Chapter 2 Information Slides

January 22, 2014

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 **normal form** *or* **strategic 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: Strategic Form for Follow-the-Leader I

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

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

Table 2: Strategic Form for Follow-the-Leader I

		Jones			
		J_1	J_2	J_3	J_4
		$L\|L,\ L\|S$	$L\|L,\ S\|S$	$S\|L,\ L\|S$	$S\|L,\ S\|S$
Smith	S_1 : *Large*	**2**, **2** (E_1)	**2**, **2** (E_2)	-1, -1	$-1, -1$
	S_2 : *Small*	$-1, -1$	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** *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.

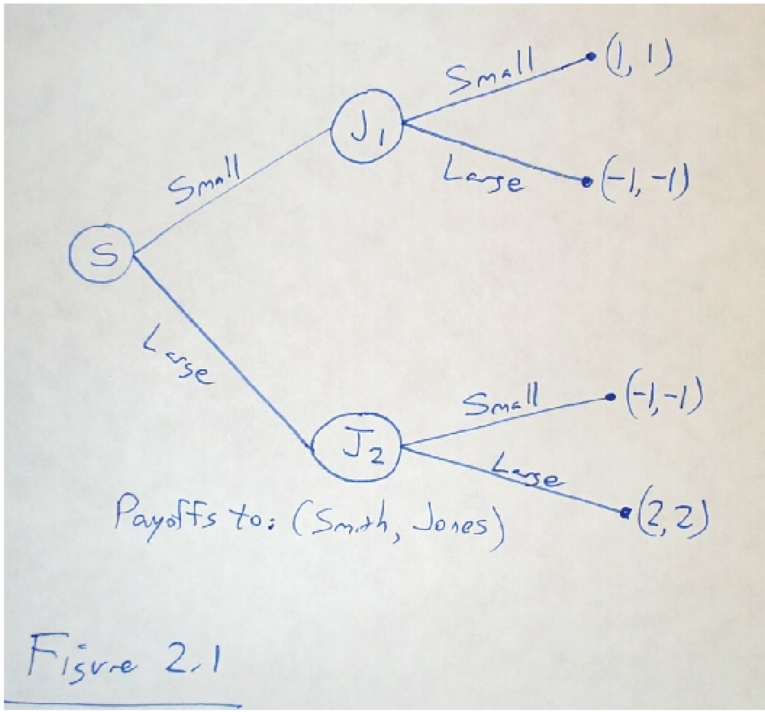

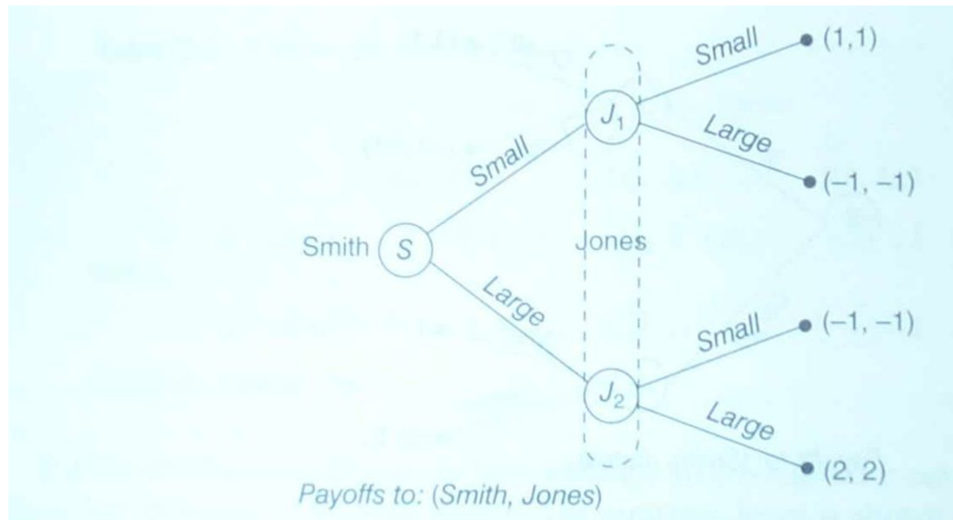

Follow-the-Leader I Ranked Coordination

The Time Line

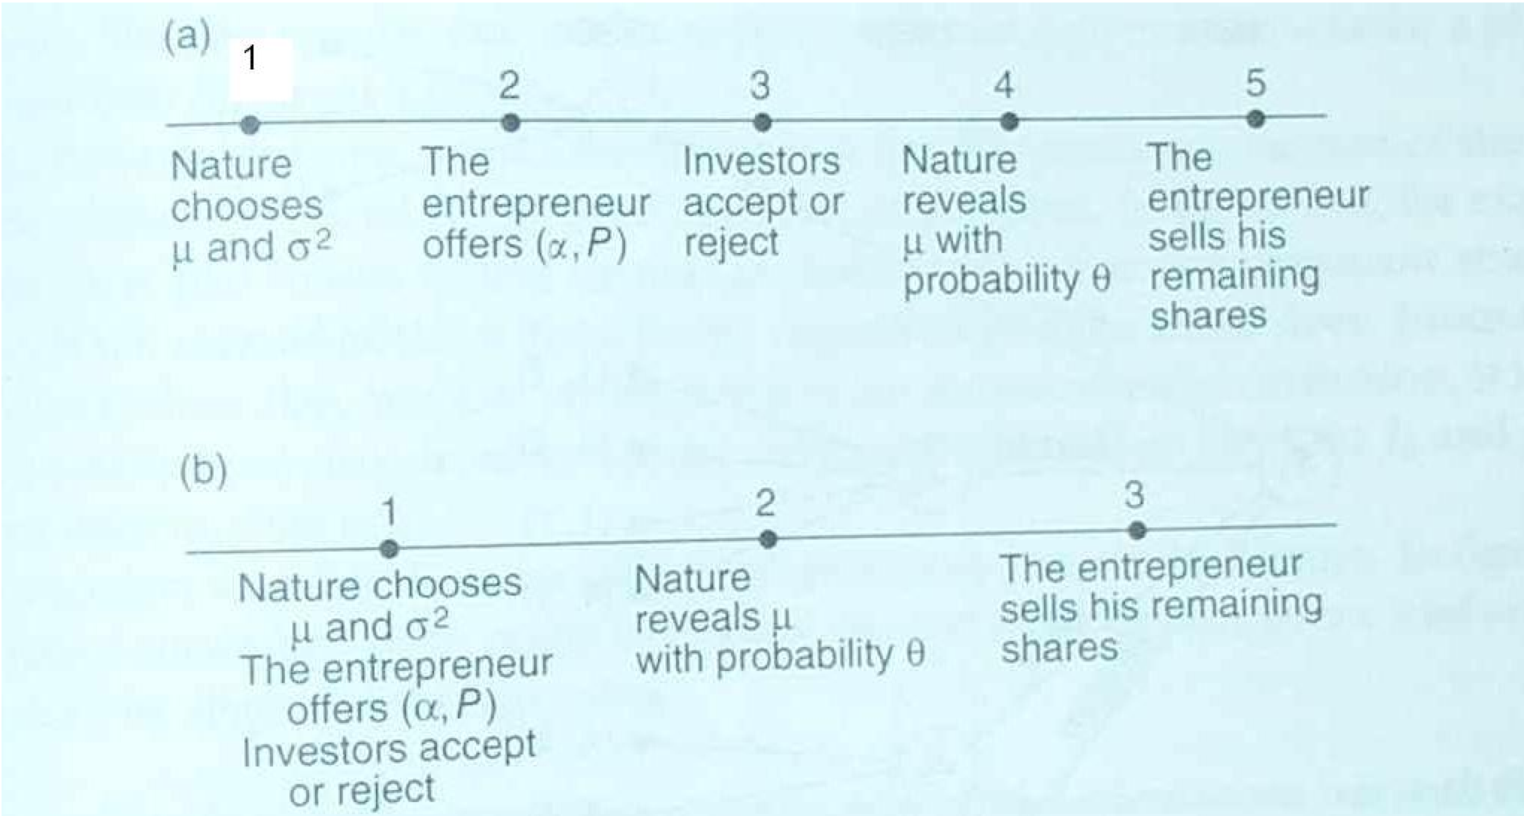

Figure 3: The Time Line for Stock Underpricing: (a) A Good Time Line; (b) A Bad Time Line

