

Game Theory: Elements of Basic Models.

Branislav L. Slantchev

Department of Political Science, University of California - San Diego

April 23, 2009

Contents.

1	The Building Blocks	2
2	Formal Definition of the Extensive Form	6
2.1	Perfect Recall	7
2.2	Finite and Infinite Games	8
2.3	Informational Categories	9
2.4	Examples of Games in Extensive Form	9
3	Pure Strategies	13
4	The Strategic (Normal) Form	19
4.1	Examples of Converting Extensive to Strategic Form	22
4.1.1	Several Chance Moves	23
4.1.2	Three Players	25
4.2	Examples in Strategic Form	26
4.3	Reduced Strategic Form	29
5	Mixed Strategies in Strategic Form Games	32
6	Mixed and Behavior Strategies in Extensive Form Games	38
6.1	Mixed Strategy Equivalent to a Behavior Strategy	39
6.2	Equivalence of Mixed and Behavior Strategies	41
6.2.1	A Mixed Strategy Can Generate Many Behavior Strategies	44
6.2.2	Different Mixed Strategies Can Generate the Same Behavior Strategy	45

We now begin the study of noncooperative game theory, the analysis of interdependent decision-making. Before we can analyze any situation, we need to describe it formally. That is, we must have the specification of the model that describes the situation, or game, that we are interested in. There are two important ways in which to do that, the *extensive form* and the *strategic form*, sometimes also called the *normal form*. Of these, the extensive form is richer and the strategic form is usually conceptualized as being derived from an extensive form. However, the strategic form is simpler and usually more convenient for analysis.

In this lecture, we shall learn how to describe all kinds of situations that we might be interested in analyzing. We shall learn to distinguish between different classes of information, when information becomes available, and how. The goal is to get a solid grasp on model description before proceeding to the study of model solutions.

1 The Building Blocks

Any situation that we wish to represent formally would have some basic elements that will be part of its description. Most often, we begin with a verbal description (that may be quite vague at times), and then distill each element from it. Let's start with a simple card game borrowed from Roger Myerson.

EXAMPLE 1. (MYERSON'S CARD GAME.) There are two players, labeled "player 1" and "player 2."¹ At the beginning of this game, each player puts a dollar in a pot. Next, player 1 draws a card from a shuffled deck of cards in which half the cards are red and half are black. Player 1 looks at his card privately and decides whether to raise or fold. If player 1 folds, then he shows his card to player 2 and the game ends; player 1 takes the money in the pot if the card is red, but player 2 takes the money if the card is black. If player 1 raises, then he adds another dollar to the pot and player 2 must decide whether meet or pass. If she passes, the game ends and player 1 takes all the money in the pot. If she meets, she puts another dollar in the pot, and then player 1 shows his card to player 2 and the game ends; if the card is red, player 1 takes all the money in the pot, but if it is black, player 2 takes all the money.

The essential elements of a game are:

1. **players:** The individuals who make decisions.
2. **rules of the game:** Who moves when? What can they do?
3. **outcomes:** What do the various combinations of actions produce?
4. **payoffs:** What are the players' preferences over the outcomes?
5. **information:** What do players know when they make decisions?
6. **chance:** Probability distribution over chance events, if any.

A **player** is a decision-maker who is participant in the game and whose goal is to choose the actions that produce his most preferred outcomes or lotteries over outcomes. We assume that *players are rational*: their preference orderings are complete and transitive. We model uncertainty over outcomes with lotteries, like we've done before. This means that preferences

¹We establish the following convention: odd-numbered players are male, and even-numbered players are female. For a generic player, we shall always use the generic male pronoun.

can be described with utility functions and rational players choose actions that maximize their expected utilities (that’s why we need the vNM theorem).

Let $\mathcal{I} = \{1, 2, \dots\}$ denote the set of players indexed by i . That is, $i \in \mathcal{I}$ is a generic member of this set. In our example, $\mathcal{I} = \{1, 2\}$, the two players labeled “player 1” and “player 2.”

We represent **chance** events by a *random move of nature*. **Nature**, denoted by N , is a pseudo-player whose actions are purely mechanical and probabilistic; that is, they determine the probability distribution over the chance events. In our example, Nature “chooses” the color of the card that player 1 randomly draws from deck. Because the number of red cards equals the number of black cards and the deck is shuffled, the probability of the randomly chosen card being red is 0.5. Fig. 1 (p. 3) shows how the random draw by player 1 can be represented as a move by Nature.




Figure 1: Move by Nature Determines Card Color.

Nature “moves” first, and so the **initial node** (or the “root node”) of the game, denoted with an empty circle, is the place where the chance event occurs. The two possible “actions” by Nature are *red* and *black*, which we represent with one **branch** each.

Each branch then leads to a **decision node** (denoted with a filled circle), where player 1 gets to make his choice between raising and folding. When player 1 gets to move, he knows the color of the card he has drawn. In our example, player 1 chooses whether to raise or fold under two distinct circumstances, depending on the color of the card. That is, he has one decision to make conditional on the card being black, and another conditional on the card being red. In both cases, the choices are between raising and folding.

We need a way to represent the fact that when player 1 gets to move, he knows the color of the card he is holding. An **information set** for some player i summarizes what the player knows when get gets to move. Player 1 has two information sets, labeled “b” and “c”. At information set “b”, player 1 knows that the card is black, and at information set “c”, he knows that the card is red. Each of these information sets contains exactly one decision node.




Figure 2: Move by Nature Followed by Choice by Player 1.

For each of his information sets, a player must choose what to do. An **action** (or move) for player i is a choice, denoted by a_i that player i can make at that information set. Let $A_i = \{a_i\}$ denote the set of choices at an information set. That is, this is the set of actions

from which the player must choose. The set of actions may be different depending on the information set. Let h denote an arbitrary information set (we shall shortly see why this letter is appropriate). Then $A_i(h)$ is the set of actions available to player i at information set h . If the player does not get to move at information set h , then $A_i(h) = \emptyset$.

In our example, player 1 always has the same two actions regardless of the color of the card: He can either raise, denoted by R , or fold, denoted by F . Thus, $A_1(b) = \{F, R\}$ and $A_1(c) = \{f, r\}$. We represent the actions available at a decision node with branches emanating from that node, as shown in Fig. 2 (p. 3). I have used upper and lower case letters to denote the actions at the different information sets to emphasize that they are, in fact, different in the sense that although the action is the same it occurs in a different context. That is, even though F and f both represent the action “fold,” the first is really “fold on black card” and the second is “fold on red card.”

Information sets that contain only one decision node are called **singletons**. Here, both information sets for player 1 are singletons. Note that we have labeled the two information sets by player 1 with “1.b” and “1.c” respectively. This is intended to convey both that player 1 gets to move and that he knows different things at the different information sets.

A **history** of the game is a sequence of actions taken by the various players at their information sets. The initial history (before the game begins) is denoted by $h^0 = \emptyset$. One history of the game is (black), that is, nature having chosen black. Another history is (black, F), that is, nature having chosen black, and player 1 having folded.

More generally, we can think of the game as a sequence of stages, where all players simultaneously choose actions from their choice sets $A_i(h)$ (remember that these choices may be “do nothing” if the player’s action set is empty at h). An **action profile** is the set of actions taken by the players at that stage. For example, h^0 is the “history” at the beginning of the game, and $a^0 = (a_1^0, \dots, a_I^0)$ is the action profile following h^0 . Then h^1 is the history identified with a^0 , and $A_i(h^1)$ is the set of actions available to player i there. Continuing iteratively in this manner, we define the history at the end stage k to be the sequence of actions in the previous stages:

$$h^{k+1} = (a^0, a^1, \dots, a^k).$$

We shall let $K + 1$ denote the total number of stages in the game, noting that for some games, we may have $K = +\infty$. In these cases, the “outcome” of the game is the infinite history h^∞ . Let $H = \{h^k\}$ denote the set of all possible histories. Since each h^{K+1} by definition describes an entire sequence of actions from the beginning of the game to its end, we shall call it a **terminal history**. The set $\mathcal{Z} = \{h^{K+1}\} \subset H$ of all terminal histories is the same as the set of outcomes when the game is played.

Returning to our example, the history (red, f) is terminal because the game ends if player 1 folds. Conversely, the histories (red) and (red, r) are not terminal because the game continues. Note that information sets are related to histories because they summarize past play and what players know about it.

For each player i , we specify a **payoff function**, $u_i : \mathcal{Z} \rightarrow \mathbb{R}$. That is, a function that maps the set of terminal histories (or **outcomes**), to real numbers. In other words, we assign numeric payoffs to the outcomes. Of course, this function must represent the preference ordering of the player over the outcomes. Since $h_1 = (\text{black}, F)$ and $h_2 = (\text{red}, f)$ are both terminal histories, the player’s (Bernoulli) payoff functions must assign numbers to these outcomes.²

²Bernoulli defined the utility function over wealth, and by convention we use the term *Bernoulli function* to refer to payoff functions defined over the outcomes. Von Neumann and Morgenstern moved away from this

Let's assume that utilities are linear in the amount of money received, or $u(z) = z$. Then:

$$\begin{aligned} u_1(h_1) &= u_2(h_2) = -1 \\ u_1(h_2) &= u_2(h_1) = 1. \end{aligned}$$

We list these payoffs below the terminal node associated with them. By convention, the order is determined by the order in which players appear in the game tree, top to bottom and left to right. In our example in Fig. 2 (p. 3), the first number is player 1's payoff and the second number is player 2's payoff.

If player 1 raises, player 2 gets to make a move. Thus, the R and r branches representing raising by player 1 lead to decision nodes for player 2. She can either meet, m , or pass, p , and so each decision node will have two branches, labeled m and p respectively, as shown in Fig. 3 (p. 5). The payoffs from the resulting terminal histories are specified in the same manner as before.




Figure 3: Myerson's Card Game in Extensive Form.

The crucial difference between the information available to player 1 and the information available to player 2 is that player 2, unlike player 1, does not know the color of player 1's card although she does observe his action (raising). In other words, when player 2 gets to move, she does not know whether player 1's card is red or black. The information set, denoted by "0" for player 2 thus includes *both* histories $h_3 = (\text{black}, R)$ and $h_4 = (\text{red}, r)$. Because each of these histories leads to a different decision node for player 2, we enclose them in a box (or connect them in some other way) to demonstrate that they belong to the same information set. We say that both h_3 and h_4 are **consistent** with the information set "0". The information set represents the fact that when player 2 gets to move, she does not know the color of the card; she only knows what she can see—namely, that player 1 has chosen to raise.³

and defined the *expected utility function* over lotteries. People sometimes call these *Von Neumann-Morgenstern Utility Function* or, simply, *vNM Utility Function*. Recall that these are subjective in the sense that preferences must be given before these utilities can be derived.

³As we shall see when we analyze the game, in equilibrium player 2 may learn about the likelihood of the card's color by using the information obtained from observing raising and knowledge of player 1's optimal strategy. In some games, the uncertainty will be fully resolved—even though player 2 cannot observe what is known to the opponent, she can *infer* that information from his observable behavior and knowledge that he, being intelligent and rational, is choosing his optimal strategy. Of course, player 1 knows all of that full well, so he may well try to obfuscate her inferences, just as he will do in this particular game. His optimal strategy is to prevent this inference. Even then player 2 will be able to learn something from the fact that he's chosen to raise. Observe, incidentally, that unless you assume that players pursue the best strategies to the best of

Player 2's information set is not a singleton because it contains two of her decision-nodes. Let $h(x)$ denote the fact that the information set h contains node x . The information set captures the idea that the player who is choosing an action at h is uncertain whether he is at x or at some other $x' \in h(x)$. We require that if $x' \in h(x)$, then the same player moves at x and x' . Otherwise, players may disagree who was supposed to move.

Information sets partition the decision-nodes such that each node belongs to exactly one information set and no more. It is in this way that information sets are related to histories. As you can see in the example, it is perfectly fine to have information sets with more than one decision node. However, it is impossible for the same decision node to appear in more than one information set.

Recall that the action sets are defined in terms of information sets. That is, $A_i(h)$ is the set of actions from which player i may choose at information set h . It is essential to realize that this implies that for all nodes in this information set, the actions available at each are the same. That is, if $x' \in h(x)$, then $A_i(x') = A_i(x)$. Thus, we can let $A_i(h)$ denote the action set at information set h .

To see why this must be the case, suppose that player 2 had another option, say "punt", at the node reached by the history $h_3 = (\text{black}, R)$ that was not available after history $h_4 = (\text{red}, r)$. This means that she could punt if and only if player 1 had a black card. But how would player 2 exercise this option if she does not know the color of the card? To represent this situation, we would have to give player 2 an action called "try to punt" and add it to both nodes in her information set. Then, if she chooses this option, she would succeed when the card is black but fail when it is red.

Note, on the other hand, that we could easily give player 1 different actions (or numbers of actions) at each of his nodes 1.b and 1.c because they belong to different information sets. It is to emphasize this that I label the actions differently in Fig. 3 (p. 5), with lowercase and uppercase letters, depending on the color.

The point is that if a player has two nodes with different sets of actions, then these nodes cannot belong to the same information set. However, one can easily have different nodes with the same sets of actions even though the nodes are not in the same information set.

This completes the extensive form representation of the card game. Note that we have specified the players, the rules of the game (who moves when and what options they have), the outcomes in terms of terminal histories, the payoffs associated with these outcomes, the information available to the players when they move, and the probability distribution of the chance events.

2 Formal Definition of the Extensive Form

In most applications, the game trees would rarely be drawn, and so one must make do with the mathematical description of the extensive form. It is necessary to go through this exercise to understand the methodology of this fundamental class of games. We shall rarely, if ever, need to resort to the finer detail, but the mathematical description allows us to define two important categories of games (perfect and imperfect recall), of which we shall only study one. The following definition follows Fudenberg & Tirole (1991).

their abilities, you cannot make such inferences, and behavior becomes unintelligible. Among other things, this would imply that we simply cannot perform any sort of meaningful analysis as social scientists.

DEFINITION 1. The extensive form of a game, $\Gamma = \{\mathcal{I}, (X, \succ), \iota(\cdot), A(\cdot), H, u\}$, contains the following elements:

1. A set of players denoted by $i \in \mathcal{I}$, with $\mathcal{I} = \{N, 1, 2, \dots\}$, with N representing the pseudo-player Nature;
2. A tree, (X, \succ) , which is a finite collection of nodes $x \in X$ endowed with the precedence relation \succ , where $x \succ x'$ means “ x is before x' .” This relation is transitive and asymmetric, and thus constitutes a *partial order*.⁴ This rules out cycles where the game may go from node x to a node x' , from x' back to x .⁵ In addition, we require that each node x has exactly one immediate predecessor, that is, one node $x' \succ x$ such that $x'' \succ x$ and $x'' \neq x'$ implies $x' \succ x''$ or $x'' \succ x'$. Thus, if x' and x'' are both predecessors of x , then either x' is before x'' or x'' is before x' .
3. A set of terminal nodes, denoted by $z \in \mathcal{Z}$ consisting of all nodes that are not predecessors of any other node. Because each z determines the path through the tree, it represents an outcome of the game. The payoffs for outcomes are assigned by the Bernoulli payoff functions $u_i : \mathcal{Z} \rightarrow \mathbb{R}$, and $u = (u_1(\cdot), \dots, u_I(\cdot))$ is the collection of these functions, one for each player.
4. A map $\iota : X \rightarrow \mathcal{I}$, with the interpretation that player $\iota(x)$ moves at node x . A function $A(x)$ that denotes the set of feasible actions at x .
5. Information sets $h \in H$ that partition the nodes of the tree such that every node is exactly in one set. The interpretation of $h(x)$ is that information set h contains the node x . We require that if $x' \in h(x)$, then $A(x') = A(x)$, and so we can let $A(h)$ denote the set of feasible actions at information set h .
6. A probability distribution over the set of alternatives for all chance nodes.

