

6 Dynamic Games with Incomplete Information

6.1 Perfect Bayesian Equilibrium: Entry Deterrence II and III

Asymmetric information, and, in particular, incomplete information, is enormously important in game theory. This is particularly true for dynamic games, since when the players have several moves in sequence, their earlier moves may convey private information that is relevant to the decisions of players moving later on. Revealing and concealing information are the basis of much of strategic behavior and are especially useful as ways of explaining actions that would be irrational in a nonstrategic world.

Chapter 4 showed that even if there is symmetric information in a dynamic game, Nash equilibrium may need to be refined using subgame perfectness if the modeller is to make sensible predictions. Asymmetric information requires a somewhat different refinement to capture the idea of sunk costs and credible threats, and Section 6.1 sets out the standard refinement of perfect bayesian equilibrium. Section 6.2 shows that even this may not be enough refinement to guarantee uniqueness and discusses further refinements based on out-of- equilibrium beliefs. Section 6.3 uses the idea to show that a player's ignorance may work to his advantage, and to explain how even when all players know something, lack of common knowledge still affects the game. Section 6.4 introduces incomplete information into the repeated prisoner's dilemma and shows the Gang of Four solution to the Chainstore Paradox of Chapter 5. Section 6.5 describes the celebrated Axelrod tournament, an experimental approach to the same paradox. Section 6.6 applies the idea of a dynamic game of incomplete information to the evolution of creditworthiness using the model of Diamond (1989).

Subgame Perfectness Is Not Enough

In games of asymmetric information, we will still require that an equilibrium be subgame perfect, but the mere forking of the game tree might not be relevant to a player's decision, because with asymmetric information he does not know which fork the game has taken. Smith might know he is at one of two different nodes depending on whether Jones has high or low production costs, but if he does not know the exact node, the "subgames" starting at each node are irrelevant to his decisions. In fact, they are not even subgames as we have defined them, because they cut across Smith's information sets. This can be seen in an asymmetric information version of Entry Deterrence I (Section 4.2). In Entry

Deterrence I, the incumbent colluded with the entrant because fighting him was more costly than colluding once the entrant had entered. Now, let us set up the game to allow some entrants to be *Strong* and some *Weak* in the sense that it is more costly for the incumbent to choose *Fight* against a *Strong* entrant than a *Weak* one. The incumbent's payoff from *Fight|Strong* will be 0, as before, but his payoff from *Fight|Weak* will be X , where X will take values ranging from 0 (Entry Deterrence I) to 300 (Entry Deterrence IV and V) in different versions of the game.

Entry Deterrence II, III, and IV will all have the extensive form shown in Figure 1. With 50 percent probability, the incumbent's payoff from *Fight* is X rather than the 0 in Entry Deterrence I, but the incumbent does not know which payoff is the correct one in the particular realization of the game. This is modelled as an initial move by Nature, who chooses between the entrant being *Weak* or *Strong*, unobserved by the incumbent.

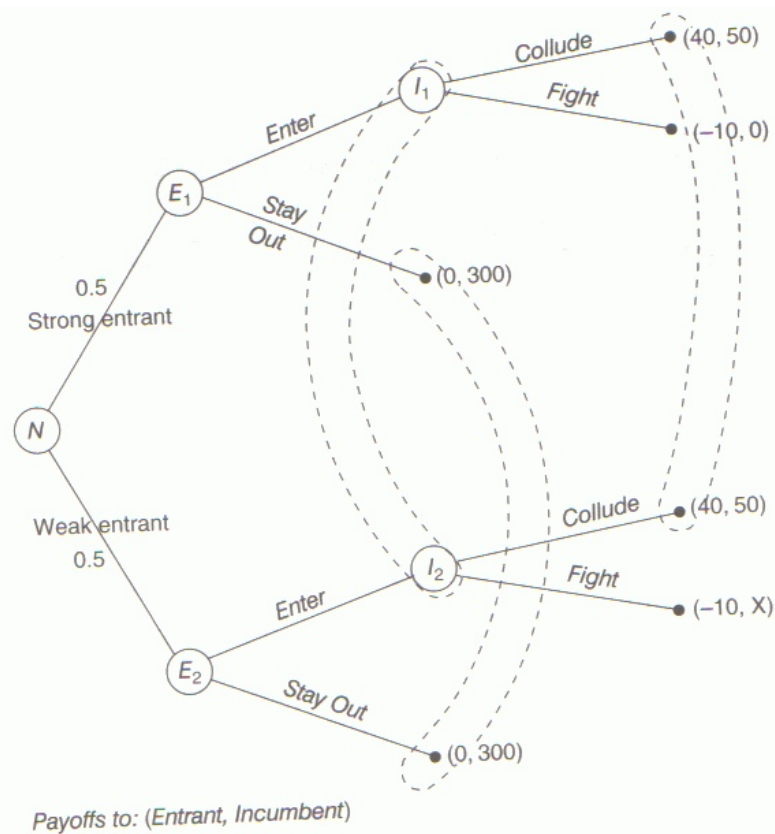


Figure 6.1 Entry Deterrence II, III, and IV

Figure 1: Entry Deterrence II, III, and IV

Entry Deterrence II: Fighting Is Never Profitable

In Entry Deterrence II, $X = 1$, so information is not very asymmetric. It is common knowledge that the incumbent never benefits from *Fight*, even though his exact payoff might be zero or might be one. Unlike in the case of Entry Deterrence I, however, subgame perfectness does not rule out any Nash equilibria, because the only subgame is the subgame

starting at node N , which is the entire game. A subgame cannot start at nodes E_1 or E_2 , because neither of those nodes are singletons in the information partitions. Thus, the implausible Nash equilibrium, (*Stay Out, Fight*), escapes elimination by a technicality.

The equilibrium concept needs to be refined in order to eliminate the implausible equilibrium. Two general approaches can be taken: either introduce small “trembles” into the game, or require that strategies be best responses given rational beliefs. The first approach takes us to the “trembling hand-perfect” equilibrium, while the second takes us to the “perfect bayesian” and “sequential” equilibrium. The results are similar whichever approach is taken.

Trembling-Hand Perfectness

Trembling-hand perfectness is an equilibrium concept introduced by Selten (1975) according to which a strategy that is to be part of an equilibrium must continue to be optimal for the player even if there is a small chance that the other player will pick an out-of-equilibrium action (i.e., that the other player’s hand will “tremble”).

Trembling-hand perfectness is defined for games with finite action sets as follows.

The strategy profile s^ is a **trembling-hand perfect equilibrium** if for any ϵ there is a vector of positive numbers $\delta_1, \dots, \delta_n \in [0, 1]$ and a vector of completely mixed strategies $\sigma_1, \dots, \sigma_n$ such that the perturbed game where every strategy is replaced by $(1 - \delta_i)s_i + \delta_i\sigma_i$ has a Nash equilibrium in which every strategy is within distance ϵ of s^* .*

Every trembling-hand perfect equilibrium is subgame perfect; indeed, Section 4.1 justified subgame perfectness using a tremble argument. Unfortunately, it is often hard to tell whether a strategy profile is trembling-hand perfect, and the concept is undefined for games with continuous strategy spaces because it is hard to work with mixtures of a continuum (see note N3.1). Moreover, the equilibrium depends on which trembles are chosen, and deciding why one tremble should be more common than another may be difficult.

