Game Theory

Author:
Todd Gaugler

Professor:
Dr. Gaoquan Liu

January 18, 2012

CONTENTS

CHAPTER

ONE

STATIC GAMES OF COMPLETE INFORMATION

In this chapter, we aim to look at games of the following simple form:

1. First the players of the game simultaneously chose actions.
2. Second, the players receive payoffs that depend on the combination of actions just selected.

We will restrict our attention to games that involve *complete information*, whose definition will appear momentarily.

1.1 Basic Theory, Normal-Form Games and Nash Equilibrium

In what we call a **normal-form** representation of a game, each player in our game choses a strategy simultaneously, and the combination of strategies determines what payoff each player receives.

We can illustrate the normal-form representation with the famous example of a game: the *Prisoner's Dilemma*.
Example. Two suspects are arrested and charged with a crime. The police lack sufficient evidence to convict the suspects, unless at least one confesses. The police hold the suspects in separate cells, and explain the consequences that will follow from the actions they could take. If neither confesses, then both will be convicted of a minor offense and sentenced to one month in jail. If both confess, both will be sentenced to jail for six months. If one confesses and the other doesn't, then the confessor will be released immediately and the other will be sentenced to nine months in jail. We can represent this dilemma with the following bi-matrix:

	Mum	Fink
Mum	-1, -1	-9.0
Fink	0, -9	-6,-6

Where each tuple (x_1, x_2) represents the outcome of prisoner 1 in x_1 and prisoner 2 in x_2.

We now turn to the general case of a normal-form game.

Definition

The normal-form representation of a game specifies:

1. The players in the game
2. The strategies available to each player
3. The payoff received by each player for each combination of strategies that could be chosen by the players

These notes have many examples in which we have an n-player game in which the players are labeled $1, ...n$ and am arbitrary player is labeled i. Let S_i represent the set of strategies available to player i (called i's strategy space) and let s_i be an arbitrary member such that $s_i \in S_i$. Let $(s_1, ...s_n)$ denote a combination of strategies, one for each player, and let u_i denote player i's payoff function: $u_i(s_1, ..s_n)$ is the payoff to player i if the players choose the strategies $(s_1, ...s_n)$.

We end up with the following definition: The **normal-form representation** of an n-player game specifies the player's strategy space $S_1, ..S_n$ and their payoff functions $u_1, ...u_n$. We denote this game by $G = \{S_1, ..., S_n; u_1, ...u_n\}$.

Notice that although each player picks a strategy simultaneously, it doesn't imply that the players need to **act** simultaneously: it's enough that each chose an action on their own without the knowledge of the other's choices, as would be the case with the prisoners if they reached decisions at arbitrary times while in their separate cells. We will eventually see examples of games in which players move sequentially, and that the normal-form games can be transfered to extensive-form representations, which is also an often more convenient framework for analyzing real world and dynamic issues.

1.2 Iterated Elimination of Strictly Dominated Strategies

Since we have now figured out a way to formally describe games, it would be a natural transition to describe how to solve game-theoretic problems. Let us start with the prisoner's dilemma, using the idea that a rational player will not play a strictly dominated strategy.

Notice that in the prisoner's dilemma, if one suspect is going to confess, then the other would prefer to confess as well and so be in jail for six months as opposed to staying quiet (being 'Mum') and having to go to jail for nine months. Similarly, if one suspect is going to play Mum, then the other would rather confess as to be released immediately rather than being sent to jail for one month. So, for prisoner i, playing Mum is dominated by playing Fink (confessing)- for each strategy that player j could chose, the payoff to prisoner i from playing Mum is less than the payoff to i from playing Fink.

This can also be represented in any bi-matrix in which the payoffs $0, -1, -6, -9$ were replaced with payoffs T, R, P and S respectively, provided that $T > R > P > S$ so as to capture the idea of 'temptation, reward, punishment, and sucker' payoffs. More generally,

Definition

In the normal form game $G = \{S_1, ...S_n; u_1, ...u_n\}$, let s_i' and s_i'' be feasible strategies for player i. Strategy s_i' is **strictly dominated** by strategy s_i'' if for each feasible combination of the other player's strategies, i's payoff from playing s_i' is strictly less than i's payoff from playing s_i'':

$$u_i(s_1, ...s_{i-1}, s_i', s_{i+1}, ..s_n) < u(s_1, ..., s_{i-1}s_i'', s_{i+1}, ...s_n)$$

For each $(s_1, ..s_{i-1}s_i, s_{i+1}, ...s_n)$ that can be constructed from the other player's strategy spaces $S_1, ..., S_{i-1}S_i, S_{i+1}, ...S_n$.

We claim that rational players do not play strictly dominated strategies, because there is no belief that a player could hold about the strategies that the other player will choose such that it would be optimal to play such a strategy. Thus, in the Prisoner's Dilemma, a rational player would confess, so $(Fink, Fink)$ would be the outcome reached by two rational players, even though $(Fink, Fink)$ results in worse payoffs for both players than would (Mum, Mum).
Example. Consider the game described in figure 1.1. Player 1 has two strategies, and player 2 has three:

$$S_1 = \{\text{ Up, Down}\} \quad S_2 = \}\text{ Left, Middle, Right }\}$$

For player 1, neither Up nor Down is strictly dominated: Up is better than down if 2 plays left (because naturally, $1 > 0$) but Down is better than Up if 2 plays Right (because $2 > 0$).

For player 2 however, Right is strictly dominated by Middle (because $2 > 1$ and $1 > 0$), so a rational player 2 will not play Right. So, if player 1 knows that player 2 is rational then player 1 can eliminate from player 2's strategy space. In other words, we have the following:

	Left	Middle	Right
Up	1,0	1,2	0,1
Down	0,3	0,1	2,0

	Left	Middle
Up	1,0	1,2
Down	0,3	0,1

Figure 1.1: Right is dominated by Middle

But now, notice that Down is strictly dominated by Up for player 1, so if player 1 is rational (and player 1 knows that player 2 is rational, so that our new game applies) then player 1 will not play down. thus, we have the following new diagram, represented in Figure 1.2. But now notice that left is strictly dominated by Middle for player 2, leaving (Up, Middle) as the outcome of this game.

	Left	Middle
Up	1,0	1,2
Down	0,3	0,1

	Left	Middle
Up	1,0	1,2

Figure 1.2: Down is Dominated by Up

This process is called *iterated elimination of strictly dominated strategies*. Although it is based on the nice idea that rational players do not play strictly dominated strategies, the process has two drawbacks.

1. Each step requires a further assumption about what the players know about each other's rationality. If we want to be able to apply the process for an arbitrary number of steps, we need to assume that it is **common knowledge** that the players are rational. That is, we need to assume not only that all the players are rational, but also that the players know that all the players are rational, and that all the players know that all the players know that all the players are rational, and so on ad infinitum.

2. The process often produces a very imprecise prediction about the play of the game. Consider the game below, for example.

   In this game, there are no strictly dominated strategies to be eliminated. Since all the strategies in the game survive the iterated elimination of strictly dominated strategies, the process produces no prediction whatsoever about the play of the game.

	L	C	R
T	0,4	4,0	5,3
M	4,0	0,4	5,3
B	3,5	3,5	6,6

Figure 1.3:

It turns out we can turn to the concept of Nash equilibrium to help us analyze some of the normal-form games.

1.2.1 Motivation and Definition of Nash Equilibrium

One can aruge that if game theory is to provide a unique solution to a game theoretic problem, then the solution must be a Nash equilibrium, in the following sense: Suppose that game theory makes a unique prediction about the strategy each player will choose. in order for this prediction to be correct, it is necessary that each player be willing to choose the strategy predicted by the theory. Thus, each player's predicted strategy must be that player's best response to the predicted strategies of the other players. Such a prediction could be called **strategically stable** or bdself-enforcing, because no single player wants to deviate from his or her predicted strategy. This leads us to the following definition:

Definition

In the n-player normal-form game $G = \{S_1, ...S_n; u_1, ...u_n\}$, the strategies $(s_1^*, ..., s_n^*)$ are a **Nash Equilibrium** if, for each player i, s_i^* is (at least tied for) player i's best response to the strategies specified for the $n-1$ other players, $(s_1^*, ...s_{i-1}^*, s_{i+1}^*, ...s_n^*)$:

$$u_i(s_1^*, ...s_{i-1}^*, s_i^* s_{i+1}^*, ...s_n) \geq u_i(s_1^*, ...s_{i-1}^*, s_i, s_{i+1}^*, ...s_n)$$

for every feasible strategy $s_i \in S_i$, that is s_i^* solves:

$$max_{s_i \in S_i} u_i(s_1^*, ...s_{i-1}^*, s_i, s_{i+1}^*, ...s_n^*)$$

Alternatively, saying that $(s_1', ..., s_n')$ is *not* a Nash equilibrium of G is equivalent to saying that there exists some player i such that s_i' is not a best response to $s_1', ...s_{i-1}', s_{i+1}', ..., s_n')$. So, if our theory offers the strategies $(s_1', ...s_n')$ as the solution to our game, but these strategies are not a Nash equilibrium, then at least one player will have an incentive to deviate from the theory's prediction, so the theory will be falsified by the actual play of the game.

A closely related motivation for Nash equilibrium involves the idea of convention: if a convention is to the develop about how to play a given game, then the strategies prescribed by the convention must b e a Nash equilibrium, else at lease one player will not abide by the convention.

Example. One can go and look at our previous solutions to games, and will notice that our solutions satisfy a Nash equilibrium. In figure 1.3, one can notice that (6,6) satisfies a Nash equilibrium, and is the only pair to do so.

Recall that our solutions to the Prisoner's dilemma and the Bi-Matrix shown in Figure 1.1 were found by iterated elimination of strictly dominated strategies. We found strategies that were the only ones that survived iterated elimination- this result can be generalized: if iterated elimination of strictly dominated strategies eliminates all but the strategies $(s_1^*, ..., s_n^*)$, then these strategies are the unique Nash equilibrium of the game.

If the strategies $(s_1^*, ..., s_n^*)$ are a Nash equilibrium then they survive iterated elimination of strictly dominated strategies, but there can be strategies that survive iterated elimination of strictly dominated strategies that are **not** part of any Nash equilibrium.

Having shown that Nash equilibrium is a stronger solution concept than iterated elimination of strictly dominated strategies, we must now ask whether Nash equilibrium is too strong a solution concept. That is, can we be sure that a Nash equilibrium exists? It was shown by Nash in 1950 that in any finite game there exists at least one Nash equilibrium. In some of the future sections, we will simply rely on Nash's Theorem (or its analog for stronger equilibrium concepts) and will simply assume that an equilibrium exists.

Example. This problem is called *The Battle of the Sexes*. This problem shows that a game can have multiple Nash equilibria, and also will be useful in the discussions of mixed strategies. In the traditional exposition of the game, a man and a woman are trying to decide on an evening's entertainment, but to avoid sexism we'll remove gender from this example. While at separate workplaces, Pat and Chris must choose to attend either the opera or a prize fight. Both players would rather spend the evening together than apart, but Pat would rather be together at the prize fight while Chris would rather be together at the opera. We have the following bi-matrix:

		Pat	
		Opera	Fight
Chris	Opera	2,1	0,0
	Fight	0,0	1,2

Figure 1.4: The Battle of the Sexes

Before, we argued that if game theory is to provide a unique solution to a game then the solution must be a Nash equilibrium. This argument ignores the possibility of games in which game theory does not provide a unique solution. We also argued that if a convention is to develop about how to play a basic game, then the strategies prescribed by the convention must be a Nash equilibrium, but this argument ignores the possibility of games for which a convention will not develop. In some games with multiple Nash equilibria one equilibrium stands out as the compelling solution to the game. Thus, the existence of multiple Nash equilibria is not a problem in and of itself. In the battle of the sexes, however, $(Opera,\ Opera)$ and $(Fight,\ Fight)$ seem equally compelling, which suggests that there may be games from which game theory does not provide a unique solution and no convention will develop. In such games, Nash equilibrium loses much of its appeal as a prediction of play.

At this point, we have the following propositions, whose proofs will be ommited since they follow pretty easily from definitions and a proofs by contradiction:

Proposition 1. *In the n-player normal form game $G = \{S_1, ..., S_n; u_1, .., , u_n\}$, if iterated elimination of strictly dominated strategies eliminates all but the strategies $(s_1^*, ..., s_n^*)$, then these strategies are the unique Nash equilibrium of the game.*

Proposition 2. *In the n-player game $G = \{S_1, ..., S_n; u_1, .., , u_n\}$, if the strategies $(s_1^*, ..., s_n^*)$ are a Nash equilibrium, then they survive iterated elimination of strictly dominated strategies.*

1.3 Applications

1.3.1 Cournot Model of Duopoly

It turns out that actually, Cournot(1838) had anticipated Nash's definition of equilibrium by over a century. Naturally it follows that some of Cournot's work is some of the classics of game theory, and is also one of the cornerstones of the theory of industrial organization. WE consider a very simple version of Cournot's model, and return to variations on the model later. We want to illustrate the following:

1. The translation of an informal statement of a problem into a normal-form representation of a game
2. The computations involved in solving for the game's Nash equilibrium
3. Iterated elimination of strictly dominated strategies.

Let q_1 and q_2 be the quantities of a homogeneous product produced by firms 1 and 2, respectively. Let $P(Q) = a - Q$ be the market-clearing price when the aggregate quantity on the market is $Q = q_1 + q_2$. Assume that the total cost to firm i of producing quantity q_i is $C_i(q_i) = cq_i$. In other words, there are no fixed costs and the marginal cost is constant at c, where we assume $c < a$. Following Cournot, suppose that the firms choose their quantities simultaneously.

In order to find the Nash equilibrium of the Cournot game, we first change the problem into a normal-form game. recall that the normal form-representation specifies:

1. The players in the game
2. The strategies available to each player
3. The payoff received by each player for each combination of strategies that could be chosen by the players

There are of course two players in any doupoly game, the two firms. In the Cournot model, the strategies available to each firm are the different quantities it might produce. We assume that output is continuously divisible. Naturally, negative outputs aren't feasible, so each firm's strategy space can be represented as $S_= [0, \infty)$, where each s_i represents a quantity choice $q_i \geq 0$. Because $P(Q) = 0$ for $Q \geq a$, neither firm will produce a quantity $q_i > a$. It remains to specify the payoff to firm i as a function of the strategies chosen by it and by the other firm, and to define and solve for equilibrium. We assume that the firm's payoff is simply its profit. Thus, the payoff $u_i(s_i, s_j)$ in a general two-player game in normal form can be written here as:

$$\pi_i(q_i, q_j) = q_i[P(q_i + q_j) - c] = q_i[a - (q_i + q_j) - c]$$

Recall that in a two player game in normal form, the strategy pair (s_i^*, s_j^*) is a Nash equilibrium if for each player i:

$$u_i(s_i^*, s_j^*) \geq u_i(s_i, s_j^*)$$

for every feasible strategy s_i in S_i. Equivalently, for each player i, s_i^* must solve the optimization problem:

$$max_{s_i \in S_i}\ u_i(s_i, s_j^*)$$

In the Cournot doupoly model, the analogous statement is that the quantity pair (q_1^*, q_2^*) is a Nash equilibrium if for each firm i q_i^* solves:

$$max_{0 \leq q_i < \infty}\ \pi_i(q_i, q_j^*) = max_{0 \leq q_i < \infty} q_i[a - (q_i + q_j^*) - c]$$

Assuming that $q_j^* < a - c$, the first order condition for firm i's optimization problem is both necessary and sufficient; it yields:

$$q_i = \frac{1}{2}(a - q_j^* - c) \tag{1.3.1}$$

So, if the quantity pair (q_1^*, q_2^*) is to be a Nash equilibrium, the firm's quantity choices must satisfy:

$$q_1^* = \frac{1}{2}(a - q_2^* - c)$$

and

$$q_2^* = \frac{1}{2}(a - q_1^* - c)$$

Solving this pair of equations yields:

$$q_1^* = q_2^* = \frac{a - c}{3}$$

which it turns out is clearly less than $a - c$, which is something we assumed.

The intuitive idea behind this equilibrium is pretty simple. Naturally each firm aims to be a monopolist in the market, in which case it would chose q_i to maximize $q_i(q_1, 0$- it would produce the monopoly quantity $(a-c)/2$ and earn the monopoly profit $\pi_i(q_m, 0) = (a-c)^2/4$. Given that there are two firms, aggregate profits for the duopoly would be maximized by setting the aggregate quantity $q_1 + q_2$ equal to the monopoly quantity q_m, as would occur if $q_i = q_m/2$ for each i, for example. The problem with this arrangement is that each firm has an incentive to deviate, because the monopoly quantity is low, the associated price $P(q_m)$ is high, and at this price each firm would like to increase its quantity in spite of the fact that such an increase in production drives down the market-clearing price. In the Cournot equilibrium, in contrast, the aggregate quantity is higher, so the associated price is lower, so the temptation to increase output is reduced- reduced by just enough that each firm is just deterred from increasing its output by the realization that the market clearing price will fall.

We could have solved this question graphically. Equation (1.3.1) gives firm i's best response to firm j's equilibrium strategy q_j^*. Analogous reasoning leads to firm 2's best response to an arbitrary strategy by firm 1, and firm 1's best response to an arbitrary strategy by firm 2. Assuming that firm 1's strategy satisfies $q_1 < a - c$, firm 2's best response is:

$$R_2(q_1) = \frac{1}{2}(a - 1_c - c)$$

And likewise,

$$R_1(q_2) = \frac{1}{2}(a - q_2 - c)$$

This can e shown through the following diagram:

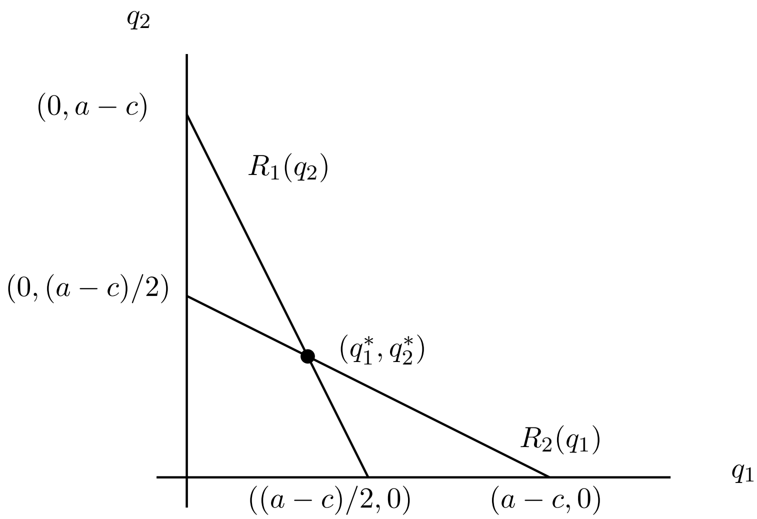

You also could have approached this question through iterated eliminations of strictly dominated strategies. First, notice that the monopoly quantity q_m dominates any higher quantity. Then, notice that the quantity $(a-c)/4$ strictly dominates any lower quantity. After proving this, one notices that the remaining quantities in each firm's strategy space lie in the interval between $(a-c)/4$ and $(a-c)/2$. Repeating these argument,s you can make this interval even smaller- repeating it infinitely many times, the intervals converge to the single point $q_i^* = (a-c)/2$.

1.3.2 Bertrand Model of Duopoly

We can now introduce a model of how two duopolists might interact, based on Bertrand's suggestion that firms actually choose prices, rather than quantities as in Cournot's model. It is important to notice that Bertrand's model is actually a different game than Cournot's model; the strategy spaces are different, the payoffs are different, and the behavior in the Nash equilibria of the two models are different. In both games however, the equilibrium concept used is the Nash equilibrium defined as previous.

We consider the case of differentiated products. If firms 1 and 2 choose prices p_1, p_2 respectively, the quantity that consumers demand from firm i is

$$q_i(p_i, p_j) = a - p_1 + bp_j$$

where $b > 0$ reflects the extent to which firm i's product is a substitute for firm j's product. We assume that there are no fixed costs of production and that marginal costs are constant at c, where $c < a$, ad that the firms act simultaneously.

As before, our first step is to translate this real world question into a normal-form game. There are again two players, but this time the strategies available to each firm are the different prices it might charge as opposed to the different quantities it might produce. Each firms strategy space can again be represented as $S_i = [0, \infty)$, and a typical strategy s_i is now a price choice, $p_1 \geq 0$.

The profit to firm i when it choses price p_i and its rival chooses price p_j is as follows:

$$\pi_i(p_i, p_j) = q_i(p_i, p_j)[p_i - c] = [a - p_i + bp_j][p_i - c]$$

So, the price pair (p_1^*, p_2^*) is a Nash equilibrium if for each firm i, p_i^* solves:

$$max_{0 \leq p_i < \infty} \pi_i(p_i, p_j^*) = max_{0 \leq p_i < \infty} [a - p_i + bp_j^*][p_i - c]$$

So, the solution to firm i's optimization problem is:

$$p_i^* = \frac{1}{2}(a + bp_j^* + c)$$

From which it follows that if the price pair (p_1^*, p_2^*) is a Nash equilibrium, then

$$p_1^* = \frac{1}{2}(a + bp_2^* + c)$$

$$p_2^* = \frac{1}{2}(a + bp_1^* + c)$$

Solving this pair of equations again, we get:

$$p_1^* = p_2^* = \frac{a + c}{2 - b}$$

1.3.3 Final-Offer Arbitration

Many public sector workers are forbidden to strike, and instead wage disputes are settled by binding arbitration. Many other disputes also involve arbitration. The two major forms of arbitration are **conventional** and *final-offer* arbitration. In final-offer arbitration, the two sides make wage offers and then the arbitrator picks one of the offers as the settlement. In conventional arbitration, the arbitrator is free to impose any wage as the settlement. We now derive the Nash equilibrium wage offers in a model of final-offer arbitration developed by Farber (1980).

Suppose that the two parties in dispute are a firm and a union,, and the dispute concerns wages. First, the firm and the union simultaneously make offers, denoted by w_f, w_u respectively. Then, the arbitrator chooses one of the two offers as the settlement. Assume that the arbitrator has an ideal settlement she would like to impose, denoted by x. Assume that further that, after observing the parties' offers, the arbitrator simply chooses the offer that is closer to x; provided that $w_f < w_u$, the arbitrator chooses w_f if $x < (w_f + w_u)/2$ and choses w_u if $x > (w_f + w_u)/2$. The arbitrator knows x, but the parties do not. The parties beleive that x is randomly distributed according to a cumulative probability distribution $F(x)$ with density $f(x)$. Given our specification of the arbitrator's behavior, if the offers are w_f, w_u, then the parties believe that the probabilities $P\{w_f \ chosen \ \}$ and $P\{w_u \ chosen \ \}$ can be expressed as follows:

$$P\{w_f \ chosen \ \} = P\left\{x < \frac{w_f + w_u}{2}\right\} = F\left(\frac{w_f + w_u}{2}\right)$$

and thus,

$$P\{w_u \ chosen \ \} = 1 - F\left(\frac{w_f + w_u}{2}\right)$$

So, based on what we know about probability and their expected values, the expected wage settlement is:

$$w_f \cdot P\{w_f \ chosen \ \} + w_u \cdot P\{w_u \ chosen \ \} = w_f \cdot F\left(\frac{w_f + w_u}{2}\right) + w_u \cdot \left[1 - F\left(\frac{w_f + w_u}{2}\right)\right]$$

We assume that the firm wants to minimize the expected wage settlement, and the union wants to maximize it. If the pair of offers (w_f^*, w_u^*) is a Nash equilibrium, then w_f^* and w_u^* must solve:

$$min_{w_f} \ w_f \cdot F\left(\frac{w_f + w_u^*}{2}\right) + w_u^* \cdot \left[1 - F\left(\frac{w_f + w_u^*}{2}\right)\right]$$

and

$$max_{w_f} \ w_f^* \cdot F\left(\frac{w_f^* + w_u}{2}\right) + w_u \cdot \left[1 - F\left(\frac{w_f^* + w_u}{2}\right)\right]$$

So, the wage pair must solve the first-order conditions for these optimization problems:

$$(w_u^* - w_f^*) \cdot \frac{1}{2} f\left(\frac{w_f^* + w_u^*}{2}\right) = F\left(\frac{w_f^* + w_u^*}{2}\right)$$

and:

$$(w_u^* - w_f^*) \cdot \frac{1}{2} f\left(\frac{w_f^* + w_u^*}{2}\right) = 1 - F\left(\frac{w_f^* + w_u^*}{2}\right)$$

Since the right hands of these first order conditions are equal, the right hands sides are also equal, which implies that:

$$F\left(\frac{w_f^* + w_u^*}{2}\right) = \frac{1}{2}$$

Substituting this into either one of the first-order conditions yields:

$$w_u^* - w_f^* = \frac{1}{f\left(\frac{w_f^* + w_u^*}{2}\right)}$$

Which can be interpreted by saying: "the gap between the offers must equal the reciprocal of the value of the density function f at the median of the arbitrator's preferred settlement."

