22 REPEATED GAMES AND REPUTATION

People often interact in ongoing relationships. For example, most employment relationships last a long time. Countries competing over tariff levels know that they will be affected by each others' policies far into the future. Firms in an industry recognize that they are not playing a static game but one in which they compete every day over time. In all of these dynamic situations, the way in which a party behaves at any given time is influenced by what this party and others did in the past. In other words, players *condition* their decisions on the history of their relationship. An employee may choose to work diligently only if his employer gave him a good bonus in the preceding month. One country may set a low import tariff only if its trading partners had maintained low tariffs in the past. A firm may elect to match its competitor's price by setting its price each day equal to the competitor's price of the preceding day.

When deciding how to behave in an ongoing relationship, one must consider how one's behavior will influence the actions of others in the future. Suppose I am one of your employees and that our history is one of cooperation: since you hired me, I have been a loyal and hard-working employee and you have given me a generous bonus each month (above my salary). Today I am considering whether to work hard or, alternatively, neglect my duties in favor of playing video games on the office computer. For me, shirking has an immediate reward—I get to avoid expending effort on the job. However, you will soon learn of my indolence, either through your monitoring activities or through an observed decrease in my productivity. *Your future behavior* (in particular, whether to give me bonuses each month) may very well be influenced by what *I do today*.

For instance, after observing that I have shirked, you might choose to discontinue my monthly bonus. You might say to yourself, "By misbehaving, Watson has lost my trust; I doubt that he will work diligently ever again and therefore I will pay him no more bonuses." Anticipating such a response, I may decide that spending the day playing chess over the Internet against someone in New Zealand is not such a good idea. Neglecting my duties may yield an immediate gain (relaxation today), but it leads to a greater loss in the future (no more bonuses each month). As this story suggests, people sometimes have an incentive to forego small immediate gains because of the threat of future retaliation by others.

The term "reputation" is often used to describe how a person's past actions affect future beliefs and behavior. If I have always worked diligently on the job, one would say that I have "established a reputation for being a hard worker." If I shirk today, then tomorrow people would say that I have "destroyed my good reputation." Often, those who nurture good reputations are trusted and rewarded; people with bad reputations are punished. As the employment story indicates, the concern for reputation may motivate parties to cooperate with one another, even if such behavior requires foregoing short-term gains. One of the great achievements of game theory is that it provides a framework for understanding how such a reputation mechanism can support cooperation.

The best way to study the interaction between immediate gains and long-term incentives is to examine a *repeated game*. A repeated game is played over discrete periods of time (period 1, period 2, and so on). We let t denote any given period and let T denote the total number of periods in the repeated game. T can be a finite number or it can be infinity, which means the players interact perpetually over time. In each period, the players play a static *stage game*, whereby they simultaneously and independently select actions. These actions lead to a stage-game payoff for the players. The stage game can be denoted by $\{A, u\}$, where $A = A_1 \times A_2 \times \cdots \times A_n$ is the set of action profiles and $u_i(a)$ is player i's stage-game payoff when profile a is played. The same stage game is played in each period. Furthermore, we assume that, in each period t, the players have observed the *history* of play—that is, the sequence of action profiles—from the first period through period $t - 1$. The payoff of the entire game is defined as the sum of the stage-game payoffs in periods 1 through T.[1]

A TWO-PERIOD REPEATED GAME

Suppose players 1 and 2 interact over two periods, called periods 1 and 2 (so $T = 2$). In each period, they play the stage game depicted in Figure 22.1 on the next page. Assume that the payoff for the entire game is the sum of the stage-game payoffs in the two periods. For instance, if (A, X) is played in the first period and (B, Y) is played in the second period, then player 1's payoff is $4 + 2 = 6$ and player 2's payoff is $3 + 1 = 4$. Figure 22.2(a) graphs the set of

[1] Sometimes a discount factor is used to represent the fact that the players discount the future. Recall the analysis of discounting in Chapter 19.

FIGURE 22.1
Stage game, repeated once ($T = 2$).

1 \ 2	X	Y	Z
A	4, 3	0, 0	1, 4
B	0, 0	2, 1	0, 0

FIGURE 22.2 Possible repeated game payoffs.

