

Strategic and extensive games

Martin J. Osborne

Department of Economics, University of Toronto

`martin.osborne@utoronto.ca`

2005-6-30. This version 2006-4-24.

Abstract The basic theory of strategic and extensive games is described. Strategic games, Bayesian games, extensive games with perfect information, and extensive games with imperfect information are defined and explained. Among the solution concepts discussed are Nash equilibrium, correlated equilibrium, rationalizability, subgame perfect equilibrium, and weak sequential equilibrium.

1. Introduction

Game theory is a collection of models designed to understand situations in which decision-makers interact. This chapter discusses models that focus on the behavior of individual decision-makers. These models are sometimes called “noncooperative”.

2. Strategic games

2.1 Definition

The basic model of decision-making by a single agent consists of a set of possible actions and a preference relation over this set. The simplest theory of the agent’s behavior is that she chooses a member of the set that is best according to the preference relation.

The model of a *strategic game* extends this model to many agents, who are referred to as *players*. Each player has a set of possible actions and a preference relation over action *profiles* (lists of actions, one for each player).

Definition 1 A *strategic game with deterministic preferences* consists of

- a set N (the set of **players**)

I am grateful to Jean Guillaume Forand for comments and to the Social Sciences and Humanities Research Council of Canada for financial support.

To appear in *The New Palgrave Dictionary of Economics, 2nd Edition* (Steven Durlauf and Lawrence Blume, eds.), Palgrave Macmillan, Basingstoke. Reproduced with the permission of Palgrave Macmillan.

and for each player $i \in N$

- a set A_i (the set of player i 's possible **actions**)
- a preference relation \succsim_i over the set $\times_{j \in N} A_j$ of action profiles.

A strategic game $\langle N, (A_i), (\succsim_i) \rangle$ is **finite** if the set N of players and the set A_i of actions of each player i are finite.

The fact that each player's preferences are defined over the set of action profiles allows for the possibility that each player cares not only about her own action but also about the other players' actions, distinguishing the model from a collection of independent single-agent decision problems.

Notice that the model does not have a temporal dimension. An assumption implicit in the solution notions applied to a game is that each player independently commits to an action before knowing the action chosen by any other player. Notice also that no structure is imposed on the players' sets of actions. In the simplest cases, a player's set of actions may consist of two elements; in more complex cases, it may consist, for example, of an interval of real numbers, a set of points in a higher dimensional space, a set of functions from one set to another, or a combination of such sets. In particular, an action may be a contingent plan, specifying a player's behavior in a variety of possible circumstances, so that the model is not limited to "static" problems (see [Section 3.1.1](#)). Thus although the model has no temporal dimension, it may be used to study "dynamic" situations under the assumption that each player chooses her plan of action once and for all.

A few examples give an idea of the range of situations that the model encompasses. The most well-known strategic game is the *Prisoner's Dilemma*. In this game, there are two players ($N = \{1, 2\}$, say), each player has two actions, *Quiet* and *Fink*, and each player's preference relation ranks the action pair in which she chooses *Fink* and the other player chooses *Quiet* highest, then $(\textit{Quiet}, \textit{Quiet})$, then $(\textit{Fink}, \textit{Fink})$, and finally the action profile in which she chooses *Quiet* and the other player chooses *Fink*. In this example, as in most examples, working with payoff representations of the players' preference relations is simpler than working with the preference relations themselves. Taking a payoff function for each player that assigns the payoffs 3, 2, 1, and 0 to the four outcomes, we may conveniently represent the game in the table in [Figure 1](#). (Any two-player strategic game in which each player has finitely many actions may be represented in a similar table.)

This game takes its name from the following scenario. The two players are suspected of joint involvement in a major crime. Sufficient evidence exists to

		Player 2	
		<i>Quiet</i>	<i>Fink</i>
Player 1	<i>Quiet</i>	2, 2	0, 3
	<i>Fink</i>	3, 0	1, 1

Figure 1. The *Prisoner's Dilemma*.

convict each one of a minor offense, but conviction of the major crime requires at least one of them to confess, thereby implicating the other (i.e. one player “finks”). Each suspect may stay quiet or may fink. If a single player finks she is rewarded by being set free, whereas the other player is convicted of the major offense. If both players fink then each is convicted but serves only a moderate sentence.

The game derives its interest not from this specific interpretation, but because the structure of the players’ preferences fits many other social and economic situations. The combination of the desirability of the players’ coordinating on an outcome and the incentive on the part of each player individually to deviate from this outcome is present in situations as diverse as duopolists setting prices and countries involved in an arms race.

Another example of a strategic game models oligopoly as suggested by [Cournot \(1838\)](#). The players are the n firms, each player’s set of actions is the set of possible outputs (the set of nonnegative real numbers), and the preference relation of player i is represented by its profit, given by the payoff function u_i defined by

$$u_i(q_1, \dots, q_n) = q_i P \left(\sum_{j=1}^n q_j \right) - C_i(q_i),$$

where q_i is player i ’s output, C_i is its cost function, and P is the inverse demand function, giving the market price for any total output. Another strategic game that models oligopoly, associated with the name of Bertrand, differs from Cournot’s model in taking the set of actions of each player to be the set of possible prices (which requires profit to be defined as a function of prices).

A strategic game that models competition between candidates for political office was suggested by [Hotelling \(1929\)](#). The set of players is a finite set of candidates; each player’s set of actions is the same subset X of the line, representing the set of possible policies. Each member of a continuum of citizens (who are not players in the game) has single-peaked preferences over X . Each citizen votes for the candidate whose position is closest to her favorite position. A density function on X represents the distribution of the citizens’ favorite policies. The total number of votes obtained by any player is the integral with respect to this den-

sity over the subset of X consisting of points closer to the player's action (chosen policy) than to the action of any other player. A player's preferences are represented by the payoff function that assigns 1 to any action profile in which she obtains more votes than every other player, $1/k$ to any action profile in which she obtains at least as many votes as any other player and $k \geq 2$ players tie for the highest number of votes, and 0 to any action profile in which she obtains fewer votes than some other player.

2.2 Nash equilibrium

Which action profile will result when a strategic game is played? Game theory provides two main approaches to answering this question. One isolates action profiles that correspond to stable "steady states". This approach leads to the notion of Nash equilibrium, discussed in this section. The other approach, discussed in [Section 2.5](#), isolates action profiles that are consistent with each player's reasoning regarding the likely actions of the other players, taking into account the other players' reasoning about each other and the player in question.

Fix an n -player strategic game and suppose that for each player in the game there exists a population of K individuals, where K is large. Imagine that in each of a long sequence of periods, K sets of n individuals are randomly selected, each set consisting of one individual from each population. In each period, each set of n individuals plays the game, the individual from population i playing the role of player i , for each value of i . The selected sets change from period to period; because K is large, the chance that an individual will play the game with the same opponent twice is low enough not to enter her strategic calculations. If play settles down to a steady state in which each individual in each population i chooses the same action, say a_i^* , whenever she plays the game, what property must the profile a^* satisfy?