1.3.4 The Problem of the Commons

Consider n farmers in a village. Each summer, all the farmers graze their goats on the village green. Denote the number of goats the i^{th} farmer owns by g_i and the total number of goats in the village by $G = g_1 + ... + g_n$. The cost of buying and caring for a goat is c, independent of how many goats a farmer owns. The value to a farmer of grazing a goat on the green when a total of G goats are grazing is $v(G)$ per goat.

Since a goat needs at least a certain amount of grass to survive, there is a maximum number of goats that can be grazed on the green,

$$G_{max} : v(g) > 0 \quad \text{for } G < G_{max} \quad \text{but } v(G) = 0 \text{ for } G \geq G_{max}$$

Also, since the first few goats have plenty of room to graze, adding one more does little harm to those already grazing- but when so many goats are grazing, then they are all just barely surviving, so adding one more drastically harms the rest.

During the spring, the farmers simultaneously choose how many goats to own. Assume goats are continuously divisible. A strategy for farmer i is the choice of a number of goats to graze on the village green g_i. Assuming that the strategy space is $[0, \infty)$ covers all the choices that could be of interest to the farmer; $[0, G_{max})$ would also suffice. The payoff to farmer i from grazing g_i goats when the numbers of goats grazed by the other farmers are $(g_1, ..., g_{i-1}, g_i, g_{i+1}, ..., g_n)$ is

$$g_i v(g_1, ..., g_{i-1}, g_i, g_{i+1}, ..., g_n) - cg_i$$

So, if $(g_1^*, ..., g_n^*)$ is to be a Nash equilibrium, then for each i, g_i^* must maximize the above equation given that the other farmers choose $(g_1^*, ...g_{i-1}^*, g_{i+1}^*, ..., g_n^*)$. The first order-condition for this optimization problem is:

$$v(g_i + g_{-i}^*) + g_i v'(g_i + g_{-i}^*) - c = 0$$

Where g_{-i}^* denotes $g_1^* + ... + g_{i-1}^* + g_{i+1}^* + ... + g_n^*$. Substituting g_i^* into our above equation and summing over all n farmers first order conditions, and then dividing by n gives us the following:

$$v(G^*) + \frac{1}{n} G^* v'(G^*) - c = 0$$

Where G^* denotes $g_1^* + ... + g_n^*$. In contrast, the social optimum denoted by G^{**} solves the following:

$$max_{0 \leq G < \infty} \; Gv(G) - Gc$$

The first order condition for which is:

$$v(G^{**}) + G^{**} v'(G^{**}) - c = 0$$

Comparing our two answers shows that $G^* > G^{**}$: too many goats are grazed in the Nash equilibrium compared to the social optimum. The first order condition reflects the incentives faced by a farmer who is already grazing g_i goats, but is considering adding one more. The value of the additional goat is $v(g_i + g^*_{-i})$ and its cost is c. The harm to the farmers existing goats is $v'(g_i + g_{-i})$ per goat, or $g_i v'(g_i + g^*_{-i})$ in total. The common resource is over utilized because each farmer considers only his or her own incentives, not the effect of his or her actions on the other farmers- hence the presence of $G^* v'(G^*)/n$ in one of our conditions but $G^{**}v'(G^{**})$ in the other.

1.4 Advanced Theory: Mixed Strategies and Existence of Equilibrium

1.4.1 Mixed Strategies

In previous sections, we defined S_i to be the set of strategies available to some player i, and the combination of strategies $(s^*_1, ..., s^*_n)$ to be a Nash equilibrium if for each player i, s^*_i is player i's best response to the strategies of the $n-1$ other players:

$$u_i(s^*_1, ..., s^*_{i-1}, s^*_i, s^*_{i+1}, ..., s^*_n) \geq u_i(s^*_1, ..., s^*_{i-1}, s_i, s^*_{i+1}, ..., s^*_n)$$

for every strategy s_i in S_i. By this definition, there is no Nash equilibrium in the following game, which is known as *Matching Pennies*:

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

In this game, each players strategy space is $\{Heads,\ Tails\}$. Suppose that each player has a penny and must choose whether to display it heads or tails up. If the two pennies match, then player 2 wins player 1's penny, if the pennies don't match then player 1 wins player 2's penny. No pair of strategies can satisfy a Nash equilibrium, since if the players strategies match $(H, H), (T, T)$, then player 1 prefers to switch strategies, while if the pairs don't match, then player 2 prefers to switch strategies.

The most distinguishing feature of Matching Pennies is that each player would like to outguess the other. Versions of this game appear in all kinds of every-day competitions, poker, baseball, war, and other games. In poker, the analogous question to this game is how often to bluff: if player i is known never to bluff then i's opponents will fold whenever i bids aggressively, and bluffing too often can be a losing strategy.

In any game where each player would like to outguess the other(s), there is no Nash equilibrium because the solution to such a game necessarily involves uncertainty about what the players will do. We now introduce the notion of a **mixed strategy** which we will interpret in terms of one player's uncertainty about what another player will do. Formally, a mixed strategy for player i is a probability distribution over the strategies in S_i. We will from here on in refer to the strategies in S_i as player i's *pure strategies*.

More generally, suppose that player i has K pure strategies, $S_i = \{s_{i1}, s_{i2}, ..., s_{iK}\}$. Then a mixed strategy for player i is a probability distribution $(p_{i1}, p_{i2}, ..., p_{iK})$ where p_{ik} is the probability that player i will play strategy s_{ik} for $k = 1, 2, ...K$. We have the following definition:

Definition

In the normal form game $G = \{S_1, ..., S_n; u_1, ..., u_n\}$, suppose that $S_i = \{s_{i1}, ..., s_{iK}\}$. Then a **mixed strategy** for player i is a probability distribution $p_i = (p_{i1}, ..., p_{iK})$, where $0 \leq p_{ik} \leq 1$ for $k = 1, ..., K$ and $p_{i1} + \ldots p_{iK} = 1$.

Recall that if a strategy s_i is strictly dominated, then there is no belief that a rational player i would find it optimal to play s_i. The converse is true, provided we allow for mixed strategies: if there is no belief that player i could hold such that it would be optimal to play the strategy s_i, then there exists a strategy that strictly dominates s_i.

		Player 2	
		L	R
Player 1	T	3, –	0, –
	M	0,–	3,–
	B	1,–	1,–

Figure 1.5:

Figure 1.5 shows that a given pure strategy may be strictly dominated by a mixed strategy, even if the pure strategy is not strictly dominated by any other pure strategy. In this game, for any belief $(q, 1-q)$ that player 1 could hold about player 2's play, player 1's best response is either T (if $q \geq 1/2$) or M (if $g \leq 1/2$), but never B. Yet, B is not strictly dominated by either T or M. The key to notice here is that B is strictly dominated by a mixed strategy, if player 1 expects T with probability 1/2 and M with probability 1/2 then 1's expected payoff is 3/2 no matter what (pure or mixed) strategy 2 plays, and 3/2 exceeds the payoff of 1 that playing B surely produces.

There exist examples of bi-matrices that show that a given pure strategy can be a best response to a mixed strategy, even if the pure strategy is not a best response to any other pure strategy.

		Player 2	
		L	R
	T	3, –	0, –
Player 1	M	0,–	3,–
	B	2,–	2,–

Figure 1.6:

Figure 1.6 demonstrates this. In this game, B is not a best response for player 1 to either L or R by player 2, but B is the best response for player 1 to the mixed strategy $(q, 1-q)$ by player 2, provided that $1/3 < q < 2/3$. This example illustrates the role of mixed strategies in the "belief that player i could hold".

1.4.2 Existence of Nash Equilibrium

Recall that the definition of Nash equilibrium given in an earlier section guarantees that each player's pure strategy is a best response to the other player's pure strategies. It would be nice if we could extend this definition to mixed strategies, and we can do so simply be requiring that each player's mixed strategy be a best response to the other player's mixed strategies. Since any pure strategy can be represented as the mixed strategy that puts zero probability on all of the players other pure strategies, this extended definition also satisfies the older one.

Computing player i's best response to a mixed strategy by player j illustrates the interpretation of player j's mixed strategy as representing player i's uncertainty about what player j will do. We begin with Matching Pennies as an example. Suppose that player 1 believes that player 2 will play Heads with probability q, and will play Tails with probability $1 - q$.

Given this belief, player 1's expected payoffs are $q(-1) + (1-q)(1) = 1 - 2q$ from playing Heads and $q \cdot 1 + (1-q) \cdot (-1) = 2q - 1$ from playing Tails. Since $1 - 2q > 2q - 1$ if and only if $q < 1/2$, player 1's best pure strategy response is Heads if $q < 1/2$ and Tails if $q > 1/2$, and indifferent between Heads and Tails if $q = 1/2$.

Let $(r, 1-r)$ denote the mixed strategy in which player 1 plays Heads with probability r. For each value of q between zero and one, we now compute the values of r, denoted $r^*(q)$ such that $(r, 1-r)$ is a best response for player 1 to $(q, 1-q)$ by player 2. This is illustrated by the following diagram:

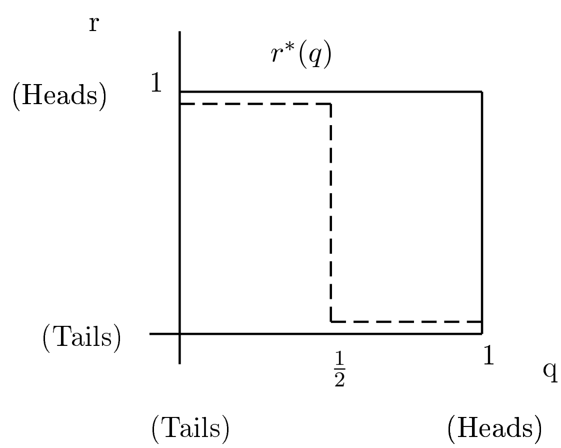

Figure 1.7:

Player 1's expected payoff from playing $(r, 1-r)$ when player 2 plays $(q, 1-q)$ is:

$$rq \cdot (-1) + r(1-q) \cdot 1 + (1-r)q \cdot 1 + (1-r)(1-q) \cdot (-1) = (2q-1) + r(2-4q) \quad (1.4.1)$$

where rq is the probability of $(Heads, Heads)$, $r(q-1)$ is the probability of $(Heads, Tails)$, and so on. Since player 1's payoff is increasing in r if $2 - 4q > 0$ and decreasing in r if $2 - 4q < 0$, player 1's best response is $r = 1$(Heads) if $q < 1/2$ and $r = 0$ (Tails) if $q > 1/2$ as indicated by the horizontal segments of $r^*(q)$ in Figure 1.7.

The nature of player 1's best response to $(q, 1-q)$ changes when $q = 1/2$. As we noted, when $q = 1/2$, player 1 is completely indifferent. Furthermore, because player 1's expected payoff (1.4.1) is independent of r when $q = 1/2$, player 1 is also indifferent among all mixed strategies $(r, 1-r)$. That is, when $q = 1/2$ the mixed strategy $(r, 1-r)$ is a best response to $(q, 1-q)$ for any value of r between zero and one. Thus $r^*(1/2)$ is the entire interval $[0, 1]$ as indicated by the vertical segment in figure 1.7 at $q = 1/2$.

To now derive player i's best response to player j's mixed strategy more generally, and to give a formal statement of the extended definition of Nash equilibrium, we restrict our attention to the two player case. let J denote the number of pure strategies in S_1 and K the number in S_2. We will write:

$$S_1 = \{s_{11}, ..., s_{1J}\} \quad S_2 = \{s_{21}, ..., s_{2K}\}$$

and we will use s_{1j} and s_{2k} to denote arbitrary strategies from S_1 and S_2. If player 1 believes that player 2 will play the strategies $(s_{21},...,s_{2K})$ with the probabilities $(p_{21},..,p_{2K})$ then player 1's expected payoff from playing the pure strategy s_{1j} is as follows:

$$\sum_{k=1}^{K} p_{2k} u_1(s_{1j}, s_{2k})$$

and player 1's expected payoff from playing the mixed strategy $p_1 = (p_{11},...,p_{1K})$ is:

$$v_1(p_1,p_2) = \sum_{j=1}^{J} p_{1j} \left[\sum_{k=1}^{K} p_{2k} u_1(s_{1j}, s_{2k}) \right] = \sum_{j=1}^{J} \sum_{k=1}^{K} p_{1j} \cdot p_{2k} u_1(s_{1j}, s_{2k})$$

where $p_{1j} \cdot p_{2k}$ is the probability that 1 plays s_{1j} and 2 plays s_{2k}. Player 1's expected payoff from the mixed strategy p_1, is the weighted sum of all the expected payoffs for each of the pure strategies $\{s_{11},...,s_{1J}\}$ where the weights are the probabilities $(p_{11},...,p_{1J})$. So, for the mixed strategy $(p_{11},...,p_{1J})$ to be a best response for player 1 to player 2's mixed strategy p_2, it must be that $p_{1j} > 0$ only if:

$$\sum_{k=1}^{K} p_{2k} u_1(s_{1j}, s_{2k}) \geq \sum_{k=1}^{K} p_{2k} u_1(s_{1j'}, s_{2k})$$

for every $s_{ij'} \in S_1$. Giving a formal statement of the extended definitino of Nash equilibrium requires computing player 2's expected payoff when players 1 and 2 play the mixed strategies p_1, p_2 respectively. If player 2 believes that player 1 will play the strategies $(s_{11},...,s_{1J})$ with probabilities $(p_{11},...,p_{1J})$ then player 2's expected payoff from playing the strategies $(s_{21},...,s_{2K})$ with probabilities $(p_{21},...,p_{2K})$ is:

$$v_2(p_1,p_2) = \sum_{k=1}^{K} p_{2k} \left[\sum_{j=1}^{J} p_{1j} u_2(s_{1j}, s_{2k}) \right] = \sum_{j=1}^{J} \sum_{k=1}^{K} p_{1j} \cdot p_{2k} u_2(s_{1j}, s_{2k})$$

So, given $v_1(p_1,p_2)$ and $v_2(p_1,p_2)$ we can restate the requirement of Nash equilibrium that each players mixed strategy be a best response to the other player's mixed strategy: for the pair of mixed strategies (p_1^*, p_2^*) to be a Nash equilibrium, p_1^* must satisfy:

$$v_1(p_1^*, p_2^*) \geq v_1(p_1, p_2^*) \tag{1.4.2}$$

for all p_1 over S_1, and p_2^* must satisfy:

$$v_2(p_1^*, p_2^*) \geq v_2(p_1^*, p_2) \tag{1.4.3}$$

for every probability distribution p_2 over S_2.

Definition

In the two player normal-form game $G = \{S_1, S_1; u_1, u_2\}$, the mixed strategies (p_1^*, p_2^*) are a **Nash equilibrium** if each player's mixed strategy is a best response to the other player's mixed strategy, and the above equations (1.4.2) and (1.4.3) must hold.

We can apply this definition to lots of games we've already seen, like Matching Pennies and the Battle of the Sexes using the graphical representation of player i's best response to player j's mixed strategy. We have already seen one example of such a representation in figure 1.7, let us now show that $q^*(r)$ should look like for the same game:

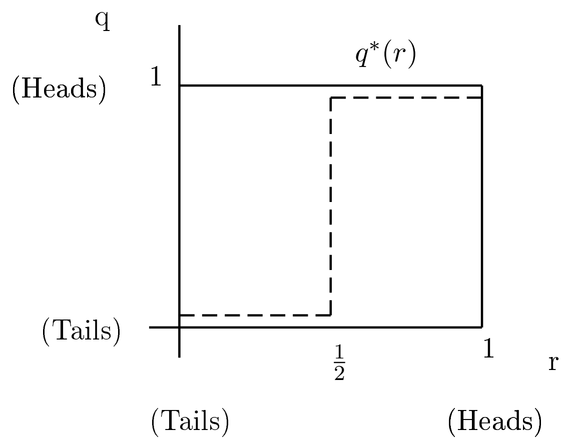

Figure 1.8:

Flipping this diagram, we get:

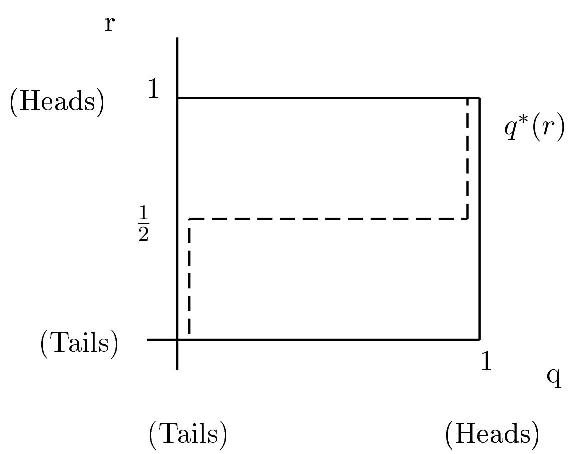

Figure 1.9:

Which isn't really the most helpful diagram, but combining it with figure 1.7 we get figure 1.10.

This particular figure is analogous to what we arrived at from the Cournot analysis in a prior section. Just as the intersection of the best-response functions $R_1(q_2)$ and $R_2(q_1)$ gave us the

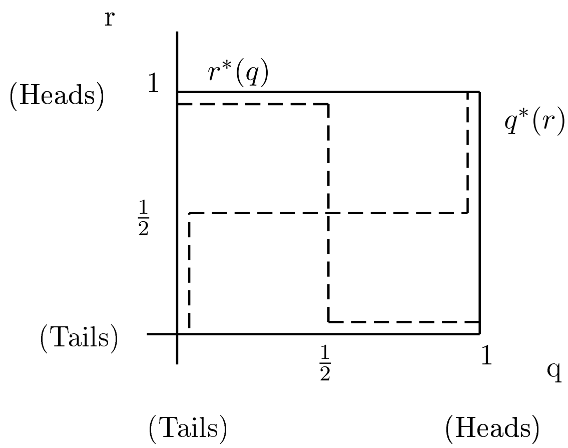

Figure 1.10:

Nash equilibrium of the Cournot game, the intersections of $r^*(q)$ and $q^*(r)$ give us the Nash equilibrium in Matching Pennies.

Another example of a mixed-strategy Nash equilibrium is the Battles of the Sexes. Let $(q, 1-q)$ be the mixed strategy in which Pat plays Opera with probability q, and let $(r, 1-r)$ be the mixed strategy in which Chris plays Opera with probability r. If Pat plays $(q, 10q)$, then Chris's expected payoffs are

$$q \cdot 2 + (1-q) \cdot 0 = 2q$$

from playing Opera and

$$q \cdot 0 + (1-q) \cdot 1 = 1 - q$$

from playing Fight. So, if $q > 1/3$ then Chris's best response is Opera ($i.e., r = 0$), if $q < 1/3$ then Chris's best response is Fight (i.e., $r = 0$) and if $q = 1/3$ then any value of r is a best response. We have the following diagram:

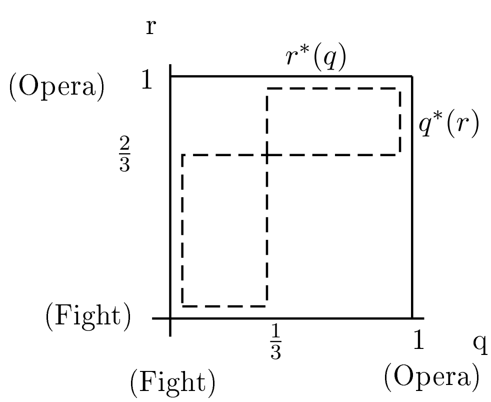

Figure 1.11:

Notice that unlike in our other diagrams, there are actually three intersections of $r^*(q)$ and $q^*(r)$. They are $(q = 0, r = 0), (q = 1, r = 1), (q = 1/3, r = 2/3)$. The other two intersections represent the pure-strategy Nash equilibria (Fight, Fight) and (Opera, Opera).

In any game, a Nash Equilibrium (involving pure or mixed strategies) appears as an intersection of the players' best response correspondences, even when there are more than two players, and even when some of the players have more than two pure strategies. Unfortunately, the only games in which the players best-response correspondences have simple graphical representations are two-player games in which each player only has two strategies. We turn next to a graphical argument that any such game has a Nash equilibrium.

Consider the following payoffs for player 1:

		Left	Right
Player 1	Up	x, –	y, –
	Down	z, –	w, –

There are two important comparisons: x versus z, and y versus w. Based on these comparisons, we can define four main cases:

1. $x > z$ and $y > w$

2. $x < z$ and $y < w$

3. $x > z$and $y < w$

4. $x < z$ and $y > w$

We first discuss these four main cases, then turn to the remaining cases involving $x = z$ or $y = w$.

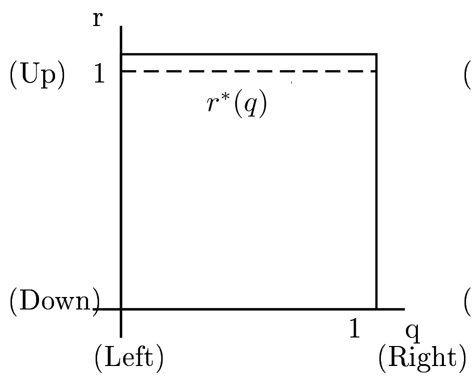

Case(i)

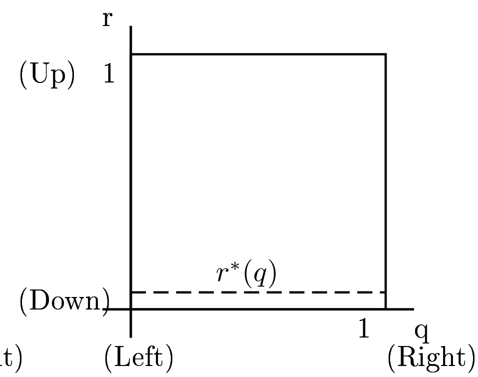

Case(ii)

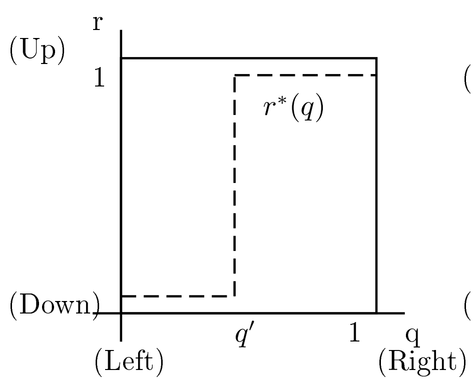

Case(iii)

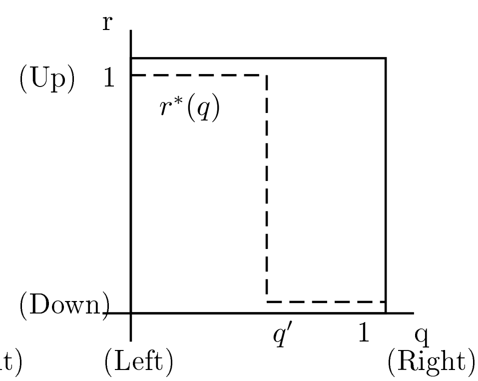

Case(iv)

- In case (i), Up strictly dominates Down for player 1, and in case (ii), Down Strictly Dominates Up for player 1

  Notice that now that if $(q, 1-q)$ is a mixed strategy for player 2, where q is the probability that player 2 will play left, then in case (i) there is no value of q such that Down is optimal for player 1, and in case (ii) there is no value of q such that Up is optimal for player 1. Letting $(r, 1-r)$ denote a mixed strategy for player 1, where r is the probability that 1 will play Up, we can represent the best-response correspondences for cases (i) and (ii) as in our figure above.

- In cases (iii) and (iv), neither Up nor Down is strictly dominated.

