

2 Information

23 November 2005. Eric Rasmusen, Erasmuse@indiana.edu. <http://www.rasmusen.org/>.

2.1 The Strategic and Extensive Forms of a Game

If half of strategic thinking is predicting what the other player will do, the other half is figuring out what he knows. Most of the games in Chapter 1 assumed that the moves were simultaneous, so the players did not have a chance to learn each other's private information by observing each other. Information becomes central as soon as players move in sequence. The important difference, in fact, between simultaneous-move games and sequential-move games is that in sequential-move games the second player acquires the information on how the first player moved before he must make his own decision.

Section 2.1 shows how to use the strategic form and the extensive form to describe games with sequential moves. Section 2.2 shows how the extensive form, or game tree, can be used to describe the information available to a player at each point in the game. Section 2.3 classifies games based on the information structure. Section 2.4 shows how to redraw games with incomplete information so that they can be analyzed using the Harsanyi transformation, and derives Bayes's Rule for combining a player's prior beliefs with information which he acquires in the course of the game. Section 2.5 concludes the chapter with the Png Settlement Game, an example of a moderately complex sequential-move game.

The Strategic Form and the Outcome Matrix

Games with moves in sequence require more care in presentation than single-move games. In Section 1.4 we used the 2-by-2 form, which for the game Ranked Coordination is shown in Table 1.

Table 1: Ranked Coordination

		Jones	
		<i>Large</i>	<i>Small</i>
	<i>Large</i>	2,2	← −1, −1
Smith		↑	↓
	<i>Small</i>	−1, −1	→ 1,1

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

Because strategies are the same as actions in Ranked Coordination and the outcomes are simple, the 2-by-2 form in Table 1 accomplishes two things: it relates strategy profiles to payoffs, and action profiles to outcomes. These two mappings are called the strategic form and the outcome matrix, and in more complicated games they are distinct from each other. The strategic form shows what payoffs result from each possible strategy profile, while the outcome matrix shows what outcome results from each possible action profile.

The definitions below use n to denote the number of players, k the number of variables in the outcome vector, p the number of strategy profiles, and q the number of action profiles.

The **strategic form** (or **normal form**) consists of

1 All possible strategy profiles s^1, s^2, \dots, s^p .

2 Payoff functions mapping s^i onto the payoff n -vector π^i , ($i = 1, 2, \dots, p$).

The **outcome matrix** consists of

1 All possible action profiles a^1, a^2, \dots, a^q .

2 Outcome functions mapping a^i onto the outcome k -vector z^i , ($i = 1, 2, \dots, q$).

Consider the following game based on Ranked Coordination, which we will call Follow-the-Leader I since we will create several variants of the game. The difference from Ranked Coordination is that Smith moves first, committing himself to a certain disk size no matter what size Jones chooses. The new game has an outcome matrix identical to Ranked Coordination, but its strategic form is different because Jones's strategies are no longer single actions. Jones's strategy set has four elements,

$$\left\{ \begin{array}{l} \text{(If Smith chose } Large, \text{ choose } Large; \text{ if Smith chose } Small, \text{ choose } Large), \\ \text{(If Smith chose } Large, \text{ choose } Large; \text{ if Smith chose } Small, \text{ choose } Small), \\ \text{(If Smith chose } Large, \text{ choose } Small; \text{ if Smith chose } Small, \text{ choose } Large), \\ \text{(If Smith chose } Large, \text{ choose } Small; \text{ if Smith chose } Small, \text{ choose } Small) \end{array} \right\}$$

which we will abbreviate as

$$\left\{ \begin{array}{l} (L|L, L|S), \\ (L|L, S|S), \\ (S|L, L|S), \\ (S|L, S|S) \end{array} \right\}$$

Follow-the-Leader I illustrates how adding a little complexity can make the strategic form too obscure to be very useful. The strategic form is shown in Table 2, with equilibria boldfaced and labelled E_1 , E_2 , and E_3 .

Table 2: Follow-the-Leader I

		Jones			
		J_1 $L L, L S$	J_2 $L L, S S$	J_3 $S L, L S$	J_4 $S L, S S$
Smith	$S_1 : Large$	[2], [2] (E_1)	[2], [2] (E_2)	[-1], [-1] -1	-1, -1
	$S_2 : Small$	-1, -1	1, [1]	[-1], [-1] -1	[1], [1] (E_3)

Payoffs to: (Smith, Jones). Best-response payoffs are boxed (with dashes, if weak)

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

Consider why E_1 , E_2 , and E_3 are Nash equilibria. In Equilibrium E_1 , Jones will respond with *Large* regardless of what Smith does, so Smith quite happily chooses *Large*. Jones would be irrational to choose *Large* if Smith chose *Small* first, but that event never happens in equilibrium. In Equilibrium E_2 , Jones will choose whatever Smith chose, so Smith chooses *Large* to make the payoff 2 instead of 1. In Equilibrium E_3 , Smith chooses *Small* because he knows that Jones will respond with *Small* whatever he does, and Jones is willing to respond with *Small* because Smith chooses *Small* in equilibrium. Equilibria E_1 and E_3 are not completely sensible, because the choices *Large*|*Small* (as specified in E_1) and *Small*|*Large* (as specified in E_3) would reduce Jones's payoff if the game ever reached a point where he had to actually play them. Except for a little discussion in connection with Figure 1, however, we will defer to Chapter 4 the discussion of how to redefine the equilibrium concept to rule them out.

The Order of Play

The “normal form” is rarely used in modelling games of any complexity. Already, in Section 1.1, we have seen an easier way to model a sequential game: the *order of play*. For it Follow the Leader I, this would be:

- 1 Smith chooses his disk size to be either *Large* or *Small*.
- 2 Jones chooses his disk size to be either *Large* or *Small*.

The reason I have retained the concept of the normal form in this edition is that it reinforces the idea of laying out all the possible strategies and comparing their payoffs. The order of play, however, gives us a better way to describe games, as I will explain next.

The Extensive Form and the Game Tree

Two other ways to describe a game are the extensive form and the game tree. First we need to define their building blocks. As you read the definitions, you may wish to refer to Figure 1 as an example.

A **node** is a point in the game at which some player or Nature takes an action, or the game ends.

A **successor** to node X is a node that may occur later in the game if X has been reached.

A **predecessor** to node X is a node that must be reached before X can be reached.

A **starting node** is a node with no predecessors.

An **end node** or **end point** is a node with no successors.

A **branch** is one action in a player's action set at a particular node.

A **path** is a sequence of nodes and branches leading from the starting node to an end node.

These concepts can be used to define the extensive form and the game tree.

The **extensive form** is a description of a game consisting of

1 A configuration of nodes and branches running without any closed loops from a single starting node to its end nodes.

2 An indication of which node belongs to which player.

3 The probabilities that Nature uses to choose different branches at its nodes.

4 The information sets into which each player's nodes are divided.

5 The payoffs for each player at each end node.

The **game tree** is the same as the extensive form except that (5) is replaced with

5' The outcomes at each end node.

“Game tree” is a looser term than “extensive form.” If the outcome is defined as the payoff profile, one payoff for each player, then the extensive form is the same as the game tree.

The extensive form for Follow-the-Leader I is shown in Figure 1. We can see why Equilibria E_1 and E_3 of Table 2 are unsatisfactory even though they are Nash equilibria. If the game actually reached nodes J_1 or J_2 , Jones would have dominant actions, *Small* at J_1 and *Large* at J_2 , but E_1 and E_3 specify other actions at those nodes. In Chapter 4 we will return to this game and show how the Nash concept can be refined to make E_2 the only equilibrium.

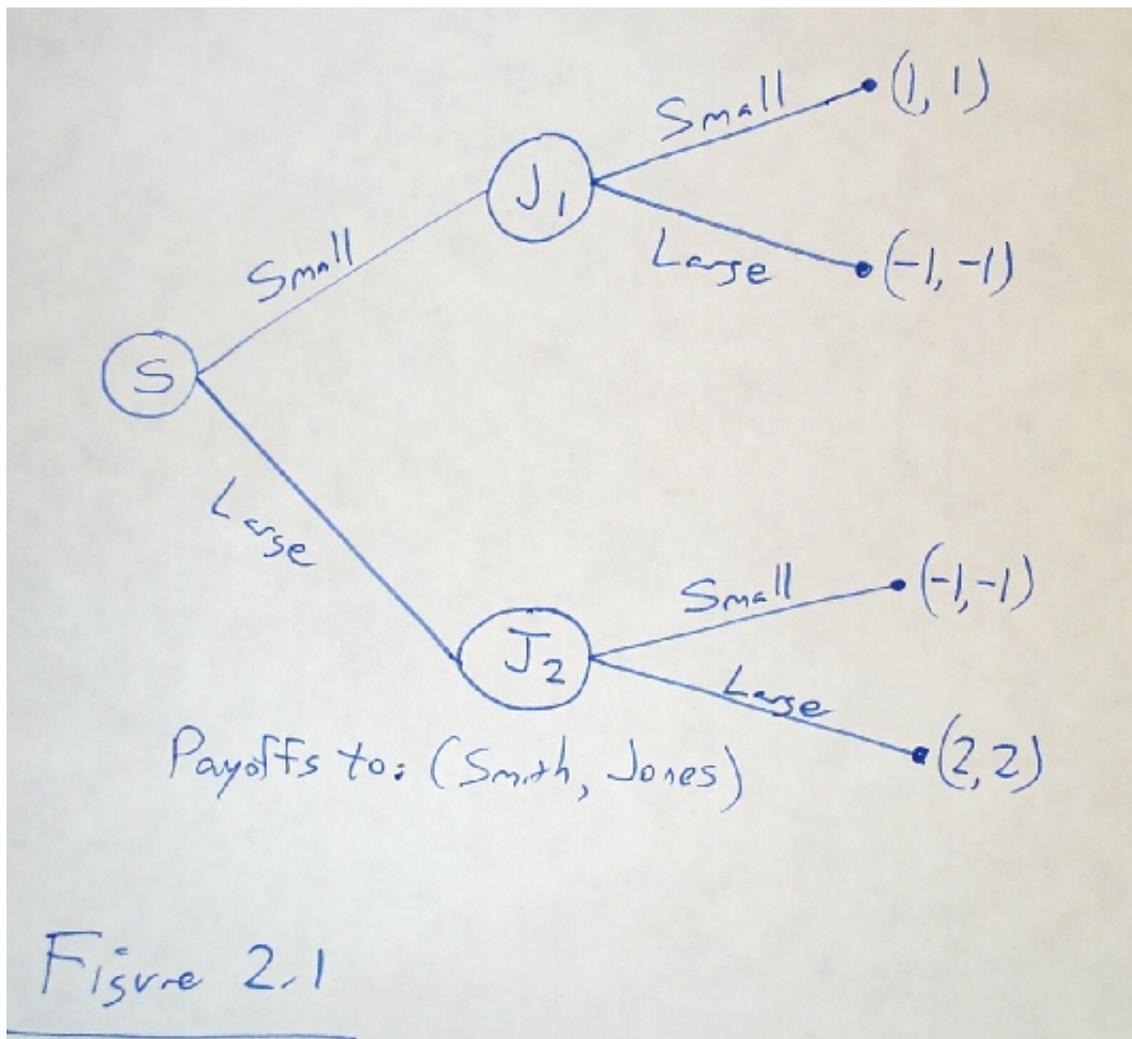


Figure 1: Follow-the-Leader I in Extensive Form

The extensive form for *Ranked Coordination*, shown in Figure 2, adds dotted lines to the extensive form for *Follow-the-Leader I*. Each player makes a single decision between two actions. The moves are simultaneous, which we show by letting Smith move first, but not letting Jones know how he moved. The dotted line shows that Jones's knowledge stays the same after Smith moves. All Jones knows is that the game has reached some node within the information set defined by the dotted line; he does not know the exact node reached.