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.*

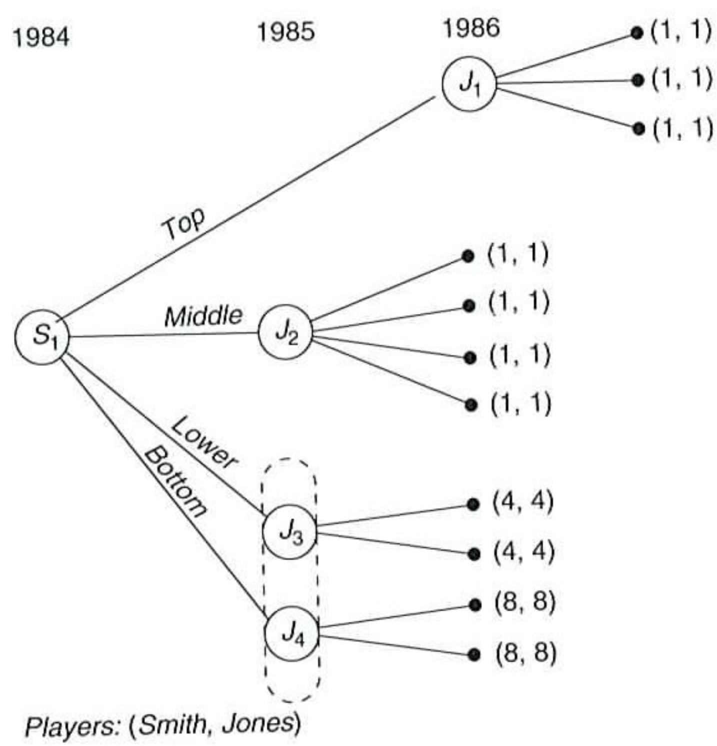

Figure 2.4 Information sets and information partitions.

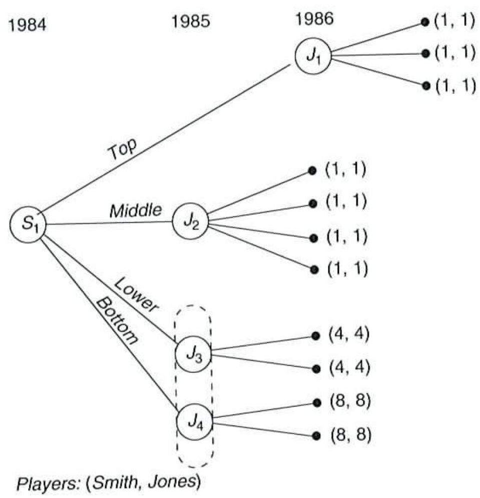

Figure 2.4 Information sets and information partitions.

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.

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.

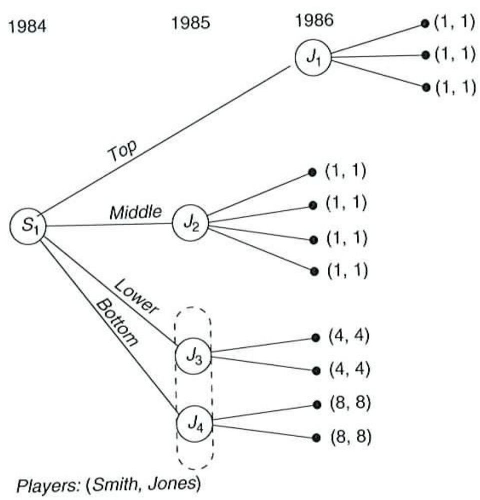

Figure 2.4 Information sets and information partitions.

Figure 4: Information Sets and Information Partitions.

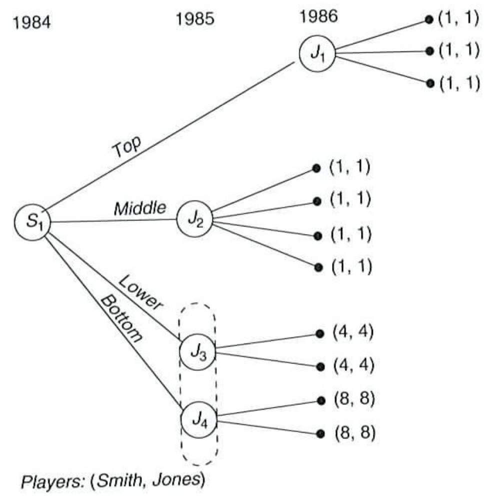

Figure 2.4 Information sets and information partitions.

