

# Strategies in Extensive-Form Games

<i>Recapitulation</i>	<u>1</u>
<i>Strategy profiles of extensive-form games</i>	<u>2</u>
<i>The strategic form of an extensive-form game</i>	<u>5</u>
<i>When Nature takes her turn at bat</i>	<u>7</u>
<i>Randomized strategies in extensive-form games</i>	<u>9</u>
Behavior strategies	<u>9</u>
Mixed strategies	<u>12</u>
Mixed strategy → behavior strategy: the role of perfect recall	<u>13</u>
Mixed strategy → behavior strategy: computing the conditional probabilities	<u>16</u>
Behavior strategy → mixed strategy	<u>18</u>
Many mixed strategies give rise to the same distribution over outcomes	<u>19</u>
Kuhn's theorem: The equivalence of behavior and mixed strategies	<u>20</u>
<i>The restriction of a strategy to a subgame</i>	<u>21</u>

## Recapitulation

We are now studying *dynamic* games—in which some decisions are made after others. We use the extensive form because it allows us to make explicit this temporal structure as well as to define the game's information structure. (The results of previous decisions by one player may not be observable immediately, if ever, to her opponents.)

We introduced the *game tree* as the supporting framework for the extensive form. Choices are made at *decision nodes*, which belong to the set  $X$ . One of these, viz.  $\mathbb{O}$ , is designated the *initial node*, at which the game begins. It is assigned either to one of the players in the player set  $I = \{1, \dots, n\}$  or to Nature; Nature would make a random choice representing any exogenous uncertainty by some or all of the players. Play progresses from node to node based on the players' decisions. The game ends when any one of the *terminal nodes*, which belong to the set  $Z$ , is reached, at which point the players are awarded the payoffs corresponding to that node. Player  $i$ 's preferences over terminal nodes are represented by her von Neumann-Morgenstern utility function  $\mu_i: Z \rightarrow \mathbb{R}$ ; namely, if terminal node  $z \in Z$  is reached, her utility is  $\mu_i(z)$ .

When called upon to move, a player might not know precisely at which of her nodes she is located. Her uncertainty is modeled by *information sets*. An information set  $h \in H$  is a set of nodes, all belonging to the same player and at all of which the same set of actions is available. When called upon to act, a

player is informed only of the information set which has been reached, i.e. that she is located at some node in that information set. We typically impose additional constraints on the structure of the information partition to assure that the game is one of *perfect recall*—one in which a player always knows what move she takes and she never forgets anything she ever knew.

We saw that an extensive game could often be decomposed into smaller parts called *subgames*, which are subtrees which respect information sets. A subgame makes sense as a game by itself, and the subgame embodies the informational conditions under which it would be played if reached in the original game. (If the subgame were reached in the larger game, it would be common knowledge that this was the game remaining to be played.)

## Strategy profiles of extensive-form games

A player's *strategy* specifies what action she would take whenever called upon to move. To be more precise: A player's strategy is a specification of what action she would implement *at each of her information sets*. We require the strategy to be in general a function of the information set because knowledge of the information set which has been reached is all the information a player has when called upon to move, and we certainly couldn't let a player's decision depend on information she doesn't have.

For example, in the game depicted in Figure 1, player 1 has four information sets:  $\alpha$ ,  $\delta_1$ ,  $\delta_2$ , and  $\delta_3$ , so  $H_1 = \{\alpha, \delta_1, \delta_2, \delta_3\}$ .<sup>1</sup> (Note that information set  $\delta_3$  contains two nodes, while the other three information sets are singletons.) Therefore a strategy for player 1 would be a 4-tuple of strategies, containing one strategy for each of her information sets, which we could write in the following form:

$$s_1 = (s_1(\alpha), s_1(\delta_1), s_1(\delta_2), s_1(\delta_3)), \quad (1)$$

where each  $s_1(h)$  is a specification of her action at the information set  $h \in H_1$ . For example consider the strategy

$$s_1 = (Y, D, D, U), \quad (2)$$

which is graphically depicted by the thick gray line segments in Figure 2. (Note that the “U” for the doubleton information set  $\delta_3$  generates two thick line segments, because there are two nodes in that information set and therefore two nodes at which that information set's action would be taken.)

---

<sup>1</sup> Note that in the previous handout: “Extensive-Form Games,” we commonly labeled nodes in a tree. Here I'm labeling information sets. E.g. “ $\alpha$ ” here is the singleton information set which contains the initial node:  $\alpha = \{\circ\}$ .




Figure 1

Similarly, player 2 has two information sets  $\beta_1$  and  $\beta_2$  (both singletons), i.e.  $H_2 = \{\beta_1, \beta_2\}$ . His strategies are of the form

$$s_2 = (s_2(\beta_1), s_2(\beta_2)). \tag{3}$$

For example the strategy

$$s_2 = (r, l), \tag{4}$$

is graphically depicted by the darker thick line segments in Figure 2.




Figure 2: The strategy profile  $(Y, D, D, U; r, l)$ .

A *strategy profile* is an  $n$ -tuple of individual strategies, one strategy for each player. Specifying a strategy profile completely determines the outcome of the game. We can write a strategy profile as a concatenation of the strategies for each player; i.e. it is of the form:

$$s = (s_1(\alpha), s_1(\delta_1), s_1(\delta_2), s_1(\delta_3); s_2(\beta_1), s_2(\beta_2)).^2 \tag{5}$$

The strategy profile corresponding to our exemplar strategies (2) and (4) would be

$$s = (Y, D, D, U; r, l). \tag{6}$$

This strategy profile is represented in Figure 2 by the union of the gray segments and the darker thick line segments. Notice that the graphical representation of this strategy profile is not connected; i.e. there are disjoint pieces. There is only one component of this union which reaches from the initial node  $\alpha$  to a terminal node.<sup>3</sup> This is the *path* which corresponds to the strategy profile  $s$ . Its outcome is the starred terminal node.

We can be more formal in defining a player’s extensive-form strategies. Recall that  $H_i$  is the set of player  $i$ ’s information sets and, for any player- $i$  information set  $h \in H_i$ ,  $A(h)$  is the set of actions available to her at that information set.<sup>4</sup> A strategy  $s_i$  for player  $i$  specifies an action for each information set; therefore for each  $h \in H_i$  we represent player  $i$ ’s choice at  $h$  by  $s_i(h) \in A(h)$ . A strategy  $s_i$  is an  $\#H_i$ -tuple of actions. The set of actions which player  $i$  has available *somewhere* is the union of the actions she has

<sup>2</sup> The semicolon serves to separate the two strategies in the profile.  
<sup>3</sup> You might find it challenging to prove this in general using graph theory.  
<sup>4</sup> See “Extensive-Form Games.”

available at each information set, viz.  $\bar{A}_i \equiv \bigcup_{h \in H_i} A(h)$ ; so we can describe her strategy as a function  $s_i: H_i \rightarrow \bar{A}_i$ . Her strategy space is the Cartesian product

$$S_i = \prod_{h \in H_i} A(h). \quad (7)$$

As we did in the case of strategic-form games, we denote the space of strategy profiles by  $S \equiv \prod_{i \in I} S_i$ .

Each strategy profile  $s \in S$  results in a particular terminal node  $z \in Z$  being reached with certainty. We denote the identity of this terminal node by  $\hat{z}(s)$ ; i.e.  $\hat{z}: S \rightarrow Z$  is a function from strategy profiles into terminal nodes. Player  $i$ 's utility associated with this strategy profile is her utility for the terminal node reached when that strategy profile is played. Therefore we define for each player  $i$  the function  $u_i: S \rightarrow \mathbb{R}$ , which specifies player  $i$ 's utility for any pure-strategy profile by:  $\forall s \in S$ ,

$$u_i(s) \equiv \mu_i(\hat{z}(s)). \quad (8)$$

As usual we let  $u: S \rightarrow \mathbb{R}^n$  be the  $n$ -tuple of individual utility functions  $u_i$ .

## The strategic form of an extensive-form game

Recall that when we studied strategic-form games we specified a game by a triple:  $(I, S, u)$ , where  $I$  was the player set,  $S$  was the space of strategy profiles, and  $u$  was an  $n$ -tuple of player utility functions  $u_i: S \rightarrow \mathbb{R}$ .<sup>5</sup> In this framework, strategy was a primitive concept: A pure-strategy  $s_i$  had no internal structure and a player's strategy space  $S_i$  was part of the specification of the game. In the extensive form, on the other hand, the game tree, the player partition  $\mathcal{X}$  (which assigns decision nodes to players), the information partition  $H$  (which assigns decision nodes to information sets), and the sets of actions  $A(x)$  available at the decision nodes  $x \in X$  are primitive specifications of the game. A player's strategy space  $S_i$  is *derived* from these according to (7).

