3 September 2009 Eric Rasmusen, Erasmuse@indiana.edu. http://www.rasmusen.org/.

Table 1: Ranked Coordination

		Jones		
		Large		*Small*
	Large	**2,2**	$\leftarrow$	$-1,-1$
Smith		$\uparrow$		$\downarrow$
	Small	$-1,-1$	$\rightarrow$	**1,1**

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

The **strategic form** *(or* **normal form***) consists of*
1 All possible strategy profiles $s^1, s^2, \ldots, s^p$.
2 Payoff functions mapping s^i *onto the payoff n-vector* π^i, $(i = 1, 2, \ldots, p)$.

Follow-the-Leader I

Smith has a strategy set of two strategies: Small or Large.

Jones has a strategy set of four different strategies:

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

Combining one strategy for each player, we get a strategy profile.

That results in an action for each player, and a payoff.

The normal form shows the strategies and payoffs, omitting the actions.

Table 2: Normal Form for 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	$S_1: Large$	$\boxed{\mathbf{2}}$, **2** (E_1)	$\boxed{\mathbf{2}}$, **2** (E_2)	$-1, -1$	$-1, -1$
	$S_2: Small$	$-1, -1$	$1, 1$	$-1, -1$	$\boxed{\mathbf{1}}$, **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 *Large*
E_2	$\{$*Large*, (*L|L, S|S*)$\}$	Both pick *Large*
E_3	$\{$*Small*,(*S|L, S|S*)$\}$	Both pick *Small*

The Extensive Form

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.*

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.