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.

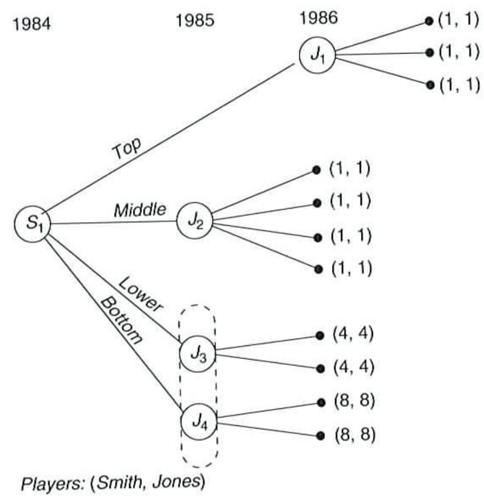

Figure 2.4 Information sets and information partitions.

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 2.3 Information partitions

<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">$\{J_1, J_2, J_3, J_4\}$</td><td rowspan="3">$\{J_1, J_2, J_3\}$</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">$\{J_3, J_4\}$</td></tr>
<tr><td>J_4</td><td>$\{J_4\}$</td><td>$\{J_4\}$</td></tr>
</table>

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.

A finer information partition is the definition of "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.

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.**

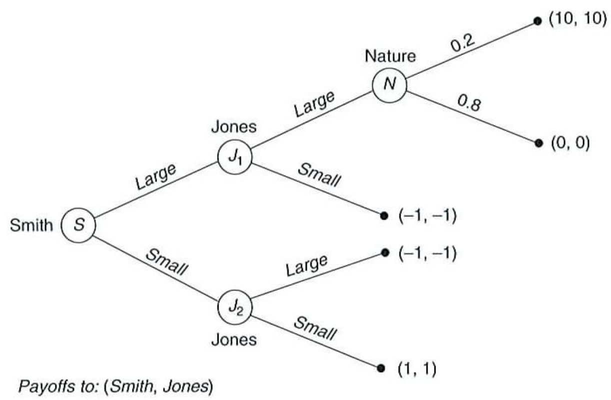

Figure 2.5 Follow-the-Leader II.

Figure 5: Follow-the-Leader II

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

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.

This is also known as a Bayesian Game.

2.4 The Harsanyi Transformation and Bayesian Games

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.)

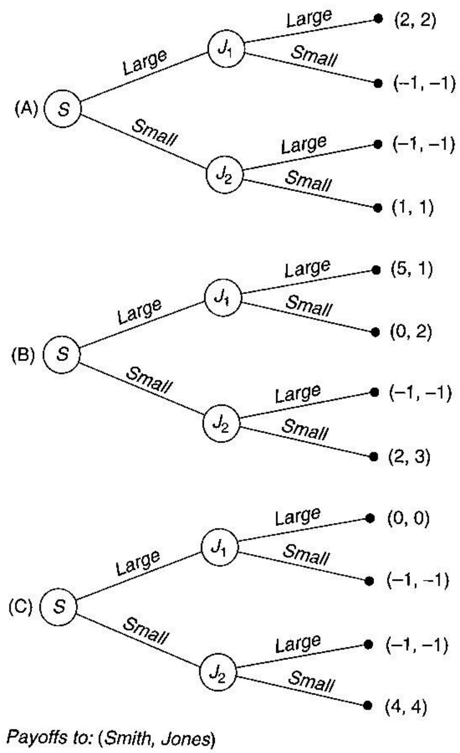

Figure 2.6 Follow-the-Leader III: original.