  Thus, Up must be optimal for some values of q and Down must be optimal for others. Let $q' = (w-y)/(x-z+w-y)$. Then in case (iii), Up is optimal for $q > q'$ and in case (iv) the reverse is true.

- Since $q' = 1$ if $x = z$ and $q' = z$ if $y = w$, the best- response correspondences for cases involving either $x - z$ or $y - w$ are $L - shaped$.

Adding arbitrary payoffs to our original payoff matrix and performing the analogous computations yield the same four best-response correspondence diagrams, except that the horizontal and vertical axes are swapped.

The crucial point to get from these diagrams is that given any of the four best-response correspondences for player 1, $r^*(q)$, and any of the four for player 2, $q^*(r)$, the pair of best-response correspondences has at least one intersection, so the game must have at least one Nash equilibrium. One can check that this is true with all sixteen possible combinations by overlapping each graph in turn. As a result, we know there There can be:

1. A single pure-strategy Nash equilibrium
2. A single mixed-strategy Nash equilibrium
3. Two pure-strategy equilibria and a single mixed-strategy equilibrium

We have already seen a few examples of these situations.
Theorem 3. *In the n-player normal-form game $G = \{S_1, ..., S_n; u_1, ..., u_n\}$, if n is finite and S_i is finite for every i then there exists at least one Nash equilibrium, possibly involving mixed strategies.*

Proof. Actually, the Proof of Nash's Equilibrium can be done using the *fixed-point theorem*. One application of the fixed-point theorem is that you can take a continuous function $f : [0, 1] \to [0, 1]$ and you are guaranteed there there exists at least one point x' such that $f(x') = x'$.

The idea is to follow two steps, using the fixed-point theorem:

1. Showing that any fixed point of a certain correspondence is a Nash Equilibrium
2. using an appropriate fixed-point theorem application to show that this correspondence must be a fixed point

The application of the fixed point theorem is due to Kakutani, who generalized Brouwer's theorem to allow for correspondences as well as functions.

The n-player best-response correspondence is computed from the n individual players' best-response correspondences as follows: Consider an arbitrary combination of mixed strategies $(p_1, ..., p_n)$. For each player i, derive i's best response(s) to the other players mixed strategies. Then construct the set of all possible combinations of one such best response for each player. A combination of mixed strategies $(p_1^*, ..., p_n^*)$is a fixed point of this correspondence if $(p_1^*, ..., p_n^*)$ belongs to the set of all possible combinations of the players' best responses to $(p_1^*, ..., p_n^*)$, but this is precisely the statement that $(p_1^*, ..., p_n^*)$ is a Nash equilibrium. This completes our first step.

Step two involves the fact that each players best response correspondence is continuous, in an appropriate sense of 'continuity'. All one needs to think about is a continuous function $f : [0, 1] \to [0, 1]$, and applying the variation on the fixed-point theorem, we can relate f to the best-response correspondences of a player and show that its fixed point is a Equilibrium. $\square$

Nash's Theorem guarantees that an equilibrium exists in a broad class of games, but none of the application analyzed in our previous sections are members of this class. This shows that the hypothesis of Nash's equilibrium are sufficient but not necessary conditions for equilibrium to exist- there are many games that do not satisfy the hypothesis of the Theorem but nonetheless have one or more Nash equilibria.

CHAPTER

TWO

DYNAMIC GAMES OF COMPLETE INFORMATION

In this chapter, we focus on dynamic games, and we again restrict our attention to games with complete information (where the player's payoff functions are common knowledge). The central issue in all dynamic games is credibility. As an example of a non-credible threat, consider the following two-move game:

1. First, player 1 chooses between giving player 2 \$1,000 and giving player 2 nothing.
2. Second, player 2 observes player 1's move then chooses whether or not to explode a grenade that will kill both players .

Suppose that player 2 threatens to explode the grenade unless player 1 pays the \$1,000. If player 1 believes the threat, player 1's best response is to pay the entire \$1,000. However, if player 1 doesn't really believe the threat, he should then pay player 2 nothing- his threat **isn't credible**.

2.1 Dynamic Games of Complete and Perfect Information

2.1.1 The Theory of Backwards Induction

The grenade game is a member of the following class of simple games of complete and perfect information:

1. Player 1 chooses an action a_1 from the feasible set A_1.
2. Player 2 observes a_1 and then chooses an action a_2 from the feasible set A_2
3. Payoffs are $u_1(a_1, a_2)$ and $u_2(a_1, a_2)$.

It actually turns out that many real life economic issues fit this discussion. Two such examples are Stackelberg's model of duopoloy and Leontief's model of wages and employment in a unionized firm. Other economic problems can be modeled by allowing for a longer sequence of actions, either by adding more players or allowing players to move more than once. The key features of a dynamic game of complete and perfect information are that:

1. The moves occur in sequence
2. All previous moves are observed before the next move is made
3. The player's payoffs from each feasible combination of moves are common knowledge

We can solve such a game by backwards induction, as follows: When player 2 gets the move at the second stage of the game, he will face the following problem given the action a_1 previously chosen by player 1:

$$max_{a_2 \in A_2} u_2(a_1, a_2)$$

Assume that for each a_1 in A_1, player 2's optimization problem has a unique solution denoted $R_2(a_1)$.This is player 2's reaction to player 1's reaction. Since player 1 can solve player 2's problem as well as 2 can, player 1 should anticipate player 2's reaction to each action $a_1 \in A_1$ that 1 might take, so 1's problem at the first stage amounts to:

$$max_{a_1 \in A_1} u_1(a_1, R_2(a_1))$$

Assume that this optimization problem for player 1 also has a unique solution, denoted a_1^*. We call $(a_1^*, R_2(a_1^*))$ the **backwards-induction outcome** of this game. The backwards-induction outcome does not involve non-credible threats: player 1 anticipates that player 2 will respond optimally to any action that 1 might choose by player $R(a_1)$; player 1 gives no credence to threats by player 2 to respond in ways that will not be in 2's self-interest when the second stage arrives.

Recall that we used normal-form representation to study static games of complete information, and we focused on the notion of Nash equilibrium as a solution concept for these games. Here, we have made no mention of either normal-form representation or Nash equilibrium. Instead, we have given an intuitive description of a game and have defined the backwards-induction outcome as the solution to that game. We can conclude this section by exploring the rationality assumption inherent in backwards-induction arguments. Consider the following three-move game:

1. Player 1 chooses L or R, where L ends the game with payoffs of 2 to player 1 and 0 to player 2
2. Player 2 observes 1's choice. If 1 chose R then 2 choses L' or R', where L' ends the game with payoffs of 1 to both players
3. Player 1 observes 2's choice, and if the earlier choices were R, R' then 1 chooses L'' or R'', both of which end the game L'' with payoffs of 3 to player 1 and 0 to player 2 and R'' with analogous payoffs of 0 and 2.

The structure of this game follows the form of the following tree:

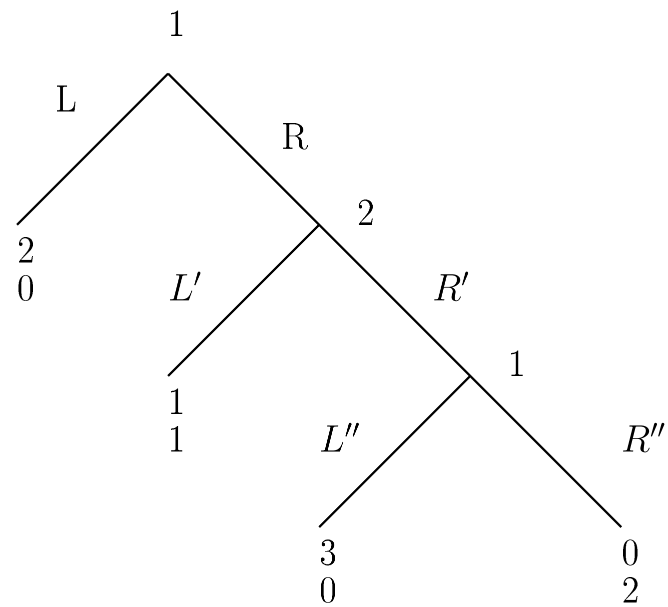

When trying to compute the backwards-induction outcome of this game, we begin at the third stage. Here player 1 faces a choice between a payoff of 3 from L'' and a payoff of 0 from R'', so clearly L'' is optimal. Thus at the second stage, player 2 would anticipate that if the game reaches the third stage then player 1 will play L'', which would yield a payoff of 0 for player 2. The second-stage choice for player 2 therefore is between a payoff of 1 from L' is optimal. Thus, at the first stage, player 1 anticipates that if the game reaches the second stage, player 2 would play L', which would yield a payoff of 1 for player 1. The first-stage choice for player 1 therefore is between a payoff of 2 from L and a payoff of 1 from R, so clearly L is optimal.

This argument overall establishes that the backwards-induction outcome of this game is for player 1 to choose L in the first stage, ending the game. Even though backwards induction predicts that the game will end in the first stage, an important part of the argument concerns what would happen if the game did not end in the first stage. In the second stage, when player 2 anticipates that the game will reach the third stage, then 1 will play L'', assuming that 1 is rational. This assumption may seem inconsistent with the fact that 2 gets to move in the second stage only if 1 deviates from the backwards-induction outcome of the game. So it may seem that if 1 plays R in the first stage then 2 can't assume that in the second stage that 1 is rational, but this is not the case: if 1 plays R in the first stage then ti cannot be common knowledge that both players are rational, but there remain reasons for 1 to have chosen R that do not contradict 2's assumption that 1 is rational. One possibility is that it is common knowledge that player 1 is rational but not that player 2 is rational: if 1 thinks 2 might not be rational, then 1 might choose R in the first stage hoping that 2 will play R' in the second stage, giving 1 the chance to play L'' in stage three.

For some games, it may be more reasonable to assume that 1 played R because 1 is indeed irrational. In these games, backwards induction loses much of its appeal as a prediction of play, just as Nash equilibrium does in games where game theory does not provide a unique solution and no convention will develop.

2.1.2 Stackelberg Model of Duopoly

In 1934, Stackelberg proposed a dynamic model of duopoly in which a dominant firm moves first an a follower firm moves second. The timing of the game is as follows:

1. Firm 1 chooses a quantity $q_1 \geq 0$
2. Firm 2 observes q_1 and then chooses a quantity $q_2 \geq 0$
3. The payoff to firm i is given by the profit function

$$\pi_i(q_i, q_j) = q_i\,[P(Q) - c]$$

   Where $P(Q) = a - Q$ is the market clearing price when the aggregate quantity on the market is $Q = q_1 + q_2$ and c is the constant marginal cost of production.

To solve for the backwards-induction outcome of this game, we have to compute firm 2's reaction to an arbitrary quantity by firm 1. In other words, $R_2(q_1)$ solves:

$$max_{q_2 \geq 0} \pi_2(q_1, q_2) = max_{q_2 \geq 0} q_2[a - q_1 - q_2 - c]$$

which yields

$$R_2(q_1) = \frac{a - q_1 - c}{2}$$

provided that $q_1 < a - c$. Interestingly, the same equation for $R_2(q_2)$ appeared in our analysis of the simultaneous move Cournot game in a previous section. Since firm 2 can solve firm 2's problem as well as firm 2 can solve it, firm 2 should anticipate that the quantity choice q_1 will be met with the reaction $R_2(q_1)$. So, firm 2's problem amounts to:

$$max_{q_1 \geq 0} \pi_1(q_1, R_2(q_1)) = max_{q_1 \geq 0} q_1[a - q_1 - R_2(q_1) - c] = max_{q_1 \geq 0} q_1 \frac{a - q_1 - c}{2}$$

Which yields the following:

$$q_1^* = \frac{a - c}{2} \quad R_2(q_1^*) = \frac{a - c}{2}$$

as the backwards-induction outcome of the Stackelberg duopoly game.

Notice that firm 2 actually does worse in the Stackelberg model than it did in the Cournot game illustrates an important difference between single and multi-person decision problems. In single-person decision theory, having more information can never make the decision maker worse of. In game theory, having more information actually **can** make the player worse off.

In the Stackelberg game, the information in question is firm 1's quantity: firm 2 knows q_1 and firm 1 knows that firm 2 knows q_1. To see the effect this information has, consider the modified sequential move game in which firm 1 chooses q_1, after which firm 2 chooses q_2 but does so without observing q_1. If firm 2 believes that firm 1 has chosen its Stackelberg quantity

$q_1^* = (a-c)/2$, then firm 2's best response is again $R_2(q_1^*) = (a-c)/4$. But if firm 2 anticipates that firm 2 will hold this belief and so choose this quantity, then firm 1 prefers to choose its best response to $(a-c)/4$, namely, $3(a-c)/8$, rather than its Stackelberg quantity $(a-c)/2$. So, firm 2 shouldn't believe that firm 1 has chosen its Stackelberg quantity. Rather, the unique Nash equilibrium of this modified sequential move game is form both firms to choose the quantity $(a-c)/3$, precisely the Nash equilibrium of the Cournot game, where the firms move simultaneously. Thus, having firm 1 knows that firm 2 knows q_1 hurts firm 2.

2.1.3 Wages and Employment in a Unionized Firm

In 1946, Leontief's model of the relationship between a firm and a monopoly union, the union has exclusive control over wages, but the firm has exclusive control over employment. The union's utility function is $U(w, L)$, where w is the wage the union demands from the firm and L is employment. Assume that $U(w, L)$ increases in both w and L The firm's profit function is $\pi(w, L) = R(L) - wL$, where $R(L)$ is the revenue the firm can earn if it employs L workers. Assume that $R(L)$ is increasing and concave.

We have the following timing of this game:

1. The union makes a wage demand
2. The firm observes and accepts w, then chooses employment, L
3. Payoffs are $U(w, L)$ and $\pi(w, L)$.

First we can characterize the firm's best response in stage 2 $L^*(w)$ to an arbitrary wage demand by the union in stage 1, w. Given w, the firm chooses $L^*(w)$ to satisfy the following:

$$max_{L\geq 0}\pi(w, L) = max_{L\geq 0}R(L) - wL$$

The first-order condition for which is:

$$R'(L) - w = 0$$

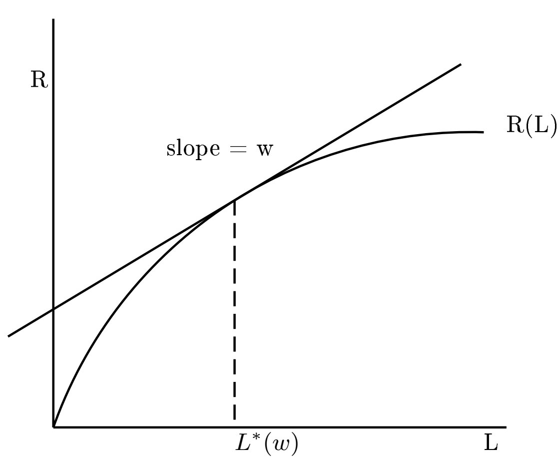

To guarantee that the first order condition $R'(L) - w = 0$ has a solution, assume that $R'(0) = \infty$ and that $R'(\infty) = 0$. Notice from our diagram that $L^*(w)$ cuts each of the firm's isoprofit curves at its maximum. Holding L fixed, the firm does better when w is lower, so lower isoprofit curves represent higher profit levels. Holding L fixed, the union does better when w is higher, so higher indifference curves represent higher utility levels for the union.

We now turn to the union's problem at stage (1). Since the union could solve the firm's second stage problem just as well as the firm can, the union should anticipate that the firm's reaction to the wage demand w will be to choose the employment level $L^*(w)$. So, the union's problem at the first stage amounts to

$$max_{w\geq 0}U(w, :^* (w))$$

in terms of the indifference curves plotted below:

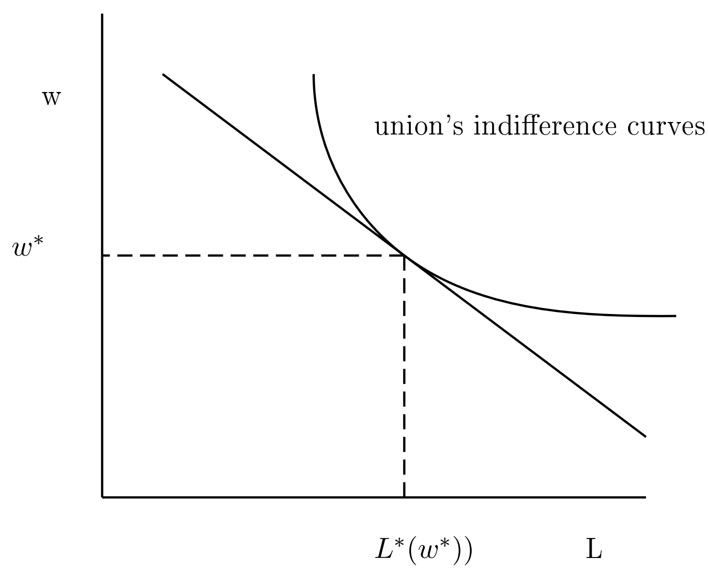

the union would like to choose the wage demand w that yields the outcome $(w, L^*(w))$ that is on the highest possible indifference curve. The solution to the union's problem is w^*, the wage demand such that the union's indifference curve through the point $(w^*, L^*(w^*))$ is tangent to $L^*(w)$ at that point.

It is straightforward to see that $(w^*, L^*(w^*))$ is inefficient: both the union's utility and the firm's profit would be increased if w and L were in the shaded portion in the figure below.

This inefficiency makes it puzzling that in practice firms seem to retain exclusive control over employment. One answer to this puzzle, based on the fact that the union and the firm negotiate repeatedly over time is proposed by Espinosa and Rhee.

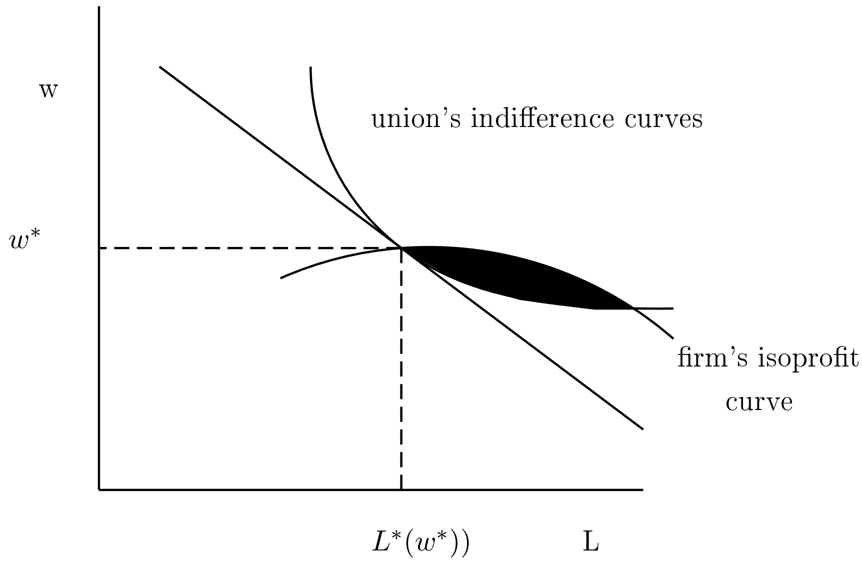

2.2 Sequential Bargaining

We now move into a three-period bargaining model. Players 1 and 2 are bargaining over a dollar. They alternate in making offers, first player 1 makes a proposal that player 2 can accept or reject; if 2 rejects then 2 makes a proposal that 1 can accept or reject, and so on. Once an offer has been rejected, it ceases to be binding and is irrelevant to the subsequent play of the game. Each offer takes one period, and the players are impatient: they discount payoffs received in later periods by the factor δ per period, where $0 < \delta < 1$. We have the following formal wording:

1. At the beginning of the first period, player 1 proposes to take a share $s_!$ of the dollar, leaving $1 - s_1$ for player 2
2. Player 2 either accepts the offer (so payoffs $s_!$ and $1 - s_1$ go to player 1 and player 2 respectively) or rejects the offer, in which case play continues.
3. At the beginning of the second period, player 2 proposes that player 1 take a share s_2 of the dollar, leaving $1 - s_2$ for player 2.
4. Player 1 either accepts the offer (so payoffs $1 - s_2$ and s_2 go to player 1 and player 2 respectively) or rejects the offer, in which case play continues. to the third period
5. At the beginning of the third period, player 1 receives a share s of the dollar, leaving $1 - s$ for player 2, where $0 < s < 1$.

To solve for the backwards-induction outcome of this three-period game, we first compute player 2's optimal offer if the second period is reached. Player 1 can receive s in the third period by rejective player 2's offer of s_2 this period, but the value this period of receiving s next period is only δs. So, player 1 will accept s_2 if and only if $s_2 \geq \delta s$. Player 2's second-period decision problem therefore amounts to choosing between receiving $1 - \delta s$ this period and receiving $1 - s$ next period. The discounted value of the latter option is $\delta(1 - s)$, which is less than the $1-\delta s$ available from the former option, so player 2's optimal second-period offer is $s_2^* = \delta s$. So, if play reaches the second period, player 2 will offer s_2^* and player 1 will accept. Since player 1 can solve player 2's second-period problem as well as player 2 can, player 1 knows that player 2 can receive $1 - s_2^*$ in the second period by rejecting player 1's offer of s_1 this period, but the value this period of receiving $1 - s_2^*$ next period is only $\delta(1 - s_2^*)$. So, player 2 will accept $1-s_1$ if and only if $1-s_1 \geq \delta(1-s_2^*)$ or $s_1 \leq 1-\delta(1-s_2^*)$. Player 1's first period decision problem therefore amounts to choosing between receiving $1 - \delta(1 - s_2^*)$ this period and receiving s_2^* next period. The discounted value of the latter option is $\delta s_2^* = \delta^2 s$, which is less than $1 - \delta(1 - \delta s_2^*) = 1 - \delta(1 - \delta s)$. So, in the backwards induction outcome, player 1 offers the settlement $(s_1^*, 1 - s_1^*)$ to player 2, who accepts.

No formal backward-induction argument was made for this game, but we do so now: Suppose that there is a backwards-induction outcome of the game as a whole in which player 1 and 2 receive the payoffs s and $1 - s$ respectively. We use these payoffs in the game beginning in the third period, then work backwards to the first period. In this new back-wards induction outcome, player 1 offers the settlement $(f(s), 1 - f(s))$ in the first period and player 2 will accept where $f(s) = 1 - \delta(1 - \delta s)$.

2.3 Two-Stage Games of Complete But Imperfect Information

2.3.1 Theory of Subgame Perfection

We now try to expand the class of games analyzed in the previous section. Just like with dynamic games of complete and perfect information, we assume that play proceeds in a sequence of stages, with the moves in all previous stages observed before the next stage begins. Unlike in the games analyzed in the previous section, we now allow there to be simultaneous moves within each stage. We can analyze the following simple game, called a **two-stage game of complete but imperfect information**:

1. Players 1 and 2 simultaneously choose actions a_1, a_2 respectively from the sets A_1, A_2.
2. Players 3 and 4 observe the outcome of the first stage (a_1, a_2) and then simultaneously choose actions a_3, a_4 from the feasible sets A_3, A_4.
3. Payoffs are $u_i(a_1, a_2, a_3, a_4)$ for $i = 1, 2, 3, 4$.

Some examples of real world economic examples include things like bank runs, tariffs and imperfect international competition. We solve a game from this class by using an approach that models the spirit of backwards induction, but this time the first step in working backwards from the end of the game involves solving a real game rather than solving a single-person optimization problem as in the previous section. To keep these things simple, we assume that for each feasible outcome of the first stage game, (a_1, a_2), the unique second stage game that remains between players 3 and 4 has a unique Nash equilibrium, denoted by $(a_3^*(a_1, a_2), a_4^*(a_1, a_2))$.

If players 1 and 2 anticipate that the second stage behavior of players 3 and 4 will be given by $(a_3^*(a_1,a_2), a_4^*(a_1,a_2))$, then the first-stage interaction between players 1 and 2 amounts to the following simultaneous move game:

1. Players 1 and 2 simultaneously choose actions a_1, a_2 from feasible sets A_1, A_2, respectively.

2. Payoffs are $u_i(a_1, a)2, a_3^*(a_1,a_2), a_4^*(a_1,a_2))$ for $i = 1, 2$.

Supposing that (a_1^*, a_2^*) is the unique Nash equilibrium of this simultaneous move game, then we will call $(a_1^*, a_2^*, a_3^*(a_1,a_2), a_4^*(a_1,a_2))$ the **subgame-perfect outcome** of this two-stage game. This outcome is the natural analog of the backwards-induction outcome in games of complete and perfect information, and the analogy applies to both the attractive and unattractive features o the latter.