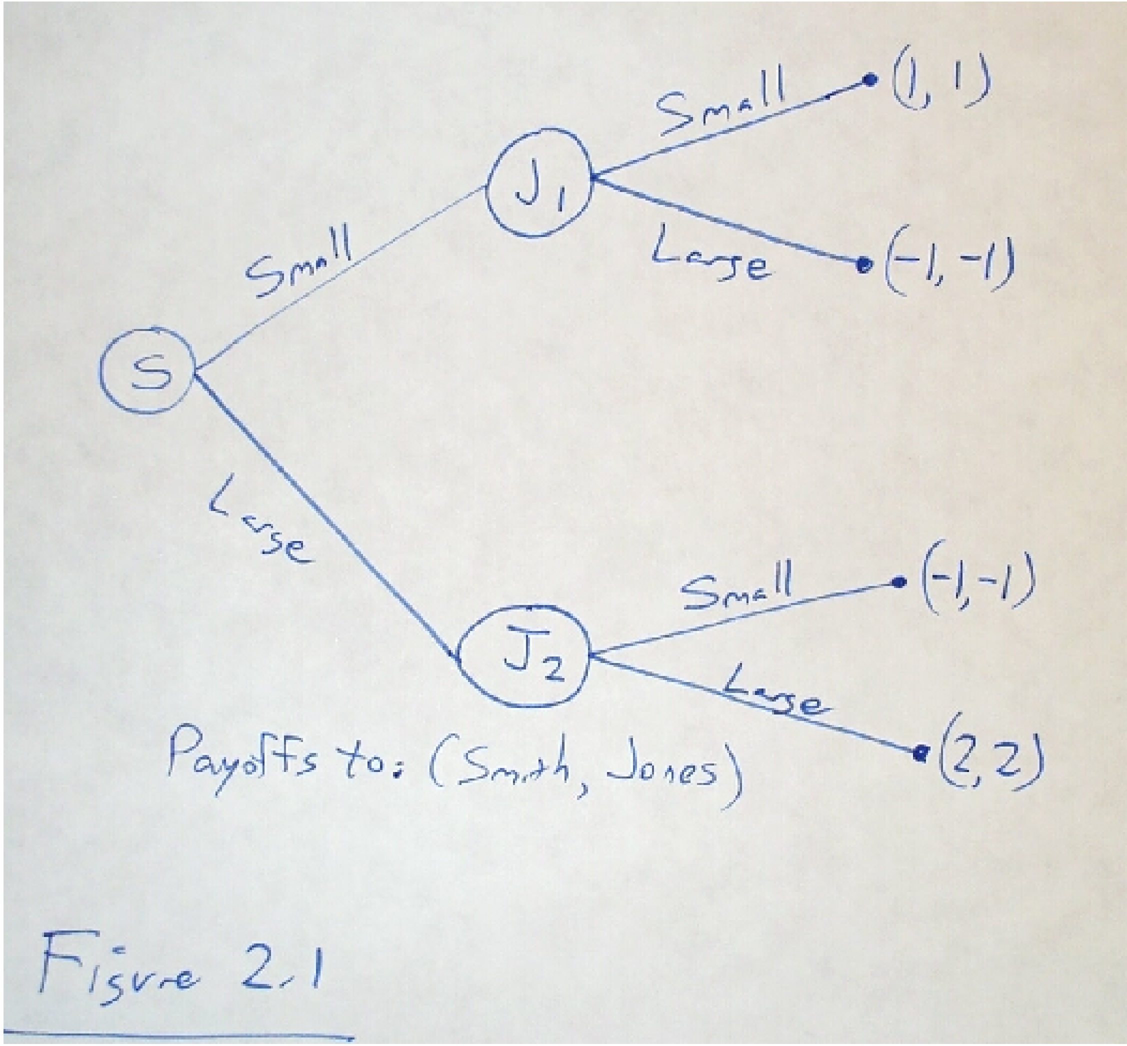

Figure 1: Follow-the-Leader I in Extensive Form

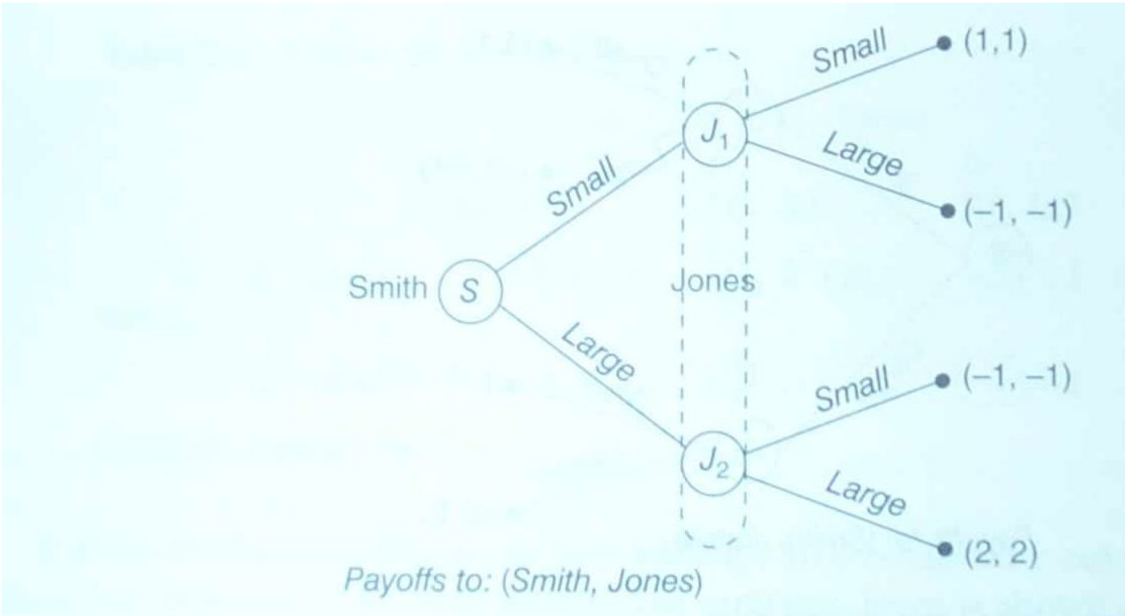

Figure 2: Ranked Coordination in Extensive Form

The Time Line

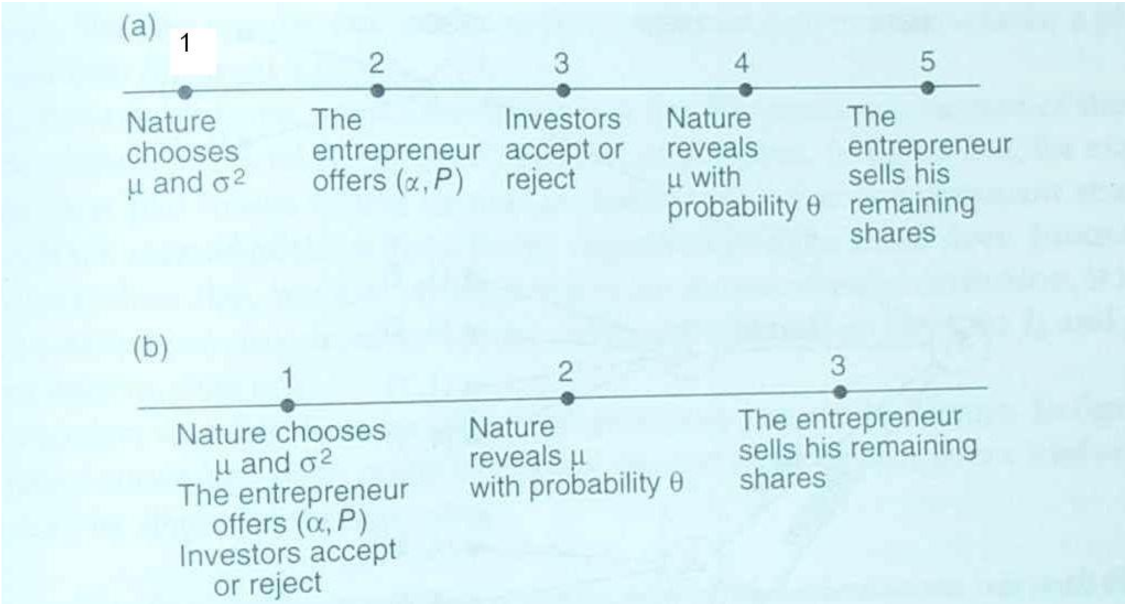

Figure 3: The Time Line for Stock Underpricing: (a) A Good Time Line; (b) A Bad Time Line