Perfect Bayesian Equilibrium and Sequential Equilibrium

The second approach to asymmetric information, introduced by Kreps & Wilson (1982b) in the spirit of Harsanyi (1967), is to start with prior beliefs, common to all players, that specify the probabilities with which Nature chooses the types of the players at the beginning of the game. Some of the players observe Nature’s move and update their beliefs, while other players can update their beliefs only by deductions they make from observing the actions of the informed players.

The deductions used to update beliefs are based on the actions specified by the equilibrium. When players update their beliefs, they assume that the other players are following the equilibrium strategies, but since the strategies themselves depend on the beliefs, an equilibrium can no longer be defined based on strategies alone. Under asymmetric information, an equilibrium is a strategy profile and a set of beliefs such that the strategies are best responses. The profile of beliefs and strategies is called an **assessment** by Kreps and Wilson.

On the equilibrium path, all that the players need to update their beliefs are their priors and Bayes' s Rule, but off the equilibrium path this is not enough. Suppose that in equilibrium, the entrant always enters. If for whatever reason the impossible happens and the entrant stays out, what is the incumbent to think about the probability that the entrant is weak? Bayes' s Rule does not help, because when $Prob(data) = 0$, which is the case for data such as *Stay Out* which is never observed in equilibrium, the posterior belief cannot be calculated using Bayes' s Rule. From section 2.4,

$$Prob(Weak|Stay Out) = \frac{Prob(Stay Out|Weak)Prob(Weak)}{Prob(Stay Out)}. \quad (1)$$

The posterior $Prob(Weak|Stay Out)$ is undefined, because (1) requires dividing by zero.

A natural way to define equilibrium is as a strategy profile consisting of best responses given that equilibrium beliefs follow Bayes' s Rule and out-of- equilibrium beliefs follow a specified pattern that does not contradict Bayes' s Rule.

A **perfect bayesian equilibrium** is a strategy profile s and a set of beliefs μ such that at each node of the game:

- (1) The strategies for the remainder of the game are Nash given the beliefs and strategies of the other players.
- (2) The beliefs at each information set are rational given the evidence appearing thus far in the game (meaning that they are based, if possible, on priors updated by Bayes' s Rule, given the observed actions of the other players under the hypothesis that they are in equilibrium).

Kreps & Wilson (1982b) use this idea to form their equilibrium concept of sequential equilibrium, but they impose a third condition, defined only for games with discrete strategies, to restrict beliefs a little further:

- (3) The beliefs are the limit of a sequence of rational beliefs, i.e., if (μ^*, s^*) is the equilibrium assessment, then some sequence of rational beliefs and completely mixed strategies converges to it:

$$(\mu^*, s^*) = \text{Lim}_{n \rightarrow \infty} (\mu^n, s^n) \text{ for some sequence } (\mu^n, s^n) \text{ in } \{\mu, s\}.$$

Condition (3) is quite reasonable and makes sequential equilibrium close to trembling-hand perfect equilibrium, but it adds more to the concept's difficulty than to its usefulness. If players are using the sequence of completely mixed strategies s^n , then every action is taken with some positive probability, so Bayes'Rule can be applied to form the beliefs μ^n after any action is observed. Condition (3) says that the equilibrium assessment has to be the limit of some such sequence (though not of every such sequence). For the rest of the book we will use perfect bayesian equilibrium and dispense with condition (3), although it usually can be satisfied.

Sequential equilibria are always subgame perfect (condition (1) takes care of that). Every trembling-hand perfect equilibrium is a sequential equilibrium, and "almost every" sequential equilibrium is trembling hand perfect. Every sequential equilibrium is perfect bayesian, but not every perfect bayesian equilibrium is sequential.

Back to Entry Deterrence II

Armed with the concept of the perfect bayesian equilibrium, we can find a sensible equilibrium for Entry Deterrence II .

Entrant: $Enter|Weak, Enter|Strong$
Incumbent: $Collude$
Beliefs: $Prob(Strong| Stay Out) = 0.4$

In this equilibrium the entrant enters whether he is *Weak* or *Strong*. The incumbent's strategy is *Collude*, which is not conditioned on Nature's move, since he does not observe it. Because the entrant enters regardless of Nature's move, an out-of-equilibrium belief for the incumbent if he should observe *Stay Out* must be specified, and this belief is arbitrarily chosen to be that the incumbent's subjective probability that the entrant is *Strong* is 0.4 given his observation that the entrant deviated by choosing *Stay Out*. Given this strategy profile and out-of-equilibrium belief, neither player has incentive to change his strategy.

There is no perfect bayesian equilibrium in which the entrant chooses *Stay Out*. *Fight* is a bad response even under the most optimistic possible belief, that the entrant is *Weak* with probability 1. Notice that perfect bayesian equilibrium is not defined structurally, like subgame perfectness, but rather in terms of optimal responses. This enables it to come closer to the economic intuition which we wish to capture by an equilibrium refinement.

Finding the perfect bayesian equilibrium of a game, like finding the Nash equilibrium, requires intelligence. Algorithms are not useful. To find a Nash equilibrium, the modeller thinks about his game, picks a plausible strategy profile, and tests whether the strategies are best responses to each other. To make it a perfect bayesian equilibrium, he notes which actions are never taken in equilibrium and specifies the beliefs that players use to interpret those actions. He then tests whether each player's strategies are best responses given his beliefs at each node, checking in particular whether any player would like to take an out-of-equilibrium action in order to set in motion the other players' out-of-equilibrium beliefs and strategies. This process does not involve testing whether a player's beliefs are beneficial to the player, because players do not choose their own beliefs; the priors and out-of-equilibrium beliefs are exogenously specified by the modeller.

One might wonder why the beliefs have to be specified in Entry Deterrence II. Does not the game tree specify the probability that the entrant is *Weak*? What difference does it make if the entrant stays out? Admittedly, Nature does choose each type with probability 0.5, so if the incumbent had no other information than this prior, that would be his belief. But the entrant's action might convey additional information. The concept of perfect bayesian equilibrium leaves the modeller free to specify how the players form beliefs from that additional information, so long as the beliefs do not violate Bayes' Rule. (A technically valid choice of beliefs by the modeller might still be met with scorn, though, as with any silly assumption.) Here, the equilibrium says that if the entrant stays out, the incumbent believes he is *Strong* with probability 0.4 and *Weak* with probability 0.6, beliefs that are arbitrary but do not contradict Bayes' s Rule.

In Entry Deterrence II the out-of-equilibrium beliefs do not and should not matter.

If the entrant chooses *Stay Out*, the game ends, so the incumbent's beliefs are irrelevant. Perfect bayesian equilibrium was only introduced as a way out of a technical problem. In the next section, however, the precise out-of- equilibrium beliefs will be crucial to which strategy profiles are equilibria.

6.2 Refining Perfect Bayesian Equilibrium: The PhD Admissions Game

Entry Deterrence III: Fighting Is Sometimes Profitable

In Entry Deterrence III, assume that $X = 60$, not $X = 1$. This means that fighting is more profitable for the incumbent than collusion if the entrant is *Weak*. As before, the entrant knows if he is *Weak*, but the incumbent does not. Retaining the prior after observing out- of-equilibrium actions, which in this game is $Prob(Strong) = 0.5$, is a convenient way to form beliefs that is called **passive conjectures**. The following is a perfect bayesian equilibrium which uses passive conjectures.