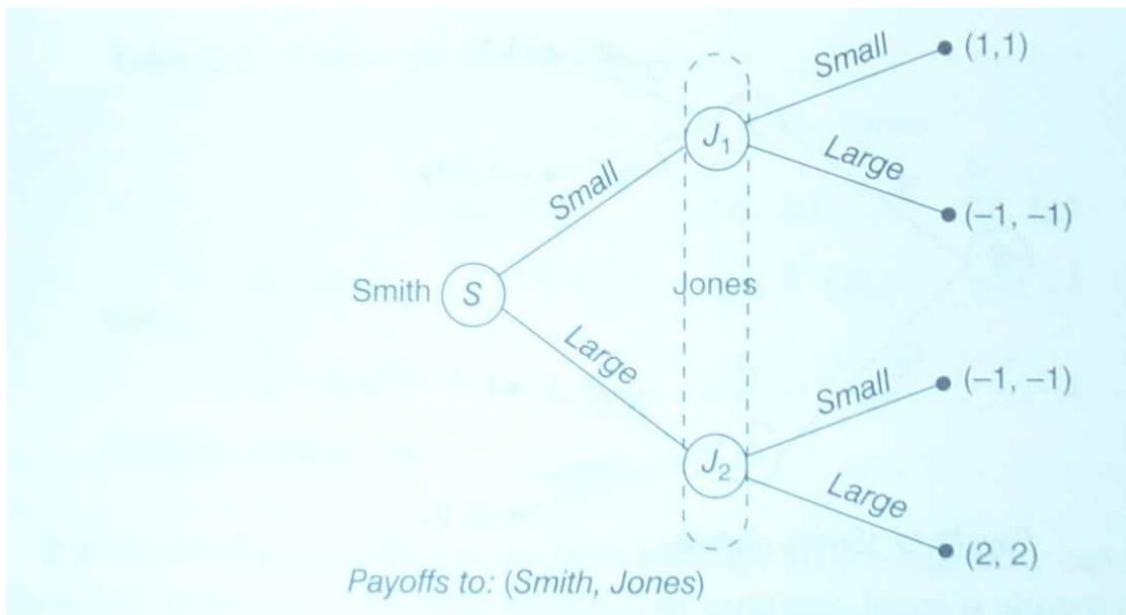


Figure 2: Ranked Coordination in Extensive Form

The Time Line

The **time line**, a line showing the order of events, is another way to describe games. Time lines are particularly useful for games with continuous strategies, exogenous arrival of information, and multiple periods, games that are frequently used in the accounting and finance literature. A typical time line is shown in Figure 3a, which represents a game that will be described in Section 11.5.

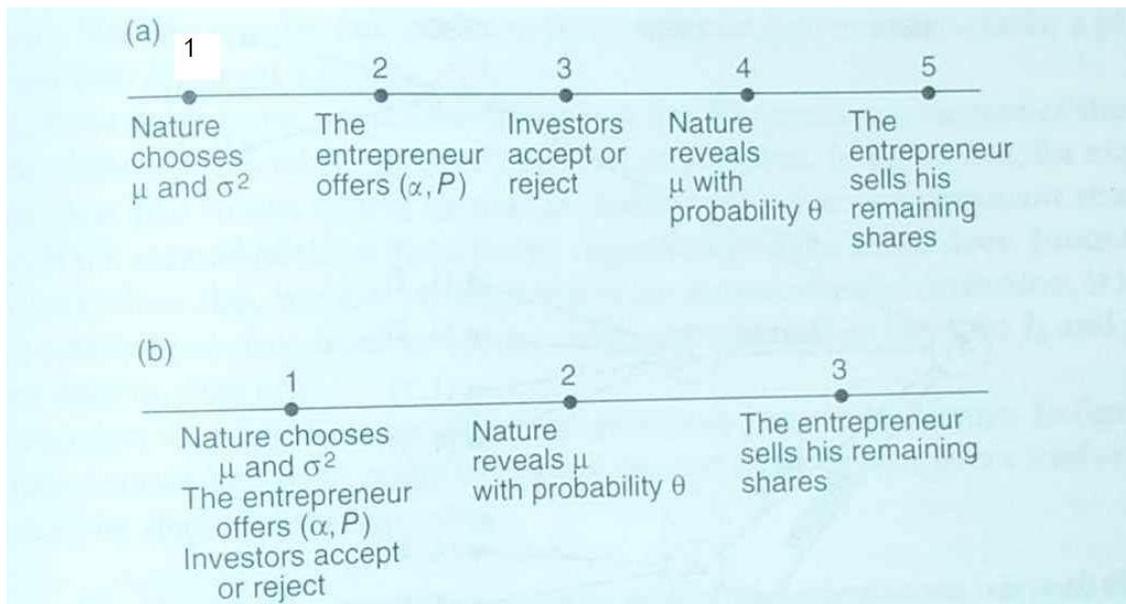


Figure 3: The Time Line for Stock Underpricing: (a) A Good Time Line; (b) A Bad Time Line

The time line illustrates the order of actions and events, not necessarily the passage of time. Certain events occur in an instant, others over an interval. In Figure 3a, events 2 and 3 occur immediately after event 1, but events 4 and 5 might occur ten years later. We sometimes refer to the sequence in which decisions are made as **decision time** and the interval over which physical actions are taken as **real time**. A major difference is that players put higher value on payments received earlier in real time because of time preference (on which see the appendix).

A common and bad modelling habit is to restrict the use of the dates on the time line to separating events in real time. Events 1 and 2 in Figure 2.3a are not separated by real time: as soon as the entrepreneur learns the project's value, he offers to sell stock. The modeller might foolishly decide to depict his model by a picture like Figure 3b in which both events happen at date 1. Figure 3b is badly drawn, because readers might wonder which event occurs first or whether they occur simultaneously. In more than one seminar, 20 minutes of heated and confusing debate could have been avoided by 10 seconds care to delineate the order of events.

2.2: Information Sets

A game's information structure, like the order of its moves, is often obscured in the strategic form. During the Watergate affair, Senator Baker became famous for the question "How much did the President know, and when did he know it?". In games, as in scandals, these are the big questions. To make this precise, however, requires technical definitions so that one can describe who knows what, and when. This is done using the "information set," the set of nodes a player thinks the game might have reached, as the basic unit of knowledge.

Player i 's information set ω_i at any particular point of the game is the set of different nodes in the game tree that he knows might be the actual node, but between which he cannot distinguish by direct observation.

As defined here, the information set for player i is a set of nodes belonging to one player but on different paths. This captures the idea that player i knows whose turn it is to move, but not the exact location the game has reached in the game tree. Historically, player i 's information set has been defined to include only nodes at which player i moves, which is appropriate for single-person decision theory, but leaves a player's knowledge undefined for most of any game with two or more players. The broader definition allows comparison of information across players, which under the older definition is a comparison of apples and oranges.

In the game in Figure 4, Smith moves at node S_1 in 1984 and Jones moves at nodes $J_1, J_2, J_3,$ and J_4 in 1985 or 1986. Smith knows his own move, but Jones can tell only whether Smith has chosen the moves which lead to $J_1, J_2,$ or "other"; he cannot distinguish between J_3 and J_4 . If Smith has chosen the move leading to J_3 , his own information set is simply $\{J_3\}$, but Jones's information set is $\{J_3, J_4\}$.

One way to show information sets on a diagram is to put dashed lines around or between nodes in the same information set. The resulting diagrams can be very cluttered, so it is often more convenient to draw dashed lines around the information set of just the

player making the move at a node. The dashed lines in Figure 4 show that J_3 and J_4 are in the same information set for Jones, even though they are in different information sets for Smith. An expressive synonym for information set which is based on the appearance of these diagrams is “cloud”: one would say that nodes J_3 and J_4 are in the same cloud, so that while Jones can tell that the game has reached that cloud, he cannot pierce the fog to tell exactly which node has been reached.

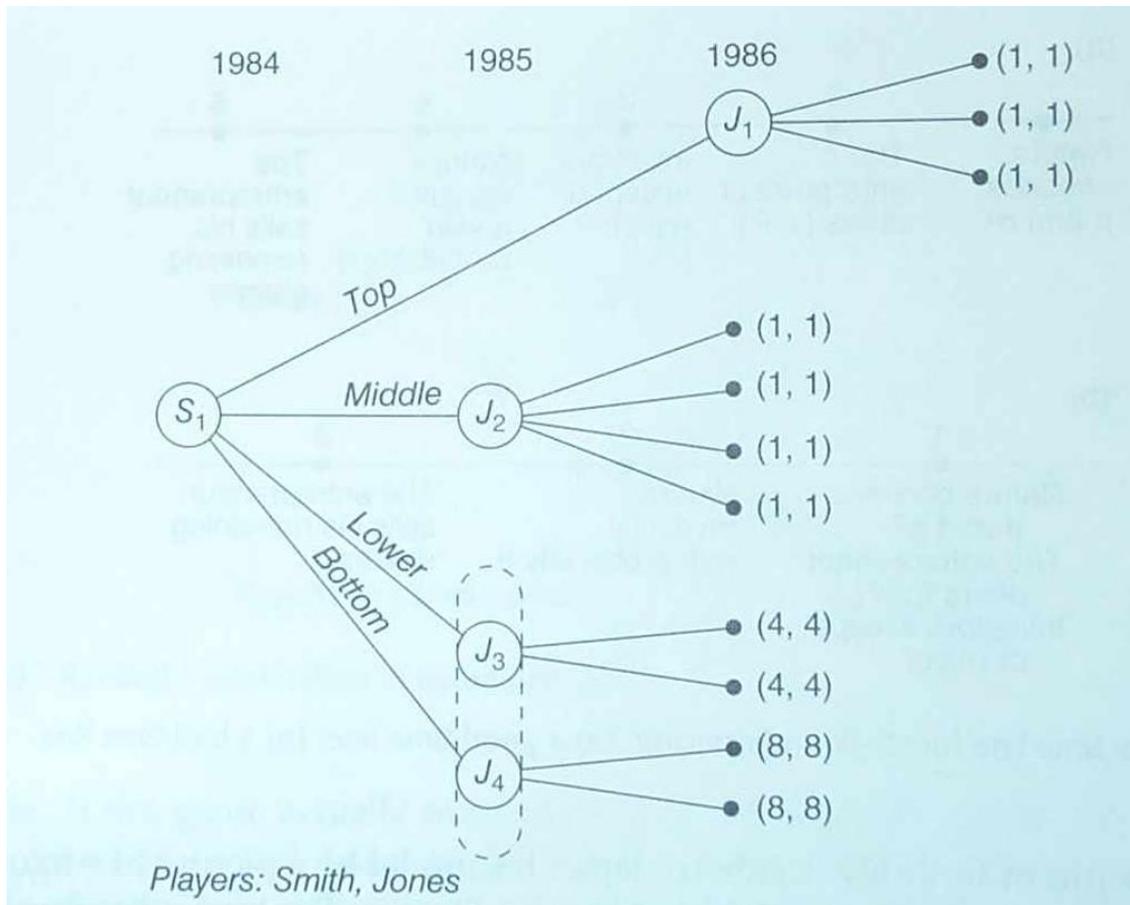


Figure 4: Information Sets and Information Partitions.

