Game Theory

Matthew Hoelle
www.matthew-hoelle.com

Spring 2014

Contents

Preface

This manuscript has been prepared for an advanced undergraduate course in game theory. This course is designed for students without any prior familiarity with game theory, or really any prior background in economics whatsoever. However, it is advised that students have the appropriate mathematical background, specifically univariate calculus and just a touch of multivariate calculus. To this end, a list of mathematical preliminaries is included at the end of the Preface.

The manuscript is divided into three parts.

Part I considers Static Games of Complete Information. Within Part I, Chapter 1 considers discrete choice games (pure strategies), Chapter 2 considers continuous choice games, and Chapter 3 considers the mixed strategies of discrete choice games.

Part II considers Dynamic Games of Complete Information. Within Part II, Chapter 4 considers games of perfect information, Chapter 5 considers games of imperfect information, and Chapter 6 considers infinitely repeated games.

Part III considers Games of Incomplete Information. Within Part III, Chapter 7 considers auctions, Chapter 8 considers the perfect Bayesian Nash equilibrium concept, Chapter 9 considers adverse selection (including signaling and screening), and Chapter 10 considers moral hazard.

This manuscript contains a large number of exercises that reinforce the material. The exercises vary in difficulty, and each has a solution provided in the Appendix. It is advised that students make an honest attempt to solve the exercises prior to looking at the solutions for guidance.

Throughout the chapters, I have included many classroom exercises that can be carried out in the classroom. In the text, I introduce and describe the exercise without providing the solution. The intent is that these games can be played in the classroom with prizes given to the students who "play the game the best". Solutions to these exercises can be provided upon request.

I would like to thank Somak Paul and the Fall 2011 class of Econ 451 at Purdue University. These students were active participants with version 1.0 of these notes and were instrumental in guiding the development of this manuscript into the finished product that you see before you right now.

Mathematical Preliminaries The following list provides the key mathematical results that are required to solve problems in this class.

1. **First derivatives**
   (a) $\frac{d}{dx}(c) = 0$ for any constant c.
   (b) $\frac{d}{dx}(cx) = c$.
   (c) $\frac{d}{dx}(cx^n) = cnx^{n-1}$ for $n \neq 0$.
   (d) $\frac{d}{dx}(c\ln(x)) = \frac{c}{x}$ for $x > 0$.
   (e) Chain Rule: If $f(x) = h(g(x))$, then $\frac{d}{dx}f(x) = \frac{d}{dy}h(y) \times \frac{d}{dx}g(x)$, where $y = g(x)$.
   (f) Product Rule: $\frac{d}{dx}(f(x) \times g(x)) = \frac{d}{dx}f(x) \times g(x) + f(x) \times \frac{d}{dx}g(x)$.
   (g) Quotient Rule: $\frac{d}{dx}\left(\frac{f(x)}{g(x)}\right) = \frac{\frac{d}{dx}f(x) \times g(x) - f(x) \times \frac{d}{dx}g(x)}{[g(x)]^2}$.

2. **Partial derivatives**
   (a) $\frac{\partial}{\partial x}f(x,y) = \frac{d}{dx}f(x;y)$, where y is held fixed and held constant. For example, $\frac{\partial}{\partial x}(x(y-x)) = y - 2x$.

3. **Property of the natural log**
   (a) $\ln(a) + \ln(b) = \ln(ab)$.
   (b) $\ln(a) - \ln(b) = \ln(\frac{a}{b})$.

4. **Summation notation**
   (a) $\sum_{i=1}^{n} x_i = x_1 + x_2 + ... + x_n$.

5. **Set notation**
   (a) The set $X = \{Larry, Curly, Moe\}$ is a set containing three elements.

(b) The set $X \backslash \{Moe\}$ equals the set X defined above, but excluding the element Moe. The set $X \backslash \{Moe\} = \{Larry, Curly\}$.

6. **Vector notation**

(a) If x_i is a one-dimensional variable and $i = 1, ..., I$, then $x = (x_1, ..., x_I)$ is an $I-$dimensional vector.

Part I

Static Games of Complete Information

Chapter 1

A Beautiful Game

1.1 Introduction to Game Theory

It is no coincidence that the scope for economic analysis has expanded in the past 60 years in unison with the growth and applicability of game theoretic reasoning. Where once 'economics' was adequately defined as the study of market prices and their effects on individual decision-making, economics as it stands today requires a broader definition. My favorite version defines 'economics' as the study of incentives in all social institutions.

Consider the interaction of individuals in the classical economic theory of markets. Individuals trade on competitive markets, meaning that they take prices as given and do not consider that either their own actions, or the actions of other individuals in the market, can result in different terms of trade. While this competitive interaction of individuals is still a good approximation for "large" markets, it is clear that many examples in life involve fewer participants. In this case, the participants have market power, so we want to be able to say something about the strategic interaction of individuals. Game theory provides the tools that allow us to predict outcomes in settings of strategic interaction.

The following are just a few real-life examples of strategic interaction:

- eBay auctions
- college admissions
- voting
- military standoffs

- marketing and advertising
- soccer penalty kicks
- marriage markets
- student movie discounts
- health insurance
- contract enforcement
- bargaining and negotiation

To understand game theory, we must first understand the meaning of a 'game.' A 'game' has a specific definition given by economists. It is important not to limit our thinking about a 'game' to its colloquial usage. While board games, card games, and games of chances are certainly examples of 'games,' they are but a small subset of the possible range of 'games' that can be addressed using the logic of game theory.

To appropriately define a 'game,' let's consider that you and your friend decide to pass the time by playing a board game. You open the closet in your dorm room and take down your favorite board game. The problem is that your friend has never played this particular game before. So you open the box and the first item that you should see is the piece of paper that details the rules of the game. A 'game' is defined according to its rules. These rules must be written so that your friend, who has never played the game before, can read them and completely understand how the game is played.

Let's break down the rules of a game into four components. These components are the four characteristics that we use to define a 'game.'

1. How many players can play the game? Depending on the context throughout this text, I will use the terms agents, households, decision-makers, and firms to describe the individuals that play the game.

2. What is the objective of players? In game theory, the objective of players is to maximize their payoff functions, where the payoff functions can be something as simple as "maximize the probability that you are the first player to successfully move all 4 of your pawns from Start to Home" as in the game of Sorry.

3. What are the players allowed to do? In game theory, each player has a set of strategies, which contains all possible strategies that the player can choose.
4. What are the effects of my decisions on other players? It is important to know how one player's action affects another player's payoff.

I have introduced some basic terminology in the above description of a game, so I want to take the time to distinguish between a 'strategy' and an 'action.'

Definition 1.1 *A strategy is a complete plan of actions for all possible contingencies.*

All possible contingencies takes into account the possible outcomes that can be reached, where an outcome is determined by the actions taken by all other players.

Definition 1.2 *An action is a single move at a single point of time.*

We are now prepared to study the first type of strategic interaction between individuals, one in which players make decisions simultaneously. Let's motivate the theory with the following example.

1.2 Motivating Example

Perhaps you have seen the movie A Beautiful Mind, which was released in 2001 and featured Ron Howard as director and Russell Crowe as the lead character John Nash. In one memorable scene of the movie (aptly called "the bar scene"), the character John Nash discovers that the best outcome for an individual is obtained not by maximizing its own payoff, but by maximizing its own payoff while accounting for the actions of the other players.

John Nash and his four colleagues (presumably the topologist John Milnor, the economist Harold Kuhn, the mathematician David Gale, and the economist Lloyd Shapley, all Princeton colleagues of Nash at the time) are sitting in a bar when five women enter. Among them is a beautiful blonde and four brunettes. The brunettes are of equal attractiveness, while the blonde is more beautiful than all of them. Nash theorizes that if all four of his friends approach the blonde, then they will block each other and no one will go home with her. If any of the men were to then approach a brunette, they would be turned down, because the brunettes do not want to be anyone's second choice. Thus, the best outcome (according to the film) is for each of the men to bypass their most preferred option (the blonde) and instead approach one of the brunettes.

We shall later see that the prediction described in the film is incorrect, using the theory of strategic interaction that was introduced by John Nash (and later coined Nash equilibrium). Nonetheless, the point that can be taken from this scene is that players need to make their decisions by taking into account the actions of other players in the game.

For our purposes, let's consider an equivalent and simpler game called the Room Selection Game. In this game, three female roommates are deciding how to allocate three bedrooms in a shared apartment. The rooms have different features and all roommates agree that Room 1 is the best, Room 2 is the second-best, and Room 3 is the worst. We suppose that the rooms are allocated according to the following game. All three roommates write down their rankings over the three rooms. These rankings are private, meaning that a player knows her own ranking, but not the rankings of the others. The rankings are then collected by an administrator.

The rankings are revealed by the administrator. If a player is the only roommate to list a particular room as her first choice, then she receives that room with probability 100%. If a player is one of two roommates to list a particular room as her first choice, then a coin is flipped and each has a 50% probability of receiving the room. If all three players list the same room as their first choice, then each has a 33.33% probability of receiving the room.

After the allocation of the rooms that are first choices, the remaining rooms are then allocated based upon the second and third choices of players. For example, suppose that two roommates submitted the rankings Room 1 - Room 2 - Room 3, while one roommate submitted the rankings Room 2 - Room 1 - Room 3. What happens to the roommate that loses the coin flip for Room 1? Well, that player cannot receive her second choice, Room 2, as this room has already been allocated. Thus, the roommate that loses the coin flip for Room 1 will receive Room 3.

If you were one of the roommates playing this game, what ranking would you write down?

More importantly, how do we predict that this game will be played? We predict that the game will be played using Nash equilibrium strategies. An equilibrium, as in chemistry, is a system at rest.

For this game, the system is the following hypothetical convergence process (the process is hypothetical, because all players have prior knowledge about the other players' payoffs and can accurately forecast how each of them will form strategies). Each of the players (in secret) writes down her ranking over the three rooms. All rankings are then placed on a table in the middle of the three players and revealed to all three. The administrator then selects a player at random. That player is given the opportunity, given what she now knows about the strategies of the other players, to change her strategy. If the ranking is changed, then the updated ranking is placed on the table for all players to see. The administrator then selects a different player at random and gives that player the same opportunity to change her ranking. The system is at rest (an equilibrium exists) when the administrator asks each of the three players in turn if she wants to change her ranking, and each one declines the offer.

What then is the equilibrium for this particular game? That analysis is delayed until after the game theory concepts have been introduced.

1.3 The Normal Form

The games considered in this chapter are static games. This simply means that all players make simultaneous decisions and then the game ends. What is important for the concept of simultaneous choice is not the timing, but rather the information possessed by the players. The players do not actually have to make their decisions simultaneously. Rather, each player only needs to make its decision without knowing what decisions have been made by the other players. Consider the Room Selection Game in which the roommates write down rankings of the rooms in secret. The roommates need not make these decisions at the exact same instant in time, but the rankings must be secret. This means that each player does not have knowledge of the actions of the other players when called upon to make its decision.

The normal form is a convenient way to depict each of the four characteristics of a game (players, payoffs, strategy sets, and effects on other players). The normal form is also commonly called the matrix form. Consider the following example (Gibbons, 1992, pgs. 5-8).

		Column Player		
		L	C	R
Row	U	2, 2	1, 4	4, 3
Player	D	1, 1	0, 2	3, 1

Figure 1.1

Figure 1.1 describes a game played between two players: Row Player and Column Player. The Row Player can either choose the strategy U or the strategy D (think Up and Down), while the Column Player can choose the strategy L, C, or R (think Left, Center, and Right). The payoffs for each of the players is depicted as elements in the 2×3 matrix. The convention is that the payoff for the Row Player is listed first, followed by a comma, followed by the payoff for the Column Player. For example, if Row Player chooses U and Column Player chooses C, then the outcome is (U,C) with a payoff of 1 for Row Player and 4 for Column Player.

As a second example, I will introduce the story behind possibly the most famous example in all of game theory: the Prisoners' Dilemma. Two male prisoners are arrested on suspicion of having jointly committed a large crime and a small crime. They are interrogated in separate rooms. The police only have evidence to convict each of the prisoners for committing the small crime. This is known by the prisoners.

Given this bind on the part of the police, they make the following proposal to Prisoner 1. If Prisoner 1 informs against Prisoner 2, then the police will have enough evidence to convict Prisoner 2 of committing the large crime. As a reward for cooperation, if Prisoner 1 informs against Prisoner 2, the police will then drop the small crime charge against Prisoner 1. The exact same proposal is made to Prisoner 2.

Suppose that a small crime conviction requires 1 year of prison time and a large crime conviction requires 3 years of prison time. The reward for cooperation is equal to a 1 year reduction in prison time (equivalent to dropping the small crime charges).

Notice that if both prisoners remain silent, then each of them will be convicted of the small crime and sentenced to 1 year in prison. If both prisoners inform against their counterpart, they are both convicted of the small and large crime (1+3=4 year prison sentence), but are rewarded with a 1 year prison reduction for having cooperated with the police.

This game is depicted in the normal form as follows (Gibbons, 1992, pgs. 2-3).

		Suspect 2	
		Mum	Fink
Suspect 1	Mum	-1, -1	-4, 0
	Fink	0, -4	-3, -3

Figure 1.2

The game is old (circa 1950's), so the different strategies are given old-fashioned names. 'Mum' means that a prisoner remains silent, while 'Fink' means that a prisoner informs against his counterpart. The payoffs listed in the 2×2 matrix reflect the story described above, where the absolute values of the payoffs are the lengths of the prison sentences and we suppose that less time in prison is better than more.

The third example of the normal form is given below (Gibbons, 1992, pgs. 5-8).

		Column Player L	C	R
Row	U	2, 2	1, 4	3, 3
Player	D	1, 3	0, 2	4, 1

Figure 1.3

The fourth example of the normal form is given below (Gibbons, 1992, pgs. 5-8).

		Column Player L	C	R
Row	U	2, 2	1, 4	3, 3
Player	D	1, 3	0, 1	4, 2

Figure 1.4

1.4 Dominant Strategies

We, as economists, want to be able to make predictions about how the games given in Figures 1.1, 1.2, 1.3, and 1.4 will be played. We begin with the first logical rule: all players will play their dominant strategy, if they have one. What is a dominant strategy? It is a strategy that provides the weakly highest payoff for a particular player, no matter what strategies are employed by the other players.

Mathematically, we define a dominant strategy using the following notation. For a game with I players, let $p_i(s)$ be the payoff to player i as a function of all players' strategies:

$$s = (s_1, s_2, ..., s_i, ..., s_I).$$

For simplicity, define s_{-i} as the strategies for all players, except player i :

$$s_{-i} = (s_1, s_2, ..., s_{i-1}, s_{i+1}, ..., s_I).$$

The strategy s_i is dominant for player i if for any s_{-i} and any alternative strategy $\tilde{s}_i$:

$$p_i(s_i, s_{-i}) \geq p_i(\tilde{s}_i, s_{-i}).$$

Let's consider some examples of dominant strategies. From Figure 1.1, let's focus only on the payoffs for the Row Player:

		Column Player		
		L	C	R
Row	U	**2**, ...	**1**, ...	**4**, ...
Player	D	1, ...	0, ...	3, ...

Figure 1.5

In doing so, we can see that the strategy U always provides a higher payoff than the strategy D, no matter what strategy is chosen by the Column Player (if Column Player chooses L, then U provides 2 and D only provides 1; if Column Player chooses C, then U provides 1 and D only provides 0; if Column Player chooses R, then U provides 4 and D only provides 3). Thus, U is a dominant strategy for the Row Player. Now, let's focus only on the payoffs for the Column Player:

		Column Player		
		L	C	R
Row	U	..., 2	..., **4**	..., 3
Player	D	..., 1	..., **2**	..., 1

Figure 1.6

As can be seen, the Column Player also has a dominant strategy, which is to play C.

Each of the two players in Figure 1.1 has a dominant strategy. Thus, we predict that the outcome of the game will have each of the players choosing this dominant strategy, so the outcome will be (U,C).

From Figure 1.2, let's focus only on the payoffs for Prisoner 1:

		Suspect 2	
		Mum	Fink
Suspect	Mum	-1, ...	-4, ...
1	Fink	**0**, ...	**-3**, ...

Figure 1.7

Prisoner 1 has a dominant strategy, which is Fink. This particular game is symmetric. A symmetric game is one in which the row player and the column player can be switched without any change to the payoff matrix. For this reason, if Fink is a dominant strategy for Prisoner 1, then Fink is a dominant strategy for Prisoner 2. If you don't believe me, consider the following matrix in which we only focus on the payoffs for Prisoner 2:

		Suspect 2	
		Mum	Fink
Suspect	Mum	..., -1	..., **0**
1	Fink	..., -4	..., **-3**

Figure 1.8

What about Figure 1.3? For this example, when we focus on the payoffs of the Row Player

		Column Player		
		L	C	R
Row	U	**2**, ...	**1**, ...	**3**, ...
Player	D	1, ...	0, ...	4, ...

Figure 1.9

we cannot find a dominant strategy. In this case, U is a better choice if the Column Player chooses L or C, but not if the Column Player chooses R. Similarly, if we focus on the payoffs of the Column Player

		Column Player		
		L	C	R
Row	U	..., **2**	..., *4*	..., 3
Player	D	..., **3**	..., *2*	..., 1

Figure 1.10

we also fail to find a dominant strategy. This suggests that we would like to have a more powerful tool to be able to make predictions not only in the games of Figures 1.1 and 1.2, but in the game in Figure 1.3 as well.

1.5 Iterated Deletion of Strictly Dominated Strategies

To introduce this new concept of Iterated Deletion of Strictly Dominated Strategies (abbreviated as IDSDS), we first define Common Knowledge. The assumption of Common Knowledge states that the Row Player knows the payoffs for the Column Player and can logically deduce its course of action. The Column Player knows that the Row Player has this information. The Row Player knows that the Column Player knows that the Row Player has this information. This higher order knowledge continues without end.

The opposite of a dominant strategy is a dominated strategy. A strategy is said to be strictly dominated if it provides strictly lower payoffs than another strategy, no matter the strategies of the other players. Mathematically, using the definitions from the previous section, the strategy s_i is strictly dominated by $\tilde{s}_i$ for player i if for any s_{-i} :

$$p_i(\tilde{s}_i, s_{-i}) > p_i(s_i, s_{-i}).$$

The logic of IDSDS begins with the realization that a player will never play a strictly dominated strategy. Consider the process of IDSDS as applied to Figure 1.3. Focusing on the payoffs for the Column Player from playing C and R,

		Column Player L	C	R
Row	U	..., xx	..., **4**	..., 3
Player	D	..., xx	..., **2**	..., 1

Figure 1.11

we see that the Column Player will never select R (it is strictly dominated by C). The Row Player knows this (Common Knowledge), so does not take into account the possible outcomes (U,R) and (D,R):

		Column Player L	C	R
Row	U	2, 2	1, 4	XXXX
Player	D	1, 3	0, 2	XXXX

Figure 1.12

With the game trimmed in the above manner, focus on the payoffs for the Row Player:

		Column Player L	C	R
Row	U	**2**, ...	**1**, ...	XXXX
Player	D	1, ...	0, ...	XXXX

Figure 1.13

The strategy D is strictly dominated and will not be played by the Row Player. The Column Player knows this (Common Knowledge) and does not take into account the possible outcomes (D, L) and (D, C):

		Column Player L	C	R
Row	U	2, 2	1, 4	XXXX
Player	D	XXXX	XXXX	XXXX

Figure 1.14

With the game logically updated for the Column Player as above, its strategy of L is strictly dominated by C. Thus, the logic of IDSDS dictates that outcome is (U, C).

Notice that the logic of IDSDS can be used to obtain a unique prediction in the games in both Figures 1.1 and 1.2. Thus, IDSDS is more general than dominant strategies as it allows for a prediction in all games that dominant strategies does, in addition to allowing for a prediction in the game of Figure 1.3. Is this the most general method that we can employ? Let's see if the method can be applied to Figure 1.4.

		Column Player L	C	R
Row	U	2, 2	1, 4	3, 3
Player	D	1, 3	0, 1	4, 2

Figure 1.4

You can stare at Figure 1.4 all day, but you won't be able to find a strategy that is strictly dominated for either player. This means that the IDSDS method cannot even begin. We seek a method that allows us to form predictions in the games in Figures 1.1, 1.2, 1.3, and 1.4.

1.6 Classroom Exercise 1: The Price is Right

Let's see how the logic of IDSDS applies to a Bizzaro "Price is Right" problem (Mas-Colell et al., 1995, pg. 263 and Dixit and Nalebuff, 2008, pgs. 315-316). Suppose that you are one of four contestants asked to bid on a Gizmo. You have no idea what the price of a Gizmo should be (as it is an imaginary object), but Drew Carey informs you that the price of a Gizmo is equal to 75% of the average bids of all four contestants. Drew then asks all contestants to write down a bid for a Gizmo (in secret) between \$1 and \$1,000. The contestant with the bid closest to the price will win (regardless if the bid is above or below the price).

You must write down a bid. The fact that the bid is made in secret means that you must write down your bid before observing what your three fellow contestants have bid. What is the bid that you will write down? If you think that this bid will win, what are you hoping about the deductive abilities of your fellow contestants?

1.7 Nash Equilibrium

The concept of a Nash equilibrium is credited to the groundbreaking papers of John Nash that are alluded to in the movie A Beautiful Mind. The concept of a Nash equilibrium follows the fundamental principle of a system at rest. Specifically, a Nash equilibrium is the set of strategies for all players such that each player is playing a best response to every other player's best response. Mathematically, using the notation introduced two sections prior, a Nash equilibrium is

$$s^* = (s_1^*, s_2^*, ..., s_i^*, ..., s_I^*)$$

with the property that for each player i :

$$p_i\left(s_i^*, s_{-i}^*\right) \geq p_i\left(s_i, s_{-i}^*\right)$$

for all other possible strategies s_i.

This equilibrium concept is more general than both dominant strategies and IDSDS. The following facts collect the logic.

1. If a player has a dominant strategy, any Nash equilibrium must contain this dominant strategy.

2. If the process of IDSDS results in a unique outcome (unique strategies for all players), then those strategies must be a Nash equilibrium. Further, this is the only Nash equilibrium.

The Nash equilibria of any two-player discrete choice game are found using the underline method. Here is how the underline method works for a 3×3 normal form game (Gibbons, 1994, pgs. 9-10).

Column Player

		L	C	R
	T	3, 3!	1, 5	3, π
Row Player	M	5, $\frac{5}{2}$	4, e	4, $\sqrt{8}$
	B	2, 1	5, 1	1, 0

Figure 1.15

Begin with the top row. Underline the largest payoff value for the Column Player.

Column Player

		L	C	R
	T	3, $\underline{3!}$	1, 5	3, π
Row Player	M	5, $\frac{5}{2}$	4, e	4, $\sqrt{8}$
	B	2, 1	5, 1	1, 0

Figure 1.16

Proceed to the second and third row. In each case, underline the largest payoff value for the Column Player.

Column Player

		L	C	R
	T	3, $\underline{3!}$	1, 5	3, π
Row Player	M	5, $\frac{5}{2}$	4, e	4, $\underline{\sqrt{8}}$
	B	2, $\underline{1}$	5, $\underline{1}$	1, 0

Figure 1.17

In the case of a tie, as in the third row in which 1 is the highest value and appears in two cells, then you need to underline each instance in which the highest value appears.

To help you avoid mistakes when using the underline method, just remember that the logic is backwards. When looking across a **row**, you are underlining the highest payoff for the **Column Player** (**row** –> **Column Player**).

The underline method continues by looking at the columns. For each column, underline the largest payoff value for the Row Player (**column** –> **Row Player**).

		Column Player		
		L	C	R
	T	3, $\underline{3!}$	1, 5	3, π
Row Player	M	$\underline{5}$, $\frac{5}{2}$	4, e	$\underline{4}$, $\underline{\sqrt{8}}$
	B	2, $\underline{1}$	$\underline{5}$, $\underline{1}$	1, 0

Figure 1.18

Any cell that contains two underlines is a Nash equilibrium. A Nash equilibrium is the set of strategies, while Nash equilibrium payoffs are the two payoff values. For Figure 1.15, the two Nash equilibria are (M,R) and (B,C), with payoffs of $\left(4, \sqrt{8}\right)$ and $(5, 1)$, respectively.

Why does the underline method work (and it always works to find all Nash equilibria in any two-player discrete choice game)? What we are underlining are actually the best responses for the players. As a Nash equilibrium is a best response to a best response, then two underlines yields a Nash equilibrium.

For the games in Figures 1.1, 1.3, and 1.4, the Nash equilibrium in all cases is (U,C), where the underline method for each is given below.

		Column Player		
		L	C	R
Row	U	$\underline{2}$,2	$\underline{1}$,$\underline{4}$	$\underline{4}$,3
Player	D	1,1	0,$\underline{2}$	3,1

Figure 1.1

		Column Player		
		L	C	R
Row	U	$\underline{2}$,2	$\underline{1}$,$\underline{4}$	3,3
Player	D	1,$\underline{3}$	0,2	$\underline{4}$,1

Figure 1.3

		Column Player		
		L	C	R
Row	U	$\underline{2}$,2	$\underline{1}$,$\underline{4}$	3,3
Player	D	1,$\underline{3}$	0,1	$\underline{4}$,2

Figure 1.4

The discrete choice games can be divided into four classes of games, where games within a particular class have similar predicted outcomes. The classes of games are (i) Collective Action Problems, (ii) Coordination Games, (iii) Anti-coordination Games, and (iv) Random Strategy Games.

1.7.1 Collective Action Problems

This class of games is characterized by the following property: the unique Nash equilibrium is not Pareto optimal. The leading example of a collective action problem is the Prisoners'

Dilemma introduced in Figure 1.2. Recalling the normal form game,

		Suspect 2 Mum	Fink
Suspect 1	Mum	-1, -1	-4, 0
	Fink	0, -4	-3, -3

Figure 1.2

the Nash equilibrium is (Fink, Fink). The equilibrium payoffs of $(-3, -3)$ are lower for both suspects compared to the outcome (Mum, Mum) with payoffs $(-1, -1)$.

A second example of a collective action problem is the Tragedy of the Commons (Baird et al., 1994, pg. 34). In this example, consider a public plot of land that two farmers may use for sheep grazing. If the resource is appropriately used, then grass is available for the sheep to eat, without a cost to the farmer. If the resource is overused, then not enough grass can be grown to feed all the sheep. Define B as the payoff to a farmer if he is the only one of the two farmers that allows its sheep to graze on the public land. Define A as the payoff to the farmer if neither of the two allow their sheep to graze on the public land. Define a as the payoff to the farmer if both allow their sheep to graze on the public land. Finally, define b as the payoff to the farmer if he does not allow his sheep to graze on the public land, while the other farmer does. The logic of the problem dictates that

$$B > A > a > b.$$

The farmer prefers to be the only one using the public land. Provided that both farmers take the same action, each would prefer to leave the land fallow and allow grass to be available next period, rather than consuming the entire resource in the current period.

In the normal form, the game is depicted as follows (Gibbons, 1992, pgs. 27-29).

		Farmer 2 Sheep graze	Sheep not graze
Farmer 1	Sheep graze	a, a	B, b
	Sheep not graze	b, B	A, A

Figure 1.19

Given the relation between payoffs, the underline method dictates that the unique Nash equilibrium is (Sheep graze, Sheep graze). The equilibrium payoffs of (a, a) are strictly lower for both players than the payoffs for the outcome (Sheep not graze, Sheep not graze). This is the tragedy, which occurs because the players are unable to commit to a collective action. Each player acts in his best interest, with the result being a suboptimal outcome.

1.7.2 Coordination Games

The two leading examples of coordination games are the Battle of the Sexes and the Stag Hunt. Here is the background for the Battle of the Sexes. A man leaves his office in a world without cell phones and e-mail and must decide to go either to the opera or to a boxing match. Meanwhile, a woman across town leaves her office at the same time and must decide to go either to the opera or to a boxing match. The man prefers boxing over opera, while the woman prefers opera over boxing, but the overarching concern of each is to be at the same event together.

Assigning reasonable payoffs to the two players, the normal form of this game can be depicted as follows (Gibbons, 1992, pgs. 11-12 and Baird et al., 1994, pgs. 41-42).

		Woman	
		Opera	Boxing
Man	Opera	40, 60	0, 0
	Boxing	0, 0	70, 30

Figure 1.20

Using the underline method, the two Nash equilibria of this game are (Opera, Opera) and (Boxing, Boxing).

Let's now consider the background for the Stag Hunt game. Hunter 1 wakes up in his castle in a world without cell phones and e-mail and must decide to either ride to the stag hunting grounds or the hare hunting grounds (a stag is a large deer). At the same time, Hunter 2 wakes up in a distant castle and must decide whether to ride to the same stag hunting grounds or to the same hare hunting grounds. The nice thing about hunting a hare is that it requires only one hunter to complete the task. Of course, a hare only has so much meat. The larger animal, the stag, is harder to hunt and requires the combined efforts of both hunters.

Assigning reasonable payoffs to the two players, the normal form of this game is given below (Baird et al., 1994, pgs. 35-36).

		Hunter 2: Stag	Hunter 2: Hare
Hunter 1	Stag	2, 2	0, 1
	Hare	1, 0	1, 1

Figure 1.21

Using the underline method, the two Nash equilibria of this game are (Stag, Stag) and (Hare, Hare).

This allows us to make the following two claims about the predicted play in Coordination Games.

1. There are multiple Nash equilibria.

2. In the case of Stag Hunt, some equilibria are "better" than others, while other equilibria provide a risk-free payoff. Notice that (Stag, Stag) provides strictly higher payoffs for both hunters compared to (Hare, Hare). However, choosing Hare provides a guaranteed safe payoff of 1 for a hunter. For this reason, the Stag Hunt game is often called the Assurance Game.

1.7.3 Classroom Exercise 2: Coordination

Your friend is returning from her study abroad in Italy. The two of you have agreed to meet as soon as she returns, but given the poor quality of Italian internet, you have been unable to reach your friend for the past several weeks. All you know is that your friend will be on campus on January 8. You must pick one location and one time as the window to meet your friend. If she picks the exact same location and the same time, then you two will meet up and both are happy. If the (time, location) pairs do not match, then you are unable to meet up.

Where would you choose to meet your friend? At what time? Suppose the stakes were raised. You are playing this game with a complete stranger, yet a student at Purdue. The administrator of the game asks you to choose one location and one time. If the other player (the stranger) chooses the exact same time and location, then the administrator provides each of you with \$100,000. If you and the stranger are unable to meet up, then both of you walk away empty handed. Do the higher stakes change your thinking about the game (or simply cause you to consider it more carefully)?

In this example of a coordination game (Mas-Colell et al., 1995, pgs. 221, 247-248), there are a large number of different strategies that you can choose. You can choose to be at the Rawls bistro at 8 am, at Rawls bistro at 9 am, and so forth; at the bookstore at 8 am, at the bookstore at 9 am, and so forth; at the bell tower at 8 am, at the bell tower at 9 am, and so forth. The payoffs for this game are 100,000 if both players choose the same strategy, and 0 otherwise. Using the underline method, we can see that for each (time, location) pair that we can imagine, a Nash equilibrium exists in which both players choose that particular pair.

This predicts a large number of Nash equilibria. Are some more likely than others? Let's use the concept of a Schelling focal point to address this question. Thomas Schelling suggested that some of the Nash equilibria are focal in the minds of the players. Focal simply means the human mind is quicker to concentrate on that particular outcome. The players are then more likely to play a strategy that is focal. What would you say is a focal time and location on Purdue's campus? A correct answer could net $100,000.

The concept of a focal point relates to the idea of a social norm. A social norm is behavior that has been adopted by society as "appropriate" without the use of laws to coordinate all individuals. Examples include walking on the right hand side of sidewalks in a U.S. city or clapping at the end of a theatre play (but not at the end of a movie or academic lecture). Consider the sidewalk example. The action that any one individual takes must be coordinated with the actions of all the other individuals in the city. Otherwise, if I walk on the left hand side and everyone else walks on the right hand side, then my travel is slowed as I am constantly running into people. Thus, it is clearly a Nash equilibrium for all individuals in a city to walk on the right hand side of sidewalks. The interesting thing is that this is not the only equilibrium. Another equilibrium would be for all individuals to walk on the right hand side on even-numbered avenues and on the left hand side on odd-numbered avenues. This is also an equilibrium as it prevents people from running into each other, but it is terribly difficult to remember. A social norm is an equilibrium that is easy to remember and apply. In essence, a social norm is a Schelling focal point (Mas-Colell et al., pgs. 247-248).

1.7.4 Anti-coordination Games

The example of an anti-coordination is the game of Chicken (Baird et al., 1994, pgs. 43-44). In this game, two drivers are driving cars toward each other at high rates of speed. Right before the instant of collision, each of the drivers can choose to either swerve to the left or go straight. If both choose to go straight, then they will crash, with this outcome resulting in the lowest possible payoffs for both drivers. If one swerves while the other goes straight, then the driver that has swerved is the 'chicken' and is mocked by his peers. The game can be depicted below (Dixit and Nalebuff, 2008, pgs. 118-119).

		Driver 2	
		Straight	Swerve
Driver 1	Straight	0, 0	10, 1
	Swerve	1, 10	5, 5

Figure 1.22

Using the underline method, there are two Nash equilibria of the Chicken game: (Straight, Swerve) and (Swerve, Straight).

1.7.5 Random Strategy Games

The examples of random strategy games are given by simple children's games. The first is called Matching Pennies. Each player has a penny. On the count of 3, each of the players must either show the penny heads side up or tails side up. Player 1 wins the game if the faces of the two pennies match: two heads face up or two tails. Player 2 wins the game if the faces of the pennies do not match. Assigning the value of 1 for a win and -1 for a loss, the game is depicted in the normal form below (Baird et al., 1994, pgs. 42-43).

		Player 2	
		Heads	Tails
Player 1	Heads	1, -1	-1, 1
	Tails	-1, 1	1, -1

Figure 1.23

Using the underline method, there are no Nash equilibria of this game.

Examine the logic for why a Nash equilibrium cannot exist. Suppose Player 1 plays Heads. Then the best thing that Player 2 can do is play Tails (no match). Well when Player 2 plays Tails, the best thing that Player 1 can do is play Tails (match). Thus, there never exists a system at rest in which both players are happy with their current strategy and its resulting payoff.

The second children's game is Rock Paper Scissors, which is played between two players. In this game, after the signal "Rock, Paper, Scissors, Shoot" each player forms their right hand into the shape of a rock (a fist), paper (a flat hand), or scissors (index and middle finger extended). Logic dictates that rock crushes scissors, scissors cuts paper, and paper covers rock. If both players throw the same shape, then the game ends in a tie. Assigning the value of 1 for a win, 0 for a tie, and -1 for a loss, the game is depicted below (Dixit and Nalebuff, 2008, pgs. 151-153).

		Player 2		
		Rock	Paper	Scissors
	Rock	0, 0	-1, 1	1, -1
Player 1	Paper	1, -1	0, 0	-1, 1
	Scissors	-1, 1	1, -1	0, 0

Figure 1.24

Using the underline method, there are no Nash equilibria of this game.

For complete disclosure, the Nash equilibria that we find in this chapter are pure-strategy Nash equilibria. Chapter 3 analyzes mixed-strategy Nash equilibria and variants of the games in Figures 1.23 and 1.24. A Nash equilibrium can be either a pure-strategy or a mixed-strategy Nash equilibrium. Theory dictates that a Nash equilibrium must exist. Thus, if a pure-strategy Nash equilibrium does not exist (as in the games of Figures 1.23 and 1.24), then a mixed-strategy Nash equilibrium must exist. Further details are postponed until Chapter 3.

1.8 Solution to the Motivating Example

Recall that three roommates are submitting rankings over three bedrooms (Room 1, Room 2, Room 3) in a shared apartment. The three rooms (Room 1, Room 2, Room 3) have values equal to $(a, b, c) = (30, 24, 6)$, respectively. Each of the roommates must decide which of the possible rankings to submit. The $3! = 6$ possible rankings are $\{123, 132, 213, 231, 312, 321\}$. It is not optimal for a roommate to list Room 3 anywhere but as her last choice. Consider that by listing Room 3 as either her first or second choice, a roommate is forfeiting the opportunity to enter a lottery with the chance to earn a strictly higher payoff. Thus, the only two possible rankings that will be chosen in equilibrium are $\{123, 213\}$.

There are four possible outcomes that can be attained for varying combinations of the strategies chosen by the three roommates.

1. If all three roommates choose 123, then each has a $\frac{1}{3}$ chance of winning its first choice (Room 1). Each has a $\frac{2}{3}$ chance of losing its first choice. Given that only two roommates remain for the two rooms, each would then have a $\frac{1}{2}$ chance of winning its second choice (Room 2). This means that the ex-ante chance of winning the second choice (Room 2) is $\frac{2}{3} \cdot \frac{1}{2} = \frac{1}{3}$. If a roommate loses both the lotteries for Room 1 and Room 2, then it is assigned the lone remaining room (Room 3). The chance of this outcome is $\frac{1}{3}$. Thus, the expected payoff for each roommate is $\frac{1}{3}(30) + \frac{1}{3}(24) + \frac{1}{3}(6) = 20$.

2. If two roommates choose 123 and one chooses 213, then the lone roommate that chooses 213 will receive Room 2 with certainty (and a payoff of 24). Each of the roommates that chooses 123 has a $\frac{1}{2}$ chance of winning Room 1 (its first choice) and a $\frac{1}{2}$ chance of winning Room 3 (the lone remaining choice). Thus, the expected payoff for each of the roommates that chooses 123 is $\frac{1}{2}(30) + \frac{1}{2}(6) = 18$.

3. If one roommate chooses 123 and two choose 213, then the lone roommate that chooses 123 will receive Room 1 with certainty (and a payoff of 30). Each of the roommates that chooses 213 has a $\frac{1}{2}$ chance of winning Room 2 (its first choice) and a $\frac{1}{2}$ chance of winning Room 3 (the lone remaining choice). Thus, the expected payoff for each of the roommates that chooses 213 is $\frac{1}{2}(24) + \frac{1}{2}(6) = 15$.

4. If all three roommates choose 213, then the same logic as Case 1 implies that the expected payoff of each is $\frac{1}{3}(30) + \frac{1}{3}(24) + \frac{1}{3}(6) = 20$.

Remember that a Nash equilibrium is the set of three strategies (rankings) such that each roommate is playing a best response to the other players' best responses. Is it then possible that $(123, 123, 123)$ is a Nash equilibrium (namely, all roommates choose the same ranking 123)? The payoff outcome is 20 for all three. Can one of the players (holding fixed the strategies of the other two) change its strategy and receive a higher expected payoff? Certainly. Consider that a roommate can receive a payoff value of 24 by using the ranking 213.

So if $(123, 123, 123)$ is not a Nash equilibrium, how can we find the Nash equilibria? The normal form game is a $2 \times 2 \times 2$ array. To depict this 3-dimensional array in 2-dimensional space, we consider the following two cross-sections. In the first, we hold fixed

Player 1's strategy at 123. In the second, we hold fixed Player 2's strategy at 213. The cells of the matrices contain three payoff values: the first belongs to roommate 1, the second to roommate 2, and the third to roommate 3.

1 plays 123		Roommate 3: 123	Roommate 3: 213
Roommate 2	123	20,20,20	18,18,24
	213	18,24,18	30,15,15

1 plays 213		Roommate 3: 123	Roommate 3: 213
Roommate 2	123	24,18,18	15,30,15
	213	15,15,30	20,20,20

Figure 1.25

Let's first find the Nash equilibria for both cross sections and then we can determine the overall Nash equilibria. Using the underline method, we obtain that the following are Nash equilibria when rooomate 1 plays 123 : $(123, 213)$ and $(213, 123)$. When roommate 1 plays 213, the Nash equilibrium is $(123, 123)$.

1 plays 123		Roommate 3: 123	Roommate 3: 213
Roommate 2	123	20,20,20	18,18,24
	213	18,24,18	30,15,15

1 plays 213		Roommate 3: 123	Roommate 3: 213
Roommate 2	123	24,18,18	15,30,15
	213	15,15,30	20,20,20

Figure 1.26

Let's now consider the Nash equilibria over all three roommates. When roommate 1 plays 123, then players 2 and 3 optimally play $(123, 213)$. Does roommate 1 have an incentive to deviate, that is, to change her ranking to 213? Holding fixed players 2 and 3 at $(123, 213)$, roommate 1 currently has a payoff of 18, while switching yields a payoff of 15 (no deviation). When roommate 1 plays 123, then players 2 and 3 also optimally play $(213, 123)$. Does roommate 1 have an incentive to deviate? Currently, roommate 1 has a payoff of 18, while switching yields a payoff of 15 (no deviation). When roommate 1 plays 213, then players 2 and 3 optimally play $(123, 123)$. Does roommate 1 have an incentive to deviate? Currently, roommate 1 has a payoff of 24, while switching yields a payoff of 20 (no deviation).

Thus, the Nash equilibria are $(123, 123, 213)$, $(123, 213, 123)$, and $(213, 123, 123)$. In each, two players have ranking 123, while the third has ranking 213. We'll return to this Room Selection Game when we discuss mixed-strategies in Chapter 3 and dynamic games in Chapter 4.

1.9 Application: Political Economy

The next game (Dixit and Nalebuff, 2008, pgs. 64-66, 283-285, 356-385) considers the interaction between two citizens and two politicians (Obama and Romney). Each citizen represents the political base of one of the two political parties (Democrat and Republican). We represent the political spectrum along the line from -1 to $+1$. The left corresponds to the political left (Democrat), while the right corresponds to the political right (Republican). All four of the players make simultaneous decisions.

The Democratic citizen prefers a platform at $-\frac{1}{2}$, with payoff function given by $-\left|-\frac{1}{2}-m\right|$, where m is the platform of the elected politician. If the elected politician has platform $m=-\frac{1}{2}$, then the Democratic citizen has the highest payoff value of 0. The Republican citizen prefers a platform at $+\frac{1}{2}$, with payoff function given by $-\left|\frac{1}{2}-m\right|$.

Obama can select a platform at either $-\frac{3}{8}$ or 0. Romney can select a platform at either 0 or $+\frac{3}{8}$ (for a fuller analysis of this game, the available platforms are $-m^*$ and 0 for Obama and 0 and m^* for Romney, where $\frac{1}{4}<m^*<\frac{1}{2}$). Each of the candidates cares only about winning and does not care about the platform of the winning candidate. The payoff from winning is 1, while the payoff from not winning is 0.

Each of the citizens can either vote or can choose to withhold its vote. The cost of voting is equal to $\frac{1}{4}$.

How are we able to find the Nash equilibria of this game? There are four players, so it is infeasible to write down the normal form in 4 dimensions. What we have to do is recognize that the process to find symmetric Nash equilibria can reduce the 4-player game to a 2-player game.

1.9.1 First equilibrium

One possible set of symmetric strategies is given below.

Democrat citizen	Not vote
Republican citizen	Not vote
Obama platform	0
Romney platform	0

Table 1.1

Let's hold fixed the strategies for (Democratic citizen, Obama) at (Note vote, 0). If we can show the optimal strategies for (Republican citizen, Romney) are (Not vote, 0), holding fixed the strategies of the other two, then the strategies in Table 1.1 are a Nash equilibrium.

Holding fixed the strategies for (Democratic citizen, Obama) at (Note vote, 0), the normal form for (Republican citizen, Romney) is given below.

		Romney	
		0	$\frac{3}{8}$
Republican citizen	Vote	$-\frac{3}{4}, \frac{1}{2}$	$-\frac{3}{8}, 1$
	Not vote	$-\frac{1}{2}, \frac{1}{2}$	$-\frac{5}{16}, \frac{1}{2}$

Figure 1.27

Let's validate the payoffs for each of the four possible outcomes. Suppose that the strategies are (Vote, 0). In this case, only one citizen votes. The citizen splits its vote between Obama and Romney, as both have the same platform. Thus, the payoff to Romney is $\frac{1}{2}$, as he has a 50% chance of winning. By voting, the Republican citizen has a payoff of $-\left|\frac{1}{2}-0\right|=-\frac{1}{2}$ from the difference in platform minus another $\frac{1}{4}$ as the cost of voting (for a total payoff of $-\frac{3}{4}$).

Suppose that the strategies are (Not Vote, 0). In this case, no one votes, meaning that the election is a tie. Each politician then has a 50% chance of winning for a payoff of $\frac{1}{2}$. By not voting, the Republican citizen has a payoff of $-\left|\frac{1}{2}-0\right|=-\frac{1}{2}$ from the difference in platform.

Suppose that the strategies are (Vote, $\frac{3}{8}$). In this case, only one citizen votes (the Republican). The citizen votes for Romney, as Romney's platform is closer to $\frac{1}{2}$, the best platform for the Republican citizen. The payoff for Romney is then 1, as Romney wins the election with certainty. By voting, the Republican citizen has a payoff of $-\left|\frac{1}{2}-\frac{3}{8}\right|=-\frac{1}{8}$ from the difference in platform minus another $\frac{1}{4}$ as the cost of voting (for a total payoff of $-\frac{3}{8}$).

Suppose that the strategies are (Not Vote, $\frac{3}{8}$). In this case, no one votes, meaning that the election is a tie. Each politician then has a 50% chance of winning for a payoff of $\frac{1}{2}$. By not voting, the Republican citizen has a payoff of $-(50\%)\left|\frac{1}{2}-0\right|-(50\%)\left|\frac{1}{2}-\frac{3}{8}\right|=-\frac{5}{16}$ from the expected difference in platform. Recall that if Obama wins (50% chance), the platform is 0, while if Romney wins (50% chance), the platform is $\frac{3}{8}$.

Let's use the underline method for Figure 1.27.

		Romney: 0	Romney: $\frac{3}{8}$
Republican citizen	Vote	$-\frac{3}{4}, \frac{1}{2}$	$-\frac{3}{8}, \underline{1}$
	Not vote	$\underline{-\frac{1}{2}}, \underline{\frac{1}{2}}$	$\underline{-\frac{5}{16}}, \underline{\frac{1}{2}}$

Figure 1.28

The two Nash equilibria of this reduced game are (Not vote, 0) and (Not vote, $\frac{3}{8}$). As (Not vote, 0) is a Nash equilibrium, given that we held fixed the strategies of the other two players at (Note vote, 0), then the strategies in Table 1.1 are in fact a Nash equilibrium.

1.9.2 Second equilibrium

A second possible set of symmetric strategies is given below.

Democrat citizen	Vote
Republican citizen	Vote
Obama platform	$-\frac{3}{8}$
Romney platform	$\frac{3}{8}$

Table 1.2

Let's hold fixed the strategies for (Democratic citizen, Obama) at (Vote, $-\frac{3}{8}$). If we can show the optimal strategies for (Republican citizen, Romney) are (Vote, $\frac{3}{8}$), holding fixed the strategies of the other two, then the strategies in Table 1.2 are a Nash equilibrium.

Holding fixed the strategies for (Democratic citizen, Obama) at (Vote, $-\frac{3}{8}$), the normal form for (Republican citizen, Romney) is given below.

		Romney: 0	Romney: $\frac{3}{8}$
Republican citizen	Vote	$-\frac{15}{16}, \frac{1}{2}$	$-\frac{3}{4}, \frac{1}{2}$
	Not vote	$-\frac{7}{8}, 0$	$-\frac{7}{8}, 0$

Figure 1.29

Let's validate the payoffs for each of the four possible outcomes. Suppose that the strategies are (Vote, 0). In this case, both citizens vote. The Democratic citizen votes for

Obama, while the Republican citizen votes for Romney. Thus, the payoff to Romney is $\frac{1}{2}$, as he has a 50% chance of winning. By voting, the Republican citizen has a payoff of $-(50\%)\left|\frac{1}{2}-\left(-\frac{3}{8}\right)\right|-(50\%)\left|\frac{1}{2}-0\right|=-\frac{11}{16}$ from the expected difference in platform minus another $\frac{1}{4}$ as the cost of voting (for a total payoff of $-\frac{15}{16}$). Recall that if Obama wins (50% chance), the platform is $-\frac{3}{8}$, while if Romney wins (50% chance), the platform is 0.

Suppose that the strategies are (Not Vote, 0). In this case, only the Democratic citizen votes. This means that Obama wins the election with certainty. The payoff to Romney is then 0. By not voting, the Republican citizen has a payoff of $-\left|\frac{1}{2}-\left(-\frac{3}{8}\right)\right|=-\frac{7}{8}$ from the difference in platform.

Suppose that the strategies are (Vote, $\frac{3}{8}$). In this case, both citizens vote. The Democratic citizen votes for Obama, while the Republican citizen votes for Romney. Thus, the payoff to Romney is $\frac{1}{2}$, as he has a 50% chance of winning. By voting, the Republican citizen has a payoff of $-(50\%)\left|\frac{1}{2}-\left(-\frac{3}{8}\right)\right|-(50\%)\left|\frac{1}{2}-\frac{3}{8}\right|=-\frac{1}{2}$ from the expected difference in platform minus another $\frac{1}{4}$ as the cost of voting (for a total payoff of $-\frac{3}{4}$).

Suppose that the strategies are (Not Vote, $\frac{3}{8}$). In this case, only the Democratic citizen votes. This means that Obama wins the election with certainty. The payoff to Romney is then 0. By not voting, the Republican citizen has a payoff of $-\left|\frac{1}{2}-\left(-\frac{3}{8}\right)\right|=-\frac{7}{8}$ from the difference in platform.

Let's use the underline method for Figure 1.29.

		Romney	
		0	$\frac{3}{8}$
Republican	Vote	$-\frac{15}{16}, \underline{\frac{1}{2}}$	$\underline{-\frac{3}{4}}, \underline{\frac{1}{2}}$
citizen	Not vote	$\underline{-\frac{7}{8}}, \underline{0}$	$-\frac{7}{8}, \underline{0}$

Figure 1.30

The two Nash equilibria of this reduced game are (Not vote, 0) and (Vote, $\frac{3}{8}$). As (Vote, $\frac{3}{8}$) is a Nash equilibrium, given that we held fixed the strategies of the other two players at (Vote, $-\frac{3}{8}$), then the strategies in Table 1.2 are a second Nash equilibrium.

1.10 Exercises

1. Find the Nash equilibria of the following normal form game (Gibbons, 1992, pgs. 9-10).

		Player 2 L	C	R
Player 1	T	3, 2	1, 6	3, 2
	M	5, 1	4, 3	3, 2
	B	2, 4	5, 4	1, 4

Figure 1.31

2. Find the Nash equilibria of the following normal form game (Gibbons, 1992, pgs. 9-10).

		Player 2 Save	Use	Delay
Player 1	Save	5, 5	-2, 10	0, 0
	Use	3, 2	1, -1	0, 0
	Delay	0, 0	0, 0	0, 0

Figure 1.32

3. Write the following games in the normal form and solve for the Nash equilibria. In some cases, multiple Nash equilibria may exist or a Nash equilibrium may not exist at all. (Note: There is some freedom in how you assign payoff values for the players. These choices may affect the Nash equilibria.)