In such a (deterministic) steady state, each individual in population i knows from her experience that every individual in every other population j chooses a_j^* . Thus we can think of each such individual as being involved in a single-person decision problem in which the set of actions is A_i and the preferences are induced by player i 's preference relation in the game when the action of every other player j is fixed at a_j^* . That is, a_i^* maximizes i 's payoff in the game given the actions of all other players. Or, looked at differently, a^* has the property that no player i can increase her payoff by changing her action a_i^* given the other players' actions. An action profile with this property is a *Nash equilibrium*. (The notion is due to [Nash 1950](#); the underlying idea goes back at least to [Cournot 1838](#)). For any action profile b , denote by (a_i, b_{-i}) the action profile in which

player i 's action is a_i and the action of every other player j is b_j .

Definition 2 A *Nash equilibrium* of the strategic game $\langle N, (A_i), (\succeq_i) \rangle$ is an action profile a^* for which

$$a^* \succeq_i (a_i, a_{-i}^*) \text{ for all } a_i \in A_i$$

for every player $i \in N$.

By inspection of the four action pairs in the *Prisoner's Dilemma* (Figure 1) we see that the action pair (*Fink*, *Fink*) is the only Nash equilibrium. For each of the three other action pairs, a player choosing *Quiet* can increase her payoff by switching to *Fink*, given the other player's action.

The games in Figure 2 immediately answer three questions: Does every strategic game necessarily have a Nash equilibrium? Can a strategic game have more than one Nash equilibrium? Is it possible that every player is better off in one Nash equilibrium than she is in another Nash equilibrium? The left-hand game, which models the game "Matching pennies", has no Nash equilibrium. The right-hand game has two Nash equilibria, (B, B) and (C, C) , and both players are better off in (C, C) than they are in (B, B) .

	B	C
B	1, -1	-1, 1
C	-1, 1	1, -1

	B	C
B	1, 1	0, 0
C	0, 0	2, 2

Figure 2. Two strategic games.

In some games, especially ones in which each player has a continuum of actions, Nash equilibria may most easily be found by first computing each player's best action for every configuration of the other players' actions. For each player i , let u_i be a payoff function that represents player i 's preferences. Fix a player i and define, for each list a_{-i} of the other players' actions, the set of actions that maximize i 's payoff:

$$B_i(a_{-i}) = \{a_i \in A_i : a_i \text{ maximizes } u_i(a_i, a_{-i}) \text{ over } a_i \in A_i\}.$$

Each member of $B_i(a_{-i})$ is a *best response* of player i to a_{-i} ; the function B_i is called player i 's *best response function*. (Note that it is set-valued.) An action profile a^* is a Nash equilibrium if and only if

$$a_i^* \in B_i(a_{-i}^*) \text{ for every player } i.$$

In some games, the set $B_i(a_{-i})$ is a singleton for every player i and every list a_{-i} . For such a game, denote the single element by $b_i(a_{-i})$. Then the condition for

the action profile a^* to be a Nash equilibrium may be written as

$$a_i^* = b_i(a_{-i}^*) \text{ for every player } i,$$

a collection of n equations in n unknowns.

Consider, for example, a two-player game in which each player's set of actions is the set of nonnegative real numbers and the preference relation of each player i is represented by the payoff function u_i defined by

$$a_i(c + a_j - a_i)$$

where $c > 0$ is a constant. In this game each player i has a unique best response to every action a_j of the other player (j), given by $b_i(a_j) = \frac{1}{2}(c + a_j)$. The two equations $a_1 = \frac{1}{2}(c + a_2)$ and $a_2 = \frac{1}{2}(c + a_1)$ immediately yield the unique solution $(a_1, a_2) = (c, c)$, which is thus the only Nash equilibrium of the game.

2.3 Mixed strategy Nash equilibrium

In a steady state modeled by the notion of Nash equilibrium, all individuals who play the role of a given player choose the same action whenever they play the game. We may generalize this notion. In a *stochastic steady state*, the rule used to select an action by individuals in the role of a given player is probabilistic rather than deterministic. In a *polymorphic steady state*, each individual chooses the same action whenever she plays the game, but different individuals in the role of a given player choose different deterministic actions.

In both of these generalized steady states an individual faces uncertainty: in a stochastic steady state because the individuals with whom she plays the game choose their actions probabilistically, and in a polymorphic steady state because her potential opponents, who are chosen probabilistically from their respective populations, choose different actions. Thus to analyze the players' behavior in such steady states, we need to specify their preferences regarding lotteries over the set of action profiles. The following extension of [Definition 1](#) assumes that these preferences are represented by the expected value of a payoff function. (The term "vNM preferences" refers to [von Neumann and Morgenstern \(1944, pp. 15–31; 1947, pp. 204–221\)](#), who give conditions on preferences under which such a representation exists.)

Definition 3 A *strategic game (with vNM preferences)* consists of

- a set N (the set of **players**)

and for each player $i \in N$

- a set A_i (the set of player i 's possible **actions**)
- a function $u_i : \times_{j \in N} A_j \rightarrow \mathbb{R}$ (player i 's **payoff function**, the expected value of which represents i 's preferences over the set of lotteries over action profiles).

A probability distribution over A_i , the set of actions of player i , is called a *mixed strategy* of player i . The notion of a *mixed strategy Nash equilibrium* corresponds to a stochastic steady state in which each player chooses her mixed strategy to maximize her expected payoff, given the other players' mixed strategies.

Definition 4 A *mixed strategy Nash equilibrium* of the strategic game $\langle N, (A_i), (u_i) \rangle$ is a profile α^* in which each component α_i^* is a probability distribution over A_i that satisfies

$$U_i(\alpha^*) \geq U_i(\alpha_i, \alpha_{-i}^*) \text{ for every probability distribution } \alpha_i \text{ on } A_i$$

for every player $i \in N$, where $U_i(\alpha)$ is the expected value of $u_i(a)$ under α .

Suppose that each player's set of actions is finite and fix the mixed strategy of every player $j \neq i$ to be α_j . Then player i 's expected payoff when she uses the mixed strategy α_i is a weighted average of her expected payoffs to each of the actions to which α_i assigns positive probability. Thus if α_i maximizes player i 's expected payoff given α_{-i} , then so too do all the actions to which α_i assigns positive probability. This observation has two significant consequences. First, a mixed strategy Nash equilibrium corresponds not only to a stochastic steady state but also to a polymorphic steady state. (The equilibrium probability $\alpha_i^*(a_i)$ is the fraction of individuals in population i that choose a_i .) Second, the fact that in a mixed strategy Nash equilibrium each player is indifferent between all the actions to which her mixed strategy assigns positive probability is sometimes useful when computing mixed strategy Nash equilibria.