One node cannot belong to two different information sets of a single player. If node J_3 belonged to information sets $\{J_2, J_3\}$ and $\{J_3, J_4\}$ (unlike in Figure 4), then if the game reached J_3 , Jones would not know whether he was at a node in $\{J_2, J_3\}$ or a node in $\{J_3, J_4\}$ — which would imply that they were really the same information set.

If the nodes in one of Jones’s information sets are nodes at which he moves, his action set must be the same at each node, because he knows his own action set (though his actions might differ later on in the game depending on whether he advances from J_3 or J_4). Jones has the same action sets at nodes J_3 and J_4 , because if he had some different action available at J_3 he would know he was there and his information set would reduce to just $\{J_3\}$. For the same reason, nodes J_1 and J_2 could not be put in the same information set; Jones must know whether he has three or four moves in his action set. We also require end nodes to be in different information sets for a player if they yield him different payoffs.

With these exceptions, we do not include in the information structure of the game any information acquired by a player’s rational deductions. In Figure 4, for example, it seems clear that Smith would choose *Bottom*, because that is a dominant strategy — his payoff is 8 instead of the 4 from *Lower*, regardless of what Jones does. Jones should be able to deduce this, but even though this is an uncontroversial deduction, it is none the less a deduction, not an observation, so the game tree does not split J_3 and J_4 into separate information sets.

Information sets also show the effects of unobserved moves by Nature. In Figure 4, if the initial move had been made by Nature instead of by Smith, Jones’s information sets would be depicted the same way.

Player i ’s information partition is a collection of his information sets such that

- 1 *Each path is represented by one node in a single information set in the partition, and*
- 2 *The predecessors of all nodes in a single information set are in one information set.*

The information partition represents the different positions that the player knows he will be able to distinguish from each other at a given stage of the game, carving up the set of all possible nodes into the subsets called information sets. One of Smith’s information partitions is $(\{J_1\}, \{J_2\}, \{J_3\}, \{J_4\})$. The definition rules out information set $\{S_1\}$ being in that partition, because the path going through S_1 and J_1 would be represented by two nodes. Instead, $\{S_1\}$ is a separate information partition, all by itself. The information partition refers to a stage of the game, not chronological time. The information partition $(\{J_1\}, \{J_2\}, \{J_3, J_4\})$ includes nodes in both 1985 and 1986, but they are all immediate successors of node S_1 .

Jones has the information partition $(\{J_1\}, \{J_2\}, \{J_3, J_4\})$. There are two ways to see that his information is worse than Smith’s. First is the fact that one of his information sets, $\{J_3, J_4\}$, contains *more* elements than Smith’s, and second, that one of his information partitions, $(\{J_1\}, \{J_2\}, \{J_3, J_4\})$, contains *fewer* elements.

Table 3 shows a number of different information partitions for this game. Partition I is Smith’s partition and partition II is Jones’s partition. We say that partition II is **coarser**, and partition I is **finer**. A profile of two or more of the information sets in a partition, which reduces the number of information sets and increases the numbers of nodes in one or more of them is a **coarsening**. A splitting of one or more of the information sets in a partition, which increases the number of information sets and reduces the number of nodes in one or more of them, is a **refinement**. Partition II is thus a coarsening of partition I, and partition I is a refinement of partition II. The ultimate refinement is for each information set to be a **singleton**, containing one node, as in the case of partition I. As in bridge, having a singleton can either help or hurt a player. The ultimate coarsening is for a player not to be able to distinguish between any of the nodes, which is partition III in Table 3.¹

A finer information partition is the formal definition for “better information.” Not all information partitions are refinements or coarsenings of each other, however, so not all information partitions can be ranked by the quality of their information. In particular,

¹Note, however, that partitions III and IV are not really allowed in this game, because Jones could tell the node from the actions available to him, as explained earlier.

just because one information partition contains more information sets does not mean it is a refinement of another information partition. Consider partitions II and IV in Figure 3. Partition II separates the nodes into three information sets, while partition IV separates them into just two information sets. Partition IV is not a coarsening of partition II, however, because it cannot be reached by combining information sets from partition II, and one cannot say that a player with partition IV has worse information. If the node reached is J_1 , partition II gives more precise information, but if the node reached is J_4 , partition IV gives more precise information.

<i>Nodes</i>	I	II	III	IV
J_1	$\{J_1\}$	$\{J_1\}$	$\left\{ \begin{array}{c} J_1 \\ J_2 \\ J_3 \\ J_4 \end{array} \right\}$	$\left\{ \begin{array}{c} J_1 \\ J_2 \\ J_3 \\ J_4 \end{array} \right\}$
J_2	$\{J_2\}$	$\{J_2\}$		
J_3	$\{J_3\}$	$\{J_3\}$		
J_4	$\{J_4\}$	$\{J_4\}$		

Table 3: Information Partitions

Information quality is defined independently of its utility to the player: it is possible for a player's information to improve and for his equilibrium payoff to fall as a result. Game theory has many paradoxical models in which a player prefers having worse information, not a result of wishful thinking, escapism, or blissful ignorance, but of cold rationality. Coarse information can have a number of advantages. (a) It may permit a player to engage in trade because other players do not fear his superior information. (b) It may give a player a stronger strategic position because he usually has a strong position and is better off not knowing that in a particular realization of the game his position is weak. Or, (c) as in the more traditional economics of uncertainty, poor information may permit players to insure each other.

I will wait till later chapters to discuss points (a) and (b), the strategic advantages of poor information (go to Section 6.3 on entry deterrence and Chapter 9 on used cars if you feel impatient), but it is worth pausing here to think about point (c), the insurance advantage. Consider the following example, which will illustrate that even when information is symmetric and behavior is nonstrategic, better information in the sense of a finer information partition, can actually reduce everybody's utility.

Suppose Smith and Jones, both risk averse, work for the same employer, and both know that one of them chosen randomly will be fired at the end of the year while the other will be promoted. The one who is fired will end with a wealth of 0 and the one who is

promoted will end with 100. The two workers will agree to insure each other by pooling their wealth: they will agree that whoever is promoted will pay 50 to whoever is fired. Each would then end up with a guaranteed utility of $U(50)$. If a helpful outsider offers to tell them who will be fired before they make their insurance agreement, they should cover their ears and refuse to listen. Such a refinement of their information would make both worse off, in expectation, because it would wreck the possibility of the two of them agreeing on an insurance arrangement. It would wreck the possibility because if they knew who would be promoted, the lucky worker would refuse to pool with the unlucky one. Each worker's expected utility with no insurance but with someone telling them what will happen is $.5 * U(0) + .5 * U(100)$, which is less than $1.0 * U(50)$ if they are risk averse. They would prefer not to know, because better information would reduce the expected utility of both of them.

Common Knowledge

We have been implicitly assuming that the players know what the game tree looks like. In fact, we have assumed that the players also know that the other players know what the game tree looks like. The term “common knowledge” is used to avoid spelling out the infinite recursion to which this leads.

*Information is **common knowledge** if it is known to all the players, if each player knows that all the players know it, if each player knows that all the players know that all the players know it, and so forth ad infinitum.*

Because of this recursion (the importance of which will be seen in Section 6.3), the assumption of common knowledge is stronger than the assumption that players have the same beliefs about where they are in the game tree. Hirshleifer & Riley (1992, p. 169) use the term **concordant beliefs** to describe a situation where players share the same belief about the probabilities that Nature has chosen different states of the world, but where they do not necessarily know they share the same beliefs. (Brandenburger [1992] uses the term **mutual knowledge** for the same idea.)

For clarity, models are set up so that information partitions are common knowledge. Every player knows how precise the other players' information is, however ignorant he himself may be as to which node the game has reached. Modelled this way, the information partitions are independent of the equilibrium concept. Making the information partitions common knowledge is important for clear modelling, and restricts the kinds of games that can be modelled less than one might think. This will be illustrated in Section 2.4 when the assumption will be imposed on a situation in which one player does not even know which of three games he is playing.

2.3 Perfect, Certain, Symmetric, and Complete Information

We categorize the information structure of a game in four different ways, so a particular game might have perfect, complete, certain, and symmetric information. The categories are summarized in Table 4.

Information category	Meaning
Perfect	Each information set is a singleton
Certain	Nature does not move after any player moves
Symmetric	No player has information different from other players when he moves, or at the end nodes
Complete	Nature does not move first, or her initial move is observed by every player

Table 4: Information Categories

The first category divides games into those with perfect and those with imperfect information.

*In a game of **perfect information** each information set is a singleton. Otherwise the game is one of **imperfect information**.*

The strongest informational requirements are met by a game of perfect information, because in such a game each player always knows exactly where he is in the game tree. No moves are simultaneous, and all players observe Nature’s moves. Ranked Coordination is a game of imperfect information because of its simultaneous moves, but Follow-the-Leader I is a game of perfect information. Any game of incomplete or asymmetric information is also a game of imperfect information.

*A game of **certainty** has no moves by Nature after any player moves. Otherwise the game is one of **uncertainty**.*

The moves by Nature in a game of uncertainty may or may not be revealed to the players immediately. A game of certainty can be a game of perfect information if it has no simultaneous moves. The notion “game of uncertainty” is new with this book, but I doubt it would surprise anyone. The only quirk in the definition is that it allows an initial move by Nature in a game of certainty, because in a game of incomplete information Nature moves first to select a player’s “type.” Most modellers do not think of this situation as uncertainty.

We have already talked about information in Ranked Coordination, a game of imperfect, complete, and symmetric information with certainty. The Prisoner’s Dilemma falls into the same categories. Follow-the-Leader I, which does not have simultaneous moves, is a game of perfect, complete, and symmetric information with certainty.

We can easily modify Follow-the-Leader I to add uncertainty, creating the game Follow-the-Leader II (Figure 5). Imagine that if both players pick *Large* for their disks, the market yields either zero profits or very high profits, depending on the state of demand, but demand would not affect the payoffs in any other strategy profile. We can quantify

this by saying that if $(Large, Large)$ is picked, the payoffs are $(10,10)$ with probability 0.2, and $(0,0)$ with probability 0.8, as shown in Figure 5.

