

Game Theory: Elements of Basic Models.

Branislav L. Slantchev

Department of Political Science, University of California – San Diego

April 19, 2005

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	8
2.4	Examples of Games in Extensive Form	8
3	Pure Strategies	12
4	The Strategic (Normal) Form	16
4.1	Examples of Converting Extensive to Strategic Form	20
4.1.1	Several Chance Moves	20
4.1.2	Three Players	21
4.2	Examples in Strategic Form	23
4.3	Reduced Strategic Form	26
5	Mixed Strategies in Strategic Form Games	28
6	Mixed and Behavior Strategies in Extensive Form Games	30

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, $A_1(b) = A_1(c) = \{F, R\}$. We represent the actions available at a decision node with branches emanating from that node, as shown in Fig. 2 (p. 3).

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 generalizations of the idea of histories because they summarize not only past play, but also 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. Let’s assume that utilities are linear in the amount of money received, or $u(x) = x$. 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 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 this 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".

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 x 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 generalizations of 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, 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. To emphasize this, I have labeled 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).

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

1. A set of players denoted by $i \in I$, with $I = \{N, 1, 2, \dots\}$, with N representing the pseudo-player Nature;
2. A tree, $(X, >)$, which is a finite collection of nodes $x \in X$ endowed with the precedence relation $>$, where $x > x'$ means “ x is before x' .” This relation is transitive and asymmetric, and thus constitutes a *partial order*.² This rules out cycles where the game may go from node x to a node x' , from x' to node x'' , and from x'' back to x .³ In addition, we require that each node x has exactly one immediate predecessor, that is, one node $x' > x$ such that $x'' > x$ and $x'' \neq x'$ implies $x' > x''$ or $x'' > x'$. 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 $z \in Z$ 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 $u_i : Z \rightarrow \mathbb{R}$, and $u = (u_1(\cdot), \dots, u_I(\cdot))$ is the collection of these functions, one for each player.

²It 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.

³To see this, suppose we constructed a game such that $x > x' > x'' > x$. By transitivity, $x'' > x'$, but since we already have $x' > x''$, this violates asymmetry.

4. A map $\iota : \mathcal{X} \rightarrow \mathcal{I}$, 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. 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. 7) illustrate some cases of imperfect recall that this requirement eliminates.




Figure 4: Games of Imperfect Recall.

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.

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, 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.

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. 9) 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. 9) 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).

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 worker's effort is unobserved, 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.

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 Nature's choice, about whether player 1 is actually playing, 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. 9), a), otherwise player 2 goes first (the play Fig. 5 (p. 9), b). This is shown in Fig. 7 (p. 10).




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.




Figure 8: A Game of Uncertainty, Variant B1.

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 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 it receives. This is shown in Fig. 11 (p. 12).




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. 12).




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

The game tree in Fig. 13 (p. 12) is equivalent to the tree in Fig. 12 (p. 12) (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 $i \in I$ 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. 12). 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, (yyn) 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):

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

Remember that a strategy is a contingent plan of action. For example, the strategy (nny) 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 4 Outcomes.

Let's examine Fig. 14 (p. 14) 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 $A_1(B, d) = \{E, F\}$. 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: $Z = \{A, (B, c), (B, d, E), (B, d, F)\}$.

Now, 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. 14), 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.⁴

Let's specify the strategies for the game in Fig. 15 (p. 15). 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

⁴It will soon become clear why we must specify strategies in this seemingly redundant way. The intuition is that what the other player chooses to do depends on what player 1 might do after his action, which in turn determines what player 1 might do initially. Thus, a strategy must specify actions at both information sets if we are to analyze the interactive component of decision-making in the game.

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:

$$S_1 = \{U, D\}$$

$$S_2 = \{AC, AE, BC, BE\}$$

$$S_3 = \{RP, RQ, TP, TQ\}$$

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.




Figure 15: An Extensive Form Game with Three Players.

Consider the game in Fig. 16 (p. 15). 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 $A_1(A) = \{W, Z\}$.




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

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 $A_2(B) = A_2(C) =$

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

$$\begin{aligned} S_1 &= \{AW, AZ, BW, BZ, CW, CZ\} \\ S_2 &= \{X, Y\} \end{aligned}$$

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, 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. 15), this calculation is $\#S_1 = \#(A_1(\emptyset)) \times \#(A_1(A)) = 3 \times 2 = 6$, while for the game in Fig. 11 (p. 12), 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.