decision time versus **real time**

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.*

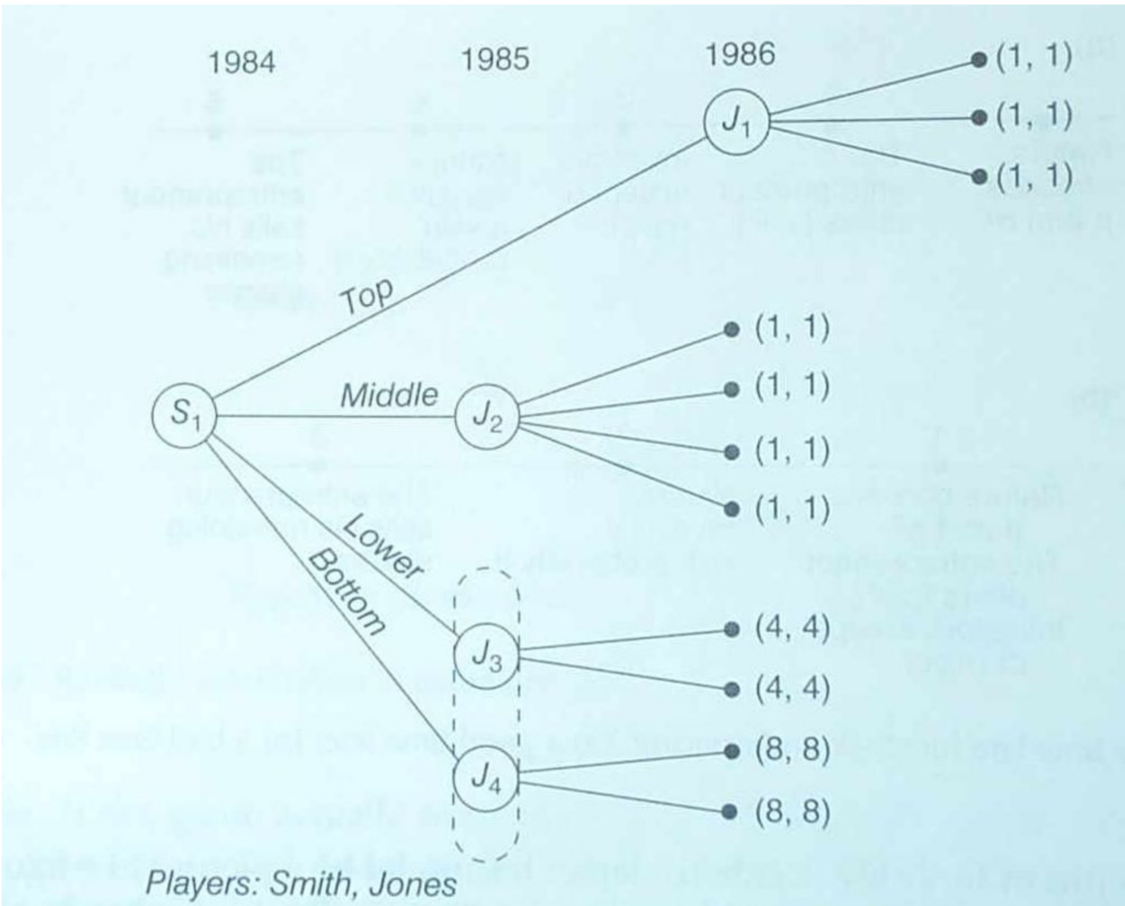

Figure 4: Information Sets and Information Partitions.

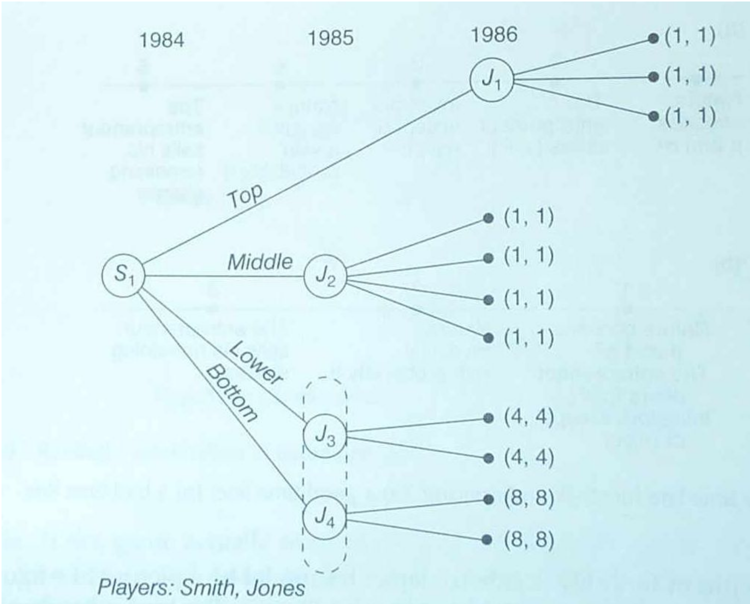

Figure 4: Information Sets and Information Partitions.