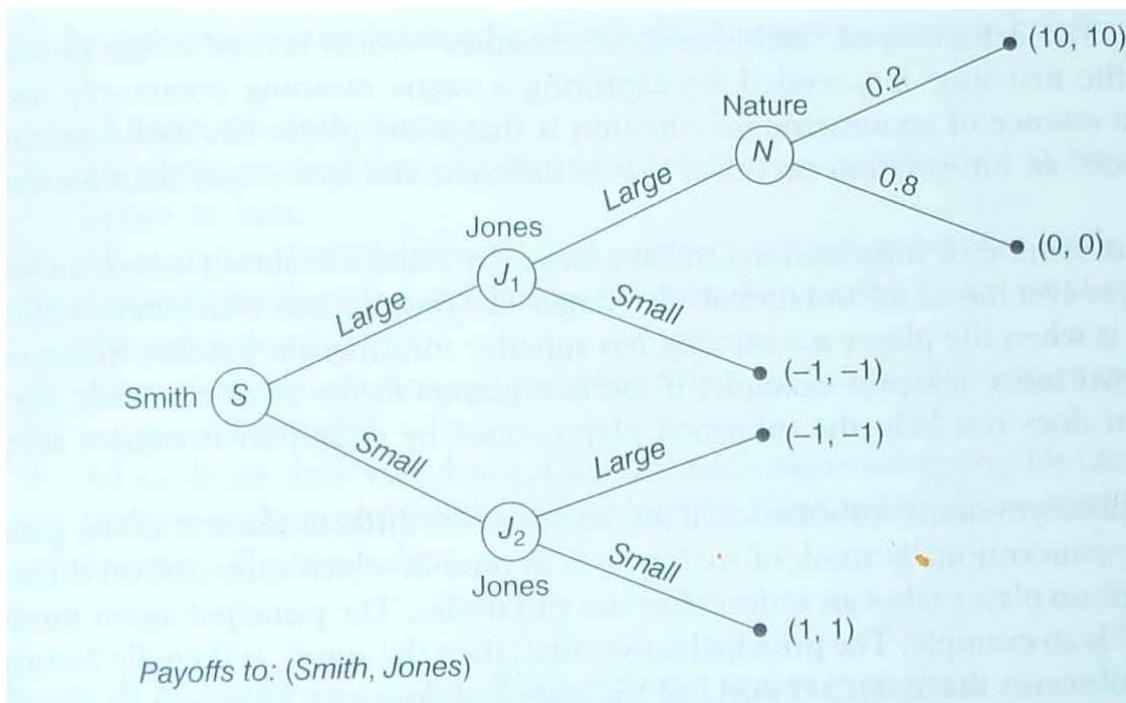


Figure 5: Follow-the-Leader II

When players face uncertainty, we need to specify how they evaluate their uncertain future payoffs. The obvious way to model their behavior is to say that the players maximize the expected values of their utilities. Players who behave in this way are said to have **von Neumann-Morgenstern utility functions**, a name chosen to underscore von Neumann & Morgenstern's (1944) development of a rigorous justification of such behavior.

Maximizing their expected utilities, the players would behave exactly the same as in Follow-the-Leader I. Often, a game of uncertainty can be transformed into a game of certainty without changing the equilibrium, by eliminating Nature's moves and changing the payoffs to their expected values based on the probabilities of Nature's moves. Here we could eliminate Nature's move and replace the payoffs 10 and 0 with the single payoff 2 ($= 0.2[10] + 0.8[0]$). This cannot be done, however, if the actions available to a player depend on Nature's moves, or if information about Nature's move is asymmetric.

The players in Figure 5 might be either risk averse or risk neutral. Risk aversion is implicitly incorporated in the payoffs because they are in units of utility, not dollars. When players maximize their expected utility, they are not necessarily maximizing their expected dollars. Moreover, the players can differ in how they map money to utility. It could be that $(0,0)$ represents $(\$0, \$5,000)$, $(10,10)$ represents $(\$100,000, \$100,000)$, and $(2,2)$, the expected utility, could here represent a non-risky $(\$3,000, \$7,000)$.

*In a game of **symmetric information**, a player's information set at 1 any node where he chooses an action, or*

2 an end node

*contains at least the same elements as the information sets of every other player. Otherwise the game is one of **asymmetric information**.*

In a game of asymmetric information, the information sets of players differ in ways relevant to their behavior, or differ at the end of the game. Such games have imperfect information, since information sets which differ across players cannot be singletons. The definition of “asymmetric information” which is used in the present book for the first time is intended for capturing a vague meaning commonly used today. The essence of asymmetric information is that some player has useful **private information**: an information partition that is different and not worse than another player’s.

A game of symmetric information can have moves by Nature or simultaneous moves, but no player ever has an informational advantage. The one point at which information may differ is when the player *not* moving has superior information because he knows what his own move *was*; for example, if the two players move simultaneously. Such information does not help the informed player, since by definition it cannot affect his move.

A game has asymmetric information if information sets differ at the end of the game because we conventionally think of such games as ones in which information differs, even though no player takes an action after the end nodes. The principal-agent model of Chapter 7 is an example. The principal moves first, then the agent, and finally Nature. The agent observes the agent’s move, but the principal does not, although he may be able to deduce it. This would be a game of symmetric information except for the fact that information continues to differ at the end nodes.

*In a game of **incomplete information**, Nature moves first and is unobserved by at least one of the players. Otherwise the game is one of **complete information**.*

A game with incomplete information also has imperfect information, because some player’s information set includes more than one node. Two kinds of games have complete but imperfect information: games with simultaneous moves, and games where, late in the game, Nature makes moves not immediately revealed to all players.

Many games of incomplete information are games of asymmetric information, but the two concepts are not equivalent. If there is no initial move by Nature, but Smith takes a move unobserved by Jones, and Smith moves again later in the game, the game has asymmetric but complete information. The principal-agent games of Chapter 7 are again examples: the agent knows how hard he worked, but his principal never learns, not even at the end nodes. A game can also have incomplete but symmetric information: let Nature, unobserved by either player, move first and choose the payoffs for *(Confess, Confess)* in the Prisoner’s Dilemma to be either $(-6, -6)$ or $(-100, -100)$.

Harris & Holmstrom (1982) have a more interesting example of incomplete but symmetric information: Nature assigns different abilities to workers, but when workers are young their ability is known neither to employers nor to themselves. As time passes, the abilities become common knowledge, and if workers are risk averse and employers are risk neutral, the model shows that equilibrium wages are constant or rising over time.

Poker Examples of Information Classification

In the game of poker, the players make bets on who will have the best hand of cards at the end, where a ranking of hands has been pre-established. How would the following rules for behavior before betting be classified? (Answers are in note N2.3)

1. All cards are dealt face up.
2. All cards are dealt face down, and a player cannot look even at his own cards before he bets.
3. All cards are dealt face down, and a player can look at his own cards.
4. All cards are dealt face up, but each player then scoops up his hand and secretly discards one card.
5. All cards are dealt face up, the players bet, and then each player receives one more card face up.
6. All cards are dealt face down, but then each player scoops up his cards without looking at them and holds them against his forehead so all the *other* players can see them (Indian poker).

2.4 The Harsanyi Transformation and Bayesian Games

The Harsanyi Transformation: Follow-the-Leader III

The term “incomplete information” is used in two quite different senses in the literature, usually without explicit definition. The definition in Section 2.3 is what economists commonly *use*, but if asked to *define* the term, they might come up with the following, older, definition.

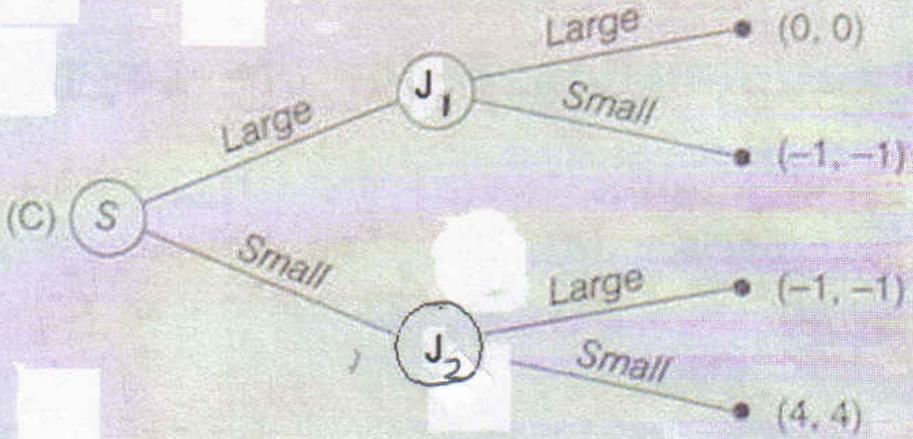
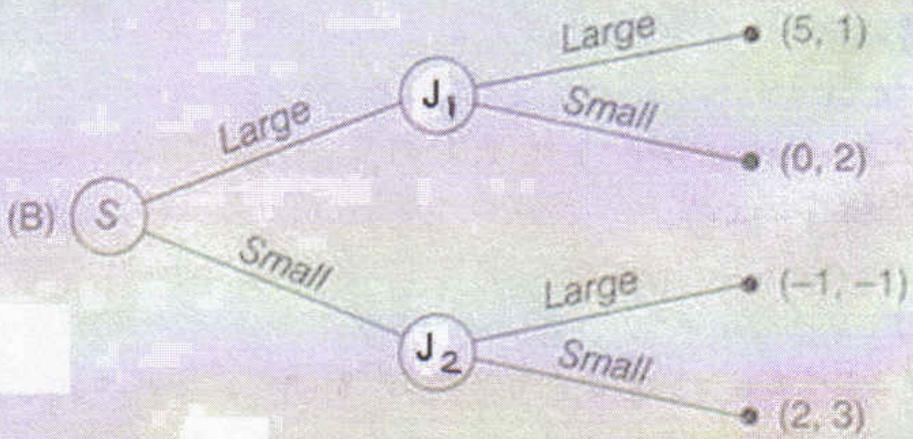
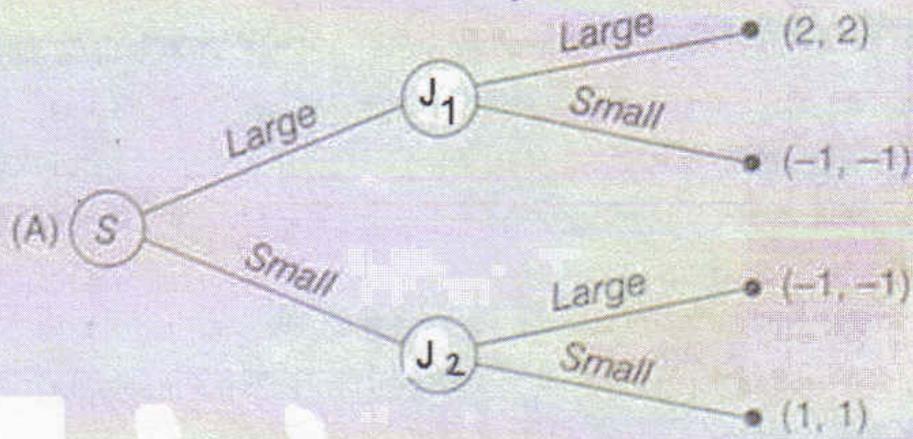
Old definition

In a game of complete information, all players know the rules of the game. Otherwise the game is one of incomplete information.

The old definition is not meaningful, since the game itself is ill defined if it does not specify exactly what the players’ information sets are. Until 1967, game theorists spoke of games of incomplete information to say that they could not be analyzed. Then John Harsanyi pointed out that any game that had incomplete information under the old definition could be remodelled as a game of complete but imperfect information without changing its essentials, simply by adding an initial move in which Nature chooses between different sets of rules. In the transformed game, all players know the new meta-rules,

including the fact that Nature has made an initial move unobserved by them. Harsanyi's suggestion trivialized the definition of incomplete information, and people began using the term to refer to the transformed game instead. Under the old definition, a game of incomplete information was transformed into a game of complete information. Under the new definition, the original game is ill defined, and the transformed version is a game of incomplete information.

Follow-the-Leader III serves to illustrate the Harsanyi transformation. Suppose that Jones does not know the game's payoffs precisely. He does have some idea of the payoffs, and we represent his beliefs by a subjective probability distribution. He places a 70 percent probability on the game being game (A) in Figure 6 (which is the same as Follow-the-Leader I), a 10 percent chance on game (B), and a 20 percent on game (C). In reality the game has a particular set of payoffs, and Smith knows what they are. This is a game of incomplete information (Jones does not know the payoffs), asymmetric information (when Smith moves, Smith knows something Jones does not), and certainty (Nature does not move after the players do.)