To illustrate the notion of a mixed strategy Nash equilibrium, consider the games in [Figure 2](#). In the game on the left, a player's expected payoff is the same (equal to 0) for her two actions when the other player chooses each action with probability $\frac{1}{2}$, so that the game has a mixed strategy Nash equilibrium in which each player chooses each action with probability $\frac{1}{2}$. The game has no other mixed strategy Nash equilibrium because each player's best response to any mixed strategy other than the one that assigns probability $\frac{1}{2}$ to each action is either the action B or the action C , and we know that the game has no equilibrium in which neither player randomizes.

The game on the right of [Figure 2](#) has three mixed strategy Nash equilibria. Two correspond to the Nash equilibria of the game in which randomization is not

allowed: each player assigns probability 1 to B , and each player assigns probability 1 to C . In the third equilibrium, each player assigns probability $\frac{2}{3}$ to B and probability $\frac{1}{3}$ to C . This strategy pair is a mixed strategy Nash equilibrium because each player's expected payoff to each of her actions is the same (equal to $\frac{2}{3}$ for both players).

The notion of mixed strategy Nash equilibrium generalizes the notion of Nash equilibrium in the following sense.

- If a^* is a Nash equilibrium of the strategic game $\langle N, (A_i), (\succsim_i) \rangle$, then the mixed strategy profile in which each player i assigns probability 1 to a_i^* is a mixed strategy Nash equilibrium of any strategic game with vNM preferences $\langle N, (A_i), (u_i) \rangle$ in which, for each player i , u_i represents \succsim_i .
- If a^* is a mixed strategy Nash equilibrium of the strategic game with vNM preferences $\langle N, (A_i), (u_i) \rangle$ in which for each player i there is an action a_i^* such that $\alpha_i^*(a_i^*) = 1$, then a^* is a Nash equilibrium of the strategic game $\langle N, (A_i), (\succsim_i) \rangle$ in which, for each player i , \succsim_i is the preference relation represented by u_i .

The following result gives a sufficient condition for a strategic game to have a mixed strategy Nash equilibrium.

Proposition 5 *A strategic game with vNM preferences $\langle N, (A_i), (u_i) \rangle$ in which the set N of players is finite has a mixed strategy Nash equilibrium if either (a) the set A_i of actions of each player i is finite or (b) the set A_i of actions of each player i is a compact convex subset of a Euclidean space and the payoff function u_i of each player i is continuous.*

Part (a) of this result is due to [Nash \(1950, 1951\)](#) and part (b) is due to [Glicksberg \(1952\)](#).

In many games of economic interest the players' payoff functions are not continuous. Several results giving conditions for the existence of a mixed strategy Nash equilibrium in such games are available; see, for example, Section 5 of [Reny \(1999\)](#).

As I have noted, in any mixed strategy Nash equilibrium in which some player chooses an action with positive probability less than 1, that player is indifferent between all the actions to which her strategy assigns positive probability. Thus she has no positive reason to choose her equilibrium strategy: any other strategy that assigns positive probability to the same actions is equally good. This fact shows that the notion of a mixed strategy equilibrium lacks robustness. A result of [Harsanyi \(1973\)](#) addresses this issue. For any strategic game G , Harsanyi

considers a game in which the players' payoffs are randomly perturbed by small amounts from their values in G . In any play of the perturbed game, each player knows her own payoffs, but not (exactly) those of the other players. (Formally the perturbed game is a Bayesian game, a model described in [Section 2.6](#).) Typically, a player has a unique optimal action in the perturbed game, and this game has an equilibrium in which no player randomizes. (Each player's equilibrium action depends on the value of her own payoffs.) Harsanyi shows that the limit of these equilibria as the perturbations go to zero defines a mixed strategy Nash equilibrium of G , and almost any mixed strategy Nash equilibrium of G is associated with the limit of such a sequence. Thus we can think of the players' strategies in a mixed strategy Nash equilibrium as approximations to collections of *strictly* optimal actions.

2.4 Correlated equilibrium

One interpretation of a mixed strategy Nash equilibrium is that each player conditions her action on the realization of a random variable, where the random variable observed by each player is independent of the random variable observed by every other player. This interpretation leads naturally to the question of how the theory changes if the players may observe random variables that are not independent.

To take a simple example, consider the game at the right of [Figure 2](#). Suppose that the players observe random variables that are perfectly correlated, each variable taking one value, say x , with some probability p , and another value, say y , with probability $1 - p$. Consider the strategy that chooses the action B if the realization of the player's random variable is x and the action C if the realization is y . If one player uses this strategy, the other player optimally does so too: if the realization is x , for example, she knows the other player will choose B , so that her best action is B . Thus the strategy pair is an equilibrium.

More generally, the players may observe random variables that are partially correlated. Equilibria in which they do so exist for the game at the right of [Figure 2](#), but the game in [Figure 3](#) is more interesting.

	B	C
B	6,6	2,7
C	7,2	0,0

Figure 3. A strategic game.

Consider the random variable that takes the values x , y , and z , each with probability $\frac{1}{3}$. Player 1 observes only whether the realization is in $\{x, y\}$ or is z

(but not, in the first case, whether it is x or y), and player 2 observes only whether it is in $\{x, z\}$ or is y . Suppose that player 1 chooses B if she observes $\{x, y\}$ and C if she observes z , and player 2 chooses B if she observes $\{x, z\}$ and C if she observes y . Then neither player has an incentive to change her action, whatever she observes. If, for example, player 1 observes $\{x, y\}$, then she infers that x and y have each occurred with probability $\frac{1}{2}$, so that player 2 will choose each of her actions with probability $\frac{1}{2}$. Thus her expected payoff is 4 if she chooses B and $\frac{7}{2}$ if she chooses C , so that B is optimal. Similarly, if player 1 observes z , she infers that player 2 will choose B , so that C is optimal for her. The outcome is (B, B) with probability $\frac{1}{3}$, (B, C) with probability $\frac{1}{3}$, and (C, B) with probability $\frac{1}{3}$, so that each player's expected payoff is 5.

An interesting feature of this equilibrium is that both players' payoffs exceed their payoffs in the unique mixed strategy Nash equilibrium (in which each player chooses B with probability $\frac{2}{3}$ and obtains the expected payoff $\frac{14}{3}$).

In general, a *correlated equilibrium* of a strategic game with vNM preferences consists of a probability space and, for each player, a partition of the set of states and a function associating an action with each set in the partition (the player's *strategy*) such that for each player and each set in the player's partition, the action assigned by her strategy to that set maximizes her expected payoff given the probability distribution over the other players' actions implied by her information. (The notion of correlated equilibrium is due to [Aumann 1974](#).)

The appeal of a correlated equilibrium differs little from the appeal of a mixed strategy equilibrium. In one respect, in fact, most correlated equilibria are more appealing: the action specified by each player's strategy for each member of her partition of the set of states is strictly optimal (she is not indifferent between that action and any others). Nevertheless, the notion of correlated equilibria has found few applications.