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).

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.

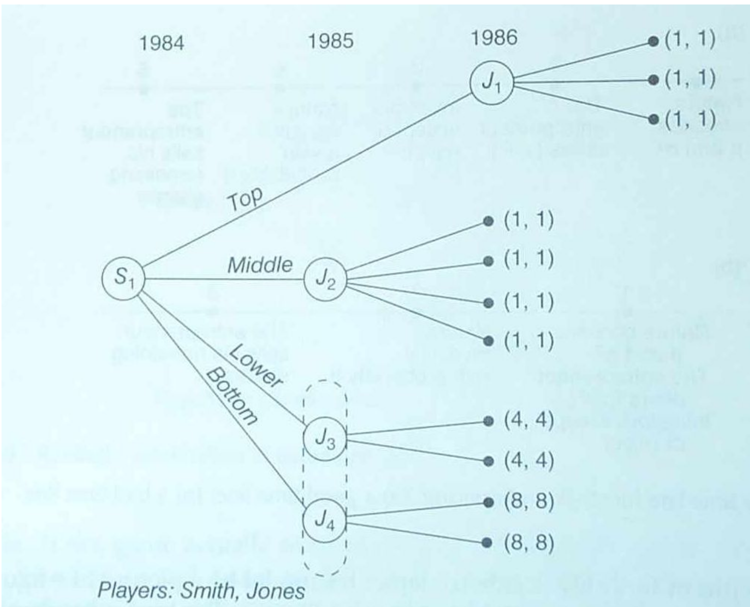

Figure 4: Information Sets and Information Partitions.

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.

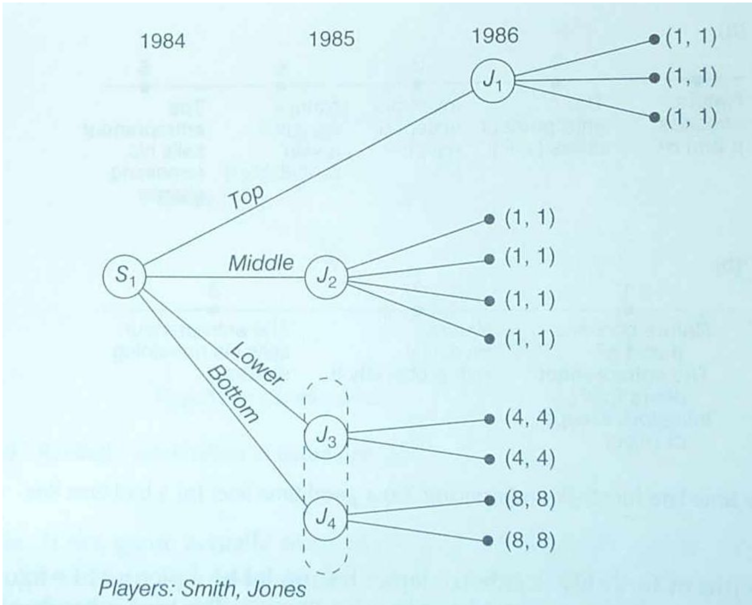

Figure 4: Information Sets and Information Partitions.

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>
<tr><th>Nodes</th><th>I</th><th>II</th><th>III</th><th>IV</th></tr>
<tr><td>J_1</td><td>$\{J_1\}$</td><td>$\{J_1\}$</td><td rowspan="4">$\left\{\begin{array}{c} J_1 \\ J_2 \\ J_3 \\ J_4 \end{array}\right\}$</td><td rowspan="3">$\left\{\begin{array}{c} J_1 \\ J_2 \\ J_3 \end{array}\right\}$</td></tr>
<tr><td>J_2</td><td>$\{J_2\}$</td><td>$\{J_2\}$</td></tr>
<tr><td>J_3</td><td>$\{J_3\}$</td><td rowspan="2">$\left\{\begin{array}{c} J_3 \\ J_4 \end{array}\right\}$</td></tr>
<tr><td>J_4</td><td>$\{J_4\}$</td><td>$\{J_4\}$</td></tr>
</table>

Table 3: Information Partitions

Partition II is **coarser**, and partition I is **finer**.

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.

A finer information partition is the formal definition for "better information."

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.

(c) Poor information may permit players to insure each other.

(c) Poor information may permit players to insure each other.

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.

Common Knowledge

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.*

Models are set up so that the extensive form is assumed to be common knowledge.

The Blue-Eyed Islander Puzzle. An island starts with 2 blue-eyed people and 48 brown-eyed, but the people do not know these numbers. If a person ever decides his eyes are blue, he must leave the island at dawn the next day. There are no mirrors and people may not talk about eye color, but they see each others' faces.

What will happen? – nobody leaves.

Now an outsider comes to the island and says, "At least one of you has blue eyes". The next dawn, nobody leaves, but on the second dawn, both blue-eyed people leave.