   (a) Two animals are fighting over some prey (this game is typically called a Hawk-Dove game). Each can be passive or aggresive. Each prefers to be aggresive if its opponent is passive, and passive if its opponent is aggresive. Given its own stance, it prefers the outcome in which its opponent is passive to that in which its opponent is aggresive. (Dixit and Nalebuff, 2008, pgs. 97-101 and Mas-Colell et al., 1995, pg. 265).

   (b) Two students wish to attend the same university. The students each receive strictly positive payoff if they attend the same university and zero payoff otherwise. The list of possible universities is Purdue University, Indiana University, and Notre Dame University. Student A prefers Purdue to Indiana and Indiana to

Notre Dame (by transitivity, he/she also prefers Purdue to Notre Dame). Student B has a scholarship at Notre Dame, so prefers Notre Dame to Purdue and Purdue to Indiana (by transitivity, he/she prefers Notre Dame to Indiana).

(c) Consider a soccer penalty kick between a Scorer and a Goalie. The Scorer is stronger kicking to his/her Right than to his/her Left. Given the 12 yards between the ball and the goal, the Goalie must pre-determine which way he/she will dive. Thus, the actions are chosen simultaneously: a Scorer can shoot either Left or Right and the Goalie can dive either Left or Right. If the Goalie guesses wrong, the Scorer always scores. If the Goalie guesses correctly with the Scorer kicking Right, then 50% of the shots are stopped. If the Goalie guesses correctly with the Scorer kicking Left, then 90% of the shots are stopped. (Dixit and Nalebuff, 2008, pgs. 143-151).

4. In the Room Assignment Game, suppose that the payoffs for the rooms are now $a = 20$, $b = 12$, and $c = 8$. What is the Nash equilibrium of the game?

5. The town of Lakesville is located next to the town of West Lakesville. The Chinese restaurant Lin's is located in Lakesville and the Chinese restaurant Wong's is located in West Lakesville. Each restaurant currently delivers take-out orders within its town only. Both restaurants are simultaneously deciding whether or not to expand their delivery service to the neighboring town (Lin's to offer delivery to West Lakesville and Wong's to Lakesville).

   A restaurant will earn $25 for selling take-out food in its own town and $15 for selling take-out food in the other town (a $10 travel cost is already included in the specified earnings). If a restaurant decides to expand its delivery service, a fixed cost of $10 must be paid (to hire an additional driver).

   If both expand their delivery service, they both maintain their current customers in their own town. If one expands and the other does not, then the one that expands sells to all consumers in both towns. The game is depicted below.

		Wong's	
		Expand	Not expand
Lin's	Expand	15, 15	30, 0
	Not expand	0, 30	25, 25

Figure 1.33

Find the Nash equilibria of this game.

6. In the previous problem, we assumed that consumers would make certain choices without actually including them in the game. Let's correct this. Suppose each town only has one consumer of Chinese take-out food. Each consumer must select a Chinese restaurant. Exactly as in the previous problem, the restaurants must decide whether or not to expand delivery service. All decisions (both those of the restaurants and those of the consumers) are made simultaneously.

   If a consumer selects a restaurant that does deliver take-out food to that consumer, then the consumer has a payoff of 1. If a consumer selects a restaurant that does not deliver take-out food to that consumer, then no sale is made and the consumer has a payoff of 0.

   The payoffs for the restaurants are determined from the second paragraph of the previous exercise.

   One of the following two sets of strategies is a Nash equilibrium. Which one is a Nash equilibrium? Justify your answer with sound reasoning.

Option 1

Lin's Restaurant	Expand
Wong's Restaurant	Expand
Lakesville consumer	Lin's
W. Lakesville consumer	Wong's

Option 2

Lin's Restaurant	Expand
Wong's Restaurant	Expand
Lakesville consumer	Wong's
W. Lakesville consumer	Lin's

Bibliography

[1] Baird, Douglas, Robert Gertner, and Randal Picker. *Game Theory and the Law*. Cambridge, MA, and London, England: Harvard University Press, 1994.

[2] Dixit, Avinash, and Barry Nalebuff. *The Art of Strategy: A Game Theorist's Guide to Success in Business and Life*. 1st ed. New York, NY: W.W. Norton & Company, Inc., 2008.

[3] Gibbons, Robert. *Game Theory for Applied Economists*. Princeton, NJ: Princeton University Press, 1992.

[4] Mas-Colell, Andrew, Michael Whinston, and Jerry Green. *Microeconomic Theory*. New York, NY: Oxford University Press, Inc., 1995.

[5] Osborne, Martin J. *An Introduction to Game Theory: International Edition*. New York, NY: Oxford University Press, 2009.

Chapter 2

Lunch Trucks at the Beach

2.1 Motivating Example

When you go to the beach, you probably enjoy the convenience of being able to obtain food and beverage at a lunch truck conveniently parked on the beach. Have you ever wondered how these lunch trucks choose where they set up shop for the day? Let's consider the strategic decisions of two lunch trucks in this Beach Game (Mas-Colell et al., 1995, pgs. 263-264).

The beach is one mile long. Evenly distributed on the beach are consumers. We assume that the consumers located immediately in front of a lunch truck will visit the lunch truck with probability 1 (for a total demand of 1). The consumers located further away from the lunch truck will visit with probability less than 1.

Vendor 1 is initially located at mile 0 of the beach, but can costlessly move to any location on the one-mile strip. Vendor 2 is initially located at mile 1 of the beach, but can costlessly move to any location.

The two players in this game are the two vendors. Each must choose its location, without knowing the location chosen by the other vendor. Suppose that the equilibrium location for vendor 1 is n_1 and the equilibrium location for vendor 2 is n_2. If so, then the consumer located at mile z on the beach goes to vendor 1 if $|n_1 - z| < |n_2 - z|$ and goes to vendor 2 if $|n_2 - z| < |n_1 - z|$ (with indifference between the two in the case of an equality). If the consumer located at mile z goes to vendor 1, then the demand equals $1 - |n_1 - z|$. If the consumer located at mile z goes to vendor 2, then the demand equals $1 - |n_2 - z|$.

Where will the lunch trucks locate? As an interesting extension (considered in Exercise 5), suppose that the lunch trucks now have to pay a cost (a fuel cost) to relocate. Where would the lunch trucks locate in this scenario?

2.2 Hotelling Model of Spatial Competition

The Hotelling model originally considered the choice by firms (stores) about where to locate, taking into account the location decisions of their competitors. The model was then adapted by Downs to analyze political competition. We will analyze this second application (Gibbons, 1992, pg. 35).

Suppose that two candidates are facing off in an election. The candidates are fickle and don't have any true beliefs. They can choose to adopt any platform along the political spectrum. The political spectrum is the one-dimensional line from -1 to $+1$, where platforms containing philosophies of the political left are located near -1 (on the left) and platforms containing philosophies of the political right are located near $+1$ (on the right). The candidates simultaneously choose the platform that they wish to adopt. We can assume that neither of the candidates is an incumbent, so neither has a fixed platform. Let's denote p_1 as the platform for candidate 1 and p_2 as the platform for candidate 2. Each candidate must take into account its oppponent's choice when making its own.

What are the effects of choosing a platform? Well, a different platform leads to a different number of citizens voting for a candidate. The citizens do have true beliefs. Each has preferences such that they want to minimize the distance between the platform of the winning politician and their preference point. The preference points differ across all citizens. An example is given in Figure 2.1 in the companion 'Figures' document. Figure 2.1 displays an economy with 5 citizens named Howard, King, Jackson, Webber, and Rose. The flags indicate the location of the preferences points for each. Roughly speaking, Howard is the most "liberal" and Rose is the most "conservative."

We assume that all citizens will vote in the election due to a strong sense of civic duty. Given this economy and the distribution of preference points, what platforms will the two candidates choose? What platform would you choose if you were one of the candidates?

We find the Nash equilibrium (a best response to a best response) using the following logical deduction. Suppose candidate 1 initially chooses its platform between Howard and King, while candidate 2 initially chooses its platform between Webber and Rose. Given this allotment of platforms, candidate 1 will receive votes from Howard, King, and Jackson and candidate 2 will receive votes from Webber and Rose. We have assumed that Jackson is closer to candidate 1 than candidate 2. Candidate 1 has three votes and wins the election.

Is this the best strategy for candidate 2, or can candidate 2 change its platform (holding fixed the platform of candidate 1) and receive a higher number of votes? Candidate 2 can move its platform to somewhere between Jackson and Webber, so that now Jackson will choose to vote for candidate 2. With Jackson's new vote, candidate 2 now has three votes and wins the election.

Is this now the best strategy for candidate 1 (given this updated strategy of candidate 2)? Of course not. Candidate 1 will move its platform to the right (holding fixed the platform of candidate 2) until it is once again the closest candidate to Jackson. In this way, candidate 1 can regain the crucial "swing vote" of Jackson and win the election.

Candidate 2 then has an incentive to undercut candidate 1 again and move closer to the position of the swing voter Jackson. We can see that this process of undercutting will continue until both candidates have their platforms located at exactly the preference point of Jackson. In this economy, Jackson is the median voter, meaning that the same number of citizens are located to the left (politically) and to the right (politically).

Is it then a Nash equilibrium of this game for both candidates to choose a platform at the preference point of the median voter, denoted $(p_1, p_2) = Jackson$? If it is a Nash equilibrium, then the platform choice for each candidate is a best response to the platform choice of the other candidate. Equivalently, this means that if candidate 1 is holding p_2 fixed, then he would have no incentive to deviate (change its platform) from p_1 (and likewise for candidate 2). Currently, candidate 1 is receiving 2.5 votes as Jackson is indifferent between the two candidates and flips a coin to decide his vote. If candidate 1 moves p_1 to the left, then he only receives two votes (Howard and King), while if he moves p_1 to the right, he only receives two votes (Webber and Rose). So candidate 1 has no incentive to deviate. The exact same logic works for candidate 2. As neither has an incentive to deviate, then we know that $(p_1, p_2) = Jackson$ is the unique Nash equilibrium of this game.

The take-away message is that for any economy with an odd number of voters and two candidates, the unique Nash equilibrium of the political competition game specifies that both candidates choose platforms located at the preference point of the median voter. What happens if three candidates are in the election? See Exercise 1.

2.3 Cournot Duopoly

The Cournot model of firm competition works as follows (Gibbons, 1992, pgs. 14-21, 49 and Mas-Colell et al., 1995, pgs. 387, 389-393). The inverse demand function for a particular market is equal to $P(Q) = a - bQ$, where P is the price of a good and Q is the total quantity of the good that is produced by all firms. The Cournot model typically considers a duopoly (two firms) in which both firms simultaneously choose a quantity to produce. Denoting the quantity q_1 as the amount produced by firm 1 and q_2 as the amount produced by firm 2, then the total quantity $Q = q_1 + q_2$.

The firm production choices are strategic substitutes. That is, an increase in the amount produced by firm 1 (holding fixed the amount produced by firm 2) will lower the market price and result in lower profits for firm 2. Thus, each firm must take into account the production decisions of the other firm when making its own production decision.

Firm 1 has a marginal cost of production equal to c_1 and firm 2 has a marginal cost of production equal to c_2. The Nash equilibrium is found using the two step process:

1. Find the best response functions for both firms $q_1 = b_1(q_2)$ and $q_2 = b_2(q_1)$.
2. Solve the two equations to find (q_1^*, q_2^*) such that $q_1^* = b_1(q_2^*)$ and $q_2^* = b_2(q_1^*)$.

The profit function for firm 1 (profit = revenue - cost) is given by:

$$\pi_1(q_1, q_2) = [a - b(q_1 + q_2)]\, q_1 - c_1 q_1.$$

The first order conditions are:

$$\frac{\partial \pi_1(q_1, q_2)}{\partial q_1} = a - 2bq_1 - bq_2 - c_1 = 0,$$

which can be solved for q_1 to yield the best response $q_1 = b_1(q_2) = \frac{a - bq_2 - c_1}{2b}$.

Similarly, the profit function and first order conditions for firm 2 are given by:

$$\begin{aligned} \pi_2(q_1, q_2) &= [a - b(q_1 + q_2)]\, q_2 - c_2 q_2. \\ \frac{\partial \pi_2(q_1, q_2)}{\partial q_2} &= a - 2bq_2 - bq_1 - c_2 = 0. \end{aligned}$$

This allows us to solve for the best response $q_2 = b_2(q_1) = \frac{a - bq_1 - c_2}{2b}$.

Solving the two equations $q_1^* = b_1(q_2^*)$ and $q_2^* = b_2(q_1^*)$ results in the following Nash equilibrium (as always, a Nash equilibrium is a best response to a best response):

$$\begin{aligned} q_1^* &= \frac{a + c_2 - 2c_1}{3b}. \\ q_2^* &= \frac{a + c_1 - 2c_2}{3b}. \end{aligned}$$

Notice that if $c_1 = c_2$, then the production decisions $q_1^* = q_2^*$.

Using this last fact, let's analyze a Cournot oligopoly with n firms, all of which have an identical marginal cost of production equal to c. The profit function for firm 1 is given by:

$$\pi_1(q_1, q_2, ..., q_n) = \left[a - b\sum_{i=1}^{n} q_i\right] q_1 - cq_1.$$

The first order conditions are:

$$\frac{\partial \pi_1(q_1, q_2, ..., q_n)}{\partial q_1} = a - 2bq_1 - b\sum_{i=2}^{n} q_i - c = 0,$$

which can be solved for q_1 to yield the best response $q_1 = b_1(q_2, ..., q_n) = \frac{a - b\left(\sum_{i=2}^{n} q_i\right) - c}{2b}$.

As all firms have the same production cost, then $q_1^* = q_2^* = ... = q_n^*$. This fact can be used to solve for the Nash equilibrium. Define the production choice of all firms simply as q^*. From the best response for firm 1:

$$q^* = b_1(q^*, ..., q^*) = \frac{a - (n-1)bq^* - c}{2b}.$$

Consequently,

$$q^* = \frac{a-c}{(n+1)b}.$$

The market price for n firms is equal to:

$$\begin{aligned} P(Q) &= a - bQ = a - nb\left(\frac{a-c}{(n+1)b}\right) \\ &= \frac{a}{n+1} + \frac{n}{n+1}c. \end{aligned}$$

The profit received by each firm is

$$\begin{aligned} \pi &= (P(Q) - c)\,q^* = \left(\frac{a}{n+1} + \frac{n}{n+1}c - c\right)\left(\frac{a-c}{(n+1)b}\right) \\ &= \left(\frac{a-c}{n+1}\right)\left(\frac{a-c}{(n+1)b}\right) = \frac{1}{b}\left(\frac{a-c}{n+1}\right)^2. \end{aligned}$$

Let's see what the model predicts when the economy only contains one firm ($n = 1$). In this case, $q^* = \frac{a-c}{2b}$. If we recall from our intermediate micro theory course (just pretend), we know that a monopolist chooses prices such that $MR = MC$. If the demand function is linear (as it is here), the marginal revenue curve has the same intercept as the inverse demand curve, but with slope twice as large (in absolute value). Thus, $MR(Q) = a - 2bQ$. The marginal cost is c. Setting marginal revenue equal to marginal cost yields:

$$\begin{aligned} a - 2bQ &= c \\ Q &= \frac{a-c}{2b}. \end{aligned}$$

This is the same prediction obtained from the Cournot model.

Let's see what the model predicts when $n \to \infty$ (an approximation of perfect competition). In this case, $P(Q) = c$. In a model with perfect competition, firms set price equal to marginal cost, which is exactly what the Cournot model has just predicted.

2.4 Bertrand Duopoly

The Bertrand model of firm competition (Mas-Colell et al., 1995, 387-389) has a different philosophy than the Cournot model. Rather than a competition in terms of quantity choices, the firms compete in terms of price choices. The consumers in the economy will purchase the good from the firm with the lowest price (with the purchases exactly split if both firms set the same price). We consider only Bertrand duopolies in which two firms simultaneously select a price, where firm 1 has marginal cost of production equal to c_1 and firm 2 has marginal cost of production equal to c_2.

If $c_1 = c_2$, then the best responses for both firms are identical, with the best response for firm 1 given by:

$$p_1 = b_1(p_2) = \left\{ \begin{matrix} p_2 - \epsilon & \text{if } p_2 > c_1 \\ c_1 & \text{if } p_2 \leq c_1 \end{matrix} \right\},$$

where $\epsilon > 0$ small. This function indicates that each firm wishes to undercut the price of the other firm, but never finds it optimal to set a price below its marginal cost (as this results in negative profit). Thus, the unique Nash equilibrium is $p_1^* = p_2^* = c_1 = c_2$. Both firms, by pricing at their marginal costs, receive zero profit.

If $c_1 < c_2$, then firm 2 will not be able to sell any output as firm 1 will simply set a price equal to $c_2 - \epsilon$, for $\epsilon > 0$ small.

This basic Bertrand model seems limited. The predictions of the model are either (i) both firms earn zero profit or (ii) one firm does not sell anything and goes out of business. To obtain a more realistic set of predictions, we incorporate transportation costs into the model. Suppose that consumers are evenly distributed on the real line from 0 to 1. We can view this as a town that is organized along a 1-mile line. At the left end of town (mile 0), firm 1 operates a store. At the right end of town (mile 1), firm 2 operates a store. Each consumer incurs a cost of 1 per mile traveled. The consumers will travel to the store that provides the lowest total cost, where the total cost is equal to price of the good plus the transportation cost.

As with all Bertrand models, the two firms compete by simultaneously choosing prices p_1 and p_2. A consumer living at location t along the 1-mile strip goes to store 1 if $p_1 + t < p_2 + (1 - t)$ and purchases one good (demand = 1). The same consumer goes to store 2 if $p_1 + t > p_2 + (1 - t)$ (again, demand = 1). If $p_1 + t^* = p_2 + (1 - t^*)$, then the consumer at location t^* is indifferent between the two stores and splits its demand. In particular, if $p_1 + t^* = p_2 + (1 - t^*)$, then all consumers to the left of t^* go to store 1 and all consumers to the right of t^* go to store 2. This is illustrated in Figure 2.2 in the companion 'Figures' document for a given choice of (p_1, p_2). The value of t^* is given by:

$$\begin{aligned} p_1 + t^* &= p_2 + (1 - t^*) \\ 2t^* &= p_2 - p_1 + 1 \\ t^* &= \left(\frac{p_2 - p_1 + 1}{2}\right). \end{aligned}$$

Assume that the consumers along the 1-mile strip have mass equal to 1 (and were previously assumed to be evenly distributed). Thus, store 1 sells to consumers in the left t^* mile of the town, which is equal to t^* consumers. Likewise, store 2 sells to consumers in the right $(1 - t^*)$ mile, which is equal to $(1 - t^*)$ consumers.

The payoff function for store 1 is equal to

$$\pi_1(p_1, p_2) = (p_1 - c_1)\, t^* = (p_1 - c_1)\left(\frac{p_2 - p_1 + 1}{2}\right),$$

after using the definition of t^* from Figure 2.2. The first order condition for store 1 is given by:

$$\frac{\partial \pi_1(p_1, p_2)}{\partial p_1} = \left(\frac{p_2 - p_1 + 1}{2}\right) - \left(\frac{p_1 - c_1}{2}\right) = 0.$$

Solving for p_1, the best response for firm 1 is then given by $p_1 = b_1(p_2) = \left(\frac{p_2 + 1 + c_1}{2}\right)$.

The payoff function and the first order condition for firm 2 are given by:

$$\begin{aligned}\pi_2(p_1,p_2) &= (p_2-c_2)(1-t^*) = (p_2-c_2)\left(\frac{p_1-p_2+1}{2}\right).\\ \frac{\partial\pi_2(p_1,p_2)}{\partial p_2} &= \left(\frac{p_1-p_2+1}{2}\right)-\left(\frac{p_2-c_2}{2}\right)=0.\end{aligned}$$

The best response for firm 2 is $p_2 = b_2(p_1) = \left(\frac{p_1+1+c_2}{2}\right)$.

The Nash equilibrium is (p_1^*, p_2^*) such that $p_1^* = b_1(p_2^*)$ and $p_1^* = b_2(p_2^*)$. Solving these two equations (you can check my algebra) yields:

$$\begin{aligned}p_1^* &= 1+\frac{1}{3}c_2+\frac{2}{3}c_1.\\ p_2^* &= 1+\frac{1}{3}c_1+\frac{2}{3}c_2.\end{aligned}$$

For the case in which $c_1 = c_2$, then both firms sell the good and do so at a price strictly greater than marginal cost (which allows for strictly positive profit). This shows that adding in a logical element like transportation costs allows for sensible predictions from the Bertrand model.

2.5 Solution to the Motivating Example

Recall the Beach Game from the beginning of the chapter. Let's initially suppose that n_1, the equilibrium location of vendor 1, and n_2, the equilibrium location of vendor 2, are such that $n_1 \leq n_2$. If this turns out not to be true, then we would need to update our solution.

The payoff for vendor 1 will be a function not just of its own choice, n_1, but also of the choice of the other vendor, n_2. Let's define this payoff function as $p_1(n_1, n_2)$. Holding n_2 fixed, let's write an expression for the solution to vendor 1's optimization problem $\max_{0\leq n_1\leq 1} p_1(n_1,n_2)$. The solution to the problem is called the best response $b_1(n_2) = \arg\max_{0\leq n_1\leq 1} p_1(n_1,n_2)$.

Likewise, the payoff for vendor 2 is a function of both locations, so we define $p_2(n_1, n_2)$. The solution to the problem is called the best response $b_2(n_1) = \arg\max_{0\leq n_2\leq 1} p_2(n_1,n_2)$. As a Nash equilibrium is a best response to a best response, then the Nash equilibrium (n_1^*, n_2^*)

must solve the following two equations:

$$\begin{aligned} n_1^* &= b_1(n_2^*). \\ n_2^* &= b_2(n_1^*). \end{aligned}$$

The difficulty in the problem is then to define $p_1(n_1, n_2)$ and $p_2(n_1, n_2)$. The payoff for any vendor is equal to the demand of the consumers that visit that particular truck (I have not mentioned any marginal or fixed costs for the trucks, so we assume that they don't have either). The total demand is equal to the area under the demand curve. Consider Figure 2.3 in the companion 'Figures' document.

The payoff functions are $p_1(n_1, n_2) = area(A) + area(B)$ and $p_2(n_1, n_2) = area(C) + area(D)$. Using the geometric rule that the area of a quadrilateral is equal to average height times average base, the areas of the four quadrilaterals in Figure 2.3 are given as follows:

$$\begin{aligned} area(A) &= \frac{1}{2}\left[(1-n_1)+1\right](n_1). \\ area(B) &= \frac{1}{2}\left[\left(1-\left(\frac{n_2-n_1}{2}\right)\right)+1\right]\left(\frac{n_2-n_1}{2}\right). \\ area(C) &= \frac{1}{2}\left[\left(1-\left(\frac{n_2-n_1}{2}\right)\right)+1\right]\left(\frac{n_2-n_1}{2}\right) = area(B). \\ area(D) &= \frac{1}{2}(n_2+1)(1-n_2). \end{aligned}$$

The areas can be equivalently written as

$$\begin{aligned} area(A) &= n_1 - \frac{1}{2}(n_1)^2. \\ area(B) &= area(C) = \left(\frac{n_2-n_1}{2}\right) - \frac{1}{2}\left(\frac{n_2-n_1}{2}\right)^2. \\ area(D) &= \frac{1}{2} - \frac{1}{2}(n_2)^2. \end{aligned}$$

Let's take the first order condition of $p_1(n_1, n_2)$ (the partial derivative of $p_1(n_1, n_2)$ with respect to n_1):

$$\frac{\partial p_1(n_1, n_2)}{\partial n_1} = 1 - n_1 - \frac{1}{2} + \frac{1}{2}\left(\frac{n_2-n_1}{2}\right) = 0.$$

Solving the equation for n_1 :

$$\begin{aligned}\frac{5}{4}n_1 &= \frac{1}{2}+\frac{1}{4}n_2.\\ n_1 &= \frac{2}{5}+\frac{1}{5}n_2.\end{aligned}$$

This is the best response for vendor 1 (as a function of what vendor 2 does).

Let's now take the first order condition of $p_2(n_1, n_2)$:

$$\frac{\partial p_2(n_1,n_2)}{\partial n_2} = \frac{1}{2} - \frac{1}{2}\left(\frac{n_2 - n_1}{2}\right) - n_2 = 0.$$

Solving the equation for n_2 :

$$\begin{aligned}\frac{5}{4}n_2 &= \frac{1}{2}+\frac{1}{4}n_1.\\ n_2 &= \frac{2}{5}+\frac{1}{5}n_1.\end{aligned}$$

Solving the two equations in two unknowns yields the equilibrium locations $(n_1, n_2) = \left(\frac{1}{2}, \frac{1}{2}\right)$.

This is remarkably similar to the prediction of the Hotelling model of spatial competition. In the Hotelling model, the demand is constant as the distance increases (a citizen votes for a candidate if that candidate has the closest platform, meaning that the demand (vote number) is not a function of the distance from that platform to the citizen's preference point). Just like with candidates, the lunch trucks will continue to undercut each other until they choose the exact same location (the median location).

Is this equilibrium optimal? The equilibrium payoffs for both are equal to $p_1(n_1, n_2) = p_2(n_1, n_2) = \frac{3}{8}$. Notice that if the vendors were to locate at $(n_1, n_2) = \left(\frac{1}{4}, \frac{3}{4}\right)$, then the payoff for each would be equal to $p_1(n_1, n_2) = p_2(n_1, n_2) = \frac{7}{16}$, where this payoff is strictly greater than the equilibrium payoff of $\frac{3}{8}$.

In fact, the Pareto optimal allocations are $(n_1, n_2) = \left(\frac{1}{4}, \frac{3}{4}\right)$. These are the locations that would be chosen to jointly maximize the payoffs of both trucks. In other words, if both trucks were owned by the same company, then the company would instruct the trucks to locate at $(n_1, n_2) = \left(\frac{1}{4}, \frac{3}{4}\right)$.

But the ability to coordinate in practice is impossible unless the two lunch trucks have the same owner. This is the reality of allowing self-interested agents to make strategic decisions:

their decisions can often result in a suboptimal outcome.

2.6 Exercises

1. Consider a political economy model where the possible policies are ordered on the line from -1 to 1. All citizens have preferences lying somewhere on the line. All citizens must vote and they vote for the candidate whose platform lies closest to their preference point. Different from above, suppose that there are 3 candidates that must simultaneously select their platform. The candidate receiving the largest number of votes is the winner. Verify that there does not exist a Nash equilibrium.

2. Consider a Cournot duopoly in which the two firms simultaneously choose a quantity to produce: q_1 for firm 1 and q_2 for firm 2. The inverse demand function for this market is $P(Q) = 48 - \frac{1}{2}Q$, where Q is the total quantity in the market: $Q = q_1 + q_2$. The marginal cost of production is $c_1 = 6$ for firm 1 and $c_2 = 12$ for firm 2.

   Solve for the Nash equilibrium of this game (the quantity choices of both firms). What are the profits for both firms?

3. Consider a Cournot duopoly in which firms compete in a single good market by simultaneously choosing an output quantity. The inverse demand function is given by

$$P(Q) = \begin{Bmatrix} 60 - 2Q & \text{if } Q \leq 30 \\ 0 & \text{if } Q > 30 \end{Bmatrix}.$$

   The marginal cost of production can be either high $c_{high} = 12$ or low $c_{low} = 6$.

   (a) Two firms

   Let the market contain two firms, the first with high marginal cost of production, $c_1 = c_{high} = 12$, and the second with low marginal cost of production, $c_2 = c_{low} = 6$. Solve for the equilibrium output decisions of the two firms. What are the profits of each firm?

   (b) $2n$ firms

   Let the number of firms be $2n$ where $n = 1, 2, ..., \infty$. For the odd-numbered firms, the cost of production is high, $c_i = c_{high} = 12$ for $i = 1, 3, ..., 2n - 1$. For the even-numbered firms, the cost of production is low, $c_i = c_{low} = 6$ for $i = 2, 4, ..., 2n$. As

n grows, do both high cost and low cost firms continue to produce (yes or no)? (Hint: the best response functions have a lower bound of $q = 0$). If one half of the firms stops producing, what is the value of n at which this happens? What is the output for each of the other half of the firms at this point? What is the equilibrium market price P at this point?

4. Consider the Bertrand model of duopoly with transportation costs. Two firms compete in prices by simultaneously selecting prices p_1 and p_2. Firm 1 is located at mile 0 and firm 2 is located at mile 1. Consumers are uniformly distributed between mile 0 and mile 1 and incur a cost of 1 for each mile traveled to reach the selected firm. Consumers purchase from the firm with the lowest total cost (price plus travel cost) and have a demand equal to 1.

   The firms have marginal costs of production given by c_1 and c_2. The analysis in class was conducted under the implicit assumption that $|c_1 - c_2| \leq 3$. What happens if $|c_1 - c_2| > 3$, say $c_1 = 8$ and $c_2 = 4$? Specifically, what are the equilibrium price choices p_1 and p_2 and the profits of each firm?

5. Consider the following update of the Beach Game. Vendor 1 is initially located at 0 and vendor 2 is initially located at 1. Both vendors incure a cost of k per mile for moving to a new location. Given the locations of the two vendors (n_1, n_2), the demand function of the consumers located at z along the beach (where $0 \leq z \leq 1$) is defined as:

$$demand_i(z) = \left\{ \begin{array}{cl} 1 - |n_i - z| & \text{if } |n_i - z| < |n_j - z| \\ \frac{1}{2}(1 - |n_i - z|) & \text{if } |n_i - z| = |n_j - z| \\ 0 & \text{if } |n_i - z| > |n_j - z| \end{array} \right\}$$

   where $(i, j) = (1, 2)$ and $(i, j) = (2, 1)$.

   Both vendors select their location simultaneously (cannot observe the location of the other vendor prior to making their decision).

   As a function of $k : 0 \leq k \leq \frac{3}{4}$, solve for the equilibrium vendor positions of this game. Verify that the result obtained for the special case of $k = 0$ is $(n_1, n_2) = \left(\frac{1}{2}, \frac{1}{2}\right)$ (the answer obtained above for the zero cost case).

Bibliography

[1] Baird, Douglas, Robert Gertner, and Randal Picker. *Game Theory and the Law.* Cambridge, MA, and London, England: Harvard University Press, 1994.

[2] Dixit, Avinash, and Barry Nalebuff. *The Art of Strategy: A Game Theorist's Guide to Success in Business and Life.* 1st ed. New York, NY: W.W. Norton & Company, Inc., 2008.

[3] Gibbons, Robert. *Game Theory for Applied Economists.* Princeton, NJ: Princeton University Press, 1992.

[4] Mas-Colell, Andrew, Michael Whinston, and Jerry Green. *Microeconomic Theory.* New York, NY: Oxford University Press, Inc., 1995.

[5] Osborne, Martin J. *An Introduction to Game Theory: International Edition.* New York, NY: Oxford University Press, 2009.

Chapter 3

One of the Classic Blunders

3.1 Motivating Example

If you have ever seen the movie The Princess Bride, then you clearly remember this famous scene. The Dread Pirate Roberts has challenged the villain Vizzini to a "Battle of Wits." Two identical cups of wine are located on the table. Placing the two cups out of view of Vizzini, the Dread Pirate Roberts inserts the odorless, colorless poison "iocane powder" into one cup. The Dread Pirate Roberts then places one cup on the table in front of himself and the second on the table in front of Vizzini.

The first move is taken by the Dread Pirate Roberts, who must decide where to place the cup containing the poison. The second move is to be taken by Vizzini, who now must decide which cup to drink from. Vizzini drinks from the cup that he chooses and the Dread Pirate Roberts from the other. The choice by Vizzini is made without knowing the choice made by the Dread Pirate Roberts. Thus, the game involves simultaneous choices that can be analyzed using our game theoretic techniques.

Define the payoff from death (the poison is deadly) as 0 and the payoff from living as 1. The matrix form for this game is then given by (Dixit and Nalebuff, 2008, pgs. 141-142):

		Vizzini	
		Cup near DPR	Cup near Vizz
The Dread	Poison near DPR	1, 0	0, 1
Pirate Roberts	Poison near Vizz	0, 1	1, 0

Figure 3.1

Let's consider the logic as quoted by Vizzini.

> Vizzini: All I have to do is divine from what I know of you: are you the sort of man who would put the poison into his own goblet or his enemy's? Now, a clever man would put the poison into his own goblet, because he would know that only a great fool would reach for what he was given. I am not a great fool, so I can clearly not choose the wine in front of you. But you must have known I was not a great fool, you would have counted on it, so I can clearly not choose the wine in front of me.

Are you following so far? What about the next logical deduction?

> Vizzini: [The poison] iocane comes from Australia, as everyone knows, and Australia is entirely peopled with criminals, and criminals are used to having people not trust them, as you are not trusted by me, so I can clearly not choose the wine in front of you. And you must have suspected I would have known the powder's origin, so I can clearly not choose the wine in front of me.

Vizzini then distracts the Dread Pirate Roberts and switches the cups on the table. When the Dread Pirate Roberts returns his attention to the table, Vizzini chooses to drink from the cup in front of him, meaning that the Dread Pirate Roberts must drink from the other cup. After drinking, the Dread Pirate Roberts informs Vizzini that he has lost. To this, Vizzini laughs and replies:

> Vizzini: You only think I guessed wrong! That's what's so funny! I switched glasses when your back was turned! Ha ha! You fool! You fell victim to one of the classic blunders - The most famous of which is "never get involved in a land war in Asia" - but only slightly less well-known is this: "Never go against a Sicilian when death is on the line"!

At the end of the sentence, Vizzini falls to the ground dead. It turns out that the Dread Pirate Roberts had placed the poison into both cups and had spent the past months building up an immunity to the poison. The moral of the story is never agree to a game which seems to be a 50/50 gamble, because the person proposing the game must have more information about the game than you do. In this case, the Dread Pirate Roberts had changed the nature of the game from Example 3.1 into a game in which he would always win.

3.2 Solving the Battle of Wits

For our purposes, let's consider if there was an optimal way to play the game as viewed by Vizzini (as given in Figure 3.1). From the quotes by Vizzini, we see the problem with finding a Nash equilibrium. Each choice that Vizzini would make would lead the Dread Pirate Roberts to make a different choice, which would then change the choice of Vizzini, and so forth. Using the underline method, we see that a Nash equilibrium does not exist (Figure 3.2).

		Vizzini	
		Cup near DPR	Cup near Vizz
The Dread	Poison near DPR	$\underline{1}$, 0	0, $\underline{1}$
Pirate Roberts	Poison near Vizz	0, $\underline{1}$	$\underline{1}$, 0

Figure 3.2

Up to this point, the Nash equilibria that we have been trying to find are Nash equilibria in terms of pure strategies. When we said that (Fink, Fink) is a Nash equilibrium of the Prisoner's Dilemma Game, we mean that each suspect is choosing Fink with probability 100%. However, there is the possibility for the existence of another type of Nash equilibrium, a type that is called mixed-strategy Nash equilibrium. In a mixed-strategy Nash equilibrium, at least one of the players is dividing the 100% probability among multiple strategies. In this case, the player is indifferent between playing the strategies and uses a random number generator to decide which strategy is played.

A Nash equilibrium can be one of two types: (i) pure-strategy Nash equilibrium (as discussed in Chapters 1 and 2) and (ii) mixed-strategy Nash equilibrium $((p^*, 1-p^*), (q^*, 1-q^*))$, where player 1 assigns probability p^* to the first strategy (with the remaining $1-p^*$ to the second strategy) and player 2 assigns probability q^* to its first strategy (with the remaining $1-q^*$ to the second strategy). An important result in economics states that a Nash equilibrium must always exist. This result was proven by John Nash and is beyond the scope of this class to discuss. If a pure-strategy Nash equilibrium does not exist (as in Figure 3.1), then a mixed-strategy Nash equilibrium must exist.

If the Dread Pirate Roberts assigns probability p^* to the first strategy and $1-p^*$ to the second, then he must be receiving the same expected payoff from both strategies. This is because the expected payoff for the Dread Pirate Roberts is inversely related to the expected payoff of Vizzini. As one goes up, the other must necessarily go down. Payoff maximization dictates that the Dread Pirate Roberts will assign 100% probability to a strategy if its expected payoff is strictly higher, resulting in a lower expected payoff for Vizzini. Thus, in order for Vizzini to maximize his own expected payoff, he must set $(q^*, 1-q^*)$ so that the Dread Pirate Roberts is indifferent between its two strategies.

Let's write down how the expected payoffs are computed for the Dread Pirate Roberts. We only consider the payoffs of the Dread Pirate Roberts:

		Vizzini	
		Cup near DPR (q^*)	Cup near Vizz $(1-q^*)$
The Dread	Poison near DPR	$1 \cdot q^*, \ldots$	$0 \cdot (1-q^*), \ldots$
Pirate Roberts	Poison near Vizz	$0 \cdot q^*, \ldots$	$1 \cdot (1-q^*), \ldots$

Figure 3.3

Using these expected payoffs, the value q^* (the probability chosen by Vizzini) is such that the Dread Pirate Roberts has the same expected payoffs for its two strategies (Poison near DPR, Poison near Vizzini):

$$1 \cdot q^* + 0 \cdot (1-q^*) = 0 \cdot q^* + 1 \cdot (1-q^*).$$

Algebra yields $q^* = \frac{1}{2}$.

Similarly, the best choice by the Dread Pirate Roberts is to choose the value p^* such that Vizzini is indifferent between its two strategies. The expected payoffs are computed by considering only the payoffs of Vizzini:

		Vizzini	
		Cup near DPR	Cup near Vizz
The Dread	Poison near DPR (p^*)	$\ldots, 0 \cdot p^*$	$\ldots, 1 \cdot p^*$
Pirate Roberts	Poison near Vizz $(1-p^*)$	$\ldots, 1 \cdot (1-p^*)$	$\ldots, 0 \cdot (1-p^*)$

Figure 3.4

Vizzini is indifferent between its two strategies when

$$0 \cdot p^* + 1 \cdot (1-p^*) = 1 \cdot p^* + 0 \cdot (1-p^*).$$

Algebra yields $p^* = \frac{1}{2}$.

Thus, the unique Nash equilibrium is a mixed-strategy Nash equilibrium in which both players perfectly randomize over their two strategies:

$$((p^*, 1 - p^*), (q^*, 1 - q^*)) = \left(\left(\frac{1}{2}, \frac{1}{2}\right), \left(\frac{1}{2}, \frac{1}{2}\right)\right).$$

As an alternative solution method, we can make a plot of the best response curves for both players. This is shown in Figure 3.5 of the companion 'Figures' document. Let's put the probabilities p^* on the x-axis and q^* on the y-axis. Given the probabilities p^* of the Dread Pirate Roberts, what is the best response for Vizzini? The expected payoffs for Vizzini are $0 \cdot p^* + 1 \cdot (1 - p^*)$ if he chooses strategy 1 ($q^* = 1$) and $1 \cdot p^* + 0 \cdot (1 - p^*)$ if he chooses strategy 2 ($q^* = 0$). The best response (as a function of p^*) is to choose the strategy that yields the highest expected payoff. The best response plot is given as follows, where the vertical line at $p^* = \frac{1}{2}$ indicates that when $p^* = \frac{1}{2}$ Vizzini receives the same expected payoff from strategy 1 and strategy 2.

Now, given the probabilities q^* of Vizzini, what is the best response for the Dread Pirate Roberts? The expected payoffs for the Dread Pirate Roberts are $1 \cdot q^* + 0 \cdot (1 - q^*)$ if he chooses strategy 1 ($p^* = 1$) and $0 \cdot q^* + 1 \cdot (1 - q^*)$ if he chooses strategy 2 ($p^* = 0$). The best response for the Dread Pirate Roberts is added on top of the best response for Vizzini in Figure 3.6 in the companion 'Figures' document.

A Nash equilibrium is defined as the intersection of b_{Vizz} and b_{DPR}, the best response curves for both players. As can be seen from Figure 3.6, the unique Nash equilibrium (any pure-strategy Nash equilibria would appear as intersections on the boundary) is given by $(p^*, q^*) = \left(\frac{1}{2}, \frac{1}{2}\right)$, which is exactly what we found from the algebra above.

3.3 Solving Soccer Penalty Kicks

As another example, let's consider a pivotal play in soccer: the penalty kick (Dixit and Nalebuff, 2008, pgs. 143-151). The penalty kick is a strategic interaction between a scorer and a goalie. The scorer is located a mere 12 yards from the goalie, who is positioned on the goal line. The scorer decides either to kick the ball to his left (shoot left) or kick the ball to his right (shoot right). Given the size of the soccer goal and the speed at which a soccer ball travels, the only way for a goalie to block a well-placed shot is to dive before knowing which direction the shot is taken. Thus, without knowing the actions of the scorer, the goalie can either dive to his left (dive left) or dive to his right (dive right).

Let's assume that the scorer is right-footed, meaning that the option to shoot left is more powerful, but less accurate. This means that if the scorer chooses shoot left, then the goal is made 70% of the time when the goalie guesses wrong (chooses dive left) and 30% of the time when the goalie guesses correctly (chooses dive right). Let's just assume that the choice of shoot right results in a binary outcome: the goal is made 100% of the time when the goalie guesses wrong (choose dive right) and 0% of the time when the goalie guesses correctly (chooses dive left).

You with me so far? All of the words are translated into the following normal form game.

			Goalie	
			Dive Left	Dive Right
			(q)	$(1-q)$
Scorer	Shoot Left	(p)	0.7, 0.3	0.3, 0.7
	Shoot Right	$(1-p)$	0, 1	1, 0

Figure 3.7

A useful practice is to write in the probabilities next to the strategies. Thus, p is the probability that the scorer chooses shoot left and q is the probability that the goalie chooses dive left. The payoffs simply represent the probability of a goal being scored (for the scorer) and the probability of a goal not being scored (for the goalie).

Using the underline method, there does not exist a pure-strategy Nash equilibrium. This is a logical outcome. By settling on one strategy with 100% probability (the definition of a pure strategy), a player would be giving a huge advantage to his opponent.

Let's take the steps required to determine the mixed-strategy Nash equilibrium. Recalling our method from the previous section, p is determined to make goalie indifferent between dive left and dive right:

$$\begin{gathered}
\text{Goalie's exp. payoff of dive left} = \text{Goalie's exp. payoff of dive right} \\
0.3 \cdot p + 1 \cdot (1-p) = 0.7 \cdot p + 0 \cdot (1-p) \\
1 - p = 0.4p \\
p = \tfrac{5}{7}.
\end{gathered}$$

Similarly, q is determined to make scorer indifferent between shoot left and shoot right:

$$\begin{aligned} \text{Scorer's exp. payoff of shoot left} &= \text{Scorer's exp. payoff of shoot right} \\ 0.7 \cdot q + 0.3 \cdot (1 - q) &= 0 \cdot q + 1 \cdot (1 - q) \\ 0.3 + 0.4q &= 1 - q \\ q &= \tfrac{1}{2}. \end{aligned}$$

Thus, the unique mixed-strategy Nash equilibrium is

$$((p, 1 - p), (q, 1 - q)) = \left(\left(\frac{2}{3}, \frac{1}{3} \right), \left(\frac{2}{5}, \frac{3}{5} \right) \right).$$

The scorer shoots left more often (his strong, but inaccurate direction), while the goalie perfectly randomizes between right and left. The probabilities for the scorer are assigned not based upon the strength of his own shot, but upon making the goalie indifferent between its two choices.

3.3.1 Classroom Exercise 3: Rock Paper Scissors Lizard Spock

You may have played the game Rock Paper Scissors as a child (Figure 1.24), but perhaps you have not heard of the variant Rock Paper Scissors Lizard Spock (Dixit and Nalebuff, 2008, pgs. 151-154). The game, as described on the television show The Big Bang Theory, is a two-player simultaneous choice game in which each player can choose among 5 options. The Rock option is chosen by forming a fist with the right hand, the Paper option with a flat hand, the Scissors option by extending the index and middle fingers, the Lizard option by forming the shape of a sock puppet, and the Spock option by forming the Vulcan hand gesture used for greeting ("Live long and prosper"). Each of the options beats two other options and loses to two other options. As always, Rock smashes Scissors, Scissors cuts Papers, and Paper covers Rock. Figure 3.8 in the companion 'Figures' document illustrates the relations.

What is the Nash equilibrium of the game Rock Paper Scissors Lizard Spock? Do you have better success playing this game as compared to Rock Paper Scissors?

3.3.2 Classroom Exercise 4: A Variant of Rock Paper Scissors

Consider the following variant of Rock Paper Scissors (Dixit and Nalebuff, 2008, pgs. 151-154). A pair of students play a match consisting of repeated plays of the game. For each game played, 3 points is awarded for a victory with Rock, 2 points is awarded for a victory with Paper, and 1 point is awarded for a victory with Scissors. No points are awarded if the game is a tie (in which both players have chosen the same strategy). The match ends when the first player reaches 10 points. We have already seen that the equilibrium strategy for the game is to assign a certain probability to Rock, a certain probability to Paper, and the remaining probability to Scissors. When the game awards the same points no matter the manner of victory, the probabilities are the same for all three strategies. Now, when the game awards different points for the different victories (in particular, a premium for victories with Rock), how do your probabilities change? Can you figure out a winning strategy that can beat a computer (a perfect randomizer)?

3.4 Symmetric Mixed Strategies of the Room Selection Game

Recall the Room Selection Game from Chapter 1 in which three players must choose either the room ranking 123 or the room ranking 213. Once all the rankings are collected, the rooms are allocated according to the rules specified in Chapter 1. In particular, the expected payoffs for all possible outcomes of the game are given in Table 3.1 below.

Outcome	Expected payoffs
3 submit 123	Submit 123 : Payoff = 20
2 submit 123, 1 submit 213	Submit 123 : Payoff = 18, Submit 231 : Payoff = 24
1 submit 123, 2 submit 213	Submit 123 : Payoff = 30, Submit 231 : Payoff = 15
3 submit 213	Submit 213 : Payoff = 20

Table 3.1

We will be looking for symmetric mixed strategy Nash equilibrium, that is, those mixed strategies in which all three players assign the same probabilities. It is possible that asymmetric mixed-strategy Nash equilibria exist, but these are illogical and much harder to find. Define p as the probability that any player chooses the strategy 123, where the remaining probability $1 - p$ is the probability that any player chooses 213.

A mixed-strategy Nash equilibrium must be one in which all three players are indifferent between their two strategies. Thus, the expected payoff from 123 equals the expected payoff from 213. For any one player, the probability that the other two players choose $(123, 123)$

equals p^2, the probability that the other two players choose either $(123, 213)$ or $(213, 123)$ equals $2p(1-p)$, and the probability that the other two players choose $(213, 213)$ equals $(1-p)^2$.

$$\begin{aligned} \text{Exp. payoff of } 123 &= \text{Exp. payoff of } 213 \\ 20p^2 + 18\{2p(1-p)\} + 30(1-p)^2 &= 24p^2 + 15\{2p(1-p)\} + 20(1-p)^2. \end{aligned}$$

Simplifying the equality results in the following expression:

$$\begin{aligned} -4p^2 + 3\{2p(1-p)\} + 10(1-p)^2 &= 0. \\ -4p^2 + 6p - 6p^2 + 10 - 20p + 10p^2 &= 0. \end{aligned}$$

The squared terms cancel and we can easily solve for the value of p :

$$\begin{aligned} 10 &= 14p \\ p &= \frac{5}{7}. \end{aligned}$$

This means that the symmetric mixed-strategy Nash equilibrium calls for each player to choose the ranking 123 with probability $\frac{5}{7}$ and ranking 213 with probability $\frac{2}{7}$.

3.5 Properties of Mixed Strategy Equilibria

3.5.1 2x2 Games

For any discrete choice game between two players, if both players have two possible strategies, then we have a 2×2 game and the game can be represented in the normal form by a 2×2 matrix. The following facts are useful when it comes to finding the mixed-strategy Nash equilibria in these games.

> Fact 1: If a pure-strategy Nash equilibrium does not exist, then there is one mixed-strategy Nash equilibrium.

Fact 1 deals with the random-strategy games. The fact follows from the result as shown by John Nash that a Nash equilibrium (either in pure strategies or in mixed strategies) must exist.

Fact 2: If multiple pure-strategy Nash equilibria exist, then there is one mixed-strategy Nash equilibrium.

Fact 2 deals with the coordination and anti-coordination games. The fact follows from the following analysis. Consider the following game in which the payoffs for the row player are (A, B, C, D) and the payoffs for the column player are (a, b, c, d) (Gibbons, 1992, pgs. 41-44).

		Column Player	
		L	R
Row Player	U	A, a	B, b
	D	C, c	D, d

Figure 3.9

Let's suppose that the multiple pure-strategy Nash equilibria are (A, a) and (D, d). By the definition of a Nash equilibrium, we have the following inequalities (assume that they are strict inequalities so that we don't have to worry about ties):

$$\begin{aligned} A &> C \text{ and } D > B \\ a &> b \text{ and } d > c \end{aligned}$$

Since $A > C$ and $D > B$, then there must exist some fraction q such that $qA + (1-q)B = qC + (1-q)D$. This fraction q will be the probability that the column player assigns to strategy L in the mixed-strategy Nash equilibrium. Similarly, as $a > b$ and $d > c$, then there must exist some fraction p such that $pa + (1-p)c = pb + (1-p)d$. This fraction p is the probability that the row player assigns to strategy U.

Fact 3: If a unique pure-strategy Nash equilibrium exists, then there does not exist a mixed-strategy Nash equilibrium.

Fact 3 deals with collective action problems. The fact follows from the following analysis. Again, consider Figure 3.9, but now suppose that the unique pure-strategy Nash equilibrium is given by (D, d). Since (D, d) is a Nash equilibrium, then by definition:

$$D > B \text{ and } d > c.$$

Further, as (A, a) is not a Nash equilibrium, then either:

$$A < C \text{ or } a < b.$$

Suppose, without loss of generality, that $A < C$. With $A < C$ and $B < D$, then no matter what probability q is chosen (where $0 \leq q \leq 1$), it can never be true that:

$$qA + (1-q)B = qC + (1-q)D.$$

Thus, we cannot find the probabilities required to make the row player indifferent, so we cannot have a mixed-strategy Nash equilibrium.

3.5.2 Games with Larger Strategy Sets

In the 2×2 games, the limited number of strategies allows us to prove some useful facts. Fact 1 continues to hold true in games with larger strategy sets.