2.5 Rationalizability

The outcome $(Fink, Fink)$ of the *Prisoner's Dilemma* is attractive not only because it is a Nash equilibrium (and hence consistent with a steady state). In addition, for each player, *Fink* is optimal and *Quiet* is suboptimal *regardless* of the other player's action. That is, we may argue solely on the basis of a player's rationality that she will select *Fink*; no reference to her belief about the other player's action is necessary.

We say that the mixed strategy α_i of player i is *rational* if there exists a probability distribution over the other players' actions to which it is a best response. (The probability distribution may entail correlation between the other players' actions; we do not require it to be derived from independent mixed strategies.)

Using this terminology, the only rational action for each player in the *Prisoner's Dilemma* is *Fink*.

This definition of rationality puts no restriction on the probability distribution over the other players' actions that justifies a player's mixed strategy. In particular, an action is rational even if it is a best response only to a belief that assigns positive probability to the other players' not being rational. For example, in the game on the left of Figure 4, *Q* is rational for player 1, but all the mixed strategies of player 2 to which *Q* is a best response for player 1 assign probability of at least $\frac{1}{2}$ to *Q*, which is not rational for player 2. Such beliefs are ruled out if we assume that each player is not only rational, but also believes that the other players are rational. In the game on the left of Figure 4 this assumption means that player 1's beliefs must assign positive probability only to player 2's action *F*, so that player 1's only optimal action is *F*. That is, in this game the assumptions that each player is rational and that each player believes the other player is rational isolate the action pair (*F*, *F*).

	<i>Q</i>	<i>F</i>
<i>Q</i>	3,2	0,3
<i>F</i>	2,0	1,1

	<i>Q</i>	<i>F</i>
<i>Q</i>	4,2	0,3
<i>X</i>	1,1	1,0
<i>F</i>	3,0	2,2

Figure 4. Two variants of the *Prisoner's Dilemma*.

We may take this argument further. Consider the game on the right of Figure 4. Player 1's action *Q* is consistent with player 1's rationality and also with a belief that player 2 is rational (because both actions of player 2 are rational). It is not, however, consistent with player 1's believing that player 2 believes that player 1 is rational. If player 2 believes that player 1 is rational, her belief must assign probability 0 to player 1's action *X* (which is not a best response to any strategy of player 2), so that her only optimal action is *F*. But if player 2 assigns positive probability only to *F*, then player 1's action *Q* is not optimal.

In all of these games—the *Prisoner's Dilemma* and the two in Figure 4—player 1's action *F* survives any number of iterations of the argument: it is consistent with player 1's rationality, player 1's belief that player 2 is rational, player 1's belief that player 2 believes that player 1 is rational, and so on. An action with this property is called *rationalizable*, a notion¹ developed independently by Bernheim (1984) and Pearce (1984).

¹Both Bernheim and Pearce discuss a slightly different notion, in which players are restricted to beliefs that are derived from *independent* probability distributions over each of the other player's actions. Their notion does not have the same properties as the notion described here.

The set of action profiles in which every player's action is rationalizable may be given a simple characterization. First define a strictly dominated action.

Definition 6 *Player i 's action a_i in the strategic game with vNM preferences $\langle N, (A_i), (u_i) \rangle$ is **strictly dominated** if for some mixed strategy α_i of player i we have*

$$U_i(\alpha_i, a_{-i}) > u_i(a_i, a_{-i}) \text{ for every } a_{-i} \in \times_{j \in N \setminus \{i\}} A_j,$$

where $U_i(\alpha_i, a_{-i})$ is player i 's expected payoff when she uses the mixed strategy α_i and the other players' actions are given by a_{-i} .

Note that the fact that α_i in this definition is a mixed strategy is essential: some strictly dominated actions are not strictly dominated by any action. For example, in the variant of the game at the left of [Figure 4](#) in which player 1 has an additional action, say Z , with $u_1(Z, Q) = 0$ and $u_1(Z, F) = 5$, the action F is not strictly dominated by any action, but is strictly dominated by the mixed strategy that assigns probability $\frac{3}{4}$ to Q and probability $\frac{1}{4}$ to Z .

We may show that an action in a finite strategic game is not rational if and only if it is strictly dominated. Given this result, it is not surprising that actions are rationalizable if they survive the iterated elimination of strictly dominated actions, defined precisely as follows.

Definition 7 *Let $G = \langle N, (A_i), (u_i) \rangle$ be a strategic game. For each $j \in N$, let $X_j^1 = A_j$, and for each $j \in N$ and each $t \geq 1$, let X_j^{t+1} be a subset of X_j^t with the property that every member of $X_j^t \setminus X_j^{t+1}$ is strictly dominated in the game $\langle N, (X_i^t), (u_i^t) \rangle$, where u_i^t denotes the restriction of the function u_i to $\times_{j \in N} X_j^t$. If no member of X_j^t for any $j \in N$ is strictly dominated, then the set $\times_{j \in N} X_j^t$ **survives iterated elimination of strictly dominated actions**.*

The procedure specified in this definition does not pin down exactly which actions are eliminated at each step. Only strictly dominated actions are eliminated, but not *all* such actions are necessarily eliminated. Thus the definition leaves open the question of the uniqueness of the set of surviving action profiles. In fact, however, this set is unique; it coincides with the set of profiles of rationalizable actions.

Proposition 8 *In a finite strategic game the set of action profiles that survives iterated elimination of strictly dominated actions is unique and is equal to the set of profiles of rationalizable actions.*

Every action of any player used with positive probability in a correlated equilibrium is rationalizable. Thus the set of profiles of rationalizable actions is the

largest “solution” for a strategic game that we have considered. In many games, in fact, it is very large. (If no player has a strictly dominated action, all actions of every player are rationalizable, for example.) However, in several of the games mentioned in the previous sections, each player has a single rationalizable action, equal to her unique Nash equilibrium action. This property holds, with some additional assumptions, for Cournot’s and Bertrand’s oligopoly games with two firms and Hotelling’s model of electoral competition with two candidates. The fact that in other games the set of rationalizable actions is large has limited applications of the notion, but it remains an important theoretical construct, delineating exactly the conclusion we may reach by assuming that the players take into account each others’ rationality.

2.6 Bayesian games

In the models discussed in the previous sections, every player is fully informed about all the players’ characteristics—their actions, payoffs, and information. In the model of a Bayesian game, players are allowed to be uncertain about these characteristics. We call each configuration of characteristics a *state*. The fact that each player’s information about the state may be imperfect is modeled by assuming that each player does not observe the state, but rather receives a *signal* that may depend on the state. At one extreme, a player may receive a different signal in every state; such a player has perfect information. At another extreme, a player may receive the same signal in every state; such a player has no information about the state. In between these extremes are situations in which a player is partially informed; she may receive the same signal in states ω_1 and ω_2 , for example, and a different signal in state ω_3 .