The reason: Both blue-eyed people realize there are either 1 or 2 blue-eyed people. When nobody leaves on the first dawn, each realizes that there must be 2– and he is one of them.

Information Categories:

Perfect: each information set is a singleton

Certain: Nature makes no moves

Symmetric: No player has information different from any other

Complete: Nature does not move first, or her initial move is public 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.**

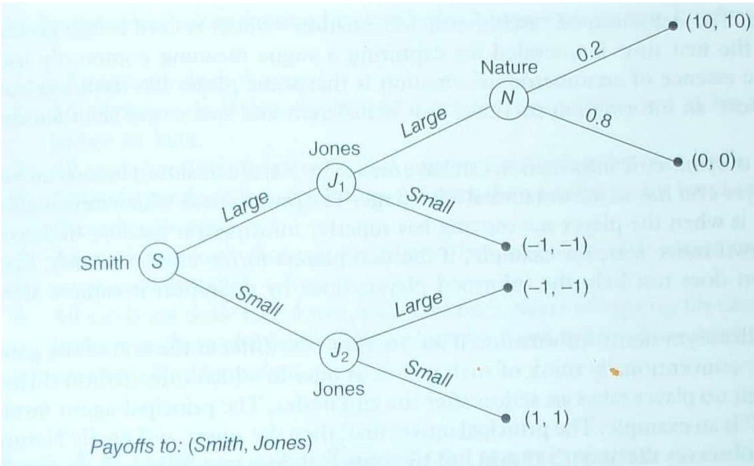

Figure 5: Follow-the-Leader II

von Neumann-Morgenstern utility functions are necessary when there is either uncertainty or random (mixed) strategies.

Maximizing their expected utilities, the players would behave exactly the same as in Follow-the-Leader I.

The players can differ in how they map money to utility– introducing risk aversion. 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.**

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.

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): Nature assigns different abilities to workers, but workers don't know them.

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 metarules, 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.)