Payoffs to: (Smith, Jones)

Figure 6: Follow-the-Leader III: Original

The game cannot be analyzed in the form shown in Figure 6. The natural way to approach such a game is to use the Harsanyi transformation. We can remodel the game to look like Figure 7, in which Nature makes the first move and chooses the payoffs of game (A), (B), or (C), in accordance with Jones's subjective probabilities. Smith observes Nature's move, but Jones does not. Figure 7 depicts the same game as Figure 6, but now we can analyze it. Both Smith and Jones know the rules of the game, and the difference between them is that Smith has observed Nature's move. Whether Nature actually makes the moves with the indicated probabilities or Jones just imagines them is irrelevant, so long as Jones's initial beliefs or fantasies are common knowledge.

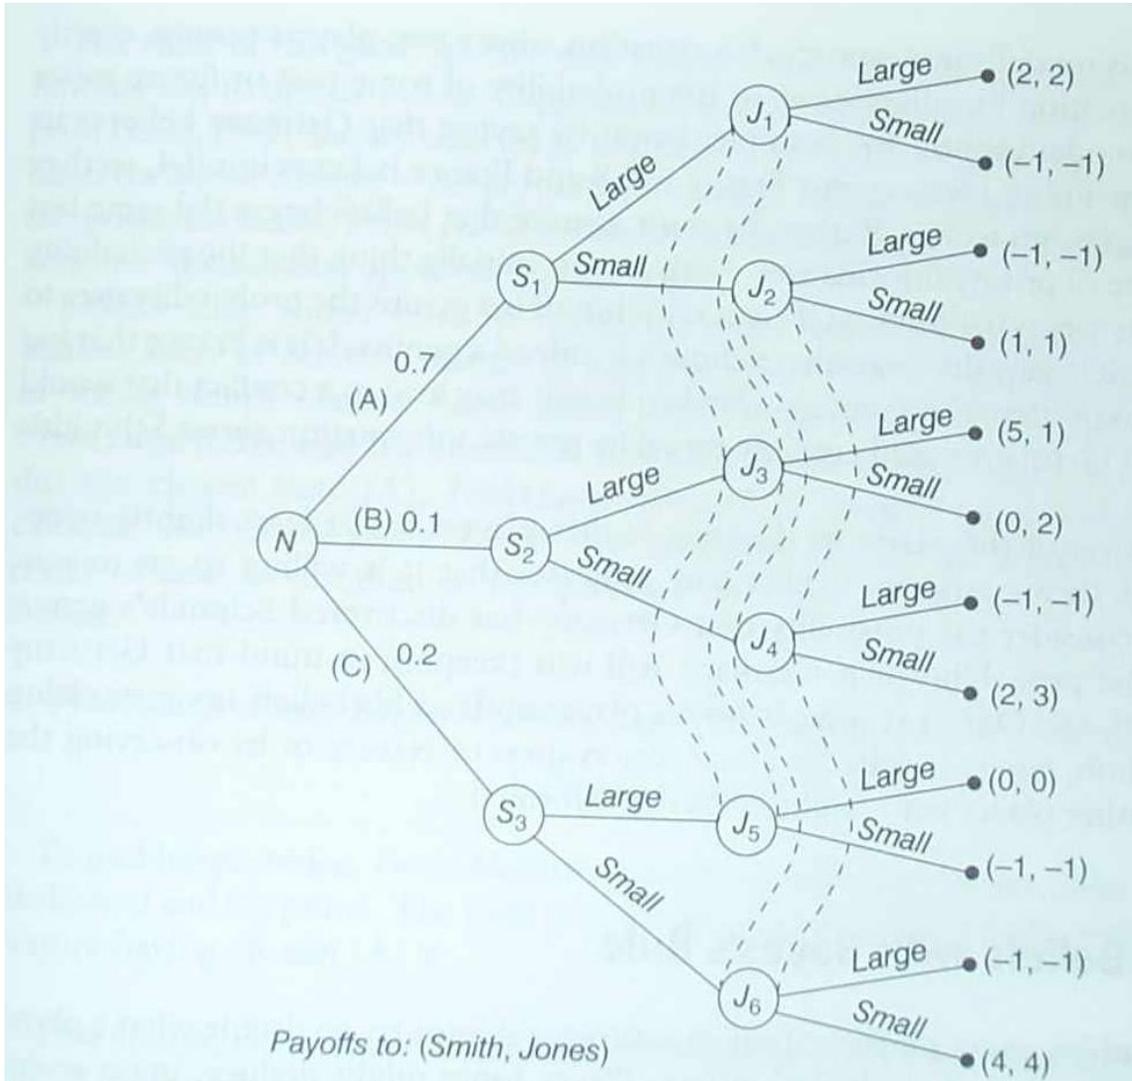


Figure 7: Follow-the-Leader III: After the Harsanyi Transformation

Often what Nature chooses at the start of a game is the strategy set, information partition, and payoff function of one of the players. We say that the player can be any

of several “types,” a term to which we will return in later chapters. When Nature moves, especially if she affects the strategy sets and payoffs of both players, it is often said that Nature has chosen a particular “state of the world.” In Figure 7 Nature chooses the state of the world to be (A), (B), or (C).

*A player’s **type** is the strategy set, information partition, and payoff function which Nature chooses for him at the start of a game of incomplete information.*

*A **state of the world** is a move by Nature.*

As I have already said, it is good modelling practice to assume that the structure of the game is common knowledge, so that though Nature’s choice of Smith’s type may really just represent Jones’s opinions about Smith’s possible type, Smith knows what Jones’s possible opinions are and Jones knows that they are just opinions. The players may have different beliefs, but that is modelled as the effect of their observing different moves by Nature. All players begin the game with the same beliefs about the probabilities of the moves Nature will make—the same priors, to use a term that will shortly be introduced. This modelling assumption is known as the **Harsanyi doctrine**. If the modeller is following it, his model can never reach a situation where two players possess exactly the same information but disagree as to the probability of some past or future move of Nature. A model cannot, for example, begin by saying that Germany believes its probability of winning a war against France is 0.8 and France believes it is 0.4, so they are both willing to go to war. Rather, he must assume that beliefs begin the same but diverge because of private information. Both players initially think that the probability of a German victory is 0.4 but that if General Schmidt is a genius the probability rises to 0.8, and then Germany discovers that Schmidt is indeed a genius. If it is France that has the initiative to declare war, France’s mistaken beliefs may lead to a conflict that would be avoidable if Germany could credibly reveal its private information about Schmidt’s genius.

An implication of the Harsanyi doctrine is that players are at least slightly open-minded about their opinions. If Germany indicates that it is willing to go to war, France must consider the possibility that Germany has discovered Schmidt’s genius and update the probability that Germany will win (keeping in mind that Germany might be bluffing). Our next topic is how a player updates his beliefs upon receiving new information, whether it be by direct observation of Nature or by observing the moves of another player who might be better informed.

Updating Beliefs with Bayes’s Rule

When we classify a game’s information structure we do not try to decide what a player can deduce from the other players’ moves. Player Jones might deduce, upon seeing Smith choose *Large*, that Nature has chosen state (A), but we do not draw Jones’s information set in Figure 7 to take this into account. In drawing the game tree we want to illustrate only the exogenous elements of the game, uncontaminated by the equilibrium concept. But to find the equilibrium we do need to think about how beliefs change over the course of the game.

One part of the rules of the game is the collection of **prior beliefs** (or **priors**) held

by the different players, beliefs that they update in the course of the game. A player holds prior beliefs concerning the types of the other players, and as he sees them take actions he updates his beliefs under the assumption that they are following equilibrium behavior.

The term **bayesian equilibrium** is used to refer to a Nash equilibrium in which players update their beliefs according to Bayes’s Rule. Since Bayes’s Rule is the natural and standard way to handle imperfect information, the adjective, “bayesian,” is really optional. But the two-step procedure of checking a Nash equilibrium has now become a three-step procedure:

- 1 Propose a strategy profile.
- 2 See what beliefs the strategy profile generates when players update their beliefs in response to each others’ moves.
- 3 Check that given those beliefs together with the strategies of the other players each player is choosing a best response for himself.

The rules of the game specify each player’s initial beliefs, and Bayes’s Rule is the rational way to update beliefs. Suppose, for example, that Jones starts with a particular prior belief, $Prob(Nature\ chose\ (A))$. In Follow-the-Leader III, this equals 0.7. He then observes Smith’s move — *Large*, perhaps. Seeing *Large* should make Jones update to the **posterior** belief, $Prob(Nature\ chose\ (A)|Smith\ chose\ Large)$, where the symbol “|” denotes “conditional upon” or “given that.”

Bayes’s Rule shows how to revise the prior belief in the light of new information such as Smith’s move. It uses two pieces of information, the likelihood of seeing Smith choose *Large* given that Nature chose state of the world (A), $Prob(Large|(A))$, and the likelihood of seeing Smith choose *Large* given that Nature did not choose state (A), $Prob(Large|(B)\ or\ (C))$. From these numbers, Jones can calculate $Prob(Smith\ chooses\ Large)$, the **marginal likelihood** of seeing *Large* as the result of one or another of the possible states of the world that Nature might choose.

$$\begin{aligned}
 Prob(Smith\ chooses\ Large) &= Prob(Large|A)Prob(A) + Prob(Large|B)Prob(B) \\
 &\quad + Prob(Large|C)Prob(C).
 \end{aligned}
 \tag{1}$$

To find his posterior, $Prob(Nature\ chose\ (A)|Smith\ chose\ Large)$, Jones uses the likelihood and his priors. The joint probability of both seeing Smith choose *Large* and Nature having chosen (A) is

$$Prob(Large, A) = Prob(A|Large)Prob(Large) = Prob(Large|A)Prob(A).
 \tag{2}$$

Since what Jones is trying to calculate is $Prob(A|Large)$, rewrite the last part of (2) as follows:

$$Prob(A|Large) = \frac{Prob(Large|A)Prob(A)}{Prob(Large)}.
 \tag{3}$$

Jones needs to calculate his new belief — his posterior — using $Prob(Large)$, which he calculates from his original knowledge using (1). Substituting the expression for $Prob(Large)$

from (1) into equation (3) gives the final result, a version of Bayes’s Rule.

$$Prob(A|Large) = \frac{Prob(Large|A)Prob(A)}{Prob(Large|A)Prob(A) + Prob(Large|B)Prob(B) + Prob(Large|C)Prob(C)}. \quad (4)$$

More generally, for Nature’s move x and the observed data,

$$Prob(x|data) = \frac{Prob(data|x)Prob(x)}{Prob(data)} \quad (5)$$

Equation (6) is a verbal form of Bayes’s Rule, which is useful for remembering the terminology,² summarized in Table 5.

$$(Posterior\ for\ Nature's\ Move) = \frac{(Likelihood\ of\ Player's\ Move) \cdot (Prior\ for\ Nature's\ Move)}{(Marginal\ likelihood\ of\ Player's\ Move)}. \quad (6)$$

Bayes’s Rule is not purely mechanical. It is the only way to rationally update beliefs. The derivation is worth understanding, because Bayes’s Rule is hard to memorize but easy to rederive.

Table 5: Bayesian Terminology

Name	Meaning
Likelihood	$Prob(data event)$
Marginal likelihood	$Prob(data)$
Conditional Likelihood	$Prob(data\ X data\ Y, event)$
Prior	$Prob(event)$
Posterior	$Prob(event data)$

Updating Beliefs in Follow-the-Leader III

Let us now return to the numbers in Follow-the-Leader III to use the belief-updating rule that was just derived. Jones has a prior belief that the probability of event “Nature picks state (A)” is 0.7 and he needs to update that belief on seeing the data “Smith picks *Large*”. His prior is $Prob(A) = 0.7$, and we wish to calculate $Prob(A|Large)$.