Updating Fact 2

Fact 2 needs to be updated. To do so, let's find the mixed-strategy Nash equilibria in a 3×3 coordination game. In particular, let's find the mixed-strategy Nash equilibria in the game from Chapter 1, Exercise 3, part (b). The normal form game that you wrote down may be different from mine, but I chose the values given in Figure 3.10 below (Gibbons, 1992, pgs. 31-40).

		Student 2		
		Purdue	Indiana	Notre Dame
	Purdue	3, 2	0, 0	0, 0
Student 1	Indiana	0, 0	2, 1	0, 0
	Notre Dame	0, 0	0, 0	1, 3

Figure 3.10

Define (p_1, p_2, p_3) as the probabilities that Student 1 assigns to its three choices (in that order), where $p_1 + p_2 + p_3 = 1$. Likewise, define (q_1, q_2, q_3) as the probabilities that Student 2 assigns to its three choices (in that order), where $q_1 + q_2 + q_3 = 1$. The formal definition of mixed-strategy equilibrium says that a player must be indifferent between all strategies that it assigns strictly positive probability to. This requires us to consider each of the four following cases.

<u>Case I</u>: Student 1 assigns $p_1 > 0$, $p_2 > 0$, and $p_3 > 0$.

In this case, it must be true that Student 1's expected payoff from Purdue equals its expected payoff from Indiana equals its expected payoff from Notre Dame. In particular, we have that

$$3q_1 = 2q_2 = q_3.$$

Given that $q_1 + q_2 + q_3 = 1$, this implies that $(q_1, q_2, q_3) = \left(\frac{2}{11}, \frac{3}{11}, \frac{6}{11}\right)$.

Now, as we just found $q_1 > 0$, $q_2 > 0$, and $q_3 > 0$, then Student 2 must be indifferent between Purdue, Indiana, and Notre Dame. This implies that

$$2p_1 = p_2 = 3p_3.$$

Given that $p_1 + p_2 + p_3 = 1$, this implies that $(p_1, p_2, p_3) = \left(\frac{3}{11}, \frac{6}{11}, \frac{2}{11}\right)$.

Case II: Student 1 assigns $p_1 > 0$, $p_2 > 0$, and $p_3 = 0$.

In this case, it must be that Student 1's expected payoff from Purdue equals its expected payoff from Indiana and is strictly greater than its expected payoff from Notre Dame. This means that

$$3q_1 = 2q_2,$$

so $q_1 = \frac{2}{3}q_2$. Let's now consider the problem of Student 2. If $q_3 > 0$, then Student 3 must be indifferent between Notre Dame and some other strategy. This requires that the expected payoff from Notre Dame is strictly positive. However, with $p_3 = 0$, the expected payoff from Notre Dame is 0. For this reason, $q_3 = 0$ and $(q_1, q_2) = \left(\frac{2}{5}, \frac{3}{5}\right)$.

As $q_1 > 0$ and $q_2 > 0$, then Student 2 must be indifferent between Purdue and Indiana. This requires that

$$2p_1 = p_2.$$

Thus, the probabilities for Student 1 are $(p_1, p_2, p_3) = \left(\frac{1}{3}, \frac{2}{3}, 0\right)$.

Case III: Student 1 assigns $p_1 > 0$, $p_2 = 0$, and $p_3 > 0$.

The same method yields the mixed-strategy Nash equilibrium $(p_1, p_2, p_3) = \left(\frac{3}{5}, 0, \frac{2}{5}\right)$ and $(q_1, q_2, q_3) = \left(\frac{1}{4}, 0, \frac{3}{4}\right)$.

Case IV: Student 1 assigns $p_1 = 0$, $p_2 > 0$, and $p_3 > 0$.

The same method yields the mixed-strategy Nash equilibrium $(p_1, p_2, p_3) = \left(0, \frac{3}{4}, \frac{1}{4}\right)$ and $(q_1, q_2, q_3) = \left(0, \frac{1}{3}, \frac{2}{3}\right)$.

We conclude that for a coordination game with n pure-strategy Nash equilibria, the number of mixed-strategy Nash equilibrium is equal to $\sum_{i=2}^{n} \binom{n}{i}$, where $\binom{n}{n} = 1$ by the definition of a combination.

Updating Fact 3

Fact 3 only holds true for 2×2 games. For larger games, it is false.[1] Figure 3.11 considers a game in which there is a unique pure-strategy Nash equilibrium and additionally two mixed-strategy Nash equilibria.

		Student 2		
		L	C	R
	U	3, 3	1, 2	2, 0
Row Player	D	2, 1	2, 2	1, 3

Example 3.11

Let's use the underline method to find that there only exists one pure-strategy Nash equilibrium: (U,L).

		Student 2		
		L	C	R
	U	$\underline{3}$, $\underline{3}$	1, 2	$\underline{2}$, 0
Row Player	D	2, 1	$\underline{2}$, 2	1, $\underline{3}$

Example 3.12

To find the mixed-strategy Nash equilibria, define $(p, 1-p)$ as the probabilities that the row player chooses for (U,D) and (q_1, q_2, q_3) as the probabilities that the column player chooses for (L,C,R), where $q_1 + q_2 + q_3 = 1$. The row player only has two strategies, so in a mixed-strategy Nash equilibrium, the row player must be indifferent between both its strategies. This means that the expected payoff of U equals the expected payoff of D:

$$3q_1 + q_2 + 2q_3 = 2q_1 + 2q_2 + q_3.$$

This implies

$$q_1 - q_2 + q_3 = 0,$$

[1] What remains true for a normal form game with any number of players and strategies is the following: if the process of IDSDS results in a unique outcome, then a mixed strategy Nash equilibrium does not exist. Recall that if all players have dominant strategies, then the process of IDSDS results in a unique outcome.

where the sum $q_1 + q_2 + q_3 = 1$ implies

$$1 - 2q_2 = 0.$$

Thus, we know that the value of $q_2 = \frac{1}{2}$, but we can't say anything about (q_1, q_3), except that $q_1 + q_3 = \frac{1}{2}$.

With $q_2 > 0$, then the expected payoff of C for the column player must be equal to the expected payoff of all strategies played with strictly positive probability. The expected payoff of C is equal to $2p + 2(1-p) = 2$. We see that for the value of $p = \frac{1}{2}$, then the expected payoff of L equals $3p + (1-p) = 2$, while the expected payoff of R is $0 + 3(1-p) = \frac{3}{2}$. When $p = \frac{1}{2}$, the column player is indifferent between L and C, but strictly prefers both to R. This require that $q_3 = 0$. So one mixed-strategy Nash equilibrium is

$$((p, 1-p), (q_1, q_2, q_3)) = \left(\left(\frac{1}{2}, \frac{1}{2}\right), \left(\frac{1}{2}, \frac{1}{2}, 0\right)\right).$$

In Exercise 3, we can find a second mixed-strategy Nash equilibrium.

3.6 Exercises

1. Find the mixed-strategy Nash equilibria in the following figures of Chapter 1: Figures 1.20, 1.21, 1.22, 1.23, and 1.24.

2. Find the mixed-strategy Nash equilibria of the normal form games that you created for Chapter 1, Question 3, parts (a) and (c) (we are not quite ready to handle part (b)).

3. There exists one more mixed-strategy Nash equilibrium in Figure 3.11. Find it.

Bibliography

[1] Baird, Douglas, Robert Gertner, and Randal Picker. *Game Theory and the Law.* Cambridge, MA, and London, England: Harvard University Press, 1994.

[2] Dixit, Avinash, and Barry Nalebuff. *The Art of Strategy: A Game Theorist's Guide to Success in Business and Life.* 1st ed. New York, NY: W.W. Norton & Company, Inc., 2008.

[3] Gibbons, Robert. *Game Theory for Applied Economists.* Princeton, NJ: Princeton University Press, 1992.

[4] Mas-Colell, Andrew, Michael Whinston, and Jerry Green. *Microeconomic Theory.* New York, NY: Oxford University Press, Inc., 1995.

[5] Osborne, Martin J. *An Introduction to Game Theory: International Edition.* New York, NY: Oxford University Press, 2009.

Part II

Dynamic Games of Complete Information

Chapter 4

Know When to Hold 'Em and When to Fold 'Em

4.1 Motivating Example

To begin our study of dynamic games, let's consider a card game played in two rounds (Dixit and Nalebuff, 2008, pgs. 23-35). There are two female players, each of whom holds a card over her head. The color of the card is either Red or Green. Each player can observe the color of her opponent's card, but not the color of her own. The object of the game is for each player to deduce the color of her card as quickly as possible. The game is played in two rounds. In the first round, if a player knows her card is red, then she exits the stage. After the round ends, the players update their beliefs about the color of their card and round two begins. In the second round, if a player knows her card is red, then she exits the stage.

The players are told at the beginning of the game that at least one of the cards is red.

If you were playing this game, how would you determine if you had a red card or not? Consider Figure 4.1 from the companion 'Figures' document. This shows the two possible things that you could be looking at: your opponent is holding a red card or your opponent is holding a green card.

If your opponent is holding a green card, then your own card must be red. This is because at least one of the cards is red. Knowing this, you would exit the stage in the first round.

If your opponent is holding a red card, then your own card can be either red or green. The only way to determine your card's color is to learn from the actions of your opponent. Your opponent is capable of using the exact same logic that you considered in the previous paragraph. That is, if your opponent sees that you have a green card, then she would exit in the first round. At the end of the first round, conditional upon seeing your opponent with a red card, then Figure 4.2 illustrates the two possible additional things that you can observe: your opponent stays or your opponent goes.

At the end of the first round, you now have enough information to know the color of your card. If your opponent exited in the first round, then your card must be green. If your card is green then you won't exit in the second round. If your opponent did not exit, then your card must be red. If your card is red, then you will exit in the second round.

Let's see if this same logic can be used in a larger version of the same game.

4.1.1 Classroom Exercise 5: Card Game

Let's consider a 5-player version of the card game played in five rounds. Each player holds a card over her head. The color of the card is either Red or Green. Each player can observe the color of the four other players, but not her own color. The game consists of five rounds. In each round, if a player knows that her card is red, then she exits the stage.

The players are told at the beginning of the game that at least one of the cards is red. Suppose that you are playing this game and observe the following colors for the other players: 3 red and 1 green. This is depicted in Figure 4.3 in the companion 'Figures' document.

How can you determine the color of your card based upon the decisions made by the other players about when to exit the stage?

4.2 Defining Contingent Actions

As evidenced in the card game in the previous section, in a dynamic game players must be able to adjust their actions based upon the previous moves played in the game. More precisely, players form one strategy, which is a collection of the contingent actions that must be taken in each round that a player is called upon to make a move. These contingent actions must take into account all possible contingencies of the game. To repeat, an action is one choice taken at one particular point in time, while a strategy is all the actions that are taken over all possible contingencies.

To illustrate the meaning of contingent actions, let's see what they would look like for the simple children's game of Tic-Tac-Toe (Mas-Colell et al., 1995, pgs. 220-222). As everyone knows, a game of Tic-Tac-Toe is played on a 3x3 board. This board game is depicted in

Figure 4.4, where I have numbered each square 1-9 for labeling purposes.

The game of Tic-Tac-Toe is played in 9 rounds, though the final rounds may be irrelevant when determining a winner. The winner is the first player to place 3 icons in a row. I don't believe there is any standard rule about which player starts, so let's suppose that the player with the 'X' icon moves first. Define the contingent action of the 'X' player in round 1 as $x_1 \in \{1, 2, ..., 9\}$. In the first round, the action is not actually contingent as no prior action has been taken. Let's suppose that $x_1 = 5$ as shown in Figure 4.5.

The second round belongs to the player with the 'O' icon. Define the contingent action of the 'O' player in round 2 as $o_2 : \{1, 2, ..., 9\} \rightarrow \{1, 2, ..., 9\} \backslash \{x_1\}$. The contingent action is a function from x_1 to the set of remaining squares (the 'O' player cannot choose x_1 as this square has already been claimed by the 'X' player in the first round). Let's suppose that $o_2(5) = 4$ as depicted in Figure 4.6 below.

For those of you that are experts at the game of Tic-Tac-Toe, you realize that 'O' has made a dreadful mistake. The 'X' player is guaranteed of victory by taking the following contingent actions.

Round 3	$x_3(\cdot) = 3$		
Round 5	$x_5(\cdot) = 7$	if $o_4(\cdot) \neq 7$	Victory!
	$x_5(\cdot) = 1$	if $o_4(\cdot) = 7$	
Round 7	$x_7(\cdot) = 2$	if $o_6(\cdot) \neq 2$	Victory!
	$x_7(\cdot) = 9$	if $o_6(\cdot) = 2$	Victory!

Table 4.1

Figures 4.7-4.11 consider the sequence of plays assuming that $o_4(\cdot) = 7$ and $o_6(\cdot) = 2$.

4.2.1 Classroom Exercise 6: Game of Hex

The game of Hex is similar to Tic-Tac-Toe, except that a winner is always guaranteed. The game of Hex is played on hexagonal graph paper and the board has dimension 7x7 (though other dimensions are possible).

With each student assigned to a 2-person team, have one member of the team learn the rules of Hex (and engage in practice games), while the other member of the team remains ignorant about the rules. The member of the team that has learned the rules is responsible for writing down an optimal strategy for the other member of the team. The other member of the team then has to implement this strategy in a competition against another team.

Given the time constraints, the strategy will likely be incomplete (and not specify actions for all contingencies), so the player that has learned the rules needs to impart in as few words as possible the "best practices" for victory.

4.3 The Logic of Backward Induction

Now that we know how to define contingent actions, we are ready to describe the form that we use to depict dynamic games, and the equilibrium concept that we use for such games.

Dynamic games are depicted in what is called the extensive form. The extensive form consists of the following four elements:

- nodes
- non-intersecting branches
- which one of the players acts at a particular node
- payoffs for all players at the end of every branch

Consider the example of the extensive form in Figure 4.12 with 2 players, 3 nodes, and 4 branch ends (Gibbons, 1992, pgs. 115-117 and Mas-Colell et al., 1995, pgs. 271-272).

Figure 4.12 specifies that Player 1 begins the game by choosing either L or R. A node is a juncture in the tree that provides all the information about the prior actions that have been taken. The game in Figure 4.12 contains 3 nodes, where the first node is at the top of the tree at the point where Player 1 must choose between L and R. If L is chosen by Player 1, then the game is at the node at the bottom of the branch for L (hereafter, the left node). If R is chosen by Player 1, then the game is at the node at the bottom of the branch for R (the right node).

Player 2 takes contingent actions, that is, Player 2 chooses either LL or RR at both the left and right node. For all possible outcomes of the game (and there are four of them in this figure), the payoffs for the players are written (Player 1 payoff, Player 2 payoff). The convention is that the payoffs are listed in the order in which the players take their initial actions. For example, the outcome in which Player 1 chooses L and Player 2 chooses LL results in a payoff of 3 for Player 1 and a payoff of 1 for Player 2.

The games that we discuss in this chapter are games of **perfect information**. This means that at each point in time at which a player is called upon to take an action, the player knows all moves that have been made previously. Chapter 5 considers dynamic games with imperfect information.

The most common game used to describe the extensive form is the Centipede game. The Centipede game is an alternating offer game between 2 players containing 100 periods. In the odd-numbered periods, Player 1 is given the opportunity to stop the game. If he stops the game, he earns $1 more in payoff than Player 2. If he continues the game, then an extra dollar is added to the pot. In the even-numbered periods, Player 2 is given the opportunity to stop the game. If he stops the game, he earns $2 more in payoff than Player 1. If he continues the game, then an extra dollar is added to the pot.

The game, as depicted in Figure 4.13 in the companion 'Figures' document (Mas-Colell et al., 1995, pgs. 281-282), continues up until Round 100 and the decision that is made by Player 2 in that round. Looking at the shape of the extensive form (or the game tree) for the game and the number of periods, we can see where the name 'centipede' comes from.

How many nodes are contained in the game in Figure 4.13? The answer is 100 (of course).

We are now prepared to introduce the equilibrium concept that is used for dynamic games. The concept is due to Reinhard Selten and is called subgame perfect Nash equilibrium (SPNE). A subgame perfect Nash equilibrium is the set of strategies for all players such that a Nash equilibrium is being played at all nodes.

We go about finding the subgame perfect Nash equilibrium using the logic of backward induction. Let's illustrate the logic by finding the subgame perfect Nash equilibrium for Figure 4.13. As the name suggests, the method of backward induction begins at all nodes in the last round of the game. The last round is Round 100, where Player 2 must choose either Stop or Continue. At this node, a Nash equilibrium must be played. This means that Player 2 chooses the action that provides the highest payoff. Player 2 chooses Stop, since the payoff of 51 (from Stop) is higher than the payoff of 50 (from Continue). To indicate this, I darken the Stop branch in Round 100 (see Figure 4.14).

Let's proceed backward to Round 99, where Player 1 must choose either Stop or Continue. Player 1 recognizes that by choosing Continue, then the game proceeds to Round 100, where Player 2 will choose Stop. Basically, Player 1 knows that the future plays of the game will proceed along the darkened branches. With this information, Player 1 chooses the action that provides the highest payoff (50 for Stop and 49 for Continue). Player 1 chooses Stop. I darken the Stop branch in Round 99 (see Figure 4.15).

The logic proceeds in a similar fashion through all the rounds from Round 98 all the way down to Round 1. In Round 1 (as indicated in Figure 4.16), Player 1 chooses Stop (and the payoff of 1) rather than Continue and following the darkened branches to a payoff of 0.

Thus, the subgame perfect Nash equilibrium payoffs are (1,0). A subgame perfect Nash equilibrium is the set of strategies, one for each of the two players. The strategies specify all actions that are taken in the game, even those actions that are not reached along the equilibrium path. The equilibrium path is reached by following the darkened branches beginning in Round 1. Notice in Figure 4.16 that the game ends in Round 1 with the choice of Stop by Player 1. The equilibrium path is simply this one branch. However, that does not lessen the importance of the strategies played off the equilibrium path. These strategies combine to entice Player 1 to make the decision to Stop in Round 1.

Thus, the subgame perfect Nash equilibrium is ((Stop, Stop, Stop, ...), (Stop, Stop, Stop, ...)), where the first vector (Stop, Stop, Stop, ...) indicates the 50 rounds in which Player 1 chooses Stop and the second vector (Stop, Stop, Stop, ...) indicates the 50 rounds in which Player 2 chooses Stop.

4.3.1 Classroom Exercise 7: Flag Game

Consider a game with two teams and 21 flags (Dixit and Nalebuff, 2008, pgs. 44-47). Taking turns, each team would choose to remove either 1, 2, or 3 flags. The team that is forced to take the last flag of the 21 is the loser.

What is the optimal strategy at any stage of the game (especially the initial stage) for each of the two teams? If a team could choose to move first, would it make this choice?

4.4 The Dynamic Room Selection Game

With the concept of a subgame perfect Nash equilibrium in our toolbox, let's return to the Room Selection Game for the final time (I promise). As before, three roommates submit their rankings over the three available rooms. The expected payoffs of all possible outcomes are listed in Table 3.1. For this version of the game, rather than the roommates simultaneously submitting their rankings, the game will be played in three rounds. In the first round, Roommate 1 submits her ranking. The other two roommates observe this choice. In the second round, Roommate 2 submits her ranking. This is observed by Roommate 3. In the third round, Roommate 3 submits her ranking.

Given that we have a dynamic game, we want to use the extensive form (Mas-Colell et al., 1995, pgs. 271-272). The extensive form is depicted in Figure 4.17 in the companion

'Figures' document.

The expected payoffs are determined for each possible outcome of the game (see Table 3.1) and are such that Roommate 1's expected payoff is listed on top, followed by Roommate 2's, and ending with Roommate 3's on the bottom.

If we want to determine how the game in Figure 4.17 will be played, we must find the subgame perfect Nash equilibrium. This is achieved using backward induction, beginning with the 4 nodes at which Roommate 3 must make a decision. The optimal actions of Roommate 3 are darkened in Figure 4.18.

Next, we proceed to the two nodes at which Roommate 2 must make a decision. Roommate 2 knows that the game will proceed along the darkened branches in the future, so takes the optimal actions as indicated in Figure 4.19.

Finally, Roommate 1 optimally chooses the ranking 213 (as seen in Figure 4.20), which results in the equilibrium payoff of (24,18,18).

Remembering that a subgame perfect Nash equilibrium must specify the complete strategies for all players (including actions taken off the equilibrium path), the subgame perfect Nash equilibrium is (213, (213,123), (213,123,123,123)) where 213 is the strategy of Roommate 1, (213,123) is the strategy of Roommate 2 specifying actions for each of her two nodes (reading from left to right), and (213,123,123,123) is the strategy of Roommate 3 specifying actions for each of her four nodes (reading from left to right).

4.5 Entry Deterrence

The game of Entry Deterrence (Mas-Colell et al., 1995, pgs. 268-290 and Baird et al., 1994, pgs. 159-186) analyzes the logic of making credible threats and determines what commitments can be chosen to make threats credible. Consider a market for gizmos. The market demand is inelastic and fixed at 100 units. That simply means that the residents in a particular town must buy 100 gizmos and will buy them at any price. There currently exists one store in the town that sells gizmos. This store will be Store 1 and we can think of it as the incumbent firm. Store 1 can sell gizmos at one of two prices: $0.50 and $2.

A potential second store (Store 2) is deciding whether or not to open a store in the town and compete with Store 1 for the gizmo market. Store 2 is not as efficient as Store 1, so could only sell the gizmos at a price of $2. If Store 2 decides to open, it must pay a $50 fixed cost.

The residents in the town will purchase the gizmos from the store with the cheapest price.

If both stores sell at the same price, then the residents will split their demand, meaning that 50 units are bought from Store 1 and 50 units are bought from Store 2.

The game is a dynamic game with two rounds. In the first round, Store 2 decides whether or not to open a store. In the second round, Store 1 chooses the price that it will sell gizmos at (recall that Store 2 must always price at $2).

Figure 4.21 depicts the game. Using the logic of backward induction, we see that the subgame perfect Nash equilibrium is (Open store, ($2, $2)). The equilibrium payoffs are $50 for Store 2 ($100 in revenue minus the $50 fixed cost) and $100 for Store 1 ($100 in revenue). Store 1 would love to be able to say, "If you enter the market, I will price so low that you will earn a negative profit," but such a threat is not credible. It is not credible, because when Store 1 is called upon to carry out this threat, its best interest is to charge $2 rather than $0.50 (the left node in Figure 4.20 after Store 2 chooses "Open store").

Is there a way that Store 1 can change the game so that it works in its favor? What if Store 1 can take a preemptive action to commit itself to the threat action? Let's consider how the game would be updated. Store 1 can place an advance order for 100 gizmos. Choosing to place the advance order will cost Store 1 $50. Moreover, let's suppose that Store 1 does not sell all 100 gizmos (due to competition from Store 2). Well, if all 100 gizmos are not sold, then Store 1 must incur a $75 storage cost.

The rest of the game remains unchanged from above, but we now have a three round dynamic game. In the first round, Store 1 must decide whether or not to make the early purchase order. In the second round, Store 2 (upon observing the order decision of Store 1) must decide whether to open a store or not. In the third round, Store 1 chooses the price that it will sell gizmos at. The game is depicted in Figure 4.22. As Store 1 now takes the first action, the payoffs are reversed and read (Store 1, Store 2).

The subgame perfect Nash equilibrium is ((Early, $0.50, $2, $2, $2), (Don't open, Open)), with payoffs equal to ($150,0). This means that even with all the costs of placing the advance order ($50 up front and a potential storage cost of $75 if all units are not sold), Store 1 has a higher payoff in Figure 4.22 compared to Figure 4.21 (in which the payoff was only $100). In Figure 4.22, Store 1 has taken an action that commits itself to carrying through on its threat to charge $0.50 per gizmo if Store 2 enters the market. This serve as the necessary deterrent to keep Store 2 out of the market.

4.6 Stackelberg Duopoly

Recall the Cournot duopoly model in which two firms simultaneously choose a quantity to produce. The Stackelberg duopoly model (Gibbons, 1992, pgs. 48, 61-64) considers the outcome when the quantity choices are made sequentially. One of the firms is the leader and chooses the quantity it wishes to produce. Suppose firm 1 is the first mover and chooses the quantity q_1. The second firm is firm 2 and is the second mover in the game. This firm observes the quantity choice made by firm 1 and then makes its own decision about the quantity to produce. This quantity is denoted q_2. The total quantity produced in the market equals $Q = q_1 + q_2$. The price at which the good can sell is given by the inverse demand function $P(Q) = a - bQ$, where P is the price of the good. Suppose that the marginal cost of production is c_1 for firm 1 and c_2 for firm 2.

This is a dynamic game with continuous choices. The strategy for the follower (firm 2) is to take an action that is contingent upon the quantity choice of firm 1. Using the method of backward induction, we start with finding the optimal contingent actions for firm 2. As we will see, this is the same as finding the best response for firm 2 as we did in Chapter 2. Contingent upon q_1, the profit function and first order condition for firm 2 are given by:

$$\begin{aligned} \pi_2(q_1, q_2) &= [a - b(q_1 + q_2)]\, q_2 - c_2 q_2. \\ \frac{\partial \pi_2(q_1, q_2)}{\partial q_2} &= a - 2bq_2 - bq_1 - c_2 = 0. \end{aligned}$$

This allows us to solve for the best response

$$q_2(q_1) = \frac{a - bq_1 - c_2}{2b}.$$

This specifies the action taken by firm 2 for all contingencies (all q_1 chosen by firm 1).

Firm 1 can solve the best response of firm 2 equally as well as we just did. Firm 1 knows that whatever quantity choice it makes, firm 2 will respond according to that best response function. Thus, firm 1 incorporates the best response function for firm 2 into its own profit function:

$$\pi_1(q_1) = [a - b(q_1 + q_2(q_1))]\, q_1 - c_1 q_1.$$

The profit function is now a function only of the quantity choice q_1. Let's insert the best response function and simplify the algebra:

$$\pi_1(q_1) = \left[a - b\left(q_1 + \frac{a - bq_1 - c_2}{2b}\right)\right] q_1 - c_1 q_1.$$
$$\pi_1(q_1) = \left[\frac{a}{2} - b\left(\frac{q_1}{2}\right) + \frac{c_2}{2}\right] q_1 - c_1 q_1.$$

Taking the first order conditions allows us to solve for the optimal quantity choice q_1 that maximizes $\pi_1(q_1)$:

$$\frac{d\pi_1(q_1)}{dq_1} = \left[\frac{a}{2} - b\left(\frac{q_1}{2}\right) + \frac{c_2}{2}\right] - \frac{b}{2} q_1 - c_1 = 0.$$
$$q_1 = \frac{a + c_2 - 2c_1}{2b}.$$

On the equilibrium path, the quantity choice for firm 2 is given by:

$$q_2 = \frac{a - b\left(\frac{a}{2b} + \frac{c_2 - 2c_1}{2b}\right) - c_2}{2b}.$$
$$q_2 = \frac{a - 3c_2 + 2c_1}{4b}.$$

As an example, if $c_1 = c_2 = c$, then the output choice of firm 1 ($q_1 = \frac{a-c}{2b}$) is the same as the monopolist output choice and exactly twice as much as the output choice of firm 2 ($q_2 = \frac{a-c}{4b}$).

The subgame perfect Nash equilibrium for the Stackelberg duopoly is $\left(\frac{a-c}{2b}, \frac{a-bq_1-c_2}{2b}\right)$. As the second mover, firm 2 adopts a strategy that is a best response function.

Any game in which the first player earns a higher profit simply by its right to take the first action has the property of the "first mover advantage." Another game with the first mover advantage is the dynamic Room Selection Game of Figure 4.17.

4.7 Application: Primaries and Runoffs

This section first shows how to find the equilibrium strategies for voters in a dynamic voting game (Dixit and Nalebuff, 2008, pgs. 359-385). While on the topic of voting, this section also demonstrates that for nearly every conceivable method to determine the winner of a voting game, the voters have an incentive to lie about their true preferences. The lone exception to this rule is the "Dictator" method in which the choice of one voter determines the winner of the election.

4.7.1 Dynamic Voting Game

Suppose that 5 voters must determine a winner among the potential candidates: Washington, Adams, Jefferson, and Madison. The voting occurs in three rounds. In the first round, the voters determine the winner in the head-to-head matchup of Jefferson and Madison. In the second round, the winner of Jefferson-Madison faces off against Adams. In the third and final round, the candidate that survives among Adams, Jefferson, and Madison faces off against Washington.

The voters are named Anne, Ben, Carl, Deb, and Ed. The voters have strict preferences over the four candidates as follows:

	Anne	Ben	Carl	Deb	Ed
1st choice	Mad	Jeff	Jeff	Wash	Wash
2nd choice	Adams	Adams	Adams	Jeff	Mad
3rd choice	Wash	Mad	Mad	Mad	Jeff
4th choice	Jeff	Wash	Wash	Adams	Adams

Table 4.1

Who will ultimately win the election? Note that the voters only care who ultimately wins the election and there is no rule requiring them to vote according to the rankings in Table 4.1.

To determine the election winner, we apply the logic of backward induction. In the third round, we can have the following matchups: (i) Washington vs. Adams, (ii) Washington vs. Jefferson, or (iii) Washington vs. Madison. As this is the final round of the game, then the voters will always choose to vote honestly; that is, they will vote according to Table 4.1.

Wash vs. Adams	3 votes for Adams (Anne, Ben, Carl)	Adams wins
Wash vs. Jeff	3 votes for Wash (Anne, Deb, Ed)	Wash wins
Wash vs. Mad	3 votes for Mad (Anne, Ben, Carl)	Mad wins

Table 4.2

After knowing what happens for each possibility in the final round, we can make two conclusions: (i) Jefferson can never win the election and (ii) Washington can only win the election if he is matched up against Jefferson. Using these facts, then we must delete Jefferson from the rankings of all 5 voters. This is shown in Table 4.3. Anyone that votes for Jefferson in either the first two rounds is actually casting a vote for Washington.

	Anne	Ben	Carl	Deb	Ed
1st choice	Mad	Adams	Adams	Wash	Wash
2nd choice	Adams	Mad	Mad	Mad	Mad
3rd choice	Wash	Wash	Wash	Adams	Adams

Table 4.3

When we delete Jefferson from the rankings, anyone ranked below Jefferson moves up one spot. In the second round, we can have the following matchups: (i) Adams vs. Jefferson (reallly, Washington) and (ii) Adams vs. Madison. Recognizing that a vote for Jefferson is a vote for Washington, then the voters will vote according to Table 4.4 below.

Adams vs. Jeff (really, Wash)	3 votes for Adams (Anne, Ben, Carl)	Adams wins
Adams vs. Mad	3 votes for Mad (Anne, Deb, Ed)	Mad wins

Table 4.4

The conclusion drawn from this second round is that Washington cannot win. Moreover, Adams can only win if he is matched up against Washington (really, Jefferson).

The first round lists a matchup between Jefferson and Madison. However, Jefferson cannot win, so any vote for Jefferson is a vote for Washington. Further, Washington cannot win, so any vote for Jefferson/Washington is a vote for Adams. Thus, the real matchup is between Adams and Madison. Adams supporters vote Jefferson, while Madison supporters vote Madison. From Table 4.3 above, we see that in a matchup between Adams and Madison, Madison has 3 votes of support (Anne, Deb, Ed). So Madison wins the election.

The winner of the election, contested in three rounds, is Madison.

Just for comparison, let's see what the outcome would be if the voters did not vote optimally (i.e., did not vote strategically), but rather just cast their votes blindly following the rankings in Table 4.1. In the first matchup, Jefferson vs. Madison, Jefferson wins with 3 votes (Ben, Carl, Deb). In the second matchup, Adams vs. Jefferson, Jefferson wins with 4 votes (Ben, Carl, Deb, Ed). In the third matchup, Washington vs. Jefferson, Washington wins with 3 votes (Anne, Deb, Ed). The victor is Washington.

4.7.2 "You Can't Handle the Truth"

What we saw with the dynamic voting game is that voters need not vote according to their true rankings. What we do in this section is consider five other methods for determining an election winner and show that only the last one provides the incentive for voters to vote truthfully. This leads to the important Gibbard-Satterthwaite result about social choice (a corollary of Arrow's impossibility theorem).

An important truth to hold onto as we proceed through this subsection is that I am not finding the equilibrium under each of the different voting methods. I am only determining whether or not the voters will vote truthfully.

Majority voting

In majority voting, all 5 voters select one candidate. The candidate with the largest number of votes wins the election. Using the preferences from Table 4.1, if the voters vote truthfully, then Madison gets 1 vote (Anne), Jefferson gets 2 votes (Ben, Carl), and Washington gets 2 votes (Deb, Ed). This means that Jefferson wins the election 50% of the time (coin flip) and Washington wins the election 50% of the time.

Does one of the voters wish to vote strategically? Consider Anne. If she votes for Washington, then he wins the election. This outcome is strictly preferred for Anne. Thus, Anne has an incentive to not vote truthfully.

Borda count

In a Borda count, all 5 voters rank the candidates from 1 to 4. The candidates receive points based upon where they are ranked by voters:

1st Choice	4 points
2nd Choice	3 points
3rd Choice	2 points
4th Choice	1 point

Table 4.5

The points don't have to be allocated in this exact fashion, but the central message is that voters will always have an incentive to not vote truthfully (no matter what point values are assigned). Use Table 4.1, consider the point totals if all the candidates did vote truthfully:

Washington	2+1+1+4+4=12
Adams	3+3+3+1+1=11
Jefferson	1+4+4+3+2=14
Madison	4+2+2+2+3=13

Table 4.6

If voters are truthful, then Jefferson would win. Does one of the voters wish to vote strategically? Suppose Ed switches his ranking to Madison - Washington - Adams - Jefferson. In this case, the Borda count updates to:

Washington	2+1+1+4+3=11
Adams	3+3+3+1+2=12
Jefferson	1+4+4+3+1=13
Madison	4+2+2+2+4=14

Table 4.7

With this change, Madison wins. This outcome is strictly preferred by Ed.

Survival of the fittest

The first survivor method asks the 5 voters to rank the candidates 1 to 4. The candidate with the fewest 1st place votes is eliminated. That candidate is removed from the rankings, meaning that any candidate ranked below him moves up one spot. With the updated rankings, the candidate with the fewest 1st place votes is eliminated. The rankings are adjusted over the final two candidates. The one with the fewest 1st place votes is eliminated, leaving the remaining candidate as the winner of the election.

Let's consider the outcome if all voters vote truthfully. The rankings would be given as in Table 4.1:

	Anne	Ben	Carl	Deb	Ed
1st choice	Mad	Jeff	Jeff	Wash	Wash
2nd choice	Adams	Adams	Adams	Jeff	Mad
3rd choice	Wash	Mad	Mad	Mad	Jeff
4th choice	Jeff	Wash	Wash	Adams	Adams

Table 4.1

As can be seen, Adams has the fewest 1st place votes (none), so he is removed from the rankings. All candidates ranked below Adams move up one spot.

	Anne	Ben	Carl	Deb	Ed
1st choice	Mad	Jeff	Jeff	Wash	Wash
2nd choice	Wash	Mad	Mad	Jeff	Mad
3rd choice	Jeff	Wash	Wash	Mad	Jeff

Table 4.8

As can be seen, Madison has the fewest 1st place votes (one), so he is removed from the rankings. The updated rankings are given by:

	Anne	Ben	Carl	Deb	Ed
1st choice	Wash	Jeff	Jeff	Wash	Wash
2nd choice	Jeff	Wash	Wash	Jeff	Jeff

Table 4.9

In the head-to-head between Washington and Jefferson, Jefferson is eliminated meaning that Washington is the winner.

Does one of the voters have an incentive to vote strategically (i.e., not according to the rankings in Table 4.1)? Consider what happens if Ben submits the ranking Madison - Jefferson - Adams - Washington. In this case, Adams is eliminated in the first round (zero first round votes). Now, unlike before, Jefferson is eliminated in the second round (only one first place vote). In the final matchup between Madison and Washington, Madison emerges victorious. Having Madison win the election, rather than Washington, is strictly preferred by Ben.

Survival of the 'not unfittest'

This second survivor method asks the 5 voters to rank the candidates from 1 to 4. The candidate with the most 4th place votes is eliminated. That candidate is removed from the rankings, meaning that any candidate ranked below him moves up one spot. With the updated rankings, the candidate with the most 3rd place votes is eliminated. The rankings are adjusted over the final two candidates. The one with the most 2nd place votes is eliminated, leaving the remaining candidate as the winner of the election.

Let's consider the outcome if all voters vote truthfully. The rankings would be as given in Table 4.1:

	Anne	Ben	Carl	Deb	Ed
1st choice	Mad	Jeff	Jeff	Wash	Wash
2nd choice	Adams	Adams	Adams	Jeff	Mad
3rd choice	Wash	Mad	Mad	Mad	Jeff
4th choice	Jeff	Wash	Wash	Adams	Adams

Table 4.1

As can be seen, Washington and Adams have the most 4th place votes. Rather than flipping a coin, let's just eliminate them both. Now, we are down to two candidates: Jefferson and Madison. The updated rankings are:

	Anne	Ben	Carl	Deb	Ed
1st choice	Mad	Jeff	Jeff	Jeff	Mad
2nd choice	Jeff	Mad	Mad	Mad	Jeff

Table 4.10

In the head-to-head between Jefferson and Madison, Madison is eliminated meaning that Jefferson is the winner.

Does one of the voters have an incentive to vote strategically? Consider what happens if Ed submits the ranking Washington - Madison - Adams - Jefferson. In this case, both Washington and Jefferson are eliminated in the first round (each has two 4th place votes). This leaves a showdown between Adams and Madison, and Madison will triumph. The outcome in which Madison wins is strictly preferred by Ed to the outcome in which Jefferson wins.

Dictator

In the dictator method, the election winner is the candidate that is chosen by Anne (we could have assigned the dictator to be any one of the 5 voters). Anne votes truthfully for Madison. The other voters don't have an incentive to do anything but vote truthfully; their votes do not matter.

This is the first voting method in which voting truthfully is an equilibrium strategy for all voters. The following result states that this is the only voting method with this property.

Theorem 4.1 *(Gibbard-Satterthwaite)*

Given the collection of all voters' rankings, we define a voting mechanism as a rule that determines one election winner. Assume that there exists at least 3 candidates and the voting mechanism allows all candidates the possibility of winning. That is, we can't simply forbid Washington and Adams from winning and just have a showdown between Jefferson and Madison. Under this assumption, then the only voting mechanism in which voters have the incentive to vote truthfully is the "Dictator" method.

Proof. Go to graduate school in economics. ■

The result has a very negative prediction. If we want voters to vote according to their true preferences, then the only voting method that we can use is the "Dictator" method, a method that I am sure most of us find completely unfair (unless you happen to be the dictator). The idea is at the heart of the movie Swing Voter starring Kevin Costner.

4.8 Exercises

1. What is the subgame perfect Nash equilibrium for Figure 4.12?

2. Suppose that two players are bargaining over $1. The game takes place in rounds, beginning with Round 1. The game ends when an offer is accepted. Player 1 makes offers in odd-numbered rounds and Player 2 makes offers in even-numbered rounds. In rounds when a player is not making an offer, it observes the offer made by the other player and decides whether to "Accept" or "Reject." At the end of each round, $0.20 is removed from the pool of money. That is, if an agreement is reached in Round 2, the total pool of money is $0.80; if agreement in Round 3, $0.60, and so forth. (Hint: The game has only 5 periods.)

   Find the subgame perfect Nash equilibrium of this bargaining game.

3. In the kingdom called Hearts, consider a game with 4 players: the King of Hearts, the Queen of Hearts, the Jack of Hearts and the Ten of Hearts. The four currently share the wealth in the kingdom:

Percent of wealth

Ten	Jack	Queen	King
6%	46%	24%	24%

Table 4.11

The law of Hearts are: (i) a proposal is voted on by all 4 players and can only be passed if it receives at least 3 votes and (ii) if a player makes a proposal that does not pass, then that player loses its entire share of the wealth (which is distributed evenly among the other players).

The voting proceeds sequentially beginning with the Ten, proceeding to the Jack, followed by the Queen, and ending with the King.

Suppose that the Ten makes the following proposal: (i) the wealth in the kingdom will be updated as

Proposed percent of wealth

Ten	Jack	Queen	King
40%	20%	20%	20%

Table 4.12

and (ii) if the proposal passes, then any player voting against the proposal will be stripped of its wealth share (which is distributed evenly among the other players).

Will the proposal by the Ten pass? To receive full credit, you must justify your answer by writing the subgame perfect Nash equilibrium strategies for all players.

4. Consider a market with two firms. The firms compete by choosing a quantity to produce, but this competition is sequential: firm 1 first chooses a quantity to produce q_1 and given this observed choice, firm 2 then chooses a quantity to produce q_2. The inverse demand function for this market is $P(Q) = 30 - 2Q$, where P is the unit price and Q is the total quantity in the market: $Q = q_1 + q_2$. The marginal cost of production is $c_1 = 4$ for firm 1 and $c_2 = 2$ for firm 2.

Solve for the subgame perfect Nash equilibrium (SPNE) of this game (the quantity choices of both firms). Which firm earns a higher profit?

5. Consider a strategic voting game with three candidates and three voters. The candidates are Romney, Perry, and Obama. The voters first select who wins between Romney and Perry. Next, the voters select who wins between the Romney-Perry winner and Obama. In terms of the winner of the Romney-Perry winner vs. Obama election, the voters have the following rankings:

	Voter 1	Voter 2	Voter 3
1st Choice	Obama	Romney	Perry
2nd Choice	Perry	Obama	Romney
3rd Choice	Romney	Perry	Obama

Table 4.13

The voters vote strategically. Which candidate wins in the Romney-Perry winner vs. Obama election?

Chapter 5

The Strategy of Law and Order

5.1 Motivating Example

Consider the following game of political intrigue and scandal. A city councilman (CC) is suspected of diverting government funds for his own personal use in remodeling his home. Such an activity is illegal, but in order to obtain a conviction, the district attorney (DA) must obtain proof. The district attorney can choose to hire a private investigator at a cost of \$150. The city councilman suspects that the district attorney may be attempting to obtain proof, so he can either continue with the remodeling or can stop. If the private investigator is hired and the city councilman continues with the remodel, then the district attorney is able to obtain proof. The cost of a trial is \$100 for both parties. If the district attorney has proof of the illegal activity, then the city councilman is convicted, and must pay a \$1000 fine plus the costs (trial cost and private investigator cost) of the district attorney. Without proof of the illegal activity, then no outcome is reached in trial and no further transfers are made.

The city councilman receives a benefit of \$10 from continuing with the illegal remodel. The district attorney places a value $V > 0$ on being able to convict the city councilman of the wrongdoing. What value must V take so that the city councilman only continues the illegal activity with probability less than 25%?

5.2 Information Sets and Subgames

In Chapter 4, we considered games of perfect information, in which the players always knew the prior actions that had been taken in the game. This chapter provides a methodology that allows us to look at games with both perfect and imperfect information. In a game with **imperfect information**, some players do not know the actions that have previously been taken by other players.

Dynamic games are always represented in the extensive form (also called the tree form). With imperfect information, these trees will include information sets. An information set is all nodes that are connected by a dotted line. The information set indicates that the player taking an action at that point of the game does not know which node has been reached. Consider Figure 5.1 in the companion 'Figures' document in which some nodes are in information sets and others are not.

A subgame is defined as all the branches that begin at a node not in an information set and continues until the end of the game. Thus, to count the number of subgames, one must simply count the number of nodes not in an information set. Confirm that the game depicted in Figure 5.1 has 4 subgames. Figure 5.2 identifies each of the 4 subgames (and yes, the entire game is a subgame).

The more complete definition of a subgame perfect Nash equilibrium states that it is the set of strategies for all players such that a Nash equilibrium is played in each subgame. If the game has perfect information, then this definition is equivalent to the one given in the previous chapter (where I stated that a subgame perfect Nash equilibrium is the set of strategies such that a Nash equilibrium is played at all nodes). For games with imperfect information, the remaining sections of the chapter offer practice in finding the subgame perfect Nash equilibria.

5.3 Relation to the Normal Form

Recall Figure 4.1 from the previous chapter. This is a different game than Figure 5.3 (Gibbons, pg. 188) in the companion 'Figures' document. Figure 5.3 has an information set that contains both the left and right nodes of Player 2. This means that when Player 2 is called upon to choose an action, he does not know whether the game is at the left node (Player 1 chose L) or the right node (Player 1 chose R).

As was found in Exercise 4.1, the subgame perfect Nash equilibrium for Figure 4.1 (found using backward induction) is (R, (RR, LL)), with equilibrium payoffs (2,1).

How about we go about finding the subgame perfect Nash equilibrium for Figure 5.3? By definition, a subgame perfect Nash equilibrium is such that a Nash equilibrium is played in all subgames. So, to apply our method of backward induction, we must fund the lowest nodes in the game tree that are not in an information set. The lowest node satisfying this requirement in Figure 5.3 is the top node, at which Player 1 must choose between L and R. The game proceeding from this node to the end is the one and only subgame.

To find the optimal strategies for this one subgame, consider that Player 2 taking an action without knowing the action taken by Player 1 is equivalent to the setting in the simultaneous games that we analyzed in Chapters 1 and 3. Thus, we have the following equivalence between the extensive form and the normal form. This is depicted in Figure 5.4.

Using this equivalence, we can find the subgame perfect Nash equilibrium of Figure 5.3 by using the underline method for the game written in the matrix form (the normal form). The subgame perfect Nash equilibrium is (L, RR), with equilibrium payoffs (1,2). Notice that when the actions of Player 1 are observed (Figure 4.1), then Player 1 has a higher payoff (equal to 2) compared to when the actions are not observed (Figure 5.3, with payoff 1). This is the concept of a "first mover advantage." In Figure 5.3, as the actions are not observed, it is irrelevant whether Player 1 is written on top or Player 2 is written on top. This is illustrated in Figure 5.5. To appreciate this figure, recall that the payoffs are always written (First mover, Second mover), so that when the order of players is switched, the payoffs must be switched as well.

5.4 Finitely Repeated Games

Suppose that the Prisoners' Dilemma in Figure 5.6 is played twice (Gibbons, 1992, pgs. 119-120).

		Suspect 2	
		Mum	Fink
Suspect 1	Mum	10, 10	0, 20
	Fink	20, 0	5, 5

Figure 5.6

That is, in the first round of play, both suspects simultaneously choose either Mum or Fink. The players then observe the outcome of Round 1 (that is, they observe their payoffs). Given this information, Round 2 begins. In this round, both suspects again must simultaneously choose either Mum or Fink. The payoffs to each player are equal to the sum of the payoffs that each receives in each of the two rounds. This game can be depicted in extensive form in Figure 5.7 (where the payoffs are written for only 2 of the 16 possible outcomes [you can fill in the rest]).

How do we solve for the subgame perfect Nash equilibrium of the 2-period version of Figure 5.6? We use the logic of backward induction. In Round 2, there are four possible outcomes that can be observed: (i) Round 1 ended in (Mum, Mum), (ii) Round 1 ended in (Mum, Fink), (iii) Round 1 ended in (Fink, Mum), and (iv) Round 1 ended in (Fink, Fink), where the first strategy in the parentheses belongs to Suspect 1.

Beginning at each of these possible outcomes, would the suspects choose to play Mum or Fink in Stage 2? The play in the final period must always be a Nash equilibrium of the stage game (the stage game is Figure 5.6). This is because each of the four possible outcomes is the beginning of a subgame and we must have a Nash equilibrium in each subgame. In the stage game (in Figure 5.6), the Nash equilibrium is (Fink, Fink).

If we wish, we can darken the branches (Fink, Fink) in Figure 5.7 for each of the lower 4 subgames. Now we carry the logic up into Round 1 of the game. Each suspect knows that when the game reaches Round 2, both suspects will play Fink. Thus, the outcome in Round 2 is independent of what happens in Round 1. So Round 1 must also follow a Nash equilibrium of the stage game: (Fink, Fink).

We see that the subgame perfect Nash equilibrium calls for both suspects to choose Fink in all subgames, with resulting payoffs equal to (10, 10), since payoffs of (5, 5) are reached in both periods.

How would this result vary if the suspects were asked to play a 3-round version of Figure 5.6? What about a 1000-round version of Figure 5.6? No matter the number of rounds, so long as that number is finite, when the suspects are called upon to make a decision in any round, they know that each of them will choose Fink in the next round and every round thereafter. Thus, it is best for each of them to choose Fink in the current round. This logic leads to the following fact.

> Fact: If the stage game has a unique Nash equilibrium, then when this stage game is repeated a finite number of times, the subgame perfect Nash equilibrium is simply to play the Nash equilibrium in all rounds.

Let's discuss what this fact does not say. It does not say anything about the case in which the stage game is repeated an infinite number of times (a topic that we postpone until

Chapter 6). Additionally, it does not say anything about the case in which the stage game has multiple Nash equilibria. See Exercise 3 in which the stage game has multiple Nash equilibria to understand the economic predictions in this setting.

5.5 The Plea Bargain

A defendant (d) is charged with a crime. The prosecutor (p) can make an initial plea bargain offer. If the defendant accepts, then the defendant does not go to trial and the game ends. If the defendant rejects, then the trial begins. Prior to the trial, both the prosecutor and the defendant simultaneously choose to hire either an expensive legal team or a cheap legal team. If only one side chooses expensive, it wins the trial.

If both sides make the same choice (either both expensive or both cheap), then the prosecutor wins with probability $\frac{3}{4}$ (meaning that the defendant wins with the remaining probability $\frac{1}{4}$).

The payoff of winning is 4. The payoff of losing is -4. The cost of a cheap legal team is 0, while the cost of an expensive legal team is 1. The plea bargain would be a transfer from the defendant to the prosecutor of size $x \geq 0$.

The game is depicted in Figure 5.8 in the companion 'Figures' document (Baird et al., 1994, pgs. 244-266). To fill in the payoffs, let's consider what each player will receive for each of the possible trial outcomes. If both players hire an expensive legal team, the expected payoff for the prosecutor is equal to:

$$\frac{3}{4}(\text{winning payoff}) + \frac{1}{4}\ (\text{losing payoff})\ \text{- cost} = \frac{3}{4}(4) + \frac{1}{4}(-4) - 1 = 1.$$

The expected payoff for the defendant (in similar fashion) is equal to -3. If both players hire a cheap legal team, then the probabilities of winning are the same for both parties, but now both sides don't have to pay the legal costs. Thus, the prosecutor's expected payoff is 2 and the defendant's expected payoff is -2.