To make a decision, given her information, a player needs to form a belief about the probabilities of the states between which she cannot distinguish. We assume that she starts with a *prior belief* over the set of states, and acts upon the posterior belief derived from this prior, given her signal, using Bayes’ Law. If, for example, there are three states, ω_1 , ω_2 , and ω_3 , to which her prior belief assigns probabilities $\frac{1}{2}$, $\frac{1}{4}$, and $\frac{1}{4}$, and she receives the same signal, say X , in states ω_1 and ω_2 , and a different signal, say Y , in state ω_3 , then her posterior belief assigns probability $\frac{2}{3}$ to ω_1 and probability $\frac{1}{3}$ to ω_2 when she receives the signal X and probability 1 to ω_3 when she receives the signal Y .

In summary, a Bayesian game is defined as follows. (The notion is due to [Harsanyi 1967/68](#).)

Definition 9 A *Bayesian game* consists of

- a set N (the set of *players*)

- a set Ω (the set of **states**)

and for each player $i \in N$

- a set A_i (the set of player i 's possible **actions**)
- a set T_i (the set of **signals** that player i may receive) and a function $\tau_i : \Omega \rightarrow T_i$ associating a signal with each state (player i 's **signal function**)
- a probability distribution p_i over Ω (player i 's **prior belief**), with $p_i(\tau_i^{-1}(t_i)) > 0$ for all $t_i \in T_i$
- a function $u_i : (\times_{j \in N} A_j) \times \Omega \rightarrow \mathbb{R}$ (player i 's **payoff function**, the expected value of which represents i 's preferences over the set of lotteries on the set $(\times_{j \in N} A_j) \times \Omega$).

This definition allows the players to hold different prior beliefs. In many applications every player is assumed to hold the same prior belief.

A widely-studied class of Bayesian games models auctions. An example is a single-object auction in which each player knows her own valuation of the object but not that of any other player and believes that every player's valuation is independently drawn from the same distribution. In a Bayesian game that models such a situation, the set of states is the set of profiles of valuations and the signal received by each player depends only on her own valuation, not on the valuation of any other player. Each player holds the same prior belief, which is derived from the assumption that each player's valuation is drawn independently from the same distribution.

The desirability for a player of each of her actions depends in general on the signal she receives. Thus a candidate for an equilibrium in a Bayesian game is a profile of functions, one for each player; the function for player i associates an action (member of A_i) with each signal she may receive (member of T_i). We refer to player i after receiving the signal t_i as *type t_i* of player i . A Nash equilibrium of a Bayesian game embodies the same principle as does a Nash equilibrium of a strategic game: each player's action is optimal given the other players' actions. Thus in an equilibrium, the action of each type of each player maximizes the payoff of that type given the action of every other type of every other player. That is, a Nash equilibrium of a Bayesian game is a Nash equilibrium of the strategic game in which the set of players is the set of pairs (i, t_i) , where i is a player in the Bayesian game and t_i is a signal that she may receive.

Definition 10 A **Nash equilibrium of a Bayesian game** $\langle N, \Omega, (a_i), (T_i), (\tau_i), (p_i), (u_i) \rangle$ is a Nash equilibrium of the following strategic game.

- The set of players is the set of all pairs (i, t_i) such that $i \in N$ and $t_i \in T_i$.
- The set of actions of player (i, t_i) is A_i .
- The payoff of player (i, t_i) when each player (j, t_j) chooses the action $a(j, t_j)$ is

$$\sum_{\omega \in \Omega} \Pr(\omega | t_i) u_i((a_i, \hat{a}_{-i}(\omega)), \omega),$$

where $\hat{a}_j(\omega) = a(j, \tau_j(\omega))$ for each $j \in N$.

To illustrate this notion, consider the two-player Bayesian game in which there are two states, each player has two actions (B and C), player 1 receives the same signal in both states, player 2 receives a different signal in each state, each player's prior belief assigns probability $\frac{1}{3}$ to state 1 and probability $\frac{2}{3}$ to state 2, and the payoffs are those shown in Figure 5. A Nash equilibrium of this Bayesian game is a Nash equilibrium of the three player game in which the players are player 1 and the two types of player 2 (one for each state). I claim that the strategy profile in which player 1 chooses B , type 1 of player 2 (i.e. player 2 after receiving the signal that the state is 1) chooses C , and type 2 of player 2 chooses B is a Nash equilibrium. The actions of the two types of player 2 are best responses to the action B of player 1. Given these actions, player 1's expected payoff to B is $\frac{2}{3}$ (because with probability $\frac{1}{3}$ the state is 1 and player 2 chooses C and with probability $\frac{2}{3}$ the state is 2 and player 2 chooses B) and her expected payoff to C is $\frac{1}{3}$. Thus player 1's action B is a best response to the actions of the two types of player 2.

	State 1 (probability $\frac{1}{3}$)		State 2 (probability $\frac{2}{3}$)	
	B	C	B	C
B	1, 0	0, 1	1, 1	0, 0
C	1, 1	1, 0	1, 0	0, 1

Figure 5. A Bayesian game.

3. Extensive games

Although situations in which players choose their actions sequentially may be modeled as strategic games, they are more naturally modeled as extensive games. In Section 3.1 I discuss a model in which each player, when choosing an action, knows the actions taken previously. In Section 3.2 I discuss a more complex model that allows players to be imperfectly informed. (The notion of an

extensive game is due to von Neumann and Morgenstern 1944 and Kuhn 1950, 1953. The formulation in terms of histories is due to Ariel Rubinstein.)

3.1 Extensive games with perfect information

An *extensive game with perfect information* describes the sequential structure of the players' actions. It does so by specifying the set of sequences of actions that may occur and the player who chooses an action after each subsequence. A sequence that starts with an action of the player who makes the first move and ends when no move remains is called a *terminal history*.

Definition 11 An *extensive game with perfect information* consists of

- a set N (the set of **players**)
- a set H of sequences (the set of **terminal histories**) with the property that no sequence is a proper subhistory of any other sequence
- a function P (the **player function**) that assigns a player to every proper subsequence of every terminal history

and for each player $i \in N$

- a preference relation \succsim_i over the set H of terminal histories.

The restriction on the set H is necessary for its members to be interpreted as terminal histories: if (x, y, z) is a terminal history then (x, y) is not a terminal history, because z may be chosen after (x, y) . We refer to subsequences of terminal histories as *histories*.

The sets of actions available to the players when making their moves, while not explicit in the definition, may be deduced from the set of terminal histories. For any history h , the set of actions available to $P(h)$, the player who moves after h , is the set of actions a for which (h, a) is a history. We denote this set $A(h)$.

Two simple examples of extensive games with perfect information are shown in Figure 6. In the game on the left, the set of terminal histories is $\{(X, w), (X, x), (Y, y), (Y, z)\}$ and the player function assigns player 1 to the empty history (a subsequence of every terminal history) and player 2 to the histories X and Y . The game begins with player 1's choosing either X or Y . If she chooses X , then player 2 chooses either w or x ; if she chooses Y , then player 2 chooses either y or z . In the game on the right, the set of terminal histories is $\{(W, x, Y), (W, x, Z), (W, y), X\}$ and the player function assigns player 1 to the empty history and the history (W, x) , and player 2 to the history W .