How then can we construct a strategic-form representation of an extensive-form game? We need to provide the constituents of the  $(I, S, u)$  triple. The player set  $I$  can obviously be supplied directly. Likewise, once we have constructed each player  $i$ 's strategy space  $S_i$  from the extensive-form primitives according to (7), we can simply supply their Cartesian product  $S$  as the space of strategy profiles for the strategic-form representation. To complete the strategic-form specification we pass along the utility function  $u: S \rightarrow \mathbb{R}^n$  constructed from the utility functions  $u_i$  from (8).

One consequence of this one-to-one correspondence between strategic-form and extensive-form strategies is that we can find the pure-strategy Nash equilibria of an extensive-form game in extensive-form strategies by finding the pure-strategy Nash equilibria of the game's strategic form. (Consider a conjectured Nash-equilibrium pure-strategy profile  $s^*$ . When viewed as extensive-form strategies, player  $i$  will have a better strategy than  $s_i^*$  against  $s_{-i}^*$  if and only if she does when they are viewed as

---

<sup>5</sup> See "Strategic-Form Games."

strategic-form strategies.)

**Example: Constructing the strategic form of an extensive-form game.**

Consider the extensive-form game in Figure 3a. Player 1's strategy space is  $S_1 = \{U, D\}$ ; player 2's is  $S_2 = \{(l, l), (l, r), (r, l), (r, r)\}$ , where we write a typical strategy for player 2 as a pair: (action if player 1 chooses  $U$ , action if player 1 chooses  $D$ ).




Figure 3: A game in (a) extensive form and (b) strategic form.

In order to cast an extensive-form game into strategic form we simply identify each extensive-form pure strategy  $s_i$  with a strategic-form pure strategy. To construct the bimatrix representation of this game we assign player 1 to rows and player 2 to columns, listing each player's strategies along her side of the bimatrix. To find the payoff vector corresponding to any given cell (strategy profile) we identify the path determined by that strategy profile; the payoff vector located at that terminal node is the one entered into the bimatrix. For example, consider the strategy profile  $(D; r, l)$ . This strategy pair is displayed in Figure 4; the payoff vector  $(4, 0)$  located at the end of its path is entered into the corresponding cell of the bimatrix of Figure 3b. The payoff vectors for the other strategy profiles are calculated similarly.




Figure 4: The strategy profile  $(D; l, r)$ .

Note that there are four payoff vectors in the extensive-form representation in the game in Figure 3a, and there are eight payoff vectors in its strategic form. For any strategy choice by player 1, one of player 2's information sets will be "off the path." For example, when player 1 chooses Up, player 2's right-hand information set is off the path. Player 2's choice at an off-the-path information set is irrelevant to

the payoff. This explains why more than one strategy profile can result in the same outcome—for example why the payoff vectors corresponding to  $(U; l, l)$  and  $(U; l, r)$  are the same.

Several extensive forms can have the same strategic form. For example, you can easily show that both extensive forms in Figure 5, which represent a simultaneous-move game, have the same strategic form.




Figure 5: Two extensive forms of the same simultaneous-move game.

## When Nature takes her turn at bat

Our game-playing agents do not always completely control their destinies; significant events can occur which are exogenously determined. If these events occur and are revealed to the players in such a way as to make them common knowledge prior to the commencement of the game, the occurrences can be built into the structure and payoffs of the extensive-form game. In this class of cases our present framework is adequate without enhancement.

However, it is possible that the players know that some event has, will, or might occur, but that one or more of the players are uncertain, when called upon to move, of exactly which event occurred. For example, when two oligopolistic firms simultaneously choose their output levels for their new no-cholesterol French fries, they are uncertain whether the University of Arizona College of Medicine will soon issue a report showing that cholesterol is actually health promoting. (See, for example, Allen [1973].) In an alternative scenario one firm, thanks to a spy in the relevant research laboratory, would know whether or not this will happen, while the other firm would not be privy to this information.

We incorporate uncertain exogenous events into the extensive form by introducing Nature as a nonstrategic player who acts randomly. By saying that Nature is nonstrategic I mean that her action is not influenced by the actions of the “real,” strategic players. Nature is not an optimizer; we don’t need to specify payoffs for her. Figure 6 represents the extensive form for the two-firm case of the above-described game in which neither firm knows whether bad news will be published.<sup>6</sup>

<sup>6</sup> Note that all four of player 2’s nodes are in one information set. He can see neither Nature’s move nor his opponent’s.




Figure 6: Cholesterol: friend or foe?

How do we construct the strategic form of an extensive-form game when Nature has a turn? In the game of Figure 6 Nature's random choice induces a probability distribution over the nodes in player 1's information set. Let  $p$  be the probability that Nature sends bad news. Each player has one information set and two actions at each information set; therefore each player's strategy space contains two strategies. Therefore the strategic form will be a  $2 \times 2$  bimatrix. Let  $p = \frac{1}{2}$  and let's compute the expected payoff vector to the pure-strategy profile  $(L; H)$ ; i.e. when firm 1 chooses the Low quantity and firm 2 chooses the High quantity. (This strategy profile is depicted in Figure 7.)

There are two terminal nodes which can be reached when player 1 chooses  $L$  and player 2 chooses  $H$ , viz. the second and sixth from the left with payoff vectors  $(-1, -3)$  and  $(5, 15)$ , respectively. The first of these is reached under the pure-strategy profile  $(L, H)$  when Nature sends bad news; the second is reached when "no news is good news." Therefore the expected payoff vector to  $(L, H)$  is

$$\frac{1}{2}(-1, -3) + \frac{1}{2}(5, 15) = (2, 6). \quad (9)$$

The expected payoff vectors for the other three pure-strategy profiles are calculated the same way. This strategic form is shown in Figure 8.




Figure 7: Firm 1 chooses Low quantity and firm 2 chooses High quantity.

	L	H
L	4,4	2,6
H	6,2	0,0

Figure 8: The strategic form of the extensive form of Figure 6, when Nature sends bad news half of the time.

## Randomized strategies in extensive-form games

### Behavior strategies

When we studied strategic-form games we saw that randomized strategies, viz. mixed strategies, were useful because randomized strategies could be Nash equilibria and because without randomized strategies Nash equilibria could fail to exist. We will find randomized strategies useful in the extensive form as well. We define a natural conception of strategic randomization in the extensive form called *behavior strategies*, in which players randomize independently at each information set.

It is easy to accept that there is a one-to-one correspondence between pure strategies in the extensive form and pure strategies in the strategic form. It is less obvious what correspondence exists between behavior strategies in the extensive form and mixed strategies in the strategic form. We will see that there is a useful correspondence, though not one-to-one, provided our game is one of perfect recall: any probability distribution over outcomes which can be achieved by a mixed strategy can be achieved by a behavior strategy and vice versa. This result gives us freedom to analyze an extensive-form game in whichever framework—strategic or extensive form—is more convenient. In particular we can find the set of Nash equilibria in an extensive-form game by computing the mixed-strategy Nash equilibria in the

strategic form.<sup>7</sup>

How should we define a randomized strategy in the extensive form? A pure strategy specifies a particular action at each of a player's information sets. One natural extension of this concept is the *behavior strategy*, which specifies at each information set a conditional probability distribution over the actions available at that information set; i.e. it specifies the probability with which each action would be taken conditional upon that information set being reached. (A pure strategy would be a degenerate case in which at each information set the probability distribution at that information set placed full weight upon a single action available at that information set.) We stipulate that the randomizations at the player's information sets are performed independently of one another.<sup>8</sup>

For example, in the game from Figure 1, a possible behavior strategy for player 1 would specify that she play  $Y$  two-thirds of the time at information set  $\alpha$ , play  $D$  with certainty at information set  $\delta_1$ , mix with equal probability at information  $\delta_2$ , and play  $U$  four-fifths of the time at each of the two nodes in information set  $\delta_3$ . Referring to her information sets in the same order as in the pure-strategy specification (1), we could write player 1's behavior strategy as

$$b_1 = (\frac{2}{3} \circ Y \oplus \frac{1}{3} \circ N, D, \frac{1}{2} \circ U \oplus \frac{1}{2} \circ D, \frac{4}{5} \circ U \oplus \frac{1}{5} \circ D). \quad (10)$$

Similarly, referring to (3), we could write a behavior strategy for player 2 as

$$b_2 = (\frac{1}{2} \circ l \oplus \frac{1}{2} \circ r, \frac{3}{4} \circ l \oplus \frac{1}{4} \circ r). \quad (11)$$

Rather than indicate a deterministic path as we did in Figure 2 for the pure-strategy profile in (2), we depict the behavior-strategy profile  $b = (b_1, b_2)$  by labeling in brackets each branch of the tree with the probability that the branch is chosen conditional on its source node being reached. See Figure 9.