2.3.2 Bank Runs

Suppose that two investors have each deposited D with a bank. The bank has invested these deposits in a long term project. If the bank is forced to liquidate its investment before the project matures, a total of $2r$ can be recovered, where $D > r > D/2$. If the bank allows the investment to reach maturity, however, the project will pay out a total of $2R$, where $R > D$. There are two dates at which the investors can make withdraws from the bank: date 1 is before the bank's investment matures; date 2 is after. If both investors make withdrawals at date 1, then each receives r and the game ends. If only one investor makes a withdrawal at date 1 then that investor receives D, the other receives $2r - D$, and the game ends. Finally, if neither investor makes a withdrawal at date 1 then the project matures and the investors make withdrawal decisions at date 2. If both investors make withdrawals at date 2 then each receives R and the game ends. If only one investor makes a withdrawal at date 2, then that investor receives $2R - D$, the other receives D, and the game ends. If neither makes a withdrawal at date 2, then the bank returns R to each investor and the game ends. This can be illustrated with figure 2.1

In analyzing this game, we work backwards. Consider the normal form game at date 2. Since $R > D$, "withdraw" strictly dominates "don't withdraw" so there is a unique Nash equilibrium in this game: both investors withdraw, leading to a payoff of (R, R). Since there is no discounting, we can simply substitute this payoff into the normal form game at date 1, as in figure 2.2.

	withdraw	don't
withdraw	r,r	D, 2r-D
don't	2r-D, D	next stage

	withdraw	don't
withdraw	R,R	D, 2R-D
don't	D, 2R-D	R,R

Figure 2.1: Bank Runs

	withdraw	don't
withdraw	r,r	D, 2R-D
don't	D, 2r-D	R,R

Figure 2.2:

Since $r < D$, this one-period version of the two-period game has two pure-strategy Nash equilibria:

1. Both investors withdraw, leading to a payoff of (r, r)
2. Both investors do not withdraw, leading to a payoff of (R, R).

So, the original two-period bank runs game has two subgame-perfect outcomes: both withdraw at date 1, or both investors withdraw at date 2.

2.4 Tariffs and Imperfect International Competition

We now look at an example from international economics. Consider two identical countries, denoted by $i = 1, 2$. Each country has a government that chooses a tariff rate, a firm that produces output for both home consumption and export, and consumers who buy on the home market from either the home firm or the foreign firm. If the total quality on the market in country i is Q_i, then the market clearing price is $P_i(Q_i) = a - Q_i$. The firm in country i produces h_i for home consumption and e_i for export. So, $Q_i = h_i + e_j$. The firms have a constant marginal cost, and no fixed costs- so the total cost of production for firm i is $C_i(h_i, e_i) = c(h_i + e_i)$. The firms also incur tariff costs on exports: if firm i exports e_i to country j when government j has set the tariff rate t_j, then firm i must pay $t_j e_i$ to government j.

First the government simultaneously choose tariff rates t_1, t_2 respectively. Then, the firms observe the tariff rates and simultaneously choose quantities for home consumption and for export, (h_1, e_1) and (h_2, e_2). Then, the payoffs are profit to firm i and total welfare to government i, where total welfare to country i is the sum of the consumers' surplus enjoyed by the consumers in country i, the profit earned by firm i, and the tariff revenue collected by government i from firm j:

$$\pi_i(t_i, t_j, h_i, e_i, h_j, e_j) = [a - (h_i + e_j)]h_i + [a - (e_i + h_j)]e_i - c(h_i + e_i) - t_j e_i$$

$$W_i(t_i, t_j, h_i, e_i, h_j, e_j) = \frac{1}{2}Q_i^2 + \pi_i(t_i, t_j, h_i, e_i, h_j, e_j) = +t_i e_j$$

Suppose the governments choose tariffs t_1 and t_2. If $(h_1^*, e_1^*, h_2^*, e_2^*)$ is a Nash equilibrium in the remaining game between firms 1 and 2, then for all i, (h_i^*, e_i^*) must solve:

$$max_{h_i, e_i \geq 0} \pi_i(t_i, t_j, h_i, e_i, h_j^*, e_j^*)$$

Since $\pi_i(t_i, t_j, h_i, e_i, h_j^*, e_j^*)$ can be written as the sum of firm i's profits on market i and firm i's profits on market j, firm i's two-market optimization problem simplifies into a pair of problems, one for each market: h_i^* must solve:

$$max_{h \geq 0} h_i[a - (h_i + e_j^*) - c]$$

and e_i^* must solve

$$max_{e_i \geq 0} e_i[a - (c_i + h_j^*) - c] - t_j e_i$$

Assuming that $e_j^* \leq a - c$, we have

$$h_i^* = \frac{1}{2}(a - e_j^* - c)$$

and assuming that $h_j^* \leq a - c - t_j$, we have

$$e_i^* = \frac{1}{2}(a - h_j^* - c - t_j)$$

Both of the best-response functions must hold for each $i = 1, 2$. Fortunately, these equations simplify to the following:

$$h_i^* = \frac{a - c + t_i}{3} \qquad e_i^* = \frac{a - c - 2t_j}{3}$$

Recall that the equilibrium quantity chosen by both firms in the Cournot game is $(a - c)/3$, but that this result was derived under the assumption of symmetric marginal costs. In the equilibrium described above, the government's tariff choices make marginal costs asymmetric marginal costs. Having solved the second stage game that remain between the two firms after the governments choose tariff rates, we can now represent the first-stage interactions between thet wo governments as the following simultaneous-move game. We can now solve for the Nash equilibrium of this game between the governments.

To simplify things, let $W_i^*(t_i, t_j)$ denote $W_i(t_i, t_j, h_i^*, e_i^*, h_j^*, e_j^*)$. If (t_1^*, t_2^*) is a Nash equilibrium of this game between the governments then, for each i, t_i^* must solve

$$max_{t_i \geq 0} W_i^*(t_i, t_j^*)$$

But, $W_i^*(t_i, t_j^*)$ equals:

$$\frac{(2(a - c) - t_i)^2}{18} + \frac{(a - c + t_i)^2}{9} + \frac{(a - c - 2t_j^*)^2}{9} + \frac{t_i(a - c - 2t)i}{3}$$

So,

$$t_i^* = \frac{a - c}{3}$$

for each i, independent of t_j^*. Substituting $t_i = t_j^* = (a - c)/3$, we have

$$h_i^* = \frac{4(a - c)}{9} \qquad e_i^* = \frac{a - c}{9}$$

as the firm's quantity choices in the second game. So, the subgame perfect outcome of this tariff game is

$$t_1^* = t_2^* = (a - c)/3 \quad h_1^* = h_2^* = 4(a - c)/9 \quad e_1^* = e_2^* = (a - c)/9$$

2.4.1 Tournaments

Suppose you have two workers and their boss. Worker i produces output $y_i = e_i + \epsilon_i$ where e_i is effort and ϵ_i is noise. Production proceeds as follows. First the workers simultaneously choose nonnegative effort levels. Second, the noise terms ϵ_1, ϵ_2 are independently drawn from a density $f(\epsilon)$ with zero mean. Third, the worker's outputs are observed but their effort choices are not. The workers' wages therefore can depend on their outputs but not on their efforts. To induce effort, the boss has them compete in a tournament. The wage earned by the winner is w_H, the wage earned by the loser is w_L. The payoff to a worker from earning wage w and expending effort e is $u(w,e) = w - g(e)$, where the disutility of effort $g(e)$ is increasing and convex. The payoff to the boss is $y_1 + y_2 - w_H - w_L$.

Making this into a more formal game, the boss is player 1, the workers are player 3 and 4 (there is no player 2) who observe the wages chosen in the first stage and then simultaneously choose actions a_3, a_4, namely the effort choices e_1, e_2. Finally, the player's payoffs are as given early. Since outputs are functions not only of the players actions but also of the noise terms ϵ_1, ϵ_2, we work with the players expected payoffs.

Suppose the boss has chosen the wages w_H, w_L. If the effort pair (e_1^*, e_2^*) is to be a Nash equilibrium of he remaining game between the workers then, for each i, e_i^* must maximize worker i's expected wage, net of the disutility of effort: e_i^* must solve

$$max_{e_i \geq 0} w_H Prob\{y_i(e_i) > y_j(e_j^*)\} + w_L Prob\{y_i(e_i) \leq y_j(e_j^*)\} - g(e_i)$$

$$= (w_H - w_L) Prob\{y_i(e_i) > y_j(e_j^*)\} + w_L - g(e_i)$$

The first order condition is then

$$(w_H - w_L) \frac{\partial Prob\{y_i(e_i) > y_j(e_j^*)\}}{\partial e_i} = e'(e_i)$$

By Baye's rule,

$$Prob\{y_i(e_i) > y_j(e_j^*)\} = Prob\{\epsilon_i > e_j^* + \epsilon_j - e_i\}$$

$$= \int_{\epsilon_j} Prob\{\epsilon_i > e_j^* + \epsilon_j - e_i | \epsilon_j\} f(\epsilon_j) d\epsilon)_j$$

$$= \int_{\epsilon_j} [1 - F(e_j^* - e_i + \epsilon_j)] f(\epsilon_j) d\epsilon_j$$

So the first order condition really becomes

$$(w_H - w_L) \int_{\epsilon_j} f(e_j^* - e_i + \epsilon_j) f(\epsilon_j) d\epsilon_j = g'(e_i)$$

In a symmetric Nash equilibrium, if ϵ is distributed normal with variance σ^2, we have

$$(w_H - w_L) \int_{\epsilon_j} f(\epsilon_j)^2 d\epsilon_j = g'(e^*) = \frac{1}{2\sigma\sqrt{\pi}}$$

Since in the symmetric Nash equilibrium each worker wins the tournament with probability one-half, if the boss intends to induce the workers to participate in the tournament then she must choose wages that satisfy

$$\frac{1}{2}w_H + \frac{1}{2}w_L - g(e^*) \geq U_a$$

Assuming that U_a is low enough that the boss wants to induce the workers to participate in the tournament, she chooses wages to maximize expected profit. At the optimum, the above equation holds with equality:

$$w_L = 2U_a + 2g(e^*) - w_H$$

So, expected profits then becomes $2e^* - 2U_a - 2g(e^*)$, so the boss wishes to choose wages such that the induced effort e^* maximizes $e^* - g(e^*)$. The optimal induced effort therefore satisfies the first-order condition $g'(e^*) = 1$. Substituting this tells us that the optimal prize $w_H - w_L$ solves

$$(w_H - w_L)\int_{\epsilon_j} f(\epsilon_J)^2 d\epsilon_j = 1$$

2.5 Repeated Games

We now look at games where we analyze whether threats and promises about future behavior can influence current behavior in repeated relationships

2.5.1 Two-Stage Repeated Games

Remember the Prisoner's dilemma given in normal form. We have the two-stage example, in the figure on the next page.

Player 2

Player 1		L_2	R_2
	L_1	1,1	5,0
	R_1	0,5	4,4

Player 2

Player 1		L_2	R_2
	L_1	2,2	6,1
	R_1	1,6	5,5

Suppose that the payoff for the entire game is simply the sum of the payoffs from the two stages. This game will be called the two-stage Prisoner's Dilemma. This is identical to what we did before, but here players 3 and 4 are identical to players 1 and 2, the action spaces A_3 and A_4 are identical to A_1 and A_2, and the payoffs $u_i(a_1, a_2, a_3, a_4)$ are simply the the sum of the payoff from the first stage outcome $(a_{1,2})$ and the payoff from the second-stage outcome (a_3, a_4). Furthermore, the second stage game that remains between players 3 and 4 has a unique Nash equilibrium, denoted by $(a_3^*(a_1, a_2), a_4^*(a_1, a_2))$. In fact, the two stage Prisoners' Dilemma satisfies this assumption in the following way: previously, we allowed for the possibility that the Nash equilibrium quantity choices in the second stage depend on the first stage outcome. In the two-stage Prisoners' Dilemma, however, the unique equilibrium of the second-stage game is (L_1, L_2), regardless of the first-stage outcome.

In computing the subgame-perfect outcome of such a game, we analyze the first stage of the two-stage Prisoners' Dilemma by taking into account that the outcome of the game remaining in the second stage will be the Nash equilibrium of that remaining game. So, the players' first stage interaction in th two stage Dilemma amounts to the one-shot game above, where the payoff pair $(1, 1)$ for the second stage has been added to each first-stage payoff pair. So, the unique subgame-perfect outcome of the two-stage Prisoners' Dilemma is $(L_{1,2})$ in the first stage, followed by (L_1, L_2) in the second stage. We end up with the following definition:

Definition

Given a stage game G, let $G(T)$ denoted the **finitely repeated game** in which G is played

T times, with the outcomes of all preceding plays observed before the next play begins. The payoffs for $G(T)$ are simply the sum of the payoffs from the T stage games.

Definition

If the stage game G has a unique Nash equilibrium then, for any finite T, the repeated game $G(T)$ has a unique subgame-perfect outcome: the Nash equilibrium of G is played in every stage.

We can now return to the two period case, but consider the possibility that the stage game G has multiple Nash equilibria, as in the above figure. The strategies labeled L_i and M_i mimic the Prisoner's Dilemma, but the strategies labeled R_i have been added to the game so that there are now two pure-strategy Nash equilibria: (L_1, L_2) as in the Prisoners' Dilemma, and now also (R_1, R_2). We end up with the following diagram:

	L_2	M_2	R_2
L_1	1,1	5,0	0,0
M_1	0,5	4,4	0,0
R_1	0,0	0,0	3,3

Supposing that the stage game above is played twice, with the first-stage outcome observed before the second stage begins, we can show that there is a subgame-perfect outcome of this repeated game in which the strategy pair (M_1, M_2) is played in the first stage. We again assume that in the first stage the players anticipate that the second-stage outcome will be a Nash equilibrium. Eventually, we can see that the players' first stage interaction then amounts to the one-shot game where $(3, 3)$ has been added to the (M_1, M_2)- cell and $(1, 1)$ has been added to the eight other cells.

There are then three pure-strategy Nash equilibria in the game above: (L_1, L_2), (M_1, M_2) and (R_1, R_2).

	L_2	M_2	R_2
L_1	2,2	6,1	1,1
M_1	1,6	7,7	1,1
R_1	1,1	1,1	4,4

So, in the figure above, the Nash equilibrium (R_1, R_2) corresponds to the subgame perfect outcome $((R_1, R_2,), (L_1, L_2))$ in the repeated game. These two subgame perfect outcomes simply concatenate Nash equilibrium outcomes from the stage game, but the third Nash equilibrium yields a quantitatively different result.

The main point to get from this example is that credible threats or promises about future behavior can influence current behavior. A second point, is that subgame perfection may not embody a strong enough definition of credibility. In deriving the subgame perfect outcome $((M_1, M_2), (R_1, R_2))$, for example, we assumed that the players anticipate that (R_1, R_2) will be the second-stage outcome if the first-stage outcome is (M_1, M_2) and that (L_1, L_2) will be the second stage outcome if any of the other eight first-stage outcomes occur.

2.6 The Theory of Infinitely Repeated Games

As in the finite-horizon case, the main theme in infinitely repeated game is that credible threats and/or promises about future behavior can influence current behavior.

We begin by studying the infinitely repeated Prisoners' Dilemma. We consider the class of infinitely repeated games analogous to the class of finitely repeated games defined in the previous section: a static game of complete information, G, is repeated infinitely, with the outcomes of all previous stages observed before the current stage begins. Suppose that the Prisoners Dilemma below is to be repeated infinitely and that, for each t, the outcomes of the $t - 1$preceding plays of the stage game are observed before the t^{th} stage begins: Simply summing the payoffs from this infinite sequence of stage games doesn't provide a useful measure of a player's payoff in the infinitely repeated game. Receiving a payoff of 4 in each period is better than receiving a payoff of 1 for each period for example, bu the sum of the payoffs is infinity in both cases. Recall that the discount factor $\delta = 1/(1 + r)$ is the value

		Player 2	
		L_2	R_2
Player 1	L_1	1,1	5,0
	R_1	0,5	4,4

today of a dollar to be received one stage later, where r is the interest rate per stage. Given a discount factor and a players payoffs sequence of stage jumps, we can compute the present value of the payoffs.

Definition

Given the discount factor δ, the **present value** of the infinite sequence of payoffs $\pi_1, \pi_2, ...$ is:

$$\pi_1 + \delta\pi_2 + \delta^2\pi_3 + ... = \sum_{t=1}^{\infty} \delta^{t-1}\pi_t$$

We can also use the factor δ to reinterpret an infinitely repeated game as a repeated game that ends after a random number of repetition.

Go back to considering the infinitely repeated Prisoner's Dilemma in which each player's discount factor is δ and each player's payoff in the repeated game is the present value of the player's payoffs from the stage games. Suppose that player i begins the infinitely repeated game by co-operating and then cooperates in each subsequent stage game if and only if both players have cooperated in every previous stage. Formally, player i's strategy is to play R_i in the first stage, and in the t^{th} stage, if the outcome of all $t-1$ preceding stages has been (R_1, R_2) then play R_i; otherwise play L_i.

This is an example of a *trigger strategy*, because player i cooperates right up until someone fails to cooperate, which triggers a witch to noncooperation ever after. If both players adopt this trigger strategy, then the outcome of the infinitely repeated game will be (R_1, R_2) in every stage. We first argue that if δ is close enough to one then it is a Nash equilibrium of the infinitely repeated game for both players to adopt this strategy. We then argue that such a Nash equilibrium is subgame-perfect, in a sense that we can now make more precise.

To show that we have a Nash equilibrium of the infinitely repeated game for both players to adopt the trigger strategy, we will assume that player i has adopted the trigger strategy and then show that provided δ is close to one, it is a best response for player j to to also adopt that strategy. Since player i will play L_i forever once one stage's outcome differs from

(R_1, R_2), player j's best response is indeed to play L_j forever once one stage differs from (R_1, R_2). Thus, player j s best response for each and every stage such that all preceding outcomes have been (R_1, R_2) will be to play R_2. We end up with the following present value of this sequence of payoffs:

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

However, if playing R_j is optimal, then

$$V = 4 + \delta V$$

where V denotes the present value of the infinite sequence of payoffs player j receives from making this choice optimally. So, R_j is optimal if and only if

$$\frac{4}{1-\delta} \geq 5 + \frac{\delta}{1-\delta}$$

Or, $\delta \geq 1/4$.

We now want to show that such a Nash equilibrium is sub-game perfect. To do this, we define a strategy in a repeated game, a subgame in a repeated game, and a subgame-perfect Nash equilibrium in a repeated game. In order to illustrate these concepts with simple examples, we will define them for both finitely and infinitely repeated games. In the previous section we said that the finitely repeated game $G(T)$ based on the stage game $G = \} A_1, ..., A_n; u_1, ..., u_n \}$. We now have the following:

Definition

Given a stage game G, let $G(\infty, \delta)$ denote the **infinitely repeated game** in which G is repeated forever and the players share the discount factor δ. For each t, the outcomes of the $t-1$ preceding plays of the stage game are observed before the t^{th} stage begins. Each player's payoff in $G(\infty, \delta)$ is the present value of the player's payoffs from the infinite sequence of stage games.

In any game, a player's strategy is a complete plan of action, it specifies a feasible action fro the player in which the players might be called upon to act. In a static game of complete information, for example, a strategy is simply an action. In a dynamic game, however, a strategy is more complicated.

Consider the two-stage Prisoners' Dilemma again. Each player acts twice, so one might think that a strategy is simply a pair of instructions (b, c) where b is the first stage action and c is the second stage action. But, there are four possible first-stage outcomes- $(L_1, L_2), (L_1, R_2), (R_1, L_2)$, and (R_1, R_2)- and these represent four separate contingencies in which each player might be called upon to act. So, each players strategy consists of five instructions, denoted (w, v, x, y, z) where v is the first stage action, and w, x, y, z are the second-stage actions to be taken following the first-stage outcomes $(L_1, L_2), (L_1, R_2), (R_1, L_2)$ and (R_1, R_2) respectively. Using this notation, we can express things like 'play b in the first stage and c in the second, unless so-and-so'. In the finitely repeated game $G(T)$ or the infinitely repeated game $G(\infty, \delta)$, the *history of play throughout stage* t is the record of the players' choices in stages 1 through t. We end up with the following definition:

Definition

In the finitely repeated game $G(T)$ or the infinitely repeated game $G(\infty, \delta)$, a players **strategy** specifies the action the player will take in each stage, for each possible history of play throughout the previous stage.

2.6.1 Subgames

We now talk quickly about subgames. A subgame is a piece of a game, the piece that remains to be played beginning at any point at which the complete history of the game thus far is common knowledge among the players. In the two-stage Prisoners' Dilemma, for example, there are four subgames, corresponding to the second-stage games that follow the four possible first-stage outcomes.

Definition

In the finitely repeated game $G(T)$, a **subgame** beginning at stage $t+1$ is the repeated game in which G is played $T-t$ times, denoted $G(T-t)$. There are many subgames that begin at stage $t+1$, one for each of the possible histories of play through stage t. In the infinitely repeated game $G(\infty, \delta)$, each subgame beginning at stage $t+1$ is identical to the original game $G(\infty, \delta)$. As in the finite-horizon case, there are as many subgames beginning at stage $t+1$ of $G(\infty, \delta)$ as there are possible histories of play through stage t.

We are now ready for the definition of subgame-perfect Nash equilibrium, which depends on the definition of Nash equilibrium.

Definition

A Nash equilibrium is **subgame-perfect** if the players' strategies constitute a Nash equilibrium in every subgame

---

One can think of Subgame-perfect Nash equilibrium as being a refinement of Nash equilibrium. To show that the trigger-strategy Nash equilibrium in the finitely repeated Prisoner's Dilemma is subgame-effect, we must show that the trigger strategies constitute a Nash equilibrium on every subgame of that infinitely repeated game. Recall that every subgame of an infinitely repeated game is identical to the game as a whole. In the trigger-strategy Nash equilibrium of the infinitely repeated Prisoner's Dilemma, these subgames can be grouped into two classes

1. Subgames in which all the outcomes of earlier stages have been (R_1, R_2)
2. Subgames in which the outcome of at least one earlier stage differs from (R_1, R_2).

If the players adopt the trigger strategy for the game as a whole, then (i) the player's strategies in a subgame in the first class are again the trigger strategy, which we have shown to be a Nash equilibrium of the game as a whole, and (ii) the players' strategies in a subgame in the second class are simply to repeat the stage-game equilibrium (L_1, L_2) forever, which in itself is also a Nash equilibrium of the game as a whole.

Applying analogous arguments in the infinitely repeated game $G(\infty, \delta)$, we end up with Friedman's Theorem. First,we call the payoffs $(x_1, ..., x_n)$ *feasible* in the stage game G if they are a convex combination of the pure strategy payoffs of G. The set of feasible payoffs for the Prisoner's Dilemma below is the shaded region;

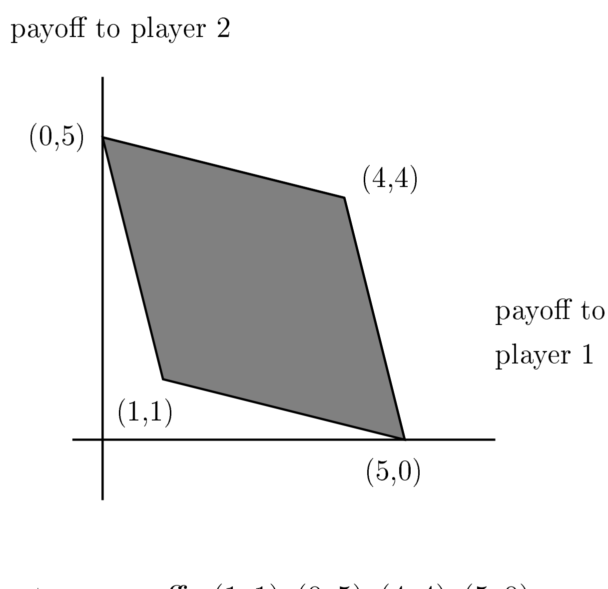

Notice that the pure-strategy payoffs $(1, 1), (0, 5), (4, 4), (5, 0)$ are all feasible. There exist other feasible payoffs, which can be modeled as ordered pairs that fit in the interior of the above figure. To get a weighted average of pure-strategy payoffs, the players could use a public randomizing device, by playing (L_1, R_2) or (R_1, L_2) depending on the flip of a coin, for example, they achieve the expected payoffs $(2.5, 2.5)$.