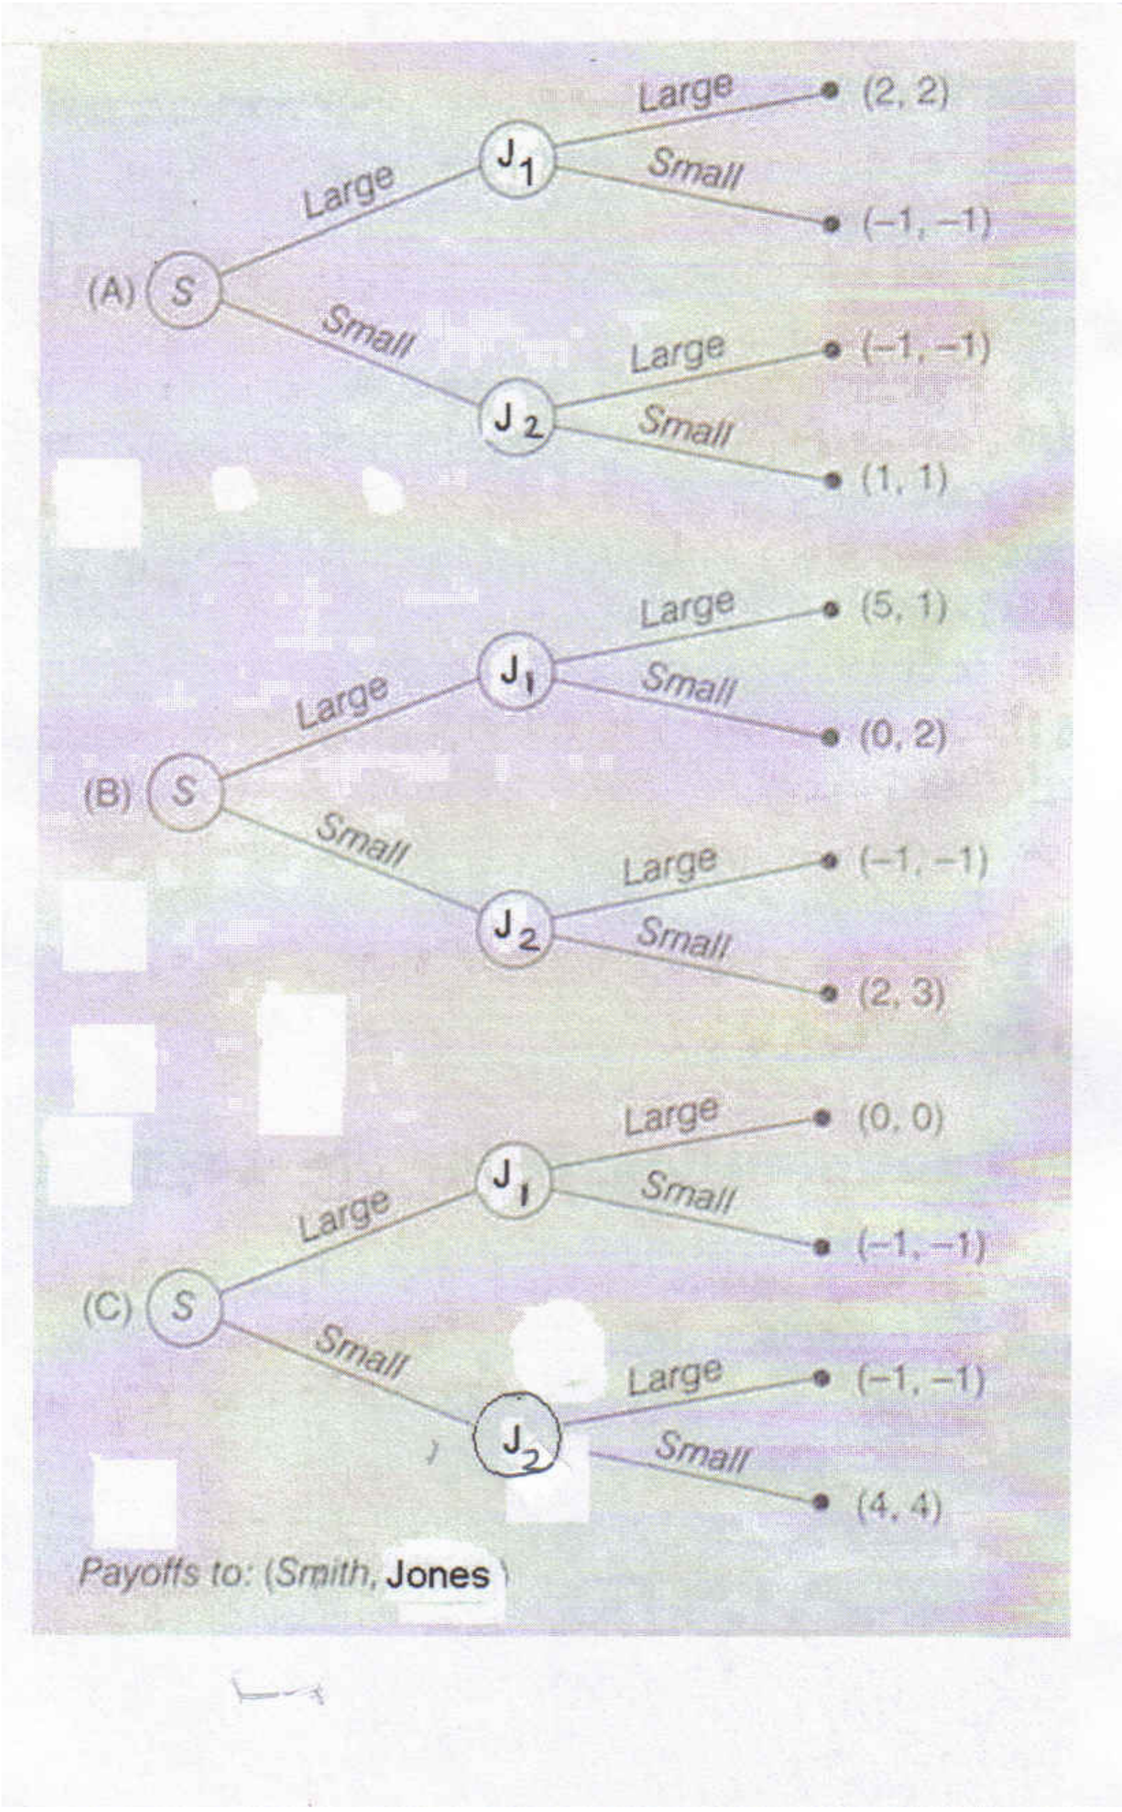
(A)
S
Large
Small
J1
J2
Large
Small
(2, 2)
(−1, −1)
(−1, −1)
(1, 1)
(B)
(5, 1)
(0, 2)
(−1, −1)
(2, 3)
(C)
(0, 0)
(−1, −1)
(−1, −1)
(4, 4)
Payoffs to: (Smith, Jones)

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.

