

Chapter 4

dynamic games with symmetric information



4.1 Subgame Perfectness

In this chapter we will make heavy use of the extensive form to study games with moves that occur in sequence. We start in section 4.1 with a refinement of the Nash equilibrium concept called perfectness that incorporates sensible implications of the order of moves. Perfectness is illustrated in section 4.2 with a game of entry deterrence. Section 4.3 expands on the idea of perfectness using the example of nuisance suits, meritless lawsuits brought in the hopes of obtaining a settlement out of court. Nuisance suits show the importance of a threat being made credible and how sinking costs early or having certain nonmonetary payoffs can benefit a player. This example will also be used to discuss the open-set problem of weak equilibria in games with continuous strategy spaces, in which a player offering a contract chooses its terms to make the other player indifferent about accepting or rejecting. The last perfectness topic will be renegotiation: the idea that when there are multiple perfect equilibria, the players will coordinate on equilibria that are Pareto optimal in subgames but not in the game as a whole.

The Perfect Equilibrium of Follow-the-Leader I

Subgame perfectness is an equilibrium concept based on the ordering of moves and the distinction between an equilibrium path and an equilibrium. The **equilibrium path** is the path through the game tree that is followed in equilibrium, but the equilibrium itself is a strategy profile, which includes the players' responses to other players' deviations from the equilibrium path. These off-equilibrium responses are crucial to decisions on the equilibrium path. A threat, for example, is a promise to carry out a certain action if another player deviates from his equilibrium actions, and it has an influence even if it is never used.

Perfectness is best introduced with an example. In section 2.1, a flaw of Nash equilibrium was revealed in the game Follow-the-Leader I, which has three pure strategy Nash equilibria of which only one is reasonable. The players are Smith and Jones, who choose disk sizes. Both their payoffs are greater if they choose the same size and greatest if they

coordinate on *Large*. Smith moves first, so his strategy set is $\{Small, Large\}$. Jones' strategy is more complicated, because it must specify an action for each information set, and Jones' information set depends on what Smith chose. A typical element of Jones' strategy set is $(Large, Small)$, which specifies that he chooses *Large* if Smith chose *Large*, and *Small* if Smith chose *Small*. From the normal form we found the following three Nash equilibria.

Equilibrium	Strategies	Outcome
E_1	$\{Large, (Large, Large)\}$	Both pick <i>Large</i> .
E_2	$\{Large, (Large, Small)\}$	Both pick <i>Large</i> .
E_3	$\{Small, (Small, Small)\}$	Both pick <i>Small</i> .

Only Equilibrium E_2 is reasonable, because the order of the moves should matter to the decisions players make. The problem with the normal form, and thus with simple Nash equilibrium, is that it ignores who moves first. Smith moves first, and it seems reasonable that Jones should be allowed – in fact should be required – to rethink his strategy after Smith moves.

Consider Jones's strategy of $(Small, Small)$ in equilibrium E_3 . If Smith deviated from equilibrium by choosing *Large*, it would be unreasonable for Jones to stick to the response *Small*. Instead, he should also choose *Large*. But if Smith expected a response of *Large*, he would have chosen *Large* in the first place, and E_3 would not be an equilibrium. A similar argument shows that it would be irrational for Jones to choose $(Large, Large)$, and we are left with E_2 as the unique equilibrium.

We say that equilibria E_1 and E_3 are Nash equilibria but not “perfect” Nash equilibria. A strategy profile is a perfect equilibrium if it remains an equilibrium on all possible paths, including not only the equilibrium path but all the other paths, which branch off into different “subgames.” A player's perfect equilibrium strategy is thus a best response to the other players' equilibrium strategies not only on the equilibrium path, but “**out of equilibrium**” or “**off the equilibrium path.**”

*A subgame is a game consisting of a node which is a singleton in every player's information partition, that node's successors, and the payoffs at the associated end nodes.*¹

A strategy profile is a subgame perfect Nash equilibrium if (a) it is a Nash equilibrium for the entire game; and (b) its relevant action rules are a Nash equilibrium for every subgame.

The extensive form of Follow the Leader I in figure 4.1 (a reprise of figure 2.1) has three subgames: (a) the entire game, (b) the subgame starting at node J_1 , and (c) the subgame starting at node J_2 . Strategy profile E_1 is not a subgame perfect equilibrium because it is only Nash in subgames (a) and (c), not in subgame (b). Strategy profile E_3 is not a subgame perfect equilibrium because it is only Nash in subgames (a) and (b), not in subgame (c). Strategy profile E_2 is perfect because it is Nash in all three subgames.

¹ Technically, this is a *proper* subgame because of the information qualifier, but no economist is so ill-bred as to use any other kind of subgame.

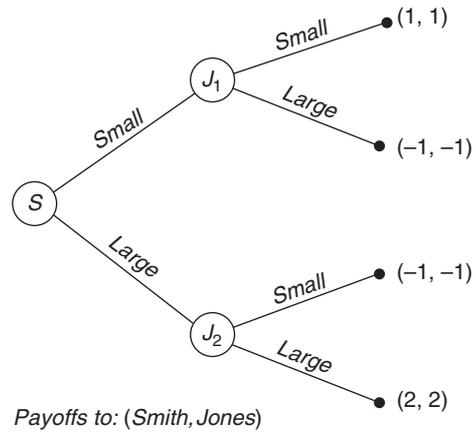


Figure 4.1 Follow-the-Leader I.

The term **sequential rationality** is often used to denote the idea that a player should maximize his payoffs at each point in the game, reoptimizing his decisions at each point and taking into account the fact that he will reoptimize in the future. This is a blend of the economic ideas of ignoring sunk costs and rational expectations. Sequential rationality is so standard a criterion for equilibrium now that often I will speak of “equilibrium” without the qualifier when I wish to refer to an equilibrium that satisfies sequential rationality in the sense of being a “subgame perfect equilibrium” or, in a game of asymmetric information, a “perfect Bayesian equilibrium.”

One reason why perfectness (the word “subgame” is usually left off) is a good equilibrium concept is because it represents the idea of sequential rationality. A second reason is that a weak Nash equilibrium is not robust to small changes in the game. So long as he is certain that Smith will not choose *Large*, Jones is indifferent between the never-to-be-used responses (*Small* if *Large*) and (*Large* if *Large*). Equilibria E_1 , E_2 , and E_3 are all weak Nash equilibria because of this. But if there is even a small probability that Smith will choose *Large* – perhaps by mistake – then Jones would prefer the response (*Large* if *Large*), and equilibria E_1 and E_3 are no longer valid. Perfectness is a way to eliminate some of these less robust weak equilibria. The small probability of a mistake is called a **tremble**, and section 6.1 returns to this **trembling hand** approach as one way to extend the notion of perfectness to games of asymmetric information.