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 is an outcome of the game.

Some people define players' preference orderings over strategy profiles, but I find this confusing. 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:

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

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

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

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. 17 (p. 18).

		Player 2	
		m	p
Player 1	Rr	0, 0	1, -1
	Rf	-0.5, 0.5	1, -1
	Fr	0.5, -0.5	0, 0
	Ff	0, 0	0, 0

Figure 17: The Strategic Form of the Game from Fig. 3 (p. 5).

The strategic game in Fig. 17 (p. 18) 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 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 $x \in \mathcal{Z}$, let $u_i(x)$ be player i 's payoff from outcome x . Then, for any strategy profile $s \in S$ and any $i \in \mathcal{I}$, let $U_i(s)$ be:

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

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. 18 (p. 19).




Figure 18: Two Simple Games.

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

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

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

		Player 2	
		U	D
Player 1	L	2, 2	4, 0
	R	1, 0	3, 1

Figure 19: The Strategic Form of the Game from Fig. 18 (p. 19) (b).

The situation in Fig. 18 (p. 19) (a) is very different. Although player 1 still has two pure strategies, $S_1 = \{L, R\}$, player 2 can condition her choice on player 1's. She has two information sets, and her strategy must specify two actions: $a_L \in A_2(L) = \{U, D\}$ is the choice after player 1 chooses L, and $a_R \in A_2(R) = \{u, d\}$ is the choice after player 1 chooses R. We shall write player 2's strategy as the ordered pair (a_L, a_R) . Hence, the strategy set for player 2 consists of four pure strategies; $S_2 = \{(U, u), (U, d), (D, u), (D, d)\}$. The game now has eight strategy profiles:

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

Because there are no moves by chance, there is no need to transform the utility functions, so the strategic form is given in Fig. 20 (p. 19).

		Player 2			
		(U, u)	(U, d)	(D, u)	(D, d)
Player 1	L	2, 2	2, 2	4, 0	4, 0
	R	1, 0	3, 1	1, 0	3, 1

Figure 20: The Strategic Form of the Game from Fig. 18 (p. 19) (a).

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. 12), 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

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

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. 21 (p. 20).

		Player 2							
		<i>yyy</i>	<i>yy n</i>	<i>y n n</i>	<i>y n y</i>	<i>n y y</i>	<i>n y n</i>	<i>n n y</i>	<i>n n n</i>
Player 1	(2, 0)	2,0	2,0	2,0	2,0	0,0	0,0	0,0	0,0
	(1, 1)	1,1	1,1	0,0	0,0	1,1	1,1	0,0	0,0
	(0, 2)	0,2	0,0	0,0	0,2	0,2	0,0	0,2	0,0

Figure 21: The Strategic Form of the Game from Fig. 11 (p. 12).

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. 12) and Fig. 13 (p. 12) described the same strategic situation using different trees. Both of these have the same strategic form representation shown in Fig. 22 (p. 20).

		Player 2	
		<i>y</i>	<i>n</i>
Player 1	(2, 0)	2,0	0,0
	(1, 1)	1,1	0,0
	(0, 2)	0,2	0,0

Figure 22: The Strategic Form of the Games from Figures 12 and 13.

Note how the game in Fig. 21 (p. 20) differs from the game in Fig. 22 (p. 20). 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. 12) and Fig. 13 (p. 12) 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. 23 (p. 21).

The first step is to calculate the probability distribution over the terminal nodes. There are actually two ways one could do that: either calculate the probability distributions induced by each strategy profile, or calculate the probability distribution over the terminal nodes directly. In the first method, we would take each strategy profile and calculate the probabilities with which it produces various outcomes. In the other method, we would take one outcome




Figure 23: A Card Game in Extensive Form With Two Chance Moves.

and calculate the probabilities with which it can be reached by various paths. These are equivalent. We'll do the first method because it is more difficult to miss profiles.

Profile	Outcome Probability, $P(x s)$								$U_1(s)$	$U_2(s)$
	o_1	o_2	o_3	o_4	o_5	o_6	o_7	o_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. Using Tab. 1 (p. 21), we can construct the strategic form representation of the game in Fig. 23 (p. 21), as shown in Fig. 24 (p. 21).

		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 24: The Strategic Form of the Game from Fig. 23 (p. 21).