This definition now allows us to make several ideas very precise.

2.1 Perfect Recall

We shall require that players have *perfect recall*. That is, a player never forgets information he once knew, and each player knows the actions he has chosen previously. (As we shall see, the fact that players may know all previous history does not force us to assume that he will take it all into account when making decisions.) This is accomplished by requiring that:

- A) if two decision nodes are in the same information set, then neither is a predecessor of the other; and
- B) if two nodes x' and x'' are in the same information set and one of them has a predecessor x , then the other one has a predecessor \hat{x} (possibly x itself) in the same information set as x and the action taken at x that leads to x' is the same as the action taken from \hat{x} that leads to the x'' .

The games in Fig. 4 (p. 8) illustrate some cases of imperfect recall that this requirement eliminates.

⁴It is not a complete order because two nodes may not be comparable. For example, consider player 2's information set in Fig. 3 (p. 5): Neither of the nodes precedes the other.

⁵To see this, suppose we constructed a game such that $x \succ x' \succ x$. By transitivity, $x \succ x$, but this violates asymmetry.

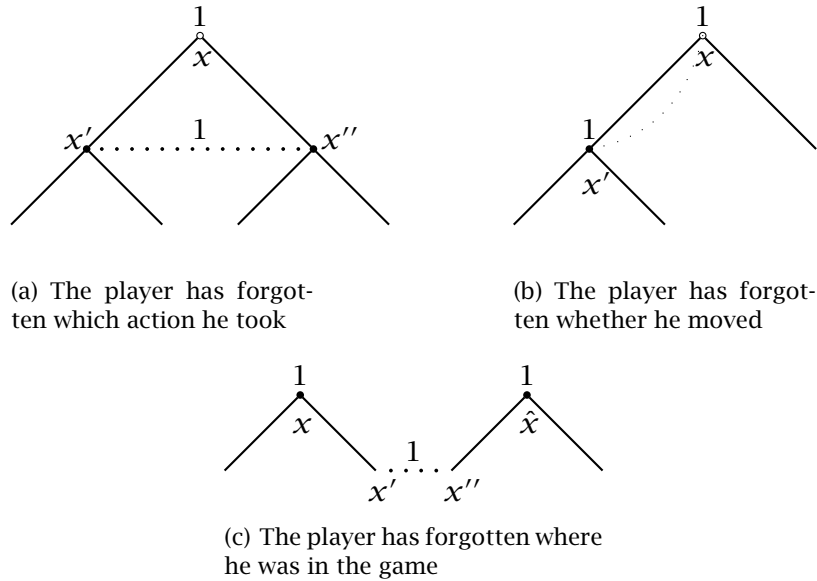


Figure 4: Games of Imperfect Recall.

The situation in Fig. 4(a) (p. 8) is ruled out by condition Condition B because even though both x' and x'' in player 1's second information set have the same predecessor, x , the actions leading from x to the information set are different. The situation in Fig. 4(b) (p. 8) is ruled out by Condition A because x and x' are in the same information set but x is a predecessor of x' . Finally, the situation in Fig. 4(c) (p. 8) is ruled out by Condition B: because x' and x'' are in the same information set and even though x is a predecessor of x' and \hat{x} is a predecessor of x'' , x and \hat{x} are not in the same information set themselves.

The literature on games with imperfect recall is very small, although there are some very interesting papers that might be worth looking at (e.g. the famous game where a drunk driver forgets whether he's been past an exit on the freeway). These games are still quite exotic and their application has been of limited usefulness. This is not to say that there are no exciting areas where these can be applied. One interesting area of research is machine game models of repeated situations: these machines have limited memory and since information is costly to acquire, a player may "forget" some of his past actions. This approach has been extensively used in low-rationality models of learning (evolutionary game theory, for example), where players look at a most recent past when forming expectations about future behavior. This course will only deal with games of perfect recall.

2.2 Finite and Infinite Games

There are three different conceptions of finiteness buried in the definition of extensive form games. The mathematical description can be easily extended to cover these as well.

DEFINITION 2. A **finite** game has (i) a finite number of players, (ii) a finite number of actions, and (iii) finite length histories. Otherwise, the game is **infinite**.

Note that relaxing any of the three requirements results in an infinite number of nodes. Thus, a game is finite if it has a finite number of nodes. Some examples of useful infinite

games that we shall encounter include games where players choose actions from some interval that is a subset of the real line; games which can be repeated indefinitely; or games involving an infinite number of players (we shall see how these games are a way to model incomplete information).

2.3 Informational Categories

We now make very precise several different informational categories. Make sure you understand the terms because we shall use them quite a bit.

DEFINITION 3. We distinguish the following informational categories:

- A game is one of **perfect information** if each information set is a singleton; otherwise it is a game of **imperfect information**.
- A game is one of **certainty** if it has no moves by Nature; otherwise it is a game of **uncertainty**.
- A game is one of **complete information** if all payoff functions are common knowledge; otherwise it is a game of **incomplete information**.
- A game is one of **symmetric information** if no player has information that is different from other players when he moves or at the terminal nodes; otherwise it is a game of **asymmetric information**.

Myerson's Card Game shown in Fig. 3 (p. 5) is a game of complete but imperfect (and therefore asymmetric) information that is also one of uncertainty. Games of imperfect recall are always games of imperfect information.

We shall see games of incomplete (asymmetric) information later in the course. We shall also see how they can be modeled (and solved) as games of imperfect information. It is worth noting that although many games of incomplete information are also games of asymmetric information, the two concepts are not equivalent. For example, the famous principal-agent problem has complete but asymmetric information: both players know all payoff functions but the principal does not observe the agent's effort, even after the end of the game.

It is also possible to have games of incomplete but symmetric information. For example, a Prisoners' Dilemma where Nature moves first and randomly assigns different payoffs to the outcomes, unknown to either player.

2.4 Examples of Games in Extensive Form

Let's now describe the extensive forms for several examples.

EXAMPLE 2. (MATCHING PENNIES.) There are two players who must each put a penny down. If the pennies match (either both heads or both tails), then player 1 pays one dollar to player 2. If they don't match, then player 2 pays one dollar to player 1.

This is a game of complete information, but as it stands, this example omits a crucial piece of information: What do players know when they get to move? After a bit of thought, it should be obvious that there are five possible ways that players can move: (i) player 1 moves first and player 2 observes his action before acting herself; (ii) player 2 moves first and player 1 observes her action before moving himself; (iii) player 1 moves first but player 2

does not observe his action before acting herself; (iv) player 2 moves first but player 1 does not observe her action before moving himself; and (v) the players move simultaneously.

Because in (i) and (ii) each player knows what the other has done in the past when it is time to move, they are games of perfect information. However, in the other three cases, neither player knows what the other has done, and so they are games of imperfect information. With more thought, it should be clear that the last three situations are equivalent from the perspective of each decision-maker: neither player 1 nor player 2 knows the other's action when they make their respective choices. It does not matter when players move if one cannot observe their actions. For example, from player 1's standpoint it does not matter whether player 2 has already made the choice which he cannot see, or is making the choice simultaneously with him, or will make the choice in the future without seeing his action.

And so, we have three different representations of the situation, depending on how we wish to specify it. Fig. 5 (p. 10) shows how the extensive form can be represented with a game-tree diagram. Note that the two variants of the imperfect information specification are equivalent.

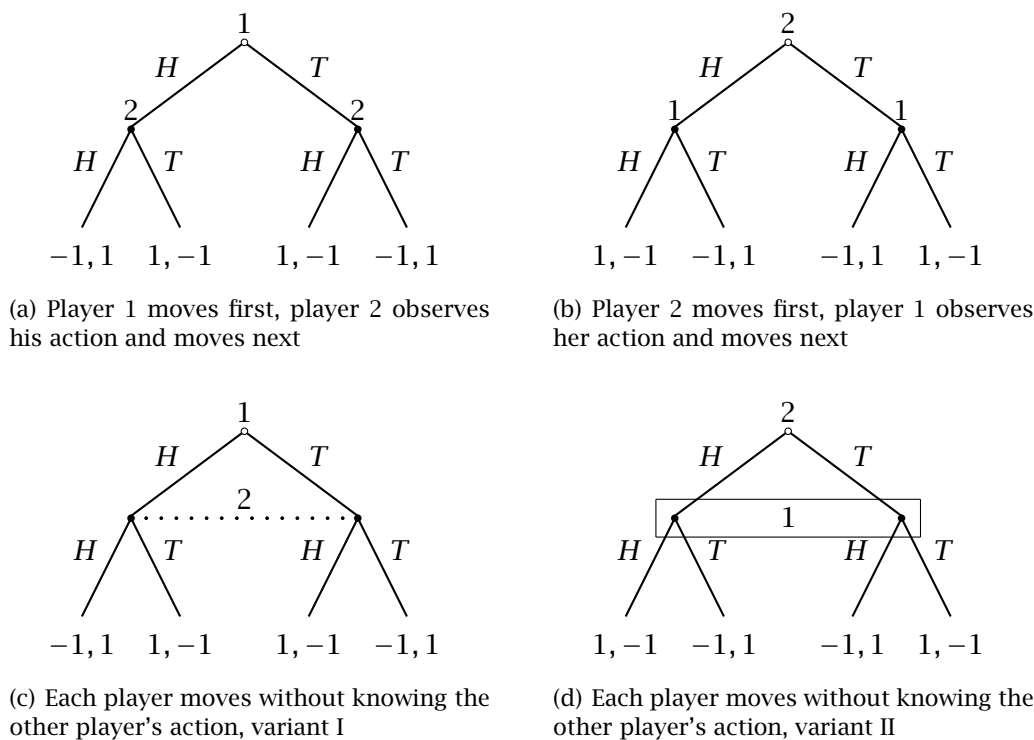


Figure 5: The Three Possible Sequences of the Matching Pennies Game.

We have assumed that players know each other's payoff functions, and so Matching Pennies is a game of complete information. However, (a) and (b) cases in Fig. 5 (p. 10) represent games of perfect information, while (c) and (d) represent the case of imperfect information. To see that (c) and (d) are equivalent representations (as claimed), just examine the information sets of the players (what they know when they get to move). Also, observe that in (c) and (d) there is an "initial" node in the game even though there really is no player that moves first.

EXAMPLE 3. (MULTIPLE SOURCES OF UNCERTAINTY.) Consider the following situation. Two players are engaged in a game where a coin is flipped, and only player 1 observes the outcome. If the outcome is heads, then the players play the Matching Pennies, and if the outcome

were tails, then only player 2 chooses H or T (recall that she does not know whether she's choosing against player 1 or against Nature's "choice" of tails).




Figure 6: The Game with Multiple Sources of Uncertainty.

In this game, player 2 is uncertain about whether player 1 is actually playing (Nature's choice), and about what his choice is if he is playing. Her information set, therefore, includes all three of her decision nodes. Of course, since player 1 observes Nature's choice, he only has one information set, and it is a singleton.

EXAMPLE 4. (MATCHING PENNIES VARIANT A.) Suppose that before playing Matching Pennies, players roll a die to determine who will go first: If the number is less than 3, then player 1 goes first (they play Fig. 5 (p. 10), a), otherwise player 2 goes first (the play Fig. 5 (p. 10), b). This is shown in Fig. 7 (p. 11).




Figure 7: A Game of Uncertainty, Variant A.

EXAMPLE 5. (MATCHING PENNIES, VARIANT B1) Before playing Matching Pennies, players roll the die to determine whether player 1 will pay 1 or 2 dollars if the pennies match. If the die shows a number less than 3, he pays 2 dollars, otherwise, he pays 1 dollar. Player 1 observes the outcome of the roll but player 2 does not, and players move simultaneously.

EXAMPLE 6. (MATCHING PENNIES, VARIANT B2) In this variant, suppose that player 2 observes the outcome of the roll but player 1 does not, and players move simultaneously.




Figure 8: A Game of Uncertainty, Variant B1.




Figure 9: A Game of Uncertainty, Variant B2.

EXAMPLE 7. (MATCHING PENNIES, VARIANT B3) In this variant, suppose that neither player observes the outcome of the roll, and players move simultaneously.




Figure 10: A Game of Uncertainty, Variant B3.

EXAMPLE 8. (TWO-WAY DIVISION.) Two people use the following procedure to share two desirable identical nondivisible objects. One of them proposes an allocation, which the other one either accepts or rejects. In the event of rejection, neither person receives either of the

objects. Each person cares only about the number of objects he receives. This is shown in Fig. 11 (p. 13).




Figure 11: The Two-Way Division Game, Perfect Information.

EXAMPLE 9. Suppose we wanted to model a situation, in which player 2 had to accept or reject the proposal without knowing what this proposal is. In effect, this transforms the game into one of imperfect information, as shown in Fig. 12 (p. 13).




Figure 12: The Two-Way Division Game, Imperfect Information, I.

The game tree in Fig. 13 (p. 13) is equivalent to the tree in Fig. 12 (p. 13) (the payoffs still specify player 1's payoff first and then player 2's payoff). This is important: a strategic situation can have more than one extensive form representation.




Figure 13: The Two-Way Division Game, Imperfect Information, II.

3 Pure Strategies

Player i 's **strategy**, s_i , is a complete rule of action that tells him which action $a_i \in A_i$ to choose at each of his information sets. That is, a strategy specifies what the player is going

to do every time it is his turn to move given what he knows. A player's **strategy space** (sometimes also called a *strategy set*), $S_i = \{s_i\}$, is the set of all possible strategies.

A strategy is a **complete contingent plan of action**. That is, a strategy in an extensive form game is a plan that specifies the action chosen by the player for *every* history after which it is his turn to move, that is, at *each* of his information sets. This is a bit counter-intuitive because it means that the strategy must specify moves at information sets that might never be reached because of actions specified by the player's strategy at earlier information sets.

DEFINITION 4. Let Γ be a game in extensive form. A **pure strategy** for player $i \in \mathcal{I}$ is a function $s_i : \mathcal{H} \rightarrow \mathcal{A}$ such that $s_i(h) \in A_i(h)$ for all $h \in \mathcal{H}$.

Let's list the strategies for the two players in Myerson's Card Game in Fig. 3 (p. 5). Player 1 has two information sets, labeled "b" and "c", with $A_1(b) = \{R, F\}$ and $A_1(c) = \{r, f\}$, so his strategy must specify two actions, $a_b \in A_1(b)$ and $a_c \in A_1(c)$. We shall write his strategy as an ordered set: $s_1 = (a_b, a_c)$, with the first element denoting the action to take at information set "b" and the second denoting the action to take at information set "c". This gives four pure strategies for player 1:

$$S_1 = \{(R, r), (R, f), (F, r), (F, f)\}.$$

For example, (R, f) is the strategy "raise if the card is black, and fold if the card is red."

Player 2 knows that she won't see the color and will only get to choose if player 1 raises, in which case she will either have to meet or pass. There is only one information set for player 2, so her pure strategy must simply specify the action, $a_0 \in A_2(0) = \{m, p\}$, she is to take at this information set. Thus,

$$S_2 = \{m, p\}.$$

The strategy m is then "meet if player 1 raises."

Let's do several other examples. Consider the game in Fig. 11 (p. 13). Player 1 takes action only after the initial history \emptyset , and so his strategy consists of only three possible actions: $S_1 = \{(2, 0), (1, 1), (0, 2)\}$. Player 2, on the other hand, gets to move after three different histories, and so her strategy must specify what to do after each of these histories. That is, a strategy for player 2 must be a *triple* where each member specifies what to do after a particular history. We can use the triple (abc) to represent player 2's strategy, with a being the action to take after history $(2, 0)$, b being the action to take after history $(1, 1)$, and c being the action to take after history $(0, 2)$. For example, $(yy n)$ is a strategy that specifies acceptance of the offers $(2, 0)$ and $(1, 1)$, and rejection of $(0, 2)$. We interpret the strategy as a *contingent plan of action*: if player 1 chooses $(2, 0)$, then player 2 will choose a ; if player 1 chooses $(1, 1)$, then player 2 will choose b , and if player 1 chooses $(0, 2)$, then player 2 will choose c .