For the moment, however, the reader should note that the tremble approach is distinct from sequential rationality. Consider figure 4.2’s Tremble Game. This game has three Nash equilibria, all weak: (*Out*, *Down*), (*Out*, *Up*), and (*In*, *Up*). Only (*Out*, *Up*) and (*In*, *Up*) are subgame perfect, because although *Down* is weakly Jones’s best response to Smith’s *Out*, it is inferior if Smith chooses *In*. In the subgame starting with Jones’s move, the only subgame perfect equilibrium is for Jones to choose *Up*. The possibility of trembles, however, rules out (*In*, *Up*) as an equilibrium. If Jones has even an infinitesimal chance of trembling and choosing *Down*, Smith will choose *Out* instead of *In*. Also, Jones will choose *Up*, not *Down*, because if Smith trembles and chooses *In*, Jones prefers *Up* to *Down*. This leaves only (*Out*, *Up*) as an equilibrium, despite the fact that it is weakly Pareto dominated by (*In*, *Up*).

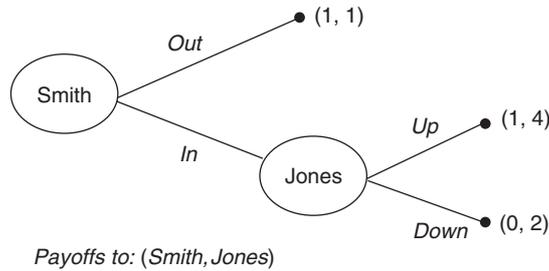


Figure 4.2 The tremble game: trembling hand versus subgame perfectness.

4.2 An Example of Perfectness: Entry Deterrence I

We turn now to a game in which perfectness plays a role just as important as in Follow-the-Leader-I but in which the players are in conflict. An old question in industrial organization is whether an incumbent monopolist can maintain his position by threatening to wage a price war against any new firm that enters the market. This idea was heavily attacked by Chicago School economists such as McGee (1958) on the grounds that a price war would hurt the incumbent more than collusion with the entrant. Game theory can present this reasoning very cleanly. Let us consider a single episode of possible entry and price warfare, which nobody expects to be repeated. We will assume that even if the incumbent chooses to collude with the entrant, maintaining a duopoly is difficult enough that market revenue drops considerably from the monopoly level.

Entry Deterrence I

PLAYERS

Two firms, the entrant and the incumbent.

THE ORDER OF PLAY

- 1 The entrant decides whether to *Enter* or *Stay Out*.
- 2 If the entrant enters, the incumbent can *Collude* with him, or *Fight* by cutting the price drastically.

PAYOFFS

Market profits are 300 at the monopoly price and 0 at the fighting price. Entry costs are 10. Duopoly competition reduces market revenue to 100, which is split evenly.

The strategy sets can be discovered from the order of play. They are $\{Enter, Stay Out\}$ for the entrant, and $\{Collude \text{ if entry occurs}, Fight \text{ if entry occurs}\}$ for the incumbent. The game has the two Nash equilibria indicated in boldface in table 4.1, $(Enter, Collude)$ and $(Stay Out, Fight)$. The equilibrium $(Stay Out, Fight)$ is weak, because the incumbent would just as soon *Collude* given that the entrant is staying out.

Table 4.1 Entry Deterrence I

		Incumbent	
		<i>Collude</i>	<i>Fight</i>
Entrant	<i>Enter</i>	40, 50 ←	-10, 0
	<i>Stay Out</i>	0, 300 ↑	0, 300 ↓

Payoffs to: (Entrant, Incumbent). Arrows show how a player can increase his payoff.

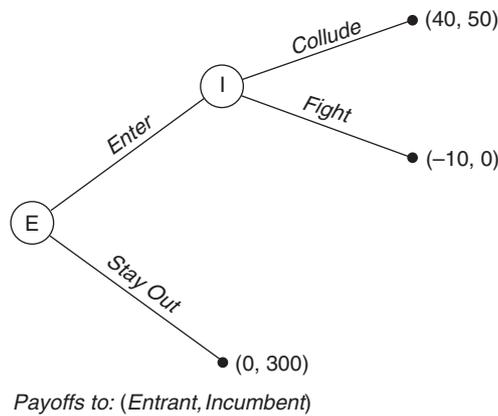


Figure 4.3 Entry deterrence I.

A piece of information has been lost by condensing from the extensive form, figure 4.3, to the normal form, table 4.1: the fact that the entrant gets to move first. Once he has chosen *Enter*, the incumbent’s best response is *Collude*. The threat to fight is not credible and would be employed only if the incumbent could bind himself to fight, in which case he never does fight, because the entrant chooses to stay out. The equilibrium (*Stay Out, Fight*) is Nash but not subgame perfect, because if the game is started after the entrant has already entered, the incumbent’s best response is *Collude*. This does not prove that collusion is inevitable in duopoly, but it is the equilibrium for Entry Deterrence I.

The trembling hand interpretation of perfect equilibrium can be used here. So long as it is certain that the entrant will not enter, the incumbent is indifferent between *Fight* and *Collude*, but if there were even a small probability of entry – perhaps because of a lapse of good judgement by the entrant – the incumbent would prefer *Collude* and the Nash equilibrium would be broken.

Perfectness rules out threats that are not credible. Entry Deterrence I is a good example because if a communication move were added to the game tree, the incumbent might tell the entrant that entry would be followed by fighting, but the entrant would ignore this noncredible threat. If, however, some means existed by which the incumbent could precommit himself to fight entry, the threat would become credible. The next section will look at one context, nuisance lawsuits, in which such precommitment might be possible.

Should the Modeller Ever Use Nonperfect Equilibria?

A game in which a player can commit himself to a strategy can be modelled in two ways:

- 1 As a game in which nonperfect equilibria are acceptable, or
- 2 By changing the game to replace the action *Do X* with *Commit To Do X* at an earlier node.

An example of (2) in Entry Deterrence I is to reformulate the game so the incumbent moves first, deciding in advance whether or not to choose *Fight* before the entrant moves. Approach (2) is better than (1) because if the modeller wants to let players commit to some actions and not to others, he can do this by carefully specifying the order of play. Allowing equilibria to be nonperfect forbids such discrimination and multiplies the number of equilibria. Indeed, the problem with subgame perfectness is not that it is too restrictive but that it still allows too many strategy profiles to be equilibria in games of asymmetric information. A subgame must start at a single node and not cut across any player's information set, so often the only subgame will be the whole game and subgame perfectness does not restrict equilibrium at all. Section 6.1 discusses perfect Bayesian equilibrium and other ways to extend the perfectness concept to games of asymmetric information.