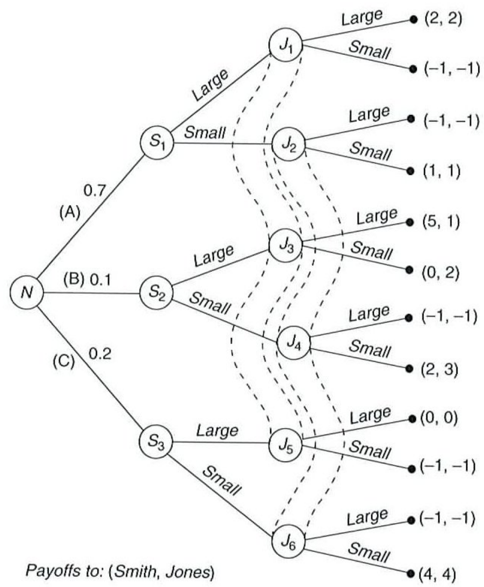

Figure 2.7 Follow-the-Leader III: after the Harsanyi transformation.

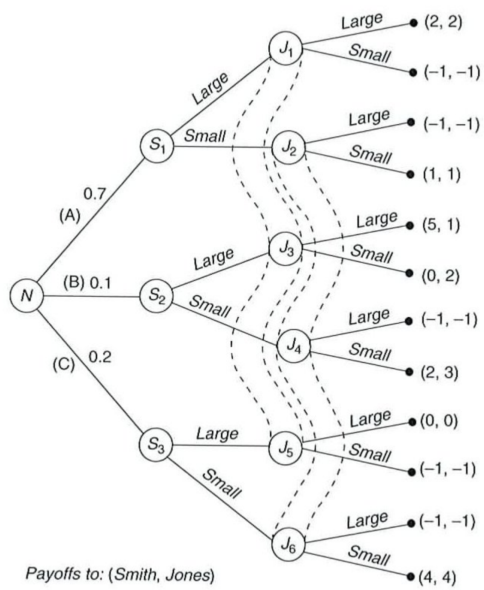

Figure 2.7 Follow-the-Leader III: after the Harsanyi transformation.

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.*

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 the Harsanyi doctrine, 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 0.6 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.

Here is way beliefs could diverge. 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 0.6, that if he isn't, it falls to .2, and that he is a genius with probability .5. Then Germany discovers that Schmidt is indeed a genius. Now, the two players have different beliefs. We model it as one of them observing a new move by Nature.

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. 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.**

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

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.

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

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

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

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

Table 5: Bayesian Terminology

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

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**

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

$$= 0.875.$$

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:

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

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.