4.1.2 Three Players

Consider the strategic form of the game in Fig. 25 (p. 22). 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 25: 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. 22).

		Player 1							
		U				D			
Player 3		Player 2				Player 2			
		(A, C)	(A, E)	(B, C)	(B, E)	(A, C)	(A, E)	(B, C)	(B, E)
	(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. The set of these action is also 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 action 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. 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 stupid 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. 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,⁵ as illustrated in Fig. 26 (p. 24).

⁵This is just like a regular matrix except each entry consists of two numbers instead of one.

		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 26: 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).⁶

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). We can represent this game using the payoff matrix in Fig. 27 (p. 24).

		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 27: Prisoner's Dilemma, I.

The best outcome for a prisoner is to defect (D) and fink on the other while the other cooperates with him (C) and says nothing. In this case he gets off and his payoff is 3. The next-best outcome is when neither defects (finks) and so they are both convicted for the minor offense, in which case the payoff is 2. The next outcome is when both defect, in which case they are convicted only for the offense and they payoff is 1. Finally, the worst outcome

⁶There 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.

is to cooperate while the other defects, in which case the player gets the full sentence plus the punishment for obstructing justice, with a payoff of 0.

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 28: Prisoner's Dilemma, II.

The situation is absolutely the same because the ordinal ranking of the payoffs is the same as in Fig. 27 (p. 24).

In fact, 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 players 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.

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. Fig. 29 (p. 25) shows the Stag Hunt strategic situation applied to the security dilemma.

		USSR	
		<i>Arm</i>	<i>Don't Arm</i>
US	<i>Arm</i>	1,1	2,0
	<i>Don't Arm</i>	0,2	3,3

Figure 29: Stag Hunt modeling the Security Dilemma.

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.

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 game in Fig. 14 (p. 14). 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. 30 (p. 26), 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	
		c	d
Player 1	AE	1,1	1,1
	AF	1,1	1,1
	BE	-1,1	3,2
	BF	-1,1	4,0

Figure 30: The Normal Form of the Game from Fig. 14 (p. 14).

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 . In our example from Fig. 30 (p. 26), 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. We can 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. 31 (p. 26). 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 31: The Reduced Normal Form of the Game from Fig. 12 (p. 12).

For practice, let’s find the reduced normal form of the extensive-form game in Fig. 32 (p. 27). Player 1 has two information sets, with $A_1(\emptyset) = \{a, b\}$ and $A_1(b) = \{c, d, e\}$. He has six pure strategies:

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




Figure 32: Another Game from Myerson (p. 55).

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:

$$\begin{aligned} U_1((b, e), x) &= (0.5)(0) + (0.5)(6) = 3 \\ U_1((b, e), y) &= (0.5)(8) + (0.5)(6) = 7. \end{aligned}$$

Analogously, for player 2:

$$\begin{aligned} U_2((b, e), x) &= (0.5)(8) + (0.5)(0) = 4 \\ U_2((b, e), y) &= (0.5)(0) + (0.5)(0) = 0. \end{aligned}$$

We are now ready to construct the strategic form representation of this extensive form game. The result is in Fig. 33 (p. 27).

		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 33: The Strategic Form of the Game from Fig. 32 (p. 27).

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. 34 (p. 28).

		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 34: The Purely Reduced Strategic Form of the Game from Fig. 32 (p. 27).

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.

Each player's randomization is statistically independent of those of his opponents,⁷ and the payoffs to the mixed strategy profile are the expected values of the corresponding pure strategy payoffs.⁸ 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

⁷That is, the joint probability equals the product of individual probabilities.

⁸In 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.

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).

For example, consider the mixed strategy profile $\sigma = \langle (1/3H, 2/3T), (1/2H, 1/2T) \rangle$ in Matching Pennies. 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 puts equal probability on H and T . There are four pure strategy profiles: $S = \{(H, H), (H, T), (T, H), (T, T)\}$ that produce the four outcomes of the game. 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 $(1/3)(1/2) = 1/6$. Analogously, the probabilities of the other pure strategy profiles being played are $\Pr(H, T) = 1/6$, $\Pr(T, H) = 1/6$, $\Pr(T, T) = 1/6$. (You should verify that these sum to 1, which they must because they are probabilities of exhaustive and mutually exclusive events.)