Thus, player 2 has 8 available strategies (2 actions to be taken at 3 possible contingencies, or $2^3 = 8$ total strategies):

$$S_2 = \{(yy y), (yy n), (y n n), (y n y), (n y y), (n y n), (n n y), (n n n)\}.$$

Remember that a strategy is a contingent plan of action. For example, the strategy $(n n y)$ reads "reject if player 1 offers $(2, 0)$, reject if player 1 offers $(1, 1)$, accept if player 1 offers $(0, 2)$." Also, remember that it is a complete plan of action, and so player 2's strategy must tell her what to do for each and every possible move by player 1.

Here's an example of a simple game where player 1 gets to move both before and after player 2 has moved. Note that you can draw game trees in just about any direction you want. Usually, left-to-right and up-to-down are the preferred directions (at least for us as people whose languages are written in these directions).




Figure 14: Extensive Form Game with Four Outcomes.

Let's examine Fig. 14 (p. 15) a little more closely. It has two players, $i \in \{1, 2\}$. The game also has seven histories: $H = \{(\emptyset), (A), (B), (B, c), (B, d), (B, d, E), (B, d, F)\}$. Recall that \mathcal{H}_i denotes the set of information sets for player i , and $A_i(h)$ denotes the set of available actions at information set h for all $h \in \mathcal{H}_i$. At the information set \emptyset , player 1 has two actions available: $A_1(\emptyset) = \{A, B\}$. At the information set (B, d) , he has two actions available $A_1(B, d) = \{F, E\}$. Player 2 only gets to move at the information set B , and has two actions available there: $A_2(B) = \{c, d\}$. There are four terminal histories: $Z = \{A, (B, c), (B, d, E), (B, d, F)\}$.

Since a strategy is a complete contingent plan of action, it must specify the actions to be taken at every information set. Player 1 has two information sets in the game, and therefore his strategy will have 2 components: an action to take at the first information set, and an action to take at the second information set. Since in both cases he has two actions available, he has a total of four different strategies:

$$S_1 = \{(AE), (AF), (BE), (BF)\}.$$

Player 2 has only one information set, with two actions there, and so she has only two possible strategies:

$$S_2 = \{c, d\}.$$

This game illustrates a point that is worth emphasizing. It is extremely important to remember that *a strategy specifies the action chosen by a player for every information set at which it is his turn to move, even for information sets that are never reached if the strategy is followed*. That is, in the game in Fig. 14 (p. 15), the first two strategies, (AE) and (AF) specify actions after the history (B, d) even though they specify action A at the initial node (which means that when the strategy is followed, history (B, d) will never be realized, and the second information set will never be reached). In this sense, a strategy differs from what we naturally consider a plan of action. In this instance, every-day language is misleading. We may say that we “plan to choose B ” and since the game will end, there is no reason to plan what to do if we played A instead. However, here we want to know whether B is better than A for player 1. To decide whether this is the case, we need to know what the consequences of choosing A are (otherwise we cannot compare the two actions). But to evaluate the consequences of A , we need to take into account what he would optimally do at his last information set and incorporate this into player 2's expectations to infer what she will do at her information set. Choosing B can only be optimal *in the context of expectations about what would happen if the*

player chose action A instead. It is because we want to find optimal strategies that we must engage in these comparisons and it is for that reason that we must specify the full strategy in what appears to be a redundant fashion. This will become clearer when we analyze some games later on.




Figure 15: An Extensive Form Game with Three Players.

Let's specify the strategies for the game in Fig. 15 (p. 16). There are three players. Player 1 has one information set following the history \emptyset and has two choices available to him there: $A_1(\emptyset) = \{U, D\}$. Player 2 has two information sets, one following the history U and another following the history D . She has two actions available at each information set, with $A_2(U) = \{A, B\}$ and $A_2(D) = \{C, E\}$. Player 3 also has two information sets: one following the history (U, A) , and another following the histories (U, B) and D, C . He also has two actions at each information set with $A_3(U, A) = \{R, T\}$ and $A_3(U, B) = A_3(D, C) = \{P, Q\}$. The strategies then are as follows:

$$\begin{aligned}
 S_1 &= \{U, D\} \\
 S_2 &= \{AC, AE, BC, BE\} \\
 S_3 &= \{RP, RQ, TP, TQ\}
 \end{aligned}$$

Note again that player 3's actions at both decision nodes in his second information set must be the same because the player does not know at which decision node he really is.

Consider the game in Fig. 16 (p. 17). In this game, player 1 has two information sets, one following the history \emptyset , and another following the history A . At the first information set, player 1 has three actions, and so $A_1(\emptyset) = \{A, B, C\}$. At the second information set, player 1 has two actions, and so $A_1(A) = \{W, Z\}$.

Player 2, on the other hand, has only one information set, following either the history B or the history C . She has two actions available at this information set, and so $A_2(B) = A_2(C) =$




Figure 16: An Extensive Form Game with Two Players and Imperfect Information.

$\{X, Y\}$. The strategies then are as follows:

$$S_1 = \{AW, AZ, BW, BZ, CW, CZ\}$$

$$S_2 = \{X, Y\}$$

Again, remember that a strategy specifies a complete plan of action for every information set, even ones that are not reached if the strategy itself is followed. Hence, the pairs of strategies for player 1 with either B or C as the action at the initial information set.

More generally, for a finite game we can determine the number of pure strategies each player has by multiplying the number of actions at each of his information sets. Letting \mathcal{H}_i denote the collection of information sets for player i , the number of pure strategies he has is

$$\#S_i = \prod_{h \in \mathcal{H}_i} \#(A_i(h))$$

In the example in Fig. 16 (p. 17), this calculation is $\#S_1 = \#(A_1(\emptyset)) \times \#(A_1(A)) = 3 \times 2 = 6$, while for the game in Fig. 11 (p. 13), the calculation is $\#S_2 = \#(A_2(2, 0)) \times \#(A_2(1, 1)) \times \#(A_2(0, 2)) = 2 \times 2 \times 2 = 8$, just as we saw.

These examples all assume finite games, so how would one go about specifying a pure strategy if a player has (a) infinitely many actions at some information set; (b) infinitely many information sets; or (c) a combination of both. Consider a game where player 1 must propose a division of some pie of size π and assume that the pie is infinitely divisible. That is, a proposal consists of division $(x, \pi - x)$ where $x \in [0, \pi]$ denotes player 1's share and $\pi - x$ denotes player 2's share. As before, player 2 observes this proposal and can either accept or reject it, with acceptance resulting in a split and rejection ending the game with both players receiving nothing. In this game, player 1 has infinitely many proposal from which to choose from, and player 2 has infinitely many information sets at which she must decide between acceptance and rejection. Clearly, it is impossible to enumerate the strategies the way we have been doing so far.

Fig. 17 (p. 18) shows an extensive form representation of this game. To indicate the fact that player 1 has infinitely many actions at his information set, we connect the branches representing the minimum, 0, and maximum, π , proposals with an arc. This represents the notion that any share between these two is admissible. Since player 2 has infinitely many information sets (one after each possible proposal), we indicate an arbitrary proposal with the branch labeled x and then show player 2's information set following that proposal. The idea is to represent her actions upon seeing player 1's action. Since in each case she can only accept or reject, we only need to specify one exemplar: the payoffs obviously depend on the precise proposal being made if she accepts. (If she had different sets of actions at




Figure 17: Division of a Pie with Infinitely Many Possible Splits.

different information sets, then this representation would not be sufficient to describe the game.) To write player 2's strategy as a function of the history of the game, we need to specify her decision rule to encompass all possible information sets. One such rule would be "accept all offers regardless of what player 1 proposes." Another would be "reject all offers regardless of what player 1 proposes." A more sophisticated rule would be "accept all offers that give player 2 at least some y share, and reject all others." Clearly, there are infinitely many strategies here that one can think of. As we shall see later on, equilibrium analysis will be able to tell us which of these we can safely ignore.

Suppose, for instance, that players only care about the size of their own share of the pie—the larger the better. Then it is trivial to prove that if some player prefers to accept some share, say, $x \in [0, \pi]$, then he will also accept any shares that are at least as large: $\hat{x} \geq x$. This means that we can ignore strategies that do not obey this rule. That is, all possibly optimal strategies will have the form "accept any offer at least as large as y ," and we shall then only have to figure out what this y is. If, on the other hand, preferences are more complicated—e.g., players care about larger shares but only up to a point, either for dietary or equity reasons—then we would have to take this into account and the strategies we have to consider will be more complicated. In that case, potentially optimal strategies may have the form "accept any offer at least as large as y but no larger than z ," and we shall then have to figure out what y and z must be. In either case, we can reduce the infinite number of strategies to something more manageable. For now, just remember that a strategy must specify what a player should do in every possible contingency. If there are infinitely many of these, the strategy would have to handle all of them. When we do the actual analysis, we shall learn how to manage that complexity.

A **strategy profile**, $s = (s_1, s_2, \dots, s_n)$, is an ordered set of strategies consisting of one strategy for each of the n players in the game. One extraordinarily useful piece of notation can let us focus on player i 's strategy s_i in the profile s . We can partition the strategy profile s as:

$$(s_i, s_{-i}) \equiv s,$$

where s_i is player i 's strategy, and s_{-i} is the set of strategies for all other players. For example, if $s = (s_1, s_2, s_3, s_4, s_5)$, and we specify (s_i, s_{-i}) for player $i = 3$, then $s_i = s_3$, and $s_{-i} = (s_1, s_2, s_4, s_5)$. Let $S = S_1 \times S_2 \times \dots \times S_n$ denote the set of strategy profiles.

Because a strategy profile specifies what each player is going to do at every point in the game where it is his turn to move, it in effect describes *how* the game will be played and what its outcome will be if the players follow the strategies in the profile. In other words, each

strategy profile will yield:

- one outcome if that there are no moves by chance; or
- a probability distribution over outcomes if there are moves by chance and the strategies are consistent with information sets where Nature moves.

Some people define players' preference orderings over strategy profiles, but I find this confusing even though it is equivalent to defining them over outcomes. It is confusing because one may think that players actually care about the strategies being played apart from the outcomes they produce. (If this is the case, then this fact must be reflected in the payoffs associated with the outcomes.) We shall define them over outcomes. A player's payoff, $u_i(s)$, is the expected utility that player i receives from the outcome produced by the strategy profile $s \in S$. Thus, each player i 's goal in a game is to choose $s_i \in S_i$ that maximizes $u_i(s_i, s_{-i})$.

4 The Strategic (Normal) Form

Every strategy profile s induces an outcome of the game: a sequence of moves actually taken as specified by the strategies and a probability distribution over the terminal nodes of the game. If the game is one of certainty (no moves by Nature), then s specifies one outcome with certainty. Otherwise, more than one outcome may occur with positive probability. The point is that we can calculate the expected payoffs of all players. Sometimes, it is useful to analyze the game in its **strategic form**, which includes only the players, their actions, and the payoffs in its description.

Putting things a little more formally, let n be the number of players. For each player i , denote the strategy space by S_i . (We shall sometimes write $s_j \in S_i$ to reflect that strategy s_j is a member of the set of strategies S_i .) Let (s_1, s_2, \dots, s_n) denote a strategy profile, where s_1 is the action of player 1, s_2 is the action of player 2, and so on. Let $S = S_1 \times S_2 \times \dots \times S_n$ denote the set of strategy profiles.

For each player i , define the vNM expected utility function $U_i : S \rightarrow \mathbb{R}$ so that for each $s \in S$ that players choose, $U_i(s)$ is player i 's expected payoff from outcome s .

DEFINITION 5. For a game with $\mathcal{I} = \{1, \dots, n\}$ players, the **strategic (normal) form representation** $G = \{\mathcal{I}, S, U\}$ specifies for each player i a set of strategies S_i and a payoff function $U_i : S \rightarrow \mathbb{R}$, where $S = \times S_i$, and $U = (U_1, \dots, U_n)$.

When we analyze these games, we often assume that players choose their strategies simultaneously, and hence we call them **simultaneous-move games**. However, this does not require that players strictly act at the same time. All that is necessary is that each player acts without knowledge of what others have done. That is, players cannot condition their strategies on observable actions of the other players.

Of course, this ignores the information about timing of moves explicitly specified by the extensive form. The question boils down to whether we think such questions are essential to the situation we are trying to analyze. If they are not, then it should not matter greatly if we simplify our description to exclude such information. In an important sense, the strategic form is a **static model** because it dispenses with the dynamics of timing of moves completely.

This may not be as controversial (or useless) as it sounds. First, as we shall see, there are great many situations that we might profitably analyze without reference to the timing of moves. Second, the simplified representation is actually considerably easier to analyze, so we

can benefit from dispensing with information that is not essential. We shall, of course, also see that there are many, many situations where ignoring timing has crucial consequences and our solutions based on the normal form will be quite suspicious precisely because they will discard such information. The question (again) will boil down to the choice of representation, which a researcher has to make based on her skill and experience.

von Neumann and Morgenstern suggested a procedure for simplifying games in extensive form by constructing the strategic form G of any Γ . This is done in an algorithmic way. First, we find all pure strategies for the players. Next, we construct the expected outcomes for all strategy profiles. Finally, we redefine the utility functions on the outcomes to be utility functions on the profiles with expected outcomes.

Consider the following scenario. The two players are going to play Myerson's Card Game in Fig. 3 (p. 5) tomorrow and today they have to plan their moves in advance. Player 1 does not know the color that he will draw but he can condition his strategy on the card color because he knows that he will see it before choosing whether to raise or fold. As we have seen, he has four pure strategies, $S_1 = \{Rr, Rf, Fr, Ff\}$. Player 2, on the other hand, will only ever get to move if player 1 raises, so her pure strategies are $S_2 = \{m, p\}$. The strategy profiles are:

$$S = S_1 \times S_2 = \left\{ \langle Rr, m \rangle, \langle Rr, p \rangle, \langle Rf, m \rangle, \langle Rf, p \rangle, \langle Fr, m \rangle, \langle Fr, p \rangle, \langle Ff, m \rangle, \langle Ff, p \rangle \right\}.$$

We now have to define the expected utility functions for the player. Recall that originally, we defined the utility functions directly in terms of the outcome. However, even if we knew here which strategy profile will be realized (that is, what strategy each player has chosen), we cannot predict the actual outcome of the game because it will depend on the color of the card, which is a chance move. For example, suppose player 1 has chosen the strategy Fr and player 2 has chosen m , and so the strategy profile is $\langle Fr, m \rangle$. The outcome will be folding by player 1 if the card is black, and raising by player 1 and meeting by player 2 if the card is red. Player 1's payoff will be -1 if the card is black, and 2 if the card is red.

So what payoff should player 1 expect from the profile $\langle Fr, m \rangle$? Its expected payoff, of course. Choosing the strategy Fr given that player 2 will be choosing m is equivalent to choosing a lottery, in which player 1 would get -1 with probability 0.5 , and 2 with probability 0.5 . We know how to compute the expected utility in this case:

$$U_1(Fr, m) = 1/2 \times u_1(\text{black}, F) + 1/2 \times u_1(\text{red}, r, m) = 1/2 \times (-1) + 1/2 \times (2) = 0.5.$$

In analogous manner, we would compute player 2's expected payoff:

$$U_2(Fr, m) = 1/2 \times u_2(\text{black}, F) + 1/2 \times u_2(\text{red}, r, m) = 1/2 \times (1) + 1/2 \times (-2) = -0.5.$$

Continuing in this way, we define the expected utility functions for the two players on all strategy profiles, and arrive at the normal form representation of this game of uncertainty shown in Fig. 18 (p. 21).