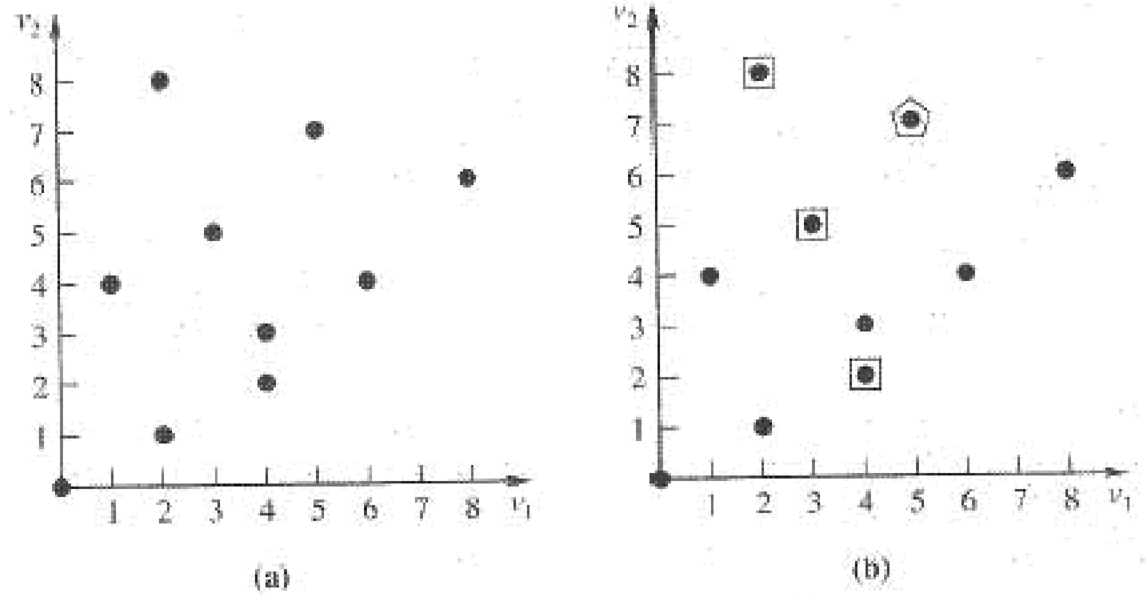

possible repeated game payoffs. Every point on the graph corresponds to the sum of two stage-game payoff vectors. For example, the payoff vector (3, 5) can be attained if (A, Z) is played in the first period and (B, Y) is played in the second period; the same payoff results if (B, Y) is played in the first period, followed by (A, Z).

This two-period repeated game has a large extensive-form representation, so I have not drawn it here. The extensive form starts with simultaneous selection of actions in the first period: player 1 chooses between A and B while player 2 selects X, Y, or Z. Then, having observed each other's first-period actions, the players again select from {A, B} and {X, Y, Z}. Because each player knows what happened in the first period, his choice in the second period can be conditioned on this information. For example, player 1 may decide to pick A in the second period if and only if (A, X) or (B, Y) was played in the first period; otherwise, he picks B. As usual, the players' information is represented by information sets. Because there are six possible outcomes of first-period interaction, each player has six information sets in the second period. In other words, each player has *six* decisions to make in period 2: what to do if the outcome of period 1 was (A, X), what to do if the outcome of period 1 was

FIGURE 22.3
The subgame following (A, Z).

1 \ 2	X	Y	Z
A	5, 7	1, 4	2, 8
B	1, 4	3, 5	1, 4

(A, Y), and so forth. It is complicated, isn't it? Do not worry—the analysis is actually very simple.

Just kidding. This game has a large set of strategies, making the analysis of rationality a bit daunting. However, the analysis is quite illuminating, so read on. Let us look for pure-strategy, subgame perfect Nash equilibria. We can simplify the search by recognizing one important thing: every equilibrium must specify that, in the *second* period, the players select an action profile that is a Nash equilibrium of the stage game.

To see why this is so, recall that subgame perfection requires equilibrium in *every subgame*. In the repeated game at hand, a different subgame is initiated following every different action profile in period 1. For example, consider what happens in the event that the players choose (A, Z) in the first period. Then, knowing that (A, Z) was the outcome of first-period interaction, the players proceed to a subgame in which they simultaneously select actions again. Their total payoff will be (1, 4) plus whatever the payoff vector is in the second play of the stage game. Thus, following the play of (A, Z) in the first period, the subgame is described by the matrix in Figure 22.3. I constructed this matrix by adding the payoff vector (1, 4) to each of the cells in the stage game (compare Figures 22.1 and 22.3). You should verify that the subgame has two Nash equilibria, (B, Y) and (A, Z). Therefore, a subgame perfect equilibrium must specify that either (B, Y) or (A, Z) be played in the second period if (A, Z) is played in the first period.