4.3 Credible Threats, Sunk Costs, and the Open-set Problem in the Game of Nuisance Suits

Like the related concepts of sunk costs and rational expectations, sequential rationality is a simple idea with tremendous power. This section will show that power in another simple game, one which models nuisance suits. We have already come across one application of game theory to law, in the Png (1983) model of section 2.5. In some ways, law is particularly well suited to analysis by game theory because the legal process is so concerned with conflict and the provision of definite rules to regulate that conflict. In what other field could an article be titled "An Economic Analysis of Rule 68," as Miller (1986) does in his discussion of the federal rule of procedure that penalizes a losing litigant who had refused to accept a settlement offer. The growth in the area can be seen by comparing the overview in the Ayres's (1990) review of the first edition of the present book with the entire book by Baird, Gertner, & Picker (1994). In law, even more clearly than in business, a major objective is to avoid inefficient outcomes by restructuring the rules, and nuisance suits are one of the inefficiencies that a good policy maker hopes to eliminate.

Nuisance suits are lawsuits with little chance of success, whose only possible purpose seems to be the hope of a settlement out of court. In the context of entry deterrence people commonly think large size is an advantage and a large incumbent will threaten a small entrant, but in the context of nuisance suits people commonly think large size is a disadvantage and a wealthy corporation is vulnerable to extortionary litigation. Nuisance Suits I models the essentials of the situation: bringing suit is costly and has little chance of success, but because defending the suit is also costly the defendant might pay a generous amount to settle it out of court. The model is similar to the Png Settlement Game of chapter 2 in many respects, but here the model will be one of symmetric information and we will

make explicit the sequential rationality requirement that was implicit in the discussion in chapter 2.

Nuisance Suits I: Simple Extortion

PLAYERS
A plaintiff and a defendant.

THE ORDER OF PLAY

- 1 The plaintiff decides whether to bring suit against the defendant at cost c .
- 2 The plaintiff makes a take-it-or-leave-it settlement offer of $s > 0$.
- 3 The defendant accepts or rejects the settlement offer.
- 4 If the defendant rejects the offer, the plaintiff decides whether to give up or go to trial at a cost p to himself and d to the defendant.
- 5 If the case goes to trial, the plaintiff wins amount x with probability γ and otherwise wins nothing.

PAYOFFS
Figure 4.4 shows the payoffs. Let $\gamma x < p$, so the plaintiff's expected winnings are less than his marginal cost of going to trial.

The perfect equilibrium is

Plaintiff: *Do nothing, Offer s, Give up*

Defendant: *Reject*

Outcome: The plaintiff does not bring a suit.

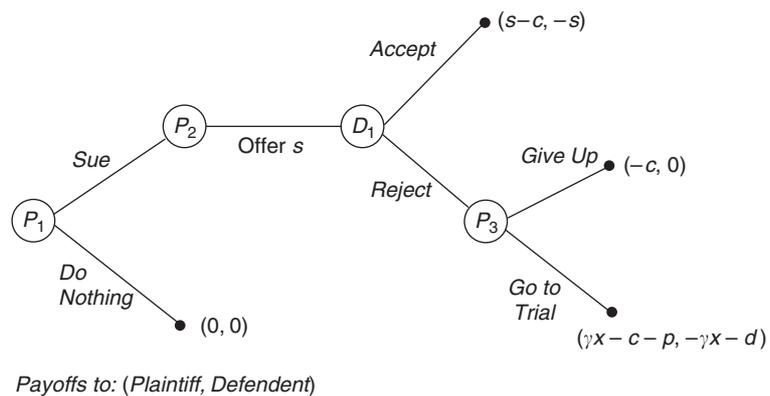


Figure 4.4 The extensive form for Nuisance Suits.

The equilibrium settlement offer s can be any positive amount. Note that the equilibrium specifies actions at all four nodes of the game, even though only the first is reached in equilibrium.

To find a perfect equilibrium the modeller starts at the end of the game tree, following the advice of Dixit & Nalebuff (1991, p. 34) to “Look ahead and reason back.” At node P_3 , the plaintiff will choose *Give Up*, since by assumption $\gamma x - c - p < -c$. This is because the suit is brought only in the hope of settlement, not in the hope of winning at trial. At node D_1 , the defendant, foreseeing that the plaintiff will give up, rejects any positive settlement offer. This makes the plaintiff’s offer at P_2 irrelevant, and, looking ahead to a payoff of $-c$ from choosing *Sue* at P_1 , the plaintiff chooses *Do Nothing*.

Thus, if nuisance suits are brought, it must be for some reason other than the obvious one, the plaintiff’s hope of extracting a settlement offer from a defendant who wants to avoid trial costs. This is fallacious because the plaintiff himself bears trial costs and hence cannot credibly make the threat. It is fallacious even if the defendant’s legal costs would be much higher than the plaintiff’s (d much bigger than p), because the relative size of the costs does not enter into the argument.

One might wonder how risk aversion affects this conclusion. Might not the defendant settle because he is more risk-averse than the plaintiff? That is a good question, but Nuisance Suits I can be adapted to risk-averse players with very little change. Risk would enter at the trial stage, as a final move by Nature to decide who wins. In Nuisance Suits I, γx represented the expected value of the award. If both the defendant and the plaintiff are equally risk-averse, γx can still represent the expected payoff from the award – one simply interprets x and 0 as the utility of the cash award and the utility of an award of 0, rather than as the actual cash amounts. If the players have different degrees of risk aversion, the expected loss to the defendant is not the same as the expected gain to the plaintiff, and the payoffs must be adjusted. If the defendant is more risk-averse, the payoffs from *Go to trial* would change to $(-c - p + \gamma x, -\gamma x - y - d)$, where y represents the extra disutility of risk to the defendant. This, however, makes no difference to the equilibrium. The crux of the game is that the plaintiff is unwilling to go to trial because of the cost to himself, and the cost to the defendant, including the cost of bearing risk, is irrelevant.

If nuisance suits are brought, it must therefore be for some more complicated reason. Already, in chapter 2, we looked at one reason for litigation to reach trial in the Png Settlement Game: incomplete information. That is probably the most important explanation and it has been much studied, as can be seen from the surveys by Cooter & Rubinfeld (1989) and Kennan & R. Wilson (1993). In this section, though, let us confine ourselves to explanations where the probability of the suit’s success is common knowledge. Even then, costly threats might be credible because of sinking costs strategically (Nuisance Suits II), or because of the nonmonetary payoffs resulting from going to trial (Nuisance Suits III).

Nuisance Suits II: Using Sunk Costs Strategically

Let us now modify the game, following the inspiration of Rosenberg & Shavell (1985), so that the plaintiff can pay his lawyer the amount p in advance, with no refund if the case settles. This inability to obtain a refund actually helps the plaintiff, by changing the payoffs from the game so his payoff from *Give Up* is $(-c - p)$, compared to $(-c - p + \gamma x)$ from *Go to Trial*. Having sunk the legal costs, he will go to trial if $\gamma x > 0$ – that is, if he has any chance of success at all.