Definition

Given a discount factor δ, the **average payoff** of the infinite sequence of payoffs $\pi_1, \pi_2, ...$ is

$$(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\pi_t$$

The advantage of being able to talk about average payoff vs. the present value is that the former is directly comparable to the payoffs from the stage game. We end up wit the following main result:

Theorem 4. *(Friedman 1971): Let G be a finite, static game of complete information. Let $(e_1, ..., e_n)$ denote the payoffs from a Nash equilibrium of G and let $(x_1, ..., x_n)$ denote any other feasible payoffs from G. If $x_i > e_i$ for every player i and if δ is sufficiently close to one, then there exists a subgame-perfect Nash equilibrium of the infinitely repeated game $G(\infty, \delta)$ that achieved $(x_1, ..., x_n)$ as the average payoff.*

2.6.2 Efficiency Wages

In efficiency-wage models, the output of a firm's work force depends on the wage the firm pays. In the context of developing countries, higher wages could lead to better nutrition. From Shapiro and Stiglitz(1984) developed the following stage game: First, the firm offers the worker a wage w. Second, the worker accepts or rejects the firm's offer. If the workers rejects w, then the worker becomes self-employed at wage w_0. If the worker accepts w, then the worker chooses either to supply effort (which entails disutility e) or to shirk (which entails no disutility). The worker's effort decision is not observed by the firm, but the worker's output is observed by both the firms and the worker. Output can either be higher or low, and for simplicity, we take low output to be zero and so write high output as $y > 0$. Suppose that if the worker supplies effort then output is sure to be high, but that if the worker shirks then output is high with probability p and low with probability $1-p$.

If the firm employs the worker at wage w, then the player's payoffs if the worker supplies effort and output is high are $y - w$ for the firm and $w - e$ for the worker. If the worker shirks, then e becomes 0; then y becomes 0. We assume that $y - e > w_0 > py$, so that it is efficient for the worker to be employed by the firm and to supply effort, and also better than the worker to be self-employed than employed by the firm and shirking.

The subgame-perfect outcome of this stage game is rather bleak, because the firm pays w in advance, so the worker has no incentive to supply effort, so the firm offers $w = 0$ and the worker chooses self-employment. In the infinitely repeated game however, the firm can induce effort by paying a wage w in excess of w_0 and threatening to fire the worker if output is ever low.

We can consider the following strategies in the infinitely repeated game, which involve the wage $w^* > w_0$ to be determined later. The firms strategy is to offer $w = w^*$ in the first period, and in each subsequent period to offer $w = w^*$ provided that the history of play is high-wage, high-output, but to offer $w = 0$ otherwise. The worker's strategy is to accept the firm's offer if $w \geq w_0$ and to supply effort if the history of play including the current offer is high-wage, high-output.

We now derive conditions under which these strategies are a subgame-perfect Nash equilibrium. Suppose that the firm offers w^* in the first period. Given the firm's strategy, it is optimal for the worker to accept. if the worker supplies effort, then the worker is sure to produce high output. If the worker supplies effort, then the worker is sure to produce high output, so the firm will again offer w^* and the worker will face the same effort-supply decision next period. So, if it is optimal for the worker to supply effort, then the present value of the worker's payoffs is

$$V_e = (w^* - e) + \delta V_e \quad \text{or} \quad V_e = (w^* - e)/(1 - \delta)$$

If the worker shirks however, then the worker will produce high output with probability p, in which case the same-effort supply decision will arise next period, but the worker will produce low output with probability $1 - p$, in which case the firm will offer $w = 0$ forever after, so the worker will be self-employed forever after. Thus, if it is optimal for the worker to shirk, then the present value of the worker's payoffs is

$$V_s = w^* + \delta \left\{ pV_s + (1 - p)\frac{w_0}{1 - \delta} \right\}$$

Since it is optimal for the worker to supply effort if $V_e \geq V_s$, or

$$w^* \geq w_0 + \frac{1 - p\delta}{\delta(1 - p)}e = w_0 + \left(1 + \frac{1 - \delta}{\delta(1 - p)}\right)e$$

So, do induce effort, the firms need to pay not only $w_0 + e$ to compensate the worker for the foregone opportunity of self-employment and for the disutility of effort, but also the wage premium. So if p is close to one, then the wage premium must be extremely high to induce effort. If $p = 0$ on the other hand, then it is optimal for the worker to supply effort if

$$\frac{1}{1 - \delta}(w^* - e) \geq w^* + \frac{\delta}{1 - \delta}w_0$$

So, given the workers strategy, the firm's problem in the first period amounts to choosing between paying $w = w^*$, thereby inducing effort by threatening to fire the worker if low output is ever observed, and so receiving the payoff $y0w^*$ each period, and paying $w - 0$, thereby inducing the worker to choose self-employment, and so receiving the payoff zero in each period. So, the firm's strategy is a best-response to the workers if

$$y - w^* \geq 0$$

Recall that we assumed that $y - e > w_0$. We require more if these strategies are to be a subgame-perfect Nash equilibrium: our two equations imply the following;

$$y - e \geq w_0 + \frac{1 - \delta}{\delta(1 - p)}e$$

Which makes logical and historical sense based on what we know about the other subgames we've examined.

2.7 Dynamic Games of Complete but Imperfect Information

2.7.1 Extensive-Form Representation of Games

Back in chapter 1, we looked at static games by representing them in a normal form. We now analyze dynamic games by representing such games in extensive form. This approach may make it seem that static games must be represented in normal form and dynamic games in extensive form, but this is not the case. Any game, it turns out, can be represented in either normal or extensive form, and we'll talk about how dynamic games can be represented in normal form.

Definition

The **extensive-form representation** of a game specifies:

1. The players in the game
2. When each player has the move
3. What each player can do at each of his or her opportunities to move
4. What each player knows at each of his or her opportunities to move
5. The payoff received by each player for each combination of moves that could be chosen by the players

---

An example of a game in extensive form is as follows:

1. Player 1 chooses an action a_1 from the feasible set $A_1 = \{L, R\}$
2. Player 2 observes a_1 and then chooses an action a_2 from the set $A_2 = \{L', R'\}$
3. Payoffs are $u_1(a_1, a_2)$ and $u_2(a_1, a_2)$ as shown in the tree in figure 2.5

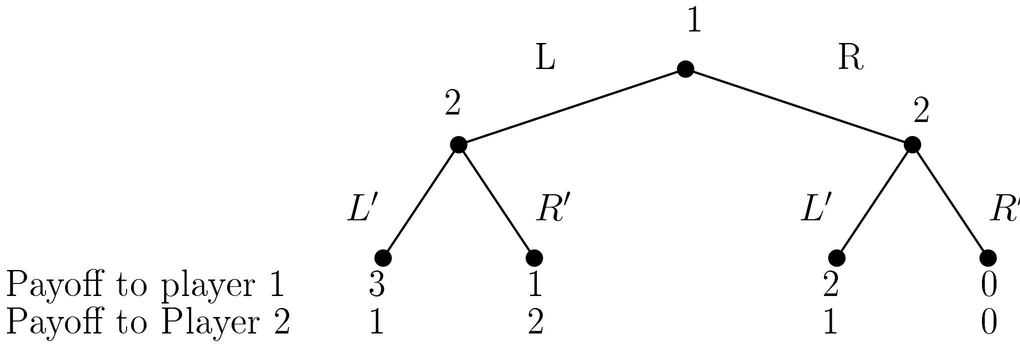

Figure 2.3:

The game begins with a *decision node* for player 1, where 1 chooses between L and R. If player 1 chooses L, then a decision node for player 2 is reached, where 2 then chooses between L' and R'. Following player 2's choices, a *terminal node* is reached, and the indicated payoffs are received.

It's actually pretty easy to extend the game shown above to represent any dynamic game of complete and imperfect information. We next derive the normal-form representation of the dynamic game as in figure 2.5. Recall the following definition:

Definition

A **strategy** for a player is a complete plan of action– it specifies a feasible action for the player in every contingency in which the player might be called on to act.

In figure 2.5 , player 2 has two actions but four strategies, because there are two different contingencies in which player 2 could be called upon to act:

1. If player 1 plays L, then play L', if player 1 plays R then play L', denoted by (L', L')
2. If player 1 plays L, then play L', if player 1 plays R then play R', denoted by (L', R')
3. If player 1 plays L, then play L', if player 1 plays R then play L', denoted by (R', L')
4. If player 1 plays L, then play R', if player 1 plays R then play R', denoted by (R', R')

Player 1 however, has two actions, but only two strategies: play L and play R. The reason player 1 has only two strategies is that there is only one contingency in which player 1 might be called upon to act, so player 1's strategy space is equivalent to the action space $A_1 = \{L, R\}$.

Given these strategy spaces for the two players, it is straightforward to derive the normal-form representation of the game from its extensive-form representation. This game is illustrated in the figure below.

		(L', L')	(L', R')	(R', L')	(R', R')
Player 1	L	3,1	3,1	1,2	1,2
	R	2,1	0,0	2,1	0,0

We now want to show how a static game can be represented in extensive form. We can represent a simultaneous-move game between player1 and 2 a follows:

1. Player 1 chooses an action a_1 from the feasible set A_1
2. Player 2 does not observe player 1's move but chooses an action a_2 from the feasible set A_2.
3. Payoffs are $u_1(a_1, a_2)$ and $u_2(a_1, a_2)$

Alternatively, player 2 could move first and player 1 could then move without observing 2's action. To represent this kind of ignorance of previous moves in an extensive-form game, we introduce the notion of a player's information set:

Definition

An **information set** for a player is a collection of decision nodes satisfying:

1. The player has the move at every node in the information set, and
2. When the play of the game reaches a node in the information set, the player with the move does not know which node in the information set has or has not been reached.

Part (ii) of this definition implies that each player must have the same set of feasible actions at each decision node in an information set, otherwise the player would be able to infer from the set of actions available that some node(s) had or had not been reached.

In an extensive-form game, we indicated that a collection of decision nodes constitutes an information set by connecting the nodes by a dotted line, as in the extensive-form representation of the Prisoner's Dilemma given in figure 2.4.

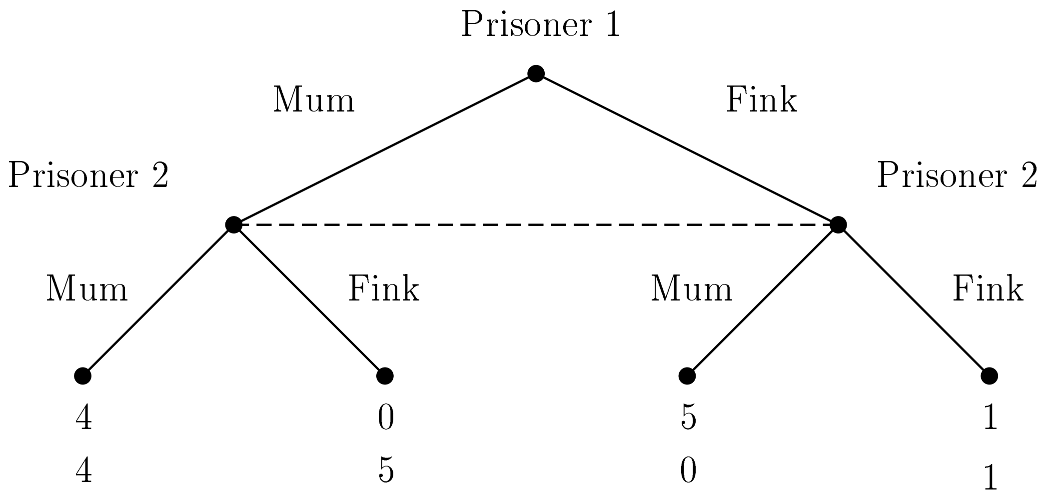

Figure 2.4:

As a second example of the use of an information set in representing ignorance of previous play, consider the following dynamic game of complete but imperfect information;

1. Player 1 chooses an action a_1 from the feasible set $A_1 = \{L, R\}$
2. Player 2 observes a_1 and then chooses an action a_2 from the feasible set $A_2 = \{L', R'\}$
3. Player 3 observes whether or not $(a_1, a_2) = (R, R')$ and then chooses an action a_3 from the feasible set $A_3 = \{L'', R''\}$

The extensive-form representation of this game is given in the following figure:

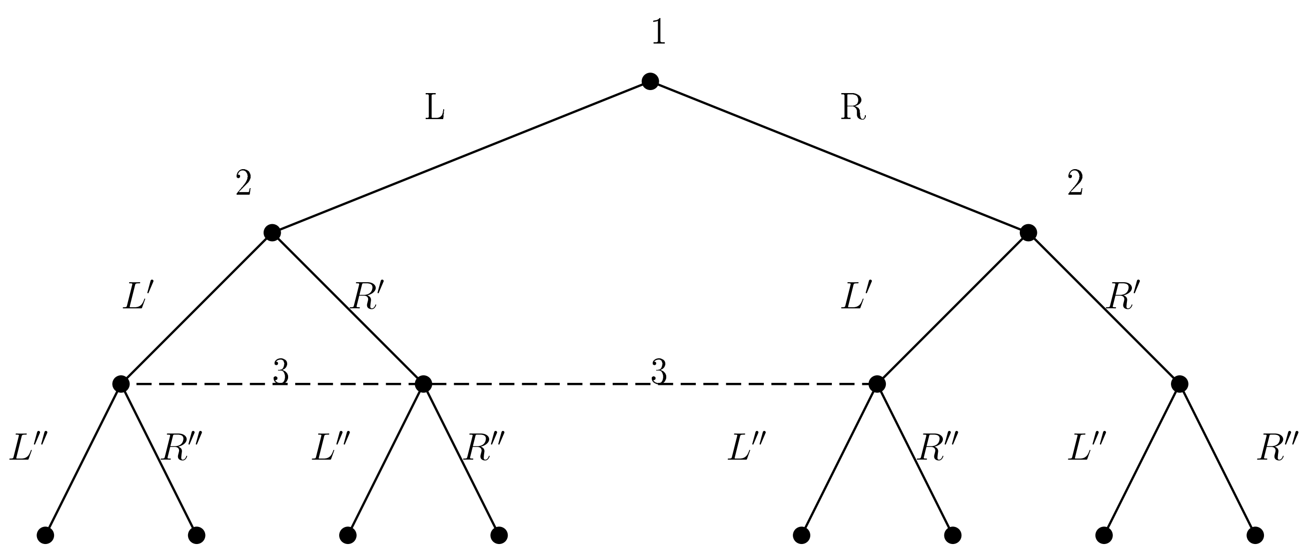

In this extensive form, player 3 has two information set: a singleton information set following R by player 1, and R' by player 2, and a non-singleton information set that includes every other node at which player 3 has the move. So, all player 3 observes is whether or not $(a_1, a_2) = (R, R')$.

Now comfortable with the notion of an information set, we can offer an alternative definition of the distinction between perfect and imperfect information. We previously defined perfect information to mean that at each move in the game, the player with the move knows the full history of the play of the game thus far. An equivalent definition of perfect information is that every information set is a singleton; imperfect information, in contrast, means that there is at least one non-singleton information set. So, the extensive-form representation of a simultaneous-move game is a game of imperfect information.

2.8 Subgame-Perfect Nash Equilibrium

Recall that we previously defined a subgame as the piece of a game that remains to be played beginning at any point at which the complete history of the game thus far is common knowledge among the players, and we gave a formal definition for the repeated games we considered there. We now have the following formal definition:

Definition

A **subgame** in an extensive-form game:

- Begins at a decision node n that is a singleton information set, but is not the game's first decision node,
- Includes all the decision and terminal nodes following n in the game tree, but no nodes that do not follow n, and
- Does not cut any information sets, i.e. if a decision node n' follows n in the game tree, then all other nodes in the information set containing n' must also follows n, and so must be included in the subgame.

So, given a general definition of a subgame, we can now apply the definition of a subgame-perfect Nash equilibrium:

Definition

A Nash equilibrium is **subgame-perfect** if the player's strategies constitute a Nash equilibrium in every subgame.

It is a fairly simple exercise to show that any finite dynamic game of complete information has a subgame-perfect Nash equilibrium, perhaps in mixed strategies. We have already encountered two ideas that are intimately related to subgame-perfect Nash equilibrium: the backwards-inductive outcome, and the subgame-perfect outcome. Put informally, the difference is that an equilibrium is a collection of strategies whereas an outcome describes what will happen only in the contingencies that are expected to arise, not in every contingency that might arise. Being more precise, we have the following definition:

Definition

In the two-stage game of complete and perfect information defined previously, the backwards-induction outcome is $(a_1^*, R_2(a_1^*))$ but the **subgame-perfect Nash Equilibrium** in Figure 2.3 is $(a_1^*, R_2(a_1))$.

In this game, the action a_1^* is a strategy for player 1 because there is only one contingency in which player 1 can be called upon to act- the beginning of the game. However, for player 2, $R_1(a_1^*)$ is an action, but not a strategy, because a strategy for player 2 must specify the action 2 will take following each of 1's possible first-stage actions. The best-response function $R_2(a_1)$, on the other hand is a strategy for player 2. In this game the subgames begin with player 2's move in the second stage. There is one subgame for each of player 1's feasible actions a_1 in A_1. To show that $(a_1^*, R_2(a_1))$ is a subgame-perfect Nash equilibrium, we have to show that $(a_1^*, R_2(a_1))$ is a Nash equilibrium and that the player's strategies constitute a Nash equilibrium in each of these subgames. Since the subgames are really just single-person decision problems, the latter reduces to requiring that player 2's action be optimal in every subgame, which is exactly the problem that the best-response function $R_2(a_1)$ solves. At last, $(a_1^*, R_2(a_1))$ is a Nash equilibrium because the player's strategies are best responses to each other: a_1^* is a best response to $R_2(a_1)$, that is, a_1^* maximizes $u_1(a_1, R_2(a_1))$ and $R_2(a_1)$ is a best response to a_1^*–that is, $R_2(a_1^*)$ maximizes $u_2(a_1^*, a_2)$.

This section concludes with an example that illustrates the main theme of the chapter: subgame-perfection eliminates Nash equilibria that rely on non-credible threats or promises. Recall the extensive-form game in figure 2.5, along with its normal-form representation.

		(L', L')	(L', R')	(R', L')	(R', R')
Player 1	L	3,1	3,1	1,2	1,2
	R	2,1	0,0	2,1	0,0

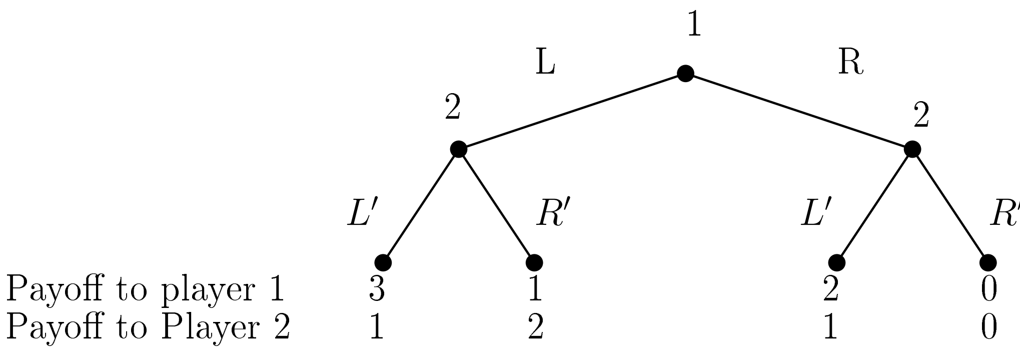

Figure 2.5:

Had we encountered the extensive-form game earlier, we would have solved it by backwards induction, getting the outcome (R, L'). Had we encountered its normal-form representation, we would have solved for its pure strategy Nash-equilibria, which are $(R, (R', L'))$, and $(L, (R', R'))$. Now comparing these Nash equilibria in the normal-form game with the results of the backwards-induction procedure, the Nash equilibrium $(R, (R', L'))$ corresponds to **all** the bold paths shown below:

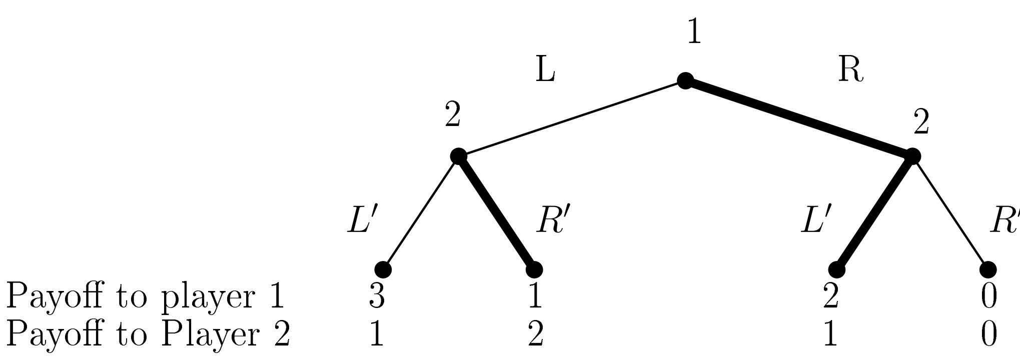

Figure 2.6:

Before, we called (R, L') the backwards-induction Nash equilibrium of the game, but we will use more general terminology and call it the subgame-perfect Nash equilibrium. The difference between the outcome and the equilibrium is that the outcome specifies only the bold path beginning at the game's first decision node and concluding at a terminal node, whereas the equilibrium also specifies the addition all bold path emanating from player 2's decision node following L from player 1.

But, what about the other Nash equilibrium $(L,(R',R'))$? In this equilibrium, player 2's strategy is to play R' not only if player 1 chooses L, but also if player 1 chooses R. Because R' leads to a payoff of 0 for player 1, player 2's best response to this strategy by player 2 is to play L, getting a better payoff of 1 for player 1, which is better than 0. One can say that 'player 2 is threatening to play R' if player 1 plays R'. If this threat works, then 2 is not given the opportunity to carry out the threat. The threat should not work though, because it isn't credible: if player 2 were given the opportunity to carry it out, then player 2 would just play L' instead of R'. More formally, the Nash equilibrium $(L,(R',R'))$ is not subgame-perfect, because the player's strategies do not constitute a Nash equilibrium in one of the subgames.

CHAPTER

THREE

STATIC GAMES OF INCOMPLETE INFORMATION

We now begin our study of games that involve something called *incomplete information*, or *Bayesian games*. Remember that in games of complete information, the player's payoff functions are all common knowledge. In games of incomplete information, at least one player is uncertain about another player's payoff function. One common example of a static game of incomplete information is a sealed-bid auction: each bidder knows his or her own valuation for the good being sold but does not know any other bidder's valuation; bids are submitted in sealed envelopes, so the player's moves can be thought of as simultaneous. Let us begin to move into methods which which we can use to analyze these games.

3.1 Static Bayesian Games and Bayesian Nash Equilibrium

3.1.1 Cournot Competition under Asymmetric Information

Consider a Cournot duopoly model wit inverse demand given by the following function:

$$P(Q) = a - Q$$

where $Q = q_1 + q_2$ is the aggregate quantity on the market. Firm 1's cost function is $C_1(q_1) = cq_1$. However, Firm 2's cost function is $C_2(q_2) = c_H q_2$ with probability θ and $C_2(q_2) = c_L q_2$ with probability $1-\theta$, where $c_L < c_H$. Furthermore, information is asymmetric: firm 2 knows its cost function and firm 1's, but firm 1 knows its cost function and only that firm 2's marginal cost is c_H with probability θ and c_L with probability $1 - \theta$. All of this is common knowledge. Firm 1 knows that firm 2 has superior information, firm 2 knows this, and so on.