The strategic game in Fig. 18 (p. 21) describes how the utilities of the players depend on the strategies they choose *at the beginning of the game*. We know from our expected utility theorem that a player would choose the strategy that yields the highest expected payoff because this would be consistent with his preferences. In other words, players will make choices that maximize their expected payoff.

In general, given any extensive form game Γ , its normal form representation G can be constructed as follows. The set of players remains the same. For any player $i \in \mathcal{I}$, let the set

		Player 2	
		<i>m</i>	<i>p</i>
Player 1	<i>Rr</i>	0, 0	1, -1
	<i>Rf</i>	-0.5, 0.5	1, -1
	<i>Fr</i>	0.5, -0.5	0, 0
	<i>Ff</i>	0, 0	0, 0

Figure 18: The Strategic Form of the Game from Fig. 3 (p. 5).

of strategies S_i in the normal form game be the same as the set of strategies in the extensive form. For any strategy profile $s \in S$ and any node x in the tree of Γ , define $P(x|s)$ to be the probability that the path of play will go through node x , when the path of play starts at the initial node, and at any decision node in the path, the next node is determined by the relevant player's strategy in s , and, at any node where nature moves, the next node is determined by the probability distribution given in Γ . At any terminal node $z \in Z$, let $u_i(z)$ be player i 's payoff from outcome z . Then, for any strategy profile $s \in S$ and any $i \in \mathcal{I}$, let $U_i(s)$ be:

$$U_i(s) = \sum_{z \in Z} P(z|s) u_i(z).$$

That is, $U_i(s)$ is player i 's expected utility if all players implement the strategies according to s . If G is derived from Γ in this way, it is called the **strategic (normal) form representation** of Γ .

To make things more concrete, let's construct the strategic form of the two extensive-form games from Fig. 19 (p. 21).




Figure 19: Two Simple Games.

In the imperfect information variant 19(a), player 1 and player 2 have two strategies each: $S_1 = \{U, D\}$, and $S_2 = \{L, R\}$. There are four outcomes,

$$S = S_1 \times S_2 = \{ \langle U, L \rangle, \langle D, L \rangle, \langle U, R \rangle, \langle D, R \rangle \}.$$

Without chance moves, there is no need to transform the utility functions. The strategic form of 19(a) is in Fig. 20 (p. 22).

The situation in Fig. 19(b) (p. 21) is very different. Although player 1 still has two pure strategies, $S_1 = \{U, D\}$, player 2 can condition her choice on player 1's. She has two information sets, and her strategy must specify two actions: $a_U \in A_2(U) = \{L, R\}$ is the choice after player 1 chooses U , and $a_D \in A_2(D) = \{l, r\}$ is the choice after player 1 chooses D . We shall write player 2's strategy as the ordered pair (a_U, a_D) . Hence, the strategy set for player

		Player 2	
		<i>L</i>	<i>R</i>
Player 1	<i>U</i>	2, 2	4, 0
	<i>D</i>	1, 0	3, 1

Figure 20: The Strategic Form of the Game from Fig. 19(a) (p. 21).

2 consists of four pure strategies; $S_2 = \{(L, l), (L, r), (R, l), (R, r)\}$. The game now has eight strategy profiles:

$$S = S_1 \times S_2 = \left\{ \langle U, (L, l) \rangle, \langle U, (L, r) \rangle, \langle U, (R, l) \rangle, \langle U, (R, r) \rangle, \right. \\ \left. \langle D, (L, l) \rangle, \langle D, (L, r) \rangle, \langle D, (R, l) \rangle, \langle D, (R, r) \rangle \right\}.$$

Because there are no moves by chance, there is no need to transform the utility functions, so the strategic form is given in Fig. 21 (p. 22).

		Player 2			
		<i>(L, l)</i>	<i>(L, r)</i>	<i>(R, l)</i>	<i>(R, r)</i>
Player 1	<i>U</i>	2, 2	2, 2	4, 0	4, 0
	<i>D</i>	1, 0	3, 1	1, 0	3, 1

Figure 21: The Strategic Form of the Game from Fig. 19(b) (p. 21).

A seemingly innocuous change in the information structure of the extensive form led to two distinct normal form representations.

4.1 Examples of Converting Extensive to Strategic Form

Going back to the extensive form game in Fig. 11 (p. 13), we can convert the game to its normal form equivalent by specifying the players' pure strategies and the payoffs. We already know the strategies:

$$S_1 = \{(2, 0), (1, 1), (0, 2)\},$$

and

$$S_2 = \{(yyy), (yyn), (ynn), (yny), (nyy), (nyn), (nny), (nnn)\},$$

which give $3 \times 8 = 24$ strategy profiles. Because there are no moves by chance, we do not have to redefine the utility functions, and so we get the strategic form in Fig. 22 (p. 22).

		Player 2							
		<i>yyy</i>	<i>yyn</i>	<i>ynn</i>	<i>yny</i>	<i>nyy</i>	<i>nyn</i>	<i>nny</i>	<i>nnn</i>
Player 1	<i>(2, 0)</i>	2, 0	2, 0	2, 0	2, 0	0, 0	0, 0	0, 0	0, 0
	<i>(1, 1)</i>	1, 1	1, 1	0, 0	0, 0	1, 1	1, 1	0, 0	0, 0
	<i>(0, 2)</i>	0, 2	0, 0	0, 0	0, 2	0, 2	0, 0	0, 2	0, 0

Figure 22: The Strategic Form of the Game from Fig. 11 (p. 13).

Although there is only one way of converting an extensive form game to a strategic form game, this does not mean that we would get different strategic games from different extensive forms. Recall that Fig. 12 (p. 13) and Fig. 13 (p. 13) described the same strategic situation

		Player 2	
		y	n
Player 1	(2, 0)	2,0	0,0
	(1, 1)	1,1	0,0
	(0, 2)	0,2	0,0

Figure 23: The Strategic Form of the Games from Figures 12 and 13.

using different trees. Both of these have the same strategic form representation shown in Fig. 23 (p. 23).

Note how the game in Fig. 22 (p. 22) differs from the game in Fig. 23 (p. 23). This is because the two describe two radically different extensive form games. In particular in the first case player 2 has three information sets, while in the second case she only has one (with three decision nodes in it). Intuitively, however, it does make sense that the two extensive-form games from Fig. 12 (p. 13) and Fig. 13 (p. 13) should have the same strategic form because they do describe equivalent strategic situations.

4.1.1 Several Chance Moves

Let's now do an example with a more than one chance move, as in Fig. 24 (p. 23).




Figure 24: A Card Game in Extensive Form With Two Chance Moves.

The first step is to calculate the probability distribution over the terminal nodes. That is, calculate the probability distributions over outcomes that is induced by each strategy profile. To do this, we take each strategy profile and calculate the probabilities with which it produces various outcomes. For example, if players use $\langle Rr, m \rangle$, then with probability $1/2$ player 1 will end up at the left information set and choose R which will then be followed by m by player 2, leading to the outcome z_1 . If, on the other hand, player 1 ends up at his second information set, which can happen with probability $1/2$ as well, then he would play r , followed by player 2's m , leading to the outcome z_5 . Hence, with this strategy profile, two outcomes are possible, z_1 and z_5 , each of which can occur with probability $1/2$.

Consider now $\langle Rf, p \rangle$. With probability $1/2$ player 1 would have to play R , which will then be followed by p , leading to the outcome z_2 . With the same probability, player would have to play f , in which case chance determines the outcome, and it will be either z_7 (with probability $3/4$) or z_8 (with probability $1/4$). Hence, this strategy profile can result in one of

three possible outcomes. The probabilities are $\Pr[z_2] = 1/2$, $\Pr[z_7] = 1/2 \times 3/4 = 3/8$, and $\Pr[z_8] = 1/2 \times 1/4 = 1/8$. Since this must be a valid probability distribution (one of these outcomes must occur for sure and they are mutually exclusive), it follows that the sum of these probabilities should be one, which we can easily verify to be the case. Continuing in this fashion, we can generate the probability distributions over outcomes for each of the strategy profiles, as shown in Tab. 1 (p. 24).

Profile	Outcome Probability, $P(z s)$								$U_1(s)$	$U_2(s)$
	z_1	z_2	z_3	z_4	z_5	z_6	z_7	z_8		
$\langle Rr, m \rangle$	$1/2$	0	0	0	$1/2$	0	0	0	0	0
$\langle Rr, p \rangle$	0	$1/2$	0	0	0	$1/2$	0	0	1	-1
$\langle Rf, m \rangle$	$1/2$	0	0	0	0	0	$3/8$	$1/8$	$5/8$	$-5/8$
$\langle Rf, p \rangle$	0	$1/2$	0	0	0	0	$3/8$	$1/8$	$1/8$	$-1/8$
$\langle Fr, m \rangle$	0	0	$1/6$	$2/6$	$1/2$	0	0	0	$-5/6$	$5/6$
$\langle Fr, p \rangle$	0	0	$1/6$	$2/6$	0	$1/2$	0	0	$2/3$	$-2/3$
$\langle Ff, m \rangle$	0	0	$1/6$	$2/6$	0	0	$3/8$	$1/8$	$-5/24$	$5/24$
$\langle Ff, p \rangle$	0	0	$1/6$	$2/6$	0	0	$3/8$	$1/8$	$-5/24$	$5/24$

Table 1: Probability Distributions Over Outcomes.

Clearly, the sum of all columns for each row should equal 1 (that is, each profile will produce some outcome with certainty). The expected utilities for each profile are calculated in the usual manner. That is, to calculate the expected payoff from a profile, we take the probabilities of the outcomes and multiply those by the corresponding payoffs, then sum over them. For example, as we have seen $\langle Rr, m \rangle$ yields z_1 or z_5 with equal probability. Hence,

$$U_1(\langle Rr, m \rangle) = 1/2 u_1(z_1) + 1/2 u_1(z_5) = 1/2(2) + 1/2(-2) = 0$$

$$U_2(\langle Rr, m \rangle) = 1/2 u_2(z_1) + 1/2 u_2(z_5) = 1/2(-2) + 1/2(2) = 0,$$

and these give us the last two columns in Tab. 1 (p. 24). Let's calculate the expected payoffs for the other profile we looked at:

$$U_1(\langle Rf, p \rangle) = 1/2 u_1(z_2) + 3/8 u_1(z_7) + 1/8 u_1(z_8) = 1/2(1) + 3/8(-1) + 1/8(0) = 1/8$$

$$U_2(\langle Rf, p \rangle) = 1/2 u_2(z_2) + 3/8 u_2(z_7) + 1/8 u_2(z_8) = 1/2(-1) + 3/8(1) + 1/8(0) = -1/8.$$

Continuing in this way, we compute the remaining expected payoffs. Using Tab. 1 (p. 24), we can now construct the strategic form representation of the game in Fig. 24 (p. 23), as shown in Fig. 25 (p. 24).

		Player 2	
		m	p
Player 1	Rr	0, 0	1, -1
	Rf	$5/8, -5/8$	$1/8, -1/8$
	Fr	$-5/6, 5/6$	$2/3, -2/3$
	Ff	$-5/24, 5/24$	$-5/24, 5/24$

Figure 25: The Strategic Form of the Game from Fig. 24 (p. 23).

4.1.2 Three Players

Consider the strategic form of the game in Fig. 26 (p. 25). We have not seen games with three players in strategic form, but the principles are the same. Player 1 has one information set with two actions, so he has two pure strategies, $S_1 = \{U, D\}$. Player 2 has two information sets with two actions at each, hence four pure strategies, $S_2 = \{(A, C), (A, E), (B, C), (B, E)\}$. Player 3 has two information sets with two actions at each, or four pure strategies, $S_3 = \{(R, P), (R, Q), (T, P), (T, Q)\}$. This gives us $2 \times 4 \times 4 = 32$ strategy profiles. Fortunately, there are no chance moves here, so we won't have to redefine the utility functions.




Figure 26: An Extensive Form Game with Three Players.

There are two ways to write the strategic form with three players. One is to write out as many separate games between players 1 and 2 as there are pure strategies for player 3. In each of these 2-player games, player 3 is choosing a particular pure strategy. In our example, this would give us 4 2-player matrices with $2 \times 4 = 8$ cells each. Of course, the total number of cells will still be 32.

The other way is to write out one big payoff matrix, as we now do here. Player 3's strategies form the rows, while player 1 and player 2's strategies jointly determine the columns. This payoff matrix has 4 rows and 8 columns. The payoffs are listed as ordered triples, (u_1, u_2, u_3) , where u_i is player i 's payoff from the relevant outcome. The strategic form is in Tab. 2 (p. 25).

		Player 1							
		U				D			
		Player 2				Player 2			
Player 3	(R, P)	(0, 1, 2)	(0, 1, 2)	(1, 0, 0)	(1, 0, 0)	(2, 2, 1)	(3, 0, 1)	(2, 2, 1)	(3, 0, 1)
	(R, Q)	(0, 1, 2)	(0, 1, 2)	(0, 2, 1)	(0, 2, 1)	(0, 0, 0)	(3, 0, 1)	(0, 0, 0)	(3, 0, 1)
	(T, P)	(1, 2, 0)	(1, 2, 0)	(1, 0, 0)	(1, 0, 0)	(2, 2, 1)	(3, 0, 1)	(2, 2, 1)	(3, 0, 1)
	(T, Q)	(1, 2, 0)	(1, 2, 0)	(0, 2, 1)	(0, 2, 1)	(0, 0, 0)	(3, 0, 1)	(0, 0, 0)	(3, 0, 1)

Table 2: The Strategic Form with Three Players.

Often we would not even have to specify the extensive form before going to the strategic form. Let's see several canonical examples of games in normal form.

4.2 Examples in Strategic Form

Let's model a situation where two players, $i \in \{1, 2\}$, want to decide between two types of entertainment to which they want to go together but the decision must be made without knowledge of what the other will do (say they are in their offices and the phones are down so they cannot communicate beforehand). The two available pieces of entertainment for the night are a boxing match (fight) and a ballet. For each player then, the set of actions consists of (1) go to the fight, or (2) go to the ballet. Note that the actions are exhaustive and mutually exclusive. This means that each player has two pure strategies, so the set is called the *strategy space* for the player.

Continuing with the example, the strategy profile then consists of one strategy for each of the two players. This gives us four different strategy profiles: (1) player 1 goes to the fight, player 2 goes to the fight; (2) player 1 goes to the fight, player 2 goes to the ballet; (3) player 1 goes to the ballet, player 2 goes to the fight; and (4) player 1 goes to the ballet, player 2 goes to the ballet. We shall specify an outcome (strategy profile) by listing first the strategy for player 1 and then the strategy for player 2. Thus, the four outcomes above can be written as (1) *(Fight, Fight)*; (2) *(Fight, Ballet)*; (3) *(Ballet, Fight)*; and (4) *(Ballet, Ballet)*.

Since each strategy profile produces a different outcome in this game, the game has 4 possible outcomes, in 2 of which the players go together to the same place, and 2 in which they fail to coordinate. Each player has (ordinal) preferences over these four outcomes. In other words, each player ranks these outcomes according to their desirability using some criterion. As we know, if preferences are rational, we can represent them numerically. Hence, we use appropriate numbers whose ordinal ranking represents the preferences as payoffs. Each outcome then consists of two elements which specify the payoff for each player for this outcome. This is often called the *payoff vector*.