This, in turn, means that the plaintiff would only prefer settlement to trial if $s > \gamma x$. The defendant would prefer settlement to trial if $s < \gamma x + d$, so there is a positive **settlement range** of $[\gamma x, \gamma x + d]$ within which both players are willing to settle. The exact amount of the settlement depends on the bargaining power of the parties, something to be examined in chapter 12. Here, allowing the plaintiff to make a take-it-or-leave-it offer means that $s = \gamma x + d$ in equilibrium, and if $\gamma x + d > p + c$, the nuisance suit will be brought even though $\gamma x < p + c$. Thus, the plaintiff is bringing the suit only because he can extort d , the amount of the defendant's legal costs.

Even though the plaintiff can now extort a settlement, he does it at some cost to himself, so an equilibrium with nuisance suits will require that

$$-c - p + \gamma x + d \geq 0. \quad (4.1)$$

If inequality (4.1) is false, then, even if the plaintiff could extract the maximum possible settlement of $s = \gamma x + d$, he would not do so, because he would then have to pay $c + p$ before reaching the settlement stage. This implies that a totally meritless suit (with $\gamma = 0$), would not be brought unless the defendant had higher legal costs than the plaintiff ($d > p$). If inequality (4.1) is satisfied, however, the following strategy profile is a perfect equilibrium:

Plaintiff: *Sue, Offer $s = \gamma x + d$, Go to Trial*

Defendant: *Accept $s \leq \gamma x + d$*

Outcome: Plaintiff sues and offers to settle, to which the defendant agrees.

An obvious counter to the plaintiff's ploy would be for the defendant to also sink his costs, by paying d before the settlement negotiations, or even before the plaintiff decides to file suit. Perhaps this is one reason why large corporations use in-house counsel, who are paid a salary regardless of how many hours they work, as well as outside counsel, hired by the hour. If so, nuisance suits cause a social loss – the wasted time of the lawyers, d – even if nuisance suits are never brought, just as aggressor nations cause social loss in the form of world military expenditure even if they never start a war.²

Two problems, however, face the defendant who tries to sink the cost d . First, although it saves him γx if it deters the plaintiff from filing suit, it also means the defendant must pay the full amount d . This is worthwhile if the plaintiff has all the bargaining power, as in Nuisance Suits II, but it might not be if s lay in the middle of the settlement range because the plaintiff was not able to make a take-it-or-leave-it offer. If settlement negotiations resulted in s lying exactly in the middle of the settlement range, so $s = \gamma x + (d/2)$, then it might not be worthwhile for the defendant to sink d to deter nuisance suits that would settle for $\gamma x + (d/2)$.

Second, there is an asymmetry in litigation: the plaintiff has the choice of whether to bring suit or not. Since it is the plaintiff who has the initiative, he can sink p and make

² Nonrefundable lawyers' fees, paid in advance, have traditionally been acceptable, but a New York court recently ruled they were unethical. The court thought that such fees unfairly restricted the client's ability to fire his lawyer, an example of how ignorance of game theory can lead to confused rule-making. See "Nonrefundable Lawyers' Fees, Paid in Advance, Are Unethical, Court Rules," *Wall Street Journal*, January 29, 1993, p. B3, citing *In the matter of Edward M. Cooperman, Appellate Division of the Supreme Court, Second Judicial Department, Brooklyn, 90-00429*.

the settlement offer before the defendant has the chance to sink d . The only way for the defendant to avoid this is to pay d well in advance, in which case the expenditure is wasted if no possible suits arise. What the defendant would like best would be to buy legal insurance which, for a small premium, would pay all defense costs in future suits that might occur. As we will see in chapters 8 and 9, however, insurance of any kind faces problems arising from asymmetric information. In this context, there is the “moral hazard” problem, in that once the defendant is insured he has less incentive to avoid causing harm to the plaintiff and provoking a lawsuit.

The Open-set Problem in Nuisance Suits II

Nuisance Suits II illustrates a technical point that arises in a great many games with continuous strategy spaces and causes great distress to novices in game theory. The equilibrium in Nuisance Suits II is only a weak Nash equilibrium. The plaintiff proposes $s = \gamma x + d$, and the defendant has the same payoff from accepting or rejecting, but in equilibrium the defendant accepts the offer with probability one, despite his indifference. This seems arbitrary, or even silly. Should not the plaintiff propose a slightly lower settlement to give the defendant a strong incentive to accept it and avoid the risk of having to go to trial? If the parameters are such that $s = \gamma x + d = 60$, for example, why does the plaintiff risk holding out for 60 when he might be rejected and most likely receive 0 at trial, when he could offer 59 and give the defendant a strong incentive to accept?

One answer is that no other equilibrium exists besides $s = 60$. Offering 59 cannot be part of an equilibrium because it is dominated by offering 59.9; offering 59.9 is dominated by offering 59.99, and so forth. This is known as the **open-set problem**, because the set of offers that the defendant strongly wishes to accept is open and has no maximum – it is bounded at 60, but a set must be bounded *and closed* to guarantee that a maximum exists. We are not just picking the Nash equilibrium in which the defendant accepts paying 60; we are picking the *only* Nash equilibrium.

A second answer is that under the assumptions of rationality and Nash equilibrium the objection’s premise is false because the plaintiff bears no risk whatsoever in offering $s = 60$. It is fundamental to Nash equilibrium that each player believe that the others will follow equilibrium behavior. Thus, if the equilibrium strategy profile says that the defendant will accept $s \leq 60$, the plaintiff can offer 60 and believe it will be accepted. This is really just to say that a weak Nash equilibrium is still a Nash equilibrium, a point emphasized in chapter 3 in connection with mixed strategies.

A third answer is that the problem is an artifact of using a model with a continuous strategy space, and it disappears if the strategy space is made discrete. Assume that s can only take values in multiples of 0.01, so it could be 59.0, 59.01, 59.02, and so forth, but not 59.001 or 59.002. The settlement part of the game will now have two perfect equilibria. In the strong equilibrium E1, $s = 59.99$ and the defendant accepts any offer $s < 60$. In the weak equilibrium E2, $s = 60$ and the defendant accepts any offer $s \leq 60$. The difference is trivial, so the discrete strategy space has made the model more complicated without any extra insight.³

³ A good example of the ideas of discrete money values and sequential rationality is in Robert Louis Stevenson’s 1893 story, “The Bottle Imp” (Stevenson [1987]). The imp grants the wishes of the bottle’s owner but will seize his soul if he dies in possession of it. Although the bottle cannot be given away, it can be sold, but only at a price less than that for which it was purchased.