If the prosecutor chooses expensive and the defendant chooses cheap, then the prosecutor wins and pays the legal cost (winning payoff - cost = 3), while the defendant loses (payoff = -4). The opposite payoffs are obtained if the prosecutor chooses a cheap legal team and the defendant chooses an expensive legal team.

I want to know the range of plea bargain offers (the smallest and the largest) such that (i) the prosecutor finds it in his best interest to make the offer and (ii) the defendant finds it in his best interest to accept. In other words, I want to find the values of x so that the subgame perfect Nash equilibrium follows the path (Offer x, Accept).

Let's begin to solve the subgame perfect Nash equilibrium. The game has 4 subgames. The bottom 2 subgames are identical and correspond to the trial. These 2 subgames are equivalent to the simultaneous game (normal form game) depicted in Figure 5.9.

		defendant	
		Expensive	Cheap
prosectuor	Expensive	1, -3	3, -4
	Cheap	-4, 3	2, -2

Figure 5.9

Using the underline method, we see that the Nash equilibrium of this stage game is (Expensive, Expensive). Let's darken the branches for (Expensive, Expensive) in the extensive form in Figure 5.10.

Let's consider the decision by the defendant to either accept or reject the plea bargain. We want to find the values such that the defendant chooses to accept the plea bargain. The defendant chooses accept when his payoff from doing so ($-x$) exceeds his payoff from going to trial (equal to -3 from Figure 5.10). Thus, we require $-x > -3$, or $x < 3$.

Let's consider the decision by the prosecutor to offer the plea bargain. The prosecutor knows that the plea bargain will be accepted by the defendant (provided $x < 3$). The prosecutor chooses to make the offer if his payoff from doing so (x) exceeds his payoff from going to trial (equal to 1 from Figure 5.10). Thus, we require that $x > 1$.

In sum, if $1 < x < 3$, then the prosecutor will make the plea bargain offer, and the defendant will accept the offer. This together with the decisions by both players to choose the expensive legal team comprise the subgame perfect Nash equilibrium.

5.6 Solution to the Motivating Example

Now that we have solved one episode of Law and Order, let's return to the original Law and Order episode that piqued our interest in the beginning of the chapter. The game has two players: the district attorney (DA) and the city councilman (CC). The third agent, the private investigator, plays a passive role and is not an actual player in the game. The game, without any information sets, is given in Figure 5.11 in the companion 'Figures' document.

The payoffs are taken exactly from the description of the game. For instance, if the private investigator is hired, the city councilman continues the remodel, and the district attorney goes to trial, then the payoff for the DA is $V > 0$ and the payoff for the CC is -$1340 ($10 benefit - $1000 fine - $100 own court costs - $100 DA court costs - $150 DA private investigator costs). The remaining payoffs can be verified by the reader.

Given the description of the game, where do we need to place information sets? First, the city councilman must decide whether or not to continue the remodel and must make this decision without knowing whether or not the private investigator has been hired (by definition, a private investigator can avoid being seen). Thus, one information set is required for the two nodes at which the city councilman must choose an action.

If the private investigator has been hired, the city councilman will soon know with certainty whether or not the city councilman is continuing with the remodel. But, if the private investigator has not been hired, then the district attorney has no way of gathering this information. Thus, we need a second information set to account for the fact that after deciding not to hire the private investigator, the district attorney must make a trial decision without knowing whether the city councilman is continuing with the illegal activity.

The complete extensive form is shown in Figure 5.12.

Now that we have translated the word problem into the extensive form, we can find the subgame perfect Nash equilibrium. To begin the method of backward induction, we look at the lower-left nodes for the district attorney. As shown in Figure 5.13, the district attorney will choose to go to trial if the city councilman continues with the remodel, and will not go to trial if the city councilman has stopped the remodel.

We now move by backward induction to the next node up the tree that is not in an information set. The next node is the top node at which the district attorney must decide whether or not to hire the private investigator. We know that this extensive form game with information sets is equivalent to a normal form game, but how many strategies do each of the players have? For the city councilman, the answer is easy. He has two strategies: continue the remodel or stop. Let's make the city councilman (the second mover) the column player in our normal form game.

How many strategies does the district attorney have? One possible strategy is to hire the private investigator. In this case, the actions for the district attorney at the bottom of the tree (to go to trial or not) are determined as in Figure 5.13. For the case in which the district attorney does not hire the private investigator, then he gains no new information prior to having to make the decision about whether to go to trial or not. Thus, the district attorney has two additional strategies: (i) don't hire the private investigator and go to trial and (ii) don't hire the private investigator and don't go to trial. The district attorney will be the row player with the three strategies as laid out in this paragraph.

The upper subgame of Figure 5.13 is then equivalent to the following 3×2 normal form game. Take care when transcribing the payoffs from the extensive form to the normal form to ensure that the payoffs correspond to the appropriate outcomes.

		CC	
		Continue	Stop
	Hire	V, -1340	-150, 0
DA	Not hire, Trial	-100, -90	-100, -100
	Not hire, No trial	0, 10	0, 0

Figure 5.14

Using the underline method, we find that the normal form game does not have any pure strategy Nash equilibria (recall that $V > 0$ by assumption). In this case, as our analysis in Chapter 3 indicates, we must look for mixed strategy Nash equilibria. Define (p_1, p_2, p_3) to be the probabilities associated with each of the three strategies (in order) of the district attorney, where $p_1 + p_2 + p_3 = 1$. Define $(q, 1 - q)$ as the probabilities associated with each of the two strategies (in order) of the city councilman.

First, the values for (p_1, p_2, p_3) are determined to make the city councilman indifferent between his two strategies.

Payoff from Continue	Payoff from Stop
$-1340p_1 - 90p_2 + 10p_3$	$-100p_2$

Table 5.1

Setting these two expected payoffs equal to each other yields:

$$10p_2 + 10p_3 = 1340p_1.$$

This implies that $p_1 > 0$. Recall that the DA must be indifferent between all the strategies that he plays with strictly positive probability. Consider the expected payoffs for the DA as a function of the probabilities $(q, 1-q)$ for the CC.

Payoff from Hire	Payoff from (Not hire, Trial)	Payoff from (Not hire, No trial)
$qV + (1-q)(-150)$	-100	0

Table 5.2

We see that the strategy (Not hire, No trial) strictly dominates (Not hire, Trial). So the DA will never play (Not hire, Trial) with strictly positive probability. This means that $p_2 = 0$. From the equation $10p_2 + 10p_3 = 1340p_1$ above, then it must be the case that:

$$p_3 = 134p_1$$

or $(p_1, p_2, p_3) = \left(\frac{1}{135}, 0, \frac{134}{135}\right)$. This states that the DA will only hire the private investigator with an extremely low probability $p_1 = \frac{1}{135}$.

The question asks us to determine what the value V needs to be so that the city councilman continues the illegal activity with probability less than 25%. Thus, we will find the smallest possible value for V so that $q = 0.25$ (probability 25% of continuing with the remodel).

The probabilities $(q, 1-q)$ are determined to make the DA indifferent between Hire and (Not hire, No Trial):

$$\begin{aligned} qV + (1-q)(-150) &= 0 \\ (0.25)V + (0.75)(-150) &= 0 \\ V &= 450. \end{aligned}$$

The DA has value V from successfully prosecuting the city councilman. Provided that $V \geq \$450$, then the DA can hire the private investigator with probability $p_1 = \frac{1}{135}$ and keep the probability of illegal activity less than or equal to 25%. So the government officials need to take into account how much benefit a DA receives from a successful prosecution when making decisions about illegal activities.

5.7 Exercises

1. Consider the following game between two investors. They have each deposited D at a bank. The bank has invested these deposits in long-term projects. If the bank is

forced to liquidate its investment before the long-term project is complete, then it can only recover $2r$, where $D > r > \frac{D}{2}$. If the bank allows the project to be completed, the total payout is higher and given by $2R$, where $R > D$.

There are two dates at which investors can make withdrawals from the bank: stage one before the long-term project is complete and stage two after the project is complete. The decisions to withdraw or not in each stage are made simultaneously by the two investors. The two-stage game described so far is equivalent to the following two normal-form games:

Stage 1

		Investor 2: Withdraw	Investor 2: Don't
1	Withdraw	r, r	$D, 2r - D$
	Don't	$2r - D, D$	Stage 2

Figure 5.15

Stage 2

		Investor 2: Withdraw	Investor 2: Don't
1	Withdraw	R, R	$2R - D, D$
	Don't	$D, 2R - D$	R, R

Figure 5.16

Find all the subgame perfect Nash equilibria of this game.

2. Consider the following game of contract negotiation between a sports team and an athlete. The negotiation takes place in 3 stages. There is no discounting of payoffs across the stages. In the first stage, both parties simultaneously choose whether the athlete receives a big contract or a small contract. If both parties agree, the contract is signed; otherwise, they proceed to the second stage. Payoffs are indicated below:

		Athlete: Big	Athlete: Small
Sports Team	Big	20, 100	Stage 2
	Small	Stage 2	100, 20

Figure 5.17

In the second stage, the two parties simultaneously decide (given the results of Stage 1) whether to walk away from the negotiation or to continue. If at least one of them wants to walk away, then the negotiation ends. If both want to continue, then the negotiation continues to the third stage. Payoffs are indicated below:

		Athlete	
		Walk	Bargain
Sports	Walk	25, 25	10, 30
Team	Bargain	30, 10	Stage 3

Figure 5.18

The third stage is depicted in Figure 5.19 in the companion 'Figures' document. In the third stage, the sports team makes either a high or a low offer. Seeing this offer, the athlete can either accept or reject it. Either way, the negotiations end. If the offer is accepted, a contract is signed. If the offer is rejected, then both parties leave with nothing.

Solve for the subgame perfect Nash equilibrium.

3. Consider two firms in an industry. The firms are attempting to maintain an uncontractable agreement between them (about prices, quality, advertising, or possibly entrance into new markets). The firms play a game in two stages. Each stage is identical, with no discounting future payoffs. In both stages, each firm simultaneously chooses one of three actions: {Cooperate, Defect, Ignore}. The payoffs are given as:

		Firm 2		
		Cooperate	Defect	Ignore
	Cooperate	10, 10	0, 13	2, 12
Firm 1	Defect	13, 0	1, 1	0, 0
	Ignore	12, 2	0, 0	5, 5

Figure 5.20

Is it possible that a subgame perfect Nash equilibrium can be constructed in which both firms "Cooperate" in the first stage? If so, write down the complete equilibrium strategy.

Bibliography

[1] Baird, Douglas, Robert Gertner, and Randal Picker. *Game Theory and the Law.* Cambridge, MA, and London, England: Harvard University Press, 1994.

[2] Dixit, Avinash, and Barry Nalebuff. *The Art of Strategy: A Game Theorist's Guide to Success in Business and Life.* 1st ed. New York, NY: W.W. Norton & Company, Inc., 2008.

[3] Gibbons, Robert. *Game Theory for Applied Economists.* Princeton, NJ: Princeton University Press, 1992.

[4] Mas-Colell, Andrew, Michael Whinston, and Jerry Green. *Microeconomic Theory.* New York, NY: Oxford University Press, Inc., 1995.

[5] Osborne, Martin J. *An Introduction to Game Theory: International Edition.* New York, NY: Oxford University Press, 2009.

Chapter 6

Cooperation through Repeated Interaction

6.1 Motivating Example

We like to view the interaction of many economic agents as belonging to the class of collective action problems (whose most famous member is the Prisoners' Dilemma). Consider a game (Dixit and Nalebuff, 2008, pgs. 161-162) between Pepsi and Coke to determine the price of soft drinks, a game between McDonald's and Burger King to offer toys with kids' meals, or a game between Marlboro and Camel about the amount to invest in cigarette advertising.

In all cases, the Prisoners' Dilemma says that the equilibrium will provide lower profit for each firm compared to the (unattainable) case in which both firms cooperate. But do we observe the Prisoners' Dilemma result in reality? Marlboro and Camel are "saved" from playing the Prisoners' Dilemma due to the U.S. government ban on smoking advertisements. McDonald's and Burger King still try to outdo each other by offering toys with the kids' meals. But in terms of Pepsi, Coke, and the price of soft drinks, the two firms do not participate in price wars. Why is this? The answer is that the firms recognize that they are playing the Prisoners' Dilemma over and over again. In fact, as corporations are infinite-lived, then both Pepsi and Coke know that they are engaged in an infinitely repeated Prisoners' Dilemma. The question for this chapter will be to determine the conditions under which players cooperate (Pepsi and Coke do not engage in price wars) in an infinitely repeated Prisoners' Dilemma.

To fix numbers, let's recall the stage game of a standard Prisoners' Dilemma that we considered in Chapter 5.

Recall Figure 5.6 from Chapter 5. If this game is played once, the unique Nash equilibrium is (Fink, Fink). We can think of (Fink, Fink) as the "defection" outcome, because both players would ideally like to find some way to cooperate and obtain the Pareto optimal payoffs associated with (Mum, Mum). The Pareto optimal strategies (Mum, Mum) can be viewed as the "cooperation" outcome.

If Figure 5.6 is repeated a finite number of times, the analysis in the previous chapter dictates that the subgame perfect Nash equilibrium is such that the set of strategies is (Fink, Fink) in all rounds of the game.

With a game played over an infinite time horizon, we must specify how "patient" the players are. Logically, it does not make sense for a player to place the same weight on payoffs received today as the payoffs received 1000 rounds from now. Mathematically, patience is required or else the payoffs are undefined infinite sums. What we show in this chapter can be stated as follows: provided that players are "patient" enough, then the cooperation outcome (Mum, Mum) can be achieved in equilibrium.

6.2 Discounting Future Payoffs

For any one player in the game, let's define the discount factor δ, where $0 < \delta < 1$. The discount factor is used to define the total payoff, which is equal to the sum of the discounted payoffs in all rounds:

$$\text{total payoff} = \sum_{t=0}^{\infty} \delta^t \text{ (payoff in round } t\text{)}.$$

Without discounting (i.e., if $\delta = 1$), then the infinite sum given above would not be well-defined. Discounting means that if the player receives the cooperation payoff of 10 next period, then the payoff only has value $10 \cdot \delta$ in the current period. As the time between the current period and the receipt of the payoff increases, then the discounted present value of that payoff decreases.

As an example, let me find the value for the total payoff to a player that receives the cooperation payoff of 10 in all periods. As this corresponds to the "Mum" action in all periods, let me define the total payoff as M :

$$M = \sum_{t=0}^{\infty} \delta^t \cdot 10 = 10 + 10\delta + 10\delta^2 + ...$$

The value for M can be found by using the following algebraic trick. The number of future periods beginning in the current period is equal to the number of future periods beginning in the next period (this is the definition of the infinite time horizon). Given that the total payoff beginning in the current period is equal to M, I recognize that when I arrive in the next period, the total payoff will be equal to M as well. This means that we can write the total payoff in the following recursive fashion:

$$\begin{aligned} M &= 10 + 10\delta + 10\delta^2 + ... \\ &= 10 + \delta\left(10 + 10\delta + 10\delta^2 + ...\right) \\ &= 10 + \delta M. \end{aligned}$$

Collecting the terms involving M and solving yields:

$$\begin{aligned} (1-\delta)M &= 10 \\ M &= \frac{10}{1-\delta}. \end{aligned}$$

Thus, whenever a player has a constant payoff K in all periods, the total payoff is given by $\frac{K}{1-\delta}$.

6.3 Infinitely Repeated Prisoners' Dilemma

When we consider an infinitely repeated Prisoners' Dilemma (Gibbons, 1992, pgs. 89-91 and 95-96), the game is played between two players with identical discount factor δ such that $0 < \delta < 1$. With an infinite number of rounds, there are many subgame perfect Nash equilibrium of this game. We, however, are only interested in the subgame perfect Nash equilibrium in which players cooperate along the equilibrium path. Thus, the equilibrium path calls for all players to play "Mum" in all rounds. This is indicated in Figure 6.1 in the companion 'Figures' document.

To ensure that this is the equilibrium path, we must specify the complete strategies for both players. Recall that a strategy must include the actions taken both on and off the equilibrium path. The strategies must specify what the punishment will be for a player that deviates from the equilibrium path.

Consider any one round of the infinitely repeated Figure 5.6. If a player only cared about its payoff in that one round, then the player would deviate from the equilibrium path and choose "Fink." By choosing "Fink" when its opponent chooses "Mum," the player receives the highest payoff of 20. This payoff of 20 exceeds the equilibrium path payoff of 10. Call the difference (20-10)=10 the "immediate gain from deviation." The gain from deviation will be the same for the remainder of this section as we will always be considering the game in Figure 5.6.

For every deviation, a punishment must take place. The punishment will depend upon the strategies chosen by both players. This section considers three different strategies. Each of the strategies is symmetric, meaning that both players choose the same strategy. Each strategy contains a different "size of punishment." If the "immediate gain from deviation" is less than or equal to the "size of punishment," then the players will not find it optimal to deviate from the equilibrium path.

6.3.1 Grim Trigger

The first punishment strategy to consider is called the "grim trigger" strategy (Gibbons, 1992, pgs. 91-92 and 95-96). This strategy is symmetric for both players and defined as in the following table:

For each player	
Round 1:	Play "Mum"
:	
Round n:	Play "Mum" if (Mum, Mum) has always been played;
	Play "Fink" otherwise

Table 6.1

If both players follow the strategies in Table 6.1, the equilibrium path is (Mum, Mum) in all periods as shown in Figure 6.1. In order to verify that the strategies in Table 6.1 are in fact a subgame perfect Nash equilibrium, we need to verify that neither player has an incentive to deviate from the equilibrium path. The game is symmetric, so if player 1 does not have an incentive to deviate, then player 2 will not have an incentive to deviate. When analyzing if player 1 has an incentive to deviate, we hold fixed the actions of player 2 according to

Table 6.1.

Let's consider what the payoff for player 1 will be if player 1 deviates from the equilibrium path. As soon as player 1 deviates from the equilibrium path, the game enters an entirely new path. You can think of this as a "Back to the Future" break in the space-time continuum. The action by player 1 open up an entirely new future. This new future is called the "off the equilibrium" path. The "off the equilibrium" path is determined according to the punishment strategies as described in Table 6.1. The punishment in words states that if the other player chooses "Fink" one time, then I will punish that player by playing "Fink" in all future periods. Compare the equilibrium path and the "off the equilibrium" path in Figure 6.2.

As indicated in Figure 6.2, the instance of deviation occurs in the round when the strategies switch from (Mum, Mum) to (Fink, Mum). Recall that player 1 is the deviator and the best deviation is to choose "Fink" in this period of deviation as this provides player 1 with an immediate payoff of 20.

The equilibrium path payoff is equal to $M = \frac{10}{1-\delta}$ as shown in the previous section. I shall denote the deviation payoff with the letter D. Beginning in the instance of deviation, that payoff is equal to:

$$D = 20 + 5\delta + 5\delta^2 + 5\delta^3 + ...$$

Beginning in the current period, the total payoff from receiving 5 every period is $\frac{5}{1-\delta}$. Thus, if we begin to receive 5 every period beginning in the next period, then the total payoff is whatever we receive in the current period plus $\left(\frac{5}{1-\delta}\right)\delta$. Therefore,

$$D = 20 + \frac{5\delta}{1-\delta}.$$

No deviation means that the total payoff from the equilibrium path (which equals M) is greater than or equal to the total payoff from the deviation path (which equals D):

$$\begin{aligned} M &\geq D \\ \frac{10}{1-\delta} &\geq 20 + \frac{5\delta}{1-\delta} \\ 10 &\geq 20(1-\delta) + 5\delta \\ 10 &\geq 20 - 20\delta + 5\delta \\ 15\delta &\geq 10 \\ \delta &\geq \frac{2}{3}. \end{aligned}$$

With all of these infinitely repeated Prisoners' dilemma problems, we are looking for the lower bound $\underline{\delta}$ such that the equilibrium path is (Mum, Mum) in all periods for all discount factors $\delta \geq \underline{\delta}$. This makes perfect sense. For high values of δ (values of δ close to 1), players place a lot of weight on future round payoffs and thus receive a greater loss from the punishment. If you consider low values of δ (values of δ close to 0), players place low weight on the future and benefit greatly from an immediate payoff gain. Recall that deviation is not optimal when the punishment loss is greater than or equal to the immediate payoff gain. So deviation is not optimal for high values of δ. How high is high? That is the task facing us in these problems.

For the present exercise, we have finished solving the problem. The grim trigger strategy (as listed in Table 6.1) is a subgame perfect Nash equilibrium for the game in Figure 5.6 repeated an infinite number of times provided that the discount factor $\delta \geq \frac{2}{3}$.

6.3.2 Tit-for-Tat

The second punishment strategy to consider is called the "tit-for-tat" strategy (Gibbons, 1992, pgs. 2-4 and 224-231). The strategy is symmetric for both players and given in the following table:

For each player	
Round 1:	Play "Mum"
:	
Round n:	Play "Mum" if the other player played "Mum" in the previous round
	Play "Fink" if the other player played "Fink" in the previous round

Table 6.2

If the players follow the strategies in Table 6.2, the equilibrium path is (Mum, Mum) in all periods as in Figure 6.1. These strategies represent different punishments compared to the grim trigger strategy. To illustrate this, let's compare the equilibrium and the "off the equilibrium" paths when player 1 decides to deviate. This is illustrated in Figure 6.3.

With the tit-for-tat strategy, we see that each player will simply mimic the prior action of its opponent (that is what tit-for-tat means). With this strategy, we have extra work to do, because we need to verify that neither player has an incentive to deviate from the "off the equilibrium" path. With the grim trigger strategy, neither player would deviate from the "off the equilibrium" path, because (Fink, Fink) is the Nash equilibrium of the stage game.

Here, it is not so clear that both players would stick to the alternating actions specified in the "off the equilibrium" path. In particular (again, we consider player 1 as the deviator), would player 1 prefer (i) the outcome (Mum, Fink) and the immediate payoff of 0 along the "off the equilibrium" path or (ii) the outcome (Fink, Fink) and the payoff of 5 in all future rounds? This latter option will be referred to as the "off the off the equilibrium" path.

As before, a deviation from the "off the equilibrium" path opens up a new future for the game. The "off the off the equilibrium" path, along with the "off the equilibrium" path and equilibrium path, are depicted in Figure 6.4 in the companion 'Figures' document.

Notice in the figure that at the second instance of deviation, player 1 chooses "Fink" instead of "Mum," so the outcome is not (Mum, Fink) as required by the "off the equilibrium" path, but rather (Fink, Fink). Given (Fink, Fink), in that second instance of deviation, the strategies in Table 6.2 dictate that (Fink, Fink) will be played in all rounds from that point forward.

Therefore, to verify that the equilibrium path is (Mum, Mum) in all rounds, we must verify two things: (i) player 1 does not move from the equilibrium path to the "off the equilibrium" path and (ii) player 1 does not move from the "off the equilibrium" path to the "off the off the equilibrium" path. Define D as the payoff along the "off the equilibrium" path beginning at the first instance of deviation, O as the payoff along the "off the equilibrium" path beginning at the second instance of deviation, and D^* as the payoff along the "off the off the equilibrium" path beginning at the second instance of deviation. Repeating steps (i) and (ii) from above, I must show that (i) $M \geq D$ and (ii) $O \geq D^*$. The equations for (D, O, D^*) are given in the table below, where I can simplify the expression for D^* using the result from the previous section.

$$D = 20 + 0 \cdot \delta + 20 \cdot \delta^2 + 0 \cdot \delta^3 + ...$$
$$O = 0 + 20 \cdot \delta + 0 \cdot \delta^2 + 20 \cdot \delta^3 + ...$$
$$D^* = 5 + 5 \cdot \delta + 5 \cdot \delta^2 + ... = \tfrac{5}{1-\delta}$$

Table 6.3

Let's use an algebraic trick as in the previous section to simplify the expressions for D and O. The total payoff D can be written recursively since the values repeat every two periods:

$$D = 20 + 0 \cdot \delta + \delta^2 \cdot D.$$

Collecting terms involving D and solving, I obtain:

$$\begin{aligned}\left(1-\delta^2\right)D &= 20+0\cdot\delta\\ (1+\delta)(1-\delta)D &= 20+0\cdot\delta\\ D &= \frac{20}{(1+\delta)(1-\delta)}.\end{aligned}$$

Using the same approach for O, I obtain:

$$O=\frac{20\delta}{(1+\delta)(1-\delta)}.$$

As mentioned above, I need to show (i) $M \geq D$ and (ii) $O \geq D^*$.

1. The inequality $M \geq D$ implies

$$\begin{aligned}\frac{10}{1-\delta} &\geq \frac{20}{(1+\delta)(1-\delta)}\\ 10(1+\delta) &\geq 20\\ 10\delta &\geq 10\\ \delta &\geq 1.\end{aligned}$$

2. The inequality $O \geq D^*$ implies

$$\begin{aligned}\frac{20\delta}{(1+\delta)(1-\delta)} &\geq \frac{5}{1-\delta}\\ 20\delta &\geq 5(1+\delta)\\ 15\delta &\geq 5\\ \delta &\geq \frac{1}{3}.\end{aligned}$$

In conclusion, both (1) and (2) hold when $\delta \geq \max\left\{1, \frac{1}{3}\right\} = 1$. Thus, there does not exist any values of $\delta < 1$ for which the tit-for-tat strategies in Table 6.2 are a subgame perfect Nash equilibrium. In other words, a player always has an incentive to deviate from the equilibrium path and follow the "off the equilibrium" path.

This is certainly not the case for all examples. In fact, for some examples, the punishment strategy contained in tit-for-tat is stronger than the punishment of the grim trigger. This means that the lower bound for grim trigger ($\underline{\delta_G}$ such that Table 6.1 are a subgame perfect

Nash equilibrium for all $\delta \geq \underline{\delta_G}$) has no relation to the lower bound for tit-for-tat ($\underline{\delta_T}$ such that Table 6.2 are a subgame perfect Nash equilibrium for all $\delta \geq \underline{\delta_T}$): $\underline{\delta_G} > \underline{\delta_T}$, $\underline{\delta_G} = \underline{\delta_T}$, or $\underline{\delta_G} < \underline{\delta_T}$. In the case of the infinitely repeated game in Figure 5.6, $\underline{\delta_G} = \frac{2}{3} < \underline{\delta_T} = 1$.

6.3.3 k-Period Punishment

The third punishment strategy to consider is called the "k-period punishment" strategy. This strategy is symmetric for both players and defined in the following table:

For each player	
Round 1:	Play "Mum"
:	
Round n:	Play "Mum" if (Mum, Mum) has always been played
	or (Fink, Fink) has been played in k previous rounds
	Play "Fink" otherwise

Table 6.4

If both players follow the strategies in Table 6.4, the equilibrium path is (Mum, Mum) in all rounds as shown in Figure 6.1. These strategies represent different punishments compared to the previous two strategies (Table 6.1 and Table 6.2). To illustrate this, let's compare the equilibrium path and the "off the equilibrium" path. These are shown in Figure 6.5. Throughout this chapter, we have considered player 1 as the deviating player and continue to do so in this figure.

As can be seen from Figure 6.5, the "off the equilibrium" path lasts for $k+1$ rounds: one round for the deviation by player 1 and k rounds for punishment. After those $k+1$ rounds, the "off the equilibrium" path reconnects with the equilibrium path. The two paths only differ in terms of these $k+1$ rounds.

The total payoff from the equilibrium path remains $M = \frac{10}{1-\delta}$. The deviation payoff is given by:

$$D = 20 + 5\delta + 5\delta^2 + ... + 5\delta^k + \frac{10\delta^{k+1}}{1-\delta}.$$

The last payoff term $\frac{10\delta^{k+1}}{1-\delta}$ is the payoff on the equilibrium path beginning $k+1$ periods from the current period (discounted back to the current period). The middle payoff terms

$$5\delta + 5\delta^2 + ... + 5\delta^k = \frac{5\delta}{1-\delta} - \frac{5\delta^{k+1}}{1-\delta}$$

as this is the difference between a payoff of 5 in all periods begining next period ($\frac{5\delta}{1-\delta}$) and a payoff of 5 in all periods begining $k+1$ periods from now ($\frac{5\delta^{k+1}}{1-\delta}$).

Player 1 remains on the equilibrium path provided that $M \geq D$:

$$\begin{aligned} \frac{10}{1-\delta} &\geq 20 + \frac{5\delta}{1-\delta} - \frac{5\delta^{k+1}}{1-\delta} + \frac{10\delta^{k+1}}{1-\delta}. \\ 10 &\geq 20\left(1-\delta\right) + 5\delta + 5\delta^{k+1}. \\ 15\delta - 5\delta^{k+1} &\geq 10. \end{aligned}$$

For instance, when $k = 1$ (the weakest of the k-period punishments), then

$$\begin{aligned} 15\delta - 5\delta^{2} &\geq 10. \\ \delta^{2} - 3\delta + 2 &\leq 0. \\ \left(\delta - 1\right)\left(\delta - 2\right) &\leq 0. \end{aligned}$$

This implies that $1 \leq \delta \leq 2$. Thus, there does not exist a discount factor $\delta < 1$ such that the strategies listed in Table 6.4 for $k = 1$ are a subgame perfect Nash equilibrium.

For values of $k > 1$, we are tasked with solving higher order polynomial systems. What about when $k \to \infty$? This corresponds to an infinite punishment. With $k \to \infty$, then $5\delta^{k+1} \to 0$, so the equation reduces to:

$$\begin{aligned} 15\delta &\geq 10. \\ \delta &\geq \frac{2}{3}. \end{aligned}$$

In words, for all discount factors $\delta \geq \frac{2}{3}$, the strategies in Table 6.4 as $k \to \infty$ are a subgame perfect Nash equilibrium. This is the exact same answer obtained for the grim trigger strategy. This makes perfect sense, since the grim trigger strategy is nothing other than an infinite period punishment (hence the name "grim").

Thus, of all the k-period punishments, grim trigger is the strongest. This means that the lower bound $\underline{\delta_G}$ is strictly lower than the lower bound for any k-period punishment ($\underline{\delta_k}$ such that Table 6.4 are a subgame perfect Nash equilibrium for all $\delta \geq \underline{\delta_k}$): $\underline{\delta_k} > \underline{\delta_G}$.

6.4 Rubinstein Bargaining

The Rubinstein bargaining game (Mas-Colell et al., 1995, pgs. 298-299, 303 and Gibbons, 1992, pgs. 70-71, 131, 136-138) is an alternating offer game played between two players over an infinite time horizon. The players are bargaining over \$1. Player 1 makes the offers in the odd-numbered rounds and has discount factor δ_1 such that $0 < \delta_1 < 1$. Upon receiving an offer from player 1 in an odd-numbered round, player 2 can either accept or reject. If the offer is accepted, the game ends. If the offer is rejected, the game moves to the next round, which is an even-numbered round.

Player 2 makes the offers in the even-numbered rounds and has discount factor δ_2 such that $0 < \delta_2 < 1$. Upon receiving an offer from player 2 in an even-numbered round, player 1 can either accept or reject. If the offer is accepted, the game ends. If the offer is rejected, the game continues.

In any bargaining game with discounting, the players' optimal strategies will allow them to reach agreement in the first round. Otherwise, resources are being wasted as both players receive a lower payoff (due to discounting) by reaching agreement in a later round.

Suppose player 1 makes an offer of $(x, 1-x)$ in the first round. Here, $x \geq 0$ is the amount that player 1 receives and $1-x \geq 0$ is the amount that player 2 receives. The total must sum to 1. Define x^* as the maximum payoff that player 1 can receive in the first round. This is the payoff x^* such that player 2 just barely accepts $1-x^*$. When does player 2 accept $1-x^*$? Define y^* as the maximum payoff that player 2 can receive in the second round. Player 2 accepts the offer of $1-x$ in the first round when:

$$1 - x \geq \delta_2 \cdot y^*.$$

Do not forget the discounting. The best strategy for player 1 is to offer $1-x^*$ so that player 2 just barely accepts:

$$1 - x^* = \delta_2 \cdot y^*.$$

This is the highest payoff for player 1.

The payoff y^* is the maximum payoff for player 2 in the second round, where player 1 accepts its payoff of $1-y^*$. Player 1 accepts $1-y$ in the second round if it is at least as high as what it can expect to receive in the third round. If x^* is the maximum payoff for player 1 in the first round, then upon arriving in the third round, the maximum payoff is still x^*.

Thus, player 1 accepts an offer $1 - y$ in the second round when:

$$1 - y \geq \delta_1 \cdot x^*.$$

The best offer for player 2 to make in the second round is such that:

$$1 - y^* = \delta_1 \cdot x^*.$$

We now have only to solve two equations in two unknowns.

$$\begin{aligned}
1 - y^* &= \delta_1 \left(1 - \delta_2 \cdot y^*\right). \\
1 - y^* &= \delta_1 - \delta_1 \delta_2 \cdot y^*. \\
1 - \delta_1 &= \left(1 - \delta_1 \delta_2\right) y^* \\
y^* &= \frac{1 - \delta_1}{1 - \delta_1 \delta_2}.
\end{aligned}$$

Given the equation for y^*, then $1 - x^* = \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2}$, meaning that x^* is given by:

$$\begin{aligned}
x^* &= 1 - \frac{\delta_2 \left(1 - \delta_1\right)}{1 - \delta_1 \delta_2} \\
&= \frac{1 - \delta_2}{1 - \delta_1 \delta_2}.
\end{aligned}$$

This means that the subgame perfect Nash equilibrium path calls for player 1 to offer $\left(\frac{1-\delta_2}{1-\delta_1\delta_2}, \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2}\right)$ in the first round and for player 2 to accept. The game ends at this point.

If the two players have identical discount factors $\delta_1 = \delta_2 = \delta$, then the first round offer simplifies to $\left(\frac{1}{1+\delta}, \frac{\delta}{1+\delta}\right)$ after using the algebraic fact $\left(1 - \delta^2\right) = \left(1 + \delta\right)\left(1 - \delta\right)$. If the common discount factor approaches 1, i.e., $\delta \to 1$ (meaning that the players become more patient), then the bargaining split approaches 50:50:

$$\left(\frac{1}{1+\delta}, \frac{\delta}{1+\delta}\right) \to (0.5, 0.5) \quad \text{as } \delta \to 1.$$

6.5 Exercises

1. Consider a infinitely repeated prisoners' dilemma with each stage composed of the following simultaneous game:

		Suspect 2 Mum	Fink
Suspect 1	Mum	15, 15	5, 20
	Fink	20, 5	10, 10

Figure 6.6

Both suspects have a discount factor of $\delta \in (0, 1)$. What is the lower bound on the discount factor such that (Mum, Mum) is played in all stages, given the "Grim trigger" strategy?

2. This question analyzes the US soft drink market in two parts.

(a) Pepsi and Coke are in competition in the soft drink market. The two firms play a game that is repeated over an infinite future. In each stage of the game, each firm simultaneously chooses either to sell at a low price or a high price. The pricing decisions of both firms determine the profits for both firms. The game played in each stage is depicted below.

		Coke Low price	High price
Pepsi	Low price	20, 20	50, 15
	High price	15, 50	30, 30

Figure 6.7

Both firms have a discount factor δ where $0 < \delta < 1$. Tell me all values of δ so that the "Grim trigger" strategy is a subgame perfect Nash equilibrium. For your convenience, the "Grim trigger" strategy is given by:

For each firm	
Stage 1:	Play "High price"
:	
Stage n:	Play "High price" if (High,High) has always been played; Play "Low price" otherwise
:	

Table 6.5

(b) Now consider that Pepsi and Coke both have a discount factor equal to $\delta = 0.5$. How much is Pepsi willing to pay in order to add a new action called "Super low price" to the set of actions in each stage? With the new action "Super low price" added, the game played in each stage is depicted below.

		Coke		
		Super low price	Low price	High price
	Super low price	5, 5	0, 0	0, 0
Pepsi	Low price	0, 0	20, 20	50, 15
	High price	0, 0	15, 50	30, 30

Figure 6.8

3. Pepsi and Coke are in competition in the soft drink market. The two firms play a game that is repeated over an infinite future. In each stage of the game, each firm simultaneously chooses either to sell at a low price or a high price. The pricing decisions of both firms determine the profits for both firms. The game played in each stage is depicted below.

		Coke	
		Low price	High price
Pepsi	Low price	20, 20	40, 5
	High price	5, 40	30, 30

.

Figure 6.9

Both firms have a discount factor δ where $0 < \delta < 1$. Tell me all values of δ so that the

"Tit-for-tat" strategy is a subgame perfect Nash equilibrium. For your convenience, the "Tit-for-tat" strategy is given by:

For each firm	
Stage 1:	Play "High price"
:	
Stage n:	Play "High price" if other firm played "High price" in Stage n-1 Play "Low price" if other firm played "Low price" in Stage n-1
:	

Table 6.6

4. Consider an infinite horizon bargaining model with two players (Rubinstein bargaining). Player 1 makes offers in the odd-numbered rounds and Player 2 makes offers in the even-numbered rounds. After each player makes an offer, the other player must decide to accept or reject. If the offer is accepted, the game ends. If the offer is rejected, the game proceeds to the next round. Both players have the discount factor $\delta = \frac{2}{3}$. The players are bargaining over \$100.

   What are the equilibrium payoffs (in dollars) for both players?

5. Consider the Rubinstein bargaining model with three agents. As in class, the bargaining occurs over an infinite time horizon. Agent 1 makes offers in periods 1, 4, 7, ...; Agent 2 makes offers in periods 2, 5, 8, ...; and Agent 3 makes offers in periods 3, 6, 9, ... The agents are bargaining over \$1. If both agents that receive an offer decide to accept that offer, then the bargaining ends. Let all three agents have the same discount factor $\delta : 0 < \delta < 1$.

   What are the equilibrium payoffs for all three agents?

Bibliography

[1] Baird, Douglas, Robert Gertner, and Randal Picker. *Game Theory and the Law.* Cambridge, MA, and London, England: Harvard University Press, 1994.

[2] Dixit, Avinash, and Barry Nalebuff. *The Art of Strategy: A Game Theorist's Guide to Success in Business and Life.* 1st ed. New York, NY: W.W. Norton & Company, Inc., 2008.

[3] Gibbons, Robert. *Game Theory for Applied Economists.* Princeton, NJ: Princeton University Press, 1992.

[4] Mas-Colell, Andrew, Michael Whinston, and Jerry Green. *Microeconomic Theory.* New York, NY: Oxford University Press, Inc., 1995.

[5] Osborne, Martin J. *An Introduction to Game Theory: International Edition.* New York, NY: Oxford University Press, 2009.

Part III

Games of Incomplete Information

Chapter 7

The Theory of Auctions

7.1 Motivating Example

Suppose that you are trying to win an item in an e-Bay auction (Dixit and Nalebuff, 2008, pgs. 301-302, 310-322). You lead a busy life, so although the auction lasts for the next 7 days, you do not have time to continually check the highest bid. Thus, you decide to use the e-Bay option to select a "maximum bid." What this option in e-Bay does is the following. Suppose that you currently have the highest bid for an object at \$2, but you have listed your maximum bid at \$10. The bid for the object will remain \$2. If between now and the end of the auction someone bids higher than your bid of \$2, then e-Bay will automatically increase your bid to be equal to the new bid plus the size of the increment. The size of the increment is typically \$0.10 for most small-value purchases. e-Bay will place this new bid for you so long as the new bid is below your maximum bid of \$10.

Suppose that all bidders for an object did not take the time to observe the current highest bid, but only set a value for their maximum bid. If you are one of the bidders in this competition, what value would you set as your maximum bid? Would you set the value that you actually want to pay for the object, or the lowest value that you think can guarantee that you win the auction?

We will see that the stylized framework for e-Bay auctions is reminiscient of a second price auction.

7.2 Types of Auctions

An auction is used to determine the allocation of a specific object to one member of a group. Each member of the group has a value for the object. Let's denote the value for player i as $v_i \geq 0$. The value is a nonnegative parameter of the model. It incorporates all aspects of winning the object in the auction (the thrill of winning, the utility of the object, the absence of loss aversion, etc.).

The type of auctions that we consider in this chapter are sealed-bid auctions. This means that each player writes down its bid for the object, and does not observe the bids of the other players. Once all bids are collected, they are revealed by the auctioneer and a winner is determined. Let's denote the bid placed by player i as b_i. As the bids are made in secret, this is equivalent to a simultaneous-action game (as in chapters 1-3), so our solution concept will be Nash equilibrium.

After all the bids are collected, the winner of the auction is the player that placed the highest bid. In the event of a tie, a fair lottery is used to determine the winner. For nearly the entire chapter, we can appreciate the theory of auctions by considering an auction between only two bidders. In the case of two bidders, a tie is broken by a coin flip.

The payoff for each of the bidders $i = 1, 2$ for each of the possible outcomes is given in the following table:

$$\begin{array}{rl} \text{payoff} = 0 & \text{if player } i \text{ loses the auction} \\ \text{payoff} = \frac{1}{2}(v_i - price) & \text{if player } i \text{ ties the auction} \\ \text{payoff} = (v_i - price) & \text{if player } i \text{ wins the auction} \end{array}$$

Table 7.1

The price that the winner has to pay will differ depending upon whether we consider a second price auction (next section) or a first price auction (two sections from now).

7.3 Second Price Auction

Auctions can be conducted under complete information or under incomplete information. Under **complete information**, each of the bidders knows the value that the other bidder has for the object. This may seem a bit unrealistic, but this is the assumption that we have been using for the first 6 chapters: all players know the payoffs of all other players.

Under **incomplete information**, each of the bidders knows its own value for the object, but does not know the value of the other bidder. The only thing that a bidder knows about the other bidder's value is from which distribution it is drawn. For simplicity, we typically consider that the values are drawn from the uniform distribution from 0 to 1, $v_i\tilde{} Unif[0,1]$.

In a second price auction (Mas-Colell et al., 1995, pgs. 262, 866 and Nalebuff, 2008, pgs. 305-307), the player with the highest bid is the winner (ties are broken with a coin flip). The price that the winner then pays for the object is equal to the bid of the *other* player. That is, the price is not equal to the highest bid (the bid from the winner), but is equal to the second-highest bid (the bid from her opponent). Hence, the name of the auction is the "second-price" auction.

Clearly you would suspect that the outcome of the game will differ depending upon whether or not the game is played under complete information or under incomplete information. What I show in this section is that the dominant strategy in a second price auction is for a player to set her bid equal to her value: $b_i = v_i$. Recall that a dominant strategy (way back in Chapter 1) is a strategy that provides the (weakly) highest payoff no matter what strategy is chosen by the other player. If a strategy is dominant, then one bidder does not need to consider the payoffs and strategies of the other bidder. All a player needs to do is place a bid equal to this dominant strategy.

Suppose that you are player 1 with a value of v_1 for the object. I claim that it is a dominant strategy to bid $b_1 = v_1$. Let's consider all possible strategies that the other bidder can choose.

1. If $b_2 < v_1$, then you have won the auction with your bid $b_1 = v_1$. The price that you pay for the object is the second price $b_2 < b_1 = v_1$. This means that your payoff is equal to $v_1 - b_2 > 0$. Is there any other strategy $b_1 \neq v_1$ that can yield a higher payoff? If $b_1 < b_2$, then you lose the auction, with payoff $= 0$. If $b_1 = b_2$, then you tie the auction, with payoff $= \frac{1}{2}(v_1 - b_2)$. If $b_1 > b_2$, then you win the auction, with payoff $= (v_1 - b_2)$. Thus, all the possible payoffs that you can receive are less than or equal to the payoff received from $b_1 = v_1$.

2. If $b_2 = v_1$, then you tie the auction with your bid $b_1 = v_1$. The price that you pay for the object is the second price $b_2 = b_1 = v_1$. This means that your payoff is equal to 0. Is there any other strategy $b_1 \neq v_1$ that can yield a higher payoff? If $b_1 < b_2$, then you lose the auction, with payoff $= 0$. If $b_1 > b_2$, then you win the auction, with payoff $= (v_1 - b_2) = 0$. Thus, all the possible payoffs that you can receive are equal to the payoff received from $b_1 = v_1$.

3. If $b_2 > v_1$, then you have lost the auction with your bid $b_1 = v_1$. The payoff is equal to 0. Is there any other strategy $b_1 \neq v_1$ that can yield a higher payoff? If $b_1 < b_2$, then you lose the auction, with payoff $= 0$. If $b_1 = b_2$, then you tie the auction, with payoff $= \frac{1}{2}(v_1 - b_2) < 0$. If $b_1 > b_2$, then you win the auction, with payoff $= (v_1 - b_2) < 0$. Thus, all the possible payoffs that you can receive are less than or equal to the payoff received from $b_1 = v_1$.

This verifies that no matter what bid player 2 makes, your highest payoff as player 1 is achieved by setting $b_1 = v_1$. The exact same argument works for any finite number of bidders. Suppose that there are n total bidders numbered 1 through n and you are bidder 1. To see that $b_1 = v_1$ is the dominant strategy, all you have to do is rewrite the above three steps where you replace b_2 with $\max\{b_2, ..., b_n\}$.

Suppose that the increment in an e-Bay auction is arbitrarily small. An auction in which the players submit their maximum bids is equivalent to a second price auction. Consider that if you have the highest maximum bid, then the automatic bidding that e-Bay follows with the maximum bid has the following results: (i) you win the object and (ii) the bid that you place to win the object is equal to the maximum bid of the next closest competitor plus the increment. If the increment is arbitrarily small, then the price of the object is simply equal to the maximum bid of the second place bidder.

Thus, in an e-Bay auction in which the players submit their maximum bids, you don't need to predict who you are bidding against and what their bidding strategies might be. Whatever you are willing to pay for an object (your true value), you simply set your maximum bid equal to this true value. If you win the object, then you likely will be a paying a price below this maximum bid (equal to the second highest bid). If you lose the object, then you would not have wanted to pay the winning bid anyway.

The second price auction has the special property that bidders have the incentive to set a bid equal to their true values. Exercise 1 asks you to verify if a similar property holds for a third price auction.

7.4 First Price Auction

We now arrive at the most common type of auction: the first price auction (Mas-Colell et al., 1995, pgs. 265, 865-866 and Dixit and Nalebuff, 2008, pgs. 315-320). In the first price auction, the winner of the auction is still the player that submits the highest bid (coin flip if tie), but now the price that the winner must pay for the object equals its own bid. Thus, the price is the highest or the "first price" bid.

With the first price auction, I claim that a bidder will never choose to place a bid above its value. Such a bid is weakly dominated. If bidder 1 were to win the auction with the bid $b_1 > v_1$, then the payoff is equal to $v_1 - b_1 < 0$. If bidder 1 were to lose the auction with that bid, then the payoff is 0. In both cases, bidder 1 is better off (strictly so in the former case) by bidding $b_1 \leq v_1$.

We consider two information structures: complete information and incomplete information.

7.4.1 Complete Information

The case with complete information is straightforward to solve. Suppose that both bidders have the same value $v_1 = v_2$. We know that the players bid less than or equal to their true value, so $b_1 \leq v_1$ and $b_2 \leq v_2$. Suppose that $b_1 < b_2 < v_1 = v_2$. Then bidder 1 has an incentive to outbid bidder 2 by placing a bid between b_2 and $v_1 = v_2$. Such an action would allow bidder 1 to win the auction and receive a strictly positive payoff. Bidder 2 would then respond by placing a bid between the new bid of player 1 and $v_1 = v_2$. The Nash equilibrium (the point at rest) in which neither bidder wishes to change its strategy is given when $b_1 = b_2 = v_1 = v_2$. The payoff for both bidders is 0.

Suppose now that bidder 1 has a larger value compared to bidder 2: $v_1 > v_2$. We know that $b_2 \leq v_2$, so we set $b_2 = v_2$ and see what the best response by bidder 1 is. Bidder 1 only seeks to win the auctions and wants to win with the lowest bid possible. Thus, bidder 1 sets $b_1 = v_2 + 0.0000001$, wins the auctions, and has payoff approximately equal to $v_1 - b_1 = v_1 - v_2 > 0$.

7.4.2 Incomplete Information

The more interesting case to consider is when the auction is conducted under incomplete information. With incomplete information, each bidder knows its own value, but not the value of the other bidder. All that is known about the other bidder's value is the distribution from which it is drawn.

The original game

The original game specifies that both bidders' values are drawn from the uniform distribution between 0 and 1: $v_i \tilde{} Unif[0,1]$. This is known by both bidders. Obviously, once a value is drawn, a bidder knows its own value. It does not observe the value drawn by its opponent.

The uniform distribution has the nice property: if $v_2 \tilde{} Unif[0,1]$, then $\Pr(v_2 \leq k) = k$ for any $0 \leq k \leq 1$. That is, the probability that your opponent has drawn a value less than or equal to k is simply k. Before drawing their random values, the two bidders are identical as they draw from the same distribution. Thus, they should have identical strategies. A strategy for a bidder is a function from the drawn value to the bid that should be placed. We will assume that this function is linear. Thus, the strategies that both of the bidders employ in equilibrium will be:

$$b_i = av_i + c \text{ for } i = 1, 2.$$

Here, the coefficients (a, c) are unknown and need to be determined.

Let's solve for the payoff maximization problem of bidder 1. This player takes the strategy $b_2 = av_2 + c$ as given. The possible payoffs for bidder 1 are as follows:

$$\begin{array}{ll} \text{payoff} = v_1 - b_1 & \text{if } b_1 > av_2 + c \\ \text{payoff} = \frac{1}{2}(v_1 - b_1) & \text{if } b_1 = av_2 + c \\ \text{payoff} = 0 & \text{if } b_1 < av_2 + c \end{array}$$

Table 7.2

Given that the bids belong to the continuum, then the probability that $b_1 = av_2 + c$ is equal to 0. Thus, the payoff for bidder 1 is given by:

$$\text{payoff} = (v_1 - b_1) \Pr(b_1 > av_2 + c).$$

Rearranging the probability, the payoff is updated as:

$$\begin{aligned} \text{payoff} &= (v_1 - b_1) \Pr\left(v_2 < \frac{b_1 - c}{a}\right) \\ &= (v_1 - b_1)\left(\frac{b_1 - c}{a}\right). \end{aligned}$$

Let's maximize the payoff for bidder 1 by taking the first order condition with respect to b_1 :

$$-\left(\frac{b_1 - c}{a}\right) + \frac{(v_1 - b_1)}{a} = 0.$$

Solve this equation for b_1 yields the strategy:

$$\begin{aligned} 2b_1 &= v_1 + c. \\ b_1 &= \frac{v_1 + c}{2}. \end{aligned}$$

Let's compare the strategy to the original linear form $b_1 = av_1 + c$. These functions must be the same and this occurs when $a = \frac{1}{2}$ and $c = \frac{c}{2}$. The only value that satisfies $c = \frac{c}{2}$ is $c = 0$. Thus, the equilibrium strategy for bidder 1 is $b_1 = \frac{v_1}{2}$. As both bidders have the same strategies, then bidder 2's strategy is given by $b_2 = \frac{v_2}{2}$.