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 $1/6(1) + 1/6(1) + 1/6(-1) + 1/6(-1) = 0$. Thus, player 1's expected payoff from the mixed strategy profile σ as specified above is 0. 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 is!

The **support** of a mixed strategy σ_i is the set of strategies to which σ_i assigns positive probability. Note that there is no requirement that a mixed strategy assign positive probability to all pure strategies. In particular, a pure strategy s_i is 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 .

As mentioned in the previous section, we can further reduce some strategic form games. Consider the game in Fig. 34 (p. 28). 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 expected payoff 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 when i uses $\hat{s}_i$ as when he 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. 32 (p. 27) (whose purely reduced normal form is in Fig. 34 (p. 28)) is given in Fig. 35 (p. 30).

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

Figure 35: The Fully Reduced Strategic Form of the Game from Fig. 32 (p. 27).

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 normal form representation* to mean the fully reduced form.

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.⁹ 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. 14) 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 $1/4$.

⁹In 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.

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. 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 $\beta = (0, 1)$.

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 distributions over outcomes for all strategies of the opponents.¹⁰ Intuitively, two strategies are equivalent if they have the same consequences regardless of what the other players do.

(Easy.) 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.¹¹

To check equivalence, we first need to specify the distribution over outcomes. For example, the game in Fig. 14 (p. 14) has four outcomes. Let the probability distribution (o_1, o_2, o_3, o_4) denote the associated probabilities for the outcomes $(1, 1)$, $(-1, 1)$, $(3, 2)$, and $(4, 0)$. Finally, let p denote the probability with which player 2 chooses *c* and $1 - p$ denote the probability with which she chooses *d*.

The behavior strategy $\beta = (1/4, 1/4)$, where player 1 chooses *A* and *E* with probability $1/4$, induces the probability distribution over outcomes given by $(1/4, 3/4p, 3/16(1 - p), 9/16(1 - p))$. (We obtained the probabilities for o_3 and o_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

¹⁰This is the same concept of equivalence we used when we discussed the reduced normal form representation of extensive games in the previous section.

¹¹This 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.

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 o_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 o_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 $3/4p$. The probability of o_3 is the probability that player 1 chooses both B and E multiplied by the probability that player 2 chooses d , which yields $\sigma(BE)(1 - p) = 3/16(1 - p)$. Finally, the probability of o_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 $9/16(1 - p)$. To summarize, the probability distribution over outcomes induced by the mixed strategy σ as defined above is $(1/4, 3/4p, 3/16(1 - p), 9/16(1 - p))$, 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!

(Not so easy.) 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.*

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 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 . 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 :

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

Note now that $\pi(h_i, a_i)$ is 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

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

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 :

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

How we define $\beta_i(h_i)(a_i)$ if $\pi_i(h_i) = 0$ is immaterial.¹² 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 an example. Consider the game in Fig. 36 (p. 34). 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 .4, (B, L) with probability .1, and (A, L) with probability .5.

We have $\sigma_1(B, R) = .4$, $\sigma_1(B, L) = .1$, $\sigma_1(A, L) = .5$, 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.

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, $\pi_1(h_1, A) = \sigma_1(A, L) + \sigma_1(A, R) = .5$. 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 $\pi_1(h_1, B) = \sigma_1(B, L) + \sigma_1(B, R) = .5$. We now have $\beta_1(h_1)(A) = \pi_1(h_1, A)/\pi_1(h_1) = .5/1 = .5$ and also $\beta_1(h_1)(B) = \pi_1(h_1, B)/\pi_1(h_1) = .5/1 = .5$.¹³ So, $\beta_1(h_1)(A) = \beta_1(h_1)(B) = .5$.

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

¹²Since 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.

¹³We verify that $\beta_1(h_1)(A) = 1 - \beta_1(h_1)(B)$, which is indeed the case.




Figure 36: A Game for Kuhn's Theorem, I.