One can also specify a more complicated bargaining game to avoid the issue of how exactly the settlement is determined. Here one could say that the settlement is not proposed by the plaintiff, but simply emerges with a value halfway through the settlement range, so $s = \gamma x + (d/2)$. This seems reasonable enough, and it adds a little extra realism to the model at the cost of a little extra complexity. It avoids the open-set problem, but only by avoiding being clear about how s is determined. I call this kind of modelling **blackboxing**, because it is as if at some point in the game, variables with certain values go into a black box and come out the other side with values determined by an exogenous process. Blackboxing is perfectly acceptable as long as it neither drives nor obscures the point the model is making. Nuisance Suits III will illustrate this method.

Fundamentally, however, the point to keep in mind is that games are models, not reality. They are meant to clear away the unimportant details of a real situation and simplify it down to the essentials. Since a model is trying to answer a question, it should focus on what answers that question. Here, the question is why nuisance suits might be brought, so it is proper to exclude details of the bargaining if they are irrelevant to the answer. Whether a plaintiff offers 59.99 or 60, and whether a rational person accepts an offer with probability 0.99 or 1.00, is part of the unimportant detail, and whatever approach is simplest should be used. If the modeller really thinks that these are important matters, they can indeed be modelled, but they are not important in this context.

One source of concern over the open-set problem, I think, is that perhaps that the payoffs are not quite realistic, because the players should derive utility from hurting “unfair” players. If the plaintiff makes a settlement offer of 60, keeping the entire savings from avoiding the trial for himself, everyday experience tells us that the defendant will indignantly refuse the offer. Guth, Schmittberger, & Schwarze (1982) have found in experiments that people turn down bargaining offers they perceive as unfair, as one might expect. If indignation is truly important, it can be explicitly incorporated into the payoffs, and if that is done, the open-set problem returns. Indignation is not boundless, whatever people may say. Suppose that accepting a settlement offer that benefits the plaintiff more than the defendant gives a disutility of x to the defendant because of his indignation at his unjust treatment. The plaintiff will then offer to settle for exactly $60 - x$, so the equilibrium is still weak and the defendant is still indifferent between accepting and rejecting the offer. The open-set problem persists, even after realistic emotions are added to the model.

I have spent so much time on the open-set problem not because it is important but because it arises so often and is a sticking point for people unfamiliar with modelling. It is not a problem that disturbs experienced modellers, unlike other basic issues we have already encountered – for example, the issue of how a Nash equilibrium comes to be common knowledge among the players – but it is important to understand why it is not important.

Nuisance Suits III: Malice

One of the most common misconceptions about game theory, as about economics in general, is that it ignores nonrational and nonmonetary motivations. Game theory does take the basic motivations of the players to be exogenous to the model, but those motivations are crucial to the outcome and they often are not monetary, although payoffs are always given numerical values. Game theory does not call somebody irrational who prefers leisure to money or who is motivated by the desire to be world dictator. It does require the players’ emotions to

be carefully gauged to determine exactly how the actions and outcomes affect the players' utility.

Emotions are often important to lawsuits, and law professors tell their students that when the cases they study seem to involve disputes too trivial to be worth taking to court, they can guess that the real motivations are emotional. Emotions could enter in a variety of distinct ways. The plaintiff might simply like going to trial, which can be expressed as a value of $p < 0$. This would be true of many criminal cases, because prosecutors like news coverage and want credit with the public for prosecuting certain kinds of crime. The Rodney King trials of 1992 and 1993 were of this variety; regardless of the merits of the cases against the policemen who beat Rodney King, the prosecutors wanted to go to trial to satisfy the public outrage, and when the state prosecutors failed in the first trial, the federal government was happy to accept the cost of bringing suit in the second trial. A different motivation is that the plaintiff might derive utility from the fact of winning the case quite separately from the monetary award, because he wants a public statement that he is in the right. This is a motivation in bringing libel suits, or for a criminal defendant who wants to clear his good name.

A different emotional motivation for going to trial is the desire to inflict losses on the defendant, a motivation we will call "malice," although it might as inaccurately be called "righteous anger." In this case, d enters as a positive argument in the plaintiff's utility function. We will construct a model of this kind, called Nuisance Suits III, and assume that $\gamma = 0.1$, $c = 3$, $p = 14$, $d = 50$, and $x = 100$, and that the plaintiff receives additional utility of 0.1 times the defendant's disutility. Let us also adopt the blackboxing technique discussed earlier and assume that the settlement s is in the middle of the settlement range. The payoffs conditional on suit being brought are

$$\pi_{\text{plaintiff}}(\text{Defendant accepts}) = s - c + 0.1s = 1.1s - 3 \quad (4.2)$$

and

$$\begin{aligned} \pi_{\text{plaintiff}}(\text{Go to Trial}) &= \gamma x - c - p + 0.1(d + \gamma x) \\ &= 10 - 3 - 14 + 6 = -1. \end{aligned} \quad (4.3)$$

Now, working back from the end in accordance with sequential rationality, note that since the plaintiff's payoff from *Give Up* is -3 , he will go to trial if the defendant rejects the settlement offer. The overall payoff from bringing a suit that eventually goes to trial is still -1 , which is worse than the payoff of 0 from not bringing suit in the first place, but if s is high enough, the payoff from bringing suit and settling is higher still. If s is greater than 1.82 ($=(-1 + 3)/1.1$, rounded), the plaintiff prefers settlement to trial, and if s is greater than about 2.73 ($=(0 + 3)/1.1$, rounded), he prefers settlement to not bringing the suit at all.

In determining the settlement range, the relevant payoff is the expected incremental payoff since the suit was brought. The plaintiff will settle for any $s \geq 1.82$, and the defendant will settle for any $s \leq \gamma x + d = 60$, as before. The settlement range is $[1.82, 60]$, and $s = 30.91$. The settlement offer is no longer the maximizing choice of a player, and hence is moved to the outcome in the equilibrium description below.

Plaintiff: *Sue, Go to Trial*

Defendant: *Accept any $s \leq 60$*

Outcome: The plaintiff sues and offers $s = 30.91$, and the defendant accepts the settlement.

Perfectness is important here because the defendant would like to threaten never to settle and be believed. The plaintiff would not bring suit given his expected payoff of -1 from bringing a suit that goes to trial, so a believable threat would be effective. But such a threat is not believable. Once the plaintiff does bring suit, the only Nash equilibrium in the remaining subgame is for the defendant to accept his settlement offer. This is interesting because the plaintiff, despite his willingness to go to trial, ends up settling out of court. When information is symmetric, as it is here, there is a tendency for equilibria to be efficient. Although the plaintiff wants to hurt the defendant, he also wants to keep his expenses low. Thus, he is willing to hurt the defendant less if it enables him to save on his own legal costs.