A plausible pooling equilibrium for Entry Deterrence III

Entrant: $Enter|Weak, Enter|Strong$

Incumbent: $Collude$, Out-of-equilibrium beliefs: $Prob(Strong| Stay Out) = 0.5$

In choosing whether to enter, the entrant must predict the incumbent's behavior. If the probability that the entrant is *Weak* is 0.5, the expected payoff to the incumbent from choosing *Fight* is 30 ($= 0.5[0] + 0.5[60]$), which is less than the payoff of 50 from *Collude*. The incumbent will collude, so the entrant enters. The entrant may know that the incumbent's payoff is actually 60, but that is irrelevant to the incumbent's behavior.

The out-of-equilibrium belief does not matter to this first equilibrium, although it will in other equilibria of the same game. Although beliefs in a perfect bayesian equilibrium must follow Bayes' s Rule, that puts very little restriction on how players interpret out-of-equilibrium behavior. Out-of- equilibrium behavior is "impossible," so when it does occur there is no obvious way the player should react. Some beliefs may seem more reasonable than others, however, and Entry Deterrence III has another equilibrium that requires less plausible beliefs off the equilibrium path.

An implausible equilibrium for Entry Deterrence III

Entrant: $Stay Out|Weak, Stay Out|Strong$

Incumbent: $Fight$, Out-of-equilibrium beliefs: $Prob(Strong|Enter) = 0.1$

This is an equilibrium because if the entrant were to deviate and enter, the incumbent would calculate his payoff from fighting to be 54 ($= 0.1[0] + 0.9[60]$), which is greater than the *Collude* payoff of 50. The entrant would therefore stay out.

The beliefs in the implausible equilibrium are different and less reasonable than in the plausible equilibrium. Why should the incumbent believe that weak entrants would enter

mistakenly nine times as often as strong entrants? The beliefs do not violate Bayes' s Rule, but they have no justification.

The reasonableness of the beliefs is important because if the incumbent uses passive conjectures, the implausible equilibrium breaks down. With passive conjectures, the incumbent would want to change his strategy to *Collude*, because the expected payoff from *Fight* would be less than 50. The implausible equilibrium is less robust with respect to beliefs than the plausible equilibrium, and it requires beliefs that are harder to justify.

Even though dubious outcomes may be perfect bayesian equilibria, the concept does have some bite, ruling out other dubious outcomes. There does not, for example, exist an equilibrium in which the entrant enters only if he is *Strong* and stays out if he is *Weak* (called a "separating equilibrium" because it separates out different types of players). Such an equilibrium would have to look like this:

A conjectured separating equilibrium for Entry Deterrence III

Entrant: *Stay Out*|*Weak*, *Enter*|*Strong*

Incumbent: *Collude*

No out-of-equilibrium beliefs are specified for the conjectures in the separating equilibrium because there is no out-of-equilibrium behavior about which to specify them. Since the incumbent might observe either *Stay out* or *Enter* in equilibrium, the incumbent will always use Bayes' s Rule to form his beliefs. He will believe that an entrant who stays out must be weak and an entrant who enters must be strong. This conforms to the idea behind Nash equilibrium that each player assumes that the other follows the equilibrium strategy, and then decides how to reply. Here, the incumbent's best response, given his beliefs, is *Collude*|*Enter*, so that is the second part of the proposed equilibrium. But this cannot be an equilibrium, because the entrant would want to deviate. Knowing that entry would be followed by collusion, even the weak entrant would enter. So there cannot be an equilibrium in which the entrant enters only when strong.

The PhD Admissions Game

Passive conjectures may not always be the most satisfactory belief, as the next example shows. Suppose that a university knows that 90 percent of the population hate economics and would be unhappy in its PhD program, and 10 percent love economics and would do well. In addition, it cannot observe the applicant's type. If the university rejects an application, its payoff is 0 and the applicant's is -1 because of the trouble needed to apply. If the university accepts the application of someone who hates economics, the payoffs of both university and student are -10 , but if the applicant loves economics, the payoffs are $+20$ for each player. Figure 2 shows this game in extensive form. The population proportions are represented by a node at which Nature chooses the student to be a *Lover* or *Hater* of economics.




Figure 2: The PhD Admissions Game

The PhD Admissions Game is a signalling game of the kind we will look at in Chapter 10. It has various perfect bayesian equilibria that differ in their out-of-equilibrium beliefs, but the equilibria can be divided into two distinct categories, depending on the outcome: the **separating equilibrium**, in which the lovers of economics apply and the haters do not, and the **pooling equilibrium**, in which neither type of student applies.

A separating equilibrium for the PhD Admissions Game

Student: *Apply* | *Lover*, *Do Not Apply* | *Hater*

University: *Admit*

The separating equilibrium does not need to specify out-of-equilibrium beliefs, because Bayes' s Rule can always be applied whenever both of the two possible actions *Apply* and *Do Not Apply* can occur in equilibrium.

A pooling equilibrium for the PhD Admissions Game

Student: *Do Not Apply* | *Lover*, *Do Not Apply* | *Hater*

University: *Reject*, Out-of-equilibrium beliefs: $Prob(Hater | Apply) = 0.9$ (passive conjectures)

The pooling equilibrium is supported by passive conjectures. Both types of students refrain from applying because they believe correctly that they would be rejected and receive a payoff of -1 ; and the university is willing to reject any student who foolishly applied, believing that he is a *Hater* with 90 percent probability.

Because the perfect bayesian equilibrium concept imposes no restrictions on out-of-equilibrium beliefs, economists have come up with a variety of exotic refinements of the equilibrium concept. Let us consider whether various alternatives to passive conjectures would support the pooling equilibrium in PhD Admissions.

Passive Conjectures. $Prob(Hater | Apply) = 0.9$

This is the belief specified above, under which out-of-equilibrium behavior leaves beliefs unchanged from the prior. The argument for passive conjectures is that the student's application is a mistake, and that both types are equally likely to make mistakes, although *Haters* are more common in the population. This supports the pooling equilibrium.

The Intuitive Criterion. $Prob(Hater | Apply) = 0$

Under the Intuitive Criterion of Cho & Kreps (1987), if there is a type of informed player who could not benefit from the out-of-equilibrium action no matter what beliefs were held by the uninformed player, the uninformed player's belief must put zero probability on that type. Here, the *Hater* could not benefit from applying under any possible beliefs of the university, so the university puts zero probability on an applicant being a *Hater*. This argument will not support the pooling equilibrium, because if the university holds this belief, it will want to admit anyone who applies.

Complete Robustness. $Prob(Hater | Apply) = m, 0 \leq m \leq 1$

Under this approach, the equilibrium strategy profile must consist of responses that are best, given any and all out-of-equilibrium beliefs. Our equilibrium for Entry Deterrence II satisfied this requirement. Complete robustness rules out a pooling equilibrium in the PhD Admissions Game, because a belief like $m = 0$ makes accepting applicants a best response, in which case only the *Lover* will apply. A useful first step in analyzing conjectured pooling equilibria is to test whether they can be supported by extreme beliefs such as $m = 0$ and $m = 1$.