To use Bayes’s Rule from equation (4), we need the values of $Prob(Large|A)$, $Prob(Large|B)$, and $Prob(Large|C)$. These values depend on what Smith does in equilibrium, so Jones’s

²The name “marginal likelihood” may seem strange to economists, since it is an unconditional likelihood and when economists use “marginal” they mean “an increment conditional on starting from a particular level”. The statisticians defined marginal likelihood this way because they start with $Prob(a, b)$, and then derive $Prob(b)$. That is like going to the margin of a graph in (a, b) -space, the b -axis, and asking how probable the value of b is integrating over all possible a ’s.

beliefs cannot be calculated independently of the equilibrium. This is the reason for the three-step procedure suggested above, for what the modeller must do is propose an equilibrium and then use it to calculate the beliefs. Afterwards, he must check that the equilibrium strategies are indeed the best responses given the beliefs they generate.

A candidate for equilibrium in Follow-the-Leader III is for Smith to choose *Large* if the state is (A) or (B) and *Small* if it is (C), and for Jones to respond to *Large* with *Large* and to *Small* with *Small*. This can be abbreviated as $(L|A, L|B, S|C; L|L, S|S)$. Let us test that this is an equilibrium, starting with the calculation of $Prob(A|Large)$.

If Jones observes *Large*, he can rule out state (C), but he does not know whether the state is (A) or (B). Bayes's Rule tells him that the posterior probability of state (A) is

$$\begin{aligned} Prob(A|Large) &= \frac{(1)(0.7)}{(1)(0.7)+(1)(0.1)+(0)(0.2)} \\ &= 0.875. \end{aligned} \tag{7}$$

The posterior probability of state (B) must then be $1 - 0.875 = 0.125$, which could also be calculated from Bayes's Rule, as follows:

$$\begin{aligned} (B|Large) &= \frac{(1)(0.1)}{(1)(0.7)+(1)(0.1)+(0)(0.2)} \\ &= 0.125. \end{aligned} \tag{8}$$

Figure 8 shows a graphic intuition for Bayes's Rule. The first line shows the total probability, 1, which is the sum of the prior probabilities of states (A), (B), and (C). The second line shows the probabilities, summing to 0.8, which remain after *Large* is observed and state (C) is ruled out. The third line shows that state (A) represents an amount 0.7 of that probability, a fraction of 0.875. The fourth line shows that state (B) represents an amount 0.1 of that probability, a fraction of 0.125.

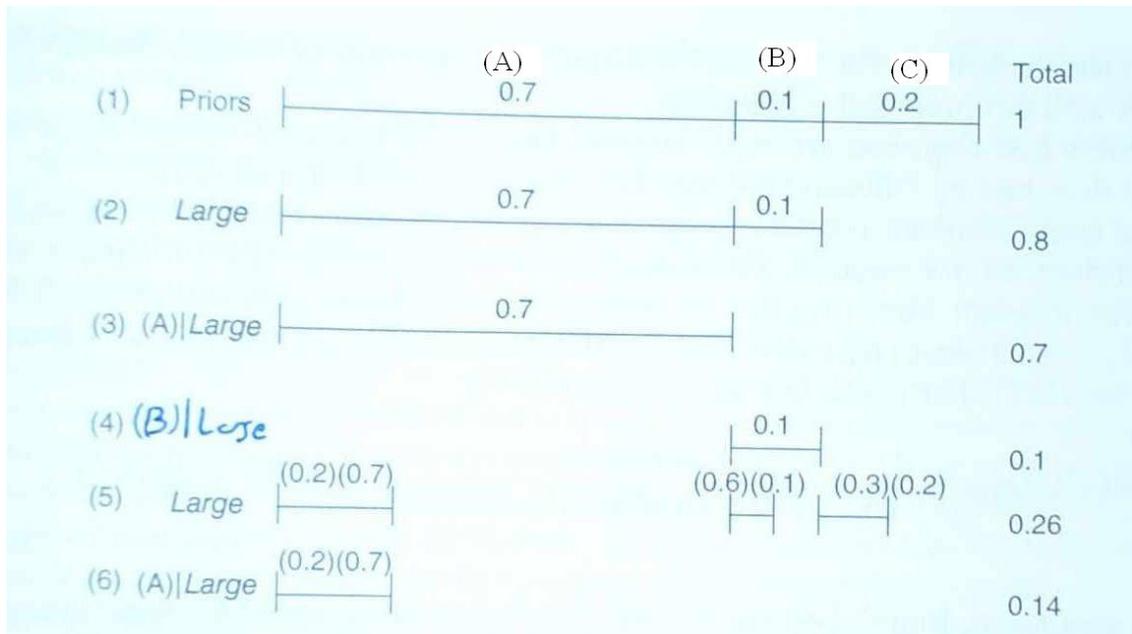


Figure 8: Bayes's Rule

Jones must use Smith's strategy in the proposed equilibrium to find numbers for $Prob(Large|A)$, $Prob(Large|B)$, and $Prob(Large|C)$. As always in Nash equilibrium, the modeller assumes that the players know which equilibrium strategies are being played out, even though they do not know which particular actions are being chosen.

Given that Jones believes that the state is (A) with probability 0.875 and state (B) with probability 0.125, his best response is *Large*, even though he knows that if the state were actually (B) the better response would be *Small*. Given that he observes *Large*, Jones's expected payoff from *Small* is -0.625 ($= 0.875[-1] + 0.125[2]$), but from *Large* it is 1.875 ($= 0.875[2] + 0.125[1]$). The strategy profile $(L|A, L|B, S|C; L|L, S|S)$ is a bayesian equilibrium.

A similar calculation can be done for $Prob(A|Small)$. Using Bayes's Rule, equation (4) becomes

$$Prob(A|Small) = \frac{(0)(0.7)}{(0)(0.7) + (0)(0.1) + (1)(0.2)} = 0. \quad (9)$$

Given that he believes the state is (C), Jones's best response to *Small* is *Small*, which agrees with our proposed equilibrium.

Smith's best responses are much simpler. Given that Jones will imitate his action, Smith does best by following his equilibrium strategy of $(L|A, L|B, S|C)$.

The calculations are relatively simple because Smith uses a nonrandom strategy in equilibrium, so, for instance, $Prob(Small|A) = 0$ in equation (9). Consider what happens if Smith uses a random strategy of picking *Large* with probability 0.2 in state (A), 0.6 in state (B), and 0.3 in state (C) (we will analyze such "mixed" strategies in Chapter 3). The equivalent of equation (7) is

$$Prob(A|Large) = \frac{(0.2)(0.7)}{(0.2)(0.7) + (0.6)(0.1) + (0.3)(0.2)} = 0.54 \text{ (rounded)}. \quad (10)$$

If he sees *Large*, Jones's best guess is still that Nature chose state (A), even though in state (A) Smith has the smallest probability of choosing *Large*, but Jones's subjective posterior probability, $Pr(A|Large)$, has fallen to 0.54 from his prior of $Pr(A) = 0.7$.

The last two lines of Figure 8 illustrate this case. The second-to-last line shows the total probability of *Large*, which is formed from the probabilities in all three states and sums to 0.26 ($= 0.14 + 0.06 + 0.06$). The last line shows the component of that probability arising from state (A), which is the amount 0.14 and fraction 0.54 (rounded).

Regression to the Mean, the Two-Armed Bandit and Cascades

Bayesian learning is important not just in modelling bayesian games, but in explaining behavior that is non-strategic, in the sense that although players may learn from the moves of other players, their payoffs are not directly affected by those moves. I will discuss three

phenomenon that give us useful explanations for behavior: regression to the mean, the bandit problem, and cascades.

Regression to the mean is an old statistical idea that has a bayesian interpretation. Suppose that each student's performance on a test results partly from his ability and partly from random error because of his mood the day of the test. The teacher does not know the individual student's ability, but does know that the average student will score 70 out of 100. If a student scores 40, what should the teacher's estimate of his ability be?

It should not be 40. A score of 30 points below the average score could be the result of two things: (1) the student's ability is below average, or (2) the student was in a bad mood the day of the test. Only if mood is completely unimportant should the teacher use 40 as his estimate. More likely, both ability and luck matter to some extent, so the teacher's best guess is that the student has an ability below average but was also unlucky. The best estimate lies somewhere between 40 and 70, reflecting the influence of both ability and luck. Of the students who score 40 on the test, more than half can be expected to score above 40 on the next test. Since the scores of these poorly performing students tend to float up towards the mean of 70, this phenomenon is called "regression to the mean." Similarly, students who score 90 on the first test will tend to score less well on the second test.

This is "regression to the mean" ("towards" would be more precise) not "regression beyond the mean." A low score does indicate low ability, on average, so the predicted score on the second test is still below average. Regression to the mean merely recognizes that both luck and ability are at work.

In bayesian terms, the teacher in this example has a prior mean of 70, and is trying to form a posterior estimate using the prior and one piece of data, the score on the first test. For typical distributions, the posterior mean will lie between the prior mean and the data point, so the posterior mean will be between 40 and 70.

In a business context, regression to the mean can be used to explain business conservatism, as I do in Rasmusen (1992b). It is sometimes claimed that businesses pass up profitable investments because they have an excessive fear of risk. Let us suppose that the business is risk neutral, because the risk associated with the project and the uncertainty over its value are nonsystematic — that is, they are risks that a widely held corporation can distribute in such a way that each shareholder's risk is trivial. Suppose that the firm will not spend \$100,000 on an investment with a present value of \$105,000. This is easily explained if the \$105,000 is an estimate and the \$100,000 is cash. If the average value of a new project of this kind is less than \$100,000 — as is likely to be the case since profitable projects are not easy to find — the best estimate of the value will lie between the measured value of \$105,000 and that average value, unless the staffer who came up with the \$105,000 figure has already adjusted his estimate. Regressing the \$105,000 to the mean may regress it past \$100,000. Put a bit differently, if the prior mean is, let us say, \$80,000, and the data point is \$105,000, the posterior may well be less than \$100,000. Regression to the mean is an alternative to strategic behavior in explaining certain odd phenomena. In analyzing test scores, one might try to explain the rise in the scores of poor students by changes in their effort level in an attempt to achieve a target grade in the course with minimum work. In analyzing business decisions, one might try to explain why apparently profitable projects

are rejected because of managers' dislike for innovations that would require them to work harder.

single decisionmaker "Two-Armed Bandit" model of Rothschild (1974). In each of a sequence of periods, a person chooses to play slot machine A or slot machine B. Slot machine A pays out \$1 with known probability 0.5 in exchange for the person putting in \$0.25 and pulling its arm. Slot machine B pays out \$1 with an unknown probability which has a prior probability density centered on 0.5. The optimal strategy is to begin by playing machine B, since not only does it have the same expected payout per period, but also playing it improves the player's information, whereas playing machine A leaves his information unchanged. The player will switch to machine A if machine B pays out \$0 often enough relative to the number of times it pays out \$1, where "often enough" depends on the particular prior beliefs he has. If the first 1,000 plays all result in a payout of \$1, he will keep playing machine B, but if the next 9,000 plays all result in a payout of \$0, he should become very sure that machine B's payout rate is less than 0.5 and he should switch to machine A. But he will never switch back. Once he is playing machine A, he is learning nothing new as a result of his wins and losses, and even if he gets a payout of \$0 ten thousand times in a row, that gives him no reason to change machines. As a result, it can happen that even if machine B actually is better, a player following the ex ante optimal strategy can end up playing machine A an infinite number of times.