One final point before leaving these models is that much of the value of modelling comes simply from setting up the rules of the game, which helps to show what is important in a situation. One problem that arises in setting up a model of nuisance suits is deciding what a “nuisance suit” really is. In the game of Nuisance Suits, it has been defined as a suit whose expected damages do not repay the plaintiff’s costs of going to trial. But having to formulate a definition brings to mind another problem that might be called the problem of nuisance suits: that the plaintiff brings suits he knows will not win unless the court makes a mistake. Since the court might make a mistake with very high probability, the games above would not be appropriate models. The parameter γ would be high, and the problem is not that the plaintiff’s expected gain from trial is low, but that it is high. This, too, is an important problem, but having to construct a model shows that it is different.

4.4 Reoordination to Pareto-dominant Equilibria in Subgames: Pareto Perfection

One simple refinement of equilibrium that was mentioned in chapter 1 is to rule out any strategy profiles that are Pareto dominated by Nash equilibria. Thus, in the game of Ranked Coordination, the inferior Nash equilibrium would be ruled out as an acceptable equilibrium. The idea behind this is that in some unmodelled way the players discuss their situation and coordinate to avoid the bad equilibria. Since only Nash equilibria are discussed, the players’ agreements are self-enforcing and this is a more limited suggestion than the approach in cooperative game theory according to which the players make binding agreements.

The coordination idea can be taken further in various ways. One is to think about coalitions of players coordinating on favorable equilibria, so that two players might coordinate on an equilibrium even if a third player dislikes it. Bernheim, Peleg, & Whinston (1987), and Bernheim & Whinston (1987) define a Nash strategy profile as a **coalition-proof Nash equilibrium** if no coalition of players could form a self-enforcing agreement to deviate from it. They take the idea further by subordinating it to the idea of sequential rationality. The natural way to do this is to require that no coalition would deviate in future subgames, a notion called by various names, including **renegotiationproofness**, **reoordination** (e.g., Laffont & Tirole [1993], p. 460), and **Pareto perfection** (e.g., Fudenberg & Tirole [1991a], p. 175). The idea has been used extensively in the analysis of infinitely repeated games, which are particularly subject to the problem of multiple equilibria; Abreu, Pearce, &

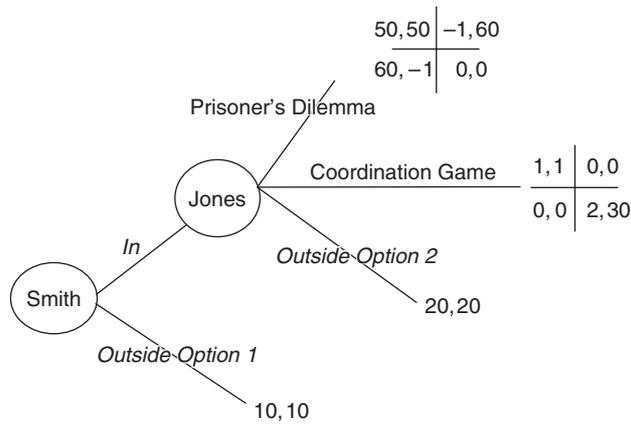


Figure 4.5 The Pareto Perfection Puzzle.

Stachetti (1986) is an example of this literature. Whichever name is used, the idea is distinct from the renegotiation problem in the principal–agent models to be studied in chapter 8, which involves the rewriting of earlier binding contracts to make new binding contracts.

The best way to demonstrate the idea of Pareto perfection is by an illustration, the Pareto Perfection Puzzle, whose extensive form is shown in figure 4.5. In this game Smith chooses *In* or *Outside Option 1*, which yields payoffs of 10 to each player. Jones then chooses *Outside Option 2*, which yields 20 to each player, or initiates either a coordination game or a prisoner’s dilemma. Rather than draw the full subgames in extensive form, figure 4.5 inserts the payoff matrix for the subgames.

The Pareto Perfection Puzzle illustrates the complicated interplay between perfectness and Pareto dominance. The Pareto-dominant strategy profile is (*In*, Prisoner’s Dilemma|*In*, any actions in the coordination subgame, the actions yielding (50, 50) in the Prisoner’s Dilemma subgame). Nobody expects this strategy profile to be an equilibrium, since it is neither perfect nor Nash. Perfectness tells us that if the Prisoner’s Dilemma subgame is reached, the payoffs will be (0, 0), and if the coordination subgame is reached they will be either (1, 1) or (2, 30). In light of this, the perfect equilibria of the Pareto Perfection Puzzle are

- E1: (*In*, *Outside option 2*|*In*, the actions yielding (1, 1) in the coordination subgame, the actions yielding (0, 0) in the Prisoner’s Dilemma subgame.) The payoffs are (20, 20).
- E2: (*Outside Option 1*, coordination game|*In*, the actions yielding (2, 30) in the coordination subgame, the actions yielding (0, 0) in the Prisoner’s Dilemma subgame.) The payoffs are (10, 10).

If one applies Pareto dominance without perfection, E1 will be the equilibrium, since both players prefer it. If the players can recoordinate at any point and change their expectations, however, then if play of the game reaches the coordination subgame, the players will recoordinate on the actions yielding (2, 30). Pareto perfection thus knocks out E1 as an equilibrium. Not only does it rule out the Pareto-dominant strategy profile that yields (50, 50) as an equilibrium, it also rules out the Pareto-dominant perfect strategy profile that yields

(20, 20) as an equilibrium. Rather, the payoff is (10, 10). Thus, Pareto perfection is not the same thing as simply picking the Pareto-dominant perfect strategy profile.

It is difficult to say which equilibrium is best here, since this is an abstract game and we cannot call upon details from the real world to refine the model. The approach of applying an equilibrium refinement is not as likely to yield results as using the intuition behind the refinement. The intuition here is that the players will somehow coordinate on Pareto-dominant equilibria, perhaps finding open discussion helpful. If we ran an experiment on student players using the Pareto Perfection Puzzle, I would expect to reach different equilibria depending on what communication is allowed. If the players are allowed to talk only before the game starts, it seems more likely that E1 would be the equilibrium, since players could agree to play it and would have no chance to explicitly re-coordinate later. If the players could talk at any time as the game proceeded, E2 becomes more plausible. Real-world situations arise with many different communications technologies, so there is no one right answer.

