotels.*

Two hotels are considering a location along a newly constructed highway through the desert. The highway is 500 miles long with an exit every 50 miles (including both ends). The hotels may choose to locate at any exit. These will be the only hotels for any traveler using the highway. Each traveler has their own most preferred location along the highway (at some exit) for a hotel, and will choose to go the hotel closest to that location. Travelers most preferred locations are distributed evenly, so that each exit has the same number of travelers who prefer that exit. If both hotels are the same distance from a traveler's most preferred location, then that traveler flips a coin to determine which hotel to stay at. A hotel would each like to maximize the number of travelers who stay at it.

If Hotel 1 locates at the 100 mile exit, where should Hotel 2 locate?

Given Hotel 2's location that you just found, where would Hotel 1 prefer to locate?

Which pairs of locations form Nash equilibria?

EXERCISE 3 *Backward Induction.*

Find the backward induction solution to Figure 1 and argue that there is a unique subgame perfect equilibrium. Provide a Nash equilibrium of that game that is not subgame perfect.

EXERCISE 4 *The Colonel Blotto Game.*

Two armies are fighting a war. There are three battlefields. Each army consists of 6 units. The armies must each decide how many units to place on each battlefield. They do this without knowing how many units the other army has committed to a given battlefield. The army who has the most units on a given battlefield, wins that battle, and the army that wins the most battles wins the war. If the armies each have the same number of units on a given battlefield then there is an equal chance that either army wins that battle. A pure strategy for an army is a list (u_1, u_2, u_3) of the number of units it places on battlefields 1, 2, and 3 respectively, where each u_k is in $\{0, 1, \dots, 6\}$ and the sum of the u_k 's is 6. For example, if army A allocates its units $(3, 2, 1)$, and army B allocates its units $(0, 3, 3)$, then army A wins the first battle, and army B wins the second and third battles and army B wins the war.

Argue that there is no pure strategy Nash equilibrium to this game.

Argue that mixing uniformly at random over all possible configurations of units is not a mixed strategy Nash equilibrium (hint - show that placing all units on one battlefield is an action that an army would not want to choose if the other army is mixing uniformly at random).

Argue that each army mixing with equal probability between $(0,3,3)$, $(3,0,3)$ and $(3,3,0)$ is not an equilibrium.¹⁴

EXERCISE 5 *Divide and Choose.*

Two children must split a pie. They are gluttons and each prefers to eat as much of the pie as they can. The parent tells one child to cut the pie into two pieces and then allows the other child to choose which piece to eat. The first child can divide the pie into any multiple of tenths (for example, splitting it into pieces that are $1/10$ and $9/10$ of the pie, or $2/10$ and $8/10$, and so forth). Show that there is a unique backward induction solution to this game.

EXERCISE 6 *Information and Equilibrium.*

Each of two players receives an envelope containing money. The amount of money has been randomly selected to be between 1 and 1000 dollars (inclusive), with each dollar amount equally likely. The random amounts in the two envelopes are drawn independently. After looking in their own envelope, the players have a chance to trade envelopes. That is, they are simultaneously asked if they would like to trade. If they both say “yes,” then the envelopes are swapped and they each go home with the new envelope. If either player says “no,” then they each go home with their original envelope.

The actions in this game are actually a full list of whether a player says yes or no for each possible amount of money he or she is initially given. To simplify things, let us write down actions in the following more limited form: an action is simply a number between 0 and 1000, meaning that if they get an envelope with more than that number, then they say “no” and otherwise they say “yes”.

¹⁴Finding equilibria to Colonel Blotto games is notoriously difficult. One exists for this particular version, but finding it will take you some time.

So, for instance, if player 1 chooses action “3”, then she says “yes” to a trade when her initial envelope has 1 or 2 or 3 dollars, but says “no” if her envelope contains 4 or more dollars.

In a pure or mixed strategy equilibrium is it possible for both players to choose action “1000” with some positive probability?

Suppose that player 2 does not play action “1000”, can a best response of player 1 involve any positive probability on the action “1000”?

Repeat the above logic to argue that neither player will ever play “999” in an equilibrium. Iterating on this logic, what is the unique Nash equilibrium of this game?