Let's say that player 1 is a man, who prefers going to the fight to seeing *Swan Lake*. And let's say that player 2 is a woman who prefers the culture of ballet to the somewhat less elegant bashing of heads. However, both prefer to go together regardless of the type of entertainment. Their worst outcome is when they end up alone at any of the places and it does not matter which place they happen to be at. Thus, the man's ordering is:

$$(F, F) \succ (B, B) \succ (F, B) \sim (B, F)$$

and the woman's ordering is:

$$(B, B) \succ (F, F) \succ (F, B) \sim (B, F)$$

Now that we have specified the ordinal rankings, we need to choose a payoff function to represent the orderings. Denote the man's utility function by u_1 , and the woman's utility function by u_2 . We need two functions such that:

$$\begin{aligned} u_1(F, F) &> u_1(B, B) > u_1(F, B) = u_1(B, F) \\ u_2(B, B) &> u_2(F, F) > u_2(F, B) = u_2(B, F). \end{aligned}$$

One possible and simple specification is

$$\begin{aligned} u_1(F, F) &= u_2(B, B) = 2 \\ u_1(B, B) &= u_2(F, F) = 1 \\ u_1(F, B) &= u_1(B, F) = u_2(F, B) = u_2(B, F) = 0. \end{aligned}$$

A convenient way of describing the (finite) strategy spaces of the players and their payoff functions for two-player games is to use a bi-matrix,⁶ as illustrated in Fig. 27 (p. 27).

		Player 2	
		<i>F</i>	<i>B</i>
Player 1	<i>F</i>	2,1	0,0
	<i>B</i>	0,0	1,2

Figure 27: Battle of the Sexes.

In this figure the two rows represent the two possible actions for player 1 (the man), and the columns represent the two possible actions for player 2 (the woman). Each box represents a possible outcome from these action, and the numbers in each box are the players' payoffs to the action profile to which the box corresponds. The first number is player 1's payoff, and the second number is player 2's payoff.

Note: the Battle of the Sexes game represents a situation where players must coordinate their actions but where they have opposed preferences over the coordinated outcomes. We shall see two other types of coordination games: pure coordination (where players only care about coordinating) and Pareto coordination (where both strictly prefer one of the coordinated outcomes to the other).⁷

Recall that although we call this a *simultaneous-moves* game, it is not necessary for players to actually act at the same time. All that is required is that each player acts with no knowledge about how the other player acts. In our BoS game, this can be achieved by requiring the players to make their choices without having access to a communication device.

Perhaps the most celebrated example of a cooperative strategic situation is the Prisoners' Dilemma. Two suspects are arrested and charged for a crime. The authorities lack enough evidence to convict them unless at least one confesses. The police put the suspects in separate cells and the DA comes to talk to them separately. The DA gives the same spiel to both: If neither suspect confesses, then both will be convicted of a minor offense and will spend 1 month in jail. If both confess, they will be sentenced to jail for 6 months. Finally, if one confesses and the other does not, the one who confesses is granted immunity and is released immediately, while the other will get a year (the 6 months for the crime and 6 more for obstructing justice).⁸

Let's say that a prisoner cooperates, *C*, with the other inmate if he remains silent, and he defects, *D*, if he spills the beans to the prosecution. Clearly, the most preferred outcome for a prisoner is to go free, which is only possible if he defects while the other cooperates. The next best outcome is for both of them to cooperate and get the shortest sentence. This is followed by the medium-length sentence resulting from both defecting. The absolute worst is for a prisoner to cooperate when the other defects. The preference ordering then is:

$$\begin{aligned} \text{Prisoner 1 : } & (D, C) \succ (C, C) \succ (D, D) \succ (C, D) \\ \text{Prisoner 2 : } & (C, D) \succ (C, C) \succ (D, D) \succ (D, C), \end{aligned}$$

⁶This is just like a regular matrix except each entry consists of two numbers instead of one.

⁷There is now a more politically-correct version of the BoS game, called *Bach or Stravinsky*, which involves two sexless players deciding between concerts of music by the two composers. Because it seems to lose some of the punch, I prefer the original formulation. If this bothers you, you can assume the woman likes boxing and the man likes ballet instead.

⁸*Law and Order's* McCoy was a master of the Prisoner's Dilemma.

where, as before, the strategy profile lists player 1's choice first and player 2's choice second; so (C, D) , for example, means prisoner 1 cooperates and prisoner 2 defects, which would result in the former getting the stiff sentence, and the latter going Scott-free. As before, the easiest way to represent these preferences is to start with a payoff of 0 for the worst outcome, and then work our way up adding 1 for each next best outcome. In this way, we can represent this game using the payoff matrix in Fig. 28 (p. 28).

		Prisoner 2	
		<i>C</i>	<i>D</i>
Prisoner 1	<i>C</i>	2,2	0,3
	<i>D</i>	3,0	1,1

Figure 28: Prisoner's Dilemma, I.

It is worth repeating that the payoffs are only meant to represent the ordinal rankings of the outcomes. For example, using the length of sentence as payoff we can construct a game that is strategically equivalent:

		Prisoner 2	
		<i>C</i>	<i>D</i>
Prisoner 1	<i>C</i>	-1, -1	-12, 0
	<i>D</i>	0, -12	-6, -6

Figure 29: Prisoner's Dilemma, II.

The situation is absolutely the same because the ordinal ranking of the payoffs is the same as in Fig. 28 (p. 28). The PD has been extensively studied in many different settings. Two of the most celebrated applications are to the arms race between the US and USSR and the "tragedy of the commons" where a common resources is overconsumed.

Another common game is the Stag Hunt suggested by Jean-Jacques Rousseau. It is very similar to PD except each player prefers the outcome in which both cooperate to the one in which one defects. The original story is as follows. There are two hunters and each has two options. He can catch a hare for sure or participate in the hunt for a stag. If both pursue the stag, they are sure to catch it and then share equally. This share is bigger than the hare. If either one goes for the hare while the other is pursuing the stag, the catcher of the hare gets to take it home while the other goes empty-handed. Each prefers to hunt for the hare alone.

The arms race situation is perhaps better modeled as a Stag Hunt instead of a Prisoners' Dilemma because acquiring arms is expensive and useless if the other one has disarmed. In this situation, arming is costly, so both countries most prefer the outcome where neither one arms. The next-best outcome is unilateral armament because it provides security (and perhaps can be used to extract concessions). The third-best outcome is for both to arm. Although this does not change the military balance, it is expensive, so both suffer the costs of doing so. The worst outcome is to fail to arm while the other arms unilaterally. In this case, the other side can extract huge concessions. Using the labels C for cooperate (meaning "do not arm") and D for defect (meaning "arm"), the preference orderings are:

$$\begin{aligned} \text{US} : (C, C) &> (D, C) > (D, D) > (C, D) \\ \text{USSR} : (C, C) &> (C, D) > (D, D) > (D, C). \end{aligned}$$

Using our now-familiar way of assigning payoffs to represent these preferences, we can construct the payoff matrix. Fig. 30 (p. 29) shows the Stag Hunt strategic situation applied to the security dilemma.

		USSR	
		C	D
US	C	3, 3	0, 2
	D	2, 0	1, 1

Figure 30: Stag Hunt modeling the Security Dilemma.

Comparing this with the Prisoner’s Dilemma in Fig. 28 (p. 28) makes the difference quite clear. We shall see later how the solutions to these two games differ and why. For now, one should keep in mind that choosing which of these two is a “better” representation of the arms race is up to the analyst. The relevant question is: which game captures the strategic situation with less distortion of reality. There is no unambiguous answer to this, especially once you get into more sophisticated models where the choices are not so simple.

The point of these examples is to convey the idea that games really describe strategic settings which may be the same across various actual applications. The abstract model can thus capture the underlying incentives in these settings. The idea here is to understand that we don’t care much about labels (we can have players, hunters, countries, prisoners). We also don’t really care about the verbal description of a particular situation. What we are interested in is the strategic environment: available actions and payoffs.

4.3 Reduced Strategic Form

Consider the (Little Horsey) game in Fig. 14 (p. 15). Player 1 has four pure strategies and player 2 has only two, resulting in a 4×2 payoff matrix. The strategic representation of this game is given in Fig. 31 (p. 29), and shows an important aspect of the definition of pure strategies: The pure strategy space may be unnecessarily large in the sense that it may contain pure strategies that are “equivalent” because they have the same consequences regardless of what the opponent does. In this example, the strategies *AE* and *AF* for player 1 are equivalent.

		Player 2	
		<i>c</i>	<i>d</i>
Player 1	<i>AE</i>	1, 1	1, 1
	<i>AF</i>	1, 1	1, 1
	<i>BE</i>	-1, 1	3, 2
	<i>BF</i>	-1, 1	4, 0

Figure 31: The Normal Form of the Game from Fig. 14 (p. 15).

Two pure strategies are *equivalent* if they induce the same probability distribution over the outcomes for all pure strategies for the opponents. Or, putting it a bit more formally:

DEFINITION 6. Given any strategic form game $G = \{I, S, U\}$, for any player i and any two strategies $s_1, s_2 \in S_i$, the strategies s_1 and s_2 are **payoff-equivalent** if, and only if,

$$U_j(s_1, s_{-i}) = U_j(s_2, s_{-i}), \quad \forall s_{-i} \in S_{-i}, \quad \forall j \in I.$$

That is, no matter what all other players do, no player cares whether i uses s_1 or s_2 . Let's parse this expression. To see whether two strategies for player 1 are payoff-equivalent, we take each strategy of player 2 in turn and compare the payoffs that player 1 obtains from playing s_1 and s_2 against that, then we compare the payoffs that player 2 obtains from player 1 playing s_1 and s_2 against her strategy. If either of these two comparisons produces a difference, stop: the two strategies are not payoff-equivalent. If, on the other hand, they yield the same payoffs in both cases, proceed to the next strategy for player 2 and repeat the process. If you exhaust all strategies for player 2 in this way and the comparisons have not yielded any differences, then the two strategies for player 1 are payoff-equivalent.

In our example from Fig. 31 (p. 29), the two pure strategies AE and AF always lead to the same outcome because the game ends when the first action is taken and so the second information set is never reached. This happens regardless of what player 2 does at her information set. That is, fix player 2's strategy to be c , then: (i) player 1's payoff from AE is 1, which is the same as his payoff from AF ; (ii) player 2's payoff from player 1 choosing AE is 1, which is the same as her payoff from him choosing AF . So neither player cares if player 1 chooses AE or AF if player 2 chooses c . Next, fix player 2's strategy to be d , then: (i) player 1's payoff from AE is 1, which is the same as his payoff from AF ; (ii) player 2's payoff from player 1 choosing AE is 1, which is the same as her payoff from him choosing AF . Hence, neither player cares if player 1 chooses AE or AF if player 2 chooses d . Since there are no more strategies for player 2 to check against, we are done: no player cares what player 1 does regardless of what player 2 chooses. Observe that in these comparisons we had to check whether player 1 himself would care, not just whether his opponent would. We can now simplify the normal form representation by removing all but one strategies from every class of equivalent strategies.

DEFINITION 7. The **purely reduced normal form** of an extensive form game is obtained by eliminating all but one member of each equivalence class of pure strategies.

Therefore, we can remove either AE or AF (but not both) to obtain the reduced normal form shown in Fig. 32 (p. 30). The "new" strategy for player 1 is called A .

		Player 2	
		c	d
Player 1	A	1, 1	1, 1
	BE	-1, 1	3, 2
	BF	-1, 1	4, 0

Figure 32: The Reduced Normal Form of the Game from Fig. 12 (p. 13).

The example we just did may be a bit misleading because the payoffs for the players are always the same in all the outcomes regardless of what player 2 chooses. This need not be the case. To see that, consider the strategic form game in Fig. 33 (p. 31).

To decide whether U and D are payoff-equivalent, we first fix player 2's strategy at L and observe that players get $(3, 1)$ no matter which of the two pure strategies under consideration player 1 chooses. We then fix player 2's strategy at R and observe that players get $(-2, 0)$ regardless of whether player 1 chooses U or D . Hence, the two are payoff-equivalent, and we can eliminate one of them. Observe that a player can get different payoffs depending on whether player 2 chooses L or R from strategies that are payoff-equivalent (i.e., player 1 can get either 3 or -2) but this is not the relevant comparison to make. For example, both players

		Player 2	
		L	R
Player 1	U	3, 1	-2, 0
	M	4, 3	4, 3
	D	3, 1	-2, 0.

		Player 2	
		L	R
Player 1	U	3, 1	-2, 0
	M	4, 3	4, 3

Figure 33: Reducing a Game with Different Payoffs.

get (4, 3) if player 1 chooses M regardless of player 2's action. However, this does not mean that L and R are payoff-equivalent (because players would get different payoffs against either one of these if player 1 chooses a different strategy.)




Figure 34: Another Game from Myerson (p. 55).

For practice, let's find the reduced normal form of the extensive-form game in Fig. 34 (p. 31). Player 1 has two information sets, with $A_1(\emptyset) = \{a, b\}$ and $A_1(b) = \{c, d, e\}$. He has six pure strategies:

$$S_1 = \{(a, c), (a, d), (a, e), (b, c), (b, d), (b, e)\}.$$

Player 2 has only one information set, and therefore just two pure strategies:

$$S_2 = \{x, y\}.$$

There are 12 pure-strategy profiles. Of these, exactly two involve chance moves: $\langle (b, e), x \rangle$ and $\langle (b, e), y \rangle$. We have to calculate the expected utilities for these:

$$U_1((b, e), x) = (0.5)(0) + (0.5)(6) = 3$$

$$U_1((b, e), y) = (0.5)(8) + (0.5)(6) = 7.$$

Analogously, for player 2:

$$U_2((b, e), x) = (0.5)(8) + (0.5)(0) = 4$$

$$U_2((b, e), y) = (0.5)(0) + (0.5)(0) = 0.$$

		Player 2	
		x	y
Player 1	(a, c)	6, 0	6, 0
	(a, d)	6, 0	6, 0
	(a, e)	6, 0	6, 0
	(b, c)	8, 0	0, 8
	(b, d)	0, 8	8, 0
	(b, e)	3, 4	7, 0

Figure 35: The Strategic Form of the Game from Fig. 34 (p. 31).

We are now ready to construct the strategic form representation of this extensive form game. The result is in Fig. 35 (p. 32).

It is fairly obvious that the strategies (a, c) , (a, d) , and (a, e) are payoff equivalent to one another because regardless of what player 2 does, the outcome from all three is the same. In other words, player 1 does not care what player 2 does if he chooses any of these three strategies. We can therefore merge these three strategies into a new one, called A , with the resulting payoff matrix in Fig. 36 (p. 32).

		Player 2	
		x	y
Player 1	A	6, 0	6, 0
	(b, c)	8, 0	0, 8
	(b, d)	0, 8	8, 0
	(b, e)	3, 4	7, 0

Figure 36: The Purely Reduced Strategic Form of the Game from Fig. 34 (p. 31).

We can reduce this game further, but to do this, we need to introduce the concept of mixed strategies.

5 Mixed Strategies in Strategic Form Games

So far, we have considered only strategies that involve playing a selected action with probability 1. We called these **pure strategies** to emphasize this. We now consider randomized choices.

DEFINITION 8. A **mixed strategy** for player i , denoted by σ_i , is a probability distribution over i 's set of pure strategies S_i . Denote the mixed strategy space for player i by Σ_i , where $\sigma_i(s_i)$ is the probability that σ_i assigns to the pure strategy $s_i \in S_i$. The space of mixed strategy profiles is denoted by $\Sigma = \Delta \Sigma_i$.