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 $\pi_1(h_2) = \sigma_1(A, L) + \sigma_1(A, R) = .5$. 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, $\pi_1(h_2, L) = \sigma_1(A, L) = .5$. 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 $\beta_1(h_2)(L) = \pi_1(h_2, L)/\pi_1(h_2) = .5/.5 = 1$, and we also have $\beta_1(h_2)(R) = \pi_1(h_2, R)/\pi_1(h_2) = 0/.5 = 0$.¹⁴

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

$$\begin{aligned}\beta_1(h_1)(A) &= .5 \\ \beta_1(h_1)(B) &= .5 \\ \beta_1(h_2)(L) &= 1 \\ \beta_1(h_2)(R) &= 0\end{aligned}$$

Let's check the equivalence claim. Let $(p, q, 1 - p - q)$ denote a mixed strategy for player 2. Using the mixed strategy σ_1 , the probabilities of reaching the outcomes are as follows:

$$\begin{aligned}o_1 : [\sigma_1(A, L) + \sigma_1(A, R)]p &= .5p \\ o_2 : \sigma_1(A, L)q &= .5q \\ o_3 : \sigma_1(A, R)q &= 0 \\ o_4 : \sigma_1(A, L)(1 - p - q) &= .5(1 - p - q) \\ o_5 : \sigma_1(A, R)(1 - p - q) &= 0 \\ o_6 : \sigma_1(B, L) + \sigma_1(B, R) &= .5\end{aligned}$$

The distribution over outcomes using σ_1 is then $(.5p, .5q, 0, .5(1 - p - q), 0, .5)$.

¹⁴We again verify that the distribution is valid, which it is because $\beta_1(h_2)(L) + \beta_1(h_2)(R) = 1$.

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

$$\begin{aligned}
o_1 : \beta_1(h_1)(A)p &= .5p \\
o_2 : \beta_1(h_1)(A)q\beta_1(h_2)(L) &= (.5)q(1) = .5q \\
o_3 : \beta_1(h_1)(A)q\beta_1(h_2)(R) &= (.5)q(0) = 0 \\
o_4 : \beta_1(h_1)(A)(1-p-q)\beta_1(h_2)(L) &= (.5)(1-p-q)(1) = .5(1-p-q) \\
o_5 : \beta_1(h_1)(A)(1-p-q)\beta_1(h_2)(R) &= (.5)(1-p-q)(0) = 0 \\
o_6 : \beta_1(h_1)(B) &= .5
\end{aligned}$$

This yields the distribution over outcomes $(.5p, .5q, 0, .5(1-p-q), 0, .5)$ that is the same as the one given by the mixed strategy. Therefore, we have shown that σ_1 and β_1 are equivalent.

Now let's illustrate the claim that a mixed strategy may generate more than one behavior strategy. Consider the same game and suppose $\sigma_1(A, L) = \sigma_1(A, R) = 0$, $\sigma_1(B, L) = 0.4$, 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:

$$\begin{aligned}
o_1 : [\sigma_1(A, L) + \sigma_1(A, R)]p &= 0 \\
o_2 : \sigma_1(A, L)q &= 0 \\
o_3 : \sigma_1(A, R)q &= 0 \\
o_4 : \sigma_1(A, L)(1-p-q) &= 0 \\
o_5 : \sigma_1(A, R)(1-p-q) &= 0 \\
o_6 : \sigma_1(B, L) + \sigma_1(B, R) &= 1.
\end{aligned}$$

Under the behavior strategy, the probability distribution is:

$$\begin{aligned}
o_1 : \beta_1(h_1)(A)p &= 0 \\
o_2 : \beta_1(h_1)(A)q\beta_1(h_2)(L) &= (0)qx = 0 \\
o_3 : \beta_1(h_1)(A)q\beta_1(h_2)(R) &= (0)q(1-x) = 0 \\
o_4 : \beta_1(h_1)(A)(1-p-q)\beta_1(h_2)(L) &= (0)(1-p-q)x = 0 \\
o_5 : \beta_1(h_1)(A)(1-p-q)\beta_1(h_2)(R) &= (0)(1-p-q)(1-x) = 0 \\
o_6 : \beta_1(h_1)(B) &= 1.
\end{aligned}$$

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.

Now let's illustrate the claim that different mixed strategies can generate the same behavioral strategy. Consider the game in Fig. 37 (p. 36). 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 37: 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 $\beta_2(h_2)(C) = \beta_2(h_2)(D) = 1/2$.¹⁵ 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 o_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 already see, there will be an infinite number of mixed strategies that generate this behavior strategy: All σ_2 such that $\sigma_2(AC) + \sigma_2(AD) = 1/2$ will do that.

Although it is important to distinguish between the two types of probabilistic strategies, in reality we shall be considering 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 .

¹⁵You 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.