Naturally, firm 2 might want to choose a different quantity if its marginal cost is high than if it is low. Firm 1 should anticipate that firm 2 may adapt its quantity to its cost in this way. Let $q_2^*(c_H)$ and $q_2^*(c_L)$ denote firm 2's quantity choice as a function of its cost, and let q_1^* denote firm 1's single quantity choice. If firm 2's cost is high, it will choose $q_2^*(c_H)$ to solve the following:

$$max_{q_2}[(a - q_1^* - q_2) - c_H]q_2$$

Similarly, if firm 2's cost is low, then $q_2^*(c_L)$ will solve

$$max_{q-2}[(a - q_2^* - q_2) - c_L]q_2$$

Finally, firm 2 knows that firm 2's cost is high with probability θ and should anticipate that firm 2's quantity choice will be $q_2^*(c_L)$, depending on firm 2's cost. So, firm 1 chooses q_1^* to solve:

$$max_{q_1}\theta[(a - q_1 - q_2^*(c_H)) - c]q_1 + (1-\theta)[(a - q_1 - q_2^*(c_L)) - c]q_1$$

The first-order conditions for these problems are the following:

$$q_2^*(c_H) = \frac{a - q_1^* - c_H}{2}$$
$$q_2^*(c_L) = \frac{a - q_1^* - c_L}{2}$$

and:

$$q_1^* = \frac{\theta[(a - q_1 - q_2^*(c_H)) - c] + (1-\theta)[(a - q_1 - q_2^*(c_L)) - c]}{2}$$

The solutions to the three first-order conditions are the following:

$$q_2^*(c_H) = \frac{a - 2c_H + c}{3} + \frac{1-\theta}{6}(c_H - c_L)$$
$$q_2^*(c_L) = \frac{a - 2c_L + c}{3} - \frac{\theta}{6}(C_H - c_L)$$
$$q_1^* = \frac{a - 2c + \theta c_H + (1-\theta)c_L}{3}$$

3.1.2 Normal Form Representation of Static Bayesian Games

Recall that when we have an n-player game of complete information, its normal-form representation is $G = \{S_1, ..., S_n; u_1, ..., u_n\}$, where S_i is player i's strategy space and $u_i(s_1, ..., s_n)$ is player i's payoff when the players choose the strategies $(s_1, s_2, ..., s_n)$. We now want to talk about developing the normal-form representation of a simultaneous move of incomplete information, also called a static Bayesian game. The first step is to represented the idea that each player knows his or her own payoff function but may be uncertain about the other player's payoff functions. Let player i's possible payoff functions be represented by $u_i(a_1, .., a_n; t_i)$ where t_i is called player i's *type* and belongs to set of possible types. Each type t_i corresponds to a different payoff function that player i might have.

As a concrete example, consider the Cournot game in the previous section. The Firm's actions are their quantity choices, q_1 and q_2. Firm 2 has two possible cost functions and thus two possible profit or payoff functions:

$$\pi_2(q_1, q_2; c_L) = [(a - q_1 - q_2) - c_L]q_2$$

and

$$\pi_2(q_1, q_2; c_H) = [(a - q_1 - q_2) - c_H]q_2$$

Firm 1 has only one possible payoff function:

$$\pi_1(q_1, q_2; c) = [(a - q_1 - q_2) - c]q_1$$

We say that firm 2's type space is $T_2 = \{c_L, c_H\}$ and that firm 1's type space is $T_1 = \{c\}$.

From this definition, saying that player i knows his or her own payoff function is equivalent to saying that player i knows his or her type. Similarly, saying that player i may be uncertain about the other player's payoff functions is equivalent to saying that player i may be uncertain about the types of the other players, denoted by

$$t_{-i} = (t_1, ..., t_{i-1}, t_{i+!}, ..., t_n)$$

We use T_{-i} to denoted the set of all possible values of t_{-i}, and we use the probability distribution $p(t_{-i}|t_i)$ to denote player i's belief about the other player's types, given player i's knowledge of his or her own type, t_i. Joining some of these concepts, we have the following definition:

Definition

The **normal-form representation** of an n-player static Bayesian game specifies the player's action spaces $A_1, ..., A_n$, their type spaces $T_1, ..., T_n$, their beliefs $p_1, ..., p_n$, and their payoff functions $u_1, ..., u_n$. Player i's **type** t_i is privately known by player i, determines player i's payoff function $u_i(a_1, ..., a_n; t_i)$, and is a member of the set of possible types T_i. Player i's **belief** $p_i(t_{-i}|t_i)$ describes i's uncertainty about the $n-1$ other player's possible types t_{-i}, given i's own type t_i. We denote this game by:

$$G = \{A_1, ..., A_n; T_1, ..., T_n; p_1, ..., p_n; u_1, ..., u_n\}$$

We will assume that the timing of a static Bayesian game is as follows:

1. Nature draws a type vector $t = (t_1, ..., t_n)$, where t_i is drawn from the set of possible types T_i;
2. Nature reveals t_i to player i but not to any other player ;

3. The players simultaneously choose actions, player i choosing a_i from the feasible set A_i, and then

4. Payoffs $u_i(a_1, ..., a_n; t_1)$ are received.

By introducing the fictional moves by nature in steps 1 and 2, we have described a game of incomplete information as a game of imperfect information, where by imperfect information we mean that at some move in the game the player with the move doesn't know the complete history of the game thus far.

Two more technical points need to be discussed to complete our thoughts on normal-from representations of static Bayesian games First, there are games in which player i has private information not only about his payoff function but also about another player's payoff function. In the n-player case we capture this possibility by allowing player i's payoff to depend not only on the actions $(a_1, ..., a_n)$ but also on the types $(t_1, ..., t_n)$. We write this payoff as $u_i(a_1, ..., a_n; t_1, ..., t_n)$.

The second technical point involves the beliefs $p_i(t_{-i} \mid t_i)$. We will assume that it is common knowledge that in step 1 of the timing of a static Bayesian game, nature draws a type vector $t = (t_1, ..., t_n)$ according to the prior probability distribution $p(t)$. When nature then reveals t_i to player i, he or she can compute the belief $p_i(t_{-i} \mid t_i)$ using Baye's rule:

$$p_i(t_{-i} \mid t_i) = \frac{p(t_{-i}, t_i)}{p(t_i)} = \frac{p(t_{-i}, t_i)}{\sum_{t_{-i} \in T_{-i}} p(t_{-i}, t_i)}$$

3.1.3 Definition of Bayesian Nash Equilibrium

We want to define an equilibrium concept for static Bayesian games. To do this we first need to define the player's strategy spaces in such a game. We have the following:

Definition

In the static Bayesian game $G + \{A_1, ..., A_n; T_1, ..., T_n'; p_1, ..., p_n; u_1, ..., u_n\}$, a **strategy** for player i is a function $s_i(t_i)$ where for each t_i in T_i, $s_i(t_i)$ specifies the action from the feasible set A_i that type t_i would choose if drawn by nature.

Unlike games of complete information, in a Bayesian game the strategy spaces are not given in the normal form representation of the game. Instead, the strategy spaces are constructed from the type and action spaces: player i's set of possible strategies S_i is the set of all possible functions with domain T_i and range A_i.

Given the definitino of a strategy in a Bayesian game, we now look at the definition of a Bayesian Nash equilibrium. In spite of the notational complexity of the definition is both simple and familiar: each player's strategy must be a best response to the other player's strategies. That is a Bayesian Nash equilibrium is simply a Nash equilibrium in a Bayesian game.

Definition

In the static Bayesian game $G = \{A_1, ..., A_n; T_1, ..., T_n; p_1, ..., p_n; u_1, ...u_n\}$ the strategies $s^* = (s_1^*, ..., s_n^*)$ are a pure strategy **Bayesian Nash Equilibrium** if for each player i and for each of i's types $t_i \in T_i$, $s_i^*(t_i)$ solves

$$max_{a_i \in A_i} \sum_{t_{-i} \in T_{-i}} u_i(s_1^*(t_1), ..., s_{i-1}^*(t_{i-1}), a_i, s_{i+1}^*(t_i + 1), ..., s_n^*(t_n); t)p_i(t_{-i}|t_i)$$

That is, no player wants to change his or her strategy, even if the change involves only one action by type.

It is easy to show that in a finite static Bayesian game, there exists a Bayesian Nash equilibrium. The proof closely looks like the proof of the existence of a mixed-strategy Nash equilibrium in finite games of complete information.

3.2 Applications

3.2.1 Battle of the Sexes

Remember that in the Battle of the Sexes, there are two pure-strategy Nash equilibria $(Opera, Opera)$ and $(Fight, Fight)$ and a mixed- strategy Nash equilibrium in which Chris plays Opera with probability 2/3 and Pat plays fight with probability 2/3.

		Pat	
		Opera	Fight
Chris	Opera	2,1	0,0
	Fight	0,0	1,2

Now suppose that although they have known each other for quite some time, Chris and Pat aren't sure of each other's payoffs. Suppose that Chris's payoff if both attend the opera is $2 + t_c$, where t_c is privately known by Chris, Pat's payoff if both attend the fight is $2 + t_p$, where t_p is privately known by Pat, and t_c and t_p are independent draws from a uniform distribution on $[0, x]$. All the other payoffs are the same. In terms of the abstract static Bayesian game in normal form, $G = \{A_c, A_p; T_c, T_p; p_c, p_p; u_c, u_p\}$, the action spaces are $A_c = A_p = \{Opera, Fight\}$, the type spaces are $T_c = T_p = [0, x]$, the beliefs are $p_c(t_p) = p_p(t_c) = 1/x$ for all t_c, t_p, and the payoffs are now as above.

		Pat	
		Opera	Fight
Chris	Opera	$2 + t_c, 1$	0,0
	Fight	0,0	$1, 2 + t_p$

We can construct a pure-strategy Bayesian Nash equilibrium of this incomplete-information version of the Battle of the Sexes in which Chris plays Opera if t_c exceeds a critical value c and plays Fight otherwise. Similarly Pat plays Fight if t_p is greater than some critical value p

In such an equilibrium, Chris plays Opera with probability $(x - c)/x$ and Oat plays Fight with probability $(x - p)/x$. For a given value of x, we will determine values c, p such that these strategies are a Bayesian Nash equilibrium. Given Pat's strategy, Chris's expected payoffs from playing Opera and from playing Fight are:

$$\frac{p}{x}(2 + t_c) + \left[1 - \frac{p}{x}\right] \cdot 0 = \frac{p}{x}(2 + t_c)$$

and

$$\frac{p}{x} \cdot 0 + \left[1 - \frac{p}{x}\right] \cdot 1 = 1 - \frac{p}{x}$$

So, playing Opera is only optimal if

$$t_c \geq \frac{x}{p} - 3 = c$$

Similarly, Pat's expected payoffs from playing Fight and Opera respectively are:

$$\frac{c}{x}(2 + t_c) + \left[1 - \frac{c}{x}\right] \cdot 0 = \frac{c}{x}(2 + t_c)$$

and

$$\frac{c}{x} \cdot 0 + \left[1 - \frac{c}{x}\right] \cdot 1 = 1 - \frac{c}{x}$$

So playing Fight is optimal if and only if

$$t_p \geq \frac{x}{c} - 3 = p$$

Solving these two optimal strategies simultaneously leads $p = c$ and $p^2 + 3p - x = 0$. Solving this quadratic gives us the following

$$1 - \frac{3 + \sqrt{9 + 4x}}{2x}$$

Which approaches $2/3$ as x goes to zero. So, as the incomplete information disappears, the player's behavior in this pure strategy Bayesian Nash equilibrium of the incomplete information game approaches their behavior in the mixed strategy Nash equilibrium in the original game of complete information.

3.2.2 An Auction

Consider the following first-price sealed-bid auction. There are two bidders, $i = 1, 2$. Bidder i has a valuation v_i for the good, that is, if bidder i gets the good and pays the price p then i's payoff is $v_i - p$. The two bidders' valuations are independently and uniformly distributed on $[0, 1]$. Bids are nonnegative, the high bidder wins the good, and in a tie, the winner is determined by the flip of a coin. All of this is common knowledge.

In formulating this problem as a static Bayesian game, we identify the action and type spaces, the beliefs, and the payoff functions. Player i's action is to submit a bid b_i, and her type is her valuation v_i. Because the valuations are independent, player i believes that v_j is uniformly distributed on $[0, 1]$ no matter what the value of v_i. Finally, player i's payoff function is:

$$u_i(b_1, b_2; v_1, v_2) = \begin{cases} v_i - b_i & \text{if } b_i > b_j \\ (v_i - b_i)/2 & \text{if } b_i = b_j \\ 0 & \text{if } b_i < b_j \end{cases}$$

We now have to construct the player's strategy spaces. Recall that in a static Bayesian game, a strategy is a function from types to actions. So, a strategy for i is a function $b_i(v_i)$ specifying the bid that each of i's types would choose. In a Bayesian Nash equilibrium, player 1's strategy $b_1(v_1)$ is a best response to player 2's strategy of $b_2(v_2)$, and vice versa. So, these pairs must satisfy the following;

$$max_{b_i}(v_i - b_i)Prob\{b_i > b_j(v_j)\} + \frac{1}{2}(v_i - b_i)Prob\{b_i = b_j(v_j)\}$$

Suppose that player j adopts the strategy $b_j(v_j) = a_j + c_j v_j$. For some given value of v_i, player i's best response solves the following:

$$max_{b_i}(b_i - b_i)Prob\{b_i > a_j + c_j v_j\}$$

where we have used the fact that $Prob\{b_i = b_j(v_j)\} = 0$, because $b_j(v_j) = a_j + c_j v_j$ and v_j is uniformly distributed. Since it is pointless for player i to bid below player j's minimum and bid, and stupid for player i to bid above player j's maximum, so we have $a_j \le b_i \le a_j + c_j$, so

$$Prob\{b_i > a_j + c_j v_j\} = Prob\{v_j < \frac{b_i - a_j}{c_j}\} = \frac{b_i - a_j}{c_j}$$

So, player i's best response is:

$$b_i(v_i) = \begin{cases} (v_i + a_j)/2 & \text{if } v_i \ge a_j \\ a_j & \text{if } v_i < a_j \end{cases}$$

3.2.3 A Double Auction

Suppose we have a buyer and seller, where each have private information about their valuations. In our case, we take a firm and a worker- the buyer is a firm, and the seller is a worker. The firm knows the worker's marginal product and the worker knows his or her outside opportunity. The seller names and asking price p_s, and the buyer simultaneously names an offer price p_b. If $p_b \geq p_s$, then trade occurs at price $p = (p + b + p_s)/2$; if $p_b < p_s$, then no trade occurs.

The buyer's valuation for the sellers' good is v_b, the seller's is independent uniform distributions on $[0, 1]$. If the buyer gets the good fro price p, then the seller's utility is $v_b - p$; if there is no trade, then the buyer's utility is zero. If the seller sells the good for price p, then the seller's utility $p - v_s$; if there is no trade, then the seller's utility is zero.

In this static Bayesian game, a strategy for the buyer is a function $p_b(v_b)$ specifying the price the buyer will offer for each of the buyer's possible valuations. Likewise, a strategy for the seller is a function $p_s(v_s)$ specifying the price the seller will demand for each of the seller's valuations. A pair of strategies $\{p_b(v_b), p_s(v_s)\}$ is a Bayesian Nash equilibrium if the following conditions hold:

For each $v_b \in [0, 1]$, $p_b(v_b)$ solves:

$$max_{p_b} \left[v_b - \frac{p_b + E[p_s(v_s)|p_b \geq p_s(v_s)]}{2} \right] Prob\{p_b \geq p_s(v_s)\}$$

And for each $v_s \in [0, 1]$, $p_s(v_s)$ solves:

$$max_{p_s} \left[\frac{p_s + E[p_b(v_b)|p_b(v_b) \geq p_s]}{2} - v_s \right] Prob\{p_b(v_b) \geq p_s\}$$

Where $E[p_b(v_b)|p_b(v_b) \geq p_s]$ is the expected price the buyer will offer, condition on the offer being greater than the seller's demand of p_s.

WE can now derive a linear Bayesian Nash equilibrium of the double auction- as previously, we aren't restricting the player's strategy spaces to include only linear strategies. Many other equilibria exist besides the one-price equilibria and the linear equilibrium, but the linear equilibrium has some interesting properties we would like to explore later.

Suppose the seller's strategy is $p_s(v_s) = a_s + c_s v_s$. Then p_s is uniformly distributed on $[a_2, a_s + c_s]$, so we have:

$$max_{p_b} \left[v_b - \frac{1}{2} \left\{ p_b + \frac{a_s + p_b}{2} \right\} \right] \frac{p_b - a_s}{c_s}$$

the first order condition for which is:

$$p_b = \frac{2}{3} v_b + \frac{1}{3} a_s$$

Similarly, suppose the buyer's strategy is $p_b(v_b) = a_b + c_b v_b$. Then, we have

$$max_{p_s} \left[\frac{1}{2} \left\{ p_s + \frac{p_2 + a_b + c_b}{2} \right\} - v_s \right] \frac{a_b + c_b - p_s}{c_b}$$

the first order condition for which is:

$$p_s = \frac{2}{3}v_s + \frac{1}{3}(a_b + c_b)$$

As shown in the figure below:

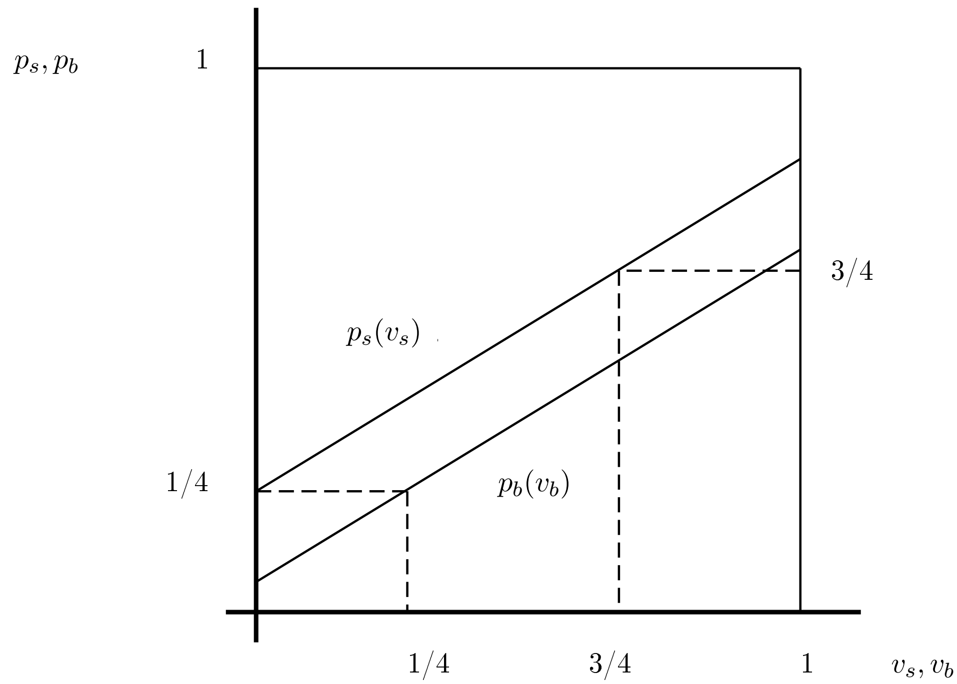

3.3 The Revelation Principle

The Revelation principle in the context of Bayesian games is an important tool for designing games when the players have private information. We can apply it to the auction and bilateral-trading problems described previously, and to lots of other Bayesian game problems.

Consider a seller that wants to design an auction to maximize his revenue. Specifying the many different auctions the seller should consider could be a big job, and overall very tedious. Fortunately, the seller can use the Revelation principle to simplify his problem in two ways. First, the seller can restrict attention to the following class of games:

1. The bidders simultaneously make claims about their types, Bidder i can claim to be any type τ_i from i's set of feasible types T_i, no matter what i's true type t_j.
2. Given the bidder's claims $(\tau_1, ..., \tau_n)$, bidder i pays $x_i(\tau_1, ..., \tau_n)$ and receives the good with probability $q_i(\tau_1, ..., \tau_n)$. For each possible combination of claims $(\tau_1, ..., \tau_n)$ the sum of the probabilities must be less than or equal to one.

Games of this kind (static Bayesian games where each player's action is to submit a claim about his or her type) are called *direct mechanisms*.

The second way the seller can use the Revelation principle is to restrict attention to those direct mechanisms in which it is a Bayesian Nash equilibrium for each bidder to tell the truth. A direct mechanism in which truth-telling is a Bayesian Nash equilibrium is called *incentive-compatible*. We conclude with the following definition:

Definition

Any Bayesian Nash equilibrium of any Bayesian game can be represented by an incentive-compatible direct mechanism.

CHAPTER

FOUR

DYNAMIC GAMES OF INCOMPLETE INFORMATION

In this chapter, we get tot about the equilibrium concept of *perfect Bayesian equilibrium*. This will make four equilibrium concepts:

1. Nash equilibrium in static games of complete information
2. Subgame perfect Nash equilibrium in dynamic games of complete information
3. Bayesian Nash equilibrium in static games of incomplete information
4. Perfect Bayesian equilibrium in dynamic games of incomplete information

4.1 Introduction to Perfect Bayesian Equilibrium

Perfect Bayesian equilibrium was invented in order to refine Bayesian Nash equilibrium. Consider the following dynamic game of complete but imperfect information. First, player 1 chooses among three actions, L,M, and R. If player 1 chooses R then the game ends without a move by player 2. If player 1 chooses either L or M, then player 2 learns that R was not chosen (but not which of L or M was chosen) and then chooses between two actions L' and R', after which the game ends. Payoffs are given in the extensive form, in figure 4.1.

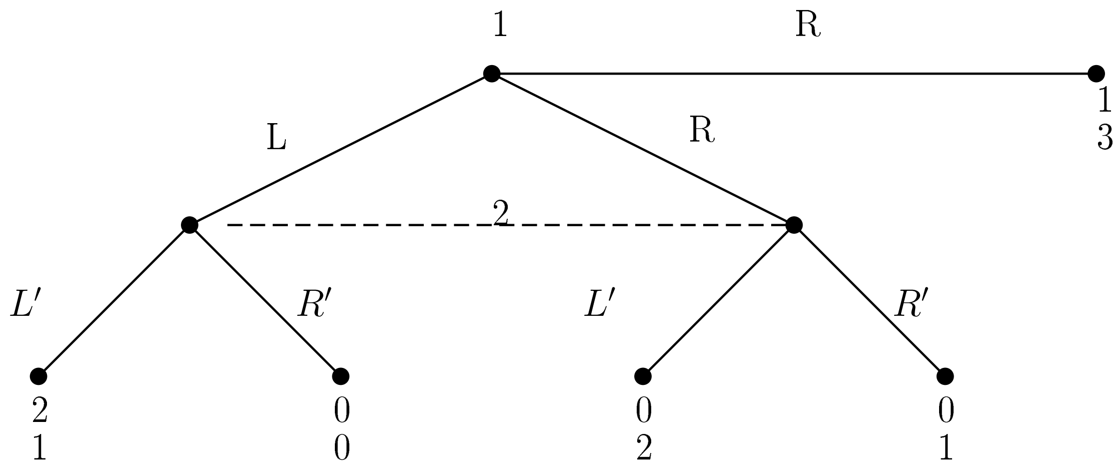

Figure 4.1:

We also have the following normal-form representation for this game:

Player 2

Player 1		L'	R'
	L	2,1	0,0
	M	0,2	0,1
	R	1,3	1,3

Looking at the normal-form representation of this game, we see that there are two pure-strategy Nash equilibria- (L, L') and (R, R'). To determined whether these Nash equilibria are subgame-effect, we use the extensive-form representation to define the game's subgames. Because a subgame is defined to begin at a decision node that is a singleton information set, the game in figure 4.1 has no subgames. If a game has no subgames, then the requirement of the subgame-perfection is trivially satisfied (namely that the player's strategies constitute a Nash equilibrium on every subgame). So, in any game that has no subgame, the definition of subgame-perfect Nash equilibrium is equivalent to the definition of Nash equilibrium, so we see that both (L, L') and (R, R') are subgame-perfect Nash equilibria. Also, (R, R')

clearly depends on a non-credible threat: if player 2 gets the move, then player L' dominates playing R', so player 1 should not be induced to play R by 2's threat to play R' if given the move. We can strengthen the equilibrium concept so as to rule out the subgame-perfect Nash equilibrium (R, R') in 4.1 is to make the following two requirements:

1. Ate each information set, the player with the move must have a belief about which node in the information set has been reached by the play of the game. For a non-singleton information set, a belief is a probability distribution over the nodes in the information set; for a singleton information set, the player's belief puts probability one on the single decision node

2. Given their beliefs, the players strategies must be sequentially rational. That is, at each information set the action taken by the player with the move must be optimal given the player's belief at that information set and the other player's subsequent strategies where a subsequent strategy is a complete plan of action covering every contingency that might arise after the given information set has been reached.