Thus, if player i has K pure strategies: $S_i = \{s_{i1}, s_{i2}, \dots, s_{iK}\}$, then a mixed strategy for player i is a probability distribution $\sigma_i = \{\sigma_i(s_{i1}), \sigma_i(s_{i2}), \dots, \sigma_i(s_{iK})\}$, where $\sigma_i(s_{ik})$ is the probability that player i will choose strategy s_{ik} for $k = 1, 2, \dots, K$. Since σ_i is a probability distribution, we require that $\sigma_i(s_{ik}) \in [0, 1]$ for all $k = 1, 2, \dots, K$ and $\sum_{k=1}^K \sigma_i(s_{ik}) = 1$. That is, the probabilities must be non-negative and not larger than 1, and should sum up to 1. You can think of a mixed strategy as a lottery whose "outcomes" are pure strategies.

Each player's randomization is statistically independent of those of his opponents,⁹ and the payoffs to the mixed strategy profile are the expected values of the corresponding pure strategy payoffs.¹⁰ You should now see why we needed Expected Utility Theory. Player i 's payoff from a mixed strategy profile $\sigma \in \Sigma$ in an n -player game is

$$U_i(\sigma) = \sum_{s \in S} \left(\prod_{j=1}^n \sigma_j(s_j) \right) u_i(s)$$

Let's parse this expression. The mixed strategy profile σ is a list of mixed strategies, one for each player: $\sigma = \{\sigma_1, \sigma_2, \dots, \sigma_n\}$. Each of these mixed strategies, e.g. σ_i , is a list of probabilities associated with player i 's set of pure strategies. To find the probability of an outcome, we need to calculate the probability that all players choose the pure strategies that produce this outcome. Thus, if the pure strategy profile $s \in S$ produces the outcome we are interested in, the probability of this outcome is the product of probabilities that each player chooses the pure strategy in this profile (because of independence).

Consider first an example from a game without chance moves, like Matching Pennies. To make things specific, let's use the mixed strategy profile $\sigma = \langle (1/3H, 2/3T), (1/4H, 3/4T) \rangle$. In this profile, player 1's mixed strategy specifies playing H with probability $1/3$ and T with probability $2/3$, and player 2's mixed strategy specifies playing H with probability $1/4$, and T with probability $3/4$. There are four pure strategy profiles: $S = \{(H, H), (H, T), (T, H), (T, T)\}$ that produce the four outcomes of the game.

As usual, the strategy profile σ induces a probability distribution over the outcomes. The probability of each outcome is the product of the probabilities that each player chooses the relevant strategy. For example, the probability of the pure strategy profile (H, H) being played is $(1/3)(1/4) = 1/12$. Analogously, the probabilities of the other pure strategy profiles being played are $\Pr(H, T) = 1/4$, $\Pr(T, H) = 1/6$, and $\Pr(T, T) = 1/2$. (You should verify that these sum to 1, which they must because they are probabilities of exhaustive and mutually exclusive events.) Fig. 37 (p. 33) shows the probability distribution over the four possible outcomes induced by the two mixed strategies.

	H	T
H	$1/12$	$1/4$
T	$1/6$	$1/2$

Figure 37: The probability distribution over outcomes induced by σ .

Player 1's payoffs from these outcomes are $u_1(H, H) = u_1(T, T) = 1$ and $u_1(H, T) = u_1(T, H) = -1$. Multiplying the payoffs by the probability of obtaining them and summing over (the expected utility calculation we have done before) yields an expected payoff of $1/12(1) + 1/2(1) + 1/4(-1) + 1/6(-1) = 1/6$. Thus, player 1's expected payoff from the mixed strategy profile σ as specified above is $1/6$. Note how we first did the multiplication term and then summed over all available pure strategy profiles, while multiplying by the utility of each. This is exactly what the expression above does. Recalling that $S = \{(H, H), (H, T), (T, H), (T, T)\}$,

⁹That is, the joint probability equals the product of individual probabilities.

¹⁰In all cases where we shall calculate mixed strategies, the space of pure strategies will be finite so we do not run into measure-theoretic problems.

we can write:

$$\begin{aligned}
U_1(\sigma) &= \sum_{s \in S} \left(\prod_{j=1}^2 \sigma_j(s_j) \right) u_1(s) \\
&= \sigma_1(H)\sigma_2(H)u_1(H,H) + \sigma_1(H)\sigma_2(T)u_1(H,T) \\
&\quad + \sigma_1(T)\sigma_2(H)u_1(T,H) + \sigma_1(T)\sigma_2(T)u_1(T,T) \\
&= (1/3)(1/4)(1) + (1/3)(3/4)(-1) + (2/3)(1/4)(-1) + (2/3)(3/4)(1) \\
&= 1/6.
\end{aligned}$$

Consider now an example from a game that does involve chance moves, like the Card Game, whose strategic form is in Fig. 18 (p. 21). Suppose we wanted to know player 2's expected payoff from the mixed strategy profile $\sigma = \langle (1/3, 1/4, 5/12, 0), (1/3, 2/3) \rangle$. That is, for player 1, $\sigma_1(Rr) = 1/3$, $\sigma_1(Rf) = 1/4$, $\sigma_1(Fr) = 5/12$, and $\sigma_1(Ff) = 0$, whereas for player 2, $\sigma_2(m) = 1/3$ and $\sigma_2(p) = 2/3$. So,

$$\begin{aligned}
U_2(\sigma) &= \sum_{s \in S} \left(\prod_{j=1}^2 \sigma_j(s_j) \right) u_2(s) \\
&= \sigma_1(Rr)\sigma_2(m)u_2(Rr,m) + \sigma_1(Rr)\sigma_2(p)u_2(Rr,p) \\
&\quad + \sigma_1(Rf)\sigma_2(m)u_2(Rf,m) + \sigma_1(Rf)\sigma_2(p)u_2(Rf,p) \\
&= \sigma_1(Fr)\sigma_2(m)u_2(Fr,m) + \sigma_1(Fr)\sigma_2(p)u_2(Fr,p) \\
&\quad + \sigma_1(Ff)\sigma_2(m)u_2(Ff,m) + \sigma_1(Ff)\sigma_2(p)u_2(Ff,p) \\
&= (1/3)(1/3)(0) + (1/3)(2/3)(-1) + (1/4)(1/3)(0.5) + (1/4)(2/3)(-1) \\
&= (5/12)(1/3)(-0.5) + (5/12)(2/3)(0) + (0)(1/3)(0) + (0)(2/3)(0) \\
&= -5/12.
\end{aligned}$$

If you wanted to compute the probability distribution over the outcomes induced by σ , you should get the result in Tab. 38 (p. 34).

	<i>m</i>	<i>p</i>
<i>Rr</i>	1/9	2/9
<i>Rf</i>	1/12	1/6
<i>Fr</i>	5/36	5/18
<i>Ff</i>	0	0

Figure 38: The probability distribution over outcomes for Fig. 18 (p. 21) induced by σ .

As the last example showed, there is no requirement that a mixed strategy puts positive probabilities on all available pure strategies. The **support** of a mixed strategy σ_i is the set of strategies to which σ_i assigns positive probability. This means that we can think of a pure strategy s_i as a **degenerate mixed strategy** that assigns probability 1 to s_i and 0 to all remaining pure strategies (i.e. the support of a degenerate mixed strategy consists of a single pure strategy). A **completely mixed strategy** assigns positive probability to every strategy in S_i .¹¹

¹¹Completely mixed strategies are important because a strategy profile of completely mixed strategies assigns positive probability to every possible outcome in the game. As we shall see later, the fundamental solution

As mentioned in the previous section, we can further reduce some strategic form games. Consider the game in Fig. 36 (p. 32). Although no other pure strategies are payoff-equivalent, the strategy (b, e) is redundant in an important sense. Suppose player 1 were to choose between the strategy A and (b, d) with a flip of a fair coin. The resulting randomized strategy can be denoted with $\sigma = 0.5[A] + 0.5[b, d]$, and would give the expected payoffs:

$$U(\sigma, x) = (0.5)(6, 0) + (0.5)(0, 8) = (3, 4)$$

$$U(\sigma, y) = (0.5)(6, 0) + (0.5)(8, 0) = (7, 0).$$

In other words, we could get the payoffs from (b, e) from randomizing between the strategies A and (b, d) . We formalize this notion as follows:

DEFINITION 9. A strategy $\hat{s}_i \in S_i$ is **randomly redundant** if and only if there exists a mixed strategy $\sigma_i \in \Sigma_i$ such that $\sigma_i(\hat{s}_i) = 0$ and

$$U_j(\hat{s}_i, s_{-i}) = \sum_{s_i \in S_i} \sigma_i(s_i) u_j(s_i, s_{-i}) \quad \forall s_{-i} \in S_{-i}, \quad \forall j \in \mathcal{I}.$$

That is each player's payoffs from the profiles involving \hat{s}_i can be expressed as the expected payoffs from a mixed strategy for player i that does not have \hat{s}_i in its support. In other words, \hat{s}_i is randomly redundant if there is some way for player i to mix his other pure strategies such that no matter what combination of strategies the other players choose, every player would get the same expected payoff whether i uses \hat{s}_i or mixes in this way.

DEFINITION 10. The **fully reduced normal form** of an extensive form game Γ is obtained from the purely reduced representation of Γ by eliminating all randomly redundant strategies.

The fully reduced normal form representation of the extensive form game from Fig. 34 (p. 31) (whose purely reduced normal form is in Fig. 36 (p. 32)) is given in Fig. 39 (p. 35).

		Player 2	
		x	y
Player 1	A	6, 0	6, 0
	(b, c)	8, 0	0, 8
	(b, d)	0, 8	8, 0

Figure 39: The Fully Reduced Strategic Form of the Game from Fig. 34 (p. 31).

Consider the example in Fig. 40 (p. 36): how are we to approach something like this to decide whether there are any strategies that are randomly redundant? Obviously, the only possibilities must involve strategies for player 1, but which one(s)? We can begin by simple elimination by asking whether any two strategies can be mixed to eliminate a third one. We cannot eliminate A by any mixture of two or more of the remaining three pure strategies because player 1's payoff against L is negative if he plays A and non-negative otherwise. Since any mixture of B , C , and D must yield a non-negative payoff against A as well, there is no way to match the payoff from A . It is also impossible to eliminate B with any combination

concept (Nash equilibrium) will not produce any odd results in that situation. Problems with Nash equilibrium (in the sense of unreasonable predictions about optimal behavior) might only occur when the strategy profile induces zero probability for one or more of the possible outcomes.

of the other three strategies: player 1's payoff against L is 3, which is strictly greater than any of the other payoffs he could get against L . This means that any mixture of A , C , and D must yield player 1 an expected payoff strictly less than 3, so they cannot match B . It is also impossible to eliminate C ; this time, note that player 2's payoff against C when she plays L is -1 , which is strictly less than any of her payoffs against the other three strategies for player 1. This means that any mixture of A , B , and D must give player 2 a payoff strictly better than -1 when she chooses L , so it will not be possible to match C .

		Player 2	
		L	R
Player 1	A	$-1, 0$	$-\frac{1}{3}, \frac{1}{2}$
	B	$3, \frac{1}{2}$	$-1, \frac{9}{8}$
	C	$0, -1$	$1, 0$
	D	$\frac{3}{4}, -\frac{1}{4}$	$0, \frac{1}{2}$

Figure 40: Less Obvious Example.

All of this means that if there is any randomly redundant strategy for player 1, it would have to be D . What mixture of some combination of A , B , and C can work? First, note that it cannot be a mixture between A and B by themselves: player 2's payoff from L would be non-negative and she must get $-\frac{1}{4}$ to match her payoff against D . Can it be a mixture between B and C by themselves? Looking at player 1's payoffs against R , we can see that he gets -1 from B and 1 from C . There is only one way to match the payoff of 0 he obtains from D : mix B and C with equal probabilities. But then player 2's payoff against the mixture would be $\frac{9}{16}$ when she chooses R , which does not match her payoff of $\frac{1}{2}$ against D . Hence, it is not possible to eliminate D with a mixture of B and C alone.

This leaves us with just one more possibility: mix A , B , and C to eliminate D . If D is randomly-redundant, then the following system of equations must have a unique solution:

$$\begin{aligned} -\sigma_1(A) + 3\sigma_1(B) &= \frac{3}{4} \\ -\frac{1}{3}\sigma_1(A) - \sigma_1(B) + \sigma_1(C) &= 0 \\ \frac{1}{2}\sigma_1(B) - \sigma_1(C) &= -\frac{1}{4} \\ \frac{1}{2}\sigma_1(A) + \frac{9}{8}\sigma_1(B) &= \frac{1}{2}, \end{aligned}$$

such that $\sigma_1(A) + \sigma_1(B) + \sigma_1(C) = 1$ and $\sigma_1(a) \in (0, 1)$ for all $a \in \{A, B, C\}$. From the last equation, we obtain $\sigma_1(A) = 1 - \frac{9}{4}\sigma_1(B)$. Plugging this into the first equation and multiplying both sides by 4 then gives us $-4 + 9\sigma_1(B) + 12\sigma_1(B) = 3$, which then yields the solution $\sigma_1(B) = \frac{7}{21} = \frac{1}{3}$. Plugging this into the third equation yields $\frac{1}{6} - \sigma_1(C) = -\frac{1}{4}$, so $\sigma_1(C) = \frac{5}{12}$. Finally, plugging these two into the second equation reduces it to $-\frac{1}{3}\sigma_1(A) - \frac{1}{3} + \frac{5}{12} = 0$, which implies $\sigma_1(A) = \frac{1}{4}$. Of course, since we know the probabilities must sum up to 1, we could have just computed $\sigma_1(A) = 1 - \sigma_1(B) - \sigma_1(C)$ to obtain the same result. This way, however, we can verify that the sum is unity, so we have not messed up any of our calculations. We now have the mixed strategy $\sigma_1 = (\frac{1}{4}, \frac{1}{3}, \frac{5}{12}, 0)$ which yields the same expected payoffs to either player as D does against L , and the same expected payoffs to either player as D does against R . Hence, D is randomly redundant and we can safely eliminate it without losing anything in the process.

One question you may have at this point is what happens if there are more than one randomly-redundant strategies: would it matter which one gets eliminated first? What if we

use some pure strategy to eliminate another and then eliminate that pure strategy itself: does that mean we have to restore the one we originally eliminated or is it possible to eliminate it without using that pure strategy? As it turns out, it does not matter which order you do the elimination in: if you can eliminate a pure strategy d by a mixed strategy that has s, s' , and s'' in its support and then s itself gets eliminated by another mixed strategy with only s' and s'' in its support, then it is possible to eliminate d with a mixed strategy that only has s' and s'' in its support. Let's see an example that illustrates this, so consider Fig. 41 (p. 37).

	L	R
A	1, 2	-2, 0
B	0, 3	$-1/2, 2$
C	-1, 4	1, 4
D	$-1/4, 13/4$	$-1/8, 5/2$

	L	R
A	1, 2	-2, 0
B	0, 3	$-1/2, 2$
C	-1, 4	1, 4

	L	R
A	1, 2	-2, 0
C	-1, 4	1, 4

Figure 41: Order of Elimination Does Not Matter.

The mixed strategy $\sigma = (1/4, 1/4, 1/2, 0)$ makes D randomly-redundant in the original game on the left, producing the reduced normal form in the middle. But then $\sigma' = (1/2, 0, 1/2)$ makes B randomly-redundant in that intermediate form, producing the fully reduced form on the right. The question then is: since we used B to eliminate D in the first step, would we still be able to eliminate D now that B itself is gone? That is, do we need B to keep D out? The claim is that since B can be eliminated by A and C , then it should be possible to eliminate D with only these two strategies as well. What is the appropriate mixture then? Since mixing A and C with equal weights eliminates B , let's distribute the weight on B in the original σ evenly to A and C and check if the result can eliminate D . That is, add $1/8$ to the probabilities σ assigns to A and C to consider $\sigma'' = (3/8, 0, 5/8, 0)$ in the original game. It is straightforward to verify that this strategy makes D randomly redundant: against L player's expected payoff is $3/8 - 5/8 = -2/8 = -1/4$ and player 2's expected payoff is $3/8(2) + 5/8(4) = 26/8 = 13/4$; analogously, against R , player 1's expected payoff is $3/8(-2) + 5/8 = -1/8$, and player 2's expected payoff is $5/8(4) = 5/2$. This means that we can use σ'' to eliminate D and then σ' to eliminate B , yielding the same fully reduced form.