What are these behavior strategies formally? Player  $i$ 's behavior strategy  $b_i$  specifies at each information set  $h \in H_i$  a probability for each action  $a \in A(h)$  available at  $h$ . We denote this probability by  $b_i(a|h)$ . The set of probability distributions over the set of actions  $A(h)$  is  $\Delta(A(h))$ .<sup>9</sup> Therefore for each  $h \in H_i$ ,  $b_i(\cdot|h) \in \Delta(A(h))$ . Therefore for every player  $i \in I$  and for every information set  $h \in H_i$ ,

$$\sum_{a \in A(h)} b_i(a|h) = 1. \quad (12)$$

Player  $i$ 's behavior strategy  $b_i$  specifies such a distribution for each information set in the finite set  $H_i$ . Therefore we can write  $B_i$ , the space of player- $i$  behavior strategies, as the Cartesian product:

<sup>7</sup> Kuhn [1953] is the seminal paper concerning these issues. See also Kreps [1990: 380–384], Fudenberg and Tirole [1991: 85–90], and Myerson [1991: 154–163] for discussions of behavior strategies and their relation to mixed strategies.

<sup>8</sup> Note that this is a stronger sense of independence than we encountered with mixed strategies of strategic-form games. There we assumed that each player randomized independently from every other player. Here we are saying in addition that each player randomizes at each of her information sets independently from her randomizations at other information sets.

<sup>9</sup> For any finite set  $T$  we denote by  $\Delta(T)$  the set of probability distributions over  $T$ .




Figure 9: The conditional probabilities specified by, and the probability distribution over outcomes determined by, the example behavior-strategy profile.

$$B_i = \prod_{h \in H_i} \Delta(A(h)). \tag{13}$$

[Compare this to (7), which describes player  $i$ 's pure-strategy space.]

When a behavior-strategy profile is played, there need not be a single terminal node reached with probability one; instead there will be a probability associated with each terminal node. The resulting probability distribution over outcomes is of interest because it determines the expected payoffs to the players. To compute the probability distribution over terminal nodes associated with the behavior-strategy profile we calculate the probability with which each path will be traversed. For a particular terminal node  $z$  we identify the unique path from the initial node to  $z$ . At each node along this path a particular action is required to keep the play on this path. We refer to the relevant player's behavior strategy to obtain the probability with which this required action is chosen conditional on reaching that node's information set. Because the randomizations at the various information sets are independent of one another, the probability that this path is traversed from start to finish is the product of the probabilities associated with the required actions at all the nodes along the path.

For example in Figure 9 to arrive at terminal node  $z_3$  player 1 must choose  $Y$  at information set  $\alpha$ , player 2 must choose  $r$  at  $\beta_1$ , and player 1 must finally choose  $U$  at  $\delta_2$ . The probabilities associated with these actions, according to the behavior-strategy profile  $b = (b_1, b_2)$  from (10) and (11), are  $\frac{2}{3}$ ,  $\frac{1}{2}$ , and  $\frac{1}{2}$ , respectively. Their product,  $1/6$ , is the probability attached to terminal node  $z_3$ . The entire probability distribution over terminal nodes is calculated in the same way. The resulting probabilities are displayed

in brackets next to each terminal node in Figure 9. Note that, as required, the sum of these terminal-node probabilities is one.

Let's formalize this computation a little. Consider any terminal node  $z$ . The path from the initial node to  $z$  encounters a set of nonterminal nodes, which we denote by  $\hat{X}(z)$ . Each nonterminal node  $x \in X$  along this path belongs to some player  $\bar{I}(x) \in I$ , and to some information set  $h_x \equiv \hat{h}(x) \in H_{\bar{I}(x)}$ . At each node  $x \in \hat{X}(z)$ , a particular action  $a \in A(h_x)$  is required to keep the play on the path to  $z$ ; call this required action  $\hat{a}(x, z) \in A(h_x)$ . The probability that player  $\bar{I}(x)$  chooses the action  $\hat{a}(x, z)$ , conditional on reaching  $x$ 's information set  $h_x$ , is  $b_{\bar{I}(x)}(\hat{a}(x, z) | h_x)$ . The probability associated with the terminal node  $z$  when the behavior-strategy profile  $b = (b_1, \dots, b_n)$  is played, then, is

$$\mathbb{P}(z | b) = \prod_{x \in \hat{X}(z)} b_{\bar{I}(x)}(\hat{a}(x, z) | \hat{h}(x)). \quad (14)$$

## Mixed strategies

Implicit within our discussion of extensive-form games has been another notion of randomized strategies. When we transformed an extensive-form game into strategic form, we identified each extensive-form pure strategy with a strategic-form pure strategy. Randomized strategies in the strategic form are *mixed* strategies: probability distributions over pure strategies. Therefore mixed strategies—in the strategic-form sense—are an alternative to behavior strategies as a vehicle to introduce randomization into extensive-form play.

At first glance, behavior and mixed strategies seem distinctly different. When a player implements a mixed strategy, she spins the roulette wheel a single time; the outcome of this spin determines which pure strategy (set of deterministic choices at each information set) she will play. When she implements a behavior strategy, she independently spins the roulette wheel every time she reaches a new information set. The space of player- $i$  behavior strategies  $B_i$  is given in (13). The space of player- $i$  mixed strategies is given by  $\Sigma_i = \Delta(S_i)$ , which we can rewrite in terms of the actions available at each of her information sets using (7) and then contrast to  $B_i$ :

$$B_i = \prod_{h \in H_i} \Delta(A(h)), \quad \Sigma_i = \Delta \left( \prod_{h \in H_i} A(h) \right). \quad (15)$$

We will see that mixed strategies permit correlations across information sets which behavior strategies cannot accommodate (because they involve only independent randomizations). So it's not at all obvious that a restriction of attention to behavior strategies would not involve a loss of generality. However, we will also see that—in games of perfect recall—the two types of randomized strategies *can* be used interchangeably in the sense that any probability distribution over outcomes which can be achieved by a mixed strategy can be achieved by a behavior strategy and vice versa. We will find as this course progresses that mixed strategies in the strategic form will provide a more useful perspective in some contexts and behavior strategies in the extensive form will be more useful in others.

First we define what it means for two strategies to be equivalent. We will see how, in games of perfect recall, a mixed strategy generates an “essentially unique” equivalent behavior strategy. Then we will see how, given a behavior strategy, we can find an equivalent mixed strategy. In fact we will see that the same behavior strategy can be generated by many different mixed strategies. We will summarize with a theorem which guarantees that, for games of perfect recall, any probability distribution over outcomes which can be achieved by a behavior strategy can be achieved by a mixed strategy and vice versa.

Consider a strategy profile  $\gamma = (\gamma_1, \dots, \gamma_n)$ , where, for each player  $i \in I$ ,  $\gamma_i$  may be either a pure, mixed, or behavior strategy. The strategy profile  $\gamma$  implies a probability distribution over terminal nodes. Consider two strategies for player  $i$ , e.g.  $\gamma_i'$  and  $\gamma_i''$ ,  $\gamma_i', \gamma_i'' \in (B_i \cup \Sigma_i)$ —perhaps one mixed and one behavior. We say that  $\gamma_i'$  and  $\gamma_i''$  are *equivalent* if, no matter what strategies  $\gamma_{-i} \in \prod_{j \in I \setminus \{i\}} (B_j \cup \Sigma_j)$  the opponents choose, the strategy profiles  $(\gamma_i', \gamma_{-i})$  and  $(\gamma_i'', \gamma_{-i})$  both induce the same probability distribution over terminal nodes.

### Mixed strategy $\rightarrow$ behavior strategy: the role of perfect recall

A behavior strategy  $b_i \in B_i$  for player  $i$  is a specification, for every player- $i$  information set  $h \in H_i$  and every action  $a \in A(h)$  feasible at  $h$ , of a conditional probability for that action  $a$  at the information set  $h$ . To say that the behavior strategy  $b_i \in B_i$  and the mixed strategy  $\sigma_i \in \Sigma_i$  are equivalent would require that, if player  $i$  has chosen the mixed strategy  $\sigma_i$ , then for every  $h \in H_i$  and every  $a \in A(h)$ ,  $b_i(a|h)$  is the probability that player  $i$  would choose action  $a$  whenever she reaches information set  $h$ .

Note that nothing in the above paragraph refers to what the other players are doing. Equivalence between a mixed and behavior strategy, then, requires that the conditional probability distribution over actions at any information set implied by a particular mixed strategy be independent of the strategy choices of the opponents. Now I'll show you an extensive-form game in which such an equivalence cannot possibly hold in general because the conditional probability distribution associated with a particular mixed strategy depends crucially on the strategy of the opponent.