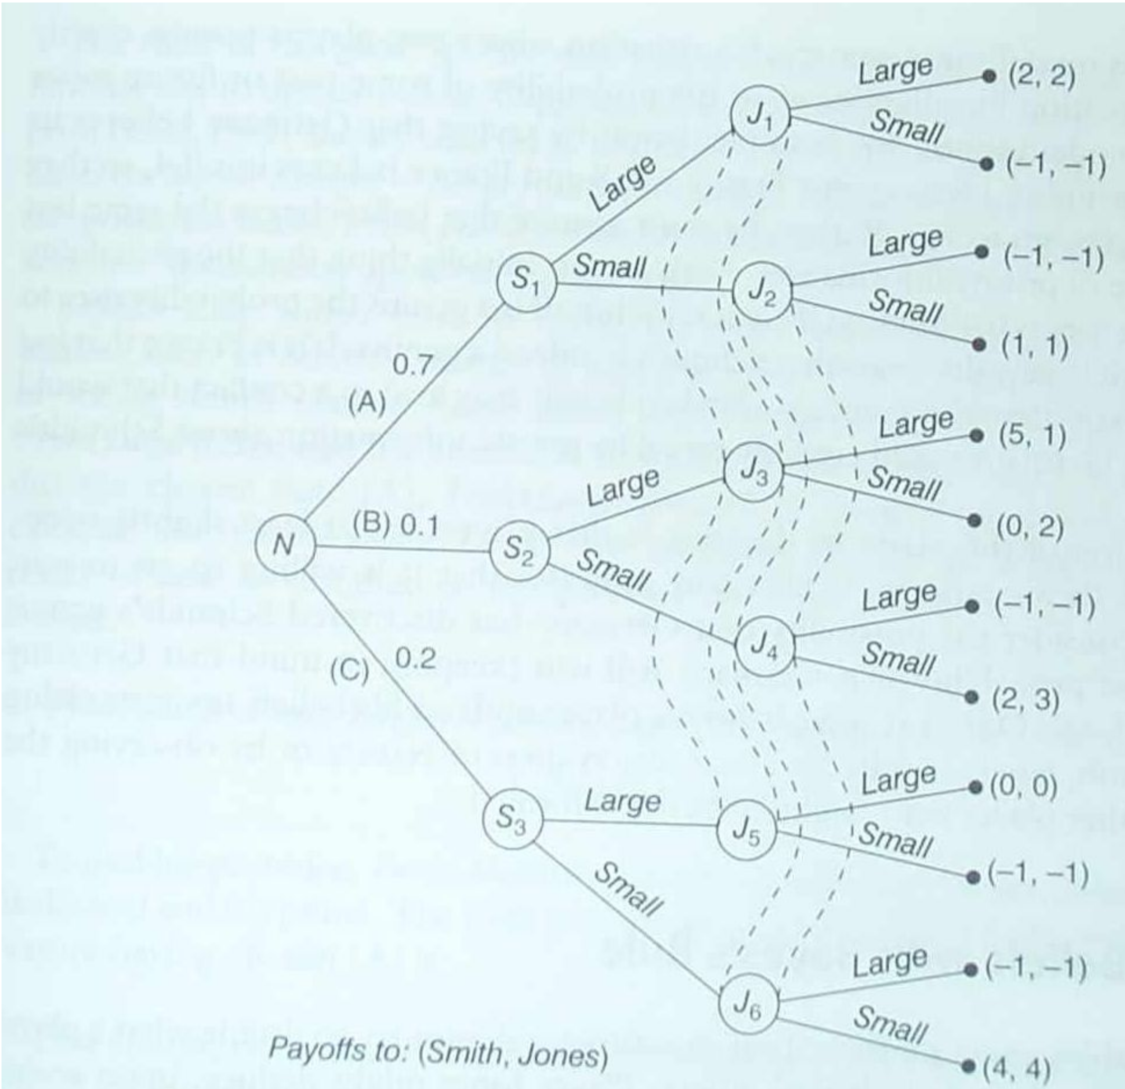

Figure 7: Follow-the-Leader III: After the Harsanyi Transformation

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, $Prob(Nature\ chose\ (A))$ 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.

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

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

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

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.

$$Prob(A|Large) = \frac{Prob(Large|A)Prob}{Prob(Large|A)Prob(A) + Prob(Large|B)Prob} \tag{4}$$

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

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

Equation (6) is a verbal form of Bayes's Rule, which is useful for remembering the terminology,[1] summarized

[1]The 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.

in Table 5.

$$(Posterior\ for\ Nature's\ \ Move) = \frac{(Likelihood\, of\ Player's\ \ Mo}{(Marginal\ likeliho} \tag{6}$$

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	$Prob(data\ \text{X}\|data\ \text{Y},\ event)$
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 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

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

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:

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

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.