An Ad Hoc Specification. $Prob(Hater | Apply) = 1$

Sometimes the modeller can justify beliefs by the circumstances of the particular game. Here, one could argue that anyone so foolish as to apply knowing that the university would

reject them could not possibly have the good taste to love economics. This supports the pooling equilibrium also.

An alternative approach to the problem of out-of-equilibrium beliefs is to remove its origin by building a model in which every outcome is possible in equilibrium because different types of players take different equilibrium actions. In the PhD Admissions Game, we could assume that there are a few students who both love economics and actually enjoy writing applications. Those students would always apply in equilibrium, so there would never be a pure pooling equilibrium in which nobody applied, and Bayes' s Rule could always be used. In equilibrium, the university would always accept someone who applied, because applying is never out-of-equilibrium behavior and it always indicates that the applicant is a *Lover*. This approach is especially attractive if the modeller takes the possibility of trembles literally, instead of just using it as a technical tool.

The arguments for different kinds of beliefs can also be applied to Entry Deterrence III, which had two different pooling equilibria and no separating equilibrium. We used passive conjectures in the “plausible” equilibrium. The intuitive criterion would not restrict beliefs at all, because both types would enter if the incumbent’s beliefs were such as to make him collude, and both would stay out if they made him fight. Complete robustness would rule out as an equilibrium the strategy profile in which the entrant stays out regardless of type, because the optimality of staying out depends on the beliefs. It would support the strategy profile in which the entrant enters and out-of-equilibrium beliefs do not matter.

6.3 The Importance of Common Knowledge: Entry Deterrence IV and V

To demonstrate the importance of common knowledge, let us consider two more versions of Entry Deterrence. We will use passive conjectures in both. In Entry Deterrence III, the incumbent was hurt by his ignorance. Entry Deterrence IV will show how he can benefit from it, and Entry Deterrence V will show what can happen when the incumbent has the same information as the entrant but the information is not common knowledge.

Entry Deterrence IV: The Incumbent Benefits from Ignorance

To construct Entry Deterrence IV, let $X = 300$ in Figure 1, so fighting is even more profitable than in Entry Deterrence III but the game is otherwise the same: the entrant knows his type, but the incumbent does not. The following is the unique perfect bayesian equilibrium in pure strategies.¹

Equilibrium for Entry Deterrence IV

Entrant: *Stay Out* |*Weak*, *Stay Out* |*Strong*

Incumbent: *Fight*, Out-of-equilibrium beliefs: $Prob(Strong|Enter) = 0.5$ (passive conjectures)

This equilibrium can be supported by other out-of-equilibrium beliefs, but no equilibrium is possible in which the entrant enters. There is no pooling equilibrium in which both

¹There exists a plausible mixed-strategy equilibrium too: *Entrant: Enter if Strong, Enter with probability $m = 0.2$ if Weak; Incumbent: Collude with probability $n = 0.2$* . The payoff from this is only 150, so if the equilibrium were one in mixed strategies, ignorance would *not* help.

types of entrant enter, because then the incumbent's expected payoff from *Fight* would be 150 ($= 0.5[0] + 0.5[300]$), which is greater than the *Collude* payoff of 50. There is no separating equilibrium, because if only the strong entrant entered and the incumbent always colluded, the weak entrant would be tempted to imitate him and enter as well.

In Entry Deterrence IV, unlike Entry Deterrence III, the incumbent benefits from his own ignorance, because he would always fight entry, even if the payoff were (unknown to himself) just zero. The entrant would very much like to communicate the costliness of fighting, but the incumbent would not believe him, so entry never occurs.

Entry Deterrence V: Lack of Common Knowledge of Ignorance

In Entry Deterrence V, it may happen that both the entrant and the incumbent know the payoff from (*Enter*, *Fight*), but the entrant does not know whether the incumbent knows. The information is known to both players, but is not common knowledge.

Figure 3 depicts this somewhat complicated situation. The game begins with Nature assigning the entrant a type, *Strong* or *Weak* as before. This is observed by the entrant but not by the incumbent. Next, Nature moves again and either tells the incumbent the entrant's type or remains silent. This is observed by the incumbent, but not by the entrant. The four games starting at nodes G_1 to G_4 represent different profiles of payoffs from (*Enter*, *Fight*) and knowledge of the incumbent. The entrant does not know how well informed the incumbent is, so the entrant's information partition is $(\{G_1, G_2\}, \{G_3, G_4\})$.

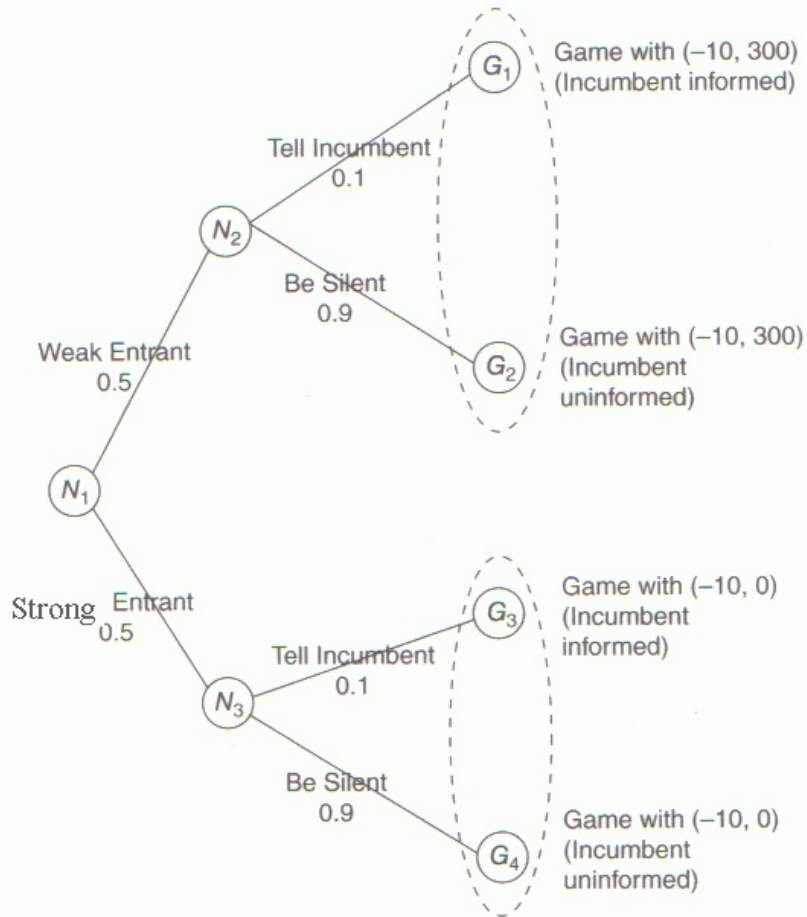


Figure 6.3 Entry Deterrence V

Figure 3 Entry Deterrence V

Equilibrium for Entry Deterrence V

Entrant: *Stay Out*|*Weak*, *Stay Out*|*Strong*

Incumbent: *Fight*|*Nature said "Weak"*, *Collude*|*Nature said "Strong"*, *Fight*|*Nature said nothing*, Out-of- equilibrium beliefs: $Prob(\text{Strong}|\text{Enter}, \text{Nature said nothing}) = 0.5$ (passive conjectures)

Since the entrant puts a high probability on the incumbent not knowing, the entrant should stay out, because the incumbent will fight for either of two reasons. With probability 0.9, Nature has said nothing and the incumbent calculates his expected payoff from *Fight* to be 150, and with probability 0.05 ($= 0.1[0.5]$) Nature has told the incumbent that the entrant is weak and the payoff from *Fight* is 300. Even if the entrant is strong and Nature tells this to the incumbent, the entrant would choose *Stay Out*, because he does not know that the incumbent knows, and his expected payoff from *Enter* would be -5 ($= [0.9][-10] + 0.1[40]$).