One nice property about this equilibrium is that the bidder with the highest valuation will always win the auction.

Scaling to the original game

Suppose that both bidders' values are drawn from the uniform distribution between A and B: $v_i \tilde{} Unif[A, B]$, where $0 < A < B$. Let's refer to the game corresponding to the $Unif[A, B]$ distribution as the true game and the game corresponding to the $Unif[0, 1]$ distribution as the updated game (i.e., the original game).

Suppose that player 1 has a value for the object in the true game equal to v_1. To get from the true game to the updated game, we have to subtract A and divide by $B - A$. Thus, the value for player 1 in the updated game is equal to $\hat{v}_1 = \frac{v_1 - A}{B - A}$. In the updated game, we know what the optimal bid should be (from the previous subsection). The optimal bid should be $\hat{b}_1 = \frac{\hat{v}_1}{2} = \frac{v_1 - A}{2(B-A)}$. To get from the updated game back to the true game, we must multiply by $B - A$ and add A (the reverse of the operations performed above). Thus, $b_1 = \hat{b}_1 (B - A) + A$.

In sum, we see that the optimal bid should be equal to:

$$\begin{aligned} b_1 &= \frac{v_1 - A}{2(B - A)} (B - A) + A \\ &= \frac{v_1 - A}{2} + A \\ &= \frac{v_1 + A}{2}. \end{aligned}$$

This is simply equal to the average between the lower bound A and the true value v_1.

Bidding with a reservation price

Suppose that both bidders' values are randomly drawn from $Unif[0,1]$. In this version of the game, a reservation price (RP) is in place. This simply means that the auction winner must have the highest bid and that bid must be greater than or equal to RP.

Let's consider the payoff maximization problem for bidder 1. Bidder 2 is assumed to bid according to the strategy $b_2 = av_2 + c$, where v_2 is the value for bidder 2. The possible payoffs for bidder 1 are as follows:

$\text{payoff} = v_1 - b_1$	if $b_1 > \max\{av_2 + c, RP\}$
$\text{payoff} = \frac{1}{2}(v_1 - b_1)$	if $b_1 = av_2 + c \geq RP$
$\text{payoff} = 0$	if $b_1 < \max\{av_2 + c, RP\}$

Table 7.3

Thus, if $v_1 < RP$, then it does not matter what your bid is, because you never want to place a bid above the reservation price RP. Thus, the bidding strategies are only valid for $v_1 \geq RP$. The strategies for both bidders are given by the linear equation:

$$b_i = av_i + c \text{ for } v_i \geq RP, \text{ where } i = 1, 2.$$

Let's consider the optimal bids for player 1 under the scenario that $v_1 \geq RP$ (under the scenario that $v_1 < RP$ then the player will never place a bid that meets the reservation price requirement). There are two cases to consider:

1. With probability $\Pr(v_2 < RP)$, player 2 will not be able to place a bid greater than or equal to RP. Thus, player 1 can win the auction with a bid equal to $b_1 = RP$. She can't win the auction with a bid lower than RP due to the reservation price requirement. So the optimal bid in this scenario is $b_1 = RP$.

2. With probability $\Pr(v_2 \geq RP)$, players 1 and 2 will be competing in a game in which valuations are drawn from $Unif[RP, 1]$. As seen from the previous subsection, after scaling the $Unif[RP, 1]$ game down to $Unif[0, 1]$, the best strategy for player 1 is given by $b_1 = \frac{1}{2}RP + \frac{1}{2}v_1$.

Players in this game maximize expected payoffs. If the optimal strategy is $b_1 = RP$ for the first scenario (with probability $\Pr(v_2 < RP) = RP$) and $b_1 = \frac{1}{2}RP + \frac{1}{2}v_1$ for the second scenario (with probability $\Pr(v_2 \geq RP) = 1 - RP$), then the optimal strategy over all possibilities is given by the expectation:

$$\begin{aligned} b_1 &= \Pr(v_2 < RP)\, RP + \Pr(v_2 \geq RP)\left(\frac{1}{2}RP + \frac{1}{2}v_1\right) \\ &= (RP)^2 + (1 - RP)\left(\frac{1}{2}RP + \frac{1}{2}v_1\right). \end{aligned}$$

More than two bidders

Let's consider a game with three bidders whose values are randomly drawn from the $Unif[0,1]$ distribution. The strategies that the bidders employ in equilibrium will be linear:

$$b_i = av_i + c \quad \text{for } i = 1, 2, 3.$$

The values of coefficients (a, c) are unknown and need to be determined.

Let's solve for the payoff maximization problem of bidder 1. This player takes the strategies $b_2 = av_2 + c$ and $b_3 = av_3 + c$ as given. The possible payoffs for bidder 1 are as follows:

$$\begin{array}{cl} \text{payoff} = v_1 - b_1 & \text{if } b_1 > \max\{av_2 + c, av_3 + c\} \\ \text{payoff} = \frac{1}{2}(v_1 - b_1) & \text{if } b_1 = \max\{av_2 + c, av_3 + c\} \\ \text{payoff} = 0 & \text{if } b_1 < \max\{av_2 + c, av_3 + c\} \end{array}$$

Table 7.5

Given that the bids belong to the continuum, then the probability that $b_1 = \max\{av_2 + c, av_3 + c\}$ is equal to 0. The payoff for bidder 1 is given by:

$$(v_1 - b_1)\Pr(b_1 > \max\{av_2 + c, av_3 + c\}).$$

The probabilities are independent, so we know that

$$\Pr(b_1 > \max\{av_2 + c, av_3 + c\}) = \Pr(b_1 > av_2 + c)\Pr(b_1 > av_3 + c).$$

Rearranging the probabilities, the payoff is updated as:

$$
\begin{aligned}
& (v_1 - b_1)\Pr\left(v_2 < \frac{b_1 - c}{a}\right)\Pr\left(v_3 < \frac{b_1 - c}{a}\right) \\
= \ & (v_1 - b_1)\left(\frac{b_1 - c}{a}\right)^2 .
\end{aligned}
$$

Let's maximize the payoff for bidder 1 by taking the first order conditions with respect to b_1 :

$$
-\left(\frac{b_1 - c}{a}\right)^2 + 2\left(v_1 - b_1\right)\left(\frac{b_1 - c}{a}\right)\frac{1}{a} = 0.
$$

From this equation, we can cancel $\left(\frac{b_1 - c}{a}\right)\frac{1}{a}$. We are left with:

$$
-\left(b_1 - c\right) + 2\left(v_1 - b_1\right) = 0.
$$

Solving this equation for b_1 yields the strategy:

$$
\begin{aligned}
3b_1 &= 2v_1 + c \\
b_1 &= \frac{2}{3}v_1 + \frac{c}{3}.
\end{aligned}
$$

Matching this equation to the original linear form $b_1 = av_1 + c$, we see that $a = \frac{2}{3}$ and $c = 0$.

Thus, the strategies for all three bidders are given by;

$$
b_i = \frac{2}{3}v_i \quad \text{for } i = 1, 2, 3.
$$

Exercise 4 asks you to extend this idea to consider the optimal bidding strategy in an auction with a total of n bidders (where $n > 3$). What is your guess for this optimal bidding strategy?

7.4.3 Classroom Exercise 8: Auctions

Thus far we have only considered scenarios in which the values for an object are drawn from a uniform distribution. In reality, we might think that values are often times more accurately described as belong to a discrete distribution. Let's consider the following game with a discrete distribution.

The game contains 2 standard decks of 52 playing cards. To refresh your memory about a deck of cards, the 52 cards include 13 numbers beginning with Ace (which equals 1) and ending with King. For each of these 13 numbers, there are four suits: Clubs, Diamonds, Hearts, Spades.

Deck 1 is used to draw the hold card. Deck 2 is used to draw the bid card. From Deck 1, each player receives a card drawn at random. This is the hold card. Players should not reveal this card to their opponents. The identity of this card is only known by the player that drew it.

Now we start the bidding portion of the game. Two players in the classroom are randomly drawn to be the two bidders in the game. From Deck 2, a random card is drawn. This card is called the bid card. The two bidders now compete in a sealed-bid first price auction. This means players would write their bids on a piece of paper, to be revealed only when both players have written bids. The player with the highest bid (ties broken with a coin flip) wins the bid card.

Let's consider the payoffs for players from losing and winning the auction. A player knows the identity of her hold card. When the bid card is drawn, there are 4 possibilities:

1. The bid card is identical to the hold card (recall that the hold card is drawn from Deck 1 and the bid card from the different Deck 2).

2. The bid card and the hold card are the same number. This is called a pair, i.e., a pair of kings or a pair of 3s.

3. The bid card and the hold card are the same suit, for instance, two clubs or two spades.

4. The bid card and the hold card are not the same number nor the same suit.

If a player is to win the auction, she now holds two cards. The value of the two cards are given by:

1. If the two cards are identical, the value equals 100.

2. If the two cards have the same number, the value equals 50.

3. If the two cards have the same suit, the value equals 25.

4. If the two cards do not match, the value equals the sum of the individual values of each card.

The winner of the auction has to pay the price for the bid card as determined from the auction. The price is equal to the highest bid in the auction (i.e., the bid that the player placed).

The loser of the auction only has one card (the hold card) and doesn't have to pay any prices. The value for the player is equal to the individual value of the card.

The individual values for the cards are ordered as follows:

Ace (any suit)	1 point
2 (any suit)	2 points
:	
:	
10 (any suit)	10 points
Jack (any suit)	10 points
Queen (any suit)	10 points
King (any suit)	10 points

Table 7.4

For instance, suppose that you are a bidder in a 2-bidder auction. You have drawn the 10 of Clubs as your hold card. The individual value of this card is equal to 10 points as seen in Table 7.4. This means that if you lose the auction, your payoff is equal to 10.

From Deck 2, the bid card is drawn. This card is the 10 of Spades. If you were to win the auction, then the value of your hand would be equal to 50 (as the cards have the same number). The value for the bid card to you is then 40 (equal to 50-10).

The payoff that you would receive as the auction winner is equal to 50 minus the bid that you have to place to win the auction. So you know that you would place a bid less than or equal to 40 in order to ensure that your payoff from winning the auction is greater than the payoff from losing the auction.

A bid of 40 is good, but perhaps you can win the auction with a lower bid. To determine what bid is optimal, you must take into account how valuable the card is to the other bidder. This requires forming the expectation over the following four outcomes:

1. The other bidder holds a 10 of Spades. The value for the other bidder from winning the auction is then 100 (the value from both cards) minus 10 (the value from just holding the 10 of Spades).

2. The other bidder holds a 10 of a different suit. The value for the other bidder from

winning the auction is then 50 (the value from both cards) minus 10 (the value of that bidder's hold card, either 10 of Diamonds or 10 of Hearts). Remember that you hold the 10 of Clubs, so the other bidder cannot hold this card.

3. The other bidder holds a card with Spade (but not a 10). The value for the other bidder form winning the auction is then 25 (the value from both cards) minus the value of that bidder's hold card (where the value ranges from 1 for the Ace of Spades to 10 for any Jack, Queen, or King of Spades).

4. The other bidder does not hold a Spade nor a 10. The value for the other bidder is then 10 (the value of a 10 of Spades by itself).

What is your bid going to be? As with the continuous distribution, the optimal bid will be equal to the expected value over all the possible outcomes for the other bidder for which you place the highest value on the object. For the continuous distribution with two-bidders and the uniform $Unif[0,1]$ distribution, suppose you are bidder 1. You only look at the values for the other bidder such that $v_2 \leq v_1$. You then take the expected value of $v_2 \tilde{} Unif[0,1]$ such that $v_2 \leq v_1$. The expected value is $\frac{1}{2}v_1$, which is the optimal bid in this auction. For the discrete auction introduced above, holding the 10 of Clubs as your hold card with a 10 of Spades as the bid card means that you only need to determine the expectation over the final three outcomes for the other bidder (10 of a different suit, Spade with a different number, and no match). The expected value for the final three outcomes is equal to 13.3. This means that your optimal bid is 13.3. As we can see, this is quite a bit less than the maximum bid of 40. This means that if the auction is won, the payoff will be equal to $50 - 13.3 = 36.7$.

Can you figure out how I arrived at the number 13.3? Can you find the optimal bid for other combinations of the hold card and the bid card? Given the time constraints (roughly 30 seconds to place a bid in this classroom auction), can you find an efficient way to approximate the optimal bid?

7.5 Revenue Equivalence

In our analysis of auctions, we have only considered the game from the point of view of the bidders (the buyers of the object). In any auction, there must also exist a seller on the other side of the market. This seller has an objective function and prefers to sell the object at a high price. An interesting property of auctions is that the expected revenue received by the seller in a first price auction is identical to the expected revenue in a second price auction (Mas-Colell et al., 1995, pgs. 889-891 and Dixit and Nalebuff, 2008, pgs. 307-308).

We will demonstrate this property in a two bidder auction in which both bidders have values drawn from the uniform distribution $Unif[0,1]$. The revenue equivalence result holds more generally.

In the first price auction, both bidders choose the bidding strategies $b_i = \frac{1}{2}v_i$ for $i = 1,2$. This means that the price that the seller receives is equal to $\max\left\{\frac{1}{2}v_1, \frac{1}{2}v_2\right\}$. The values v_1 and v_2 are drawn from the uniform distribution $Unif[0,1]$. The probability density function is equal to $f(v) = 1$ for all values of $v \in [0,1]$. The cumulative density function is $F(v) = \int_0^v f(v)dv = v$. Now consider the distribution $w = \max\{v_1, v_2\}$. The cumulative density function is $G(w) = F(w)^2 = w^2$. Why is this? The probability that $y \leq \max\{x_1, x_2\}$ is equal to the product $\Pr(y \leq x_1)\Pr(y \leq x_2)$. Thus, the probability density function is equal to $g(w) = 2w$ after taking the derivative of $G(w) = w^2$.

The formula to compute the expected payoff is given by $\int wg(w)dw$. So the expected payoff to the seller is equal to:

$$\begin{aligned}\frac{1}{2}\int_0^1 w(2w)dw &= \frac{1}{2}\left[\frac{2}{3}w^3\right]_0^1 \\ &= \frac{1}{2}\left[\frac{2}{3}\right] = \frac{1}{3}.\end{aligned}$$

In the second price auction, both bidders choose the bidding strategies $b_i = v_i$ for $i = 1,2$ and the price that the seller receives is equal to $\min\{v_1, v_2\}$, where v_1 and v_2 are drawn from the uniform distribution $Unif[0,1]$. Consider the distribution $w = \min\{v_1, v_2\}$. The cumulative density function is $H(w) = 1 - (1 - F(w))^2$. Why is this? This is because the probability that $y \geq \min\{x_1, x_2\}$ is equal to the product $\Pr(y \geq x_1)\Pr(y \geq x_2)$, where

$$\Pr(y \leq \min\{x_1, x_2\}) = 1 - \Pr(y \geq \min\{x_1, x_2\})$$

and

$$\Pr(y \geq x_1)\Pr(y \geq x_2) = (1 - \Pr(y \leq x_1))(1 - \Pr(y \leq x_2)).$$

The cumulative density function is given by:

$$\begin{aligned} H(w) &= 1 - (1-w)^2 \\ &= 1 - \left(1 - 2w + w^2\right) = 2w - w^2. \end{aligned}$$

Thus, the probability density function is equal to $h(w) = 2 - 2w$ after taking the derivative of $H(w) = 2w - w^2$.

The formula to compute the expected payoff is given by $\int wh(w)dw$. Thus, the expected payoff to the seller is equal to:

$$\begin{aligned} \int_0^1 w(2-2w)dw &= \left[w^2 - \frac{2}{3}w^3\right]_0^1 \\ &= 1 - \frac{2}{3} = \frac{1}{3}. \end{aligned}$$

Hence, we have verified that the revenues are equivalent for the two auctions: first price and second price auction (with two bidders and uniform value distributions).

7.6 Exercises

1. Consider the auction for a single good between three bidders. For simplicity, each bidder has value $v_1 = v_2 = v_3 = v$ for the object. Consider a third price sealed-bid auction. That is, the winner of the auction is the one with the highest bid (ties broken randomly) and the winner must pay the third highest bid for the object (with only three bidders, the third highest bid is also the lowest submitted bid).

   Is it a dominant strategy for a bidder to bid its value v? Why or why not?

2. Suppose you are one of two bidders in a first price sealed bid auction for Super Bowl tickets. Each bidder (including you) knows its own value for the Super Bowl tickets, but not the values of the other bidders. All a bidder knows about the values of the other bidders is that they are uniformly distributed between \$200 and \$300. Your value for the tickets is \$280. What is the optimal bid that you should place?

3. Suppose you are one of two bidders in a first price sealed bid auction. The values that each of the bidders (including you) place on the object are uniformly distributed between \$50 and \$150. You know your own value (which is \$130), but not the value of

the other bidder. In this game, the seller has already decided to impose a reservation price of \$80. A reservation price means that your bid can only win if it is both (i) greather than or equal to the reservation price and (ii) greater than or equal to the bid of the other bidder.

Given that your value for the object is \$130, what is the optimal bid that you should place?

4. Consider a first price sealed-bid auction between n different bidders, where $n > 3$ The bidders' valuations, v_i, are uniformly distributed between 0 and 1. Each bidder knows its true value, but does not know the values of the other bidders. The equilibrium bidding strategy for all bidders is given by $b_i = av_i$, where a is some unknown constant.

   Find the equilibrium bidding strategy for all the bidders (that is, determine the value for the constant a).

5. Consider Exercise 7.2 from above. Suppose the auction now has three bidders and your own value is still \$280. What will you decide to bid for the Super Bowl tickets?

6. Consider Exercise 7.3 from above. Suppose the auction now has 3 bidders and your own value is still \$130. What will you decide to bid for the object?

Bibliography

[1] Baird, Douglas, Robert Gertner, and Randal Picker. *Game Theory and the Law*. Cambridge, MA, and London, England: Harvard University Press, 1994.

[2] Dixit, Avinash, and Barry Nalebuff. *The Art of Strategy: A Game Theorist's Guide to Success in Business and Life*. 1st ed. New York, NY: W.W. Norton & Company, Inc., 2008.

[3] Gibbons, Robert. *Game Theory for Applied Economists*. Princeton, NJ: Princeton University Press, 1992.

[4] Mas-Colell, Andrew, Michael Whinston, and Jerry Green. *Microeconomic Theory*. New York, NY: Oxford University Press, Inc., 1995.

[5] Osborne, Martin J. *An Introduction to Game Theory: International Edition*. New York, NY: Oxford University Press, 2009.

Chapter 8

Forming Beliefs ahead of a Date

8.1 Motivating Example

A man and a woman are approaching the evening of their first date. The woman is preparing a home-cooked meal, while the man is bringing a bottle of wine. The meal being prepared is a surprise to the man, who must decide whether to bring a red wine (better with pasta and meat dishes) or a white wine (better with fish and salad dishes). The man and woman leave the office at 5 pm (yes, they work in the same office). The woman can either go East (where the fish and pasta stores are located) or can go West (where the meat and salad stores are located). The man observes whether the woman goes East or West, but not the store that the woman ultimately ends up at.

Given this information, the man forms beliefs about what type of meal the woman is preparing. The man then chooses to bring either red or white wine. The man and the woman have different payoff values for each of the 4 dinner choices, but each wants the wine and the food to pair well together (red wine with pasta or meat and white wine with fish or salad).

What is the optimal behavior for the man and the woman? Optimal behavior would include the man forming correct beliefs about the action taken by the woman. Correct beliefs are the next step as we introduce the third equilibrium concept of this manuscript. Only with correct beliefs can the system be at rest (our informal notion of equilibrium). Importantly, the notion of equilibrium in this chapter, unlike Nash equilibrium in Part I and subgame perfect Nash equilibrium in Part II, contains both the strategies of players and the beliefs that they form about the actions of the other players.

8.2 Information and Beliefs

This chapter serves as a prerequisite for the material in Chapters 9 and 10 on games of incomplete information.

> Complete information: All players know the payoff values of all other players (they know who they are playing against).

Games with incomplete information were introduced in the previous chapter (Chapter 7) in the context of auctions. However, these games were static as they were one-shot auctions. Every other interesting game with incomplete information is dynamic and requires the tools that we introduce in this chapter. Specifically, the games in this chapter have **complete**, but **imperfect** information.

> Perfect information: All players know the complete history of actions taken at any point in time.

If you recall, the dynamic games introduced in Chapter 4, without the presence of information sets, were games with perfect information. The dynamic games in Chapters 5 and 6, which included information sets, were games with imperfect information.

What is the connection between the two? In this chapter, we introduce the equilibrium concept called a "perfect Bayesian Nash equilibrium," or PBNE for short. This is an equilibrium concept that is used for dynamic games of imperfect information. In the following two chapters, we will translate the games of incomplete information into games of imperfect information. After the translation, we then apply the equilibrium concept of PBNE to form our theoretical predictions.

A perfect Bayesian Nash equilibrium is players' strategies and beliefs such that:

1. given beliefs, players' strategies are optimal.

2. at each information set reached along the equilibrium path, the beliefs must satisfy Bayes' rule.

We will see how to find the perfect Bayes Nash equilibrium in the following two examples, before I step aside to allow you to practice on your own. The beliefs in the PBNE are conditional probabilities. Even though they are probabilities, there is no relation to the mixed strategy Nash equilibria in Chapter 3. With mixed strategies, players assign certain probabilities to each of their strategies. With the PBNE beliefs, players assign certain probabilities to what they believe another player's action to have been. The action itself is pure (only one action is chosen), but the beliefs need not be.

Consider Figure 8.1 in the companion 'Figures' document (Gibbons, 1992, pgs. 175-178). As we can see, this game is played between two players: Player 1 and Player 2. Player 2 has an information set over the actions D1 and D2 of Player 1. The beliefs are included in Figure 8.1:

- p is the belief of Player 2 that Player 1 has played D1, conditional on either D1 or D2 having been played.

- $1 - p$ is the belief of Player 2 that Player 1 has played D2, conditional on either D1 or D2 having been played.

If we adopt the interpretation that an information set is a "cloud of uncertainty," then the beliefs indicate where Player 2 believes he is located within this cloud.

There are three strategies that can be chosen by Player 1, corresponding to the actions D1, D2, and T. Therefore, we have at most 3 PBNE. The method to find the PBNE will be the time-honored guess-and-check method:

1. Guess that Player 1 plays D1. Find the beliefs and actions of Player 2. Check to see if Player 1 then has an incentive to deviate (given the actions of Player 2).

2. Guess that Player 1 plays D2. Find the beliefs and actions of Player 2. Check to see if Player 1 then has an incentive to deviate (given the actions of Player 2).

3. Guess that Player 1 plays T. Find the beliefs and actions of Player 2. Check to see if Player 1 then has an incentive to deviate (given the actions of Player 2).

Player 1 plays D1 If Player 1 plays D1, then the information set in Figure 8.1 is reached with strictly positive probability. The information set is said to be on the equilibrium path. Consequently, the definition of a PBNE requires that the beliefs $(p, 1 - p)$ satisfy Bayes' rule. This means that the beliefs must be consistent with the action just taken by Player 1, i.e., beliefs are correct. So, if Player 1 plays D1, then Player 2 must believe that Player 1 has played D1 and all probability needs to be massed on the left: $p = 1$.

Given the belief $p = 1$, Player 2 takes the action from the left node in the information set. As the payoff from R (which equals 2) exceeds the payoff from L (which equals 1), Player 2 will choose R. Thus, the strategy and belief of Player 2 is given by $(p = 1, R)$.

As we are dealing with an information set, Player 2 can only take one action at both nodes and this action is R. This is held fixed and we want to see if Player 1 is really best off choosing D1. The payoff from D1 (which equals 1) is greater than the payoff from D2 (which is 0, as Player 2 plays R), but is strictly less than the payoff from T (which equals 2).

As Player 1 finds it optimal to deviate, then we cannot have a PBNE in which Player 1 plays D1.

Player 1 plays D2 The information set is reached with strictly positive probability, so the beliefs must be consistent. If Player 1 plays D2, then the probability needs to be massed on the right: $1 - p = 1$ or $p = 0$.

Given the belief $p = 0$, Player 2 takes the action from the right node in the information set. As the payoff from L (which equals 2) exceeds the payoff from R (which equals 1), Player 2 will choose L. Thus, the strategy and belief of Player 2 is given by $(p = 0, L)$.

As we are dealing with an information set, Player 2 can only take one action at both nodes and this action is L. We want to see if Player 1 is really best off choosing D2. The payoff from D2 (which equals 3) is greater than the payoff from D1 (which equals 2, as Player 2 plays L) and greater than the payoff from T (which equals 2).

Thus, one PBNE is given by $(D2, (p = 0, L))$, where D2 is the strategy of Player 1 (no beliefs are necessary) and $(p = 0, L)$ are the belief and strategy of Player 2.

Player 1 plays T The information set is not reached with strictly positive probability, meaning that the beliefs are not determined by the actions of Player 1. Player 2 is free to form whatever beliefs it desires. This means that we, as the game theorists, will find the values for p that are required to obtain a PBNE.

Consider that if Player 1 plays T, then its payoff is known to be equal to 2. Can Player 1 do better by deviating to either D1 or D2? The answer is that it depends upon what Player 2 does. The only deviation that leads to a strictly higher payoff for Player 1 occurs when Player 1 plays D2 and Player 2 plays L. Thus, provided that the Player 2 beliefs are such that Player 2 plays R, then there is no profitable deviation for Player 1. Without a a profitable deviation, Player 1 playing T can be part of a PBNE.

The expected payoffs to Player 2 are given by:

$$\begin{aligned} E_p(L) &= 1 \cdot p + 2 \cdot (1-p). \\ E_p(R) &= 2 \cdot p + 1 \cdot (1-p). \end{aligned}$$

The probabilities required for Player 2 to choose R are such that:

$$\begin{aligned} E_p(R) &\geq E_p(L) \\ 2 \cdot p + 1 \cdot (1-p) &\geq 1 \cdot p + 2 \cdot (1-p) \\ p + 1 &\geq 2 - p \\ 2p &\geq 1 \\ p &\geq \frac{1}{2}. \end{aligned}$$

It does not matter whether we allow for a weak inequality (tie goes to R) or a strict inequality. For the purposes of this course, we suppose that if the player is indifferent, then she will behave in the way that we, the game theorists, require.

Thus, a second PBNE is given by $\left(T, \left(p \geq \frac{1}{2}, R\right)\right)$, where T is the strategy of Player 1 (no beliefs are necessary) and $\left(p \geq \frac{1}{2}, R\right)$ are the belief and strategy of Player 2.

8.3 Comparing Equilibrium Concepts

One reason that we consider the new equilibrium concept of PBNE is to rule out "implausible" strategies that satisfy the requirements for Nash equilibrium and subgame perfect Nash equilibrium. This is not to say that this is the only reason to introduce PBNE. The key addition with a PBNE is the beliefs, which are useful when we consider games of incomplete information in the following two chapters. Additionally, this is not to suggest that the number of PBNE is always less than or equal to the number of Nash equilibria. It only suggests that in many examples, out of the multiple Nash equilibria, the "implausible" Nash equilibria are ruled out by introducing the PBNE concept.

Consider Figure 8.2 in the companion 'Figures' document (Gibbons, 1992, pgs. 175-178). If we tried to apply the solution concept of subgame perfect Nash equilibrium to this game, we soon realize that there is only one subgame: the entire game itself. Thus, the subgame perfect Nash equilibria are the same as the Nash equilibria. To find these, we write the game in the normal form.

Player 2

		L	R
	T	2, 1	2, 1
Player 1	D1	2, 1	1, 0
	D2	3, 2	0, 1

Table 8.1

Using the underline method, there are two SPNE: (T, R) and $(D2, L)$ (those same two sets of strategies are also Nash equilibria).

Exercise 1 asks us to show that the only PBNE is $(D2, (p = 0, L))$. There are no beliefs p that can support the strategies (T, R) as equilibrium strategies.

8.4 Solution to the Motivating Example

The date night game between the woman and the man is depicted in Figure 8.3 in the companion 'Figures' document. We know that there are 4 possible strategies for the woman, one for each of the possible dinner choices (Fish, Pasta, Meat, Salad). This means that there are at most 4 PBNE. To find the PBNE, we proceed through each of the 4 possible strategies with the following logic:

1. assume the woman plays that strategy,
2. find the beliefs and actions of the man, and then
3. determine if the woman is best served by choosing the assumed strategy.

Woman chooses Fish In this case, the left information set is reached on the equilibrium path. The beliefs are consistent when $p = 1$. The beliefs in the right information set are off the equilibrium path, so the probability q can take on any value. Given $p = 1$, the man chooses either red wine or white wine from the left information set. Knowing that he is at the left node of the left information set, the man chooses white wine (as the payoff of 4 is higher than the payoff of 3). The man prefers to pair white wine with fish (instead of red wine).

The belief and action of the man ($p = 1$, white wine if left information set) dictate that the payoff for the woman is equal to 6. Does the woman ever have an incentive to deviate?

Comparing the payoff values for the woman, we see that 6 is the highest. Thus, no matter what the man believes in the right information set, the woman will not deviate.

To be thorough, let us specify how the beliefs by the man in the right information set dictate the action taken at that information set. The expected payoffs are given by:

$$\begin{aligned} E_q\,(\text{Red Wine}) &= 6 \cdot q + 0 \cdot (1-q) \\ E_q\,(\text{White Wine}) &= 2 \cdot q + 3 \cdot (1-q) \end{aligned}$$

meaning that the man chooses red wine in the right information set whenever the probability q satisfies:

$$\begin{aligned} E_q\,(\text{Red Wine}) &\geq E_q\,(\text{White Wine}) \\ 6 \cdot q + 0 \cdot (1-q) &\geq 2 \cdot q + 3 \cdot (1-q) \\ q &\geq \frac{3}{7}. \end{aligned}$$

Thus, one PBNE is given by:

$$\left(Fish, \left(p = 1,\ \text{White Wine},\ \begin{array}{cc} q \geq \frac{3}{7} & \text{Red Wine} \\ q < \frac{3}{7} & \text{White Wine} \end{array} \right) \right).$$

Woman chooses Pasta In this case, the left information set is reached on the equilibrium path. The beliefs are consistent when $p = 0$. The beliefs in the right information set are off the equilibrium path, so the probability q can take any value. Given $p = 0$, the man chooses red wine in the left information set (as the payoff of 3 is higher than the payoff of 2).

The belief and action of the man ($p = 0$, red wine if left information set) dictate that the payoff for the woman equals 4. Does the woman have a profitable deviation? If she chooses fish, then her payoff is equal to 2 (as the man is bringing red wine in the left information set). Her payoffs that she receives by switching to either meat or salad depend upon what the man's beliefs are in the right information set. If the man brings white wine in the right information set, then the woman can do better than pasta by choosing salad (and receiving a payoff of 5). Thus, it is only optimal for the woman to choose pasta (no incentive to deviate) if the man has beliefs q such that he brings red wine in the right information set. When the man brings red wine in the right information set, the woman only has a possible payoff of 3 from meat and 2 from salad, both of which are less than the current payoff of 4 from pasta.

From the analysis above, we know that the man brings red wine in the right information set when $q \geq \frac{3}{7}$. Thus, the second PBNE is given by:

$$\left(Pasta, \left(p = 0, \text{ Red Wine}, q \geq \frac{3}{7}, \text{ Red Wine}\right)\right).$$

Woman chooses Meat In this case, the right information set is reached on the equilibrium path. The beliefs are consistent when $q = 1$. The beliefs in the left information set are off the equilibrium path, so the probability p can take on any value. Given $q = 1$, the man chooses red wine (as the payoff of 6 is higher than the payoff of 2).

The beliefs and action of the man ($q = 1$, red wine if right information set) dictate that the payoff for the woman equals 3. Does the woman ever have an incentive to deviate? She could choose salad and get a payoff of 2 (as the man chooses red wine in the right information set). The payoffs that she can receive for the other dinner options depends upon what beliefs the man holds in the left information set. Let's see how the probability p dictates what action is taken by the man at the left information set. The expected payoffs are given by:

$$\begin{aligned} E_p\,(\text{Red Wine}) &= 3 \cdot p + 3 \cdot (1 - p) \\ E_p\,(\text{White Wine}) &= 4 \cdot p + 2 \cdot (1 - p) \end{aligned}$$

meaning that the man chooses red wine in the left information set whenever the probability p satisfies:

$$\begin{aligned} E_p\,(\text{Red Wine}) &\geq E_p\,(\text{White Wine}) \\ 3 \cdot p + 3 \cdot (1 - p) &\geq 4 \cdot p + 2 \cdot (1 - p) \\ p &\leq \frac{1}{2}. \end{aligned}$$

Let's consider the two cases:

1. If $p \leq \frac{1}{2}$, the man chooses red wine in the left information set meaning that the woman can choose pasta and earn a payoff of 4 (which is strictly higher than the current payoff of 3 from meat).

2. If $p > \frac{1}{2}$, the man chooses white wine in the left information set meaning that the woman can choose fish and earn a payoff of 6 (which is strictly higher than the current payoff of 3 from meat).

This means that no matter what beliefs are held by the man in the left information set, the woman will always have an incentive to deviate from meat. Thus, we cannot have a PBNE in which meat is chosen by the woman.

Woman chooses Salad In this case, the right information set is reached on the equilibrium path. The beliefs are consistent when $q = 0$. The beliefs in the left information set are off the equilibrium path, so the probability p can take any value. Given $q = 0$, the man chooses white wine in the right information set (as the payoff of 3 is higher than the payoff of 0).

The belief and action of the man ($q = 0$, white wine if right information set) dictate that the payoff for the woman equals 5. Does the woman have a profitable deviation? If she chooses meat, then her payoff is equal to 1 as the man brings white wine in the right information set. Her payoffs that she gets by switching to either fish or pasta will depend upon what the man's beliefs are in the left information set. If the man brings white wine in the left information set, then the woman can do better than salad by choosing fish (and receiving the payoff of 6). Thus, it is only optimal for the woman to choose salad (no incentive to deviate) if the man has beliefs p such that he brings red wine in the left information set. When the man brings red wine in the left information set, the woman only has a payoff of 2 from fish and 4 from pasta, both of which are less than the current payoff of 5 from salad.

From the analysis above, we know that the man brings red wine in the left information set when $p \leq \frac{1}{2}$. Thus, the third and final PBNE is given by

$$\left(Salad, \left(p \leq \frac{1}{2}, \text{ Red Wine}, q = 0, \text{ White Wine}\right)\right).$$

8.5 Exercises

1. Show that the unique PBNE of Figure 8.2 is (D2,(p=0,L)).

2. Find the perfect Bayesian Nash equilibrium of the game in Figure 8.4, which is referred to as Selten's horse (Osborne, 2009, pg. 331).

3. A father is purchasing a Christmas present for his son. The father can buy one of three gifts: {Lego set, Nerf gun, video game}. The father selects one of the presents and then wraps it the day before Christmas and places it under the tree. The Lego set and the Nerf gun are the same size and shape, while the video game is noticeably smaller. The

son observes the size and shape of the wrapped present the day before Christmas and forms beliefs about what the present might be. The son is told that any bad behavior at the Christmas Eve mass will result in the loss of the Christmas present, but the son finds it very costly to avoid bad behavior during the long Christmas Eve mass.

The dynamic game is depicted in Figure 8.5. Find all perfect Bayesian Nash equilibria (PBNE) of this game.

Bibliography

[1] Baird, Douglas, Robert Gertner, and Randal Picker. *Game Theory and the Law.* Cambridge, MA, and London, England: Harvard University Press, 1994.

[2] Dixit, Avinash, and Barry Nalebuff. *The Art of Strategy: A Game Theorist's Guide to Success in Business and Life.* 1st ed. New York, NY: W.W. Norton & Company, Inc., 2008.

[3] Gibbons, Robert. *Game Theory for Applied Economists.* Princeton, NJ: Princeton University Press, 1992.

[4] Mas-Colell, Andrew, Michael Whinston, and Jerry Green. *Microeconomic Theory.* New York, NY: Oxford University Press, Inc., 1995.

[5] Osborne, Martin J. *An Introduction to Game Theory: International Edition.* New York, NY: Oxford University Press, 2009.

Chapter 9

Beer or Quiche?

9.1 Motivating Example

Let's consider how the actions of a player can be chosen to signal some unknown information. A student walks into a diner. Sitting in the diner is a bully. The bully receives utility from beating up students, but finds it costly to beat up the students that are "tough guys," while it is fairly costless to beat up the students that are "wimps." Just by looking at a student, the bully cannot tell whether the student is a tough guy or a wimp. All that the bully knows is that out of all the students at the school, 50% are tough guys and 50% are wimps.

The student knows whether he is a tough guy or a wimp. This is the incomplete information setting of this game: the student knows his own type, but the bully does not know the type of the student. The student has more information than the bully.

The student must select one of two items from the menu: either beer or quiche. The student receives payoff from both the menu item selected and from the ability to avoid a fight with the bully. The stereotypical assumption is that beer is a tough guy's drink and quiche (a baked egg-based entree) is not a common selection for a tough guy. Thus, we can imagine that the students who are tough guys receive a higher payoff from beer (compared to the students who are wimps) and a relatively lower payoff from quiche (compared to the students who are wimps).

The bully observes what item the student has ordered. Using this observation, he can update his beliefs about the type of student. Just because the students who are wimps prefer quiche, such students may choose to order beer just to mimic the students who are tough guys and confuse the bully. If the bully observes that all students order beer, then he has no new information. Any student that orders beer still has the probability 50% of being a tough guy and 50% of being a wimp. Students that are wimps may lose utility from choosing beer (instead of their preferred quiche), but are gaining utility by avoiding a fight with the bully (potentially).

What if the bully has played the diner game many times and observes that the only type of students that order quiche are the wimps? In this case, the bully needs to update beliefs using Bayes' rule.

This inference problem where an observed action is used to infer some piece of missing information is a classic problem that is found in many economic contexts. To begin to envision all possible applications, we need look no further than your decision to enroll in a university to signal your high ability level to future employers. We discuss this application later in the chapter.

9.2 Adverse Selection

Games of incomplete information are divided into two classes:

1. adverse selection and
2. moral hazard.

In games with adverse selection, one player fully knows the characteristics of an economic object or an economic player, and the other does not. Examples include the purchase of a used car (quality known by the dealer, but not by the purchaser) or online dating (a two-sided information problem in which each individual knows its true motives for posting an online profile, but not the motives of the second person).

Adverse selection is further divided into two subclasses:

1. signaling and
2. screening.

In signaling games, the leading example is the Beer-Quiche game in which the player with the missing information attempts to fill this gap by observing an action taken by its opponent.

In screening games, the player with the missing information chooses the possible actions that can be taken by its opponent. The plan is to choose the possible actions (a so called "menu" of choices) such that all opponents of a certain type choose Action A and all opponents of a different type choose Action B. In this way, the opponent will self-select into one of the possible actions and in doing so will reveal its unobservable type. This separation method is referred to as "screening." This chapter considers screening games in the context of an insurance company selling insurance to one of two types of drivers: safe (lower probability of accidents) and risky (higher probability of accidents).

Games with moral hazard are also referred to as games of "hidden action." In these games, one player is unable to observe the actions of the other player. We can view this in the context of an effort game. Two players, a boss and a worker, must agree on the terms of a labor contract. After agreement, the worker chooses how much effort to put into its job. The boss cannot observe the effort level of the worker, so cannot write a contract that conditions on effort. The only observable that the boss has is the productivity level of the worker. The effort level and the productivity level are correlated, but not perfectly correlated. You can imagine scenarios in which a random component of production causes a hard worker to reap low production, or vice versa.

Moral hazard games are considered in Chapter 10. This chapter considers games of adverse selection, starting with signal and moving onto screening.

9.3 Signaling I: Beer-Quiche

This section considers the canonical game in which both the unknown information about a player and the possible actions it can take belong to a finite set. The game is the Beer-Quiche game as described in the opening of this chapter. From the description of the game, we notice that the bully has less information than the student. The student knows whether he is a tough guy or a wimp, but the bully can only infer this missing information by observing the student's menu choice. The student has two possible actions: beer or quiche. The bully has two possible actions: fight the student or not fight.

We can equivalently formulate this game in terms of imperfect information. Suppose that a fictitious player called Nature (without payoffs) chooses at the beginning of the game whether to send a student that is a tough guy into the diner or whether to send a student that is a wimp into the diner. The student knows which action has been taken by Nature (it knows its type), but the bully does not observe this action by Nature. The bully's initial beliefs are that the student is a tough guy 50% of the time and a wimp 50% of the time. These beliefs can be updated based upon what the bully observes as the menu order of the student. The game is depicted in Figure 9.1 in the companion 'Figures' document (Baird et al., 1994, pgs. 157-158), where as always the payoff of the first player (the student) are listed on the left and the payoffs of the second player (the bully) are listed on the right.

For our purposes, it is much more convenient to view the game using an "inside-out" tree diagram, rather than the conventional "linear" tree diagram. Consider Figure 9.2 in the companion 'Figures' document. The first move is made by nature in the center of the figure. The next actions are taken by the two types of students (both the tough guy and the wimp) about what to order, either beer or quiche. The final actions are taken by the bully who makes a fight decision upon observing the order choice of the student. Naturally, the information sets denote the imperfect information held by the bully. The bully only observes the action of the students, and not the action of nature.

The student payoffs are always the left term in the parentheses, and the bully payoffs are the right term. The prior beliefs held by the bully are depicted at the point of time when nature takes its actions: 0.5 chance of choosing the tough guy type and 0.5 chance of choosing the wimp type.

The updated beliefs of the bully must satisfy the Bayes' rule requirement of the perfect Bayesian Nash equilibrium. That is, the beliefs are updated according to Bayes' rule whenever the information set is reached on the equilibrium path. We can view this game as being played by a large number of students, where 50% of the students are the tough guy type and 50% of the students are the wimp type.

With two types of students (two possible actions by nature) and two actions for each type of student, there are four possible strategies for the student:

1. tough guy type is the only one to select beer,

2. tough guy type is the only one to select quiche,

3. both types select beer, and

4. both types select quiche.

As in the previous chapter, with a total of four possible strategies, there are no more than four perfect Bayesian Nash equilibria. Our task is to consider each of the four possibilities and determine if we can find strategies and beliefs that constitute a perfect Bayesian Nash equilibrium. The first two equilibria are referred to as separating equilibria, as each type selects a different action. The final two equilibria are referred to as pooling equilibria as both types select the same action.

Let's proceed through all four of the possible PBNE and take the steps required to determine if there exists strategies and beliefs that satisfy the PBNE requirements.

Tough Guy chooses Beer; Wimp chooses Quiche If the tough guy type selects beer and the wimp type selects quiche, then the actions perfectly reveal the unobservable type of the student. The requirement that Bayes' rule must be satisfied implies that the bully now has correct beliefs about a student's type.

Define p as the probability that a student is of the tough guy type given that beer has been ordered. Define q as the probability that a student is of the tough guy type given that quiche has been ordered. Correctly updated beliefs require that $p = 1$ and $q = 0$.

Given these beliefs, what are the optimal actions of the bully? With $p = 1$, the bully knows that the game is at the top of the left information set in Figure 9.2, meaning that the payoff from fighting strictly exceeds the payoff from not fighting (1 is strictly greater than 0). Thus, the bully fights if he observes that beer has been ordered. With $q = 0$, the bully knows that the game is at the bottom of the information set, meaning that the payoff from not fighting strictly exceeds the payoff from fighting (1 is strictly greater than 0). Thus, the bully does not fight if he observes quiche.

Is there a profitable deviation for either the tough guy type or the wimp type? Currently, the tough guy type is earning a payoff of 0 from beer, but can earn a payoff of 2 from quiche (since the bully's action is to not fight upon observing quiche). This profitable deviation means that we cannot have a PBNE in Case 1.

Tough Guy chooses Quiche; Wimp chooses Beer If the tough guy type selects quiche and the wimp type selects beer, then correctly updated beliefs require that $p = 0$ and $q = 1$ (using the definitions of p and q as above).

Given these beliefs, what are the optimal actions for the bully? With $p = 0$, the bully knows that the game is at the bottom of the left information set in Figure 9.2. This means that the payoff from not fighting strictly exceeds the payoff from fighting (2 is strictly greater than 0). Thus, the bully does not fight if he observes that beer has been ordered. With $q = 1$, the bully knows that the game is at the top of the right information set, meaning that the payoff from fighting strictly exceeds the payoff from not fighting (2 is strictly greater than 0). Thus, the bully fights if he observes that quiche has been ordered.

Is there a profitable deviation for either the tough guy type or the wimp type? Currently, the tough guy type is earning a payoff of 1 from ordering quiche. The tough guy type can earn a payoff of 3 from beer, since the bully's action is to not fight when observing an order of beer. This profitable deviation means that we cannot have a PBNE in Case 2.

Both choose Beer In this case, only the left information set is reached on the equilibrium path. As both types take the same action, the bully does not have any new information above and beyond his initial beliefs. Thus, the correctly updated beliefs are such that the probabilities of tough guy and wimp after observing beer are both 50%. These are the initial beliefs. So $p = 0.5$.

Importantly, the right information set is not on the equilibrium path. The beliefs q are not pinned down and can take any value. We, as the economists, will choose the values for q to obtain a PBNE.

Given $p = 0.5$, the expected payoffs for the bully are either $E_p(\text{Fight}) = 1 \cdot p + 0 = 0.5$ from fighting and $E_p(\text{Not Fight}) = 0 + 2 \cdot (1 - p) = 1$ from not fighting. Thus, the bully will choose not to fight upon observing beer. This results in a payoff of 3 for the tough guy type and a payoff of 2 for the wimp type.

Does either type have an incentive to switch its action from beer to quiche? Notice that the tough guy type is currently receiving its highest possible payoff of 3.

The wimp type could either receive a payoff of 0 or 3 by switching to quiche. If the beliefs by the bully after observing quiche are such that the bully chooses to fight, then the wimp type would not have a profitable deviation (the beer payoff of 2 exceeds the quiche payoff of 0). By definition, $E_q(\text{Fight}) = 2 \cdot q + 0 = 2q$ from fighting and $E_p(\text{Not Fight}) = 0 + 1 \cdot (1 - q) = 1 - q$ from not fighting. The optimal action by the bully is to fight when

$$\begin{aligned} E_q(\text{Fight}) &\geq E_p(\text{Not Fight}) \\ 2q &\geq 1 - q \\ q &\geq \frac{1}{3}. \end{aligned}$$

Thus, we have found our first PBNE. I express the PBNE as the actions of the two types of students (the tough guy type and the wimp type), followed by the beliefs and actions of the bully:

$$\left(\text{Beer, Beer, } \left(p = 0.5, q \geq \frac{1}{3}, \text{ Not Fight if Beer, Fight if Quiche}\right)\right).$$

Both choose Quiche The correctly updated beliefs are such that the probabilities of tough guy and wimp are 50% after observing quiche. Thus, $q = 0.5$.

The left information set is not reached on the equilibrium path, so the belief p can take any value.

Given $q = 0.5$, we have found from the previous case that the bully will choose to fight upon observing quiche. This results in a payoff of 1 for the tough guy and a payoff of 0 for the wimp.

No matter what action is taken by the bully upon observing beer (a function of the belief p), the wimp type will always have a profitable deviation. Observe that the quiche payoff for the wimp type is equal to 0. This payoff is strictly dominated by both possible payoffs from beer: when the bully fights (payoff equals 1) and when the bully does not fight (payoff equals 2). This profitable deviation means that we cannot have a PBNE in Case 4.

Exercises 1 and 2 allow you to practice solving signaling games. The figures in these exercises provide nearly identical payoff values as those from Figure 9.2. And yet, where Figure 9.2 only has one PBNE (a pooling PBNE), Exercises 1 and 2 have either two pooling PBNE or one separating PBNE.

9.4 Signaling II: Job Market

This section considers a game (Mas-Colell et al., 1995, pgs. 450-460) in which the possible actions of a player continue to belong to a finite set, but now the unknown information about that player is described by a continuous distribution. For example, in the context of Beer-Quiche, the possible types of students could lie in the continuum between $[0, 1]$, where 0 indicates the weakest of wimps and 1 indicates the strongest of tough guys. Based upon this continuous distribution of types, the students will continue to select one of two items: either beer or quiche.

Another popular game to consider is one in which the possible actions of players can be taken continuously, but the possible types of players continue to be finite. We will not consider such a game in these notes, but in the context of a food order at a diner, you can imagine that a student can be either a tough guy or a wimp, and can signal this unknown information to the bully by choosing how much Tabasco hot sauce to put on his eggs. The amount of hot sauce is a continuous function.

The motivation for this game in this section will be job market signaling. We make a few standard assumptions to simplify the analysis and to make you feel better about all of the time and effort that you have invested to earn your college degree.

The natural ability of a student will be denoted a and will be evenly distributed along the interval $[1, 2]$, so $1 \leq a \leq 2$. The ability is known to the student, but not by the firms that wants to hire the student. We assume that the marginal profit to firms from hiring a student with ability a is equal to a. The marginal profit only depends upon natural ability, and does not depend upon the level of educational attainment. This is not to suggest that you give up your educational pursuits; keep reading.

The firms cannot observe the natural ability of a student, but can observe his or her highest level of educational attainment. The possible education levels are high school (H), bachelor's degree (B), or graduate degree (G). The student chooses its highest level of educational attainment, but must pay a cost as given in the following table.

$$c(H) = 0$$
$$c(B) = 2 - a$$
$$c(G) = 2(2 - a)$$

Table 9.1

What do we observe from these different cost functions? The first observation is that a graduate degree is more costly (in terms of time, effort, and lost opportunities) than a bachelor's degree, which is more costly than a high school degree. The second observation is that students with higher ability incur lower costs for each level of education. The highest ability student ($a = 2$) actually incurs zero cost for all 3 levels of education.

The firms are assumed to choose wages in a competitive setting. Recall that in a competitive setting, many firms compete for the same students. Consequently, the entire surplus is given to the students. Firms must pay out wages that are exactly equal to the marginal profits of the firm. Recall that the firms can only observe the highest education level of the students: H, B, or G. Thus, there are 3 possible wages that will be offered: $w(H)$, $w(B)$, and $w(G)$. These competitive market wages are equal to the expected marginal profit of the firms, as defined in the following table.

$$w(H) = \int_1^2 [p(a|H) \cdot a]\, da$$
$$w(B) = \int_1^2 [p(a|B) \cdot a]\, da$$
$$w(G) = \int_1^2 [p(a|G) \cdot a]\, da$$

Table 9.2

The wage is equal to the expected marginal profit conditional upon the highest level of educational attainment. Let's consider the high school wage. This is equal to the integral over the marginal profit a per student multiplied by the probability that a student has ability a given that the educational choice was high school. The integral is taken over all possible ability levels from 1 to 2.