We can say more. Because the subgame matrix is formed by adding the same payoff vector to each cell of the stage-game matrix, the players' preferences over action profiles in the subgame *are exactly the same* as their preferences in an isolated play of the stage game. In other words, the subgame and the stage game have exactly the same Nash equilibria.[2] In fact, every subgame starting in period 2 has the same set of Nash equilibria, because these games have matrices like that in Figure 22.3—where the stage-game payoff in the first period is added to every cell of the stage game. Another way of

[2] You should verify that (A, Z) and (B, Y) are the Nash equilibria of the stage game.

thinking about this is that, once the first period is over, the payoffs from the first period are *sunk*. Whatever the players obtain in the second period, it is in addition to what they have already received in the first. Thus a subgame perfect equilibrium must specify that the players select a Nash equilibrium in the stage game in period 2, whatever happens in period 1. To save ink, I use the phrase "stage Nash profile" to refer to a Nash equilibrium in the stage game.

Knowing that a stage Nash profile will be played in the second period, we can turn our attention to two other matters: (1) action choices in the first period and (2) how behavior in the first period determines *which* of the two stage Nash equilibria will be played in the second period.

First consider subgame perfect equilibria that specify that a stage Nash profile be played in the first period (as well as in the second). Here is one such strategy profile: the players are instructed to choose action profile (A, Z) in the first period and then, regardless of the outcome of the first period, they are to choose (A, Z) in the second period. Because each player has six information sets in the second period (six potential decisions to make), the phrase "regardless of the outcome of the first period" is crucial; it means that, even if one or both of the players deviates from (A, Z) in the first period, they are supposed to play (A, Z) in the second. You can easily verify that this strategy profile is a subgame perfect equilibrium—neither player can gain by deviating in either or both periods, given the other player's strategy. In this equilibrium, player 1 obtains $1 + 1 = 2$ and player 2 gets $4 + 4 = 8$. This payoff vector is one of those boxed in Figure 22.2(b).

Any combination of stage Nash profiles can be supported as a subgame perfect equilibrium outcome. For example, "choose (A, Z) in the first period and then, regardless of the first-period outcome, choose (B, Y) in the second period" is a subgame perfect equilibrium; it yields the payoff vector (3, 5). The payoffs of equilibria that specify stage Nash profiles in both periods are all boxed in Figure 22.2(b). I recommend reviewing the various combinations of stage Nash profiles and verifying that the associated equilibrium payoffs are boxed in Figure 22.2(b). As the example intimates, the following general result holds:

> **Result:** Consider any repeated game. Any sequence of stage Nash profiles can be supported as the outcome of a subgame perfect Nash equilibrium.

It probably does not surprise you that stage Nash profiles can be supported as equilibrium play. A more interesting question is whether there are equilibria stipulating actions that are *not* stage Nash profiles. In fact, the answer is "yes," as the two-period example at hand illustrates. Consider the following strategy profile:

> Select (A, X) in the first period and then, as long as player 2 does not deviate from X, select (A, Z) in the second period; if player 2 deviated by playing Y or Z in the first period, then play (B, Y) in the second period.

This strategy profile prescribes that the players' second-period actions depend on what player 2 did in period 1. By playing X in the first period, player 2 establishes a reputation for cooperating; in this case, he is rewarded in the second period as the players coordinate on the stage Nash profile that is more favorable to him. On the other hand, if player 2 deviates by, say, choosing Z in the first period, then he is branded a "cheater." In this case, his punishment is that the players coordinate on (B, Y) in the second period.

To verify that this strategy profile is a subgame perfect equilibrium, we must check each player's incentives. Suppose player 1 behaves as prescribed and consider the incentives of player 2. If player 2 goes along with the strategy prescription, he obtains 3 in the first period and 4 in the second period. If player 2 deviates in the first period, he can increase his first-period payoff to 4 (by picking Z). However, this choice induces player 1 to select B in the second period, where player 2 then best responds with Y. Thus, although a first-period deviation yields an immediate gain of *1*, it costs *3* in the second period (4 − 1). This shows that player 2 prefers to behave as prescribed. For his part, player 1 has the incentive to go along with the prescription for play; deviating in either period reduces player 1's payoff. The payoff vector for this subgame perfect equilibrium is enclosed by a pentagon in Figure 22.2(b).

Although the equilibrium construction is a bit complicated, it really is intuitive. Player 2's concern about his reputation and what it implies for his second-period payoff gives him the incentive to forego a short-term gain. If he misbehaves in the first period, his reputation is destroyed and he then suffers in the second period.

Any two-period repeated game can be analyzed as has been done here.[3] Only stage Nash profiles can be played in the second period. However, sometimes reputational equilibria exist whereby the players select nonstage Nash profiles in the first period. These selections are supported by making the second-period actions contingent on the outcome in the first period (in particular, whether the players cheat or not). The exercises at the end of this chapter will help you better explore the reputation phenomenon.

---

[3] A general analysis of finitely repeated games is reported in "Finitely Repeated Games," by J. P. Benoit and V. Krishna, *Econometrica* 53(1985):905–922.