It is sometimes quite tricky to identify randomly redundant strategies. It may be worth your while to try anyway because by reducing the number of strategies to consider for the analysis, you will greatly simplify your task (you will see what I mean when we begin solving the games next time). Unless we explicitly state otherwise, we shall take the *reduced strategic form representation* to mean the fully reduced form.

You might wonder why we are eliminating redundant strategies: after all, the ones we remove from considerations do, in fact, specify ways to play the game and reach possibly different outcomes. For instance, in the reduced strategic form in Fig. 41 (p. 37), there are no outcomes $\langle D, L \rangle$ or $\langle D, R \rangle$, which were both available in the original specification. Aren't we losing something when we do not consider them? If there are several redundant strategies, does it not matter which ones we eliminate? The answer is that for the *analysis* of the game, it will not matter. When we find solutions that involve a strategy that has other payoff-equivalent ones in the original game, then we will immediately know that the original game has more solutions: we would obtain those by replacing the strategy with the payoff-equivalent ones we eliminated. Thus, suppose for instance that in the reduced form we found solutions in which A and C are played with probability $1/2$ each. Because we know that this mixed strategy is payoff equivalent to the pure strategy B , we immediately know that

there are solutions in which player 2's strategy is the same but player 1 plays B instead of that particular mixed strategy. If, however, the solution involved A and C with some other probabilities, then there will be no solutions that involve B . Thus, when we want to provide a substantive interpretation for the solution, we have to remember the payoff-equivalent strategies.

6 Mixed and Behavior Strategies in Extensive Form Games

Unlike strategic form games, extensive form games admit two distinct types of randomization: a player can either randomize over his pure strategies or he can randomize over the actions at each of his information sets.

As in the normal form game, a **mixed strategy** for player i is a probability distribution over i 's set of pure strategies.¹² That is, a mixed strategy specifies the probabilities with which pure strategies are played but each pure strategy specifies a definite action at each information set.

The other type of randomizing strategy is the **behavior strategy**, which specifies a probability distribution over actions at each information set. *These distributions are independent.* That is, a behavior strategy specifies the probabilities with which actions are chosen at every information set. Thus, a pure strategy is a special kind of behavior strategy where the distribution at each information set is degenerate.

To help illustrate the difference between the two types of randomization, Luce and Raiffa (1957) offer the following analogy: A pure strategy is a book of instructions, where each page tells how to play at a particular information set. The space of pure strategies is a library of these books. A mixed strategy is a probability distribution over this library (i.e. it specifies the probability with which books are chosen). A behavior strategy is a single book where each page prescribes a random action. Thus, a player may randomly select a pure strategy or he might plan a set of randomizations, one for every point at which he has to take action.

An example may be helpful. Consider the game in Fig. 14 (p. 15) and recall that player 1 has four pure strategies: (AE) , (AF) , (BE) , and (BF) . A mixed strategy is a probability distribution over these four strategies. For example, a mixed strategy $\sigma = (1/4, 1/4, 1/4, 1/4)$ specifies that player 1 will play each of his pure strategies with equal probability of $1/4$. Another mixed strategy might be $\sigma = (1/3, 0, 1/6, 1/2)$, which specifies that player 1 should play AE with probability $1/3$, AF with probability 0, BE with probability $1/6$, and BF with probability $1/2$. You can see the close correspondence with mixed strategies in normal form games.

On the other hand, a behavior strategy for player 1 would specify probabilities for actions at all information sets. Because player 1 has two information sets, the strategy must specify two probability distributions, one for each information set. For example, $\beta = (1/4, 1/4)$ means that player 1 will choose A at his first information set with probability $1/4$ (and choose B with complementary probability $3/4$), and he will choose E with probability $1/4$ at the second information set. Another behavior strategy might be $\beta = (0, 1/2)$, which specifies that player 1 should choose B with probability 1 at the first information set and play E and F with equal probability at the second information set. Just like a pure strategy will have as many elements as there are information sets at which the player must move, the behavior strategy will also have as many elements as there are information sets. The difference is that the pure

¹²In extensive form games of perfect information little is added by considering mixed strategies. We will not see them until later, when we learn about games of incomplete information.

strategy will prescribe a certain action for each information set whereas the behavior strategy prescribes a probability distribution over the actions at this set. (Of course, the number of elements in a mixed strategy equals the number of pure strategies.) As we noted, a pure strategy is a behavior strategy with degenerate distributions at each information set. So, for example, the pure strategy BE is the behavior strategy $\beta = (0, 1)$ just as it is the degenerate mixed strategy $\sigma = (0, 0, 1, 0)$.

As you probably already suspect, the two types of randomizing strategies are closely related. We shall call two strategies *equivalent* if they induce the same probability distributions over outcomes for all strategies of the opponents.¹³ Intuitively, two strategies are equivalent if they have the same consequences regardless of what the other players do.

6.1 Mixed Strategy Equivalent to a Behavior Strategy

Let's see how we can generate a mixed strategy that is equivalent to some arbitrary behavior strategy β_i for player i . Let $\beta_i(h_i)(a_i)$ denote the probability with which action $a_i \in A_i(h_i)$ is taken (that is the probability with which an action is chosen from the set of actions available after history h_i). Let $s_i(h_i)$ denote the action specified by the pure strategy s_i at the information set h_i (and so s_i specifies one action for all information sets where player i gets to move). Define the mixed strategy σ_i to assign the following probability to each pure strategy s_i :

$$\sigma_i(s_i) = \prod_{h_i \in H} \beta_i(h_i)(s_i(h_i)). \quad (1)$$

That is, the probability with which the pure strategy is chosen is simply the product of probabilities assigned by the behavior strategy to the action the pure strategy prescribes at each information set. Note that we made use of the assumption that the behavior randomizations are independent across information sets.¹⁴

Let's ask ourselves about the intuition behind this. Essentially, a pure strategy, s_i , gives a "path" of play through the game: given what other players are doing, this strategy tells i what to choose at each of his information sets until the game tree reaches a terminal node. This means that σ_i would have to assign to that "path" a probability that equals the probabilities with which each of its separate components is taken by i 's choice. Since β_i gives the probability of the action prescribed by s_i for each information set, the probability of the entire "path" is just the product of the probabilities that i picks the relevant actions that constitute that path.

Consider the (Little Horsey) game in Fig. 14 (p. 15). A behavior strategy for player 1 has two elements, a probability distribution over his two actions $\{A, B\}$ at his first information set, and another probability distribution over the actions $\{E, F\}$ at his second information set. Consider some fixed (possibly mixed) strategy for player 2, σ_2 such that $\sigma_2(d) > 0$, and consider the outcome after history (B, d, F) . Denote this outcome by z_4 . The only pure strategy for player 1 that can produce this with positive probability is $s_1 = (B, F)$. That is $\Pr[z_4 | s_1] = \sigma_2(d)$. Observe now that a (non-degenerate) behavior strategy will put positive probabilities on both B and F but will not choose them with certainty. Hence, the probability of z_4 will be $\Pr[z_4 | \beta_1] = \beta_1(\emptyset)(B) \times \sigma_2(d) \times \beta_1(Bd)(F)$. That is, it multiplies the probabilities

¹³This is the same concept of equivalence we used when we discussed the reduced normal form representation of extensive games in the previous section.

¹⁴This holds for all games of perfect recall. In games of imperfect recall, it is possible to have behavior strategies that cannot be duplicated by any mixed strategy.

it assigns to the actions specified by s_1 at each information set: $\Pr[z_4|\beta_1] = \beta_1(\emptyset)(s_1(\emptyset)) \times \sigma_2(d) \times \beta_1(Bd)(s_1(Bd))$, where we note that $s_1 = (B, F)$ is, if we were to use the full definition of a pure strategy as a function that takes an information set and returns an action, equivalent to $s_1(\emptyset) = B$ and $s_1(Bd) = F$. Now, a mixed strategy for player 1 can also produce z_4 with positive probability as long as $\sigma_1(BF) > 0$. In particular, since we want σ_1 to produce z_4 with the same probability as β_1 , it must be the case that in that mixed strategy the probability of player 1 choosing *both* B and F at the respective information sets must be the same under σ_1 as it is under β_1 . Under β_1 , we have seen that the probability of choosing B and F is $\beta_1(\emptyset)(B) \times \beta_1(Bd)(F)$, which would give z_4 with probability $\sigma_2(d)$. Since the only way to reach this outcome must involve playing s_1 , the mixed strategy must assign this exact probability to that pure strategy: $\sigma_1(s_1) = \beta(\emptyset)(s_1(\emptyset)) \times \beta_1(Bd)(s_1(B, d))$, that is, exactly as in (1). The probability of reaching z_4 using σ_1 is also $\sigma_2(d)$.

Let's now consider a specific example. To check equivalence, we first need to specify the distribution over outcomes. The Little Horse game in Fig. 14 (p. 15) has four outcomes. Let the probability distribution (z_1, z_2, z_3, z_4) denote the associated probabilities for the outcomes $(1, 1)$, $(-1, 1)$, $(3, 2)$, and $(4, 0)$. Finally, let $\sigma_2(c)$ denote the probability with which player 2 chooses c and $\sigma_2(d) = 1 - \sigma_2(c)$ denote the probability with which she chooses d . The behavior strategy $\beta = ((1/4, 3/4), (1/4, 3/4))$, where player 1 chooses A and E with probability $1/4$, induces the probability distribution over outcomes $(1/4, 3/4\sigma_2(c), 3/16\sigma_2(d), 9/16\sigma_2(d))$. (We obtained the probabilities for z_3 and z_4 by multiplying the the probability of each action specified by the behavior strategy by the probability that the initial action is B . You should verify that the distribution over outcomes is valid: i.e. all probabilities sum to 1.) Now, using our Equation 1, we can define the mixed strategy σ as follows:

$$\begin{aligned}\sigma(AE) &= \beta(\emptyset)(A) \times \beta(Bd)(E) = 1/4 \times 1/4 = 1/16 \\ \sigma(AF) &= \beta(\emptyset)(A) \times \beta(Bd)(F) = 1/4 \times 3/4 = 3/16 \\ \sigma(BE) &= \beta(\emptyset)(B) \times \beta(Bd)(E) = 3/4 \times 1/4 = 3/16 \\ \sigma(BF) &= \beta(\emptyset)(B) \times \beta(Bd)(F) = 3/4 \times 3/4 = 9/16\end{aligned}$$

(We again verify that this is a valid probability distribution by noting that the probabilities all sum to 1.) Is this mixed strategy equivalent to the original behavior strategy? That is, does it induce the same probability over outcomes regardless of what the other player does? The probability of outcome z_1 equals the probability that player 1 chooses A , which he does in two of his strategies, and so it is $\sigma(AE) + \sigma(AF) = 1/4$. The probability of z_2 is the probability that player 1 will choose B , which is $\sigma(BE) + \sigma(BF) = 3/4$, multiplied by the probability that player 2 chooses c . This yields $3/4\sigma_2(c)$. The probability of z_3 is the probability that player 1 chooses both B and E multiplied by the probability that player 2 chooses d , which yields $\sigma(BE)\sigma_2(d) = 3/16\sigma_2(d)$. Finally, the probability of z_4 is the probability that player 1 chooses both B and F , $\sigma(BF)$, multiplied by the probability that player 2 chooses d , which yields $9/16\sigma_2(d)$. To summarize, the probability distribution over outcomes induced by the mixed strategy σ as defined above is $(1/4, 3/4\sigma_2(c), 3/16\sigma_2(d), 9/16\sigma_2(d))$, which is the same as the probability distribution induced by the behavior strategy β . We have now seen how to generate an equivalent mixed strategy from an arbitrary behavior strategy. But there is more to equivalence than this!

6.2 Equivalence of Mixed and Behavior Strategies

An important result is that in a game of perfect recall, mixed and behavior strategies are equivalent.

THEOREM 1 (Kuhn 1953). *In a game of perfect recall,*

- *every behavior strategy is equivalent to every mixed strategy that generates it;*
- *every mixed strategy is equivalent to the unique behavior strategy it generates.*

That is, different mixed strategies can generate the same behavior strategy even though each mixed strategy either generates exactly one behavior strategy or else infinitely many behavior strategies. To make this a bit more concrete, two different mixed strategies can generate the same behavior strategy (we shall see an example below). The first part of the claim is that this behavior strategy is going to be equivalent to each of the two different mixed strategies that generate it. The two mixed strategies are *behaviorally equivalent*.

Further, every mixed strategy has at least one behavioral representation, and it may have many. It may have many if there are information sets that the mixed strategy does not reach with positive probability: In this case it does not matter what probability distribution the behavior strategy specifies for that information set. If, however, the mixed strategy reaches all information sets with positive probability, then it will generate a unique behavior strategy. The second part of the claim states the these will be equivalent.

Finally, note that we can generate a mixed strategy σ_i from a behavior strategy β_i as shown above in (1). In this case, σ_i is the mixed representation of β_i , and they are equivalent. Further, it is not hard to show that if σ_i is the mixed representation of β_i , then β_i is the behavioral representation of σ_i .

To see how the theorem works, let's derive a behavior strategy for some given mixed strategy. Let σ_i be a mixed strategy for player i . For any history h_i , let $R_i(h_i)$ denote the set of player i 's pure strategies that are consistent with h_i . That is, for all $s_i \in R_i(h_i)$, there is a profile s_{-i} for the other players that reaches h_i . We shall call the strategies in $R_i(h_i)$ *consistent* with the history h_i . For example, in the Little Horsey game from Fig. 14 (p. 15), all four pure strategies for player 1 are consistent with his first information set, \emptyset for the simple reason that the initial information set is always reached regardless of what players are going to do from that point on. On the other hand, the information set (Bd) can only be reached for some strategy by player 2 (in this case, d) provided player 1 chooses B at his first information set. There are only two pure strategies that involve such a choice: (BE) and (BF) . Therefore, $R_1(Bd) = \{BE, BF\}$, and neither AE nor AF is consistent with the history Bd .

Now let $\pi_i(h_i)$ be the sum of probabilities according to σ_i of all the pure strategies that are consistent with h_i :

$$\pi_i(h_i) = \sum_{s_i \in R_i(h_i)} \sigma_i(s_i).$$

Intuitively, this is the probability with which the game will reach h_i provided i (and the other players) choose actions consistent with this history. It answers the question: "Suppose all other players use pure strategies that are on the path toward h_i . What is the probability of reaching h_i if player i uses σ_i ?" In our example, $\pi_1(\emptyset) = \sigma_1(AE) + \sigma_1(AF) + \sigma_1(BE) + \sigma_1(BF) = 1$, and $\pi_1(Bd) = \sigma_1(BE) + \sigma_1(BF)$. In either case, we are supposing that player 2

is choosing d in the sense that she is not playing a strategy that would make reaching Bd impossible no matter what player 1 does.