In figure 4.1, our first requirement implies that if the play of the game reaches player 2's nonsingleton information set, then player 2 must have a belief about which node has been reached. This belief is represented by the probabilities p and $1-p$ attached to the relevant nodes in the tree, as we have in the following figure:

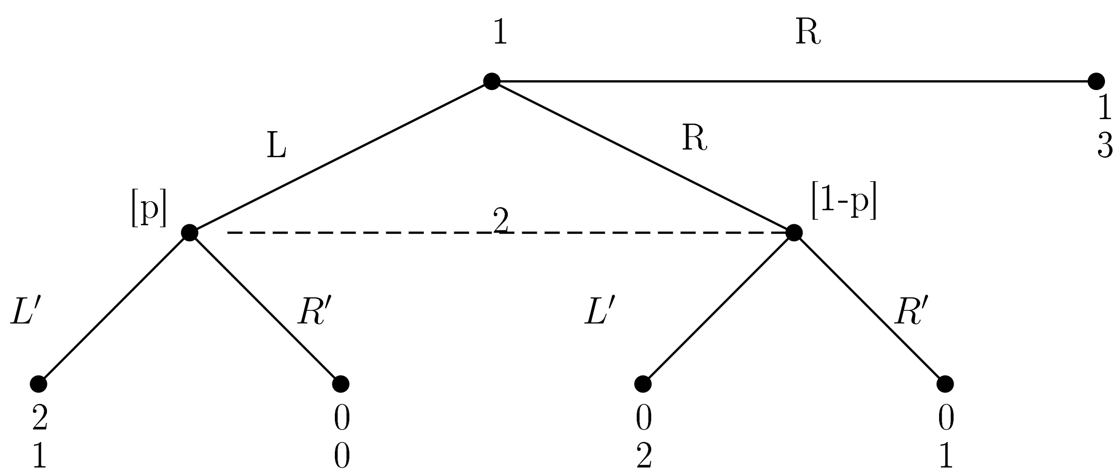

Given player 2's belief, the expected payoff from playing R' is

$$p \cdot 0 + (1-p) \cdot 1 = 1 - p$$

and the expected payoff from playing L' is

$$p \cdot 1 + (1-p) \cdot 2 = 2 - p$$

Since $2 - p > 1 - p$ for any value of p, requirement 2 prevents player 2 from choosing R'. So, this requires that each player have a belief and act optimally given this belief suffices to eliminated the implausible equilibrium (R, R').

Our two requirements insist that the players have beliefs and act optimal given these beliefs, but not that these beliefs have to be reasonable. In order to impose further requirements on the player's beliefs, we distinguish between information sets that are on the equilibrium path and those that are off the equilibrium path:

Definition

For a given equilibrium in a given extensive-form game, an information set is **on the equilibrium path** if it will be reached with positive probability if the game is played according to the equilibrium strategies, and is **off the equilibrium path** if it is certain not to be reached if the game is played according to the equilibrium strategies (where "equilibrium" can mean Nash, subgame-perfect, Bayesian, or perfect Bayesian equilibrium).

This also gives us a third requirements:

3 At information sets on the equilibrium path, beliefs are determined by Baye's rule and the player's equilibrium strategies.

Requirements 1 through 3 are really what we want to focus on when we're talking about perfect Bayesian equilibrium. The important new feature of this equilibrium is as follows: beliefs are elevated to the level of importance of strategies in the definition of equilibrium. Formally, an equilibrium no longer consists of just a strategy for each player but now also includes a belief for each player at each information set at which the player has the move. The advantage of making the player's belief explicit in this way is that just as in earlier chapters, we insisted that the players choose credible strategies, we can now also insist that they hold reasonable beliefs, both on the equilibrium path and off the equilibrium path. We now add a fourth requirement:

4 At information sets off the equilibrium path, beliefs are determined by Baye's rule and the player's equilibrium strategies where possible.

Definition

A **perfect Bayesian equilibrium** consists of strategies and beliefs satisfying requirements 1 through 4.

To motivate requirement four, we look at the three player games in **??** and 4.3.

This game has one subgame: it begins at playe 2's singleton information set. the unique Nash equilibrium in this subgame between player 2 and 3 is (L, R'), so the unique subgame-perfect Nash equilibrium of the entire game is (D, L, R'). These strategies and the beleif $p = 1$ for

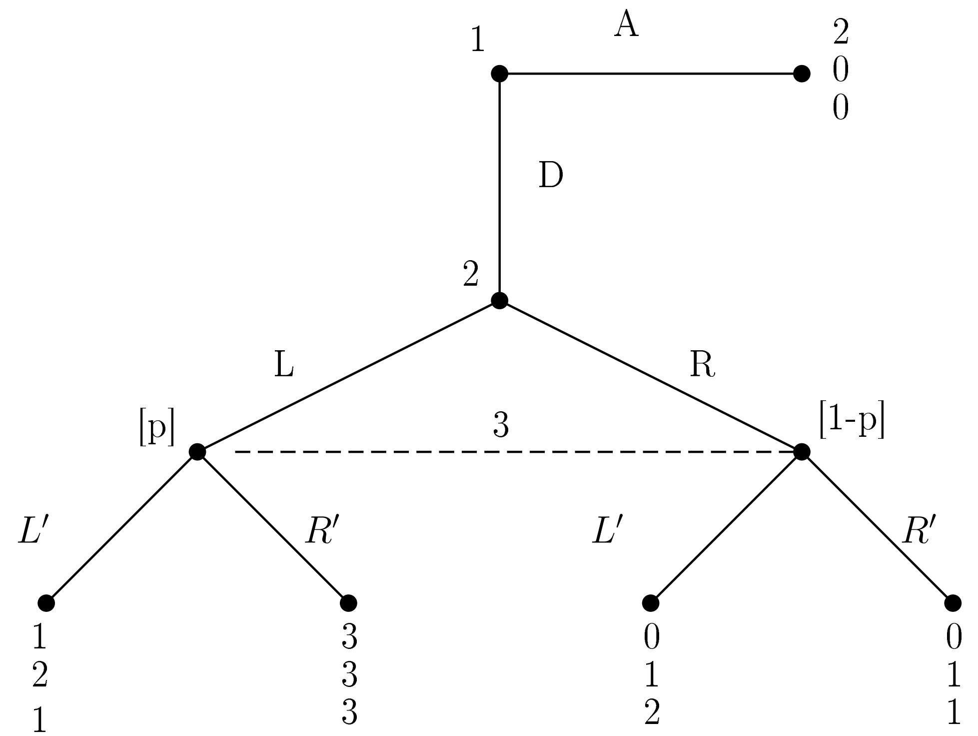

Figure 4.2:

hte player 3 satisfy Requirements 1 through 3. They also trivially satisfy requirement 4, since there is no information set off this equilibrium path, and so constittue a perfect Bayesian equilibrium.

Now consider the strategies (A, L, L') together with the belief $p = 0$. These strategies are a Nash equilibrium- no player wants to deviate unilaterally. These strategies and belief also satisfy Requirements 1 through 3 has a belief and acts optimally given it, and player 1 and 2 act optimally given the subsequent strategies of the other players. But this Nash equilibrium is not subgame-perfect, because the unique Nash equilibrium of the games only subgame is (L, R'). So, requirements 1-3 do not guarantee that the player's strategies are a subgame perfect Nash equilibrium. The problem is that player 3's belief $(p = 0)$ is inconsistent with player 2's strategy (L), but requirements 1 through 3 impose no restrictions on 3's belief because 4's informations et is not reached if the game is played according to the specified strategies. Requirement 3, however, forces player 3's belief to be determined by player 2's strategy: if 2's strategy is L then 3's belief must be $p = 1$; if 2's strategy is R then 3's belief must be $p = 0$. But if 3's belief is $p = 1$, then requirement 2 forces 3's strategy to be R', so the strategies (A, L, L') and the belief $p = 0$ do not satisfy our four requirements.

As another illustration of this 4^{th} requirement, look at figure 4.3. AS before, if player 1's equilibrium strategy is A then player 3's information set is off the equilibrium path, but now Requirement 4 may not determined 3's belief from 2's strategy. If 2's strategy is A' then Requirement 4 puts no restrictions on 3's belief, but if 2's strategy is to play L with probability q_1, R with probability q_2, and A' with probability $1 - q_1 - q_2$, where $q_1 + q_2 > 0$,

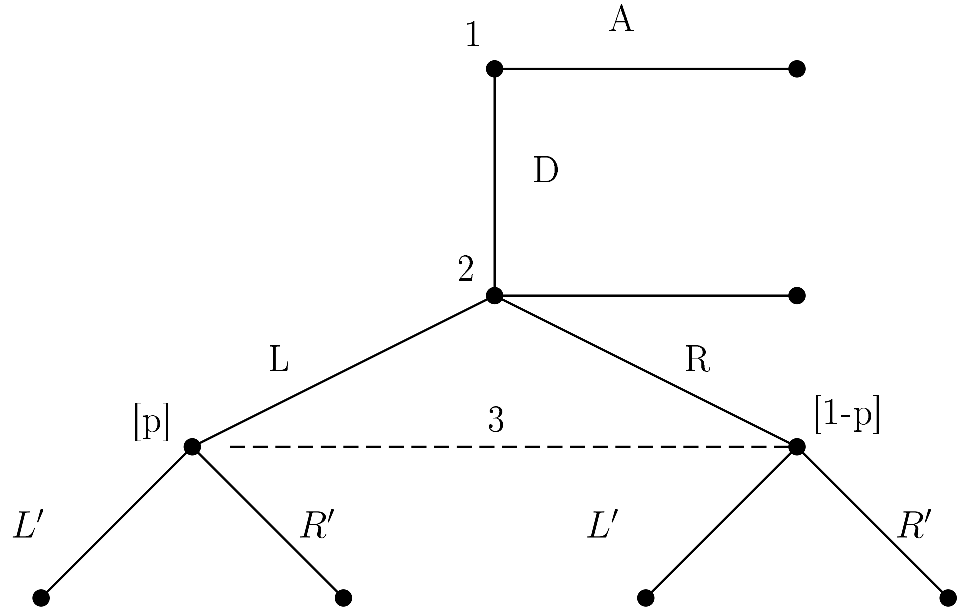

Figure 4.3:

then Requirement 4 forces 3's belief to be $p = q_1/(q_1 + q_2)$.

To finis up this section, we note a connection between perfect Bayesian equilibrium to the equilibrium concepts introduced in earlier chapters. In a Nash equilibrium, each player's strategy must be a best response to the other player's strategies, so no player chooses a strictly dominated strategy. In a perfect Bayesian equilibrium, Requirements 1 and 2 are equivalent to insisting that no player's strategy be strictly dominated beginning at any information set. Nash and Bayesian Nash equilibrium do not share this feature at information sets off the equilibrium path; even subgame-perfect Nash equilibrium does not share this feature at some information sets off the equilibrium path. Perfect Bayesian equilibrium closes this loophole: players can't threaten to play strategies that are strictly dominated beginning at any information set off the equilibrium path.

4.2 Signaling Games

4.2.1 Perfect Bayesian Equilibrium in Signaling Games

A signaling game is a dynamic game of incomplete information involving two players: a sender (S) and a receiver (R). The game looks as follows:

1. Nature draws a type t_i for the sender from a set of feasible types $T = \{t_1, ..., t_I\}$ according to a probability distribution $p(t_i)$, where $p(t_i) > 0$ for every i and $p(t_1) + ... + p(t_I) = 1$.

2. The sender observes t_i and then chooses a message m_j from a set of feasible messages $M = \{m_1, ..., m_J\}$
3. The receiver observes m_j (but not t_i) and then chooses an action a_k from a set of feasible actions $A = \{a_1, ..., a_K\}$
4. Payoffs are given by $U_S(t_i, m_j, a_k)$ and $U_R(t_i, m_j, a_k)$

In many real applications, the sets T, M and A are intervals on the real line rather than the finite sets we're going to consider here. It is straightforward to allow the set of feasible message to depend on the type nature draws, and the set of feasible actions to depend on the message the Sender chooses. For the rest of this section, we analyze the abstract signaling game rather than applications. Figure 4.4 gives an extensive form representation (without payoffs) of a simple case: $T = \{t_1, t_2\}, M = \{m_1, m_2\}, A = \{a_1, a_2\}$, and $Prob\{t_1\} = p$.

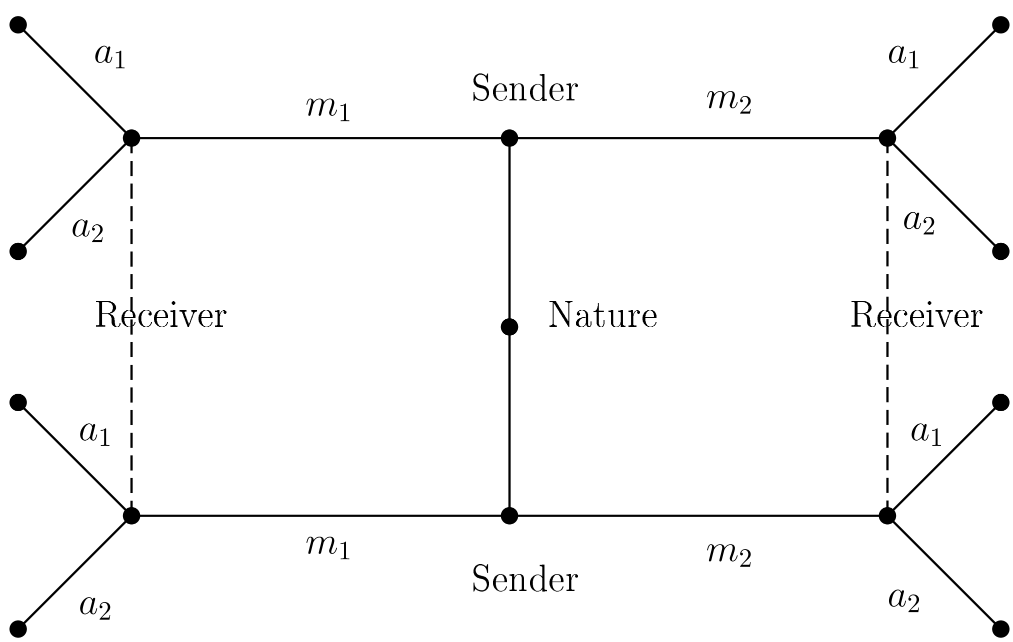

Figure 4.4:

Recall that in any game, a player's strategy is a complete plan of action– a strategy specifies a feasible action in every contingency in which the player might be called upon to act. In a signaling game, therefore, a pure strategy for the Sender is a function $m(t_i)$ specifying which message will be chosen for each type that nature might draw, and a pure strategy for the Receiver is a function $a(m_j)$ specifying which action will be chosen for each message

The sender then has the following four strategies:

1. Play m_1 if nature draws t_1 and play m_1 if nature draws t_2
2. Play m_1 if nature draws t_1 and play m_2 if nature draws t_2
3. Play m_2 if nature draws t_1 and play m_1 if nature draws t_2

4. Play m_2 if nature draws t_1 and play m_2 if nature draws t_2

The receiver has the following four strategies:

1. Play a_1 if the Sender chooses m_1, and play a_1 if the Sender chooses m_2
2. Play a_1 if the Sender chooses m_1, and play a_2 if the Sender chooses m_2
3. Play a_2 if the Sender chooses m_1, and play a_1 if the Sender chooses m_2
4. Play a_2 if the Sender chooses m_1, and play a_2 if the Sender chooses m_2

We call the sender's first and fourth strategies *pooling* because each type sends the same message, and the second and third *separating* because each type sends a different message. In a model with more than two types there are also *partially pooling* (or *semi-separating*) strategies in which all the types in a given set of types send the same message but different sets of types send different messages. In the two-type game in figure **??**, there are analogous mixed strategies, called *hybrid strategies* in which we say t_1 plays m_1 but t_2 randomizes between m_1 and m_2.

We now end up translating the informal statements of Requirements 1 through 3 into a formal definition of a perfect Bayesian equilibrium in a signaling game. To keep things simple, we only pay attention to pure strategies; hybrid strategies will come in the next section. Because the sender knows the full history of the game when choosing a message, this choice occurs at a singleton information set. Thus, requirement 1 is trivial when applied to the sender. The receiver in contrast, chooses an action after observing the sender's message but without knowing the sender's type, so the receiver's choice occurs at a nonsingleton information set. We get the following requirement:

1. After observing any message m_j from M, the receiver must have a belief about which types could have set m_j. Denote this belief by the probability distribution $\mu(t_i|m_j)$ where $\mu(t_i|m_j \geq 0$ for each $t)i$ in T, and

$$\sum_{t_i \in T} \mu(t_i|m_j) = 1$$

Applying requirement 2 to the receiver yields the following for the receiver:

2 For each m_j in M, the receiver's actions $a^*(m_j)$ must maximize the receiver's expected utility, given the belief $\mu(t_i|m_j)$ about which types could have sent m_j. That is, $a^*(m_j)$ solves

$$max_{a_k \in A} \mu(t_i \mid m_j) U_R(t_i, m_j, a_k)$$

Requirement 2 also applies to the sender, but the sender has complete information and only moves at the beginning of the game, so we have the next requirement:

3 For each $t_i \in T$, the Sender's message $m^*(t_i)$ must maximize the sender's utility, given the Receiver's strategy $a^*(m_j)$. That is, $m^*(t_i)$ solves

$$max_{m_j \in M} U_S(t_i, m_j, a^*(m_j))$$

Finally, for messages on the equilibrium path, applying the third requirement to the receiver's beliefs gives us:

4 For each $m_j \in M$, if there exists t_i
inT such that $m^*(t_i) = m_j$, then the receiver's belief at the information set corresponding to m_j must follow from Baye's rule and the Sender's strategy:

$$\mu)t_i \mid m_j) = \frac{p(t_i)}{\sum_{t_i \in T_i} p(t_i)}$$

Definition

A pure strategy **Bayesian equilibrium** in a signaling game is a pair of strategies $m^*(t_i)$ and $a^*(m_j)$ and a belief $\mu(t_i|m_j)$ satisfying Signaling requirements 1-4.

4.3 Job–Market Signaling

We now turn our attention to a game with the following timing:

1. Nature determines a worker's productivity ability ρ, which can be either high H or low L. The probability that $\rho = H$ is q.
2. The worker learns his own ability and then chooses a level of education, $e \geq 0$.
3. Two firms observe the workers education and then simultaneously make wage offers to the worker
4. The worker accepts the higher of the two wage offers, flipping a coin in case of a tie. Let w denote the wage the worker accepts.

The payoffs are then $w - c(\rho, e)$ to the worker, where $c(\rho, e)$ is the cost to a worker with ability ρ of obtaining education e; and $y(\rho, e) - w$ to the firm that employs the worker, where $y(\rho, e)$ is the output of a worker with ability ρ that has obtained education e, and zero to the firm that doesn't employ the worker.

We can make lots of assumptions based on other peoples models- we will use Spence's model, which assumes critically that low-ability workers find signaling more costly than do high-ability workers. More precisely, the marginal cost of education is higher for low than for high-ability workers: for every e,

$$c_e(L, e) > c_e(H, e)$$

where $c_e(\rho, e)$ denotes the marginal cost of education for a worker of ability ρ at education e. To interpret this assumption, consider a worker with education e_1, who is paid w_1, as depicted in figure 4.5.

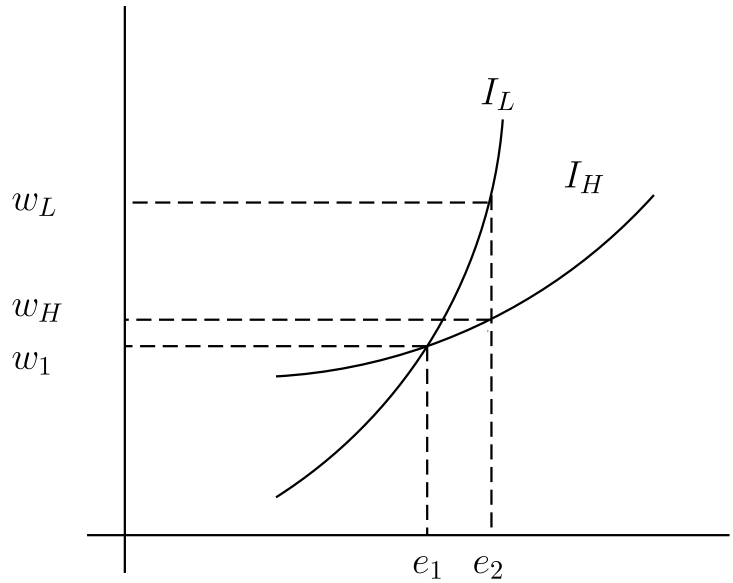

Figure 4.5:

An calculate the increase in wags that would be necessary to compensate this worker for an increase in education from e_1 to e_2. The answer depends on the worker's ability: low ability workers find it more difficult to acquire the extra education, and so require a larger in crease in wages.

Spence also assumes that competition among firms will drive expected profits to zero. One way to build this assumption into our model would be to replace the two firms in stage (3) with a single player called the market, that makes a single wage offer w and has the payoff $-[y(\rho, e) - w]^2$. To maximize its expected payoff, as required by our signaling requirements, the market would offer a wage equal to the expected output of a worker with education e, given the market's belief about the worker's ability after observing e:

$$w(e) = \mu(H|e) \cdot y(H, e) + [1 - \mu(H|e)] \cdot y(L, e)$$

To prepare for the analysis of the perfect Bayesian equilibria of this signaling game, we first consider the complete-information analog of the game. That is, we assume that the worker's ability is common knowledge among all the players, rather than privately known by the worker. In this case, competition between the two firms implies that a worker of ability ρ with education e picks the wage $w(e) = y(\rho, e)$, and a worker with ability ρ therefore picks e to solve

$$max_e y(\rho, e) - c(\rho, e)$$

We denote this solution by $e^*(\rho)$, as shown in figure 4.6 and let $w^*(\rho) = y[\rho, e^*(\rho)]$.

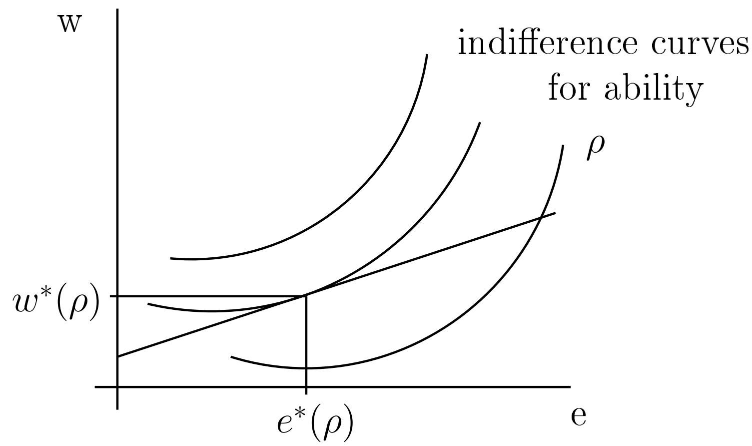

Figure 4.6:

We now return to the assumption that the worker's ability is private information. This opens the possibility that a low-ability worker could try to masquerade as a high-ability worker. Similarly, the low-ability worker could be said to envy the high-ability worker's complete-information wage and education level. The latter case is both more realistic and interesting. In a model with more than two values of the worker's ability, the former case arises only if each possible value of ability is sufficiently different from the adjacent possible values. If ability is a continuous variable, then the latter case applies.

As described, three kinds of perfect Bayesian equilibria can exist in this model: pooling, separating, and hybrid. Each kind typically exists in profusion, we restrict our attention to just a few examples. In a pooling equilibrium both worker types choose a single level of education, say e_p. Signaling requirement 3 then implies that the firm's belief after observing p_e must be the prior belief $\mu(H|e_p) = q$, which in turn implies that the wage offered after observing e_p is:

$$w_p = q \cdot y(H, e_p) + (1 - q) \cdot y(L, e_p)$$

To complete the description of a pooling perfect Bayesian equilibrium, it remains to

1. Specify the firm's belief for out of equilibrium education choices
2. Show that both worker types best response to the firm's strategy $w(e)$ is to choose $e = e_p$

One possibility is that the firms belief that any education level other than e_p implies that the worker has low ability. Although this sounds strange, nothing in the definition of Bayesian equilibrium rules it out. If the firms belief is:

$$\mu(H|e) = \begin{cases} 0 & \text{for } e \neq e_p \\ q & \text{for } e = e_p \end{cases}$$

then this implies that the firm's strategy is

$$w(e) = \begin{cases} y(L, e) & \text{for } e \neq e_p \\ w_p & \text{for } e = e_p \end{cases}$$

A worker of ability ρ then chooses e to solve

$$max_e w(e) - c(\rho, e)$$

solving this equation is actually pretty simple, a worker of ability ρ chooses either e_p or the level of education that maximizes $y(, e) - c(\rho, e)$.