The term $p(a|H)$ is the conditional probability that a student has ability a after choosing high school. These beliefs of the firms are updated according to Bayes' rule. The is identical to how the bully updated beliefs conditional upon observing the menu choice of the student.

With job market signaling, a student's payoff is equal to the wage it receives minus any costs of education.

I have now specified all the relevant information to solve this problem. I will first consider the wage offers and education choices when the firm has complete information about a student's natural ability (which is of course an outrageous assumption). I then consider the differences when the information is incomplete. For the remainder of this manuscript, I refer to incomplete information as asymmetric information. The information in this context is asymmetric because a student would know its natural ability, while a firm would not.

9.4.1 Complete Information

With complete information, a firm will simply pay a student according to its ability level (which is known). Thus a student of ability a will receive wage a. Consequently, the payoffs for the students are

$$a - \text{Cost of Education.}$$

This means that all students will choose to not receive any education (the highest ability student $a = 2$ is indifferent between all levels of education).

The prediction with complete information is that all students decide to stop with a high school education and not attend college.

9.4.2 Asymmetric Information

Let me specify the endogenous cutoffs as b^* and g^*. The values for these cutoffs will be found such that the students make educational decisions using the cutoff rule in the following table.

$a < b^*$	choose high school
$b^* \leq a < g^*$	choose bachelor's degree
$a \geq g^*$	choose graduate degree

Table 9.3

Using this cutoff rule, let's solve for the wages using the integrals specified previously. We begin with the high school wage:

$$w(H) = \int_1^{b^*} \left[p\left(a|H\right) \cdot a\right] da.$$

The abilities are distributed uniformly. The probability density function (pdf) of a continuous distribution must satisfy:

$$\int_1^{b^*} p\left(a|H\right) da = 1.$$

The pdf of the uniform distribution is a constant function $p\left(a|H\right) = \frac{1}{b^*-1}$. To see why, we must verify that the pdf property above is satisfied:

$$\int_1^{b^*} \frac{1}{b^*-1} da = \left[\frac{a}{b^*-1}\right]_1^{b^*} = \frac{b^*-1}{b^*-1} = 1.$$

Let's return to the wage integral:

$$\begin{aligned}
w(H) &= \int_1^{b^*} \left[\frac{1}{b^*-1} \cdot a\right] da \\
&= \left[\frac{a^2}{2\left(b^*-1\right)}\right]_1^{b^*} \\
&= \frac{b^{*2}}{2\left(b^*-1\right)} - \frac{1}{2\left(b^*-1\right)} \\
&= \frac{\left(b^*+1\right)\left(b^*-1\right)}{2\left(b^*-1\right)} = \frac{b^*+1}{2}.
\end{aligned}$$

Thus, the wage is equal to the average ability of all students to choose the educational level H.

In similar fashion, we can solve for the wage for a bachelor's degree:

$$w(B) = \int_{b^*}^{g^*} \left[p\left(a|B\right) \cdot a \right] da.$$

The probability $p\left(a|B\right) = \frac{1}{g^*-b^*}$ using the same argument as above. Thus, the wage integral can be evaluated as:

$$\begin{aligned} w(B) &= \int_{b^*}^{g^*} \left[\frac{1}{g^* - b^*} \cdot a \right] da \\ &= \left[\frac{a^2}{2\left(g^* - b^*\right)} \right]_{b^*}^{g^*} \\ &= \frac{g^{*2}}{2\left(g^* - b^*\right)} - \frac{b^{*2}}{2\left(g^* - b^*\right)} \\ &= \frac{\left(g^* + b^*\right)\left(g^* - b^*\right)}{2\left(g^* - b^*\right)} = \frac{g^* + b^*}{2}. \end{aligned}$$

The wage is equal to the average ability between b^* and g^*.

And finally, using the same methods, we can solve for the wage for a graduate degree:

$$w(G) = \int_{g^*}^{2} \left[p\left(a|G\right) \cdot a \right] da.$$

The probability $p\left(a|G\right) = \frac{1}{2-g^*}$ using the same argument as above. Thus, the wage integral can be evaluated as:

$$\begin{aligned} w(G) &= \int_{g^*}^{2} \left[\frac{1}{2 - g^*} \cdot a \right] da \\ &= \left[\frac{a^2}{2\left(2 - g^*\right)} \right]_{g^*}^{2} \\ &= \frac{2^2}{2\left(2 - g^*\right)} - \frac{g^{*2}}{2\left(2 - g^*\right)} \\ &= \frac{\left(2 + g^*\right)\left(2 - g^*\right)}{2\left(2 - g^*\right)} = \frac{2 + g^*}{2}. \end{aligned}$$

The wage is equal to the average ability from $a = g^*$ all the way up to $a = 2$.

The ability level $a = b^*$ is defined so that a student with this exact ability is indifferent between a high school education and a bachelor's degree. The payoffs for ability $a = b^*$ for these two choices are given in the following table.

High school payoff	Bachelor's degree payoff
$w(H) - 0 = \frac{b^*+1}{2}$	$w(B) - (2 - b^*) = \frac{g^*+b^*}{2} - (2 - b^*)$

Table 9.4

The payoffs are equal when:

$$\begin{aligned} \frac{b^*+1}{2} &= \frac{g^*+b^*}{2} - (2 - b^*) \\ (2 - b^*) &= \frac{g^*-1}{2}. \end{aligned}$$

The ability level $a = g^*$ is defined so that a student with this exact ability is indifferent between a bachelor's degree and a graduate degree. The payoffs for ability $a = g^*$ for these two choices are given in the following table.

Bachelor's degree payoff	Graduate degree payoff
$w(B) - (2 - g^*) = \frac{g^*+b^*}{2} - (2 - g^*)$	$w(G) - 2(2 - b^*) = \frac{2+g^*}{2} - 2(2 - g^*)$

Table 9.5

The payoffs are equal when:

$$\begin{aligned} \frac{g^*+b^*}{2} - (2 - g^*) &= \frac{2+g^*}{2} - 2(2 - g^*) \\ (2 - g^*) &= \frac{2-b^*}{2}. \end{aligned}$$

Let's solve the following two equations for the two unknowns (b^*, g^*) :

1. $(2 - b^*) = \frac{g^*-1}{2}$.
2. $(2 - g^*) = \frac{2-b^*}{2}$.

The solution is $(b^*, g^*) = (1.6, 1.8)$. This means that the wages are given by $w(H) = 1.3$, $w(B) = 1.7$, and $w(G) = 1.9$. In this model with asymmetric information, 40% of the students obtain at least a bachelor's degree, and 20% receive a graduate degree. This is in stark contrast to the complete information case in which 0% of the students obtain at least a bachelor's degree.

Another way to analyze our solution is in terms of the welfare considerations. With complete information, the payoffs for any student is simply equal to its ability level a. With asymmetric information, the students $a < 1.3$ are the only ones that can achieve a payoff strictly higher than their ability level:

High school payoff	Ability level
$w(H) - 0 = 1.3$	a

Table 9.6

The payoffs for the high school graduates are always equal to $w(H) = 1.3$. The payoffs for those with a bachelor's degree (ability levels $1.6 \leq a \leq 1.8$) range from 1.3 up to 1.5. The payoffs for those with a graduate degree (ability levels $1.8 \leq a \leq 2$) range from 1.5 up to 1.9. For all students with a bachelor's degree or a graduate degree, students receive a lower payoff than they would receive under the complete information case. These students are required to obtain an education (at a cost) in order to signal their high ability and separate themselves from lower ability students (and receive a higher wage). Additionally, students with the same educational level are pooled together, and those in the same pool with the lower ability level pull down the wage of those in that pool with the higher ability level.

This gives you something to look forward to while looking for a job in the coming years. Keep in mind, however, that if we had enough time to formulate this model correctly, we would allow a student's educational attainment to increase the firm's marginal profit (and thus increase the wage that the student receives).

9.5 Screening: Insurance

The following game (Mas-Colell et al., 1995, pgs. 460-467) is played between two groups of players: insurance companies and automobile drivers. There are assumed to be at least two risk-neutral insurance companies. With two companies, we can think back to a simple Bertrand competition game to know that the competition between the companies will drive profits down to zero.

There are two types of drivers in the economy: high type (H) and low type (L). The high type are high risk and have a probability of an accident equal to $p_H = \frac{1}{2}$. The low type are low risk and have a probability of an accident equal to $p_L = \frac{1}{3}$.

The insurance companies offer contracts to the drivers. The contracts are composed of two parts: a premium is a payment that is always made from the driver to the insurance company, and the deductible is the amount that the driver is responsible for in the case of an accident (the insurance company is responsible for the remainder of the repair costs). Let us denote the premium by M and the deductible by D. As there are two types of drivers (H and L), the insurance companies offer two possible insurance contracts: (M_H, D_H) and (M_L, D_L).

The drivers begin with an initial wealth equal to $W = 10$. In the event of an accident, the cost to repair the vehicle equals $R = 12$. Thus, the drivers must buy insurance or else they would be bankrupt and unable to pay for the repair costs in the case of an accident. Unlike the insurance companies, the drivers are risk averse. The expected utility of a high type driver that accepts the insurance contract (M_H, D_H) is given by:

$$\begin{aligned}\text{Payoff }(H;(M_H, D_H)) &= (1-p_H)\ln(W-M_H) + p_H\ln(W-M_H-D_H)\\ &= \frac{1}{2}\ln(10-M_H) + \frac{1}{2}\ln(10-M_H-D_H).\end{aligned}$$

Likewise, we can defined the expected utility for each of the three possibilities (high type accept (M_L, D_L), low type accept (M_H, D_H), and low type accept (M_L, D_L)).

Under complete information, the insurance companies can identify the type of driver. They can then set the contract accordingly. With asymmetric information, the insurance companies cannot observe what type of driver someone is. Consequently, whatever contracts are offered can by accepted by any type of driver. Accounting for this, the insurance companies will cleverly choose the two possible insurance contracts so that all high types select (M_H, D_H) and all low types accept (M_L, D_L).

9.5.1 Complete Information

The insurance companies are zero profit risk-neutral firms. The profit when the insurance companies offer the contract (M_H, D_H) to the high risk type is given by:

$$\pi_H = M_H - p_H(R - D_H) = 0.$$

This is equal to the collected amount M_H minus the expected payouts that occur in the case of an accident. The insurance company is responsible for the repair costs $R = 12$ minus the deductible paid by the driver D_H. Under zero profit, we can solve for the premium in terms of the deductible:

$$M_H = \frac{1}{2}(12 - D_H) = 6 - \frac{1}{2}D_H.$$

A similar story holds when the insurance companies offers the contract (M_L, D_L) to the low risk types:

$$\pi_L = M_L - p_L(R - D_L) = 0.$$

This implies that:

$$M_L = \frac{1}{3}(12 - D_L) = 4 - \frac{1}{3}D_L.$$

In addition to the insurance companies earning zero profit, the competition motive ensures that the contracts will maximize the expected utility of the drivers that accept them. The insurance companies know who the high type drivers are and offer them the contract (M_H, D_H) such that:

$$\max_{D_H} \frac{1}{2}\ln\left(10 - \left[6 - \frac{1}{2}D_H\right]\right) + \frac{1}{2}\ln\left(10 - \left[6 - \frac{1}{2}D_H\right] - D_H\right).$$

Let's simplify this optimization problem:

$$\max_{D_H} \frac{1}{2}\ln\left(4 + \frac{1}{2}D_H\right) + \frac{1}{2}\ln\left(4 - \frac{1}{2}D_H\right).$$

Taking the first order condition reveals:

$$\frac{1}{2}\left(\frac{0.5}{4 + \frac{1}{2}D_H}\right) + \frac{1}{2}\left(\frac{-0.5}{4 - \frac{1}{2}D_H}\right) = 0.$$

This implies that

$$4 + \frac{1}{2}D_H = 4 - \frac{1}{2}D_H,$$

so the deductible $D_H = 0$. With zero deductible, the drivers have perfect insurance and receive the same consumption whether an accident occurs or not.

Similarly, the contract (M_L, D_L) is offered to the low type drivers such that:

$$\max_{D_L} \frac{2}{3}\ln\left(10 - \left[4 - \frac{1}{3}D_L\right]\right) + \frac{1}{3}\ln\left(10 - \left[4 - \frac{1}{3}D_L\right] - D_L\right).$$

Let's simplify this optimization problem:

$$\max_{D_L} \frac{2}{3} \ln \left(6 + \frac{1}{3} D_L\right) + \frac{1}{3} \ln \left(6 - \frac{2}{3} D_L\right).$$

Taking the first order condition reveals:

$$\frac{2}{3} \left(\frac{1/3}{6 + \frac{1}{3} D_L}\right) + \frac{1}{3} \left(\frac{-2/3}{6 - \frac{2}{3} D_L}\right) = 0.$$

This implies that

$$6 + \frac{1}{3} D_L = 6 - \frac{2}{3} D_L,$$

so the deductible $D_L = 0$.

The two equilibrium contracts are $(M_H, D_H) = (6, 0)$ and $(M_L, D_L) = (4, 0)$. These are depicted in Figure 9.3 in the companion 'Figures' document. The x-axis is the amount of consumption in the good state (no accident) and the y-axis is the amount of consumption in the bad state (accident). The purple line is the 45-degree line. The blue curve is the indifference curve for the high risk type and the red curve is the indifference curve for the low risk type. In equilibrium, the high risk types consumes at the intersection of the purple and blue curves at $(4, 4)$. The low risk type consumes at the intersection of the purple and red curves at $(6, 6)$. The low risk type are less likely to have an accident, so will pay a lower insurance premium. This allows for a higher consumption level.

9.5.2 Asymmetric Information

With asymmetric information, the insurance companies cannot assign a certain insurance contract to a certain type of driver, because the driver types are unobserved by the insurance companies. Thus, when the companies offer the two possible contracts (M_H, D_H) and (M_L, D_L), the drivers can select which contract provides them with the highest expected payoff. Consider Figure 9.3 and which contracts the drivers would accept if they were able to accept either $(M_H, D_H) = (6, 0)$ or $(M_L, D_L) = (4, 0)$. The low type would continue to accept the contract (M_L, D_L) with consumption $(6, 6)$. However, the high type would not choose the contract (M_H, D_H), but would prefer the contract (M_L, D_L) with the lower premium. With the high type drivers accepting the wrong contract, the insurance companies are looking at negative profits. How do they adjust?

They must select the contracts (M_H, D_H) and (M_L, D_L) such that only the high type accepts the contract (M_H, D_H) and only the low type accepts the contract (M_L, D_L).

This leads to the notion of incentive compatibility. A contract is incentive compatible if it induces the desired type of driver (and only that type) to accept the contract. In this case, the contracts (M_H, D_H) and (M_L, D_L) are incentive compatible if high accepts high and low accepts low.

We know that the low type would never try to lie and pretend to be the high type. Why would a good driver pretend to be a bad driver in order to pay a higher premium? Thus, the contract (M_H, D_H) remains unchanged, because the low type would never want to accept that one. Recall that $(M_H, D_H) = (6, 0)$.

The incentive compatibility condition for the high type can be written as the following inequality:

$$\text{Payoff } (H; (M_H, D_H)) \geq \text{Payoff } (H; (M_L, D_L)).$$

We know that $(M_H, D_H) = (6, 0)$ and zero profit requires $M_L = 4 - \frac{1}{3}D_L$, so the only unknown variable in the incentive compatibility condition is D_L :

$$\begin{aligned} & \frac{1}{2}\ln\left(10 - \left[6 - \frac{1}{2}D_H\right]\right) + \frac{1}{2}\ln\left(10 - \left[6 - \frac{1}{2}D_H\right] - D_H\right) \\ \geq \; & \frac{1}{2}\ln\left(10 - \left[4 - \frac{1}{3}D_L\right]\right) + \frac{1}{2}\ln\left(10 - \left[4 - \frac{1}{3}D_L\right] - D_L\right). \end{aligned}$$

After plugging in $D_H = 0$, the constraint is given by:

$$\frac{1}{2}\ln(4) + \frac{1}{2}\ln(4) \geq \frac{1}{2}\ln\left(6 + \frac{1}{3}D_L\right) + \frac{1}{2}\ln\left(6 - \frac{2}{3}D_L\right).$$

We cancel out the $\frac{1}{2}$ from all terms. Let's use the property $\ln(a) + \ln(b) = \ln(ab)$:

$$\ln(16) \geq \ln\left\{\left(6 + \frac{1}{3}D_L\right)\left(6 - \frac{2}{3}D_L\right)\right\}.$$

Since the ln function is strictly increasing, then the inequality is equivalent to:

$$16 \geq \left(6 + \frac{1}{3}D_L\right)\left(6 - \frac{2}{3}D_L\right).$$

The remaining algebraic steps are:

$$\begin{aligned} 16 &\geq 36 - 2D_L - \frac{2}{9}D_L^2 \\ \frac{2}{9}D_L^2 + 2D_L - 20 &\geq 0 \\ D_L^2 + 9D_L - 90 &\geq 0 \\ (D_L - 6)(D_L + 15) &\geq 0 \end{aligned}$$

This implies that $D_L \geq 6$. The insurance companies offer contracts that provide the highest expected payoffs for the drivers. The best deductible for the drivers is the lowest $D_L = 6$. The incentive compatibility condition holds with equality. The contract is $(M_L, D_L) = \left(4 - \frac{1}{3}D_L, D_L\right) = (2, 6)$.

Consider Figure 9.4 from the companion 'Figures' document. The purple, red, and blue curves are exactly the same as in Figure 9.3. For this figure, I have added the green curve, which is the indifference curve for the low type under the new contract $(M_L, D_L) = (2, 6)$. Under this new contract, the consumption of the low type is given by $(8, 2)$, located at the intersection of the blue and green curves in the figure. The high type is exactly indifferent between the consumption $(8, 2)$ from the contract (M_L, D_L) and the consumption $(4, 4)$ from the contract (M_H, D_H). The low risk type obviously prefers the consumption $(8, 2)$ from (M_L, D_L) to the consumption $(4, 4)$ from (M_H, D_H). The incentive compatibility constraint holds with equality.

9.6 Exercises

1. Solve for the perfect Bayesian Nash equilibria of the Beer-Quiche game depicted in Figure 9.5. The game is identical to Figure 9.2 with two exceptions (which are indicated in **bold** in the figure).

2. Solve for the perfect Bayesian Nash equilibria of the Beer-Quiche game depicted in Figure 9.6. The game is identical to Figure 9.2 with one exception (which is indicated in **bold** in the figure).

3. Consider a game of asymmetric information played between a student and a firm. The student can have one of two ability levels: high or low. The student knows its own ability level, while the firm does not. All the firm knows is that 75% of students are

of high ability and 25% of students are of low ability.

The student sends a signal by deciding whether or not to attend college. The cost of college for a high ability student is 4, while the cost of college for a low ability student is 8. There is no cost associated with the decision not to attend college.

The firm forms beliefs about a student's ability level and then must decide to either offer the student high wages or low wages. High wages are value 10, while low wages are value 1.

The payoffs to the firm are as follows:

$$\text{Payoff} = 1 \text{ if offer} \quad \begin{array}{l} \text{High wages to High ability OR} \\ \text{Low wages to Low ability} \end{array} .$$

$$\text{Payoff} = 0 \text{ if offer} \quad \begin{array}{l} \text{High wages to Low ability OR} \\ \text{Low wages to High ability} \end{array} .$$

The payoff for the student is simply Payoff = Wages - Cost of Education.

Draw the game tree. Does there exist a perfect Bayesian Nash equilibrium (PBNE) in which both ability levels choose to attend college? Show all work to justify your answer.

4. Consider a game exactly as in the previous question, except now that 50% of students are of High ability and 50% of students are of Low ability. An additional difference is that now the cost of college for a high ability student is 5, while the cost of college fo ra low ability student is 10. All other facts are identical to those described in the previous question.

   Draw the game tree. Does there exist a perfect Bayesian Nash equilibrium (PBNE) in which only students of High ability go to college, while students of Low ability do not? Show all work to justify your answer.

5. Consider an insurance market with two risk-neutral insurance companies and risk-averse individuals. The individuals are of two types: High (H) and Low (L). Both types have wealth $W = 6$ and in the bad states (b) of the world suffer a loss of $R = 8$ (in the good states (g), there is no loss). The high type suffers a loss with probability $p_H = \frac{1}{2}$ and the low type suffers a loss with probability $p_L = \frac{2}{5}$. The individuals have

utility $u(\cdot) = \ln(\cdot)$, so the payoff functions are:

$$\begin{aligned} \text{Payoff}\,(H) &= p_H \cdot u(c_{H,b}) + (1 - p_H) \cdot u(c_{H,g}) \\ \text{Payoff}\,(L) &= p_L \cdot u(c_{L,b}) + (1 - p_L) \cdot u(c_{L,g}) \end{aligned}$$

The insurance companies offer two contracts: (M_H, D_H) and (M_L, D_L), where M is the premium and D is the deductible. With asymmetric information, solve for the optimal contracts that are offered by the insurance companies.

6. Consider an identical problem as above, except now both types of individuals have wealth $W = \frac{5}{2}$ and in the bad states (b) of the world suffer a loss of $R = 3$,where the high type suffers a loss with probability $p_H = \frac{1}{2}$ and the low type suffers a loss with probability $p_L = \frac{1}{3}$.

   With asymmetric information, solve for the optimal contracts that are offered by the insurance companies.

Bibliography

[1] Baird, Douglas, Robert Gertner, and Randal Picker. *Game Theory and the Law.* Cambridge, MA, and London, England: Harvard University Press, 1994.

[2] Dixit, Avinash, and Barry Nalebuff. *The Art of Strategy: A Game Theorist's Guide to Success in Business and Life.* 1st ed. New York, NY: W.W. Norton & Company, Inc., 2008.

[3] Gibbons, Robert. *Game Theory for Applied Economists.* Princeton, NJ: Princeton University Press, 1992.

[4] Mas-Colell, Andrew, Michael Whinston, and Jerry Green. *Microeconomic Theory.* New York, NY: Oxford University Press, Inc., 1995.

[5] Osborne, Martin J. *An Introduction to Game Theory: International Edition.* New York, NY: Oxford University Press, 2009.

Chapter 10

Writing Incentives into Contracts

10.1 Motivating Example

Let's pose the problem in terms of a principal (a fancy name for a boss) making a contract offer to an agent (a fancy name for a worker). The game (Mas-Colell et al., 1995, pgs. 477-506) that I describe is referred to as a principal-agent problem. The asymmetric information is that the principal cannot observe the action taken by the agent. Suppose that the principal cannot observe how much effort the agent is putting into the job. The principal wants the agents to put in a high level of effort, while the agent is only willing to put in a high level of effort if she is appropriately compensated. The principal cannot simply say, "If you work hard, you will receive a high wage," because the effort is unobservable.

What the principal can observe is the outcome of a risky production process. The outcome is correlated, but not perfectly correlated, with the amount of effort supplied by the agent. Thus, if the agent puts in a high level of effort, then the project is more likely to be successful, but the randomly determined outcome might still be an unsuccessful one. The principal can do one of two things: (i) write the contracts so that the agent only exerts a low level of effort or (ii) write the contracts so that the agent exerts a high level of effort. In case (i), the principal, as the owner of the firm's profits, would receive lower expected revenue, but would also have lower wage costs. The exact opposite holds for case (ii).

The task in this chapter is to solve for the optimal trade-off for the principal. The analysis in this chapter will determine the optimal wage contracts that should be offered by the principal in the presence of asymmetric information.

10.2 Kuhn-Tucker Conditions

Before starting with the economic analysis, we need to first review the equations associated with the solutions to constrained optimization problems. The solutions to constrained optimization problems satisfy the Kuhn-Tucker conditions. A constrained optimization problem takes the following form:

$$\begin{array}{cc} \max & f(x) \\ \text{subject to} & g(x) \geq 0 \\ \text{and} & h(x) \geq 0 \end{array} .$$

Notice that the constraints $g(x) \geq 0$ and $h(x) \geq 0$ are always written as nonnegativity constraints. An example of a constrained optimization problem is a utility maximization problem in which $f(x) = u(x_1, x_2)$ and $g(x) = w - p_1 x_1 - p_2 x_2 \geq 0$ (no constraint h).

Let λ be the Lagrange multiplier associated with the constraint $g(x) \geq 0$ and μ be the Lagrange multiplier associated with the constraint $h(x) \geq 0$. The Kuhn-Tucker conditions have two components:

1. First order conditions

$$Df(x) + \lambda Dg(x) + \mu Dh(x) = 0.$$

2. Complimentary slackness conditions

$$\begin{array}{lll} g(x) \geq 0. & \lambda \geq 0. & \lambda g(x) = 0. \\ h(x) \geq 0. & \mu \geq 0. & \mu h(x) = 0. \end{array}$$

From the complimentary slackness conditions, if $\lambda > 0$, then $g(x) = 0$. Likewise, if $\mu > 0$, then $h(x) = 0$.

10.3 Principal-Agent Setup

In the principal-agent problem, the agent can choose one of two effort levels: $e = e_L$ or $e = e_H$. The principal cannot observe the effort level. The possible profits that the principal can receive from a risky project are of two values: $\pi = 1$ or $\pi = 10$. The profits are randomly determined as a functions of the effort chosen by the agent.

$$prob\,(\pi = 10|e = e_H) = p_H = \tfrac{3}{4}$$
$$prob\,(\pi = 10|e = e_L) = p_L = \tfrac{1}{4}$$

Table 10.1

Putting forth high effort is costly to the agent. The costs for effort are given in the following table.

$$c\,(e = e_L) = 0$$
$$c\,(e = e_H) = c = \tfrac{1}{2}$$

Table 10.2

The agent has an outside option, meaning that it can quit its job at any time, and receive a payoff equal to $\bar{u} = 1$. Without the labor input from the agent, we assume that the principal earns a profit $\pi = 0$ (though we could imagine $\pi < 0$ in the presence of sunk costs).

Given that there are two profit levels, the principal can offer the agent a menu of two wages: $w(10)$ is the wage received when $\pi = 10$ and $w(1)$ is the wage received when $\pi = 1$.

The principal is risk neutral, so will maximize the expected profit as seen in the following table.

$$p_L\,(10 - w(10)) + (1 - p_L)\,(1 - w(1)) \quad \text{if effort } e = e_L$$
$$p_H\,(10 - w(10)) + (1 - p_H)\,(1 - w(1)) \quad \text{if effort } e = e_H$$

Table 10.3

The agent is risk averse and maximizes the expected utility as given in the following table.

$$p_L \ln{(w(10))} + (1 - p_L) \ln{(w(1))} \quad \text{if effort } e = e_L$$
$$p_H \ln{(w(10))} + (1 - p_H) \ln{(w(1))} \quad \text{if effort } e = e_H$$

Table 10.4

10.4 Complete Information

Under complete information, the effort level of the agent is observable. This is a strong assumption, but we use it just now to find the benchmark outcome for later comparison.

If observe low effort The wages $(w(10), w(1))$ are chosen to solve the following constrained optimization problem of the principal:

$$\begin{array}{rl} \max\limits_{w(10),w(1)} & p_L\,(10 - w(10)) + (1 - p_L)\,(1 - w(1)) \\ \text{subject to} & p_L \ln(w(10)) + (1 - p_L)\ln(w(1)) \geq \bar{u} \end{array}$$

The problem states that the principal is maximizing expected profit by offering a contract that is accepted by the agent.

The first order condition with respect to $w(10)$ is given by:

$$-p_L + \lambda p_L \left(\frac{1}{w(10)}\right) = 0.$$

This means that $\lambda = w(10)$. The first order condition with respect to $w(1)$ is given by:

$$-(1 - p_L) + \lambda(1 - p_L)\left(\frac{1}{w(10)}\right) = 0.$$

This means that $\lambda = w(1)$.

Given $w(10) = w(1)$ and $\lambda > 0$, the wages are chosen such that the constraint $p_L \ln(w(10)) + (1 - p_L)\ln(w(1)) \geq \bar{u}$ holds with equality:

$$\begin{aligned} \frac{1}{4}\ln(w) + \frac{3}{4}\ln(w) &= 1. \\ w &= w(10) = w(1) = \exp(1). \end{aligned}$$

The principal's expected profit is equal to:

$$\frac{1}{4}(10 - \exp(1)) + \frac{3}{4}(1 - \exp(1)) \cong 0.532.$$

If observe high effort The wages $(w(10), w(1))$ are chosen to solve the following constrained optimization problem of the principal (don't forget the effort cost in the expected utility for the agent):

$$\begin{array}{rl} \max\limits_{w(10),w(1)} & p_H\,(10 - w(10)) + (1 - p_H)\,(1 - w(1)) \\ \text{subject to} & p_H \ln(w(10)) + (1 - p_H)\ln(w(1)) - c \geq \bar{u} \end{array}$$

The first order condition with respect to $w(10)$ is given by:

$$-p_H + \lambda p_H \left(\frac{1}{w(10)}\right) = 0.$$

This means that $\lambda = w(10)$. The first order condition with respect to $w(1)$ is given by:

$$-(1 - p_H) + \lambda(1 - p_H)\left(\frac{1}{w(10)}\right) = 0.$$

This means that $\lambda = w(1)$.

Given $w(10) = w(1)$ and $\lambda > 0$, the wages are chosen such that the constraint $p_H \ln(w(10)) + (1 - p_H)\ln(w(1)) - c \geq \bar{u}$ holds with equality:

$$\begin{aligned} \frac{3}{4}\ln(w) + \frac{1}{4}\ln(w) - \frac{1}{2} &= 1. \\ w &= w(10) = w(1) = \exp\left(\frac{3}{2}\right). \end{aligned}$$

The principal's expected profit is equal to:

$$\frac{3}{4}\left(10 - \exp\left(\frac{3}{2}\right)\right) + \frac{1}{4}\left(1 - \exp\left(\frac{3}{2}\right)\right) \cong 3.268.$$

10.5 Asymmetric Information

Under asymmetric information, the principal can no longer observe the effort level of the agent. If the principal offers the wage menu $(w(10), w(1)) = (\exp(1), \exp(1))$, then the agent would choose the low effort level. This is one possible option for the principal: to design the wage contract so that the effort level is low.

If the principal offers the wage menu $(w(10), w(1)) = \left(\exp\left(\frac{3}{2}\right), \exp\left(\frac{3}{2}\right)\right)$, then the agent would continue to put forth low effort. The principal is worse off with $\left(\exp\left(\frac{3}{2}\right), \exp\left(\frac{3}{2}\right)\right)$ compared to $(\exp(1), \exp(1))$: same effort level and high wage. The second possible option for the principal is to design the wage contract so that the effort level is high. Clearly the wage menu $(w(10), w(1)) = \left(\exp\left(\frac{3}{2}\right), \exp\left(\frac{3}{2}\right)\right)$ does not suffice. In order to induce high effort, the wages must be chosen to satisfy the following constrained optimization problem of the principal:

$$\begin{array}{cc} \max\limits_{w(10),w(1)} & p_H\left(10-w(10)\right)+\left(1-p_H\right)\left(1-w(1)\right) \\ \text{subject to} & p_H\ln\left(w(10)\right)+\left(1-p_H\right)\ln\left(w(1)\right)-c\geq \bar{u} \\ \text{and} & p_H\ln\left(w(10)\right)+\left(1-p_H\right)\ln\left(w(1)\right)-c\geq p_L\ln\left(w(10)\right)+\left(1-p_L\right)\ln\left(w(1)\right) \end{array}$$

The first constraint was already considered in the complete information case: the outside utility constraint, also known as the individual rationality constraint. The second constraint is new. It is the incentive compatibility constraint stating that the agent's expected utility from high effort (including the cost of effort) is at least as large as its expected utility from low effort.

Let λ be the Lagrange multiplier associated with the first constraint (the individual rationality (IR) constraint) and μ be the Lagrange multiplier associated with the second constraint (the incentive compatibility (IC) constraint). From the algebraic appendix at the end of this chapter, I know that both of the constraints will be equalities (see Section 10.6).

With both constrained as equalities, we have two equations that we can use to solve for the only two variables $(w(10), w(1))$:

$$\begin{array}{rr} \tfrac{3}{4}\ln\left(w(10)\right)+\tfrac{1}{4}\ln\left(w(1)\right)-\tfrac{1}{2}=1 & (IR) \\ \tfrac{3}{4}\ln\left(w(10)\right)+\tfrac{1}{4}\ln\left(w(1)\right)-\tfrac{1}{2}=\tfrac{1}{4}\ln\left(w(10)\right)+\tfrac{3}{4}\ln\left(w(1)\right) & (IC) \end{array}$$

From the (IC) equation, we obtain:

$$\begin{aligned} \frac{1}{2}\ln\left(w(10)\right)-\frac{1}{2}\ln\left(w(1)\right) &= \frac{1}{2}. \\ \ln\left(w(10)\right)-\ln\left(w(1)\right) &= 1. \end{aligned}$$

Combining this fact with the (IR) equation, we obtain:

$$\begin{aligned} \frac{3}{4}\left\{\ln\left(w(10)\right)-\ln\left(w(1)\right)\right\}+\ln\left(w(1)\right)-\frac{1}{2} &= 1. \\ \frac{3}{4}+\ln\left(w(1)\right)-\frac{1}{2} &= 1. \\ w(1) &= \exp\left(\frac{3}{4}\right). \end{aligned}$$

Notice that $w(1)=\exp\left(\frac{3}{4}\right)$ is a smaller wage than that which is offered by the principal to induce low effort.

We can also solve for $w(10)$:

$$\begin{aligned} \ln(w(10)) &= \ln(w(1)) + 1. \\ w(10) &= \exp\left(\frac{7}{4}\right). \end{aligned}$$

Notice that $w(10) = \exp\left(\frac{7}{4}\right)$ is a higher wage than that received in the complete information case when high effort is observed. This risk-reward payoff is what entices the agent to put forth high effort.

The expected profit of the principal is given by:

$$\frac{3}{4}\left(10 - \exp\left(\frac{7}{4}\right)\right) + \frac{1}{4}\left(1 - \exp\left(\frac{3}{4}\right)\right) \cong 2.905.$$

This results in a higher expected profit than what the principal would receive by setting the wage contracts to induce low effort.

Figure 10.1 in the companion 'Figures' document displays the outcome. The x-axis lists the wage of failure $w(1)$ and the y-axis lists the wage of success $w(10)$. The green line is the 45-degree line. The red curve is the indifference curve for the agent under low effort and the blue curve is the indifference curve for the agent under high effort. Under complete information, if high effort is observed, the wages are given by the intersection of the green and bluecurves at $(w(1), w(10)) = \left(\exp\left(\frac{3}{2}\right), \exp\left(\frac{3}{2}\right)\right) = (4.48, 4.48)$.

Under asymmetric information, if the principal chooses the contracts to induce low effort, then the outcome occurs at the intersection of the green and red curves at $(w(1), w(10)) = (\exp(1), \exp(1)) = (2.72, 2.72)$. If the principal chooses the contracts to induce high effort, then the outcome is at the intersection of the red and blue curves at $(w(1), w(10)) = \left(\exp\left(\frac{3}{4}\right), \exp\left(\frac{7}{4}\right)\right) = (2.12, 5.75)$. At this point, the agent is indifferent between high and low effort. This is the implication of the incentive compatibility constraint holding with equality.

For these chosen numbers, we see that the principal can earn a higher expected profit under asymmetric information by writing the contracts to induce high effort. Just for kicks, let's look at the comparative statics associated with changing the values of our four main parameters: $(\bar{u}, c, p_H, p_L)$. In all cases, if it becomes "more costly" for the principal to write wage contracts to induce high effort, then the expected profit will decrease.

Recall that the principal can earn 0 by offering contracts that violate the agent's IR constraints. So if you see any negative values in the following tables, then you know that the principal will offer contracts that won't be accepted.

1. Changing $\bar{u}$ and holding fixed $(c, p_H, p_L) = \left(\frac{1}{2}, \frac{3}{4}, \frac{1}{4}\right)$:

	$\bar{u} = 1$	$\bar{u} = 2$	$\bar{u} = 3$
Principal profit (induce low)	0.532	-4.139	-16.836
Principal profit (induce high)	2.905	-5.421	-28.051

Table 10.5

2. Changing c and holding fixed $(\bar{u}, p_H, p_L) = \left(1, \frac{3}{4}, \frac{1}{4}\right)$:

	$c = \frac{1}{2}$	$c = 1$	$c = 2$
Principal profit (induce low)	0.532	0.532	0.532
Principal profit (induce high)	2.905	-1.799	-33.449

Table 10.6

3. Changing p_H and holding fixed $(\bar{u}, c, p_L) = \left(1, \frac{1}{2}, \frac{1}{4}\right)$:

	$p_H = \frac{3}{4}$	$p_H = \frac{5}{8}$	$p_H = \frac{1}{2}$
Principal profit (induce low)	0.532	0.532	0.532
Principal profit (induce high)	2.905	1.276	-1.416

Table 10.7

4. Changing p_L and holding fixed $(\bar{u}, c, p_H) = \left(1, \frac{1}{2}, \frac{3}{4}\right)$:

	$p_L = \frac{1}{4}$	$p_L = \frac{3}{8}$	$p_L = \frac{1}{2}$
Principal profit (induce low)	0.532	1.657	2.782
Principal profit (induce high)	2.905	2.647	1.958

Table 10.8

See if you can generate the expected profits for the principal in Tables 10.1-10.4 on your own. As additional practice, complete the exercise at the end of the chapter.

10.6 Algebraic Appendix

10.6.1 Maximization Problem

In the prior sections, we wrote down the maximization problem for the principal under the specification that the wage are set to induce high effort:

$$\begin{array}{lcr} \max\limits_{w(10),w(1)} & p_H\left(10-w(10)\right)+\left(1-p_H\right)\left(1-w(1)\right) & \\ \text{subject to} & p_H\ln\left(w(10)\right)+\left(1-p_H\right)\ln\left(w(1)\right)-c\geq\bar{u} & (\lambda) \\ \text{and} & p_H\ln\left(w(10)\right)+\left(1-p_H\right)\ln\left(w(1)\right)-c\geq p_L\ln\left(w(10)\right)+\left(1-p_L\right)\ln\left(w(1)\right) & (\mu) \end{array}.$$

Here, we let λ be the Lagrange multiplier associated with the first constraint and μ be the Lagrange multiplier associated with the second constraint.

10.6.2 Both Constraints Bind

Our task is to show that both $\lambda > 0$ and $\mu > 0$. If $\lambda > 0$, then the complimentary slackness condition implies that:

$$p_H\ln\left(w(10)\right)+\left(1-p_H\right)\ln\left(w(1)\right)-c=\bar{u}.$$

If $\mu > 0$, then the complimentary slackness condition implies that:

$$p_H\ln\left(w(10)\right)+\left(1-p_H\right)\ln\left(w(1)\right)-c=p_L\ln\left(w(10)\right)+\left(1-p_L\right)\ln\left(w(1)\right).$$

To show $\lambda > 0$ and $\mu > 0$, take the first order conditions with respect to both $w(10)$ and $w(1)$. The first order condition with respect to $w(10)$ is given by:

$$-p_H+\lambda p_H\left(\frac{1}{w(10)}\right)+\mu\left(p_H-p_L\right)\left(\frac{1}{w(10)}\right)=0.$$

The first order condition with respect to $w(1)$ is given by:

$$-(1-p_H)+\lambda(1-p_H)\left(\frac{1}{w(1)}\right)+\mu\left((1-p_H)-(1-p_L)\right)\left(\frac{1}{w(1)}\right)=0.$$

Lambda

From the two first order conditions, what happens if $\lambda = 0$? When $\lambda = 0$, the $w(1)$ first order condition becomes:

$$-(1 - p_H) + \mu (p_L - p_H) \left(\frac{1}{w(1)} \right) = 0,$$

or equivalently

$$\mu \left(\frac{1}{w(1)} \right) = \frac{1 - p_H}{p_L - p_H}.$$

However, this is not possible as the left-hand side is nonnegative (since $\mu \geq 0$ and $w(1) \geq 0$), while the right-hand side is strictly negative (since $p_L - p_H < 0$). Thus, we can never have $\lambda = 0$. The only possibility is $\lambda > 0$.

Mu

From the two first order conditions, what happens if $\mu = 0$? When $\mu = 0$, the first order conditions become:

$$\begin{aligned} -p_H + \lambda p_H \left(\frac{1}{w(10)} \right) &= 0. \\ -(1 - p_H) + \lambda(1 - p_H) \left(\frac{1}{w(1)} \right) &= 0. \end{aligned}$$

This implies that $w(10) = w(1) = \lambda$. With equal wages $w = w(10) = w(1)$, the incentive compatibility condition

$$p_H \ln (w(10)) + (1 - p_H) \ln (w(1)) - c \geq p_L \ln (w(10)) + (1 - p_L) \ln (w(1))$$

reduces to

$$\ln(w) - c \geq \ln(w).$$

Obviously, this inequality cannot hold as $c > 0$. Thus, we can never have $\mu = 0$. The only possibility is $\mu > 0$.

10.7 Exercises

1. Consider a game of moral hazard between a risk-neutral owner and a risk-averse manager. A project is available with random returns: either of value 200 or of value 100. The owner can offer a wage pair $(w(200), w(100))$, where $w(200)$ is the wage when the project has return 200 (success) and $w(100)$ is the wage when the project has return 100 (fail).

   The payoff function for the owner is then given by:

   $$p \cdot (200 - w(200)) + (1 - p)(100 - w(100)).$$

   The payoff function for the managers is given by:

   $$p \cdot \sqrt{w(200)} + (1 - p)\sqrt{w(100)} - \text{cost of effort}.$$

   The probability of success p is determined based upon the effort decision of the manager:

   $$\begin{array}{ll} p = p_H = 0.75 & \text{if manager chooses High effort} \\ p = p_L = 0.25 & \text{if manager chooses Low effort} \end{array}.$$

   The cost of effort for the manager is 0 for Low effort and $c = 4$ for High effort. Additionally, the manager has an outside option with payoff $\bar{u} = 4$.

   In this game with asymmetric information in which the owner cannot observe the effort decision of the manager, what wage pair $(w(200), w(100))$ provides the incentives for the manager to choose High effort? What are the payoffs for both the manager and the owner?

Bibliography

[1] Baird, Douglas, Robert Gertner, and Randal Picker. *Game Theory and the Law.* Cambridge, MA, and London, England: Harvard University Press, 1994.

[2] Dixit, Avinash, and Barry Nalebuff. *The Art of Strategy: A Game Theorist's Guide to Success in Business and Life.* 1st ed. New York, NY: W.W. Norton & Company, Inc., 2008.

[3] Gibbons, Robert. *Game Theory for Applied Economists.* Princeton, NJ: Princeton University Press, 1992.

[4] Mas-Colell, Andrew, Michael Whinston, and Jerry Green. *Microeconomic Theory.* New York, NY: Oxford University Press, Inc., 1995.

[5] Osborne, Martin J. *An Introduction to Game Theory: International Edition.* New York, NY: Oxford University Press, 2009.

Appendix A

Solutions to Exercises

A.1 Chapter 1 Solutions

1. Find the Nash equilibria of the following normal form game (Gibbons, 1992, pgs. 9-10).

		Player 2		
		L	C	R
	T	3, 2	1, 6	3, 2
Player 1	M	5, 1	4, 3	3, 2
	B	2, 4	5, 4	1, 4

Solution: Using the underline method, the Nash equilibrium of the game is (B, C):

		Player 2		
		L	C	R
	T	3, 2	1, $\underline{6}$	$\underline{3}$, 2
Player 1	M	$\underline{5}$, 1	4, $\underline{3}$	$\underline{3}$, 2
	B	2, $\underline{4}$	$\underline{5}$, $\underline{4}$	1, $\underline{4}$

2. Find the Nash equilibria of the following normal form game (Gibbons, 1992, pgs. 9-10).

		Player 2		
		Save	Use	Delay
	Save	5, 5	-2, 10	0, 0
Player 1	Use	3, 2	1, -1	0, 0
	Delay	0, 0	0, 0	0, 0

Solution: Using the underline method, the Nash equilibrium of the game is (Delay, Delay):

		Player 2		
		Save	Use	Delay
	Save	$\underline{5}$, 5	-2, $\underline{10}$	$\underline{0}$, 0
Player 1	Use	3, $\underline{2}$	$\underline{1}$, -1	$\underline{0}$, 0
	Delay	0, $\underline{0}$	0, $\underline{0}$	$\underline{0}$, $\underline{0}$

3. Write the following games in the normal form and solve for the Nash equilibria. In some cases, multiple Nash equilibria may exist or a Nash equilibrium may not exist at all. (Note: There is some freedom in how you assign payoff values for the players. These choices may affect the Nash equilibria.)

   (a) Two animals are fighting over some prey (this game is typically called a Hawk-Dove game). Each can be passive or aggresive. Each prefers to be aggresive if its opponent is passive, and passive if its opponent is aggresive. Given its own stance, it prefers the outcome in which its opponent is passive to that in which its opponent is aggresive. (Dixit and Nalebuff, 2008, pgs. 97-101 and Mas-Colell et al., 1995, pg. 265).

   Solution: The payoffs can vary, but one set of possible payoff values is:

		Animal 2	
		Passive	Aggresive
Animal 1	Passive	0, 0	1, 3
	Aggresive	3, 1	-2, -2

   Using the underline method, there are two Nash equilibria (pure-strategy) of the above game:

		Animal 2	
		Passive	Aggresive
Animal	Passive	0, 0	1, 3
1	Aggresive	3, 1	-2, -2

The Nash equilibria are (Aggresive, Passive) and (Passive, Aggresive).[1]

(b) Two students wish to attend the same university. The students each receive strictly positive payoff if they attend the same university and zero payoff otherwise. The list of possible universities is Purdue University, Indiana University, and Notre Dame University. Student A prefers Purdue to Indiana and Indiana to Notre Dame (by transitivity, he/she also prefers Purdue to Notre Dame). Student B has a scholarship at Notre Dame, so prefers Notre Dame to Purdue and Purdue to Indiana (by transitivity, he/she prefers Notre Dame to Indiana).

Solution: The payoffs can vary, but one set of possible payoff values is:

		Student 2		
		Purdue	Indiana	Notre Dame
	Purdue	3, 2	0, 0	0, 0
Student 1	Indiana	0, 0	2, 1	0, 0
	Notre Dame	0, 0	0, 0	1, 3

Using the underline method, there are three Nash equilibria (pure-strategy) of the above game:

		Student 2		
		Purdue	Indiana	Notre Dame
	Purdue	3, 2	0, 0	0, 0
Student 1	Indiana	0, 0	2, 1	0, 0
	Notre Dame	0, 0	0, 0	1, 3

The Nash equilibria are (Purdue, Purdue), (Indiana, Indiana), and (Notre Dame, Notre Dame).

(c) Consider a soccer penalty kick between a Scorer and a Goalie. The Scorer is

[1] When writing down an equilibrium, always write the strategies of Player 1 (Row Player) first.

stronger kicking to his/her Right than to his/her Left. Given the 12 yards between the ball and the goal, the Goalie must pre-determine which way he/she will dive. Thus, the actions are chosen simultaneously: a Scorer can shoot either Left or Right and the Goalie can dive either Left or Right. If the Goalie guesses wrong, the Scorer always scores. If the Goalie guesses correctly with the Scorer kicking Right, then 50% of the shots are stopped. If the Goalie guesses correctly with the Scorer kicking Left, then 90% of the shots are stopped. (Dixit and Nalebuff, 2008, pgs. 143-151).

Solution: The payoffs can vary, but one set of possible payoff values is:

		Goalie	
		Dive Left	Dive Right
Scorer	Shoot Left	1, 0	0.1, 0.9
	Shoot Right	0.5, 0.5	1, 0

Using the underline method, there does not exist a Nash equilibrium (pure-strategy) of the above game:

		Goalie	
		Dive Left	Dive Right
Scorer	Shoot Left	$\underline{1}$, 0	0.1, $\underline{0.9}$
	Shoot Right	0.5, $\underline{0.5}$	$\underline{1}$, 0

4. In the Room Assignment Game, suppose that the payoffs for the rooms are now $a = 20$, $b = 12$, and $c = 8$. What is the Nash equilibrium of the game?

Solution: First observe that all students will either select the ranking 123 or 213. Listing Room 3 any higher than the third choice is pointless as the student is guaranteed a room at least that good.

Is it optimal for any of the students to list ranking 213? If one were to do so, she would be guaranteed a payoff of $b = 12$ (no lottery for Room 2). However, by submitting the ranking 123, even when the other two roommates do the same, each has an expected payoff of

$$\frac{1}{3}(20) + \frac{1}{3}(12) + \frac{1}{3}(8) = \frac{40}{3}.$$

As the expected payoff from ranking 123 exceeds that from 213, then it is never best for any student to list ranking 213.

Thus, the Nash equilibrium (the strategies for all three roommates) is (123, 123, 123).

5. The town of Lakesville is located next to the town of West Lakesville. The Chinese restaurant Lin's is located in Lakesville and the Chinese restaurant Wong's is located in West Lakesville. Each restaurant currently delivers take-out orders within its town only. Both restaurants are simultaneously deciding whether or not to expand their delivery service to the neighboring town (Lin's to offer delivery to West Lakesville and Wong's to Lakesville).

   A restaurant will earn \$25 for selling take-out food in its own town and \$15 for selling take-out food in the other town (a \$10 travel cost is already included in the specified earnings). If a restaurant decides to expand its delivery service, a fixed cost of \$10 must be paid (to hire an additional driver).

   If both expand their delivery service, they both maintain their current customers in their own town. If one expands and the other does not, then the one that expands sells to all consumers in both towns. The game is depicted below.

		Wong's	
		Expand	Not expand
Lin's	Expand	15, 15	30, 0
	Not expand	0, 30	25, 25

   Find the Nash equilibria of this game.

   Solution: Using the underline method, the unique pure-strategy Nash equilibrium is (Expand, Expand).

		Wong's	
		Expand	Not expand
Lin's	Expand	$\underline{15}$, $\underline{15}$	$\underline{30}$, 0
	Not expand	0, $\underline{30}$	25, 25

   We know that there does not exist a mixed-strategy Nash equilibrium for this game (games with a unique pure-strategy Nash equilibrium do not have a mixed-strategy Nash equilibrium). The game is akin to Prisoners' Dilemma.

6. In the previous problem, we assumed that consumers would make certain choices without actually including them in the game. Let's correct this. Suppose each town only has one consumer of Chinese take-out food. Each consumer must select a Chinese restaurant. Exactly as in the previous problem, the restaurants must decide whether or not to expand delivery service. All decisions (both those of the restaurants and those of the consumers) are made simultaneously.