If it were common knowledge that the entrant was strong, the entrant would enter and the incumbent would collude. If it is known by both players, but not common knowledge,

the entrant stays out, even though the incumbent would collude if he entered. Such is the importance of common knowledge.

6.4 Incomplete Information in the Repeated Prisoner's Dilemma: The Gang of Four Model

Chapter 5 explored various ways to steer between the Scylla of the Chainstore Paradox and the Charybdis of the Folk Theorem to find a resolution to the problem of repeated games. In the end, uncertainty turned out to make little difference to the problem, but incomplete information was left unexamined in Chapter 5. One might imagine that if the players did not know each others' types, the resulting confusion might allow cooperation. Let us investigate this by adding incomplete information to the finitely repeated Prisoner's Dilemma (whose payoffs are repeated in Table 1) and finding the perfect bayesian equilibria.

Table 1: The Prisoner's Dilemma

		Column	
		<i>Silence</i>	<i>Blame</i>
Row:	<i>Silence</i>	5,5	-5,10
	<i>Blame</i>	10,-5	0,0

Payoffs to: (Row, Column)

One way to incorporate incomplete information would be to assume that a large number of players are irrational, but that a given player does not know whether any other player is of the irrational type or not. In this vein, one might assume that with high probability Row is a player who blindly follows the strategy of Tit-for-Tat. If Column thinks he is playing against a Tit-for-Tat player, his optimal strategy is to *Silence* until near the last period (how near depending on the parameters), and then *Blame*. If he were not certain of this, but the probability were high that he faced a Tit-for-Tat player, Row would choose that same strategy. Such a model begs the question, because it is not the incompleteness of the information that drives the model, but the high probability that one player blindly uses Tit-for-Tat. Tit-for-Tat is not a rational strategy, and to assume that many players use it is to assume away the problem. A more surprising result is that a small amount of incomplete information can make a big difference to the outcome.²

The Gang of Four Model

One of the most important explanations of reputation is that of Kreps, Milgrom, Roberts & Wilson (1982), hereafter referred to as the Gang of Four. In their model, a few players are genuinely unable to play any strategy but Tit-for-Tat, and many players pretend to be of that type. The beauty of the model is that it requires only a small amount of incomplete information, and a low probability γ that player Row is a Tit-for-Tat player.

²Begging the question is not as illegitimate in modelling as in rhetoric, however, because it may indicate that the question is a vacuous one in the first place. If the payoffs of the Prisoner's Dilemma are not those of most of the people one is trying to model, the Chainstore Paradox becomes irrelevant.

It is not unreasonable to suppose that the world contains a few mildly irrational tit-for-tat players, and such behavior is especially plausible among consumers, who are subject to less evolutionary pressure than firms.

It may even be misleading to call Tit-for-Tat “irrational”, because they may just have unusual payoffs, particularly since we will assume that they are rare. The unusual players have a small direct influence, but they matter because other players imitate them. Even if Column knows that with high probability Row is just pretending to be a Tit-for-Tat player, Column does not care what the truth is so long as Row keeps on pretending. Hypocrisy is not only the tribute vice pays to virtue; it can be just as good for deterring misbehavior.

Theorem 6.1: The Gang of Four Theorem

Consider a T -stage, repeated Prisoner’s Dilemma, without discounting but with a probability γ of a Tit-for-Tat player. In any perfect bayesian equilibrium, the number of stages in which either player chooses Blame is less than some number M that depends on γ but not on T .

The significance of the Gang of Four theorem is that while the players do resort to *Blame* as the last period approaches, the number of periods during which they *Blame* is independent of the total number of periods. Suppose $M = 2,500$. If $T = 2,500$, there might be *Blame* every period. But if $T = 10,000$, there are 7,500 periods without a *Blame* move. For reasonable probabilities of the unusual type, the number of periods of cooperation can be much larger. Wilson (unpublished) has set up an entry deterrence model in which the incumbent fights entry (the equivalent of *Silence* above) up to seven periods from the end, although the probability the entrant is of the unusual type is only 0.008.

The Gang of Four Theorem characterizes the equilibrium outcome rather than the equilibrium. Finding perfect bayesian equilibria is difficult and tedious, since the modeller must check all the out-of-equilibrium subgames, as well as the equilibrium path. Modellers usually content themselves with describing important characteristics of the equilibrium strategies and payoffs.

To get a feeling for why Theorem 6.1 is correct, consider what would happen in a 10,001 period game with a probability of 0.01 that Row is playing the Grim Strategy of *Silence* until the first *Blame*, and *Blame* every period thereafter. Using Table 1’s payoffs, a best response for Column to a known Grim player is (*Blame* only in the last period, unless Row chooses *Blame* first, in which case respond with *Blame*). Both players will choose *Silence* until the last period, and Column’s payoff will be 50,010 ($= (10,000)(5) + 10$). Suppose for the moment that if Row is not Grim, he is highly aggressive, and will choose *Blame* every period. If Column follows the strategy just described, the outcome will be (*Blame, Silence*) in the first period and (*Blame, Blame*) thereafter, for a payoff to Column of $-5 (= -5 + (10,000)(0))$. If the probabilities of the two outcomes are 0.01 and 0.99, Column’s expected payoff from the strategy described is 495.15. If instead he follows a strategy of (*Blame* every period), his expected payoff is just 0.1 ($= 0.01(10) + 0.99(0)$). It is clearly in Column’s advantage to take a chance by cooperating with Row, even if Row has a 0.99 probability of following a very aggressive strategy.

The aggressive strategy, however, is not Row’s best response to Column’s strategy. A better response is for Row to choose *Silence* until the second-to-last period, and then to

choose *Blame*. Given that Column is cooperating in the early periods, Row will cooperate also. This argument has not described what the Nash equilibrium actually is, since the iteration back and forth between Row and Column can be continued, but it does show why Column chooses *Silence* in the first period, which is the leverage the argument needs: the payoff is so great if Row is actually the grim player that it is worthwhile for Column to risk a low payoff for one period.

The Gang of Four Theorem provides a way out of the Chainstore Paradox, but it creates a problem of multiple equilibria in much the same way as the infinitely repeated game. For one thing, if the asymmetry is two-sided, so both players might be unusual types, it becomes much less clear what happens in threat games such as Entry Deterrence. Also, what happens depends on which unusual behaviors have positive, if small, probability. Theorem 6.2 says that the modeller can make the average payoffs take any particular values by making the game last long enough and choosing the form of the irrationality carefully.