Figure 6. Two extensive games with perfect information. Player 1's payoff is the first number in each pair.

Another example of an extensive game with perfect information is a sequential variant of Cournot's model of oligopoly in which firm 1 chooses an output, then firm 2 chooses an output, and so on. In this game, the set of terminal histories is the set of all sequences (q_1, \dots, q_n) of outputs for the firms; the player function assigns player 1 to the empty history and, for $k = 1, \dots, n - 1$, player $k + 1$ to every sequence (q_1, \dots, q_k) . (Because a continuum of actions is available after each nonterminal history, this game cannot easily be represented by a diagram like those in [Figure 6](#).)

A further example is the bargaining game of alternating offers. This game has terminal histories of infinite length (those in which every offer is rejected).

3.1.1 Strategies A key concept in the analysis of an extensive game is that of a *strategy*. The definition is very simple: a strategy of any player j is a function that associates with *every* history h after which player j moves a member of $A(h)$, the set of actions available after h .

Definition 12 A *strategy* of player j in an extensive game with perfect information $\langle N, H, P, (\succsim_i) \rangle$ is a function that assigns to every history h (subsequence of H) for which $P(h) = j$ an action in $A(h)$.

In the game at the left of [Figure 6](#), player 1 has two strategies, X and Y . Player 2 has *four* strategies, which we may represent by wy , wz , xy , and xz , where the first component in each pair is the action taken after the history X and the second component is the action taken after the history Y . This example illustrates that a strategy is a complete plan of action, specifying the player's action in every eventuality. Before the game begins, player 2 does not know whether player 1 will choose X or Y ; her strategy prepares her for both eventualities.

The game at the right of [Figure 6](#) illustrates another aspect of the definition. Player 1 in this game has *four* strategies, WY , WZ , XY , and XZ . In particular, XY and XZ are distinct strategies. (Remember that a player's strategy assigns an action to *every* history after which she moves.) I discuss the interpretation of strategies like these in [Section 3.1.3](#).

3.1.2 Nash equilibrium A Nash equilibrium of an extensive game with perfect information is defined in the same way as a Nash equilibrium of a strategic game: it is a strategy profile with the property that no player can increase her payoff by changing her strategy, given the other players' strategies. Precisely, first define the *outcome* $O(s)$ of a strategy profile s to be the terminal history that results when the players use s . (The outcome $O(X, wy)$ of the strategy pair (X, wy) in the game on the left of [Figure 6](#), for example, is the terminal history (X, w) .)

Definition 13 A *Nash equilibrium* of the extensive game with perfect information $\langle N, H, P, (\succsim_i) \rangle$ is a strategy profile s^* for which

$$O(s^*) \succsim_i O(s_i, s_{-i}^*) \text{ for all } s_i \in S_i$$

for every player $i \in N$, where S_i is player i 's set of strategies.

As an example, the game on the left of [Figure 6](#) has three Nash equilibria, (X, wy) , (X, wz) , and (Y, xy) . (One way to find these equilibria is to construct a table like the one in [Figure 1](#) in which each row is a strategy of player 1 and each column is a strategy of player 2.)

For each of the last two equilibria, there exists a history h such that the action specified by player 2's strategy after h is not optimal for her in the rest of the game. For example, in the last equilibrium, player 2's strategy specifies that she will choose x after the history X , whereas only w is optimal for her after this history. Why is such a strategy optimal? Because player 1's strategy calls for her to choose Y , so that the action player 2 plans to take after the history X has no effect on the outcome: the terminal history is (Y, y) regardless of player 2's action after the history X .

I argue that this feature of the strategy pair (Y, xy) detracts from its status as an equilibrium. Its equilibrium status depends on player 1's believing that if she deviates to X then player 2 will choose x . Given that only w is optimal for player 2 after the history X , such a belief seems unreasonable.

Suppose that player 1 forms her belief on the basis of her experience. If she *always* chooses Y , then no amount of experience will enlighten her regarding player 2's choice after the history X . However, in a slightly perturbed steady state in which she very occasionally erroneously chooses X at the start of the game

and player 2 chooses her optimal action whenever called upon to move, player 1 knows that player 2 chooses w , not x , after the history X .

If player 1 bases her belief on her reasoning about player 2's rational behavior (in the spirit of rationalizability), she reaches the same conclusion. (Note, however, that this reasoning process is straightforward in this game only because the game has a finite horizon and one player is indifferent between two terminal histories if and only if the other player is also indifferent.)

In either case, we conclude that player 1 should believe that player 2 will choose w , not x , after the history X . Similarly, the Nash equilibrium (X, wz) entails player 1's unreasonable belief that player 2 will choose z , rather than y , after the history Y . We now extend this idea to all extensive games with perfect information.

3.1.3 Subgame perfect equilibrium A *subgame perfect equilibrium* is a strategy profile in which each player's strategy is optimal not only at the start of the game, but also after every history. (The notion is due to [Selten 1965](#).)

Definition 14 A *subgame perfect equilibrium* of the extensive game with perfect information $\langle N, H, P, (\succsim_i) \rangle$ is a strategy profile s^* for which

$$O_h(s^*) \succsim_i O_h(s_i, s_{-i}^*) \text{ for all } s_i \in S_i$$

for every player $i \in N$ and every history h after which it is player i 's turn to move (i.e. $P(h) = i$), where S_i is player i 's set of strategies and $O_h(s)$ is the terminal history consisting of h followed by the sequence of actions generated by s after h .

For any nonterminal history h , define the *subgame* following h to be the part of the game that remains after h has occurred. With this terminology, we have a simple result: a strategy profile is a subgame perfect equilibrium if and only if it induces a Nash equilibrium in every subgame. Note, in particular, that a subgame perfect equilibrium is a Nash equilibrium of the whole game. (The function O in [Definition 13](#) is the same as the function O_\emptyset in [Definition 14](#), where \emptyset denotes the empty history.) The converse is not true, as we have seen: in the game at the left of [Figure 6](#), player 2's only optimal action after the history X is w and her only optimal action after the history Y is y , so that the game has a single subgame perfect equilibrium, (X, wy) , whereas it has three Nash equilibria.

Now consider the game at the right of [Figure 6](#). Player 1's only optimal action after the history (W, x) is Y ; given that player 1 chooses Y after (W, x) , player 2's only optimal action after the history W is x ; and given that player 2 chooses x after the history W , player 1's only optimal action at the start of the game is X . Thus the game has a unique subgame perfect equilibrium, (XY, x) .

Note, in particular, that player 1's strategy XZ , which generates the same outcome as does her strategy XY regardless of player 2's strategy, is *not* part of a subgame perfect equilibrium. That is, the notion of subgame perfect equilibrium differentiates between these two strategies even though they correspond to the same "plan of action". This observation brings us back to a question raised in [Section 3.1.1](#): how should the strategies XZ and XY be interpreted?