   If a consumer selects a restaurant that does deliver take-out food to that consumer, then the consumer has a payoff of 1. If a consumer selects a restaurant that does not deliver take-out food to that consumer, then no sale is made and the consumer has a payoff of 0.

   The payoffs for the restaurants are determined from the second paragraph of the previous exercise.

   One of the following two sets of strategies is a Nash equilibrium. Which one is a Nash equilibrium? Justify your answer with sound reasoning.

Option 1

Lin's Restaurant	Expand
Wong's Restaurant	Expand
Lakesville consumer	Lin's
W. Lakesville consumer	Wong's

Option 2

Lin's Restaurant	Expand
Wong's Restaurant	Expand
Lakesville consumer	Wong's
W. Lakesville consumer	Lin's

Solution: The problem can be solved in two ways. First, you can show that none of the players have an incentive to deviate from the above strategies. The other option provides a more complete solution. This second option requires recognizing that the above strategies are symmetric.

Let's consider Option 1 first. We fix the strategies (Lin's = Expand, Lakesville consumer = Lin's) and attempt to verify that the following is optimal for Wong's and the West Lakesville consumer: (Wong's = Expand, West Lakesville consumer = Wong's). The 2-player game between Wong's and the West Lakesville consumer is given by:

		Wong's: Expand	Wong's: Not expand
West Lakesville consumer	Lin's	1, -10	1, 0
	Wong's	1, 15	1, 25

.

Using the underline method, we see that the following is NOT optimal for Wong's and West Lakesville: (Wong's = Expand, West Lakesville consumer = Wong's):

		Wong's	
		Expand	Not expand
West Lakesville consumer	Lin's	$\underline{1}$, -10	$\underline{1}$, $\underline{0}$
	Wong's	$\underline{1}$, 15	$\underline{1}$, $\underline{25}$

Thus, Option 1 is NOT a Nash equilibrium.

Let's now consider Option 2. We fix the strategies (Lin's = Expand, Lakesville consumer = Wong's) and verify that the following is optimal for Wong's and the West Lakesville consumer: (Wong's = Expand, West Lakesville consumer = Lin's). The 2-player game between Wong's and the West Lakesville consumer is given by:

		Wong's	
		Expand	Not expand
West Lakesville consumer	Lin's	1, 5	1, 0
	Wong's	1, 30	1, 25

Using the underline method, we have verified that the following is optimal for Wong's and West Lakesville: (Wong's = Expand, West Lakesville consumer = Lin's):

		Wong's	
		Expand	Not expand
West Lakesville consumer	Lin's	$\underline{1}$, $\underline{5}$	$\underline{1}$, 0
	Wong's	$\underline{1}$, $\underline{30}$	$\underline{1}$, 25

Thus, Option 2 is a Nash equilibrium.

A.2 Chapter 2 Solutions

1. Consider a political economy model where the possible policies are ordered on the line from -1 to 1. All citizens have preferences lying somewhere on the line. All citizens must vote and they vote for the candidate whose platform lies closest to their preference point. Different from above, suppose that there are 3 candidates that must

simultaneously select their platform. The candidate receiving the largest number of votes is the winner. Verify that there does not exist a Nash equilibrium.

Solution: Implicit in the statement of the problem should be that fewer than $\frac{1}{3}$ of the citizens have their favorite position equal to the median position. The proof that a Nash equilibrium (pure-strategy) does not exist follows by considering several cases.

<u>Case I:</u> All candidates select the same platform, different from the median position.

In this case, all candidates receive $\frac{1}{3}$ of the citizens' votes. This cannot be an equilibrium. Each candidate has an incentive to move slightly closer to the median position. This allows that candidate to receive more than $\frac{1}{2}$ of the citizens' votes.

<u>Case II:</u> All candidates select the same platform, equal to the median position.

In this case, all candidates receive $\frac{1}{3}$ of the citizens' votes. Given the assumption that fewer than $\frac{1}{3}$ of the citizens have favorite position equal to the median position, then more than $\frac{2}{3}$ of the citizens have favorite positions that lie away from the median. Thus, one side of the median must be such that more than $\frac{1}{3}$ of the citizens have favorite positions on that side of the median. Each candidate then has an incentive to move slightly away from the median to that side. This allows that candidate to receive more than $\frac{1}{3}$ of the citizens' votes.

<u>Case III:</u> Two candidates have the same platform; the third selects a different one.

Suppose that candidates 1 and 2 have the same platform, while candidate 3 has a different one. In this case, candidate 3 always has an incentive to move closer to the platform of candidates 1 and 2 as this vote-stealing action nets more votes for candidate 3.

<u>Case IV:</u> All three candidates have different platforms.

The selected platforms may be ordered so that the platform for candidate 1 is left of the platform for candidate 2, which is left of the platform for candidate 3. Both candidate 1 and candidate 3 have an incentive to move closer to the platform of candidate 2. This vote-stealing action nets more votes for these candidates.

2. Consider a Cournot duopoly in which the two firms simultaneously choose a quantity to produce: q_1 for firm 1 and q_2 for firm 2. The inverse demand function for this market is $P(Q) = 48 - \frac{1}{2}Q$, where Q is the total quantity in the market: $Q = q_1 + q_2$. The marginal cost of production is $c_1 = 6$ for firm 1 and $c_2 = 12$ for firm 2.

Solve for the Nash equilibrium of this game (the quantity choices of both firms). What are the profits for both firms?

Solution: We solve the problem in three steps. First, we write down the profit functions for both firms:

$$\begin{aligned}
\pi_1(q_1, q_2) &= \left(48 - \frac{1}{2}q_1 - \frac{1}{2}q_2\right) q_1 - 6q_1. \\
\pi_2(q_1, q_2) &= \left(48 - \frac{1}{2}q_1 - \frac{1}{2}q_2\right) q_2 - 12q_2.
\end{aligned}$$

The second step is to find the best response functions for both firms. This involves taking the first order conditions:

$$\begin{aligned}
\frac{\partial \pi_1(q_1, q_2)}{\partial q_1} &= \left(48 - \frac{1}{2}q_1 - \frac{1}{2}q_2\right) - \frac{1}{2}q_1 - 6 = 0. \\
\frac{\partial \pi_2(q_1, q_2)}{\partial q_2} &= \left(48 - \frac{1}{2}q_1 - \frac{1}{2}q_2\right) - \frac{1}{2}q_2 - 12 = 0.
\end{aligned}$$

Solving the initial first order condition for q_1 yields the best response function for firm 1:

$$q_1 = 42 - \frac{1}{2}q_2.$$

Solving the second first order condition for q_2 yields the best response function for firm 2:

$$q_2 = 36 - \frac{1}{2}q_1.$$

The third step is to solve for the Nash equilibrium. This involves solving a system of two equations in two unknowns.

$$\begin{aligned}
q_1 &= 42 - \frac{1}{2}\left(36 - \frac{1}{2}q_1\right) \\
\frac{3}{4}q_1 &= 24 \\
q_1 &= 32 \\
q_2 &= 36 - \frac{1}{2}(32) = 20.
\end{aligned}$$

The Nash equilibrium is $(q_1 = 32, q_2 = 20)$. The profits are given by:

$$\begin{aligned}\pi_1(q_1,q_2) &= \left(48 - \frac{1}{2}(32) - \frac{1}{2}(20)\right)(32) - 6(32) = 16 \cdot 32 = 512.\\ \pi_2(q_1,q_2) &= \left(48 - \frac{1}{2}(32) - \frac{1}{2}(20)\right)(20) - 12(20) = 10 \cdot 20 = 200.\end{aligned}$$

3. Consider a Cournot duopoly in which firms compete in a single good market by simultaneously choosing an output quantity. The inverse demand function is given by

$$P(Q) = \left\{\begin{array}{cc} 60 - 2Q & \text{if } Q \leq 30 \\ 0 & \text{if } Q > 30 \end{array}\right\}.$$

The marginal cost of production can be either high $c_{high} = 12$ or low $c_{low} = 6$.

(a) Two firms

Solution: Let the market contain two firms, the first with high marginal cost of production, $c_1 = c_{high} = 12$, and the second with low marginal cost of production, $c_2 = c_{low} = 6$. Solve for the equilibrium output decisions of the two firms. What are the profits of each firm?

We begin with the profit function for firm 1 :

$$\pi_1(q_1) = (60 - 2q_1 - 2q_2)\, q_1 - 12q_1.$$

The first order conditions are given by:

$$\begin{aligned}\tfrac{\partial \pi_1(q_1)}{\partial q_1} &= 60 - 2q_1 - 2q_2 - 2q_1 - 12 = 0.\\ 48 - 4q_1 - 2q_2 &= 0.\end{aligned}$$

Solving for q_1 as a function of q_2 yields the best response for firm 1 :

$$q_1 = \frac{48 - 2q_2}{4}.$$

For firm 2, the profit function is:

$$\pi_2(q_2) = (60 - 2q_1 - 2q_2)\, q_2 - 6q_2.$$

The first order conditions are given by:

$$\frac{\partial \pi_2(q_2)}{\partial q_2} = 60 - 2q_1 - 2q_2 - 2q_2 - 6 = 0.$$
$$54 - 2q_1 - 4q_2 = 0.$$

Solving for q_2 as a function of q_1 yields the best response for firm 2 :

$$q_2 = \frac{54 - 2q_1}{4}.$$

An equilibrium (best response to a best response) is (q_1, q_2) so that the following two equations hold:

$$\begin{aligned} q_1 &= \frac{48 - 2q_2}{4}. \\ q_2 &= \frac{54 - 2q_1}{4}. \end{aligned}$$

Substituting the second equation into the first and solving for q_1 :

$$\begin{aligned} q_1 &= \frac{48 - 2\left(\frac{54-2q_1}{4}\right)}{4}. \\ q_1 &= \frac{48 - 27 + q_1}{4}. \\ \frac{3}{4}q_1 &= \frac{21}{4}. \\ q_1 &= 7. \end{aligned}$$

Solving for q_2 :

$$\begin{aligned} q_2 &= \frac{54 - 2(7)}{4} \\ q_2 &= 10. \end{aligned}$$

The profits for each firm are given by:

$$\begin{aligned} \pi_1(q_1) &= (60 - 2 \cdot 7 - 2 \cdot 10) \cdot 7 - 12 \cdot 7. \\ &= (26) \cdot 7 - 12 \cdot 7 = 98. \\ \pi_2(q_2) &= (26) \cdot 10 - 6 \cdot 10 = 200. \end{aligned}$$

(b) $2n$ firms

Let the number of firms be $2n$ where $n = 1, 2, ..., \infty$. For the odd-numbered firms, the cost of production is high, $c_i = c_{high} = 12$ for $i = 1, 3, ..., 2n-1$. For the even-numbered firms, the cost of production is low, $c_i = c_{low} = 6$ for $i = 2, 4, ..., 2n$. As n grows, do both high cost and low cost firms continue to produce (yes or no)? (Hint: the best response functions have a lower bound of $q = 0$). If one half of the firms stops producing, what is the value of n at which this happens? What is the output for each of the other half of the firms at this point? What is the equilibrium market price P at this point?

Solution: Using the symmetry of the problem, the equilibrium output of the odd-numbered firms is identical, $q_1 = q_3 = ... = q_{2n-1}$ and the equilibrium output of the even-numbered firms is identical, $q_2 = q_4 = ... = q_{2n}$. We begin with the profit function of an odd-numbered firm, for simplicity let this firm be firm 1 :

$$\pi_1(q_1) = \left(60 - 2\sum_{i=1}^{2n} q_i\right) q_1 - 12q_1.$$

The first order conditions are given by:

$$\frac{\partial \pi_1(q_1)}{\partial q_1} = 60 - 2\sum_{i=1}^{2n} q_i - 2q_1 - 12 = 0.$$

$$48 - 4q_1 - 2\sum_{i=2}^{2n} q_i = 0.$$

Solving for q_1 as a function of the other output choices yields the best response for firm 1 :

$$q_1 = \frac{48 - 2\sum_{i=2}^{2n} q_i}{4}.$$

For an even-numbered firm, for simplicity let this firm be firm 2, the profit function is:

$$\pi_2(q_2) = \left(60 - 2\sum_{i=1}^{2n} q_i\right) q_2 - 6q_2.$$

The first order conditions are given by:

$$\frac{\partial \pi_2(q_2)}{\partial q_2} = 60 - 2\sum_{i=1}^{2n} q_i - 2q_2 - 6 = 0.$$
$$54 - 2q_1 - 4q_2 - 2\sum_{i=3}^{2n} q_i = 0.$$

Solving for q_2 as a function of the other output choies yields the best response for firm 2 :

$$q_2 = \frac{54 - 2q_1 - 2\sum_{i=3}^{2n} q_i}{4}.$$

Using the symmetry, the best responses can be simplified to:

$$\begin{aligned} q_1 &= \frac{48 - 2\left[nq_2 + (n-1)q_1\right]}{4}. \\ q_2 &= \frac{54 - 2\left[nq_1 + (n-1)q_2\right]}{4}. \end{aligned}$$

From the best response for firm 1:

$$\begin{aligned} &4q_1 = 48 - 2nq_2 - 2\left(n-1\right)q_1 \\ &4q_1 + 2\left(n-1\right)q_1 + 2nq_2 = 48 \ . \\ &2\left(n+1\right)q_1 + 2nq_2 = 48 \end{aligned}$$

From the best response for firm 2:

$$\begin{aligned} &4q_2 = 54 - 2nq_1 - 2\left(n-1\right)q_2 \\ &4q_2 + 2\left(n-1\right)q_2 + 2nq_1 = 54 \ . \\ &2\left(n+1\right)q_2 + 2nq_1 = 54 \end{aligned}$$

An equilibrium (best response to a best response) is (q_1, q_2) so that the following two equations hold:

$$\begin{aligned} 2\left(n+1\right)q_1 + 2nq_2 = 48 \\ 2\left(n+1\right)q_2 + 2nq_1 = 54 \end{aligned} .$$

Subtracting the second equation from the first equation, we arrive at:

$$2q_1 - 2q_2 = -6,$$

so $q_2 = q_1 + 3$. Substituting this into the equation $2(n+1)q_1 + 2nq_2 = 48$ and solving for q_1 :

$$\begin{aligned} 2(n+1)q_1 + 2n(q_1+3) &= 48 \\ q_1(2(n+1)+2n) &= 48-6n \\ q_1 &= \frac{48-6n}{2(n+1)+2n} \end{aligned}$$

or (after factoring out 2 from both numerator and denominator):

$$q_1 = \frac{24-3n}{2n+1}.$$

Solving for q_2 :

$$\begin{aligned} q_2 &= \frac{24-3n}{2n+1} + 3 \\ q_2 &= \frac{24-3n}{2n+1} + \frac{3(2n+1)}{2n+1} \\ q_2 &= \frac{3n+27}{2n+1}. \end{aligned}$$

Notice that $q_1 \geq 0$ only if $n \leq 8$. Thus, when $n = 8$, the odd-numbered firms (those with the high costs) cease producing. The output for the even-numbered firms when $n = 8$ is given by:

$$q_2 = \frac{3(8)+27}{2(8)+1} = 3.$$

When $n = 8$, the total quantity produced is $Q = 8q_2 = 24$, so the equilibrium market price is:

$$P = 60 - 2Q = 60 - 48 = 12,$$

which is exactly the marginal cost for the odd-numbered firms who have chosen to cease production when $n = 8$.

4. Consider the Bertrand model of duopoly with transportation costs. Two firms compete in prices by simultaneously selecting prices p_1 and p_2. Firm 1 is located at mile 0 and firm 2 is located at mile 1. Consumers are uniformly distributed between mile 0 and mile 1 and incur a cost of 1 for each mile traveled to reach the selected firm. Consumers

purchase from the firm with the lowest total cost (price plus travel cost) and have a demand equal to 1.

The firms have marginal costs of production given by c_1 and c_2. The analysis in class was conducted under the implicit assumption that $|c_1 - c_2| \leq 3$. What happens if $|c_1 - c_2| > 3$, say $c_1 = 8$ and $c_2 = 4$? Specifically, what are the equilibrium price choices p_1 and p_2 and the profits of each firm?

Solution: From the solutions in the chapter, the optimal choices for prices p_1 and p_2 are:

$$\begin{aligned} p_1 &= 1 + \frac{2}{3}c_1 + \frac{1}{2}c_2. \\ p_2 &= 1 + \frac{2}{3}c_2 + \frac{1}{2}c_1. \end{aligned}$$

At these prices, the value for t is given by:

$$\begin{aligned} t &= \frac{p_2 - p_1 + 1}{2}. \\ &= \frac{\left(1 + \frac{2}{3}c_2 + \frac{1}{2}c_1\right) - \left(1 + \frac{2}{3}c_1 + \frac{1}{2}c_2\right) + 1}{2}. \\ &= \frac{1}{6}c_2 - \frac{1}{6}c_1 + \frac{1}{2}. \end{aligned}$$

But the analysis has required that $t : 0 \leq t \leq 1$. This is only the case when:

$$\begin{aligned} 0 &\leq \frac{1}{6}c_2 - \frac{1}{6}c_1 + \frac{1}{2} \leq 1. \\ -\frac{1}{2} &\leq \frac{1}{6}c_2 - \frac{1}{6}c_1 \leq \frac{1}{2}. \\ -3 &\leq c_2 - c_1 \leq 3. \end{aligned}$$

In the specified case, $c_2 - c_1 = -4$, so we are outside the scope of the analysis laid out in class. That means we have to return to the logic.

Firm 1 only makes nonnegative profit when $p_1 \geq c_1 = 8$. This constraint is not considered during the analysis from class. The first inclination is that firm 2 will set a price to ensure that it sells to the whole market. The highest price at which firm 2 can ensure all the sales is $p_2 = 7 - 0.0000001$. With this price, even the consumers at mile 0 prefer to pay the price p_2 with travel cost 1 than pay p_1. The profit to firm 2 from this strategy is

$$\pi_2 = (p_2 - c_2) \cdot 1 \cong (7-4) \cdot 1 = 3.$$

Thus, I am claiming the equilibrium is $p_1 = 8$ and $p_2 = 7$. Is each firm, given the price choice of the other firm, behaving optimally? From Firm 1's point of view, it earns zero profit with $p_1 \geq 8$ and strictly negative profit with $p_1 < 8$, so Firm 1 is behaving optimally. From Firm 2's point of view, it is selling to all consumers at $p_2 = 7$. By lowering the price to $p_2 < 7$, firm 2 still sells to all consumers, but at a lower profit per unit. By raising the price to $p_2 > 7$, then the profit of firm 2 would be given by:

$$\pi_2 = (p_2 - c_2) \cdot (1-t).$$

They would now not sell to all consumers, but only $1-t$ of the consumers where (holding the choice of firm 1 fixed at $p_1 = 8$):

$$\begin{aligned} 1-t &= 1 - \frac{p_2 - p_1 + 1}{2} \\ &= 1 - \frac{p_2 - 8 + 1}{2} \\ &= \frac{9}{2} - \frac{p_2}{2}. \end{aligned}$$

The profit is then:

$$\pi_2 = (p_2 - 4) \cdot \left(\frac{9}{2} - \frac{p_2}{2}\right).$$

Taking the first order condition yields:

$$\frac{\partial \pi_2}{\partial p_2} = \left(\frac{9}{2} - \frac{p_2}{2}\right) - \left(\frac{p_2 - 4}{2}\right) = 0.$$

Solving for p_2 yields $p_2 = \frac{13}{2}$. But recall we are only considering if firm 2 finds it optimal to sell at price $p_2 > 7$. Thus, firm 2 wishes to sell at a price p_2 as low as possible, but not lower than $p_2 = 7$ (as previously argued). Thus, it is optimal (given firm 1 chooses $p_1 = 8$) for firm 2 to choose $p_2 = 7$. These choices constitute an equilibrium.

5. Consider the following update of the Beach Game. Vendor 1 is initially located at 0 and vendor 2 is initially located at 1. Both vendors incure a cost of k per mile for moving to a new location. Given the locations of the two vendors (n_1, n_2), the demand function of the consumers located at z along the beach (where $0 \leq z \leq 1$) is defined

as:

$$demand_i(z) = \begin{Bmatrix} 1 - |n_i - z| & \text{if } |n_i - z| < |n_j - z| \\ \frac{1}{2}(1 - |n_i - z|) & \text{if } |n_i - z| = |n_j - z| \\ 0 & \text{if } |n_i - z| > |n_j - z| \end{Bmatrix}$$

where $(i, j) = (1, 2)$ and $(i, j) = (2, 1)$.

Both vendors select their location simultaneously (cannot observe the location of the other vendor prior to making their decision).

As a function of $k : 0 \leq k \leq \frac{3}{4}$, solve for the equilibrium vendor positions of this game. Verify that the result obtained for the special case of $k = 0$ is $(n_1, n_2) = \left(\frac{1}{2}, \frac{1}{2}\right)$ (the answer obtained above for the zero cost case).

Solution: The solution method is to take the payoff functions of each vendor and then take first order conditions to find the best response functions. A best response function for vendor i specifies, as a function of vendor $j's$ location, the best location for vendor i. An equilibrium is defined as each vendor choosing a best response to the other's best response. Mathematically, this is simply that both best response functions need to be satisfied. This provides us with two equations in the two unknowns, n_1 and n_2.

The math is as follows. First, the payoff function for vendor 1 :

$$\begin{aligned} &p_1(n_1, n_2) = A + B - kn_1. \\ &= \tfrac{1}{2}(1 - n_1 + 1)\, n_1 + \tfrac{1}{2}\left(1 + 1 - \tfrac{(n_2 - n_1)}{2}\right)\left[\tfrac{(n_2 - n_1)}{2}\right] - kn_1. \\ &= n_1 - \tfrac{1}{2}(n_1)^2 + \left[\tfrac{(n_2 - n_1)}{2}\right] - \tfrac{1}{2}\left[\tfrac{(n_2 - n_1)}{2}\right]^2 - kn_1. \end{aligned}$$

The first order condition of the payoff function is:

$$\begin{aligned} \tfrac{\partial p_1(n_1, n_2)}{\partial n_1} &= 1 - n_1 - \tfrac{1}{2} + \tfrac{1}{2}\left[\tfrac{(n_2 - n_1)}{2}\right] - k = 0. \\ \tfrac{\partial p_1(n_1, n_2)}{\partial n_1} &= \tfrac{1}{2} - \tfrac{5}{4}n_1 + \tfrac{1}{4}n_2 - k = 0. \end{aligned}$$

Solving for n_1 provides the best response function:

$$n_1 = \frac{1}{5}n_2 + \frac{2}{5} - \frac{4}{5}k.$$

Now, the problem for vendor 2. First, the payoff function:

$$\begin{aligned}
&p_2(n_1,n_2) = C + D - k(1-n_2).\\
&= \tfrac{1}{2}\left(1+1-\tfrac{(n_2-n_1)}{2}\right)\left[\tfrac{(n_2-n_1)}{2}\right] + \tfrac{1}{2}(1+n_2)(1-n_2) - k(1-n_2).\\
&= \left[\tfrac{(n_2-n_1)}{2}\right] - \tfrac{1}{2}\left[\tfrac{(n_2-n_1)}{2}\right]^2 + \tfrac{1}{2}\left(1-(n_2)^2\right) - k(1-n_2).
\end{aligned}$$

The first order condition for the payoff function is:

$$\begin{aligned}
\tfrac{\partial p_2(n_1,n_2)}{\partial n_2} &= \tfrac{1}{2} - \tfrac{1}{2}\left[\tfrac{(n_2-n_1)}{2}\right] - n_2 + k = 0.\\
\tfrac{\partial p_2(n_1,n_2)}{\partial n_2} &= \tfrac{1}{2} - \tfrac{5}{4}n_2 + \tfrac{1}{4}n_1 + k = 0.
\end{aligned}$$

Solving for n_2 provides the best response function:

$$n_2 = \frac{1}{5}n_1 + \frac{2}{5} + \frac{4}{5}k.$$

We have two equations in two unknowns:

$$\begin{aligned}
n_1 &= \frac{1}{5}n_2 + \frac{2}{5} - \frac{4}{5}k.\\
n_2 &= \frac{1}{5}n_1 + \frac{2}{5} + \frac{4}{5}k.
\end{aligned}$$

Substituting the second equation into the first:

$$\begin{aligned}
n_1 &= \frac{1}{5}\left(\frac{1}{5}n_1 + \frac{2}{5} + \frac{4}{5}k\right) + \frac{2}{5} - \frac{4}{5}k.\\
&= \frac{1}{25}n_1 + \frac{2}{25} + \frac{4}{25}k + \frac{2}{5} - \frac{4}{5}k.
\end{aligned}$$

Solving for n_1 :

$$\begin{aligned}
\frac{24}{25}n_1 &= \frac{12}{25} - \frac{16}{25}k.\\
n_1 &= \frac{1}{2} - \frac{2}{3}k.
\end{aligned}$$

Solving for n_2 :

$$\begin{aligned} n_2 &= \frac{1}{5}\left(\frac{1}{2} - \frac{2}{3}k\right) + \frac{2}{5} + \frac{4}{5}k \\ &= \frac{1}{10} - \frac{2}{15}k + \frac{2}{5} + \frac{4}{5}k \\ &= \frac{1}{2} + \frac{2}{3}k. \end{aligned}$$

The equilibrium locations, as a function of $k : 0 \leq k \leq \frac{3}{4}$, are:

$$\begin{aligned} n_1 &= \frac{1}{2} - \frac{2}{3}k. \\ n_2 &= \frac{1}{2} + \frac{2}{3}k. \end{aligned}$$

A.3 Chapter 3 Solutions

1. Find the mixed-strategy Nash equilibria in the following figures of Chapter 1: Figures 1.20, 1.21, 1.22, 1.23, and 1.24.

Figure 1.20 Solution: Define p as the probability that the man chooses opera and q as the probability that the woman chooses opera. Then, p is determined to make the woman indifferent between opera and boxing:

$$\begin{aligned} \text{Woman's opera payoff} &= \text{Woman's boxing payoff} \\ 60p + 0\,(1-p) &= 0p + 30(1-p) \\ 60p &= 30(1-p) \\ p &= \frac{1}{3}. \end{aligned}$$

Similarly, q is determined to make the man indifferent between opera and boxing:

$$\begin{aligned} \text{Man's opera payoff} &= \text{Man's boxing payoff} \\ 40q + 0\,(1-q) &= 0q + 70(1-q) \\ 40q &= 70(1-q) \\ q &= \frac{7}{11}. \end{aligned}$$

The mixed-strategy Nash equilibrium is

$$((p, 1-p), (q, 1-q)) = \left(\left(\frac{1}{3}, \frac{2}{3}\right), \left(\frac{7}{11}, \frac{4}{11}\right)\right).$$

Figure 1.21 Solution: Define p as the probability that hunter 1 chooses stag and q as the probability that hunter 2 chooses stag. Then, p is determined to make hunter 2 indifferent between stag and hare:

$$\begin{aligned} \text{Hunter 2 stag payoff} &= \text{Hunter 2 hare payoff} \\ 2p + 0(1-p) &= 1p + 1(1-p) \\ 2p &= 1 \\ p &= \frac{1}{2}. \end{aligned}$$

Similarly, q is determined to make hunter 1 indifferent between stag and hare:

$$\begin{aligned} \text{Hunter 1 stag payoff} &= \text{Hunter 1 hare payoff} \\ 2q + 0(1-q) &= 1q + 1(1-q) \\ 2q &= 1 \\ q &= \frac{1}{2}. \end{aligned}$$

The mixed-strategy Nash equilibrium is

$$((p, 1-p), (q, 1-q)) = \left(\left(\frac{1}{2}, \frac{1}{2}\right), \left(\frac{1}{2}, \frac{1}{2}\right)\right).$$

Figure 1.22 Solution: Define p as the probability that driver 1 chooses straight and q as the probability that driver 2 chooses straight. Then, p is determined to make driver 2 indifferent between straight and swerve:

$$\begin{aligned} \text{Driver 2 straight payoff} &= \text{Driver 2 swerve payoff} \\ 0p + 10(1-p) &= 1p + 5(1-p) \\ 10 - 10p &= 5 - 4p \\ p &= \frac{5}{6}. \end{aligned}$$

Similarly, q is determined to make driver 1 indifferent between straight and swerve:

$$\begin{aligned}\text{Driver 1 straight payoff} &= \text{Driver 1 swerve payoff} \\ 0q + 10\,(1-q) &= 1q + 5(1-q) \\ 10 - 10q &= 5 - 4q \\ q &= \frac{5}{6}.\end{aligned}$$

The mixed-strategy Nash equilibrium is

$$((p, 1-p), (q, 1-q)) = \left(\left(\frac{5}{6}, \frac{1}{6}\right), \left(\frac{5}{6}, \frac{1}{6}\right)\right).$$

Figure 1.23 Solution: Define p as the probability that player 1 chooses heads and q as the probability that player 2 chooses heads. Then, p is determined to make player 2 indifferent between heads and tails:

$$\begin{aligned}\text{Player 2 heads payoff} &= \text{Player 2 tails payoff} \\ -1p + 1\,(1-p) &= 1p - 1(1-p) \\ 1 - 2p &= 2p - 1 \\ p &= \frac{1}{2}.\end{aligned}$$

Similarly, q is determined to make player 1 indifferent between heads and tails:

$$\begin{aligned}\text{Player 1 heads payoff} &= \text{Player 1 tails payoff} \\ 1q - 1\,(1-q) &= -1q + 1(1-q) \\ 2q - 1 &= 1 - 2q \\ q &= \frac{1}{2}.\end{aligned}$$

The mixed-strategy Nash equilibrium is

$$((p, 1-p), (q, 1-q)) = \left(\left(\frac{1}{2}, \frac{1}{2}\right), \left(\frac{1}{2}, \frac{1}{2}\right)\right).$$

Figure 1.24 Solution: Define $(p_1, p_2, 1-p_1-p_2)$ as the probabilities that player 1 assigns to (Rock, Paper, Scissors), respectively, and define $(q_1, q_2, 1-q_1-q_2)$ as the probabilities that player 2 assigns to (Rock, Paper, Scissors), respectively. The values (p_1, p_2) are set to make player 2 indifferent between all three: rock, paper, and scissors:

Payoff of rock	$-p_2 + (1 - p_1 - p_2)$
Payoff of paper	$p_1 - (1 - p_1 - p_2)$
Payoff of scissors	$-p_1 + p_2$

The equality between paper and scissors leads to the following equation:

$$\begin{aligned} p_1 - (1 - p_1 - p_2) &= -p_1 + p_2 \\ 2p_1 + p_2 - 1 &= -p_1 + p_2 \\ p_1 &= \frac{1}{3}. \end{aligned}$$

The equality between rock and scissors leads to the following equation:

$$\begin{aligned} -p_2 + (1 - p_1 - p_2) &= -p_1 + p_2 \\ 1 - p_1 - 2p_2 &= -p_1 + p_2 \\ p_2 &= \frac{1}{3}. \end{aligned}$$

By symmetry, the same values are found for $(q_1, q_2, 1 - q_1 - q_2)$. The mixed-strategy Nash equilibrium is

$$\begin{aligned} (p_1, p_2, 1 - p_1 - p_2) &= \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right). \\ (q_1, q_2, 1 - q_1 - q_2) &= \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right). \end{aligned}$$

2. Find the mixed-strategy Nash equilibria of the normal form games that you created for Chapter 1, Question 3, parts (a) and (c) (we are not quite ready to handle part (b)).

Part (a) Solution: Recall the game that I wrote down is as follows (yours will likely be different):

		Animal 2	
		Passive	Aggresive
Animal	Passive	0, 0	1, 3
1	Aggresive	3, 1	-2, -2

To find the mixed-strategy Nash equilibrium, let p be the probability with which Animal 1 plays "Passive." Given the symmetry of the payoff values, if q is the probability with which Animal 2 plays "Passive," then $q = p$ in equilibrium.

The equilibrium value for p is found so that Animal 2 is indifferent between playing "Passive" and "Aggresive." The expected payoffs of the two options are:

$$\text{Expected payoff of "Passive"} = \text{Expected payoff of "Aggresive"}$$
$$0 \cdot p + 1 \cdot (1-p) = 3 \cdot p + (-2) \cdot (1-p)$$
$$1 - p = 3p - 2(1-p)$$

Solving for p yields:

$$\begin{aligned} 1 - p &= 3p - 2 + 2p \\ 3 &= 6p \\ p &= \frac{1}{2}. \end{aligned}$$

Thus, the mixed-strategy Nash equilibrium is $\left(\left(\frac{1}{2}, \frac{1}{2}\right), \left(\frac{1}{2}, \frac{1}{2}\right)\right)$.

Part (c) Solution: Recall the game that I wrote down is as follows (yours will likely be different):

		Goalie	
		Dive Left	Dive Right
Scorer	Shoot Left	1, 0	0.1, 0.9
	Shoot Right	0.5, 0.5	1, 0

To find the mixed-strategy Nash equilibrium, let p be the probability with which Scorer plays "Shoot Left."

The equilibrium value for p is found so that Goalie is indifferent between playing "Dive

Left" and "Dive Right." The expected payoffs for Goalie of the two options are:

$$\text{Expected payoff of "Dive Left"} = \text{Expected payoff of "Dive Right"}$$
$$0 \cdot p + 0.5 \cdot (1-p) = 0.9 \cdot p + 0 \cdot (1-p)$$
$$0.5 - 0.5p = 0.9p$$

Solving for p yields:

$$\begin{aligned} 1.4p &= 0.5 \\ p &= \frac{5}{14}. \end{aligned}$$

Let q be the probability with which Goalie plays "Dive Left."

The equilibrium value for q is found so that Scorer is indifferent between playing "Shoot Left" and "Shoot Right." The expected payoffs for Scorer of the two options are:

$$\text{Expected payoff of "Shoot Left"} = \text{Expected payoff of "Shoot Right"}$$
$$1 \cdot q + 0.1 \cdot (1-q) = 0.5 \cdot q + 1 \cdot (1-q)$$
$$0.9q + 0.1 = -0.5q + 1$$

Solving for q yields:

$$\begin{aligned} 1.4q &= 0.9 \\ q &= \frac{9}{14}. \end{aligned}$$

Thus, the mixed-strategy Nash equilibrium is $\left(\left(\frac{5}{14}, \frac{9}{14}\right), \left(\frac{9}{14}, \frac{5}{14}\right)\right)$.

3. There exists one more mixed-strategy Nash equilibrium in Figure 3.11. Find it.

Solution: As in the text, with $q_2 > 0$, then the expected payoff of C for the column player must be equal to the expected payoff of all strategies played with strictly positive probability. The expected payoff of C is equal to $2p + 2(1-p) = 2$. We see that for the value of $p = \frac{1}{3}$, then the expected payoff of L equals $3p + (1-p) = \frac{4}{3}$, while the expected payoff of R is $0 + 3(1-p) = 2$. When $p = \frac{1}{3}$, the column player is indifferent between C and R, but strictly prefers both to L. This requires that $q_1 = 0$. As $q_1 + q_3 = \frac{1}{2}$, then it must be that $q_3 = \frac{1}{2}$. So the second mixed-strategy Nash equilibrium is:

$$(p, 1-p) = \left(\frac{1}{3}, \frac{2}{3}\right).$$
$$(q_1, q_2, q_3) = \left(0, \frac{1}{2}, \frac{1}{2}\right).$$

A.4 Chapter 4 Solutions

1. What is the subgame perfect Nash equilibrium for Figure 4.12?

   Solution: The subgame perfect Nash equilibrium of Figure 4.12 is (R, (RR,LL)), where R is the strategy of Player 1 and (RR,LL) is the strategy (both contingent actions) of Player 2.

2. Suppose that two players are bargaining over $1. The game takes place in rounds, beginning with Round 1. The game ends when an offer is accepted. Player 1 makes offers in odd-numbered rounds and Player 2 makes offers in even-numbered rounds. In rounds when a player is not making an offer, it observes the offer made by the other player and decides whether to "Accept" or "Reject." At the end of each round, $0.20 is removed from the pool of money. That is, if an agreement is reached in Round 2, the total pool of money is $0.80; if agreement in Round 3, $0.60, and so forth. (Hint: The game has only 5 periods.)

   Find the subgame perfect Nash equilibrium of this bargaining game.

   Solution: This game has a definite ending point. At the end of Round 5, no money is remaining ($0.00). Begin at this round and employ backward induction.

   In Round 5, Player 1 is making the offer. Player 1 can offer $0.00 to Player 2 in Round 5. This offer is accepted by Player 2 as he/she cannot do any better by rejecting the offer. This leaves $0.20 for Player 1.

   Moving to Round 4, Player 2 is making the offer. Player 2 can offer $0.20 to Player 1 in Round 4. This offers is accepted by Player 1, because he/she cannot do any better by rejecting the offer and waiting for Round 5. This leaves $0.20 for Player 2.

   Proceeding by backward induction:

   - Player 1 offers $0.20 to Player 2 in Round 3 (accepted).

- Player 2 offers \$0.40 in Player 1 in Round 2 (accepted).
- Player 1 offers \$0.40 to Player 2 in Round 1. This offer is accepted. This leaves Player 1 with the remaining \$0.60.

The subgame perfect Nash equilibrium must specify the complete strategies for both players (actions in all 5 rounds):

Player 1: (Offer \$0.40 to P2, Accept if offered at least \$0.40, Offer \$0.20 to P2, Accept if offered at least \$0.20, Offer \$0.00 to P2).

Player 2: (Accept if offered at least \$0.40, Offer \$0.40 to P1, Accept if offered at least \$0.20, Offer \$0.20 to P1, Accept if offered at least \$0.00).

3. In the kingdom called Hearts, consider a game with 4 players: the King of Hearts, the Queen of Hearts, the Jack of Hearts and the Ten of Hearts. The four currently share the wealth in the kingdom:

Percent of wealth

Ten	Jack	Queen	King
6%	46%	24%	24%

The law of Hearts are: (i) a proposal is voted on by all 4 players and can only be passed if it receives at least 3 votes and (ii) if a player makes a proposal that does not pass, then that player loses its entire share of the wealth (which is distributed evenly among the other players).

The voting proceeds sequentially beginning with the Ten, proceeding to the Jack, followed by the Queen, and ending with the King.

Suppose that the Ten makes the following proposal: (i) the wealth in the kingdom will be updated as

Proposed percent of wealth

Ten	Jack	Queen	King
40%	20%	20%	20%

and (ii) if the proposal passes, then any player voting against the proposal will be stripped of its wealth share (which is distributed evenly among the other players).

Will the proposal by the Ten pass? To receive full credit, you must justify your answer by writing the subgame perfect Nash equilibrium strategies for all players.

Solution: We know that Ten will always vote "Yes" for its own proposal. We can now solve the game by backward induction. The last vote is cast by the King. The King can observe (a) three "Yes" votes for the proposal, (b) two "Yes" votes for the proposal, and (c) one "Yes" vote for the proposal. For (a), with three "Yes" votes, the proposal has already passed, so the King prefers to get 20% rather than 0%, so will vote "Yes." For (b), with two "Yes" votes, the King is the decisive voter. By voting "Yes," then the proposal passes and the King receives 20% plus 1/3 of the share of the dissenting voter (which is also 20%). This amount is higher than what the King receives by voting "No" (keeping its original share of 24% plus 1/3 of the original share of Ten). For (c), with one "Yes" vote, proposal fails, so the King's vote doesn't matter.

The second to last vote is cast by the Queen. The Queen can observe (a') two "Yes" votes for the proposal, and (b') one "Yes" vote for the proposal. Under (a'), with two "Yes" votes, then the Queen knows that the King will always follow with a "Yes" vote. Thus, the proposal will pass, so it is optimal for the Queen to vote "Yes" as she prefers to get 20% rather than 0%. Under (b'), with one "Yes" vote, the Queen is the decisive voter. By voting "Yes," then the proposal passes and the Queen receives 20% plus 1/3 of the share of the dissenting voter (that share being the 20% share of the Jack). This amount is higher than what the Queen receives by voting "No" (keeping her original share of 24% plus 1/3 of the original share of Ten).

The third to last vote is cast by the Jack. The Jack is in a tough spot. The Jack observes the one "Yes" vote by the Ten. No matter what the Jack does, the proposal will pass (the King and the Queen will both vote "Yes"). Thus, it is optimal for the Jack to vote "Yes" as he prefers to get 20% rather than 0%.

In the end, the proposal passes with four "Yes" votes. The strategies of the players have been explained in the previous paragraphs.

4. Consider a market with two firms. The firms compete by choosing a quantity to produce, but this competition is sequential: firm 1 first chooses a quantity to produce q_1 and given this observed choice, firm 2 then chooses a quantity to produce q_2. The inverse demand function for this market is $P(Q) = 30 - 2Q$, where P is the unit price and Q is the total quantity in the market: $Q = q_1 + q_2$. The marginal cost of production is $c_1 = 4$ for firm 1 and $c_2 = 2$ for firm 2.

   Solve for the subgame perfect Nash equilibrium (SPNE) of this game (the quantity choices of both firms). Which firm earns a higher profit?

Solution: We solve the problem by backward induction. First, we write down the profit functions for the second player (firm 2):

$$\pi_2(q_1, q_2) = (30 - 2q_1 - 2q_2)\, q_2 - 2q_2.$$

Next, we find the best response function for firm 2. This involves taking the first order condition:

$$\frac{\partial \pi_2(q_1, q_2)}{\partial q_2} = (30 - 2q_1 - 2q_2) - 2q_2 - 2 = 0.$$

Solving the first order condition for q_2 yields the best response function for firm 2:

$$\begin{aligned} q_2 &= \frac{28 - 2q_1}{4} \\ &= \frac{14 - q_1}{2}. \end{aligned}$$

The first player (firm 1) is equally capable of solving for this best response function. Thus, when firm 1 makes it decision, it internalizes the best response function for firm 2 into its profit maximization problem:

$$\begin{aligned} \pi_1(q_1, q_2) &= \left(30 - 2q_1 - 2\left[\frac{14 - q_1}{2}\right]\right) q_1 - 4q_1. \\ &= (16 - q_1)\, q_1 - 4q_1. \end{aligned}$$

To maximize profit, firm 1 takes the first order condition:

$$\begin{aligned} \frac{\partial \pi_1(q_1, q_2(q_1))}{\partial q_1} &= (16 - q_1) - q_1 - 4 = 0. \\ 2q_1 &= 12. \\ q_1 &= 6. \end{aligned}$$

The best response choice by firm 2 is then $q_2(q_1) = \frac{14-6}{2} = 4$.

Firm 1 produces more, but which firm has the higher profit? The profits for the two firms are given by:

$$\begin{aligned} \pi_1(q_1, q_2) &= (30 - 2(6) - 2(4))(6) - 4(6) = 60 - 24 = 36. \\ \pi_2(q_1, q_2) &= (30 - 2(6) - 2(4))(4) - 2(4) = 40 - 8 = 32. \end{aligned}$$

Thus, firm 1 has the higher profit (despite having the higher marginal cost of production).

5. Consider a strategic voting game with three candidates and three voters. The candidates are Romney, Perry, and Obama. The voters first select who wins between Romney and Perry. Next, the voters select who wins between the Romney-Perry winner and Obama. In terms of the winner of the Romney-Perry winner vs. Obama election, the voters have the following rankings:

	Voter 1	Voter 2	Voter 3
1st Choice	Obama	Romney	Perry
2nd Choice	Perry	Obama	Romney
3rd Choice	Romney	Perry	Obama

The voters vote strategically. Which candidate wins in the Romney-Perry winner vs. Obama election?

Solution: We solve by backward induction. In the Romney-Perry winner vs. Obama election, all voters vote according to their rankings. In an election between Romney and Obama, the winner is Romney (receiving votes from Voters 2 and 3). In an election between Perry and Obama, the winner is Obama (receiving votes from Voters 1 and 2). Knowing this, then Perry can never win the final election, and Obama can only win if Perry is his opponent.

Proceed by backward induction to the Romney vs. Perry election. Perry cannot win the final election, so the Romney vs. Perry election is actually Romney vs. Obama. All voters vote according to their rankings for Romney and Obama. In this case, Romney wins (receiving votes from Voters 2 and 3). Thus, the Romney-Perry winner is Romney, and Romney defeats Obama in the final election.

A.5 Chapter 5 Solutions

1. Consider the following game between two investors. They have each deposited D at a bank. The bank has invested these deposits in long-term projects. If the bank is forced to liquidate its investment before the long-term project is complete, then it can only recover $2r$, where $D > r > \frac{D}{2}$. If the bank allows the project to be completed, the

total payout is higher and given by $2R$, where $R > D$.

There are two dates at which investors can make withdrawals from the bank: stage one before the long-term project is complete and stage two after the project is complete. The decisions to withdraw or not in each stage are made simultaneously by the two investors. The two-stage game described so far is equivalent to the following two normal-form games:

Stage 1

		Investor 2	
		Withdraw	Don't
1	Withdraw	r, r	$D, 2r - D$
	Don't	$2r - D, D$	Stage 2

Stage 2

		Investor 2	
		Withdraw	Don't
1	Withdraw	R, R	$2R - D, D$
	Don't	$D, 2R - D$	R, R

.

Find all the subgame perfect Nash equilibria of this game.

Solution: We solve for the subgame perfect Nash equilibria by using backward induction. The Nash equilibria of the Stage 2 game are found using the underline method:

		Inv. 2	
		Withdraw	Don't
Inv. 1	Withdraw	$\underline{R}$,$\underline{R}$	$\underline{2R - D}, D$
	Don't	D,$\underline{2R - D}$	R, R

.

Thus, we know that (R, R) will be the payoffs in Stage 2. Insert these values into the Stage 1 game and find the Nash equilibria of that (again, using the underline method):

		Inv. 2	
		Withdraw	Don't
Inv. 1	Withdraw	$\underline{r}$,$\underline{r}$	$D, 2r - D$
	Don't	$2r - D, D$	$\underline{R}$,$\underline{R}$

.

We see that there are two Nash equilibria of Stage 1 (plus one mixed-strategy Nash equilibrium that we choose to ignore). Thus, there are two subgame perfect Nash equilibria:

1. (Withdraw, Withdraw) , (Withdraw, Withdraw) .
2. (Don't, Don't) , (Withdraw, Withdraw) .

2. Consider the following game of contract negotiation between a sports team and an athlete. The negotiation takes place in 3 stages. There is no discounting of payoffs across the stages. In the first stage, both parties simultaneously choose whether the athlete receives a big contract or a small contract. If both parties agree, the contract is signed; otherwise, they proceed to the second stage. Payoffs are indicated below:

		Athlete	
		Big	Small
Sports	Big	20, 100	Stage 2
Team	Small	Stage 2	100, 20

In the second stage, the two parties simultaneously decide (given the results of Stage 1) whether to walk away from the negotiation or to continue. If at least one of them wants to walk away, then the negotiation ends. If both want to continue, then the negotiation continues to the third stage. Payoffs are indicated below:

		Athlete	
		Walk	Bargain
Sports	Walk	25, 25	10, 30
Team	Bargain	30, 10	Stage 3

The third stage is depicted in Figure 5.19 in the companion 'Figures' document. In the third stage, the sports team makes either a high or a low offer. Seeing this offer, the athlete can either accept or reject it. Either way, the negotiations end. If the offer is accepted, a contract is signed. If the offer is rejected, then both parties leave with nothing.

Solve for the subgame perfect Nash equilibrium.

Solution: The game is solved using backward induction. The solution to Stage 3 is given in Figure A.1 in the companion 'Figures' document.

We have solved the final stage of the game by backward induction. Using this result, we can fill in the payoff matrix of Stage 2 and solve for the Nash equilibrium in that stage:

		Athlete	
		Walk	Bargain
Sports	Walk	25, 25	10, 30
Team	Bargain	30, 10	80, 15

Using this result, we can fill in for the payoff matrix of Stage 1 and solve for the Nash equilibrium in that stage:

		Athlete	
		Big	Small
Sports	Big	20, 100	80, 15
Team	Small	80, 15	100, 20

Thus, the subgame perfect Nash equilibrium is:

Sports Team: (Small, Bargain, Low)

Athlete: (Small, Bargain, Accept if Low, Accept if High)

The equilibrium outcome (payoffs) is (100, 20).

3. Consider two firms in an industry. The firms are attempting to maintain an uncontractable agreement between them (about prices, quality, advertising, or possibly entrance into new markets). The firms play a game in two stages. Each stage is identical, with no discounting future payoffs. In both stages, each firm simultaneously chooses one of three actions: {Cooperate, Defect, Ignore}. The payoffs are given as:

		Firm 2		
		Cooperate	Defect	Ignore
	Cooperate	10, 10	0, 13	2, 12
Firm 1	Defect	13, 0	1, 1	0, 0
	Ignore	12, 2	0, 0	5, 5

Is it possible that a subgame perfect Nash equilibrium can be constructed in which both firms "Cooperate" in the first stage? If so, write down the complete equilibrium strategy.

Solution: First, the Nash equilibria of the stage game need to be identified (we are only interested in finding the pure-strategy Nash equilibria). They are (Defect, Defect) and (Ignore, Ignore):

		Firm 2		
		Cooperate	Defect	Ignore
	Cooperate	10, 10	0, <u>13</u>	2, 12
Firm 1	Defect	<u>13</u>, 0	<u>1</u>, <u>1</u>	0, 0
	Ignore	12, 2	0, 0	<u>5</u>, <u>5</u>

By the definition of a subgame perfect Nash equilibrium, one of these Nash equilibria must be selected in the final stage (Stage 2). Given that, is it possible for (Cooperate, Cooperate) to be played in Stage 1. Notice that both firms have an incentive to deviate in Stage 1. But, could punishment in Stage 2 effectively deter such deviation?

Consider the following symmetric strategies:

- Stage 1: The firm selects Cooperate.
- Stage 2:
  - If both firms have selected "Cooperate" in Stage 1, the firm selects "Ignore."
  - For all other Stage 1 outcomes, the firm selects "Defect."

The proposed strategy satisfies the definition of subgame perfect Nash equilibrium in Stage 2 as a Nash equilibrium is played in Stage 2 (either (Defect, Defect) or (Ignore, Ignore)).

Is the strategy optimal beginning in Stage 1? Consider the payoffs. From the strategy proposed above, the total payoffs are $10 + 5 = 15$. If one firm, say Firm 1, were to deviate in Stage 1, it would deviate by selecting "Defect" and receiving a payoff of 13. This leads to a higher payoff in Stage 1, but the payoff in Stage 2 is now only 1. The total deviation payoff is $13 + 1 = 14$.