Theorem 6.2: The Incomplete Information Folk Theorem(Fudenberg & Maskin [1986] p. 547)

For any two-person repeated game without discounting, the modeller can choose a form of irrationality so that for any probability $\epsilon > 0$ there is some finite number of repetitions such that with probability $(1 - \epsilon)$ a player is rational and the average payoffs in some sequential equilibrium are closer than ϵ to any desired payoffs greater than the minimax payoffs.

6.5 The Axelrod Tournament

Another way to approach the repeated Prisoner's Dilemma is through experiments, such as the round robin tournament described by political scientist Robert Axelrod in his 1984 book. Contestants submitted strategies for a 200-repetition Prisoner's Dilemma. Since the strategies could not be updated during play, players could precommit, but the strategies could be as complicated as they wished. If a player wanted to specify a strategy which simulated subgame perfectness by adapting to past history just as a noncommitted player would, he was free to do so, but he could also submit a non-perfect strategy such as Tit-for-Tat or the slightly more forgiving Tit-for-Two-Tats. Strategies were submitted in the form of computer programs that were matched with each other and played automatically. In Axelrod's first tournament, 14 programs were submitted as entries. Every program played every other program, and the winner was the one with the greatest sum of payoffs over all the plays. The winner was Anatol Rapoport, whose strategy was Tit-for-Tat.

The tournament helps to show which strategies are robust against a variety of other strategies in a game with given parameters. It is quite different from trying to find a Nash equilibrium, because it is not common knowledge what the equilibrium is in such a tournament. The situation could be viewed as a game of incomplete information in which Nature chooses the number and cognitive abilities of the players and their priors regarding each other.

After the results of the first tournament were announced, Axelrod ran a second tournament, adding a probability $\theta = 0.00346$ that the game would end each round so as to avoid the Chainstore Paradox. The winner among the 62 entrants was again Anatol Rapoport, and again he used Tit-for-Tat.

Before choosing his tournament strategy, Rapoport had written an entire book on The Prisoner's Dilemma in analysis, experiment, and simulation (Rapoport & Chammah [1965]). Why did he choose such a simple strategy as Tit-for-Tat? Axelrod points out that Tit-for-Tat has three strong points.

1. It never initiates blaming (**niceness**);
2. It retaliates instantly against blaming (**provocability**);
3. It forgives someone who plays *Blame* but then goes back to cooperating (it is **forgiving**).

Despite these advantages, care must be taken in interpreting the results of the tournament. It does not follow that Tit-for-Tat is the best strategy, or that cooperative behavior should always be expected in repeated games.

First, Tit-for-Tat never beats any other strategy in a one-on-one contest. It won the tournament by piling up points through cooperation, having lots of high score plays and very few low score plays. In an elimination tournament, Tit-for-Tat would be eliminated very early, because it scores *high* payoffs but never the *highest* payoff.

Second, the other players' strategies matter to the success of Tit-for-Tat. In neither tournament were the strategies submitted a Nash equilibrium. If a player knew what strategies he was facing, he would want to revise his own. Some of the strategies submitted in the second tournament would have won the first, but they did poorly because the environment had changed. Other programs, designed to try to probe the strategies of their opposition, wasted too many (*Blame, Blame*) episodes on the learning process, but if the games had lasted a thousand repetitions they would have done better.

Third, in a game in which players occasionally blaméd because of trembles, two Tit-for-Tat players facing each other would do very badly. The strategy instantly punishes a blaming player, and it has no provision for ending the punishment phase.

Optimality depends on the environment. When information is complete and the payoffs are all common knowledge, blaming is the only equilibrium outcome. In practically any real-world setting, however, information is slightly incomplete, so cooperation becomes more plausible. Tit-for-Tat is suboptimal for any given environment, but it is robust across environments, and that is its advantage.

6.6 Credit and the Age of the Firm: The Diamond Model

An example of another way to look at reputation is Diamond's model of credit terms, which seeks to explain why older firms get cheaper credit using a game similar to the Gang of Four model. Telser (1966) suggested that predatory pricing would be a credible threat if the incumbent had access to cheaper credit than the entrant, and so could hold out for more periods of losses before going bankrupt. While one might wonder whether this is effective protection against entry—what if the entrant is a large old firm from another industry?—we shall focus on how better-established firms might get cheaper credit.

D. Diamond (1989) aims to explain why old firms are less likely than young firms to default on debt. His model has both adverse selection, because firms differ in type, and moral hazard, because they take hidden actions. The three types of firms, R, S, and RS, are “born” at time zero and borrow to finance projects at the start of each of T periods. We must imagine that there are overlapping generations of firms, so that at any point in time a variety of ages are coexisting, but the model looks at the lifecycle of only one generation. All the players are risk neutral. Type RS firms can choose independently risky projects with negative expected values or safe projects with low but positive expected values. Although the risky projects are worse in expectation, if they are successful the return is much higher than from safe projects. Type R firms can only choose risky projects, and type S firms only safe projects. At the end of each period the projects bring in their profits and loans are repaid, after which new loans and projects are chosen for the next period. Lenders cannot tell which project is chosen or what a firm's current profits are, but they can seize the firm's assets if a loan is not repaid, which always happens if the risky project was chosen and turned out unsuccessfully.

This game foreshadows two other models of credit that will be described in this book, the Repossession Game of section 8.4 and the Stiglitz-Weiss model of section 9.6. Both will be one-shot games in which the bank worried about not being repaid; in the Repossession Game because the borrower did not exert enough effort, and in the Stiglitz-Weiss model because he was of an undesirable type that could not repay. The Diamond model is a mixture of adverse selection and moral hazard: the borrowers differ in type, but some borrowers have a choice of action.

The equilibrium path has three parts. The RS firms start by choosing risky projects. Their downside risk is limited by bankruptcy, but if the project is successful the firm keeps large residual profits after repaying the loan. Over time, the number of firms with access to the risky project (the RS's and R's) diminishes through bankruptcy, while the number of S's remains unchanged. Lenders can therefore maintain zero profits while lowering their interest rates. When the interest rate falls, the value of a stream of safe investment profits minus interest payments rises relative to the expected value of the few periods of risky returns minus interest payments before bankruptcy. After the interest rate has fallen enough, the second phase of the game begins when the RS firms switch to safe projects at a period we will call t_1 . Only the tiny and diminishing group of type R firms continue to choose risky projects. Since the lenders know that the RS firms switch, the interest rate can fall sharply at t_1 . A firm that is older is less likely to be a type R, so it is charged a lower interest rate. Figure 4 shows the path of the interest rate over time.




Figure 6.4 The interest rate over time

Figure 4 The Interest Rate over Time

Towards period T , the value of future profits from safe projects declines and even with a low interest rate the RS's are again tempted to choose risky projects. They do not all switch at once, however, unlike in period t_1 . In period t_1 , if a few RS's had decided to switch to safe projects, the lenders would have been willing to lower the interest rate, which would have made switching even more attractive. If a few firms switch to risky projects at some time t_2 , on the other hand, the interest rate rises and switching to risky projects becomes more attractive—a result that will also be seen in the Lemons model in Chapter 9. Between t_2 and t_3 , the RS's follow a mixed strategy, an increasing number of them choosing risky projects as time passes. The increasing proportion of risky projects causes the interest rate to rise. At t_3 , the interest rate is high enough and the end of the game is close enough that the RS's revert to the pure strategy of choosing risky projects. The interest rate declines during this last phase as the number of RS's diminishes because of failed risky projects.