Another model with a similar flavor is the **cascade model**. Consider a simplified version of the first example of a cascade in Bikchandani, Hirshleifer & Welch (1992) (who with Bannjerjee [1992] originated the idea; see also Hirshleifer [1995]). A sequence of people must decide whether to Adopt at cost 0.5 or Reject a project worth either 0 or 1 with equal prior probabilities, having observed the decisions of people ahead of them in the sequence plus an independent private signal that takes the value High with probability $p > 0.5$ if the project's value is 1 and with probability $(1 - p)$ if it is 0, and otherwise takes the value Low.

The first person will simply follow his signal, choosing Adopt if the signal is High and Reject if it is Low. The second person uses the information of the first person's decision plus his own signal. One Nash equilibrium is for the second person to always imitate the first person. It is easy to see that he should imitate the first person if the first person chose Adopt and the second signal is High. What if the first person chose Adopt and the second signal is Low? Then the second person can deduce that the first signal was High, and choosing on the basis of a prior of 0.5 and two contradictory signals of equal accuracy, he is indifferent— and so will not deviate from an equilibrium in which his assigned strategy is to imitate the first person when indifferent. The third person, having seen the first two choose Adopt, will also deduce that the first person's signal was High. He will ignore the second person's decision, knowing that in equilibrium that person just imitates, but he, too will imitate. Thus, even if the sequence of signals is (High, Low, Low, Low, Low...), everyone will choose Adopt. A "cascade" has begun, in which players later in the sequence ignore their own information and rely completely on previous players. Thus, we have a way to explain fads and fashions as Bayesian updating under incomplete information, without any strategic behavior.

2.5: An Example: The Png Settlement Game

The Png (1983) model of out-of-court settlement is an example of a game with a fairly complicated extensive form.³ The plaintiff alleges that the defendant was negligent in providing safety equipment at a chemical plant, a charge which is true with probability q . The plaintiff files suit, but the case is not decided immediately. In the meantime, the defendant and the plaintiff can settle out of court.

What are the moves in this game? It is really made up of two games: the one in which the defendant is liable for damages, and the one in which he is blameless. We therefore start the game tree with a move by Nature, who makes the defendant either liable or blameless. At the next node, the plaintiff takes an action: *Sue* or *Grumble*. If he decides on *Grumble* the game ends with zero payoffs for both players. If he decides to *Sue*, we go to the next node. The defendant then decides whether to *Resist* or *Offer* to settle. If the defendant chooses *Offer*, then the plaintiff can *Settle* or *Refuse*; if the defendant chooses to *Resist*, the plaintiff can *Try* the case or *Drop* it. The following description adds payoffs to this model.

The Png Settlement Game

Players

The plaintiff and the defendant.

The Order of Play

0 Nature chooses the defendant to be Liable for injury to the plaintiff with probability $q = 0.13$ and Blameless otherwise. The defendant observes this but the plaintiff does not.

1 The plaintiff decides to Sue or just to Grumble.

2 The defendant Offers a settlement amount of $S = 0.15$ to the plaintiff, or Resist, setting $S = 0$.

3a If the defendant offered $S = 0.15$, the plaintiff agrees to Settle or he Refuses and goes to trial.

3b If the defendant offered $S = 0$, the plaintiff Drops the case, for legal costs of $P = 0$ and $D = 0$ for himself and the defendant, or chooses to Try it, creating legal costs of $P = 0.1$ and $D = 0.2$

4 If the case goes to trial, the plaintiff wins damages of $W = 1$ if the defendant is Liable and $W = 0$ if the defendant is Blameless. If the case is dropped, $W = 0$.

Payoffs

The plaintiff's payoff is $S + W - P$. The defendant's payoff is $-S - W - D$.

We can also depict this on a game tree, as in Figure 9.

³“Png,” by the way, is pronounced the same way it is spelt.

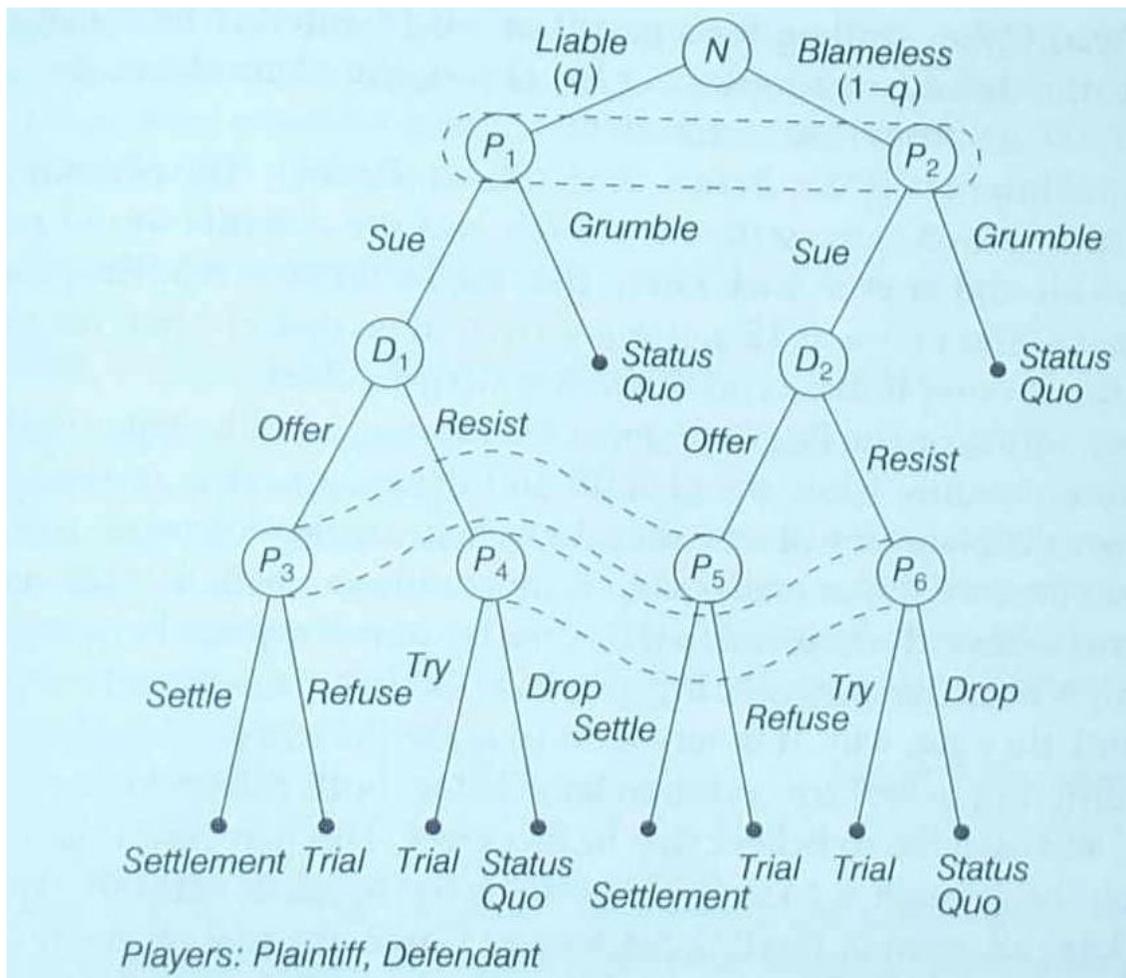


Figure 9: The Game Tree for The Png Settlement Game

This model assumes that the settlement amount, $S = 0.15$, and the amounts spent on legal fees are exogenous. Except in the infinitely long games without end nodes that will appear in Chapter 5, an extensive form should incorporate all costs and benefits into the payoffs at the end nodes, even if costs are incurred along the way. If the court required a \$100 filing fee (which it does not in this game, although a fee will be required in the similar game of Nuisance Suits in Section 4.3), it would be subtracted from the plaintiff's payoffs at every end node except those resulting from his choice of *Grumble*. Such consolidation makes it easier to analyze the game and would not affect the equilibrium strategies unless payments along the way revealed information, in which case what matters is the information, not the fact that payoffs change.

We assume that if the case reaches the court, justice is done. In addition to his legal fees D , the defendant pays damages $W = 1$ only if he is liable. We also assume that the players are risk neutral, so they only care about the expected dollars they will receive, not the variance. Without this assumption we would have to translate the dollar payoffs into utility, but the game tree would be unaffected.

This is a game of certain, asymmetric, imperfect, and incomplete information. We

have assumed that the defendant knows whether he is liable, but we could modify the game by assuming that he has no better idea than the plaintiff of whether the evidence is sufficient to prove him so. The game would become one of symmetric information and we could reasonably simplify the extensive form by eliminating the initial move by Nature and setting the payoffs equal to the expected values. We cannot perform this simplification in the original game, because the fact that the defendant, and only the defendant, knows whether he is liable strongly affects the behavior of both players.

Let us now find the equilibrium. Using dominance we can rule out one of the plaintiff's strategies immediately — *Grumble* — which is dominated by $(Sue, Settle, Drop)$.

Whether a strategy profile is a Nash equilibrium depends on the parameters of the model— S, W, P, D and q , which are the settlement amount, the damages, the court costs for the plaintiff and defendant, and the probability the defendant is liable. Depending on the parameter values, three outcomes are possible: settlement (if the settlement amount is low), trial (if expected damages are high and the plaintiff's court costs are low), and the plaintiff dropping the action (if expected damages minus court costs are negative). Here, I have inserted the parameter values $S = 0.15, D = 0.2, W = 1, q = 0.13$, and $P = 0.1$. Two Nash equilibria exist for this set of parameter values, both weak.

One equilibrium is the strategy profile $\{(Sue, Settle, Try), (Offer, Offer)\}$. The plaintiff sues, the defendant offers to settle (whether liable or not), and the plaintiff agrees to settle. Both players know that if the defendant did not offer to settle, the plaintiff would go to court and try the case. Such **out-of-equilibrium** behavior is specified by the equilibrium, because the threat of trial is what induces the defendant to offer to settle, even though trials never occur in equilibrium. This is a Nash equilibrium because given that the plaintiff chooses $(Sue, Settle, Try)$, the defendant can do no better than $(Offer, Offer)$, settling for a payoff of -0.15 whether he is liable or not; and, given that the defendant chooses $(Offer, Offer)$, the plaintiff can do no better than the payoff of 0.15 from $(Sue, Settle, Try)$.

The other equilibrium is $\{(Sue, Refuse, Try), (Resist, Resist)\}$. The plaintiff sues, the defendant resists and makes no settlement offer, the plaintiff would refuse any offer that was made, and goes to trial. Since he foresees the plaintiff will refuse a settlement offer of $S = 0.15$, the defendant is willing to resist, because his action makes no difference.

One final observation on the Png Settlement Game: the game illustrates the Harsanyi doctrine in action, because while the plaintiff and defendant differ in their beliefs as to the probability the plaintiff will win, they do so because the defendant has different information, not because the modeller assigns them different beliefs at the start of the game. This seems awkward compared to the everyday way of approaching this problem in which we simply note that potential litigants have different beliefs, and will go to trial if they both think they can win. It is very hard to make the story consistent, however, because if the differing beliefs are common knowledge, both players know that one of them is wrong, and each has to believe that he is correct. This may be fine as a “reduced form,” in which the attempt is to simply describe what happens without explaining it in any depth. After all, even in the Png Settlement Game, if a trial occurs it is because the players differ in their beliefs, so one could simply chop off the first part of the game tree. But that is also the problem with violating the Harsanyi doctrine: one cannot analyze how the players react to each other's

moves if the modeller simply assigns them inflexible beliefs. In the Png Settlement Game, a settlement is rejected and a trial can occur under certain parameters because the plaintiff weighs the probability that the defendant knows he will win versus the probability that he is bluffing, and sometimes decides to risk a trial. Without the Harsanyi doctrine it is very hard to evaluate such an explanation for trials.