4.4 Corporate Investment and Capital Structure

Suppose you have an entrepreneur that has a new company but needs outside funding to actually undergo a new project. The entrepreneur has private information about the profitability of the existing company, but the payoff of the new project can't be removed from the payoff of the existing company- all that can be observed is the aggregate profit of the firm. Suppose the entrepreneur offers a potential investor an equity stake in the firm in exchange for the necessary financing. Under what circumstances will the new project be undertaken, and what will this equity stake be?

We can turn this game into a signaling game in the following way:

1. Nature determines the profit of the existing company. The probability that $\pi = L$ is p.
2. The entrepreneur learns π and then offers the potential investor and equity stake s, where $0 \leq s \leq 1$.
3. The investor observes s and then decides whether or not to accept the offer.
4. If the investor rejects the offer then the investor's payoff is $I(1+r)$ and the entrepreneur's payoff is π. If the investor accepts s, then the investor's payoff is $s(\pi R)$ and the entrepreneur's is $(1-s)(\pi + R)$.

This is a nice game from two aspects: The receiver's set of feasible actions is extremely limited, and the sender's set of feasible signals is larger but still rather ineffective. suppose that after getting the offer s the investor thinks that the probability that $\pi = L$ is q. Then the investor will accept s if and only if

$$s[q(1-q)H + R] \geq I(1+r)$$

As for the entrepreneur, suppose that the profit of the existing company is π and consider whether the entrepreneur prefers t receive the financing at the cost of an equity stake of s or to leave the project. The former is preferable if and only if

$$s \leq \frac{R}{\pi + R}$$

In a pooling perfect Bayesian equilibrium, the investor's belief must be $q = p$ after getting this offer. Since the participation constraint is harder to satisfy for $\pi = H$ than for $\pi = L$, a pooling equilibrium exists if and only if

$$\frac{I(1+r)}{pL + (1-p)H + R} \leq \frac{R}{H + R}$$

It first looks as though the high-profit type must subsidize the low-profit type. We can set $q = p$ in our first relation and get

$$s \geq I(1+r)/[pL + (1-p)H + R]$$

In which case if the investor were certain that $\pi = H$ he would take the smaller equity stake $s \geq I(1+r)/(H+R)$. The larger stake in a pooling equilibrium is very expensive for the high-profit firm, and is perhaps so expensive that it makes the firm forego the new project. If this equation fails however, then a pooling equilibrium does not exist.

4.5 Other Applications of Perfect Bayesian Equilibrium

4.5.1 Cheap-Talk Games

Cheap-talk games are analogous to signaling games, but in cheap talk games the Sender's messages are just talk. Cheap talk can't be informative in Spence's model because all the Sender's types have the same preferences over the Receiver's possible actions: all workers prefer higher wages, independent of ability. To see why such uniformity of preferences across Sender-types vitiates cheap talk, suppose that there were a pure-strategy equilibrium in which one subset of Sender-types T_1 sends one message m_1 while another subset of types T_2 sends another message m_2. In equilibrium, the receiver will interpret m_i as coming from T_i and will take the optimal action given this belief, denote this action by a_i. Since all Sender-types have the same preferences over actions, if one prefers a_1 to a_2, then all types have this preference and will send m_1 instead of m_2, destroying the putative equilibrium.

So, for these cheap-talk situations, information transmission involves not only a simple cheap-talk game but also a more complicated version of an economic environment. The timing of the simplest cheap–talk game is identical to the timing of the simplest signaling game, with a few tweaks:

1. Nature draws a type t_i for the Sender from a set of feasible types $T = \{t_1, ..., t_I\}$ according to a probability distribution $p(t_i)$, where $p(t_i) > 0$ for each i and $p(t_1) + ... + p(t_I) = 1$

2. The sender observes t_i, and then chooses a message m_j from the set of feasible messages $M = \{m_1, ..., m_j\}$,

3. The receiver observes m_j and then chooses an action a_k from a set of feasible actions $A = \{a_1, ..., a_K\}$

4. Payoffs are given by $U_S(t_i, a_k)$ and $U_R(t_i, a_k)$

One of the key features of such a game is that the message doesn't have a direct effect on either the Sender or Receiver's payoff. The only way the message can matter is through its information content: by changing the receiver's belief about the Sender's type, a message can change the Receiver's action, and then indirectly affect both player's payoffs.

Because the simplest cheap-talk and signaling games have the same timing, the definitions of perfect Bayesian equilibrium in the two games are identical as well:

Definition

A pure-strategy Bayesian equilibrium in a cheap-talk game is a pair of strategies $m^*(t_i)$ and $a^*(m_j)$ and a belief $\mu(t_i|m_j)$ satisfying the signaling requirements. One difference between signaling and cheap-talk games however is that in the latter a pooling equilibrium always exists. Because messages have no effect (directly) on the Sender's payoff, if the receiver will ignore all messages then pooling is a best response for the sender; if the sender is pooling then the best response fro the receiver is to ignore all messages. Formally, let a^* denote the Receiver's optimal action in a pooling equilibrium, that is, a^* solves

$$max_{a_k \in A} \sum_{t_i \in T_i} p(t_i) U_R(t_i, a_k)$$

As an example, consider the following game: let $T = \{t_L, t_H\}$, $Prob(t_L) = p$, and $A = \{a_L, a_H\}$. The payoffs from the type action pair (t_i, a_k) are independent of which message is chosen, so we can describe the payoffs of this game with the figure 4.7.

	t_L	t_H
a_L	x,1	y,0
a_H	z,0	w,1

Figure 4.7:

The first payoff in each cell is the Sender's, and the second the Receiver's, but this figure i not a normal form game- rather it simply lists the player's payoffs from each type-action pair. We have chosen the Receiver's payoffs so that the Receiver prefers the low action (a_L) when the Sender's type is low, and high action when the type is high. To illustrate the first necessary condition, suppose that both Sender-types have the same preferences of interaction $x > z$ and $y > w$, for example so that both types prefer a_L to a_H. Then both types would like the receiver to think that $t = t_L$, so the Receiver can't believe such a claim. To illustrate the third condition, suppose that the player's preferences are completely opposed: $z > x$ and $y > w$. Then t_L would like the receiver to think that $t = t_H$ and t_H would like the Receiver to think that $t = t_L$, so the Receiver can't believe either of these claims. In this two-type two-action game, the only case that satisfies both the first and the third necessary conditions is $x \geq z$ and $y \geq w$ – the player's interests are perfectly aligned, in the sense that given the Sender's type the players agree on which action should be taken. Formally, in a separating perfect Bayesian equilibrium in this cheap-talk game, the Sender's strategy is $[m(t_L) = t_L, m(t_H) = t_H]$, the Receiver's beliefs are $\mu(t_L|t_L) = 1$ and $\mu(t_L|t_H) = 0$, and the Receiver's strategy is $[a(t_L) = a_L, a(t_H) = a_H]$. For these strategies and beliefs to be an equilibrium, each Sender type t_i must prefer to announce the truth, thereby inducing the action $a)i$ rather than to lie, thereby inducing a_j. Thus, a separating equilibrium exists if and only if $x \geq z$ and $y \geq w$.

4.6 Sequential Bargaining under Asymmetric Information

Think of a firm and union bargaining over wages. Assume that employment is fixed. The union's reservation wage, the amount that union members earn if not employed by the firm, is w_r. The firm's profit, π, is uniformly distributed on $[\pi_L, \pi_H]$, but the true value of π is privately known by the firm. Such private information might reflect the firm's superior knowledge about new products in the planning stage, for example. To make things simple, we assume that $w_r = \pi_L = 0$.

The bargaining game has two periods. In the first period, the union makes a wage offer, w_1. If the firm accepts this offer, then the game ends, the union's payoff is w_1 and the firm's is $\pi - w_1$. If the firm rejects this offer, then the union makes a second offer w_2. If the firm accepts this offer then the present value of he player's payoffs are δw_2 for the union and $\delta(\pi - w_2)$ for the firm, where δ is a constant reflecting both the discounting and reduced life of the contract remaining after the first period.

Defining a perfect Bayesian equilibrium is hard in this model, but we can start by looking at the union's first-period wage, offer, which is:

$$w_1^* = \frac{(2-\delta)^2}{2(4-3\delta)}\pi_H$$

If the firm's profit π exceeds the following:

$$\pi_1^* = \frac{2_{w_1}}{2-\delta} = \frac{2-\delta}{4-3\delta}\pi_H$$

then the firm accepts w_1^*, otherwise the firm rejects w_1^*. If its first-period offer is rejected, then the union changes its belief about the firm's profit: the union believes that π is uniformly distributed on $[0, \pi_!^*]$. The union's second-period wage offer is then

$$w_2^* = \frac{\pi_1^*}{2} = \frac{2-\delta}{2(4-3\delta)}\pi_H < w_1^*$$

If the firm's profit π is greater than w_2^*, then the firm accepts the offer- otherwise, it rejects it.

Figure 4.8 provides an extensive-form representation of a simplified version of this game.

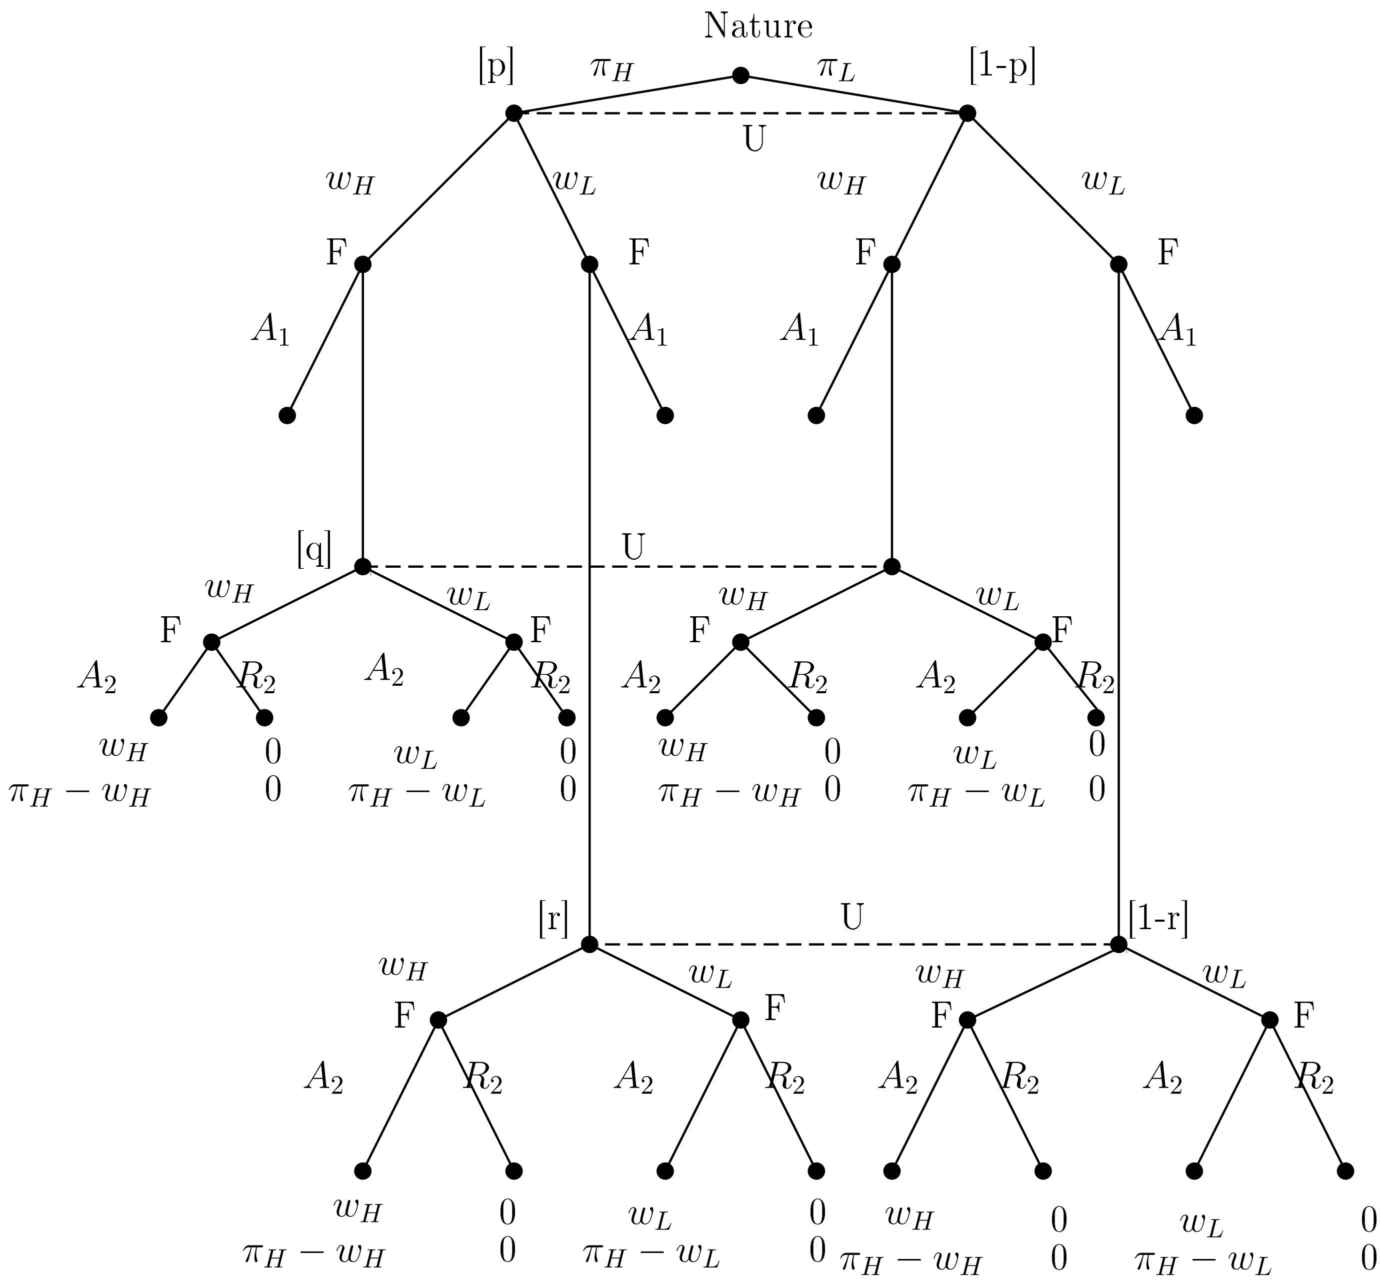

Figure 4.8:

In this game, the union has the move at three information sets, so the union's strategy consists of three wage offers: the first is w_1, the other two come in response after $w_1 = w_H$ and $w_1 = w_L$ are rejected. These three moves occur at three non-singleton information sets, at which the union's beliefs are denoted $(p, 1-p)$ $(q, 1-q)$, and $(r, 1-r)$ respectively.

In the full game, the a strategy for the union is a first-period offer w_1 and a second-period information set for each different first-period wage offer the union might take. Within both the lone first-period and the continuum of second-period information sets, there is one decision node for each possible value of $]pi$. At each information set, the union's belief is a probability distribution over these nodes. In the full game, we denote the union's first period belief by $\mu_1(\pi)$, and the union's second period belief by $\mu_2(\pi|w_1)$. A strategy for the firm involves two decisions. Let $A_1(w_1|\pi)$ be 1 if the firm would accept the first-period offer w_1 when its profit is π, and zero if it would reject it. Likewise, we do the same for $A_2(w_2|\pi, w_1)$. A strategy for the firm is then a pair of functions:

$$[A_1(w_1|\pi), A_2(w_2|\pi, w_1)]$$

and the beliefs $[\mu_1(\pi), \mu_2(\pi|w_1)]$ are a perfect Bayesian equilibrium if they satisfy the requirements given in an earlier section. We can show that there is a unique perfect Bayesian equilibrium, and we begin by considering the following one-period problem: suppose that the union thinks that the firm's profit is uniform over $[0, \pi_1]$, where for the moment π_1 is arbitrary. If the union offers w then the firm's best response is clear: accept w if and only if $\pi \geq w$. So, we have the following:

$$max_w w \cdot Prob\{ \text{ firm accepts } w \ \} + 0 \cdot Prob\{ \text{ firm rejects } w \ \}$$

Where $Prob\{\text{firm accepts } w\} = (\pi_1 - w)/\pi_1$. Going back to the two-period problem, for arbitrary values of w_1 and w_2 it turns out that if the union offers w_1 in the first period and the firm expects the union to offer w_2 in the second period, then all firms with sufficiently high profit will accept w_1 and all others will reject it. The firm's possible payoffs would then be $\pi - w_1$ from accepting w_1, and $\delta(\pi - w_2)$ from rejecting w_1 and accepting w_2, and zero from denying both offers. The firm then prefers accepting w_1 to w_2 if $\pi - w_1 > \delta(\pi - w_2)$ or

$$\pi > \frac{w_1 - \delta w_2}{1 - \delta} \equiv \pi^*(w_1, w_2)$$

From this one can ten derive that the union's optimal second-period offer must be $w^*(\pi_1) = \pi_1/2$, which yields an implicit equation for π_1 as a function w_1:

$$\pi_1 = max\{\pi^*(w_1, \pi_1/2), w\}$$

which when solved, we eventually see that

$$Prob\{\text{firm accepts } w_1\} = \frac{\pi_H - \pi_1(w_1)}{\pi_H}$$

4.7 Reputation in the Finitely Repeated Prisoner's Dilemma

A simple example of a reputation equilibrium is in the finitely repeated Prisoner's Dilemma Rather than assume that one player has private information about his payoffs, we can assume that the player has private information about his feasible strategies. In looking at a simple example, we have the following timing:

1. Nature draws a type for the row player. With probability p, Row has only the 'Tit-for-Tat' strategy available; with probability $1 - p$ row can play any strategy. Row learns his or her type, but Column doesn't learn row's type.
2. Row and Column play the prisoner's dilemma. The player's choices in this stage game become common knowledge.
3. Row and Column play the prisoner's dilemma for a second and las time
4. Payoffs are received. The payoffs to the rational Row and to column are the sums of their stage game payoffs.

To make this into a Prisoner's Dilemma, we have figure 4.9.

		Column: Cooperate	Column: Fink
Row	Cooperate	1,1	b,a
	Fink	a,b	0,0

Figure 4.9:

Just like in the last period of a finitely repeated Prisoner's Dilemma under complete information, finking strictly dominates cooperating the last stage of this two period game of incomplete information both for the rational Row and for Column. Column will surely fink in the last stage, there is no reason why the rational Row should cooperate in the final stage. Finally, tit-for-tat begins the game by cooperating. Thus, the only move to be determined is Column's first-period move which is then mimicked in the second period.

By letting column's first-period move be cooperate,, column receives the expected payoff of $p \cdot 1 + (1-p) \cdot b$ in the first period and $p \cdot a$ in the second. In contrast, supposing he Finks, then Column receives $p \cdot a$ in the first period and zero in the second. So, Column will cooperate in the first period if and only if

$$p + (1 - p)b \geq 0$$

4.8 Refinements of Perfect Bayesian Equilibrium

In earlier sections, we defined a perfect Bayesian equilibrium to be strategies and beliefs that satisfied requirements 1-4, and we observed that in such an equilibrium no player's strategy can be strictly dominated beginning at any information set. We now consider two further requirements , the first of which formalizes the idea that since perfect Bayesian equilibrium prevents player i from playing a strategy that is strictly dominated beginning at any information set, it is not reasonable for player j to believe that i would play such a strategy.

Consider the game in figure 4.10. The key feature of this example is that

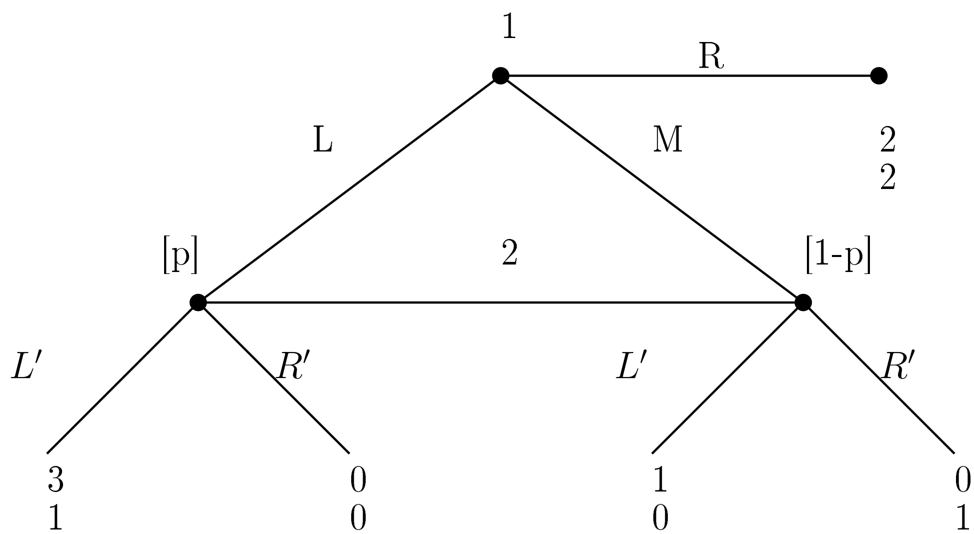

Figure 4.10:

M is a strictly dominated strategy for player 1- the payoff of 2 from R exceeds both of the payoffs that player 1 could receive from playing $M--0$ and 1. The two other things to notice with this example is that although M is strictly dominated, L is not. Also, this example doesn't illustrate the requirement described initially, because M is not just strictly dominated beginning at an information set but also strictly dominated. Since M is strictly dominated, it is certainly not reasonable for player 2 to believe that 1 might have played M, but strict dominance is too strong a test, and thus yields too weak a requirement. We end up with the following definition:

Definition

Consider an information set at which player i has the move. The strategy s_i^prime is **strictly dominated beginning at this information set** if there exists another strategy s_i such that for each belief that i could hold at the given information set, and for each possible combination of the other player's strategies, i's expected payoff from taking the action specified by s_i at the given information set and playing the subsequent strategy specified by s_i is strictly greater than the expected payoff from taking the action and playing the subsequent strategy specified by s_i'.

We thus end up at our fifth requirement:

5 If possible, each player's beliefs off the equilibrium path should place zero probability on nodes that are reached only if another player plays a strategy that is strictly dominated beginning at some information set.

An equivalent way to look at this requirement is by use of the following definition:

Definition

In a signaling game, the message m_j from M is **dominated for type** t_i from T if there is another message $m_{j'}$ from M such that t_i's lowest possible payoff from $m_{j'}$ is greater than t_i's highest possible pay off from $m_{j'}$:

$$min_{a_k \in A} U_S(t_i, m_{j'}, a_k) > max_{a_k \in A} U_S(t_i, m_j, a_k)$$

Which leads us to the following reworking of our fifth requirement:

5 If the information set following m_j is off the equilibrium path and m_j is dominated for type t_i, then the receiver's belief $\mu(t_i|m_j)$ should place zero probability on type t_i

Continuing in this line of thought, we can construct the following definition:

Definition

Given a perfect Bayesian equilibrium in a signaling game, the message m_j from M if **equilibrium-dominated for type** t_i from T if t_i's equilibrium payoff denoted $U^*(t_i)$ is greater than t_i's highest possible payoff from m_j:

$$U^*(t_i) > max_{a_k \in A} U_S(t_i, m_j, a_k)$$

Which brings us to our sixth and final requirement for perfect Bayesian equilibrium, with which we conclude this section:

6 If the information set following m_j is off the equilibrium path and m_j is equilibrium dominated for type t_i then (if possible) the Receiver's belief$\mu(t_i|m)$ should place zero probability on type t_i.