Consider the two-player game in Figure 10. Player 1's strategy space is  $S_1 = \{(l, a), (l, b), (r, a), (r, b)\}$ . Let  $h$  be player 1's two-node information set. Consider the player-1 mixed strategy  $\sigma_1 = \frac{1}{2} \circ (l, a) \oplus \frac{1}{2} \circ (r, b)$ . If player 2 chooses  $D$ , one-half the time player 1 will be playing the pure strategy  $(l, a)$  and the left-hand node of information set  $h$  will be reached, at which player 1 will then choose  $a$ . The other half of the time player 1 is playing the pure strategy  $(r, b)$ , in which case the right-hand node of  $h$  is reached, at which player 1 chooses  $b$ . Therefore if player 2 chooses  $D$ , the probability of player 1 choosing  $a$  conditional on reaching  $h$  is one-half; i.e.  $\mathbb{P}(a|h) = \mathbb{P}(b|h) = \frac{1}{2}$ .

However, if player 2 chooses  $U$ , information set  $h$  is reached only half of the time, viz. when player 1's randomization results in the pure strategy  $(r, b)$ . So, if player 2 chooses  $U$  then, *conditional on information set  $h$  being reached*, player 1 will always choose  $b$  at  $h$ ; i.e.  $\mathbb{P}(a|h) = 0$  and  $\mathbb{P}(b|h) = 1$ . Therefore the conditional probabilities at information set  $h$  are *not* independent of player 2's strategy.




Figure 10: The conditional probability distribution over actions at an information set implied by a mixed strategy is not necessarily independent of the opponent's strategy.

Player 1's mixed strategy allowed her to direct probability from information set  $h$  in different directions depending on whether player 2's strategy was relevant. (When player 1 chose  $r$ , player 2's strategy was irrelevant, and player 1 chose  $b$  at information set  $h$ . When player 1 chose  $l$ , player 2's strategy was relevant, and player 1 chose  $a$  at information set  $h$ .) As a result, the direction in which probability flows out of information set  $h$ , and thus the conditional probability distribution over actions at  $h$ , depends on the action of player 2.

Even when the probability distribution over actions implied by a mixed strategy is independent of the opponents' strategies, it is possible that the behavior strategy which implements those conditional probabilities at each information set will generate a terminal-node distribution which differs from the distribution produced by the mixed strategy; i.e. the behavior and mixed strategies would not be equivalent.

To see this consider the one-player game in Figure 11.<sup>10</sup> Consider the mixed strategy  $\sigma_1 = \frac{1}{2} \circ (U, l) \oplus \frac{1}{2} \circ (D, r)$ . This induces the probability distribution over terminal nodes  $\frac{1}{2} \circ z_1 \oplus \frac{1}{2} \circ z_4$ . The information set  $\alpha$  is always reached; in one-half the cases player 1 chooses  $U$ ; in the others she chooses  $D$ . Therefore the conditional probability distribution over actions at information set  $\alpha$  is given by  $\mathbb{P}(U|\alpha) = \mathbb{P}(D|\alpha) = \frac{1}{2}$ . The information set  $\beta$  is always reached; in one-half the cases player 1 chooses  $l$ ; in the others she chooses  $r$ . Therefore the conditional probability distribution over actions at information set  $\beta$  is given by  $\mathbb{P}(l|\beta) = \mathbb{P}(r|\beta) = \frac{1}{2}$ . The behavior strategy  $b_1$  which implements these conditional probabilities is defined by  $b_1(U|\alpha) = b_1(D|\alpha) = b_1(l|\beta) = b_1(r|\beta) = \frac{1}{2}$ . However, this behavior strategy results in the probability distribution over terminal nodes in which each node is reached  $\frac{1}{4}$  of the time. In fact, no behavior strategy can duplicate the distribution achieved by the mixed strategy. Any such behavior strategy would have to play each of  $U$  and  $D$  and each of  $l$  and  $r$  with nonzero probability at  $\alpha$  and  $\beta$ , respectively. However, since the behavior strategy cannot distinguish between the two nodes in information set  $\beta$ , such a behavior strategy must put nonzero probability upon reaching each of  $z_2$  and  $z_3$ .

<sup>10</sup> Of course it seems unreasonable to call this a game, since there is only one player. One player suffices to make the point, and I want to keep the situation as simple as possible in order not to cloud the issue with needless complexities.




Figure 11: No behavior strategy can replicate the terminal-node distribution of the mixed strategy  $\sigma_1 = \frac{1}{2} \circ (U, l) \oplus \frac{1}{2} \circ (D, r)$ .

In the game of Figure 11 player 1's mixed strategy allows her to create a different conditional probability distribution over actions at each of the nodes in information set  $\beta$ : At the left-hand node she chooses  $l$  with probability one; at the right-hand node she chooses  $l$  with probability zero. A behavior strategy on the other hand imposes the same conditional probability distribution over all nodes within the same information set. It is not surprising then that the added flexibility of mixed strategies can generate probability distributions over terminal nodes which cannot be duplicated by behavior strategies.

The games in Figure 10 and Figure 11 were troublesome for finding a behavior strategy which was equivalent to some mixed strategy. A common feature of both games is that they do not satisfy perfect recall. (In both cases, at player 1's second information set, she has forgotten the action she took at the initial node.) A mixed strategy specifies a probability distribution over pure strategies, and each pure strategy can stipulate a particular pair of actions at any pair of information sets. When two player- $i$  nodes within the same information set are distinguished by an earlier action of player  $i$ 's, a mixed strategy can effectively dictate a different action at each node, even though those nodes are in the same information set. In games of perfect recall this is not possible: If one player- $i$  information set follows another, then every node of the later information set must be reached by the same action at the earlier information set.

Requiring that our games satisfy perfect recall certainly squashes the counterexamples of Figure 10 and Figure 11. It turns out that requiring perfect recall is sufficient to squash all other counterexamples as well; it is sufficient to guarantee that mixed and behavior strategies are equivalent in the sense we have defined.

Given that mixed strategies can correlate behavior across information sets, why should behavior strategies be able to duplicate every terminal-node distribution which can be generated by mixed strategies? Consider two information sets  $h$  and  $h'$  for a player. There are only two classes of cases of interest: either 1  $h$  and  $h'$  cannot be reached along the same path [i.e. every pair of nodes  $(x, x') \in h \times h'$  is unordered by precedence] or 2 there is a path which encounters  $h$  and later encounters  $h'$  (or vice versa). If there is no path which reaches both  $h$  and  $h'$ , a behavior strategy's inability to correlate its randomizations at these two information sets does not limit the distributions over terminal nodes it can achieve, because at most one of these information sets will be reached in any single play of the game,

and the randomization at the unreached information set has no influence on the outcome. If there is a path which encounters first  $h$ , say, and then  $h'$ , then a behavior strategy could effectively correlate its actions at the two information sets by conditioning its  $h'$  action on its knowledge that  $h$  was reached first. But this correlation would require that the player remember she had reached  $h$ , and this recollection is guaranteed when the game is one of perfect recall.

### Mixed strategy $\rightarrow$ behavior strategy: computing the conditional probabilities

Assume now, and for the rest of the semester (unless directed otherwise on an exam or problem set) that our game satisfies perfect recall. I will show heuristically how to take a given mixed strategy for a player and derive an equivalent behavior strategy; then I will formalize, generalize, and qualify the presentation. (For a more rigorous, nonheuristic demonstration, see the Appendix.) Consider the game in Figure 12. The players' strategy spaces are  $S_1 = \{(L, U), (L, D), (R, U), (R, D)\}$  and  $S_2 = \{g, s\}$ . Consider any mixed-strategy profile  $\sigma = (\sigma_1, \sigma_2)$ , where, for any player  $i \in \{1, 2\}$  and any player- $i$  pure strategy  $s_i \in S_i$ ,  $\sigma_i(s_i)$  is the probability with which player  $i$  chooses the pure strategy  $s_i$ . What are the corresponding behavior strategies  $b_1$  and  $b_2$ ?




Figure 12: Generating a behavior strategy from a mixed strategy.

We can think of probability as a flow that goes into a node and then is parceled out among its outgoing branches. Computing the conditional probability of a particular action at a node is the same as asking what fraction of the entering probability exits along that action's branch. We first seek  $b_1(L|\alpha)$ , the probability with which player 1 chooses  $L$  at the initial node  $\alpha$  conditional upon her reaching node  $\alpha$ . (Obviously the conditioning is ineffectual at this node, since player 1 always reaches it.) The only probability which exits node  $\alpha$  along the  $L$  branch is that associated with pure strategies which specify that player 1 choose  $L$  at  $\alpha$ , viz.  $(L, U)$  and  $(L, D)$ . Therefore

$$b_1(L|\alpha) = \frac{\sigma_1((L, U)) + \sigma_1((L, D))}{\sigma_1((L, U)) + \sigma_1((L, D)) + \sigma_1((R, U)) + \sigma_1((R, D))} = \sigma_1((L, U)) + \sigma_1((L, D)). \quad (16)$$

The denominator is the sum of the probabilities associated with the pure strategies entering  $\alpha$ . Because  $\alpha$  has no predecessors belonging to player 1, this includes all of player 1's pure strategies; hence the denominator sums to unity.