If we view a subgame perfect equilibrium as a model of a perturbed steady state in which every player occasionally makes mistakes, the interpretation of player 1's strategy XY is that she chooses X at the start of the game, but if she erroneously chooses W and player 2 subsequently chooses x , she chooses Y . More generally, a component of a player's strategy that specifies an action after a history h precluded by the other components of the strategy is interpreted to be the action the player takes if, after a series of mistakes, the history h occurs. Note that this interpretation is strained in a game in which some histories occur only after a long series of mistakes, and thus are extremely unlikely.

In some finite horizon games, we may alternatively interpret a subgame perfect equilibrium to be the outcome of the players' calculations about each other's optimal actions. If no player is indifferent between any two terminal histories, then every player can deduce the actions chosen in every subgame of length 1 (at the end of the game); she can use this information to deduce the actions chosen in every subgame of length 2; and she can similarly work back to the start of every subgame at which she has to choose an action. Under this interpretation, the component Y of the strategy XY in the game at the right of [Figure 6](#) is player 2's belief about player 1's action after the history (W, x) and also player 1's belief about the action player 2 believes player 1 will choose after the history (W, x) . (This interpretation makes sense also under the weaker condition that whenever one player is indifferent between the outcomes of two actions, every other player is also indifferent (a sufficient condition for each player to be able to deduce her payoff when the other players act optimally, even if she cannot deduce the other players' strategies).)

This interpretation, like the previous one, is strained in some games. Consider the game that differs from the one at the right of [Figure 6](#) only in that player 1's payoff of 3 after the history (W, y) is replaced by 1. The unique subgame perfect equilibrium of this game is (XY, x) (as for the original game). The equilibrium entails player 2's belief that player 1 will choose Y if player 2 chooses x after player 1 chooses W . But choosing W is inconsistent with player 1's acting rationally: she guarantees herself a payoff of 2 if she chooses X , but can get at most 1 if she chooses W . Thus it seems that player 2 should either take player 1's action W as an indication that player 1 believes the game to differ from the game that player 2 perceives, or view the action as a mistake. In the first case the way

in which player 2 should form a belief about player 1's action after the history (W, x) is unclear. The second case faces difficulties in games with histories that occur only after a long series of mistakes, as for the interpretation of a subgame perfect equilibrium as a perturbed steady state.

The subgame perfect equilibria of the games in Figure 6 may be found by working back from the end of the game, isolating the optimal action after any history given the optimal actions in the following subgame. This procedure, known as *backward induction*, may be used in any finite horizon game in which no player is indifferent between any two terminal histories. A modified version that deals appropriately with indifference may be used in any finite horizon game.

3.2 Extensive games with imperfect information

In an extensive game with perfect information, each player, when taking an action, knows all actions chosen previously. To capture situations in which some or all players are not perfectly informed of past actions we need to extend the model. A general extensive game allows arbitrary gaps in players' knowledge of past actions by specifying, for each player, a partition of the set of histories after which the player moves. The interpretation of this partition is that the player, when choosing an action, knows only the member of the partition in which the history lies, not the history itself. Members of the partition are called *information sets*. When choosing an action, a player has to know the choices available to her; if the choices available after different histories in a given information set were different, the player would know the history that had occurred. Thus for an information partition to be consistent with a player's not knowing which history in a given information set has occurred, for every history h in any given information set, the set $A(h)$ of available actions must be the same. We denote the set of actions available after the information set I_i by $A(I_i)$.

Definition 15 An *extensive game* consists of

- a set N (the set of **players**)
- a set H of sequences (the set of **terminal histories**) with the property that no sequence is a proper subhistory of any other sequence
- a function P (the **player function**) that assigns a player to every proper subsequence of every terminal history

and for each player $i \in N$

- a partition \mathcal{I}_i of the set of histories assigned to i by the player function (player i 's **information partition**) such that for every history h in any given member of the partition (**information set**), the set $A(h)$ of actions available is the same
- a preference relation \succsim_i over the set H of terminal histories.

(A further generalization of the notion of an extensive game allows for events to occur randomly during the course of play. This generalization involves no significant conceptual issue, and I do not discuss it.)

An example is shown in Figure 7. The dotted line indicates that the histories X and Y are in the same information set: player 2, when choosing between x and y , does not know whether the history is X or Y . (Formally, player 2's information partition is $\{\{X, Y\}, \{Z\}\}$. Notice that $A(X) = A(Y) (= \{x, y\})$, as required by the definition.)




Figure 7. An extensive game with imperfect information. The dotted line indicates that the histories X and Y are in the same information set.

A strategy for any player j in an extensive game associates with each of her information sets I_j a member of $A(I_j)$.

Definition 16 A **strategy** of player j in an extensive game $\langle N, H, P, (\mathcal{I}_i), (\succsim_i) \rangle$ is a function that assigns to every information set $I_j \in \mathcal{I}_j$ of player j an action in $A(I_j)$.

Given this definition, a Nash equilibrium is defined exactly as for an extensive game with perfect information (Definition 13)—and, as before, is not a satisfactory solution. Before discussing alternatives, we need to consider the possibility of players' randomizing.

In an extensive game with perfect information, allowing players to randomize does not significantly change the set of equilibrium outcomes. In an extensive game with imperfect information, the same is not true. A straightforward way of incorporating the possibility of randomization is to follow the theory of strategic

games and allow each player to choose her strategy randomly. That is, we may define a mixed strategy to be a probability distribution over (pure) strategies. An approach more directly suited to the analysis of an extensive game is to allow each player to randomize independently at each information set. This second approach involves the notion of a behavioral strategy, defined as follows.

Definition 17 A *behavioral strategy* of player j in an extensive game $\langle N, H, P, (\mathcal{I}_i), (\zeta_i) \rangle$ is a function that assigns to each information set $I_j \in \mathcal{I}_j$ a probability distribution over the actions in $A(I_j)$, with the property that each probability distribution is independent of every other distribution.

For a large class of games, mixed strategies and behavioral strategies are equivalent: for every mixed strategy there exists a behavioral strategy that yields the same outcome regardless of the other players' strategies, and vice versa. (This result is due to [Kuhn 1950, 1953](#).) In the discussion that follows, I work with behavioral strategies.

The notion of subgame perfect equilibrium for an extensive game with perfect information embodies two conditions: whenever a player takes an action, (a) this action is optimal given her belief about the other players' strategies and (b) her belief about the other players' strategies is correct. In such a game, each player needs to form a belief only about the other players' future actions. In an extensive game with imperfect information, players need also to form beliefs about the other player's *past* actions. Thus in order to impose condition *b* on a strategy profile in an extensive game with imperfect information, we need to consider how a player choosing an action at an information set containing more than one history forms a belief about which history has occurred and what it means for such a belief to be correct.