One might ask, in the spirit of modelling by example, why the model contains three types of firms rather than two. Types S and RS are clearly needed, but why type R? The little extra detail in the game description allows simplification of the equilibrium, because with three types bankruptcy is never out-of-equilibrium behaviour, since the failing firm might be a type R. Bayes's Rule can therefore always be applied, eliminating the problem of ruling out peculiar beliefs and absurd perfect bayesian equilibria.

This is a Gang of Four model but differs from previous examples in an important respect: the Diamond model is not stationary, and as time progresses, some firms of types R and RS go bankrupt, which changes the lenders' payoff functions. Thus, it is not, strictly speaking, a repeated game.

Notes

N6.1 Perfect Bayesian Equilibrium: Entry Deterrence I and II

- Section 4.1 showed that even in games of perfect information, not every subgame perfect equilibrium is trembling-hand perfect. In games of perfect information, however, every subgame perfect equilibrium is a perfect bayesian equilibrium, since no out-of-equilibrium beliefs need to be specified.

N6.2 Refining Perfect Bayesian Equilibrium: The PhD Admissions Game

- Fudenberg & Tirole (1991b) is a careful analysis of the issues involved in defining perfect bayesian equilibrium.
- Section 6.2 is about debatable ways of restricting beliefs such as passive conjectures or equilibrium dominance, but less controversial restrictions are sometimes useful. In a three-player game, consider what happens when Smith and Jones have incomplete information about Brown, and then Jones deviates. If it was Brown himself who had deviated, one might think that the other players might deduce something about Brown's type. But should they update their priors on Brown because Jones has deviated? Especially, should Jones update his beliefs, just because he himself deviated? Passive conjectures seems much more reasonable.

If, to take a second possibility, Brown himself does deviate, is it reasonable for the out-of-equilibrium beliefs to specify that Smith and Jones update their beliefs about Brown in different ways? This seems dubious in light of the Harsanyi doctrine that everyone begins with the same priors.

On the other hand, consider a tremble interpretation of out-of- equilibrium moves. Maybe if Jones trembles and picks the wrong strategy, that really does say something about Brown's type. Jones might tremble more often, for example, if Brown's type is strong than if it is weak. Jones himself might learn from his own trembles. Once we are in the realm of non-bayesian beliefs, it is hard to know what to do without a real-world context.

Dominance and tremble arguments used to rule out Nash equilibria apply to past, present (in simultaneous move games), and future actions of the other player. Belief arguments only depend on past actions, because they rely on the uninformed player observing behavior and interpreting it. Thus, for example, a tremble or weak dominance argument might say a player should take action 1 instead of 2 because although their payoffs are equal, action 2 would lead to a very low payoff if the other player later trembled and chose an unintended action that hurt both of them. An argument based on beliefs would not work in such a game.

- For discussions of the appropriateness of different equilibrium concepts in actual economic models see Rubinstein (1985b) on bargaining, Shleifer & Vishny (1986) on greenmail and D. Hirshleifer & Titman (1990) on tender offers.
- **Exotic refinements.** Binmore (1990) and Kreps (1990b) are booklength treatments of rationality and equilibrium concepts. Van Damme (1989) introduces the curious "money burning" idea of "forward induction."
- **The Beer-Quiche Game** of Cho & Kreps (1987). To illustrate their "intuitive criterion", Cho and Kreps use the Beer-Quiche Game. In this game, Player I might be either weak or

strong in his duelling ability, but he wishes to avoid a duel even if he thinks he can win. Player II wishes to fight a duel only if player I is weak, which has a probability of 0.1. Player II does not know player I's type, but he observes what player I has for breakfast. He knows that weak players prefer quiche for breakfast, while strong players prefer beer. The payoffs are shown in Figure 5.

Figure 5 illustrates a few twists on how to draw an extensive form. It begins with Nature's choice of *Strong* or *Weak* in the middle of the diagram. Player I then chooses whether to breakfast on *beer* or *quiche*. Player II's nodes are connected by a dotted line if they are in the same information set. Player II chooses *Duel* or *Don't*, and payoffs are then received.

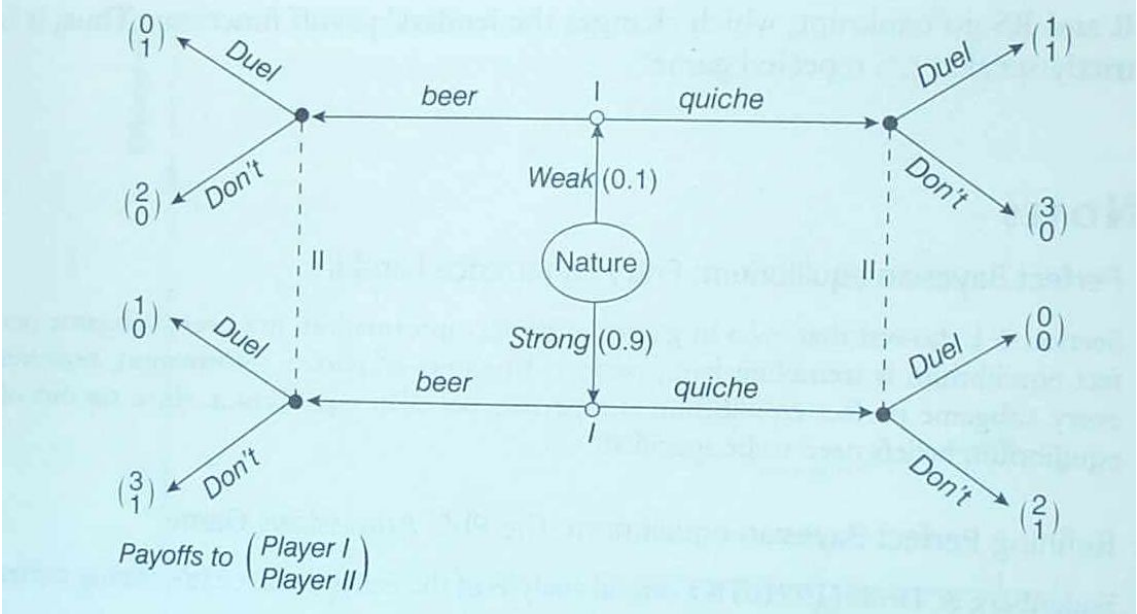


Figure 5 The Beer-Quiche Game

This game has two perfect bayesian equilibrium outcomes, both of which are pooling. In E_1 , player I has beer for breakfast regardless of type, and Player II chooses not to duel. This is supported by the out-of-equilibrium belief that a quiche-eating player I is weak with probability over 0.5, in which case player II would choose to duel on observing quiche. In E_2 , player I has quiche for breakfast regardless of type, and player II chooses not to duel. This is supported by the out-of-equilibrium belief that a beer-drinking player I is weak with probability greater than 0.5, in which case player II would choose to duel on observing beer. Passive conjectures and the intuitive criterion both rule out equilibrium E_2 . According to the reasoning of the intuitive criterion, player I could deviate without fear of a duel by giving the following convincing speech,

I am having beer for breakfast, which ought to convince you I am strong. The only conceivable benefit to me of breakfasting on beer comes if I am strong. I would never wish to have beer for breakfast if I were weak, but if I am strong and this message is convincing, then I benefit from having beer for breakfast.