Similarly, player 2's conditional probability of choosing  $g$  at node  $\beta$  is

$$b_2(g|\beta) = \frac{\sigma_2(g)}{\sigma_2(g) + \sigma_2(s)} = \sigma_2(g). \quad (17)$$

Note that we ignored the extent to which player 1 might be choosing  $L$  or  $R$  at  $\alpha$  when calculating the denominator of (17), which I said was the probability entering  $\beta$ . So I'll now be more precise: the denominator is the probability which enters  $\beta$  over which player 2 has control.<sup>11</sup>

Calculating player 1's local randomization toward  $U$  at  $\delta$ , viz.  $b_1(U|\delta)$ , is more interesting. In this case we must take more care when identifying the probability which enters  $\delta$ . We use the notion of "a strategy which does not preclude a particular information set." For example, the pure strategies for player 1 ( $R, U$ ) and ( $R, D$ ) preclude play ever reaching  $\delta$ . ( $L, U$ ) and ( $L, D$ ) do not preclude information set  $\delta$ .<sup>12</sup> In this case the denominator is the sum of the mixing probabilities associated with the pure strategies which do not preclude  $\delta$  being reached. The numerator is the sum of the probabilities of the subset of these strategies which require player 1 to choose  $U$ , viz. ( $L, U$ ). Therefore

$$b_1(U|\delta) = \frac{\sigma_1((L, U))}{\sigma_1((L, U)) + \sigma_1((L, D))}. \quad (18)$$

Note a problem with the formula in (18): What if the mixed strategy  $\sigma_1$  we are considering puts no weight upon either ( $L, U$ ) or ( $L, D$ )? Then the denominator would vanish, leaving the conditional probability undefined. In this case our specification of  $b_1(U|\delta)$  and  $b_1(D|\delta)$  is arbitrary, subject to the constraint that  $b_1(\cdot|\delta)$  constitutes a probability distribution over the actions  $A(\delta)$  available at  $\delta$ . This requires only that  $b_1(U|\delta) + b_1(D|\delta) = 1$ . When the mixed strategy we are considering precludes an information set being reached, it is conventional that we determine the local behavior there by summing the probabilities associated with all the pure strategies which specify the desired action at the information set in question (whether any one of them precludes the information set or not).<sup>13</sup> In this example, the pure strategies ( $L, U$ ) and ( $R, U$ ) both specify  $U$  at information set  $\delta$ ; therefore, if  $\sigma_1((L, U)) = \sigma_1((L, D)) = 0$ , we would define

$$b_1(U|\delta) = \sigma_1((L, U)) + \sigma_1((R, U)) = \sigma_1((R, U)). \quad (19)$$

Now that we have seen in a simple example how a mixed strategy  $\sigma_i$  generates a behavior strategy  $b_i$ , we can formalize the process. For every information set  $h \in H_i$  which belongs to player  $i$ , let  $\hat{S}_i(h) \in S_i$  be the subset of player  $i$ 's pure-strategy space  $S_i$  whose strategies do not preclude  $h$  from being reached. In

<sup>11</sup> We can ignore the extent to which player 1 played  $L$  at  $\alpha$  because that multiplicative effect on the incoming probability in the numerator is exactly the effect on the outgoing probability, and therefore the two effects divide out. This is shown explicitly in the Appendix.

<sup>12</sup> Choosing one of these strategies does not guarantee that play will reach  $\delta$ , because player 2 could choose  $s$ . We are only interested in partitioning player 1's pure-strategy space into 1 strategies which by themselves (i.e. regardless of her opponent's play) rule out a particular information set being reached and 2 the remainder, for each of which there exists a strategy by her opponent such that the resulting strategy profile reaches the information set in question.

<sup>13</sup> This is a simple way to guarantee that  $b_i(\cdot|h)$  is a probability distribution over the actions at  $h$ .

other words, if  $s_i \in \hat{S}_i(h)$ , then there exists a deleted pure-strategy profile  $s_{-i} \in S_{-i}$  such that the strategy profile  $(s_i, s_{-i})$  reaches information set  $h$ . (I.e.  $\hat{X}(\hat{z}(s_i, s_{-i})) \cap h \neq \emptyset$ .)

Consider now some player  $i \in I$  and one of her mixed strategies  $\sigma_i \in \Sigma_i$ . For any information set  $h \in H_i$  we want to compute a conditional probability distribution  $b_i(\cdot | h)$  over the actions  $A(h)$  available at  $h$ . There are two cases: either 1 the mixed strategy  $\sigma_i$  is *compatible* with reaching  $h$ —it puts positive weight on at least one of the pure strategies which do not preclude  $h$ ; i.e.  $\exists s_i \in \hat{S}_i(h)$  such that  $\sigma_i(s_i) > 0$  or 2  $\sigma_i$  puts no weight on these strategies; i.e.  $\text{supp } \sigma_i \cap \hat{S}_i(h) = \emptyset$ . When the mixed strategy is compatible with reaching  $h$ , the conditional probability of choosing some action  $a \in A(h)$  is the ratio of the probability leaving along the branch  $a$  to the incoming probability to  $h$  over which  $i$  has control; otherwise we make the aforementioned conventional, arbitrary assignment when  $h$  is never reached by  $\sigma_i$ . In other words,

$$b_i(a | h) = \begin{cases} \frac{\sum_{\substack{s_i \in \hat{S}_i(h) \\ \text{s.t. } s_i(h)=a}} \sigma_i(s_i)}{\sum_{s_i \in \hat{S}_i(h)} \sigma_i(s_i)}, & \text{supp } \sigma_i \cap \hat{S}_i(h) \neq \emptyset, \\ \sum_{\substack{s_i \in S_i \\ \text{s.t. } s_i(h)=a}} \sigma_i(s_i), & \text{supp } \sigma_i \cap \hat{S}_i(h) = \emptyset. \end{cases} \quad (20)$$

You can verify that  $\sum_{a \in A(h)} b_i(a | h) \equiv 1$  for every  $i \in I$ ,  $h \in H$ , and  $\sigma_i \in \Sigma_i$ .<sup>14</sup>

We say that the mixed strategy  $\sigma_i \in \Sigma_i$  determines an “essentially unique” behavior strategy  $b_i \in B_i$  because the determination is unique at information sets which are compatible with  $\sigma_i$ . Although there is a high degree of arbitrariness in the definition at other information sets, the assignments at those are inconsequential because they do not affect the probability distribution over terminal nodes.

## Behavior strategy $\rightarrow$ mixed strategy

Now we take a given behavior strategy  $b_i$  for player  $i$  and construct an equivalent mixed strategy  $\sigma_i$ . The mixed strategy we construct will be but one of many equivalent mixed strategies. We’ll look for a mixed strategy  $\sigma_i$  such that 1 at any information set  $h$ , the player’s choice of action from  $A(h)$  has the marginal probability distribution specified by the behavior strategy, viz.  $b_i(\cdot | h)$  and 2 the player’s choice at any information set is made independently of her choice at any other information set.

Remember that a pure strategy  $s_i \in S_i$  in the extensive form is a specification at every information set

<sup>14</sup> Implicit in the lower branch of the definition of  $b_i(a | h)$ , viz. the  $\text{supp } \sigma_i \cap \hat{S}_i(h) = \emptyset$  case, is a denominator which sums  $\sigma_i(s_i)$  over all pure strategies  $s_i \in S_i$ . This sum is equal to one, so it is omitted.