NOTES

N2.1 The strategic and extensive forms of a game

- The term “outcome matrix” is used in Shubik (1982, p. 70), but never formally defined there.
- The term “node” is sometimes defined to include only points at which a player or Nature makes a decision, which excludes the end points.

N2.2 Information Sets

- If you wish to depict a situation in which a player does not know whether the game has reached node A_1 or A_2 and he has different action sets at the two nodes, restructure the game. If you wish to say that he has action set (X,Y,Z) at A_1 and (X,Y) at A_2 , first add action Z to the information set at A_2 . Then specify that at A_2 , action Z simply leads to a new node, A_3 , at which the choice is between X and Y .
- The term “common knowledge” comes from Lewis (1969). Discussions include Brandenburger (1992) and Geanakoplos (1992). For rigorous but nonintuitive definitions of common knowledge, see Aumann (1976) (for two players) and Milgrom (1981a) (for n players).

N2.3 Perfect, Certain, Symmetric, and Complete Information

- Tirole (1988, p. 431) (and more precisely Fudenberg & Tirole [1991a, p. 82]) have defined games of *almost perfect* information. They use this term to refer to repeated simultaneous-move games (of the kind studied here in Chapter 5) in which at each repetition all players know the results of all the moves, including those of Nature, in previous repetitions. It is a pity they use such a general-sounding term to describe so narrow a class of games; it could be usefully extended to cover all games which have perfect information except for simultaneous moves.
- **Poker Classifications.** (1) Perfect, certain. (2) Incomplete, symmetric, certain. (3) Incomplete, asymmetric, certain. (4) Complete, asymmetric, certain. (5) Perfect, uncertain. (6) Incomplete, asymmetric, certain.
- For explanation of von Neumann-Morgenstern utility, see Varian (1992, chapter 11) or Kreps (1990a, Chapter 3). For other approaches to utility, see Starmer (2000). Expected utility and Bayesian updating are the two foundations of standard game theory, partly because they seem realistic but more because they are so simple to use. Sometimes they do not explain people’s behavior well, and there exist extensive literatures (a) pointing out anomalies, and (b) suggesting alternatives. So far no alternatives have proven to be big enough improvements to justify replacing the standard techniques, given the tradeoff between descriptive realism and added complexity in modelling. The standard response is to admit and ignore the anomalies in theoretical work, and to not press any theoretical models too hard in situations where the anomalies are likely to make a significant difference. On anomalies, see Kahneman, Slovic & Tversky (1982) (an edited collection); Thaler (1992) (essays from his *Journal of Economic Perspectives* column); and Dawes (1988) (a good mix of psychology and business).

- Mixed strategies (to be described in Section 3.1) are allowed in a game of perfect information because they are an aspect of the game’s equilibrium, not of its exogenous structure.
- Although the word “perfect,” appears in both “perfect information” (Section 2.3) and “perfect equilibrium” (Section 4.1), the concepts are unrelated.
- An unobserved move by Nature in a game of symmetric information can be represented in any of three ways: (1) as the last move in the game; (2) as the first move in the game; or (3) by replacing the payoffs with the expected payoffs and not using any explicit moves by Nature.

N2.4 The Harsanyi Transformation and Bayesian Games

- Mertens & Zamir (1985) probes the mathematical foundations of the Harsanyi transformation. The transformation requires the extensive form to be common knowledge, which raises subtle questions of recursion.
- A player always has some idea of what the payoffs are, so we can always assign him a subjective probability for each possible payoff. What would happen if he had no idea? Such a question is meaningless, because people always have some notion, and when they say they do not, they generally mean that their prior probabilities are low but positive for a great many possibilities. You, for instance, probably have as little idea as I do of how many cups of coffee I have consumed in my lifetime, but you would admit it to be a nonnegative number less than 3,000,000, and you could make a much more precise guess than that. On the topic of subjective probability, the classic reference is Savage (1954).
- The term “marginal likelihood” is confusing for economists, since it refers to an unconditional likelihood. Statisticians came up with it because they start with $Prob(a, b)$ and then move to $Prob(a)$. That is like going to the margin of a graph— the a -axis— and asking how probable each value of a is.
- If two players have common priors and their information partitions are finite, but they each have private information, iterated communication between them will lead to the adoption of a common posterior. This posterior is not always the posterior they would reach if they directly pooled their information, but it is almost always that posterior (Geanakoplos & Polemarchakis [1982]).

Problems

2.1. The Monty Hall Problem (easy)

You are a contestant on the TV show, “Let’s Make a Deal.” You face three curtains, labelled A, B and C. Behind two of them are toasters, and behind the third is a Mazda Miata car. You choose A, and the TV showmaster says, pulling curtain B aside to reveal a toaster, “You’re lucky you didn’t choose B, but before I show you what is behind the other two curtains, would you like to change from curtain A to curtain C?” Should you switch? What is the exact probability that curtain C hides the Miata?

2.2. Elmer’s Apple Pie (hard)

Mrs Jones has made an apple pie for her son, Elmer, and she is trying to figure out whether the pie tasted divine, or merely good. Her pies turn out divinely a third of the time. Elmer might be ravenous, or merely hungry, and he will eat either 2, 3, or 4 pieces of pie. Mrs Jones knows he is ravenous half the time (but not which half). If the pie is divine, then, if Elmer is hungry, the probabilities of the three consumptions are (0, 0.6, 0.4), but if he is ravenous the probabilities are (0, 0, 1). If the pie is just good, then the probabilities are (0.2, 0.4, 0.4) if he is hungry and (0.1, 0.3, 0.6) if he is ravenous.

Elmer is a sensitive, but useless, boy. He will always say that the pie is divine and his appetite weak, regardless of his true inner feelings.

- (a) What is the probability that he will eat four pieces of pie?
- (b) If Mrs Jones sees Elmer eat four pieces of pie, what is the probability that he is ravenous and the pie is merely good?
- (c) If Mrs Jones sees Elmer eat four pieces of pie, what is the probability that the pie is divine?

2.3. Cancer Tests (easy) (adapted from McMillan [1992, p. 211])

Imagine that you are being tested for cancer, using a test that is 98 percent accurate. If you indeed have cancer, the test shows positive (indicating cancer) 98 percent of the time. If you do not have cancer, it shows negative 98 percent of the time. You have heard that 1 in 20 people in the population actually have cancer. Now your doctor tells you that you tested positive, but you shouldn’t worry because his last 19 patients all died. How worried should you be? What is the probability you have cancer?

2.4. The Battleship Problem (hard) (adapted from Barry Nalebuff, “Puzzles,” *Journal of Economic Perspectives*, 2:181-82 [Fall 1988])

The Pentagon has the choice of building one battleship or two cruisers. One battleship costs the same as two cruisers, but a cruiser is sufficient to carry out the navy’s mission — if the cruiser survives to get close enough to the target. The battleship has a probability of p of carrying out its mission, whereas a cruiser only has probability $p/2$. Whatever the outcome, the war ends and any surviving ships are scrapped. Which option is superior?

2.5. Joint Ventures (medium)

Software Inc. and Hardware Inc. have formed a joint venture. Each can exert either high or

low effort, which is equivalent to costs of 20 and 0. Hardware moves first, but Software cannot observe his effort. Revenues are split equally at the end, and the two firms are risk neutral. If both firms exert low effort, total revenues are 100. If the parts are defective, the total revenue is 100; otherwise, if both exert high effort, revenue is 200, but if only one player does, revenue is 100 with probability 0.9 and 200 with probability 0.1. Before they start, both players believe that the probability of defective parts is 0.7. Hardware discovers the truth about the parts by observation before he chooses effort, but Software does not.

- (a) Draw the extensive form and put dotted lines around the information sets of Software at any nodes at which he moves.
- (b) What is the Nash equilibrium?
- (c) What is Software's belief, in equilibrium, as to the probability that Hardware chooses low effort?
- (d) If Software sees that revenue is 100, what probability does he assign to defective parts if he himself exerted high effort and he believes that Hardware chose low effort?

2.6. California Drought (hard)

California is in a drought and the reservoirs are running low. The probability of rainfall in 1991 is $1/2$, but with probability 1 there will be heavy rainfall in 1992 and any saved water will be useless. The state uses rationing rather than the price system, and it must decide how much water to consume in 1990 and how much to save till 1991. Each Californian has a utility function of $U = \log(w_{90}) + \log(w_{91})$. Show that if the discount rate is zero the state should allocate twice as much water to 1990 as to 1991.

2.7. Smith's Energy Level (easy)

The boss is trying to decide whether Smith's energy level is high or low. He can only look in on Smith once during the day. He knows if Smith's energy is low, he will be yawning with a 50 percent probability, but if it is high, he will be yawning with a 10 percent probability. Before he looks in on him, the boss thinks that there is an 80 percent probability that Smith's energy is high, but then he sees him yawning. What probability of high energy should the boss now assess?

2.8. Two Games (medium)

Suppose that Column gets to choose which of the two payoff structures in Tables 6 and 7 applies to the simultaneous-move game he plays with Row. Row does not know which of these Column has chosen.

Table 6: Payoffs (A), The Prisoner's Dilemma

		Column	
		<i>Deny</i>	<i>Confess</i>
Row:	<i>Deny</i>	-1,-1	-10, 0
	<i>Confess</i>	0,-10	-8,-8

Payoffs to: (Row, Column).

Table 7: Payoffs (B), A Confession Game

		Column	
		<i>Deny</i>	<i>Confess</i>
Row:	<i>Deny</i>	-4,-4	-12, -200
	<i>Confess</i>	-200,-12	-10,-10

Payoffs to: (Row,Column).

- (a) What is one example of a strategy for each player?
- (b) Find a Nash equilibrium. Is it unique? Explain your reasoning.
- (c) Is there a dominant strategy for Column? Explain why or why not.
- (d) Is there a dominant strategy for Row? Explain why or why not.
- (e) Does Row's choice of strategy depend on whether Column is rational or not? Explain why or why not.

Bayes Rule at the Bar: A Classroom Game for Chapter 2

I have wandered into a dangerous bar in Jersey City. There are six people in there. Based on past experience, I estimate that three are cold-blooded killers and three are cowardly bullies. I know that $2/3$ of killers are aggressive and $1/3$ reasonable; but $1/3$ of cowards are aggressive and $2/3$ are reasonable. Unfortunately, I spill my drink on a mean-looking rascal, who asks me if I want to die.

In crafting my response in the two seconds I have to think, I would like to know the probability I have offended a killer. Give me your estimate.

The story continues. A friend of the wet rascal comes in from the street outside the bar and learns what happened. He, too, turns aggressive. I know that the friend is just like the first rascal— a killer if the first one was a killer, a coward otherwise. Does this extra trouble change your estimate that the two of them are killers?

This game is a descendant of the game in Charles Holt & Lisa R. Anderson. “Classroom Games: Understanding Bayes Rule,” *Journal of Economic Perspectives*, 10: 179-187 (Spring 1996), but I use a different heuristic for the rule and a barroom story instead of urns. Psychologists have found that people can solve logical puzzles better if the puzzles are associated with a story involving social interactions. See Chapter 7 of Robin Dunbar’s *The Trouble with Science*, which explains experiments and ideas from Cosmides & Toobey (1993).