Therefore, there does not exist a profitable deviation in Stage 1. The subgame perfect Nash equilibrium is as specified in the bullet points above and both firms "Cooperate" in the first stage.

A.6 Chapter 6 Solutions

1. Consider a infinitely repeated prisoners' dilemma with each stage composed of the following simultaneous game:

		Suspect 2 Mum	Fink
Suspect	Mum	15, 15	5, 20
1	Fink	20, 5	10, 10

Both suspects have a discount factor of $\delta \in (0,1)$. What is the lower bound on the discount factor such that (Mum, Mum) is played in all stages, given the "Grim trigger" strategy?

Solution: The payoff (same for both prisoners) along the equilibrium path (from playing "Mum" in all stages) is equal to:

$$\begin{aligned} M &= 15 + 15 \cdot \delta + 15 \cdot \delta^2 - ... \\ M &= 15 + \delta M \\ M(1-\delta) &= 15 \\ M &= \frac{15}{1-\delta}. \end{aligned}$$

The payoff from deviating is given by:

$$\begin{aligned} D &= 20 + 10 \cdot \delta + 10 \cdot \delta^2 - .... \\ D &= 20 + \delta\,(D - 10) \\ D(1-\delta) &= 20 - 10\delta \\ D &= \frac{20 - 10\delta}{1-\delta}. \end{aligned}$$

To rule out a profitable deviation, we must have $M \geq D$:

$$\begin{aligned} \frac{15}{1-\delta} &\geq \frac{20-10\delta}{1-\delta} \\ 15 &\geq 20-10\delta \\ 10\delta &\geq 5 \\ \delta &\geq \frac{5}{10}=\frac{1}{2}. \end{aligned}$$

2. This question analyzes the US soft drink market in two parts.

   (a) Pepsi and Coke are in competition in the soft drink market. The two firms play a game that is repeated over an infinite future. In each stage of the game, each firm simultaneously chooses either to sell at a low price or a high price. The pricing decisions of both firms determine the profits for both firms. The game played in each stage is depicted below.

		Coke	
		Low price	High price
Pepsi	Low price	20, 20	50, 15
	High price	15, 50	30, 30

Both firms have a discount factor δ where $0 < \delta < 1$. Tell me all values of δ so that the "Grim trigger" strategy is a subgame perfect Nash equilibrium. For your convenience, the "Grim trigger" strategy is given by:

For each firm
Stage 1: Play "High price"
:
Stage n: Play "High price" if (High,High) has always been played; Play "Low price" otherwise
:

Solution: Suppose that Coke decides to deviate at some stage and play "Low price." From that point of deviation forward, the discounted future payoffs for Coke are given by:

$$D = 50 + \frac{20\delta}{1-\delta}.$$

By sticking with "High price," the discounted future payoffs for Coke would be:

$$M = \frac{30}{1-\delta}.$$

We need to ensure $M \geq D$:

$$\begin{aligned}
\frac{30}{1-\delta} &\geq 50 + \frac{20\delta}{1-\delta}. \\
30 &\geq 50(1-\delta) + 20\delta. \\
30 &\geq 50 - 50\delta + 20\delta. \\
30\delta &\geq 20. \\
\delta &\geq \frac{2}{3}.
\end{aligned}$$

(b) Now consider that Pepsi and Coke both have a discount factor equal to $\delta = 0.5$. How much is Pepsi willing to pay in order to add a new action called "Super low price" to the set of actions in each stage? With the new action "Super low price" added, the game played in each stage is depicted below.

		Coke		
		Super low price	Low price	High price
	Super low price	5, 5	0, 0	0, 0
Pepsi	Low price	0, 0	20, 20	50, 15
	High price	0, 0	15, 50	30, 30

Solution: Without the action "Super low price" added, the only subgame perfect Nash equilibrium with a discount factor of $\delta = 0.5$ is to play "Low price" in all stages. The total discounted payoff of this strategy is:

$$\frac{20}{1-\delta} = \frac{20}{1-0.5} = 40.$$

If the action "Supper low price" is added, then perhaps there can be a subgame perfect Nash equilibrium in which "High price" is played in all stages. We consider a strong punishment that is also easy to work with (the "Grim trigger" strategy). We use the following "Grim trigger" strategy, which dictates that the off-the-equilibrium path consists of "Super low" being played in all stages following a deviation:

For each firm	
Stage 1:	Play "High price"
:	
Stage n:	Play "High price" if (High,High) has always been played; Play "Super low price" otherwise
:	

Suppose that Coke decides to deviate at some stage and play "Low price." From that point of deviation forward, the discounted future payoffs for Coke are given by:

$$D = 50 + \frac{5\delta}{1-\delta}.$$

By sticking with "High price," the discounted future payoffs for Coke would be:

$$M = \frac{30}{1-\delta}.$$

We need to ensure $M \geq D$:

$$\begin{aligned}
\frac{30}{1-\delta} &\geq 50 + \frac{5\delta}{1-\delta}. \\
30 &\geq 50(1-\delta) + 5\delta. \\
30 &\geq 50 - 50\delta + 5\delta. \\
45\delta &\geq 20. \\
\delta &\geq \frac{4}{9}.
\end{aligned}$$

As $\delta = 0.5 \geq \frac{4}{9}$, then the "Grim trigger" strategy above is a subgame perfect Nash equilibrium, with a total discounted payoff equal to:

$$\frac{30}{1-\delta} = \frac{30}{1-0.5} = 60.$$

Thus, Pepsi (actually, each firm is willing to pay this amount) is willing to pay any amount up to $(60-40) = 20$ in order to add the action "Super low price" to the game.

3. Pepsi and Coke are in competition in the soft drink market. The two firms play a game that is repeated over an infinite future. In each stage of the game, each firm simultaneously chooses either to sell at a low price or a high price. The pricing decisions of both firms determine the profits for both firms. The game played in each stage is depicted below.

		Coke	
		Low price	High price
Pepsi	Low price	20, 20	40, 5
	High price	5, 40	30, 30

Both firms have a discount factor δ where $0 < \delta < 1$. Tell me all values of δ so that the "Tit-for-tat" strategy is a subgame perfect Nash equilibrium. For your convenience, the "Tit-for-tat" strategy is given by:

For each firm
Stage 1: Play "High price"
:
Stage n: Play "High price" if other firm played "High price" in Stage n-1
Play "Low price" if other firm played "Low price" in Stage n-1
:

Solution: We have to ensure that two possible deviations do not occur. The first deviation would take a firm (suppose that Coke is the deviator) from the equilibrium path

(High, High) - (High, High) - (High, High) - ......

to the off the equilibrium path

(High, Low) - (Low, High) - (High, Low) - (Low, High) -......

The total discounted payoff (for Coke) for the equilibrium path is equal to:

$$H = \frac{30}{1-\delta}.$$

The total discounted payoff (for Coke) for the off the equilibrium path is equal to:

$$D_1 = \frac{40 + 5\delta}{1 - \delta^2}.$$

To ensure that the equilibrium path is followed, then $H \geq D_1$:

$$\begin{aligned} \frac{30}{1-\delta} &\geq \frac{40 + 5\delta}{1 - \delta^2} \\ 30\,(1+\delta) &\geq 40 + 5\delta \\ 25\delta &\geq 10 \\ \delta &\geq \frac{2}{5}. \end{aligned}$$

The second deviation would take a firm (again, suppose that Coke is the deviator) from the off the equilibrium path

(Low, High) - (High, Low) - (Low, High) - (High, Low) -......

to the off the off the equilibrium path

(Low, Low) - (Low, Low) - (Low, Low) - (Low, Low) -......

Notice that the off the equilibrium path for this second deviation begins one stage past the initial deviation (we are considering a deviation within a deviation). The total discounted payoff (for Coke) for the off the equilibrium path is equal to:

$$P = \frac{5 + 40\delta}{1 - \delta^2}.$$

The total discounted payoff (for Coke) for the off the off the equilibrium path is equal to:

$$D_2 = \frac{20}{1 - \delta}.$$

To ensure that the off the equilibrium path is followed, then $P \geq D_2$:

$$\begin{aligned}
\frac{5+40\delta}{1-\delta^2} &\geq \frac{20}{1-\delta} \\
5+40\delta &\geq 20(1+\delta) \\
20\delta &\geq 15 \\
\delta &\geq \frac{3}{4}.
\end{aligned}$$

In conclusion, the "Tit-for-tat" strategy is a subgame perfect Nash equilibrium whenever:

$$\delta \geq \max\left\{\frac{2}{5}, \frac{3}{4}\right\} = \frac{3}{4}.$$

Just for comparison, the "Grim trigger" strategy is a subgame perfect Nash equilibrium whenever $\delta \geq \frac{1}{2}$.

4. Consider an infinite horizon bargaining model with two players (Rubinstein bargaining). Player 1 makes offers in the odd-numbered rounds and Player 2 makes offers in the even-numbered rounds. After each player makes an offer, the other player must decide to accept or reject. If the offer is accepted, the game ends. If the offer is rejected, the game proceeds to the next round. Both players have the discount factor $\delta = \frac{2}{3}$. The players are bargaining over \$100.

   What are the equilibrium payoffs (in dollars) for both players?

   Solution: The equilibrium strategies of Rubinstein bargaining (with a common discount factor δ) is for Player 1 to offer to split the \$100 according to the fractions $\left(\frac{1}{1+\delta}, \frac{\delta}{1+\delta}\right)$ in Round 1 and Player 2 to accept. Thus, the equilibrium payoffs are

$$\begin{aligned}
\left(\frac{1}{1+\frac{2}{3}}, \frac{\frac{2}{3}}{1+\frac{2}{3}}\right) \cdot 100 &= \\
\left(\frac{3}{5}, \frac{2}{5}\right) \cdot 100 &= (60, 40).
\end{aligned}$$

5. Consider the Rubinstein bargaining model with three agents. As in class, the bargaining occurs over an infinite time horizon. Agent 1 makes offers in periods 1, 4, 7, ...; Agent 2 makes offers in periods 2, 5, 8, ...; and Agent 3 makes offers in periods 3, 6, 9, ... The agents are bargaining over \$1. If both agents that receive an offer decide to

accept that offer, then the bargaining ends. Let all three agents have the same discount factor $\delta : 0 < \delta < 1$.

What are the equilibrium payoffs for all three agents?

Solution: As all agents have the same discount factor, then the amount in an agent's offer that the agent gets to keep for itself is identical no matter which agent is making the offer (a longer argument justifying this is contained at the end). Let w be the amount in an agent's offer that the agent gets to keep for itself. It is best for all agents to agree in the first period. Let the first period payoffs be given by (w, y_1, z_1), the second period payoffs be (x_2, w, z_2), and the third period payoffs be (x_3, y_3, w). The period 1 offer by Agent 1 is only accepted by Agent 2 if

$$y_1 \geq \delta w.$$

The offer is only accepted by Agent 3 if

$$z_1 \geq \delta z_2.$$

The period 2 offer by Agent 2 is only accepted by Agent 3 if

$$z_2 \geq \delta w.$$

All the optimal offers have binding inequalities:

$$\begin{aligned} y_1 &= \delta w. \\ z_1 &= \delta z_2 = \delta^2 w. \end{aligned}$$

The total payoffs must sum to \$1:

$$\begin{aligned} w + y_1 + z_1 &= 1 \\ w + \delta w + \delta^2 w &= 1 \\ w &= \frac{1}{1+\delta+\delta^2}. \end{aligned}$$

Thus, the equilibrium payoffs are $(w, y_1, z_1) = \left(\frac{1}{1+\delta+\delta^2}, \frac{\delta}{1+\delta+\delta^2}, \frac{\delta^2}{1+\delta+\delta^2}\right)$.

Longer argument:

Let the offers be (x_1, y_1, z_1) such that $x_1 + y_1 + z_1 = 1$ in the first period, (x_2, y_2, z_2) such that $x_2 + y_2 + z_2 = 1$ in the second, and (x_3, y_3, z_3) such that $x_3 + y_3 + z_3 = 1$ in the third. Each offer is made so that it is accepted by the other agents. The first period offer is accepted when:

$$\begin{aligned} y_1 &\geq \delta y_2. \\ z_1 &\geq \delta z_2. \end{aligned}$$

The second period offer is accepted when:

$$\begin{aligned} x_2 &\geq \delta x_3. \\ z_2 &\geq \delta z_3. \end{aligned}$$

The third period offer is accepted when:

$$\begin{aligned} x_3 &\geq \delta x_4. \\ y_3 &\geq \delta y_4. \end{aligned}$$

All the offers are optimal, so the above inequalities are binding. As the game is repeating, then it must be that $(x_1, y_1, z_1) = (x_4, y_4, z_4)$ (same agent making the offer in both periods). Thus, I have a subset of equations:

$$\begin{aligned} x_2 &= \delta^2 x_1. \\ y_1 &= \delta y_2. \\ z_1 &= \delta z_2. \end{aligned}$$

Using the equation $x_1 + y_1 + z_1 = 1$, then we obtain:

$$\frac{x_2}{\delta^2} + \delta\left(y_2 + z_2\right) = 1.$$

Using the equation $y_2 + z_2 = 1 - x_2$, then we obtain:

$$\frac{x_2}{\delta^2} + \delta\left(1 - x_2\right) = 1.$$

After collecting terms:

$$x_2\left(\frac{1}{\delta^2}-\delta\right)=1-\delta,$$

and simplifying:

$$x_2\left(\frac{1-\delta^3}{\delta^2}\right)=1-\delta.$$

Solving for x_2 :

$$x_2=\frac{(1-\delta)\,\delta^2}{\left(1-\delta^3\right)}.$$

We can factor $\left(1-\delta^3\right)=(1-\delta)\left(1+\delta+\delta^2\right)$. Thus,

$$x_2=\frac{\delta^2}{\left(1+\delta+\delta^2\right)}.$$

Using our above equations for (x_1,x_2,x_3), we obtain:

$$x_1=\tfrac{1}{\left(1+\delta+\delta^2\right)}\qquad x_3=\tfrac{\delta}{\left(1+\delta+\delta^2\right)}\ .$$

We can proceed in similar fashion and use the equation $x_2+y_2+z_2=1$ to solve for (y_1,y_2,y_3) and the equation $x_3+y_3+z_3=1$ to solve for (z_1,z_2,z_3). We arrive at the exact same answer as above.

A.7 Chapter 7 Solutions

1. Consider the auction for a single good between three bidders. For simplicity, each bidder has value $v_1=v_2=v_3=v$ for the object. Consider a third price sealed-bid auction. That is, the winner of the auction is the one with the highest bid (ties broken randomly) and the winner must pay the third highest bid for the object (with only three bidders, the third highest bid is also the lowest submitted bid).

   Is it a dominant strategy for a bidder to bid its value v? Why or why not?

   Solution: Let the bids for all three bidders be given by b_1, b_2, and b_3. A dominant strategy means that for all strategies by the other bidders, no other strategy does strictly better. Consider whether $b_1=v$ is a dominant strategy. Let $b_2<v$ and $b_3>v$. The strategy $b_1=v$ is worse than $b_1^*>b_3$. With $b_1^*>b_3$, bidder 1 wins the auction and has payoff $v-p=v-b_2>0$. The strategy $b_1=v$ results in a payoff of 0 as bidder 1 loses the auction.

2. Suppose you are one of two bidders in a first price sealed bid auction for Super Bowl tickets. Each bidder (including you) knows its own value for the Super Bowl tickets, but not the values of the other bidders. All a bidder knows about the values of the other bidders is that they are uniformly distributed between \$200 and \$300. Your value for the tickets is \$280. What is the optimal bid that you should place?

   Solution: Suppose that you are bidder 1. The optimal bid is such that

$$\begin{aligned} b_1 &= \frac{1}{2}v_1 + \frac{1}{2}(\$200) \\ &= \frac{1}{2}(\$280) + \$100 \\ &= \$240. \end{aligned}$$

3. Suppose you are one of two bidders in a first price sealed bid auction. The values that each of the bidders (including you) place on the object are uniformly distributed between \$50 and \$150. You know your own value (which is \$130), but not the value of the other bidder. In this game, the seller has already decided to impose a reservation price of \$80. A reservation price means that your bid can only win if it is both (i) greather than or equal to the reservation price and (ii) greater than or equal to the bid of the other bidder.

   Given that your value for the object is \$130, what is the optimal bid that you should place?

   Solution: The optimal bid if $v_2 < \$80$ is given by:

$$b_1 = \$80.$$

   The optimal bid if $v_2 \geq \$80$ is

$$b_1 = \frac{1}{2}(\$80) + \frac{1}{2}(\$130) = \$105.$$

   Since the probability $\Pr(v_2 < \$80) = 0.3$ and $\Pr(v_2 \geq \$80) = 0.7$, then the optimal bid is the expected bid:

$$b_1 = (0.3)(\$80) + (0.7)(\$105) = \$97.5.$$

4. Consider a first price sealed-bid auction between n different bidders, where $n = 2, 3, 4, 5, ...$ The bidders' valuations, v_i, are uniformly distributed between 0 and 1. Each bidder knows its true value, but does not know the values of the other bidders. The equilibrium bidding strategy for all bidders is given by $b_i = av_i$, where a is some unknown constant.

Find the equilibrium bidding strategy for all the bidders (that is, determine the value for the constant a).

Solution: Consider the problem from the viewpoint of bidder 1. All other $n-1$ bidders bid according to the strategies $b_i = av_i$ for $i \geq 2$. The payoff for bidder 1 is given by:

$$\left\{ \begin{array}{cc} v_1 - b_1 & \text{if } b_1 > a \cdot \max_{i \geq 2} v_i \\ \frac{1}{2}(v_1 - b_1) & \text{if } b_1 = a \cdot \max_{i \geq 2} v_i \\ 0 & \text{if } b_1 < a \cdot \max_{i \geq 2} v_i \end{array} \right\}.$$

The payoff for bidder 1 is equal to:

$$\begin{aligned} & (v_1 - b_1) \cdot \Pr(b_1 > av_2) \cdot ... \cdot \Pr(b_1 > av_n) \\ = \; & (v_1 - b_1) \cdot \Pr\left(v_2 < \frac{b_1}{a}\right) \cdot ... \cdot \Pr\left(v_n < \frac{b_1}{a}\right) \\ = \; & (v_1 - b_1) \cdot \left(\frac{b_1}{a}\right) \cdot ... \cdot \left(\frac{b_1}{a}\right) \\ = \; & (v_1 - b_1) \cdot \left(\frac{b_1}{a}\right)^{n-1}. \end{aligned}$$

Taking the first order conditions with respect to b_1 allows us to solve for the optimal bidding strategy:

$$\begin{aligned} \frac{\partial}{\partial b_1} &= (n-1)(v_1 - b_1) \cdot \left(\frac{b_1}{a}\right)^{n-2} \cdot \left(\frac{1}{a}\right) - \left(\frac{b_1}{a}\right)^{n-1} = 0 \\ \frac{\partial}{\partial b_1} &= \frac{(n-1)v_1}{a}\left(\frac{b_1}{a}\right)^{n-2} - (n-1)\left(\frac{b_1}{a}\right)^{n-1} - \left(\frac{b_1}{a}\right)^{n-1} = 0 \\ \frac{\partial}{\partial b_1} &= \frac{(n-1)v_1}{a}\left(\frac{b_1}{a}\right)^{n-2} - n\left(\frac{b_1}{a}\right)^{n-1} = 0. \end{aligned}$$

Solving for b_1 yields the optimal bidding strategy:

$$\begin{aligned} \frac{(n-1)\,v_1}{a}\left(\frac{b_1}{a}\right)^{n-2} &= n\left(\frac{b_1}{a}\right)^{n-1} \\ \frac{(n-1)\,v_1}{a} &= n\frac{b_1}{a} \\ b_1 &= \frac{n-1}{n}v_1. \end{aligned}$$

Recall that the optimal strategies are of the form $b_i = av_i$ for $i \geq 2$. Since the game is symmetric, $b_1 = av_1$. Thus, $a = \frac{n-1}{n}$ and all bidders have bidding strategies:

$$b_i = \frac{n-1}{n}v_i \text{ for } i = 1, 2, ..., n.$$

Notice that this fits with our class analysis when $n = 2$.

5. Consider Exercise 7.2 from above. Suppose the auction now has three bidders and your own value is still \$280. What will you decide to bid for the Super Bowl tickets?

   Solution: The optimal bidding strategy is given by:

   $$b_1 = \frac{2}{3}v_1 + \frac{1}{3}(\$200).$$

   Given that your value is equal to \$280, then your optimal bid should be $b_1 = \frac{2}{3}(\$280) + \frac{1}{3}(\$200) = \$253.33$.

6. Consider Exercise 7.3 from above. Suppose the auction now has 3 bidders and your own value is still \$130. What will you decide to bid for the object?

   Solution: With three bidders, you will find it optimal to bid more than what you bid with only two bidders.

   The optimal bid if both $v_2 < \$80$ and $v_3 < \$80$ is given by:

   $$b_1 = \$80.$$

   The optimal bid if either $v_2 \geq \$80$ or $v_3 \geq \$80$ (but not both) is given by:

   $$b_1 = \frac{1}{2}(\$80) + \frac{1}{2}(\$130) = \$105.$$

The optimal bid if both $v_2 \geq \$80$ and $v_3 \geq \$80$ is given by:

$$b_1 = \frac{1}{3}(\$80) + \frac{2}{3}(\$130) = \$113.33.$$

The probabilities are given by:

$$\begin{aligned}
\Pr(v_2 < \$80)\Pr(v_3 < \$80) &= (0.3)^2 \\
\Pr(v_2 < \$80)\Pr(v_3 \geq \$80) + \Pr(v_2 \geq \$80)\Pr(v_3 < \$80) &= 2(0.7)(0.3) \\
\Pr(v_2 \geq \$80)\Pr(v_3 \geq \$80) &= (0.7)^2
\end{aligned}$$

The optimal bid is given by:

$$\begin{aligned}
(0.3)^2(\$80) + 2(0.7)(0.3)(\$105) + (0.7)^2(\$113.33) &= \\
(0.09)(\$80) + (0.42)(\$105) + (0.49)(\$113.33) &= \$106.83.
\end{aligned}$$

A.8 Chapter 8 Solutions

1. Show that the unique PBNE of Figure 8.2 is (D2,(p=0,L)).

   Solution: The PBNE is found by considering all possible strategies for Player 1.

   (a) Player 1 plays D1

   If Player 1 plays D1, then the information set is reached with strictly positive probability. So, if Player 1 plays D1, then Player 2 must believe that Player 1 has played D1 and all the probability needs to be massed on the left: $p = 1$.

   Given the belief $p = 1$, Player 1 takes the action from the left node in the information set. As the payoff from L (which equals 1) exceeds the payoff from R (which equals 0), Player 2 will choose L. Thus, the strategy and belief of Player 2 is given by $(p = 1, L)$.

   As we are dealing with an information set, Player 2 can only take one action and this action is L. This is held fixed and we want to see if Player 1 is really best off choosing D1. The payoff from D1 (which equals 2) is strictly less than the payoff from D2 (which equals 3, as Player 2 plays L).

   As Player 1 finds it optimal to deviate, then we cannot have a PBNE in which Player 1 plays D1.

(b) Player 1 plays D2

The information set is reached with strictly positive probability, so the beliefs must be consistent. If Player 1 plays D2, then the probability needs to be massed on the right: $1 - p = 1$ or $p = 0$.

Given the belief $p = 0$, Player 2 takes the action from the right node in the information set. As the payoff from L (which equals 2) exceeds the payoffs from R (which equals 1). Player 2 will choose L. Thus, the strategy and belief of Player 2 is given by $(p = 0, L)$.

As we are dealing with an information set, Player 2 can only take one action and this action is L. This is held fixed and we want to see if Player 1 is really best off choosing D2. The payoff from D2 (which equals 3) is greater than the payoff from D1 (which equals 1, as Player 2 plays L) and greater than the payoff from T (which equals 2).

Thus, one PBNE is given by $(D2, (p = 0, L))$.

(c) Player 1 plays T

The information set is not reached with strictly positive probability, meaning that the beliefs are not determined by the actions of Player 1. Player 2 is free to form whatever beliefs it chooses. This means that we, as the game theorists, will find the values $(p, 1 - p)$ that are required to obtain a PBNE.

Consider that if Player 1 plays T, then its payoff is equal to 2. Can Player 1 do better by deviating to either D1 or D2? The answer is that it depends upon what Player 2 does. The only deviation that leads to a strictly higher payoff occurs when Player 1 plays D2 and Player 2 plays L. Thus, provided that Player 2 beliefs are specified such that Player 2 plays R, then there is no profitable deviation for Player 1. Without a profitable deviation, then playing T can be part of a PBNE.

The expected payoffs to Player 2 are given by:

$$
\begin{aligned}
E_p(L) &= 1p + 2(1 - p) \\
E_p(R) &= 0p + 1(1 - p)
\end{aligned}
$$

The probabilities required for Player 2 to choose R are:

$$\begin{aligned} E_p(R) &\geq E_p(L) \\ (1-p) &\geq p + 2(1-p) \\ -p &\geq 1-p. \end{aligned}$$

Thus, there does not exist any values for p such that Player 2 chooses to play R. The strategy L is a dominant strategy: strictly higher payoffs for both possibilities. For this reason, there does not exist a PBNE in which Player 1 plays T.

2. Find the perfect Bayesian Nash equilibrium of the game in Figure 8.4, which is referred to as Selten's horse (Osborne, 2009, pg. 331).

   Solution: In the information set for Player 3, let p be the probability that Player 1 selects D and $1-p$ be the probability that Player 1 selects C and Player 2 selects d. Let's consider the possible strategies of Players 1 and 2.

   Suppose Player 1 selects D. Then $p = 1$. For Player 3, $E_p(L) = 2 \cdot p + 0 \cdot (1-p)$ and $E_p(R) = 0 \cdot p + 1 \cdot (1-p)$. Thus, Player 3 selects L. The payoffs are $(3, 3, 2)$. Given that Player 3 plays L, then it is optimal for Player 2 to play d (payoff of 4 exceeds payoff of 1). So, is it optimal for Player 1 to play D initially? No, by deviating to C, Player 1 can receive a payoff of 4 (as Player 2 plays d and Player 3 plays L). No equilibrium.

   Suppose Player 1 selects C. Then Player 2 can select d or c.

   (a) If Player 2 plays d, then $p = 0$. For Player 3, $E_p(L) = 2 \cdot p + 0 \cdot (1-p)$ and $E_p(R) = 0 \cdot p + 1 \cdot (1-p)$. Thus, Player 1 selects R. Player 2 has a profitable deviation by selecting c. No equilibrium.

   If Player 2 plays c, then $p : 0 \leq p \leq 1$. The payoffs are $(1, 1, 1)$. There is no profitable deviation for either Player 1 or Player 2 when Player 3 selects R. Player 3 selects R if $E_p(R) \geq E_p(L)$. This requires:

$$\begin{aligned} 0 \cdot p + 1 \cdot (1-p) &\geq 2 \cdot p + 0 \cdot (1-p) \\ 1-p &\geq 2p \\ p &\leq \frac{1}{3}. \end{aligned}$$

   Thus, the one perfect Bayesian Nash equilibrium is $\left(\text{C, c, } \left(p \leq \frac{1}{3}, R\right)\right)$.

3. A father is purchasing a Christmas present for his son. The father can buy one of three gifts: {Lego set, Nerf gun, video game}. The father selects one of the presents and then wraps it the day before Christmas and places it under the tree. The Lego set and the Nerf gun are the same size and shape, while the video game is noticeably smaller. The son observes the size and shape of the wrapped present the day before Christmas and forms beliefs about what the present might be. The son is told that any bad behavior at the Christmas Eve mass will result in the loss of the Christmas present, but the son finds it very costly to avoid bad behavior during the long Christmas Eve mass.

   The dynamic game is depicted in Figure 8.5. Find all perfect Bayesian Nash equilibria (PBNE) of this game.

   Solution: First, define p as the probability that the son believes that the gift is a Lego set given that it is the size and shape of either the Lego set or the Nerf gun. We know by backward induction that if the gift is the size and shape of the video game, then the son will choose 'Not behave' in mass, so the payoffs on the right-hand side are (3, 2). Then we consider each of the strategies by the father in order:

   <u>Father chooses Lego set</u>

   In this case, $p = 1$ and the son chooses 'Behave'. This provides a payoff of 4 to the father. The father cannot deviate and receive a strictly higher payoff. Thus, one PBNE is given by (Lego set, $p = 1$, Behave).

   <u>Father chooses Nerf gun</u>

   In this case, $p = 0$ and the son chooses 'Not behave'. This provides a payoff of 2 to the father. The father can deviate to Video game and receive a strictly higher payoff. Thus, this cannot be a PBNE.

   <u>Father chooses Video game</u>

   In this case, p can be anything. The payoff to the father is 3. There is no incentive for the father to deviate if the beliefs p are such that the son chooses 'Not behave' at the information set. The son chooses 'Not behave' when:

$$
\begin{aligned}
E_p\,(Not) &\geq E_p\,(Behave) \\
p \cdot 2 + (1-p) \cdot 2 &\geq p \cdot 3 + (1-p) \cdot 0 \\
2 &\geq 3p \\
p &\leq \frac{2}{3}.
\end{aligned}
$$

Thus, the second PBNE is given by (Video game, $p \leq \frac{2}{3}$, Not behave).

A.9 Chapter 9 Solutions

1. Solve for the perfect Bayesian Nash equilibria of the Beer-Quiche game depicted in Figure 9.5. The game is identical to Figure 9.2 with two exceptions (which are indicated in **bold** in the figure).

   Solution: Let's analyze each of the four possible strategies for the Student.

   Tough Guy orders Beer; Wimp orders Quiche

   Define $p = \Pr(Tough|Beer)$. Then $p = 1$. For the Bully, $E_p(Fight) = 1 \cdot p + 0 \cdot (1-p)$ and $E_p(Not) = 0 \cdot p + 2 \cdot (1-p)$. Thus Bully "Fight" when student orders Beer.

   Define $q = \Pr(Tough|Quiche)$. Then $q = 0$. For the Bully, $E_q(Fight) = 1 \cdot q + 0 \cdot (1-q)$ and $E_q(Not) = 0 \cdot q + 2 \cdot (1-q)$. Thus Bully "Not" when student orders Quiche.

   Payoff to Tough Guy is 0. Does there exist a profitable deviation? If select Quiche rather than Beer, Bully's action fixed at "Not," so Tough Guy receives payoff of 2.

   No equilibrium.

   Tough Guy orders Quiche; Wimp orders Beer

   Then $p = 0$. For the Bully, $E_p(Fight) = 1 \cdot p + 0 \cdot (1-p)$ and $E_p(Not) = 0 \cdot p + 2 \cdot (1-p)$. Thus Bully "Not" when student orders Beer.

   Then $q = 1$. For the Bully, $E_q(Fight) = 1 \cdot q + 0 \cdot (1-q)$ and $E_q(Not) = 0 \cdot q + 2 \cdot (1-q)$. Thus Bully "Fight" when student orders Quiche.

   Payoff to Tough Guy is 1. Does there exist a profitable deviation? If select "Beer" rather than "Quiche," Bully's action fixed at "Not," so Tough Guy receives payoff of 3.

   No equilibrium.

   Tough Guy orders Beer; Wimp orders Beer

   Then $p = \frac{1}{2}$. For the Bully, $E_p(Fight) = 1 \cdot p + 0 \cdot (1-p)$ and $E_p(Not) = 0 \cdot p + 2 \cdot (1-p)$. Thus Bully "Not" when student orders Beer.

   Then q can take any value $0 \leq q \leq 1$.

Payoff for Tough Guy is 3. Payoff for Wimp is 2. Neither has an incentive to deviate if Bully "Fight" after observing Quiche.

For the Bully, $E_q(Fight) = 1 \cdot q + 0 \cdot (1-q)$ and $E_q(Not) = 0 \cdot q + 2 \cdot (1-q)$. Thus,

$$\begin{aligned} E_q(Fight) &\geq E_q(Not) \\ 1 \cdot q + 0 \cdot (1-q) &\geq 0 \cdot q + 2 \cdot (1-q) \\ q &\geq 2(1-q) \\ q &\geq \frac{2}{3}. \end{aligned}$$

Thus, a perfect Bayesian Nash equilibrium is $\left((Beer, Beer), \left(p = \frac{1}{2}, q \geq \frac{2}{3}, Not, Fight\right)\right)$.

<u>Tough Guy orders Quiche; Wimp orders Quiche</u>

Then $q = \frac{1}{2}$. For the Bully, $E_q(Fight) = 1 \cdot q + 0 \cdot (1-q)$ and $E_q(Not) = 0 \cdot q + 2 \cdot (1-q)$. Thus Bully "Not" when student orders Quiche.

Then p can take any value $0 \leq p \leq 1$.

Payoff for Tough Guy is 2. Payoff for Wimp is 3. Neither has an incentive to deviate if Bully "Fight" when observes Beer.

For the Bully, $E_p(Fight) = 1 \cdot p + 0 \cdot (1-p)$ and $E_p(Not) = 0 \cdot p + 2 \cdot (1-p)$. Thus,

$$\begin{aligned} E_p(Fight) &\geq E_p(Not) \\ 1 \cdot p + 0 \cdot (1-p) &\geq 0 \cdot p + 2 \cdot (1-p) \\ p &\geq 2(1-p) \\ p &\geq \frac{2}{3}. \end{aligned}$$

Thus, a perfect Bayesian Nash equilibrium is $\left((Quiche, Quiche), \left(p \geq \frac{2}{3}, q = \frac{1}{2}, Fight, Not\right)\right)$.

In summary, there exist two pooling perfect Bayesian Nash equilibria.

2. Solve for the perfect Bayesian Nash equilibria of the Beer-Quiche game depicted in Figure 9.6. The game is identical to Figure 9.2 with one exception (which is indicated in **bold** in the figure).

   Solution: Let's analyze each of the four possible strategies for the Student.

   <u>Tough Guy orders Beer; Wimp orders Quiche</u>

Define $p = \Pr(Tough|Beer)$. Then $p = 1$. For the Bully, $E_p(Fight) = 1 \cdot p + 0 \cdot (1-p)$ and $E_p(Not) = 0 \cdot p + 2 \cdot (1-p)$. Thus Bully "Fight" when student orders Beer.

Define $q = \Pr(Tough|Quiche)$. Then $q = 0$. For the Bully, $E_q(Fight) = 2 \cdot q + 0 \cdot (1-q)$ and $E_q(Not) = 0 \cdot q + 1 \cdot (1-q)$. Thus Bully "Not" when student orders Quiche.

Payoff to Tough Guy is 2.5. Does there exist a profitable deviation? No, because if select "Quiche" rather than "Beer," Bully's action fixed at "Not," so Tough Guy receives payoff of 2.

Payoff to Wimp is 3. Does there exist a profitable deviation? No, because if select "Beer" rather than "Quiche," Bully's action fixed at "Fight," so Wimp receives a payoff of 1.

Thus, a perfect Bayesian Nash equilibrium is $((Beer, Quiche), (p = 1, q = 0, Fight, Not))$.

Tough Guy orders Quiche; Wimp orders Beer

Then $p = 0$. For the Bully, $E_p(Fight) = 1 \cdot p + 0 \cdot (1-p)$ and $E_p(Not) = 0 \cdot p + 2 \cdot (1-p)$. Thus Bully "Not" when student orders Beer.

Then $q = 1$. For the Bully, $E_q(Fight) = 2 \cdot q + 0 \cdot (1-q)$ and $E_q(Not) = 0 \cdot q + 1 \cdot (1-q)$. Thus Bully "Fight" when student orders Quiche.

Payoff to Tough Guy is 1. Does there exist a profitable deviation? If select "Beer" rather than "Quiche," Bully's action fixed at "Not," so Tough Guy receives payoff of 3.

No equilibrium.

Tough Guy orders Beer; Wimp orders Beer

Then $p = \frac{1}{2}$. For the Bully, $E_p(Fight) = 1 \cdot p + 0 \cdot (1-p)$ and $E_p(Not) = 0 \cdot p + 2 \cdot (1-p)$. Thus Bully "Not" when student orders Beer.

Then q can take any value $0 \leq q \leq 1$.

Payoff for Tough Guy is 3. Payoff for Wimp is 2. Neither has an incentive to deviate if Bully "Fight" when observes Quiche.

For the Bully, $E_q(Fight) = 2 \cdot q + 0 \cdot (1-q)$ and $E_q(Not) = 0 \cdot q + 1 \cdot (1-q)$. Thus,

$$\begin{aligned} E_q(Fight) &\geq E_q(Not) \\ 2 \cdot q + 0 \cdot (1-q) &\geq 0 \cdot q + 1 \cdot (1-q) \\ 2q &\geq (1-q) \\ q &\geq \frac{1}{3}. \end{aligned}$$

Thus, a perfect Bayesian Nash equilibrium is $\left((Beer, Beer), \left(p = \frac{1}{2}, q \geq \frac{1}{3}, Not, Fight\right)\right)$.

Tough Guy orders Quiche; Wimp orders Quiche

Then $q = \frac{1}{2}$. For the Bully, $E_q(Fight) = 2 \cdot q + 0 \cdot (1-q)$ and $E_q(Not) = 0 \cdot q + 1 \cdot (1-q)$. Thus Bully "Fight" when student orders Quiche.

Then p can take any value $0 \leq p \leq 1$.

Payoff for Tough Guy is 1. Payoff for Wimp is 0. Both have an incentive to deviate. Regardless of what the Bully selects after observing Beer, both the Tough Guy and the Wimp receive a higher payoff by selecting "Beer."

No equilibrium.

In summary, there exists one separating perfect Bayesian Nash equilibrium and one pooling perfect Bayesian Nash equilibrium.

3. Consider a game of asymmetric information played between a student and a firm. The student can have one of two ability levels: high or low. The student knows its own ability level, while the firm does not. All the firm knows is that 75% of students are of high ability and 25% of students are of low ability.

   The student sends a signal by deciding whether or not to attend college. The cost of college for a high ability student is 4, while the cost of college for a low ability student is 8. There is no cost associated with the decision not to attend college.

   The firm forms beliefs about a student's ability level and then must decide to either offer the student high wages or low wages. High wages are value 10, while low wages are value 1.

The payoffs to the firm are as follows:

$$\text{Payoff} = 1 \text{ if offer} \quad \begin{array}{l} \text{High wages to High ability OR} \\ \text{Low wages to Low ability} \end{array}.$$

$$\text{Payoff} = 0 \text{ if offer} \quad \begin{array}{l} \text{High wages to Low ability OR} \\ \text{Low wages to High ability} \end{array}.$$

The payoff for the student is simply Payoff = Wages - Cost of Education.

Draw the game tree. Does there exist a perfect Bayesian Nash equilibrium (PBNE) in which both ability levels choose to attend college? Show all work to justify your answer.

Solution: The game tree is depicted in Figure A.2 of the companion 'Figures' document.

The question asks us to verify if a pooling PBNE exists in which both ability levels choose to attend college. Define p as the belief by the firm that a student attending college is of high ability. Define q as the belief by the firm that a student not attending college is of high ability. If both ability levels attend college, then $p = 0.75$. Given this belief, the firm chooses to offer high wages to any student that attends college. Both types of students are happy to stay at their current choice of college so long as the beliefs q are such that the firm chooses to offer low wages to any student that does not attend college. For the values $q \leq 0.5$, the firm will choose to offer low wages to any student that does not attend college. Thus, we have discovered the following PBNE:

$$(College, College), (p = 0.75, q \leq 0.5, High, Low).$$

4. Consider a game exactly as in the previous question, except now that 50% of students are of High ability and 50% of students are of Low ability. An additional difference is that now the cost of college for a high ability student is 5, while the cost of college fo ra low ability student is 10. All other facts are identical to those described in the previous question.

   Draw the game tree. Does there exist a perfect Bayesian Nash equilibrium (PBNE) in which only students of High ability go to college, while students of Low ability do not? Show all work to justify your answer.

Solution: The game tree is depicted in Figure A.3 of the companion 'Figures' document.

The question asks us to verify if a separating PBNE exists in which only students of high ability attend college. Define p as the belief by the firm that a student attending college is of high ability. Define q as the belief by the firm that a student not attending college is of high ability. If only high ability students attend college, then $p = 1$. Given this belief, the firm chooses to offer high wages to any student that attends college. If only low ability students do not attend college, then $q = 0$. Given this belief, the firm chooses to offer low wages to any student that does not attend college. If the high ability student switched (from college to no college), then their payoff would drop from 5 to 1. If the low ability student switched (from no college to college), then their payoff would drop from 1 to 0. Thus, we have discovered the following PBNE:

$$(College, No), (p = 1, q = 0, High, Low).$$

5. Consider an insurance market with two risk-neutral insurance companies and risk-averse individuals. The individuals are of two types: High (H) and Low (L). Both types have wealth $W = 6$ and in the bad states (b) of the world suffer a loss of $R = 8$ (in the good states (g), there is no loss). The high type suffers a loss with probability $p_H = \frac{1}{2}$ and the low type suffers a loss with probability $p_L = \frac{2}{5}$. The individuals have utility $u(\cdot) = \ln(\cdot)$, so the payoff functions are:

$$\begin{aligned} \text{Payoff } (H) &= p_H \cdot u(c_{H,b}) + (1 - p_H) \cdot u(c_{H,g}) \\ \text{Payoff } (L) &= p_L \cdot u(c_{L,b}) + (1 - p_L) \cdot u(c_{L,g}) \end{aligned}$$

The insurance companies offer two contracts: (M_H, D_H) and (M_L, D_L), where M is the premium and D is the deductible. With asymmetric information, solve for the optimal contracts that are offered by the insurance companies.

Solution: The contract (M_H, D_H) will be selected by the high risk types and (M_L, D_L) will be selected by the low risk types. Both contracts are determined so that the insurance companies earn zero profit:

$$\begin{aligned} \pi_H &= M_H - p_H (R - D_H) = 0 \\ \pi_L &= M_L - p_L (R - D_L) = 0 \end{aligned}$$

This implies $M_H = p_H(8 - D_H)$ and $M_L = p_L(8 - D_L)$. The contract offered to the High type individuals under complete information remains unchanged in this setting with incomplete information. The reason is that the low risk types have no interest in purchasing this contract. We just need to make sure that contract (M_H, D_H) provides the optimal payoff for the high risk types:

$$\max_{D_H} p_H \cdot u(W - [p_H(8 - D_H)] - D_H) + (1 - p_H) \cdot u(W - [p_H(8 - D_H)]).$$

Given the risk aversion ($u(\cdot) = \ln(\cdot)$ is strictly concave), then

$$\begin{aligned} W - [p_H(8 - D_H)] - D_H &= W - [p_H(8 - D_H)] \\ D_H &= 0. \end{aligned}$$

The contract (M_L, D_L) cannot provide the optimal payoff for the low risk type, because the incentive compatibility condition must be satisfied for the high risk type (the high risk type selects the contract (M_H, R_H) over (M_L, R_L)). Thus, we have

$$\frac{1}{2}\ln(2) + \frac{1}{2}\ln(2) \geq \frac{1}{2}\ln(W - [p_L(8 - D_L)] - D_L) + \frac{1}{2}\ln(W - [p_L(8 - D_L)]).$$

Cancel $\frac{1}{2}$ from all terms and use the natural log property $\ln(a) + \ln(b) = \ln(ab)$:

$$\ln(4) \geq \ln\left\{\left(\frac{14}{5} - \frac{3}{5}D_L\right)\left(\frac{14}{5} + \frac{2}{5}D_L\right)\right\}.$$

As $\ln(\cdot)$ is always increasing,

$$\begin{aligned} 4 &\geq \left(\frac{14}{5} - \frac{3}{5}D_L\right)\left(\frac{14}{5} + \frac{2}{5}D_L\right) \\ 100 &\geq 196 - 14D_L - 6D_L^2. \end{aligned}$$

The algebra yields:

$$\begin{aligned} 6D_L^2 + 14D_L - 96 &\geq 0 \\ 3D_L^2 + 7D_L - 48 &\geq 0 \\ (3D_L + 16)(D_L - 3) &\geq 0. \end{aligned}$$

This implies that $D_L \geq 3$ and the deductible that maximizes the payoff for the driver is given by $D_L = 3$.

Thus, the contracts are: $(M_H, D_H) = (4, 0)$ and $(M_L, D_L) = (2, 3)$.

6. Consider an identical problem as above, except now both types of individuals have wealth $W = \frac{5}{2}$ and in the bad states (b) of the world suffer a loss of $R = 3$,where the high type suffers a loss with probability $p_H = \frac{1}{2}$ and the low type suffers a loss with probability $p_L = \frac{1}{3}$.

   With asymmetric information, solve for the optimal contracts that are offered by the insurance companies.

   Solution: The contract (M_H, D_H) will be selected by the high risk types and (M_L, D_L) will be selected by the low risk types. Both contracts are determined so that the insurance companies earn zero profit:

$$\begin{aligned} \pi_H &= M_H - p_H (R - D_H) = 0 \\ \pi_L &= M_L - p_L (R - D_L) = 0 \end{aligned}$$

   This implies $M_H = p_H (3 - D_H)$ and $M_L = p_L (3 - D_L)$. The contract offered to the High type individuals under complete information remains unchanged in this setting with incomplete information. The reason is that the low risk types have no interest in purchasing this contract. We just need to make sure that contract (M_H, D_H) provides the optimal payoff for the high risk types:

$$\max_{D_H} p_H \cdot u\left(\frac{5}{2} - [p_H (3 - D_H)] - D_H\right) + (1 - p_H) \cdot u\left(\frac{5}{2} - [p_H (3 - D_H)]\right).$$

   Given the risk aversion ($u(\cdot) = \ln(\cdot)$ is strictly concave), then

$$\begin{aligned} \frac{5}{2} - [p_H (3 - D_H)] - D_H &= \frac{5}{2} - [p_H (3 - D_H)] \\ D_H &= 0. \end{aligned}$$

   The contract (M_L, D_L) cannot provide the optimal payoff for the low risk type, because the incentive compatibility condition must be satisfied for the high risk type (the high risk type selects the contract (M_H, R_H) over (M_L, R_L)). Thus, we have

$$\frac{1}{2}\ln(1)+\frac{1}{2}\ln(1)\geq\frac{1}{2}\ln\left(\frac{5}{2}-\left[p_L\left(3-D_L\right)\right]-D_L\right)+\frac{1}{2}\ln\left(\frac{5}{2}-\left[p_L\left(3-D_L\right)\right]\right).$$

Cancel $\frac{1}{2}$ from all terms and use the natural log property $\ln(a)+\ln(b)=\ln(ab)$:

$$\ln(1)\geq\ln\left\{\left(\frac{3}{2}-\frac{2}{3}D_L\right)\left(\frac{3}{2}+\frac{1}{3}D_L\right)\right\}.$$

As $\ln(\cdot)$ is always increasing,

$$\begin{aligned} 1 &\geq \left(\frac{9}{6}-\frac{4}{6}D_L\right)\left(\frac{9}{6}+\frac{2}{6}D_L\right) \\ 36 &\geq 81-18D_L-8D_L^2. \end{aligned}$$

The algebra yields:

$$\begin{aligned} 8D_L^2+18D_L-45 &\geq 0 \\ (4D_L+15)(2D_L-3) &\geq 0. \end{aligned}$$

This implies that $D_L \geq \frac{3}{2}$ and the deductible that maximizes the payoff for the driver is given by $D_L = \frac{3}{2}$.

Thus, the contracts are: $(M_H, D_H) = \left(\frac{3}{2}, 0\right)$ and $(M_L, D_L) = \left(\frac{1}{2}, \frac{3}{2}\right)$.

A.10 Chapter 10 Solutions

1. Consider a game of moral hazard between a risk-neutral owner and a risk-averse manager. A project is available with random returns: either of value 200 or of value 100. The owner can offer a wage pair $(w(200), w(100))$, where $w(200)$ is the wage when the project has return 200 (success) and $w(100)$ is the wage when the project has return 100 (fail).

   The payoff function for the owner is then given by:

$$p\cdot(200-w(200))+(1-p)(100-w(100)).$$

The payoff function for the managers is given by:

$$p \cdot \sqrt{w(200)} + (1-p)\sqrt{w(100)} - \text{cost of effort}.$$

The probability of success p is determined based upon the effort decision of the manager:

$$\begin{array}{ll} p = p_H = 0.75 & \text{if manager chooses High effort} \\ p = p_L = 0.25 & \text{if manager chooses Low effort} \end{array}.$$

The cost of effort for the manager is 0 for Low effort and $c = 4$ for High effort. Additionally, the manager has an outside option with payoff $\bar{u} = 4$.

In this game with asymmetric information in which the owner cannot observe the effort decision of the manager, what wage pair $(w(200), w(100))$ provides the incentives for the manager to choose High effort? What are the payoffs for both the manager and the owner?

Solution: In order to induce High effort, the owner must solve the following constrained maximization problem:

$$\begin{array}{ll} \max & (0.75) \cdot (200 - w(200)) + (0.25)(100 - w(100)) \\ \text{subject to} & (0.75) \cdot \sqrt{w(200)} + (0.25)\sqrt{w(100)} - 4 \geq 4 \\ & (0.75) \cdot \sqrt{w(200)} + (0.25)\sqrt{w(100)} - 4 \geq (0.25) \cdot \sqrt{w(200)} + (0.75)\sqrt{w(100)} \end{array}.$$

We know from class that both of the constraints bind. Thus, we can solve for two equations in two unknowns:

$$\begin{aligned} (0.75) \cdot \sqrt{w(200)} + (0.25)\sqrt{w(100)} &= 8. \\ (0.5) \cdot \sqrt{w(200)} - (0.5)\sqrt{w(100)} &= 4. \end{aligned}$$

Multiply the first equation by 4 and the second equation by 2:

$$\begin{aligned} 3\sqrt{w(200)} + \sqrt{w(100)} &= 32. \\ \sqrt{w(200)} - \sqrt{w(100)} &= 8. \end{aligned}$$

Adding the equations together yields:

$$4\sqrt{w(200)} = 40,$$

so $\sqrt{w(200)} = 10$. This leaves $\sqrt{w(100)} = 2$.

If you had remembered the formulas from class, then you can simply insert the values for the parameters:

$$\begin{aligned}\sqrt{w(200)} &= \bar{u} + c\left(\frac{1-p_L}{p_H - p_L}\right) = 10. \\ \sqrt{w(100)} &= \bar{u} + c\left(\frac{-p_L}{p_H - p_L}\right) = 2.\end{aligned}$$

Thus, the wage pair is given by $(w(200), w(100)) = (100, 4)$. The payoff for the manager is equal to $\bar{u} = 4$ and the payoff for the owner is equal to:

$$(0.75) \cdot (200 - 100) + (0.25)(100 - 4) = 75 + 24 = 99.$$