$h \in H_i$  of a feasible action  $s_i(h) \in A(h)$ . What is the probability that player  $i$ 's randomizations according to the behavior strategy  $b_i$  result in a realization of  $s_i$ ? Such a realization would require that at every player- $i$  information set  $h \in H_i$  the realized action was that specified by the pure strategy  $s_i$ , viz.  $s_i(h)$ . The probability that the action  $s_i(h)$  is the realized action at  $h$  is  $b_i(s_i(h) | h)$ . Because the randomizations are independent across information sets, the probability that at every information set  $h$  the realized action is  $s_i(h)$  is the product of the probabilities  $b_i(s_i(h) | h)$ . In other words symbols, the probability that  $s_i$  is realized by the behavior strategy  $b_i$  is

$$\mathbb{P}(s_i | b_i) = \prod_{h \in H_i} b_i(s_i(h) | h).$$

A mixed strategy  $\sigma_i \in \Sigma_i$  would be equivalent to the behavior strategy  $b_i$  if it specified that each player- $i$  pure strategy  $s_i \in S_i$  was chosen with the same probability with which the pure strategy  $s_i$  would be realized by the behavior strategy  $b_i$ , i.e. if, for all  $s_i \in S_i$ ,  $\sigma_i(s_i) = \mathbb{P}(s_i | b_i)$  and therefore if

$$\sigma_i(s_i) = \prod_{h \in H_i} b_i(s_i(h) | h). \quad (21)$$

To verify that  $\sigma_i$  as defined in (21) really is equivalent to the behavior strategy  $b_i$  you can substitute (21) into (20) to obtain an identity in  $b_i(a | h)$ .<sup>15</sup>

## Many mixed strategies give rise to the same distribution over outcomes

Consider the game in Figure 13. Consider two mixed strategies for player 2:

$$\sigma_2 = \frac{1}{4} \circ (L, U) \oplus \frac{1}{4} \circ (L, D) \oplus \frac{1}{4} \circ (R, U) \oplus \frac{1}{4} \circ (R, D), \quad (22)$$

$$\sigma_2' = \frac{1}{2} \circ (L, U) \oplus \frac{1}{2} \circ (R, D). \quad (23)$$

You can easily verify that both mixed strategies give exactly the same behavior strategy— $b_1(L | \alpha) = b_1(U | \delta) = \frac{1}{2}$ , as computed by (16) and (18).

What is the origin of this degeneracy? Consider two information sets  $h$  and  $h'$  for player  $i$  on distinct paths (i.e. so that no node of  $h$  precedes or succeeds a node of  $h'$ ). A mixed strategy allows player  $i$  to correlate his randomizations at the two information sets. For example, one-half the time she may choose the pure strategy that plays  $a$  at  $h$  and  $a'$  at  $h'$ , and the other half of the time she plays the pure strategy which chooses  $b$  and  $b'$  at  $h$  and  $h'$ , respectively. Never will she play, for example, a pure strategy which prescribes  $a$  at  $h$  and  $b'$  at  $h'$ . However, since  $h$  and  $h'$  are on distinct paths, player  $i$  will never reach both information sets in a single play of the game. How she correlates her randomizations across these two information sets is irrelevant to the distribution over terminal nodes. What's important for behavior strategies is only the marginal distribution over actions at each information set. Many different mixed strategies can result in identical marginal distributions.

<sup>15</sup> You might also find it challenging to show directly that the sum of  $\sigma_i(s_i)$ , as defined in (17), over all  $s_i \in S_i$  is unity.




Figure 13

For example, the matrix in Figure 14a represents mixed strategies by player 2 in the game of Figure 13. The probability in each cell is the value of  $\sigma_2(s_i)$  for one of the four pure strategies  $(L, U)$ ,  $(L, D)$ ,  $(R, U)$ , and  $(R, D)$ . The sum of each row is the marginal probability of that action at node  $\alpha$ . The sum of each column is the marginal probability of that action being chosen at  $\beta$ . The marginal probabilities  $p_L$ ,  $p_R$ ,  $p_U$ , and  $p_D$  define a behavior strategy. In Figure 14b the probability of  $(L, U)$  is set to  $t$  and the entries in the remaining cells are constructed so as to satisfy the given marginal probabilities. We see that the parameter  $t$  is free to vary somewhat while still satisfying the marginal probabilities: there are many ways to construct mixed strategies which yield these marginal probabilities. For example, let all four marginal probabilities equal one-half. Then any mixed strategy of the form

$$\sigma_i = t \circ (L, U) \oplus (\frac{1}{2} - t) \circ (L, D) \oplus (\frac{1}{2} - t) \circ (R, U) \oplus t \circ (R, D), \tag{24}$$

where  $t \in [0, \frac{1}{2}]$ , will generate these marginal probabilities. The choices for  $\sigma_2$  and  $\sigma_2'$  in (22) and (23) were but two examples from this family.

	$U$	$D$	
$L$	$\sigma_2((L, U))$	$\sigma_2((L, D))$	$p_L$
$R$	$\sigma_2((R, U))$	$\sigma_2((R, D))$	$1 - p_L$
	$p_U$	$1 - p_U$	$1$

(a)

	$U$	$D$	
$L$	$t$	$p_L - t$	$p_L$
$R$	$p_U - t$	$1 - p_L - p_U + t$	$1 - p_L$
	$p_U$	$1 - p_U$	$1$

(b)

Figure 14: (a) Representing player 2's mixed strategy with one information set's action on each dimension. (b) The one-parameter family of mixed strategies corresponding to particular marginal probabilities.

### Kuhn's theorem: The equivalence of behavior and mixed strategies

We defined two strategies to be equivalent if, for any set of strategies by the other players, they result in the same distribution over terminal nodes. We just saw that, given a behavior strategy, we could construct many mixed strategies which are equivalent. Before that, for a given mixed strategy, we constructed an equivalent behavior strategy; however, we were unable to do so when the game was not one of perfect recall.

This exposition was merely suggestive and does not constitute a proof of anything. Kuhn [1953] proves the following theorem: Every behavior strategy is generated by some mixed strategy and every mixed strategy generates an essentially unique behavior strategy. If the game has perfect recall, then

every behavior strategy is equivalent to any mixed strategy which generates it.<sup>16,17</sup>

A consequence of this theorem is that we can find the Nash equilibria of a game in whichever form—extensive or strategic—is more convenient. Consider a conjectured Nash-equilibrium mixed-strategy profile  $\sigma^*$ . If  $\sigma_i^*$  is not a best response to  $\sigma_{-i}^*$ , then there exists a  $\bar{\sigma}_i$  such that  $(\bar{\sigma}_i, \sigma_{-i}^*)$  yields player  $i$  a better distribution over terminal nodes (and hence a higher expected payoff). Now consider the behavior-strategy profile  $b^*$  such that, for all  $j \in I$ ,  $\sigma_j^*$  and  $b_j^*$  are equivalent strategies. There exists a behavior strategy  $\bar{b}_i$  which is equivalent to  $\bar{\sigma}_i$ . Therefore  $(\bar{b}_i, b_{-i}^*)$  yields player  $i$  a higher expected payoff than  $b^*$ . So we see that a mixed-strategy profile is a Nash equilibrium if and only if an equivalent behavior-strategy profile is.

## The restriction of a strategy to a subgame

Often when we analyze subgames we want to know what a strategy profile in the original game implies about play in the subgame. Let's discuss the subgame  $\Gamma'$ , of the game  $\Gamma$  from Figure 1, which is indicated by the shaded box in Figure 15. The strategy profile  $s$  of  $\Gamma$  from (6) is *not* a strategy profile for this subgame  $\Gamma'$ , because it specifies actions at information sets which don't even exist in the subgame. To make sense of what it would mean to discuss the strategy profile  $s$  with regard to this subgame we first *restrict*  $s$ , creating a new strategy profile  $s'$ , by throwing away all those actions which correspond to information sets that don't belong to the subgame. (We also say that  $s'$  is the *restriction of  $s$  to the subgame*.) In other words,  $s'$  is of the form

$$s' = (s_1(\delta_1), s_1(\delta_2); s_2(\beta_1)),^{18} \quad (25)$$

and corresponding to our particular example in (6) we write

$$s' = (D, D; r). \quad (26)$$

More generally.... Recall that, when we decompose a game in order to form a subgame, we create a new information partition  $\tilde{H} \subset H$  by restriction of the original partition  $H$  to the residual set of nodes  $\tilde{V}$ ; in other words  $\tilde{H}$  contains those information sets of the original game which are also in the subgame.<sup>19</sup>  $\tilde{H}_i \subset \tilde{H}$  is the set of information sets in the subgame which belong to player  $i \in I$ .

<sup>16</sup> See also Myerson [1991: 202–204] for another proof.

<sup>17</sup> Kuhn uses a slightly weaker definition of perfect recall than we (and others) have used—and which is actually a misnomer because it tolerates some amnesia—and proves that his flavor of perfect recall is necessary and sufficient, rather than just sufficient as we state here, for the equivalence of behavior and mixed strategies. See Kreps [1990: 374–375] for a discussion.

<sup>18</sup> This may seem an odd order of actions, because  $\beta_1$  comes before  $\delta_1$  and  $\delta_2$  in the game tree. However, this is consistent with specifying player 1's strategy followed by 2's strategy.

<sup>19</sup> See “Extensive-Form Games.”




Figure 15. Strategy profile restricted to a subgame.

A pure strategy  $s_i$  for player  $i$  is a map from player- $i$  information sets to feasible actions. Given any player- $i$  pure strategy  $s_i: H_i \rightarrow A_i$  in the original game, we form  $\tilde{s}_i$ , the restriction of  $s_i$  to the subgame, by restricting the domain of  $s_i$  to the new set of player- $i$  information sets  $\tilde{H}_i$ ; i.e.  $\tilde{s}_i: \tilde{H}_i \rightarrow \bar{A}_i$  and  $\forall h \in \tilde{H}_i, \tilde{s}_i(h) = s_i(h)$ . Player  $i$ 's pure-strategy space in the subgame then is

$$\tilde{S}_i = \prod_{h \in \tilde{H}_i} A(h). \tag{27}$$

[Compare this to (7).]