Notes

N4.1 Subgame perfectness

- The terms “perfectness” and “perfection” are used synonymously. Selten (1965) proposed the equilibrium concept in an article written in German. “Perfectness” is used in Selten (1975) and conveys an impression of completeness more appropriate to the concept than the goodness implied by “perfection.” “Perfection,” however, is more common.
- It is debatable whether the definition of subgame ought to include the original game. Gibbon (1992, p. 122) does not, for example, and modellers usually do not in their conversation.
- Perfectness is not the only way to eliminate weak Nash equilibria like (*Stay Out, Collude*). In Entry Deterrence I, (*Enter, Collude*) is the only iterated dominance equilibrium, because *Fight* is weakly dominated for the incumbent.
- The distinction between perfect and nonperfect Nash equilibria is like the distinction between **closed loop** and **open loop** trajectories in dynamic programming. Closed loop (or **feedback**) trajectories can be revised after they start, like perfect equilibrium strategies, while open loop trajectories are completely prespecified (though they may depend on state variables). In dynamic programming the distinction is not so important, because prespecified strategies do not change the behavior of other players. No threat, for example, is going to alter the pull of the moon’s gravity on a rocket.
- A subgame can be infinite in length, and infinite games can have nonperfect equilibria. The infinitely repeated Prisoner’s Dilemma is an example; here every subgame looks exactly like the original game but begins at a different point in time.
- **Sequential rationality in macroeconomics.** In macroeconomics the requirement of **dynamic consistency** or **time consistency** is similar to perfectness. These terms are less precisely defined than perfectness, but they usually require that strategies need only be best responses in subgames starting from nodes on the equilibrium path, instead of all subgames. Under this interpretation, time consistency is a less stringent condition than perfectness.

The Federal Reserve, for example, might like to induce inflation to stimulate the economy, but the economy is stimulated only if the inflation is unexpected. If the inflation is expected, its effects are purely bad. Since members of the public know that the Fed would like to fool them, they disbelieve its claims that it will not generate inflation (see Kydland & Prescott [1977]). Likewise, the government would like to issue nominal debt, and promises lenders that it will keep inflation

low, but once the debt is issued, the government has an incentive to inflate its real value to zero. One reason the US Federal Reserve Board was established to be independent of Congress in the United States was to diminish this problem.

- Often, irrationality – behavior that is automatic rather than strategic – is an advantage. The Domsday Machine in the movie *Dr Strangelove* is one example. The Soviet Union decides that it cannot win a rational arms race against the richer United States, so it creates a bomb which automatically blows up the entire world if anyone explodes a nuclear bomb. The movie also illustrates a crucial detail without which such irrationality is worse than useless: you have to tell the other side that you have the Domsday Machine.

President Nixon reportedly told his aide H. R. Haldeman that he followed a more complicated version of this strategy: “I call it the Madman Theory, Bob. I want the North Vietnamese to believe that I’ve reached the point where I might do *anything* to stop the war. We’ll just slip the word to them that ‘for God’s sake, you know Nixon is obsessed about Communism. We can’t restrain him when he’s angry – and he has his hand on the nuclear button’ – and Ho Chi Minh himself will be in Paris in two days begging for peace” (H. R. Haldeman & Joseph DiMona, *The Ends of Power*, 1978, p. 83). The Gang of Four model in section 6.4 tries to model a situation like that.

- The “lock-up agreement” is an example of a credible threat: in a takeover defense, the threat to destroy the firm is made legally binding. See Macey & McChesney (1985), p. 33.
- A famous paradox relating to sequential rationality is the “Quiz on Friday Paradox.” You are going to have a quiz next week, but I am going to surprise you with my choice of day. If we reach Thursday without a quiz, however, you will know the quiz must be on Friday and not be surprised, so the quiz must be earlier. The same argument applies to looking ahead to Thursday on Wednesday, however, and by iteration all days can be ruled out. Philosophers from Quine (1953) to Schick (2003, ch. 5) have puzzled over this.

N4.3 An example of perfectness: Entry Deterrence I

- The Stackelberg equilibrium of a duopoly game (section 3.4) can be viewed as the perfect equilibrium of a Cournot game modified so that one player moves first, a game similar to Entry Deterrence I. The player moving first is the Stackelberg leader and the player moving second is the Stackelberg follower. The follower could threaten to produce a high output, but he will not carry out his threat if the leader produces a high output first.
- Perfectness is not so desirable a property of equilibrium in biological games. The reason the order of moves matters is because the rational best reply depends on the node at which the game has arrived. In many biological games the players act by instinct and unthinking behavior is not unrealistic.
- Reinganum & Stokey (1985) is a clear presentation of the implications of perfectness and commitment illustrated with the example of natural resource extraction.

Problems

4.1: Repeated Entry Deterrence (easy)

Consider two repetitions without discounting of the game Entry Deterrence I from section 4.2. Assume that there is one entrant, who sequentially decides whether to enter two markets that have the same incumbent.

- Draw the extensive form of this game.
- What are the 16 elements of the strategy sets of the entrant?

- (c) What is the subgame perfect equilibrium?
- (d) What is one of the nonperfect Nash equilibria?

4.2: The Three-way Duel (medium) (after Shubik [1954])

Three gangsters armed with pistols, Al, Bob, and Curly, are in a room with a suitcase containing 120,000 dollars. Al is the least accurate, with a 20 percent chance of killing his target. Bob has a 40 percent probability. Curly is slow but sure; he kills his target with 70 percent probability. For each, the value of his own life outweighs the value of any amount of money. Survivors split the money.

- (a) Suppose each gangster has one bullet and the order of shooting is first Al, then Bob, then Curly. Assume also that each gangster must try to kill another gangster when his turn comes. What is an equilibrium strategy profile and what is the probability that each of them dies in that equilibrium? Hint: Do not try to draw a game tree.
- (b) Suppose now that each gangster has the additional option of shooting his gun at the ceiling, which may kill somebody upstairs but has no direct effect on his payoff. Does the strategy profile that you found was an equilibrium in part (a) remain an equilibrium?
- (c) Replace the three gangsters with three companies, Apex, Brydox, and Costco, which are competing with slightly different products. What story can you tell about their advertising strategies?
- (d) In the United States, before the general election a candidate must win the nomination of his party. It is often noted that candidates are reluctant to be seen as the frontrunner in the race for the nomination of their party, Democrat or Republican. In the general election, however, no candidate ever minds being seen to be ahead of his rival from the other party. Why?
- (e) In the 1920s, several men vied for power in the Soviet Union after Lenin died. First Stalin and Zinoviev combined against Trotsky. Then Stalin and Bukharin combined against Zinoviev. Then Stalin turned on Bukharin. Relate this to Curly, Bob, and Al.

4.3: Heresthetics in Pliny and the Freedmen's Trial (easy) (Pliny [105] "To Aristo," Riker [1986, pp. 78–88])

Afranius Dexter died mysteriously, perhaps dead by his own hand, perhaps killed by his freedmen (servants a step above slaves), or perhaps killed by his freedmen by his own orders. The freedmen went on trial before the Roman Senate. Assume that 45 percent of the senators favor acquittal, 35 percent favor banishment, and 20 percent favor execution, and that the preference rankings in the three groups are $A > B > E$, $B > A > E$, and $E > B > A$. Also assume that each group has a leader and votes as a bloc.