N6.5 The Axelrod tournament

- Hofstadter (1983) is a nice discussion of the Prisoner's Dilemma and the Axelrod tournament by an intelligent computer scientist who came to the subject untouched by the preconceptions or training of economics. It is useful for elementary economics classes. Axelrod's 1984 book provides a fuller treatment.

Problems

6.1. Cournot Duopoly under Incomplete Information about Costs (hard)

This problem introduces incomplete information into the Cournot model of Chapter 3 and allows for a continuum of player types.

- (a) Modify the Cournot Game of Chapter 3 by specifying that Apex's average cost of production be c per unit, while Brydnox's remains zero. What are the outputs of each firm if the costs are common knowledge? What are the numerical values if $c = 10$?
- (b) Let Apex's cost c be c_{max} with probability θ and 0 with probability $1 - \theta$, so Apex is one of two types. Brydnox does not know Apex's type. What are the outputs of each firm?
- (c) Let Apex's cost c be drawn from the interval $[0, c_{max}]$ using the uniform distribution, so there is a continuum of types. Brydnox does not know Apex's type. What are the outputs of each firm?
- (d) Outputs were 40 for each firm in the zero-cost game in chapter 3. Check your answers in parts (b) and (c) by seeing what happens if $c_{max} = 0$.
- (e) Let $c_{max} = 20$ and $\theta = 0.5$, so the expectation of Apex's average cost is 10 in parts (a), (b), and (c). What are the average outputs for Apex in each case?
- (f) Modify the model of part (b) so that $c_{max} = 20$ and $\theta = 0.5$, but somehow $c = 30$. What outputs do your formulas from part (b) generate? Is there anything this could sensibly model?

Problem 6.2. Limit Pricing (medium) (see Milgrom and Roberts [1982a])

An incumbent firm operates in the local computer market, which is a natural monopoly in which only one firm can survive. The incumbent knows his own operating cost c , which is 20 with probability 0.2 and 30 with probability 0.8.

In the first period, the incumbent can price *Low*, losing 40 in profits, or *High*, losing nothing if his cost is $c = 20$. If his cost is $c = 30$, however, then pricing *Low* he loses 180 in profits. (You might imagine that all consumers have a reservation price that is *High*, so a static monopolist would choose that price whether marginal cost was 20 or 30.)

A potential entrant knows those probabilities, but not the incumbent's exact cost. In the second period, the entrant can enter at a cost of 70, and his operating cost of 25 is common knowledge. If there are two firms in the market, each incurs an immediate loss of 50, but one then drops out and the survivor earns the monopoly revenue of 200 and pays his operating cost. There is no discounting: $r = 0$.

- (a) In a perfect bayesian equilibrium in which the incumbent prices *High* regardless of its costs (a pooling equilibrium), about what do out-of-equilibrium beliefs have to be specified?
- (b) Find a pooling perfect bayesian equilibrium, in which the incumbent always chooses the same price no matter what his costs may be.
- (c) What is a set of out-of-equilibrium beliefs that do not support a pooling equilibrium at a *High* price?

- (d) What is a separating equilibrium for this game?

6.3. Symmetric Information and Prior Beliefs (medium)

In the Expensive-Talk Game of Table 2, the Battle of the Sexes is preceded by a communication move in which the man chooses *Silence* or *Talk*. *Talk* costs 1 payoff unit, and consists of a declaration by the man that he is going to the prize fight. This declaration is just talk; it is not binding on him.

Table 2: Subgame Payoffs in the Expensive-Talk Game

		Woman	
		<i>Fight</i>	<i>Ballet</i>
Man:	<i>Fight</i>	3,1	0,0
	<i>Ballet</i>	0,0	1,3

Payoffs to: (Man, Woman)

- (a) Draw the extensive form for this game, putting the man's move first in the simultaneous-move subgame.
- (b) What are the strategy sets for the game? (Start with the woman's.)
- (c) What are the three perfect pure-strategy equilibrium outcomes in terms of observed actions? (Remember: strategies are not the same thing as outcomes.)
- (d) Describe the equilibrium strategies for a perfect equilibrium in which the man chooses to talk.
- (e) The idea of "forward induction" says that an equilibrium should remain an equilibrium even if strategies dominated in that equilibrium are removed from the game and the procedure is iterated. Show that this procedure rules out SBB as an equilibrium outcome. (See Van Damme [1989]. In fact, this procedure rules out TFF (*Talk, Fight, Fight*) also.)

6.4. Lack of Common Knowledge (medium)

This problem looks at what happens if the parameter values in Entry Deterrence V are changed.

- (a) Why does $Pr(\text{Strong}|\text{Enter}, \text{Nature said nothing}) = 0.95$ not support the equilibrium in Section 6.3?
- (b) Why is the equilibrium in Section 6.3 not an equilibrium if 0.7 is the probability that Nature tells the incumbent?
- (c) Describe the equilibrium if 0.7 is the probability that Nature tells the incumbent. For what out-of-equilibrium beliefs does this remain the equilibrium?

The Repeated Prisoner's Dilemma under Incomplete Information: A Classroom Game for Chapter 6

Consider the Prisoner's Dilemma in Table 3, obtained by adding 8 to each payoff in Chapter 1's Table 2, and identical to Chapter 5's Table 10:

Table 3: The Prisoner's Dilemma

		Column	
		<i>Silence</i>	<i>Blame</i>
Row	<i>Silence</i>	7,7	→ -2, 8
	<i>Blame</i>	8,-2	→ 0,0

Payoffs to: (Row, Column)

This game will be repeated five times, and your objective is to get as high a summed, undiscounted, payoff as possible (*not* just to get a higher summed payoff than anybody else). Remember, too, that there are lots of pairing of Row and Column in the class, so to just beat your immediate opponent would not even be the right tournament strategy.

The instructor will form groups of three students each to represent *Row*, and groups of one student each to represent *Column*. Each *Row* group will play against multiple *Columns*.

The five-repetition games will be different in how *Column* behaves.

Game (i) Complete Information: Column will seek to maximize his payoff according to Table 3.

Game (ii) 80% Tit-for-Tat: With 20% probability, Column will seek to maximize his payoff according to Table 3. With 80% probability, Column is a "Tit-for-Tat Player" and must use the strategy of "Tit-for-Tat," starting with *Silence* in Round 1 and after that imitating what Row did in the previous round.

Game (iii) 10% Tit-for-Tat: With 90% probability, Column will seek to maximize his payoff according to Table 3. With 10% probability, Column is a "Tit-for-Tat Player" and must use the strategy of "Tit-for-Tat," starting with *Silence* in Round 1 and after that imitating what Row did in the previous round. The identities of the Game (ii).

The probabilities are independent, so although in Game(ii) the most likely outcome is that 8 of 10 Column players use tit-for-tat, it is possible that 7 or 9 do, or even (improbably) 0 or 10.