A behavior strategy for player  $i$  is a map from player- $i$  information sets to probability distributions over feasible actions; i.e.  $\forall h \in H_i, b_i(\cdot | h) \in \Delta(A(h))$ . Given any player- $i$  behavior strategy  $b_i$ , we form its restriction to the subgame  $\tilde{b}_i$  by restricting the domain of player- $i$  information sets over which it is defined to the set  $\tilde{H}_i$  of player- $i$  information sets in the subgame:  $\forall h \in \tilde{H}_i, \tilde{b}_i(\cdot | h) = b_i(\cdot | h) \in \Delta(A(h))$ . Player  $i$ 's space of subgame behavior strategies then is

$$\tilde{B}_i = \prod_{h \in \tilde{H}_i} \Delta(A(h)). \tag{28}$$

[Compare this to (13).]

Return to the example behavior profile  $b = (b_1, b_2)$  defined in (10) and (11) for the game of Figure 15 (as depicted in Figure 9). Similar to what we did in the pure-strategy case, to restrict the behavior strategy  $b$  to the subgame  $\Gamma'$  we simply throw out those probability distributions over actions which are

defined at information sets of the original game which are not present in the subgame. The restricted strategies are:

$$\tilde{b}_1 = (D, \frac{1}{2} \circ U \oplus \frac{1}{2} \circ D), \quad (29)$$

$$\tilde{b}_2 = (\frac{1}{2} \circ l \oplus \frac{1}{2} \circ r). \quad (30)$$

Note, when comparing (27) to (7) and comparing (28) to (13), that the restricted strategy spaces  $\tilde{S}_i$  and  $\tilde{B}_i$  are simply projections of the original spaces  $S_i$  and  $B_i$ , respectively, onto the restricted set of player- $i$  information sets  $\tilde{H}_i$ .

Now I'll explain how to restrict a mixed strategy to a subgame—NOT! This is why we use behavior strategies in the first place. It's hard to make sense of what it would mean to restrict a mixed strategy to a subgame.

# Appendix: Mixed strategy → Behavior strategy

Consider player  $i$ 's mixed strategy  $\sigma_i \in \Sigma_i$  in an extensive-form game with perfect recall. Our task is to compute the conditional probability  $b_i(a|h)$  with which player  $i$  will choose action  $a$  at information set  $h$  for every player- $i$  information set  $h \in H_i$  at which this is defined and action  $a \in A(h)$  which is feasible at  $h$ . (Implicit in the statement of the problem is the belief that these conditional probabilities will be independent of the deleted strategy profile  $\sigma_{-i} \in \Sigma_{-i}$  played by player  $i$ 's opponents.) If any behavioral strategy is equivalent to the mixed strategy  $\sigma_i$ , it must be consistent with the  $b_i(a|h)$  so defined.<sup>20</sup> In fact, Kuhn's [1953] theorem guarantees that this behavior strategy is indeed equivalent to the mixed strategy  $\sigma_i$ .

For any node  $v \in V$ , let  $\hat{X}(v) \subset X$  be the set of decision nodes on the unique path from the initial node to  $v$ .<sup>21</sup> Let  $\hat{H}(v) \subset H$  be the set of information sets encountered on the path  $\hat{X}(v)$ ; i.e.  $\hat{H}(v) \equiv \hat{h}(\hat{X}(v)) = \bigcup_{x \in \hat{X}(v)} \hat{h}(x)$ . For any pure-strategy profile  $s \in S$ , let  $\tilde{X}(s) \subset X$  be the set of decision nodes on the path of  $s$ ; i.e.  $\tilde{X}(s) \equiv \hat{X}(\hat{z}(s))$ .<sup>22</sup> Let  $\tilde{H}(s) \subset H$  be the set of information sets encountered on the path  $\tilde{X}(s)$ ; i.e.  $\tilde{H}(s) \equiv \hat{h}(\tilde{X}(s))$ .

Our event space is the space  $S$  of pure-strategy profiles. Let  $\hat{S}(x) \subset S$  be the event (i.e. set of strategy profiles such) that node  $x \in X$  is reached; i.e.

$$\hat{S}(x) = \{s \in S: x \in \tilde{X}(s)\}.^{23} \tag{A.1}$$

Let  $\hat{S}(h) \subset S$  be the event that information set  $h$  is reached, i.e.

$$\hat{S}(h) = \{s \in S: h \in \tilde{H}(s)\} = \{s \in S: \tilde{X}(s) \cap h \neq \emptyset\} = \bigsqcup_{x \in h} \hat{S}(x).^{24} \tag{A.2}$$

Because of perfect recall, all the nodes in any information set are unordered by precedence and therefore no strategy profile can generate a path which contains two distinct nodes of the same information set. I.e.  $\forall h \in H, \forall x, x' \in h$  such that  $x \neq x'$ ,

$$\hat{S}(x) \cap \hat{S}(x') = \emptyset.^{25} \tag{A.3}$$

For any decision node  $x \in X$ , let  $\hat{S}_i(x) \subset S_i$  be the projection of  $\hat{S}(x) \subset S$  onto player  $i$ 's strategy space

<sup>20</sup> A behavioral strategy  $b_i'$  which is equivalent to  $\sigma_i$  only needs to agree with  $b_i$  at information sets at which the conditional probability is defined.

<sup>21</sup> Recall that  $V$  is the set of all nodes and is partitioned into decision nodes  $X$  and terminal nodes  $Z$ .

<sup>22</sup> Recall that  $\hat{z}(s) \in Z$  is the terminal node reached by strategy profile  $s \in S$ .

<sup>23</sup> Therefore  $s \in \hat{S}(x) \Leftrightarrow x \in \tilde{X}(s)$ ; i.e. strategy profile  $s$  encounters node  $x$  if and only if node  $x$  is encountered by strategy profile  $s$ .

<sup>24</sup> Therefore  $\hat{S}(h)$  is just the image of  $h$  under  $\hat{S}$ , so there is no abuse of notation here.

<sup>25</sup> If to the contrary  $\exists s \in (\hat{S}(x) \cap \hat{S}(x'))$ , then  $x, x' \in \tilde{X}(s)$ ; i.e.  $x$  and  $x'$  would be on the same path and therefore—since they are distinct nodes—ordered by precedence.

$S_i$ ; i.e.

$$\hat{S}_i(x) = \{s_i \in S_i : \exists s_{-i} \in S_{-i}, (s_i, s_{-i}) \in \hat{S}(x)\} = \{s_i \in S_i : \exists s_{-i} \in S_{-i}, x \in \tilde{X}((s_i, s_{-i}))\}. \quad (\text{A.4})$$

These are the player- $i$  pure strategies which do not preclude node  $x \in X$ . For any decision node  $x \in X$  and any player- $i$  pure strategy  $s_i \in \hat{S}_i(x)$  which does not preclude  $x$ , let

$$\tilde{S}_{-i}(x, s_i) = \{s_{-i} \in S_{-i} : (s_i, s_{-i}) \in \hat{S}(x)\} = \{s_{-i} \in S_{-i} : x \in \tilde{X}((s_i, s_{-i}))\}. \quad (\text{A.5})$$

This is the set of deleted strategy profiles which can be combined with  $s_i \in \hat{S}_i(x)$  in order to reach  $x$ .

### Lemma 1

Consider any decision node  $x \in X$  which is encountered by a strategy profile  $s \in S$ . The strategy profile  $s' \in S$  encounters node  $x$  if and only if the two strategy profiles agree at every previously encountered information set. I.e.  $\forall x \in X, \forall s, s' \in S$ , such that  $x \in \tilde{X}(s)$ , we have  $x \in \tilde{X}(s')$  if and only if  $\forall h \in \hat{H}(x) \setminus \{\hat{h}(x)\}, s_{i(h)}(h) = s_{i(h)}'(h)$ .

### Proof

Omitted.

### Lemma 2

For all player- $i$  information sets  $h \in H_i$  and for all pairs  $x, x' \in h$  of nodes in this information set, the set of player- $i$  strategies which do not preclude the first node is exactly the set of strategies which do not preclude the second. I.e.  $\forall i \in I, \forall h \in H_i, \forall x, x' \in h, \hat{S}_i(x) = \hat{S}_i(x')$ .