Let $\pi(h_i, a_i)$ denote the sum of probabilities according to σ_i of all pure strategies that are consistent with h_i followed by action $a_i \in A_i(h_i)$. So we have

$$\pi_i(h_i, a_i) = \sum_{s_i \in R_i(h_i) \wedge s_i(h_i) = a_i} \sigma_i(s_i).$$

Intuitively, this is very similar to $\pi_i(h_i)$ except that it asks “What is the probability of reaching h_i and choosing a_i at that information set?” (Again, provided the other players use strategies that do not preclude reaching that point in the game.) In our example, $\pi_1(\emptyset, A) = \sigma_1(AE) + \sigma_1(AF)$ because each of AE and AF is both consistent with the initial history \emptyset and prescribes A as the action at that set. Similarly, $\pi_1(\emptyset, B) = \sigma_1(BE) + \sigma_1(BF)$. At the second information set, we have $\pi_1(Bd, E) = \sigma_1(BE)$ because even though both BE and BF are consistent with this history, only BE involves choosing E at the second information set. Analogously, $\pi_1(Bd, F) = \sigma_1(BF)$.

We now have the two components we need. Observe that $\pi_i(h_i, a_i)$ is the probability of reaching h_i and playing a_i . However, to define $\beta(h_i)(a_i)$, we need to find the probability of playing a_i provided h_i has been reached. This requires us to condition $\pi_i(h_i, a_i)$ on the probability of reaching h_i , which is $\pi_i(h_i)$. If σ_i assigns positive probability to some $s_i \in R_i(h_i)$, define the probability that the behavior strategy β_i assigns to $a_i \in A_i(h_i)$ as the probability of taking action a_i conditional on reaching the information set h_i :

$$\beta_i(h_i)(a_i) = \frac{\pi_i(h_i, a_i)}{\pi_i(h_i)}.$$

Intuitively, the probability of picking a_i at the information set h_i is the probability of reaching h_i and picking a_i conditioned on the probability of reaching h_i . In our example, $\beta_1(\emptyset)(A) = \sigma_1(AE) + \sigma_1(AF)$ and $\beta_1(\emptyset, B) = \sigma_1(BE) + \sigma_1(BF)$. At the second information set, $\beta_1(Bd, E) = \sigma_1(BE) / [\sigma_1(BE) + \sigma_1(BF)]$; that is, the probability the behavior strategy must assign to the action E is the probability σ_1 assigns to it conditional on reaching this information set if σ_1 is followed. Finally, $\beta_1(Bd, F) = \sigma_1(BF) / [\sigma_1(BE) + \sigma_1(BF)]$. How we define $\beta_i(h_i)(a_i)$ if $\pi_i(h_i) = 0$ is immaterial.¹⁵ One possible specification is to assign the probabilities given by the mixed strategy: $\beta_i(h_i)(a_i) = \sum_{s_i(h_i) = a_i} \sigma_i(s_i)$, but anything will do. In either case, the $\beta_i(\cdot)(\cdot)$ are nonnegative, and

$$\sum_{a_i \in A_i(h_i)} \beta_i(h_i)(a_i) = 1,$$

because each s_i specifies an action for player i at the information set h_i . In other words, β_i specifies a valid distribution for each information set h_i . If $\pi_i(h_i) > 0$ for all histories, then the mixed strategy will generate a unique behavior strategy.

Let’s look at concrete example. Consider the game in Fig. 42 (p. 43). We want to find the behavior strategy for player 1 that is equivalent to his mixed strategy in which he plays (B, R) with probability 0.4, (B, L) with probability 0.1, and (A, L) with probability 0.5.

We have $\sigma_1(B, R) = 0.4$, $\sigma_1(B, L) = 0.1$, $\sigma_1(A, L) = 0.5$, and (since the mixed strategy is a probability distribution), $\sigma_1(A, R) = 0$. Player 1 has two information sets: one after the \emptyset history, and another after the histories (A, M) and (A, D) . The behavior strategy will thus specify two probability distributions, one for each information set.

¹⁵Since h_i cannot be reached under σ_i , the behavior strategies at h_i are arbitrary in the same sense that Bayes’ Rule does not determine posterior probabilities after 0-probability events.




Figure 42: A Game for Kuhn's Theorem, I.

Since $h_1 = \emptyset$ is the initial history, all pure strategies are consistent with it. (This is trivially true: there is no pure strategy for player i such that this history cannot be reached.) Thus, $R_1(h_1) = \{(A, L), (A, R), (B, L), (B, R)\}$, which also means $\pi_1(h_1) = 1$. Since there are two possible actions player 1 can take at h_1 , we must calculate $\pi_1(h_1, A)$ and $\pi_1(h_1, B)$. There are two pure strategies s_1 such that $s_1 \in R_1(h_1) \wedge s_1(h_1) = A$, and these are (A, L) and (A, R) . Therefore, $\pi_1(h_1, A) = \sigma_1(A, L) + \sigma_1(A, R) = 0.5$. Also, there are two pure strategies such that $s_1 \in R_1(h_1) \wedge s_1(h_1) = B$, and these are (B, L) and (B, R) . This means $\pi_1(h_1, B) = \sigma_1(B, L) + \sigma_1(B, R) = 0.5$. We now have $\beta_1(h_1)(A) = \pi_1(h_1, A)/\pi_1(h_1) = 0.5/1 = 0.5$ and also $\beta_1(h_1)(B) = \pi_1(h_1, B)/\pi_1(h_1) = 0.5/1 = 0.5$.¹⁶ So, $\beta_1(h_1)(A) = \beta_1(h_1)(B) = 0.5$.

Now consider $h_2 = \{(A, M), (A, D)\}$. The only pure strategies for player 1 that are consistent with this history are the ones that specify A for the move at the first information set. (That is, there exists no strategy for player 2 such that h_2 is reached if player 1 chooses B at the first information set.) Therefore, $R_1(h_2) = \{(A, L), (A, R)\}$, which means that $\pi_1(h_2) = \sigma_1(A, L) + \sigma_1(A, R) = 0.5$. Since player 1 has two possible actions at h_2 , we must also calculate $\pi_1(h_2, L)$ and $\pi_1(h_2, R)$. There is only one pure strategy such that $s_1 \in R_1(h_2) \wedge s_1(h_2) = L$, and it is (A, L) . Therefore, $\pi_1(h_2, L) = \sigma_1(A, L) = 0.5$. Also, there is only one pure strategy such that $s_1 \in R_1(h_2) \wedge s_1(h_2) = R$, and it is (A, R) , which means $\pi_1(h_2, R) = \sigma_1(A, R) = 0$. We now have $\beta_1(h_2)(L) = \pi_1(h_2, L)/\pi_1(h_2) = 0.5/0.5 = 1$, and we also have $\beta_1(h_2)(R) = \pi_1(h_2, R)/\pi_1(h_2) = 0/0.5 = 0$.¹⁷

We conclude that the mixed strategy σ_1 has an equivalent behavior strategy β_1 , which is as follows:

$$\begin{aligned}\beta_1(h_1)(A) &= 0.5 \\ \beta_1(h_1)(B) &= 0.5 \\ \beta_1(h_2)(L) &= 1 \\ \beta_1(h_2)(R) &= 0\end{aligned}$$

Let's check the equivalence claim. Let σ_2 denote a mixed strategy for player 2. Using the

¹⁶We verify that $\beta_1(h_1)(A) = 1 - \beta_1(h_1)(B)$, which is indeed the case.

¹⁷We again verify that the distribution is valid, which it is because $\beta_1(h_2)(L) + \beta_1(h_2)(R) = 1$.

mixed strategy σ_1 , the probabilities of reaching the outcomes are as follows:

$$\begin{aligned} z_1 &: [\sigma_1(A, L) + \sigma_1(A, R)]\sigma_2(U) = 0.5\sigma_2(U) \\ z_2 &: \sigma_1(A, L)\sigma_2(M) = 0.5\sigma_2(M) \\ z_3 &: \sigma_1(A, R)\sigma_2(M) = 0 \\ z_4 &: \sigma_1(A, L)\sigma_2(D) = 0.5\sigma_2(D) \\ z_5 &: \sigma_1(A, R)\sigma_2(D) = 0 \\ z_6 &: \sigma_1(B, L) + \sigma_1(B, R) = 0.5 \end{aligned}$$

The distribution over outcomes using σ_1 is then $(0.5\sigma_2(U), 0.5\sigma_2(M), 0, 0.5\sigma_2(D), 0, 0.5)$.

Using the behavior strategy β_1 , the probabilities of reaching the outcomes are as follows.

$$\begin{aligned} z_1 &: \beta_1(h_1)(A)\sigma_2(U) = 0.5\sigma_2(U) \\ z_2 &: \beta_1(h_1)(A)\sigma_2(M)\beta_1(h_2)(L) = (0.5)\sigma_2(M)(1) = 0.5\sigma_2(M) \\ z_3 &: \beta_1(h_1)(A)\sigma_2(M)\beta_1(h_2)(R) = (0.5)\sigma_2(M)(0) = 0 \\ z_4 &: \beta_1(h_1)(A)\sigma_2(D)\beta_1(h_2)(L) = (0.5)\sigma_2(D)(1) = 0.5\sigma_2(D) \\ z_5 &: \beta_1(h_1)(A)\sigma_2(D)\beta_1(h_2)(R) = (0.5)\sigma_2(D)(0) = 0 \\ z_6 &: \beta_1(h_1)(B) = 0.5 \end{aligned}$$

This yields the distribution over outcomes $(0.5\sigma_2(U), 0.5\sigma_2(M), 0, 0.5\sigma_2(D), 0, 0.5)$ that is the same as the one given by the mixed strategy. Therefore, we have shown that σ_1 and β_1 are equivalent.

6.2.1 A Mixed Strategy Can Generate Many Behavior Strategies

Now let's illustrate the claim that a mixed strategy may generate more than one behavior strategy. Consider the same game and suppose $\sigma_1(A, L) = \sigma_1(A, R) = 0$, $\sigma_1(B, L) = 0.5$, and $\sigma_1(B, R) = 0.5$. As before, we have $R_1(h_1) = \{(A, L), (A, R), (B, L), (B, R)\}$, and $\pi_1(h_1) = 1$. Further, we have $\pi_1(h_1, A) = 0$ (because the mixed strategy assigns probability zero to all pure strategies with $s_1(h_1) = A$), and $\pi_1(h_1, B) = 1$. Thus, we get $\beta_1(h_1)(A) = 0$ and $\beta_1(h_1)(B) = 1$.

We now have to specify the probability distribution for the information set following $h_2 = \{(A, M), (A, D)\}$. Note that $R_1(h_2) = \{(A, L), (A, R)\}$ and $\pi_1(h_2) = 0$. Further, $\pi_1(h_2, L) = \sigma_1(A, L) = 0$ and $\pi_1(h_2, R) = \sigma_1(A, R) = 0$. Hence, we cannot use the conditional formula to define $\beta_1(h_2)(L)$. As noted before, in this case we could use any probability distribution, so let's say $\beta_1(h_2)(L) = x$ and $\beta_1(h_2)(R) = 1 - x$, with $x \in [0, 1]$. Clearly, there is an infinite number of possible specifications here.

Let's check equivalence. Under the mixed strategy, the probability distribution over outcomes is:

$$\begin{aligned} z_1 &: [\sigma_1(A, L) + \sigma_1(A, R)]\sigma_2(U) = 0 \\ z_2 &: \sigma_1(A, L)\sigma_2(M) = 0 \\ z_3 &: \sigma_1(A, R)\sigma_2(M) = 0 \\ z_4 &: \sigma_1(A, L)\sigma_2(D) = 0 \\ z_5 &: \sigma_1(A, R)\sigma_2(D) = 0 \\ z_6 &: \sigma_1(B, L) + \sigma_1(B, R) = 1. \end{aligned}$$

Under the behavior strategy, the probability distribution is:

$$\begin{aligned}
z_1 &: \beta_1(h_1)(A)\sigma_2(U) = 0 \\
z_2 &: \beta_1(h_1)(A)\sigma_2(M)\beta_1(h_2)(L) = (0)\sigma_2(M)x = 0 \\
z_3 &: \beta_1(h_1)(A)\sigma_2(M)\beta_1(h_2)(R) = (0)\sigma_2(M)(1-x) = 0 \\
z_4 &: \beta_1(h_1)(A)\sigma_2(D)\beta_1(h_2)(L) = (0)\sigma_2(D)x = 0 \\
z_5 &: \beta_1(h_1)(A)\sigma_2(D)\beta_1(h_2)(R) = (0)\sigma_2(D)(1-x) = 0 \\
z_6 &: \beta_1(h_1)(B) = 1.
\end{aligned}$$

That is, the two distributions are the same. Note that this holds for any value of x we might have chosen. Thus, one mixed strategy can generate more than one behavior strategy. It should be obvious, however, that if the mixed strategy reaches all information sets with positive probability, then it must necessarily generate a unique behavior strategy. Hence, a mixed strategy either generates a unique behavior strategy or else generates an infinite number of behavior strategies.

6.2.2 Different Mixed Strategies Can Generate the Same Behavior Strategy

Now let's illustrate the claim that different mixed strategies can generate the same behavioral strategy. Consider the game in Fig. 43 (p. 45). Let h_1 denote the history following action U by player 1, let h_2 denote the history following D . Since there are two information sets, with two actions at each, player 2 has four pure strategies: (A, C) , (A, D) , (B, C) , and (B, D) .




Figure 43: A Game for Kuhn's Theorem, II.

Now consider two mixed strategies $\sigma_2 = (1/4, 1/4, 1/4, 1/4)$ and $\hat{\sigma}_2 = (1/2, 0, 0, 1/2)$. Both of these generate the behavior strategy β_2 , where $\beta_2(h_1)(A) = \beta_2(h_1)(B) = 1/2$ and $\beta_2(h_2)(C) = \beta_2(h_2)(D) = 1/2$.¹⁸ To see that σ_2 , $\hat{\sigma}_2$, and β_2 are equivalent, note that they all yield the same distribution over the terminal nodes for any arbitrary mixed strategy for player 1. For example, the probability of reaching z_1 equals $\sigma_1(U)/2$ regardless of whether we calculate it under σ_2 , where it equals $\sigma_1(U)[\sigma_2(A, C) + \sigma_2(A, D)]$, or under $\hat{\sigma}_2$, where it equals $\sigma_1(U)[\hat{\sigma}_2(A, C) + \hat{\sigma}_2(A, D)]$, or under β_2 , where it equals $\sigma_1(U)\beta_2(h_1)(A)$. As you probably

¹⁸You should verify this. In our notation, $R_2(h_1) = R_2(h_2) = \{AC, AD, BC, BD\}$. That is, all strategies for player 2 are consistent with these histories. This is trivially true because she has no move to determine which of these histories is reached. We then calculate the probability associated with each history, which, given that all strategies are consistent with it, is simply $\pi_2(h_1) = \sum_{s_2 \in R_2(h_1)} \sigma_2(s_2) = 1$. Next, we calculate the probability of taking action A after h_1 : $\pi(h_1, A) = \sum_{s_2 \in R_2(h_1) \wedge s_2(h_1)=A} \sigma_2(s_2) = \sigma_2(AC) + \sigma_2(AD) = 0.5$. Finally, we calculate the behavior strategy $\beta_2(h_1)(A) = \pi_2(h_1, A)/\pi_2(h_1) = (0.5)/(1) = 0.5$. We can generate the other strategy in a similar way.

already see, there will be an infinite number of mixed strategies that generate this behavior strategy: All σ_2 such that $\sigma_2(A, C) + \sigma_2(A, D) = 1/2$ and $\sigma_2(A, C) + \sigma_2(B, C) = 1/2$ will do that.

Although it is important to distinguish between the two types of probabilistic strategies, in practice we shall use behavior strategies throughout the rest of this class. Because it is cumbersome to refer to them as such all the time, whenever we refer to a mixed strategy of an extensive form game, we shall always mean a behavior strategy (unless explicitly noted otherwise). To this end, we shall also retain our σ -notation for mixed strategies: Let $\sigma_i(a_i|h_i)$ denote the probability with which player i chooses action a_i at the information set h_i .