- (a) Modern legal procedure requires the court to decide guilt first and then assign a penalty if the accused is found guilty. Draw a tree to represent the sequence of events (this will not be a game tree, since it will represent the actions of groups of players, not of individuals). What is the outcome in a perfect equilibrium?
- (b) Suppose that the acquittal bloc can pre-commit to how they will vote in the second round if guilt wins in the first round. What will they do, and what will happen? What would the execution bloc do if they could control the second-period vote of the acquittal bloc?
- (c) The normal Roman procedure began with a vote on execution versus no execution, and then voted on the alternatives in a second round if execution failed to gain a majority. Draw a tree to represent this. What would happen in this case?

- (d) Pliny proposed that the Senators divide into three groups, depending on whether they supported acquittal, banishment, or execution, and that the outcome with the most votes should win. This proposal caused a roar of protest. Why did he propose it?
- (e) Pliny did not get the result he wanted with his voting procedure. Why not?
- (f) Suppose that personal considerations made it most important to a senator that he show his stand by his vote, even if he had to sacrifice his preference for a particular outcome. If there were a vote over whether to use the traditional Roman procedure or Pliny's procedure, who would vote with Pliny, and what would happen to the freedmen?

4.4: Garbage entry (medium)

Mr Turner is thinking of entering the garbage collection business in a certain large city. Currently, Cutright Enterprises has a monopoly, earning 40 million dollars from the 40 routes the city offers up for bids. Turner thinks he can take away as many routes as he wants from Cutright, at a profit of 1.5 million per route for him. He is worried, however, that Cutright might resort to assassination, killing him to regain their lost routes. He would be willing to be assassinated for profit of 80 million dollars, and assassination would cost Cutright 6 million dollars in expected legal costs and possible prison sentences.

How many routes should Turner try to take away from Cutright?

4.5: Voting cycles (medium)

Uno, Duo, and Tres are three people voting on whether the budget devoted to a project should be Increased, kept the Same, or Reduced. Their payoffs from the different outcomes, given in table 4.2, are not monotonic in budget size. Uno thinks the project could be very profitable if its budget were increased, but will fail otherwise. Duo mildly wants a smaller budget. Tres likes the budget as it is now.

Table 4.2 Payoffs from different policies

	<i>Uno</i>	<i>Duo</i>	<i>Tres</i>
Increase	100	2	4
Same	3	6	9
Reduce	9	8	1

Each of the three voters writes down his first choice. If a policy gets a majority of the votes, it wins. Otherwise, *Same* is the chosen policy.

- (a) Show that $(\textit{Same}, \textit{Same}, \textit{Same})$ is a Nash equilibrium. Why does this equilibrium seem unreasonable to us?
- (b) Show that $(\textit{Increase}, \textit{Same}, \textit{Same})$ is a Nash equilibrium.
- (c) Show that if each player has an independent small probability ϵ of “trembling” and choosing each possible wrong action by mistake, $(\textit{Same}, \textit{Same}, \textit{Same})$ and $(\textit{Increase}, \textit{Same}, \textit{Same})$ are no longer equilibria.
- (d) Show that $(\textit{Reduce}, \textit{Reduce}, \textit{Same})$ is a Nash equilibrium that survives each player has an independent small probability ϵ of “trembling” and choosing each possible wrong action by mistake.
- (e) Part (d) showed that if Uno and Duo are expected to choose *Reduce*, then Tres would choose *Same* if he could hope they might tremble – not *Increase*. Suppose, instead, that Tres votes

first, and publicly. Construct a subgame perfect equilibrium in which Tres chooses *Increase*. You need not worry about trembles now.

- (f) Consider the following voting procedure. First, the three voters vote between *Increase* and *Same*. In the second round, they vote between the winning policy and *Reduce*. If, at that point, *Increase* is not the winning policy, the third vote is between *Increase* and whatever policy won in the second round.
- What will happen? (Watch out for the trick in this question!)
- (g) Speculate about what would happen if the payoffs are in terms of dollar willingness to pay by each player and the players could make binding agreements to buy and sell votes. What, if anything, can you say about which policy would win, and what votes would be bought at what price?

US Air for Sale: A Classroom Game for Chapter 4

On October 2, 1995, US Air, at the time the nation's fifth largest airline, announced that it had approached United Airlines and American Airlines about a possible buyout. United and American are the two largest U.S. airlines, and both are interested in growth in the highly competitive airline market. United and American must now consider what to do.

Financial analysts for both United and American have made projections for all the possible scenarios, on the request of the strategic bidding consultants. They say that United and American start in equally strong positions, but that can change depending on who wins the auction.

The three possible outcomes for an airline are:

Neither firm bids enough for US Air to accept. Both United and American will maintain their strong positions, and they can expect future profits of \$50 billion each.

Our airline wins. The winning airline will become the dominant firm in the market, and can expect future profits of \$80 billion. The purchase cost of US Air, however, must be subtracted from this in order to calculate net profit. Therefore, the payoff for the firm with the winning bid is

$$\text{Payoff}_{\text{Winner}} = 80 - B_{\text{winning}}. \quad (4.4)$$

The minimum price US Air will accept is \$10 billion.

Our airline loses. The losing firm may have trouble competing with the larger network of the winner because of the greater variety of flights a large airline can offer. The more, however, that the winner pays for US Air, the better off is the loser, because the cash flow the winner uses to pay for the acquisition is no longer available for other investments in new equipment and lenders will be more reluctant to lend to the newly enlarged firm. The analysts suggest that a good approximation to the payoff of the losing firm will be

$$\text{Payoff}_{\text{Loser}} = 30 + 0.25 * B_{\text{winning}}. \quad (4.5)$$

US Air for Sale: A Classroom Game for Chapter 4 (Continued)

In particular, if the winner pays just \$10 billion for US Air, the loser's profit will be \$32.5 billion.

We will look at two kinds of auction rules: ascending, and first price.

First-price Auction Rules: US Air will solicit bids in writing from anyone who cares to bid. Once all bids are received, they will be announced publicly. If there is just one bid, it wins if it is \$10 billion or above. If there are no bids, US Air will continue as an independent firm. If there are two bids, and at least one is \$10 billion or above, the winner has purchased US Air, ties being broken by the flip of a coin.

Ascending Auction Rules: US Air will solicit initial bids in writing from United and American, choosing the first bidder randomly and announcing its bid publicly before asking the other for a written bid. If there are no written bids of at least 10 billion dollars, US Air will continue as an independent firm. If there is at least one bid of 10 billion dollars or above, US Air will allow counteroffers, and an ascending auction commences. The winner has purchased US Air.

Students playing the game will be put into groups of three representing either United or American. Groups will be paired up to play out the auction under the ascending or first-price rules. First, play all the pairs by the first-price rules, then by the ascending rules.