Consider the game in [Figure 7](#). If player 1's strategy is X or Y , then the requirement that player 2's belief about the history be correct is easy to implement: if player 1's strategy specifies X then she believes X has occurred, whereas if player 1's strategy specifies Y then she believes Y has occurred. If player 1's strategy is Z , however, this strategy gives player 2 no basis on which to form a belief—we cannot derive from player 1's strategy a belief of player 2 about player 1's action. The main approach to defining equilibrium avoids this difficulty by specifying player 1's belief as a component of an equilibrium. Precisely, we define a belief system and an assessment as follows.

Definition 18 A *belief system* is a function that assigns to every information set a probability distribution over the set of histories in the set. An *assessment* is a pair consisting of a profile of behavioral strategies and a belief system.

We may now define an equilibrium to be an assessment satisfying conditions a and b . To do so, we need to decide exactly how to implement b . One option is to require consistency of beliefs with strategies only at information sets reached if the players follow their strategies, and to impose no conditions on beliefs at information sets not reached if the players follow their strategies. The resulting notion of equilibrium is called a weak sequential equilibrium. (The name “perfect Bayesian equilibrium” is sometimes used, although the notion with this name defined by [Fudenberg and Tirole \(1991\)](#) covers a smaller class of games and imposes an additional condition on assessments.)

Definition 19 *An assessment (β, μ) , where β is a behavioral strategy profile and μ is a belief system, is a **weak sequential equilibrium** if it satisfies the following two conditions.*

Sequential rationality *Each player's strategy is optimal in the part of the game that follows each of her information sets, given the other players' strategies and her belief about the history in the information set that has occurred. Precisely, for each player i and each information set I_i of player i , player i 's expected payoff to the probability distribution over terminal histories generated by her belief μ_i at I_i and the behavior prescribed subsequently by the strategy profile β is at least as large as her expected payoff to the probability distribution over terminal histories generated by her belief μ_i at I_i and the behavior prescribed subsequently by the strategy profile (γ_i, β_{-i}) , for each of her behavioral strategies γ_i .*

Weak consistency of beliefs with strategies *For every information set I_i reached with positive probability given the strategy profile β , the probability assigned by the belief system to each history h in I_i is the probability of h occurring conditional on I_i being reached, as given by Bayes' law.*

Consider the game in [Figure 7](#). Notice that player 2's action x yields her a higher payoff than does y regardless of her belief. Thus in any weak sequential equilibrium she chooses x with probability 1. Given this strategy, player 1's only optimal strategy assigns probability 1 to Y . Thus the game has a unique weak sequential equilibrium, in which player 1's strategy is Y , player 2's strategy is x , and player 2's belief assigns probability 1 to the history Y .

Now consider the game in [Figure 8](#). I claim that the assessment in which player 1's strategy is $(\frac{1}{2}, \frac{1}{2}, 0)$, player 2's strategy is $(\frac{1}{2}, \frac{1}{2})$, and player 2's belief assigns probability $\frac{1}{2}$ to X and probability $\frac{1}{2}$ to Y is a weak sequential equilibrium. Given her beliefs, player 2's expected payoffs to x and y are both 2, and given player 2's strategy, player 1's expected payoffs to X and Y are both $\frac{5}{2}$ and her payoff to Z is 2. Thus each player's strategy is sequentially rational. Further, player 2's

belief is consistent with player 1's strategy. This game has an additional weak sequential equilibrium in which player 1's strategy is Z , player 2's strategy is y , and player 2's belief assigns probability 1 to the history Y . Note that the consistency condition does not restrict player 2's belief in this equilibrium, because player 1 chooses neither X nor Y with positive probability.




Figure 8. An extensive game with imperfect information.

In some games the notion of weak sequential equilibrium yields sharp predictions, but in others it is insufficiently restrictive. Some games, for example, have weak sequential equilibria that do not satisfy a natural generalization of the notion of subgame perfect equilibrium. In response to these problems, several “refinements” of the notion of a weak sequential equilibrium have been studied, including sequential equilibrium (due to [Kreps and Wilson 1982](#)) and perfect Bayesian equilibrium (due to [Fudenberg and Tirole 1991](#)).

References

- Aumann, Robert J. (1974), “Subjectivity and correlation in randomized strategies.” *Journal of Mathematical Economics*, 1, 67–96. [[10](#)]
- Bernheim, B. Douglas (1984), “Rationalizable strategic behavior.” *Econometrica*, 52, 1007–1028. [[11](#)]
- Cournot, Antoine A. (1838), *Recherches sur les principes mathématiques de la théorie des richesses*. Hachette, Paris. English translation: *Researches into the mathematical principles of the theory of wealth*, New York: Macmillan, 1897. [[3](#), [4](#)]
- Fudenberg, Drew and Jean Tirole (1991), “Perfect Bayesian equilibrium and sequential equilibrium.” *Journal of Economic Theory*, 53, 236–260. [[24](#), [25](#)]

- Glicksberg, I. L. (1952), "A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points." *Proceedings of the American Mathematical Society*, 3, 170–174. [8]
- Harsanyi, John C. (1967/68), "Games with incomplete information played by 'Bayesian' players, Parts I, II, and III." *Management Science*, 14, 159–182, 320–334, and 486–502. [13]
- Harsanyi, John C. (1973), "Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points." *International Journal of Game Theory*, 2, 1–23. [8]
- Hotelling, Harold (1929), "Stability in competition." *Economic Journal*, 39, 41–57. [3]
- Kreps, David M. and Robert B. Wilson (1982), "Sequential equilibria." *Econometrica*, 50, 863–894. [25]
- Kuhn, Harold W. (1950), "Extensive games." *Proceedings of the National Academy of Sciences of the United States of America*, 36, 570–576. [16, 23]
- Kuhn, Harold W. (1953), "Extensive games and the problem of information." In *Contributions to the theory of games, Volume II (Annals of Mathematics Studies, 28)* (Harold W. Kuhn and A. W. Tucker, eds.), 193–216, Princeton University Press, Princeton, NJ. [16, 23]
- Nash, Jr., John F. (1950), "Equilibrium points in N -person games." *Proceedings of the National Academy of Sciences of the United States of America*, 36, 48–49. [4, 8]
- Nash, Jr., John F. (1951), "Non-cooperative games." *Annals of Mathematics*, 54, 286–295. [8]
- Pearce, David G. (1984), "Rationalizable strategic behavior and the problem of perfection." *Econometrica*, 52, 1029–1050. [11]
- Reny, Philip J. (1999), "On the existence of pure and mixed strategy Nash equilibria in discontinuous games." *Econometrica*, 67, 1029–1056. [8]
- Selten, Reinhard (1965), "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit." *Zeitschrift für die gesamte Staatswissenschaft*, 121, 301–324 and 667–689. [19]

von Neumann, John and Oskar Morgenstern (1944), *Theory of games and economic behavior*. Princeton University Press, Princeton, NJ. [[6](#), [16](#)]

von Neumann, John and Oskar Morgenstern (1947), *Theory of games and economic behavior (2nd edition)*. Princeton University Press, Princeton, NJ. [[6](#)]