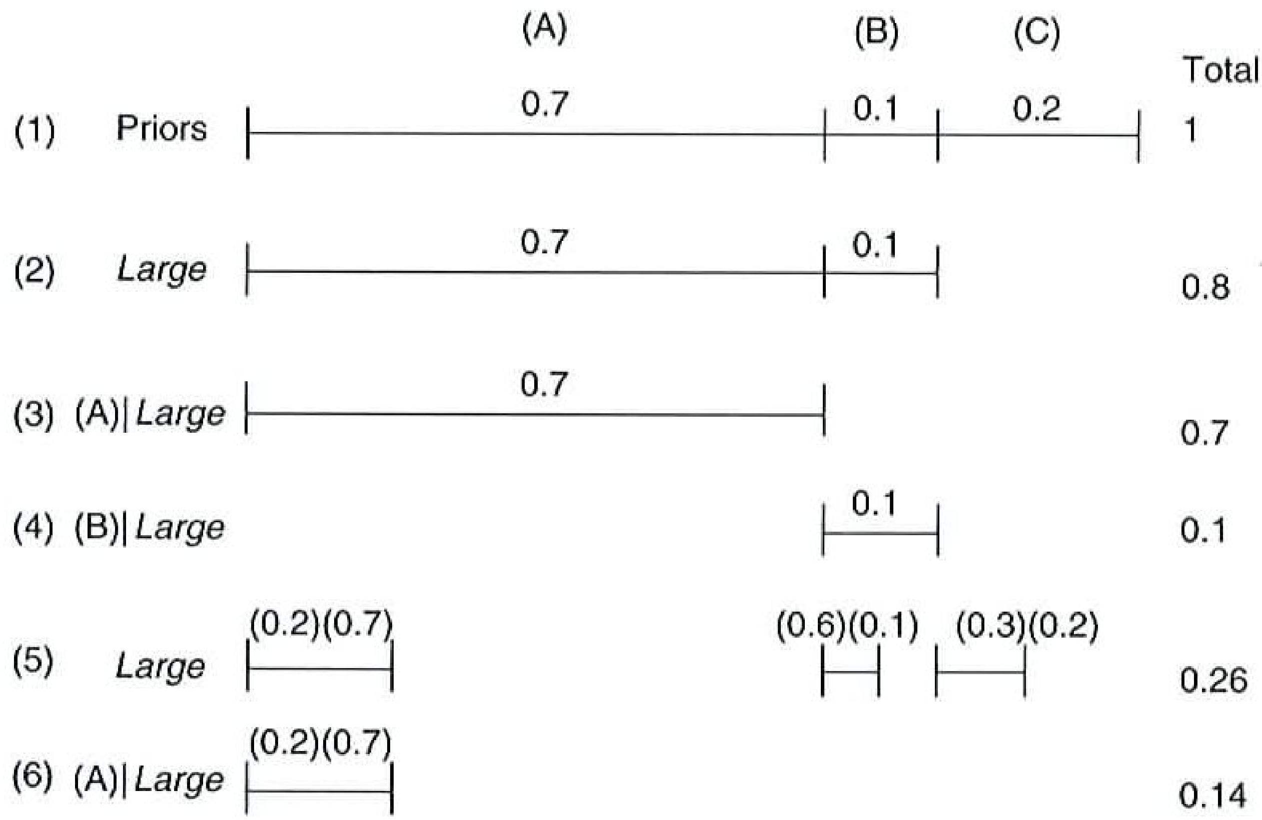

Figure 2.8 Bayes' Rule.

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

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

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

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

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 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.

A group of 30 people is told, "At least one of you has a white hat. How many of you have white hats? I will ask you several times, with a pause in between. If anybody knows, he should raise his hand. " It turns out that they can all deduce how many have white hats.

Let w be the number of people with white hats. A player's information partition at the start is that he can see how many other people have white hats, but he cannot tell if he himself has a white hat or not. Suppose he sees m white hats. His information partition has eliminated most states as a result. His information set by observation contains two kinds of states of the world:

(States of the world in which m other people have white hats and I have a black hat so $w = m$, States of the world in which m other people have white hats, and I have a white hat so $w = m + 1$.)

Together, these represent the event that "I saw m other people with white hats."

Let w be the number of people with white hats. Suppose our player sees m white hats.

Each player has a different partition, because "I" is different. Depending on what the truth is, m will vary. Each of the players can immediately deduce that either $w = m$ or $w = m + 1$ from the m that he observes.

But each player has extra information: that at least one person has a white hat. That rules out the single state of the world in which nobody has a white hat. Notice that that rules out the state "States of the world in which $m = 0$ other people have white hats, and I have a black hat so $w = 0$." So for $m = 0$, he puts 100% probability on:

(States of the world in which $m = 0$ other people have white hats and I have a white hat)

So he knows that exactly 1 person has a white hat, himself.

Okay— so that says that in the case that just one person has a white hat, that person will tell the announcer: "I know that $w = 1$; just one

person has a white hat!"

But what if the number is greater?

Well, in that case, nobody knows in the first round. But by the second round, they have acquired information from the silence of everyone else— that $w = 1$ is impossible, as well as $w = 0$. Suppose for someone that $m = 1$. That person had the information set in the first round:

(States of the world in which $m = 1$ other people have white hats and I have a black hat so $w = 1$, States of the world in which $m = 1$ other people have white hats, and I have a white hat so $w = 2$.)

That person will speak up at the second round and say, "I know there are two white hats!" He can rule out the $w = 1$ state of the world by deduction.

We'll go one more round. If that does NOT happen, and nobody speaks up in the second round, then they have learned that $w = 2$ is impossible also. Suppose someone had the information set

(States of the world in which $m = 2$ other people have white hats and I have a black hat so $w = 2$, States of the world in which $m = 2$ other people have white hats, and I have a white hat so $w = 3$.)

That person can rule out the $w = 2$ event, so he can conclude that 3 people have white hats. And if nobody sees $m = 2$ we can continue to $m = 3$ and beyond, until eventually even if $w = 30$ the people will realize it.

The value of this example is in showing how the modeller starts by narrowing down players' information to what they might be seeing from direct information, and then goes on to see what they can deduce from extra information.