### Proof

Assume to the contrary that, for some  $i \in I, h \in H_i$ , and some node pair  $x, x' \in h$ , there exists a player- $i$  pure strategy  $s_i \in \hat{S}_i(x)$  such that  $s_i \notin \hat{S}_i(x')$ . Therefore there exists a  $s_{-i} \in S_{-i}$  such that  $x \in \tilde{X}((s_i, s_{-i}))$ . Because  $\hat{S}_i(x') \neq \emptyset, \exists s_i' \in \hat{S}_i(x')$  such that  $s_i' \neq s_i$ . Consider the player- $i$  information sets encountered on the path to  $x$ , viz.  $(\hat{H}(x) \cap H_i) \setminus \{h\}$ . The two strategies  $s_i$  and  $s_i'$  cannot agree at all of these information sets because, if they did, the two strategy profiles  $(s_i, s_{-i})$  and  $(s_i', s_{-i})$  would agree at all the on-the-path information sets preceding  $x$ , viz.  $\hat{H}(x) \setminus \{h\}$ , and therefore  $(s_i', s_{-i})$  would also encounter  $x$  by Lemma 1, and this would violate perfect recall.

So  $s_i$  and  $s_i'$  must differ at some previous information set  $h' \in (\hat{H}(x) \cap H_i) \setminus \{h\}$ . But this would also violate perfect recall because  $x$  and  $x'$  share an information set,  $x$  ultimately succeeds  $h'$  via  $s_i(h')$ , and  $x'$  ultimately succeeds  $h'$  via a *different* action  $s_i'(h')$  at  $h'$ . ☺

We use the result of Lemma 2 to justify writing

$$\forall h \in H_i, \forall x \in h, \hat{S}_i(h) = \hat{S}_i(x). \quad (\text{A.6})$$

I.e. we can meaningfully talk about the player- $i$  pure strategies which do not preclude an information set  $h \in H_i$ .

**Lemma 3**

If neither of two player- $i$  strategies preclude a given player- $i$  node  $x \in X_i$ , then the two strategies agree at all on-the-path information sets prior to that node. I.e. if  $s_i, s_i' \in \hat{S}_i(x)$ , then  $\forall h \in [(H_i \cap \hat{H}(x)) \setminus \{\hat{h}(x)\}]$ ,  $s_i(h) = s_i'(h)$ .

**Proof**

From (A.4) we know that there exist  $s_{-i}, s_{-i}' \in S_{-i}$  such that  $x \in (\tilde{X}((s_i, s_{-i})) \cap \tilde{X}((s_i', s_{-i}')))$ . Therefore  $(s_i, s_{-i})$  and  $(s_i', s_{-i}')$  agree for all on-the-path information sets preceding  $x$  by Lemma 1, and therefore  $s_i$  and  $s_i'$  agree for all information sets  $h \in [(H_i \cap \hat{H}(x)) \setminus \{\hat{h}(x)\}]$ . ☺

**Lemma 4**

Let  $x \in X_i$  be a player- $i$  node. If neither of two player- $i$  strategies preclude node  $x$ , then the deleted strategy profiles which encounter  $x$  are the same for the two player- $i$  strategies. I.e.

$$\forall s_i, s_i' \in \hat{S}_i(x), \tilde{S}_{-i}(x, s_i) = \tilde{S}_{-i}(x, s_i').$$

**Proof**

I will show that  $s_{-i} \in \tilde{S}_{-i}(x, s_i) \Rightarrow s_{-i} \in \tilde{S}_{-i}(x, s_i')$ . We have  $x \in \tilde{X}((s_i, s_{-i}))$ . From Lemma 3,  $s_i$  and  $s_i'$  must agree at all player- $i$  on-the-path information sets encountered before  $x$ . Therefore  $(s_i, s_{-i})$  and  $(s_i', s_{-i})$  agree at all on-the-path information sets encountered before  $x$ . Therefore from Lemma 1  $(s_i', s_{-i})$  must encounter  $x$  and therefore  $s_{-i} \in \tilde{S}_{-i}(x, s_i')$ . ☺

Lemma 4 allows us to define for all  $x \in X$

$$\check{S}_{-i}(x) = \tilde{S}_{-i}(x, s_i), \tag{A.7}$$

where  $s_i \in \hat{S}_i(x)$  is any player- $i$  strategy which does not preclude node  $x$ .

For all information sets  $h \in H$  and actions  $a \in A(h)$ , let  $\dot{S}(a, h) \subset S$  be the event that  $s_{I(h)}(h) = a$ ; i.e.

$$\dot{S}(a, h) = \{s \in S: s_{I(h)}(h) = a\}. \tag{A.8}$$

This is the set of strategy profiles in which the owner of information set  $h$  chooses action  $a$  if play reaches  $h$ .

The conditional probability we seek is:

$$b_i(a|h) = \frac{\mathbb{P}(\dot{S}(a, h) \cap \hat{S}(h))}{\mathbb{P}(\hat{S}(h))}, \tag{A.9}$$

where we restrict attention to information sets  $h$  which are reached with positive probability by the mixed strategy  $\sigma_i$ . (I.e. we restrict attention to strategy/information-set pairs such that the mixed strategy is *compatible* with the information set.)

First we calculate the denominator. From (A.2) and (A.3), the probability that play reaches  $h$  is

$$\begin{aligned}
\mathbb{P}(\hat{S}(h)) &= \prod_{x \in h} \mathbb{P}(\hat{S}(x)) = \prod_{x \in h} \mathbb{P}\left(\prod_{s \in \mathcal{S}(x)} s\right) = \prod_{x \in h} \prod_{s \in \mathcal{S}(x)} \sigma(s) \\
&= \prod_{x \in h} \prod_{s_i \in \hat{S}_i(x)} \sigma_i(s_i) \prod_{s_{-i} \in \tilde{\mathcal{S}}_{-i}(x, s_i)} \prod_{j \in I \setminus \{i\}} \sigma_j(s_j).
\end{aligned} \tag{A.10}$$

From (A.6) and (A.7) we change the ranges of two summations according to

$$s_i \in \hat{S}_i(x) \quad \rightarrow \quad s_i \in \hat{S}_i(h), \tag{A.11}$$

$$s_{-i} \in \tilde{\mathcal{S}}_{-i}(x, s_i) \quad \rightarrow \quad s_{-i} \in \check{\mathcal{S}}_{-i}(x). \tag{A.12}$$

These changes allow a rearrangement and separation of the terms in (A.8), becoming

$$\mathbb{P}(\hat{S}(h)) = \left( \prod_{s_i \in \hat{S}_i(h)} \sigma_i(s_i) \right) \left( \prod_{x \in h} \prod_{s_{-i} \in \check{\mathcal{S}}_{-i}(x)} \prod_{j \in I \setminus \{i\}} \sigma_j(s_j) \right). \tag{A.13}$$

Now to calculate the numerator.... We transform the event in the numerator to become

$$\dot{S}(a|h) \cap \hat{S}(h) = \dot{S}(a|h) \cap \left( \prod_{x \in h} \hat{S}(x) \right) = \prod_{x \in h} (\dot{S}(a|h) \cap \hat{S}(x)). \tag{A.14}$$

Then

$$\mathbb{P}(\dot{S}(a|h) \cap \hat{S}(h)) = \prod_{x \in h} \prod_{s \in (\dot{S}(a|h) \cap \hat{S}(h))} \sigma(s) = \prod_{x \in h} \prod_{\substack{s \in \mathcal{S}(h) \\ \text{s.t. } s_i(h)=a}} \sigma(s). \tag{A.15}$$

A very similar chain of reasoning (with an added finesse or two) to what we used in the calculation of the denominator yields

$$\mathbb{P}(\dot{S}(a|h) \cap \hat{S}(h)) = \left( \prod_{\substack{s_i \in \hat{S}_i(h) \\ \text{s.t. } s_i(h)=a}} \sigma_i(s_i) \right) \left( \prod_{x \in h} \prod_{s_{-i} \in \check{\mathcal{S}}_{-i}(x)} \prod_{j \in I \setminus \{i\}} \sigma_j(s_j) \right). \tag{A.16}$$

Combining (A.14) and (A.11) yields

$$b_i(a|h) = \frac{\prod_{\substack{s_i \in \hat{S}_i(h) \\ \text{s.t. } s_i(h)=a}} \sigma_i(s_i)}{\prod_{s_i \in \hat{S}_i(h)} \sigma_i(s_i)}. \tag{A.17}$$

This is exactly the branch of (20) which corresponds to the compatible strategy case.

## References

Allen, Woody (director) [1973] “Sleeper.”

Fudenberg, Drew and Jean Tirole [1991] *Game Theory*, MIT Press.

Kreps, David M. [1990] *A Course in Microeconomic Theory*, Princeton University Press.

Kuhn, Harold W. [1953] “Extensive Games and the Problem of Information,” in *Contributions to the Theory of Games*, eds. Harold W. Kuhn and A. Tucker, Vol. 2, Princeton University Press, pp. 193–216.

Myerson, Roger B. [1991] *Game Theory: Analysis of Conflict*, Harvard University Press.