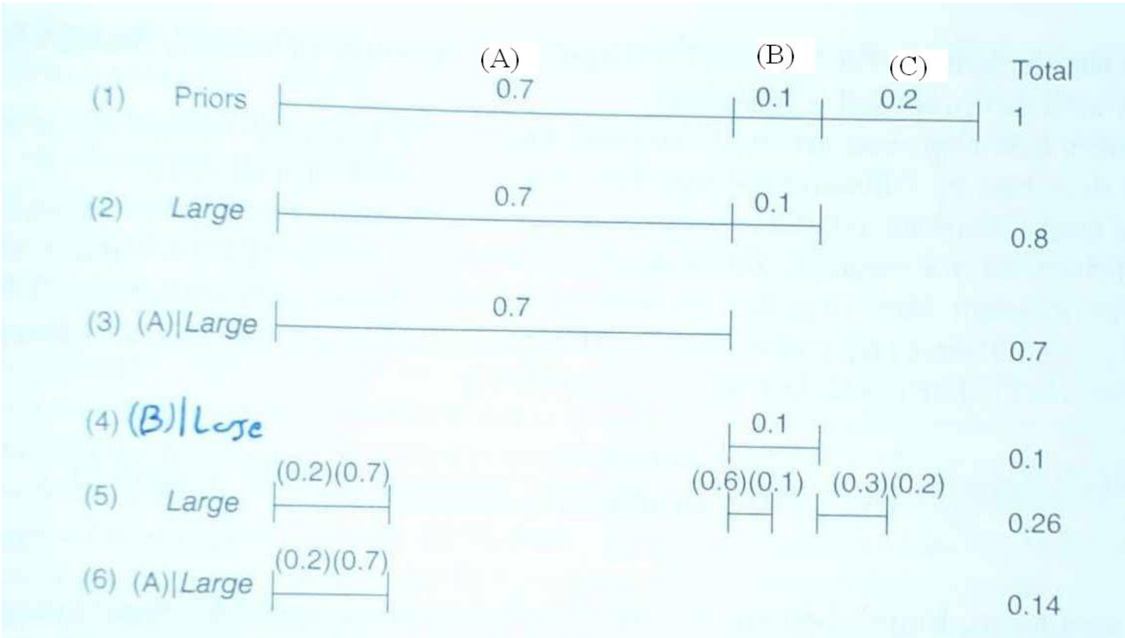

Figure 8: Bayes's Rule

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

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

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 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.

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.

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.

The "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.

Cascade model. Bikchandani, Hirshleifer & Welch (1992).

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.

Each has 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 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 $-S - W - D$.

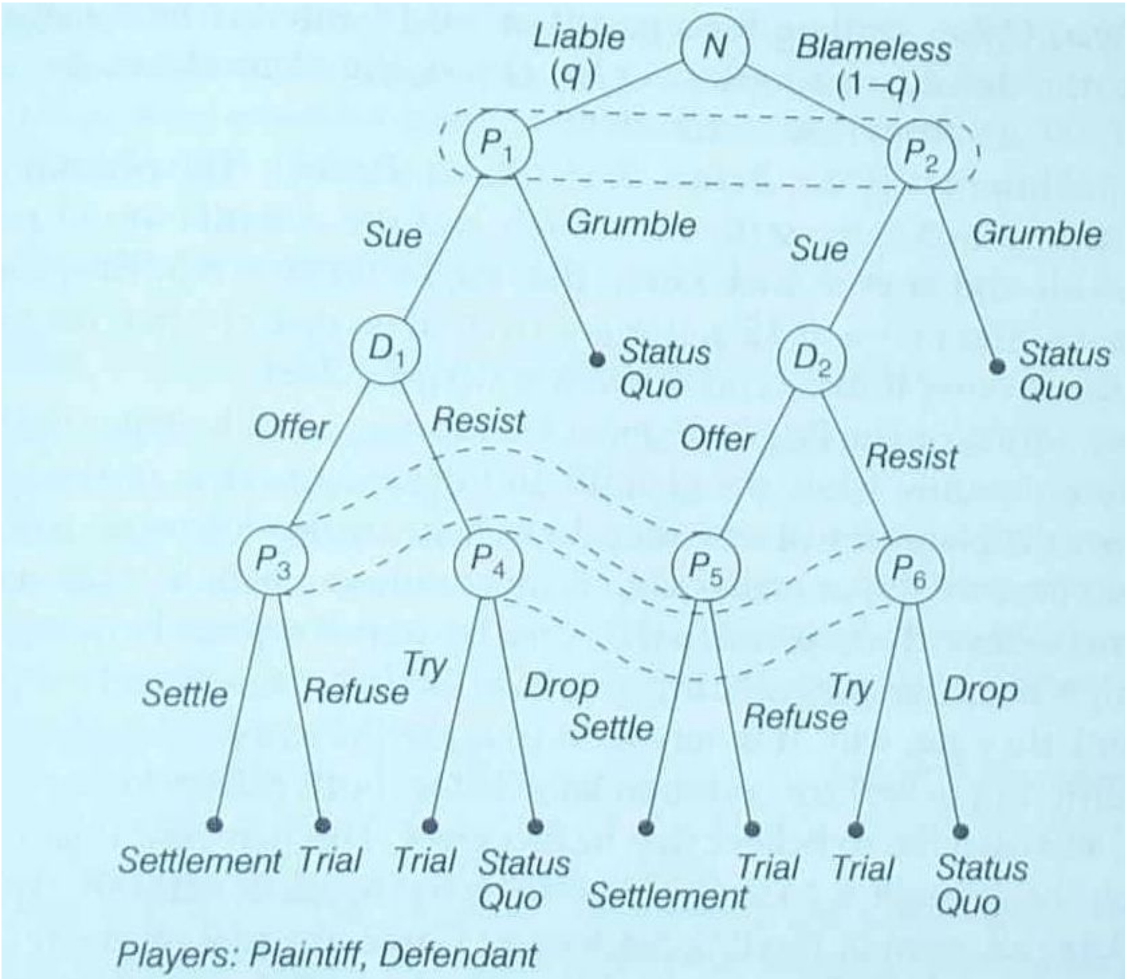

Figure 9: The Game Tree for The Png Settlement Game

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.