FOUNDATIONS OF NON-COOPERATIVE GAMES

Ron Holzman, Department of Mathematics, Technion-Israel Institute of Technology, Israel

Keywords: Bayesian game, Bayesian Nash equilibrium, Chess-like game, Consistency, Determinacy, Extensive form, Information, Nash equilibrium, Non-cooperative game, Non-zero-sum game, Payoff function, Rationality, Strategic form, Strategy, Subgame perfect equilibrium, Value, Zero-sum game.

Contents

Glossary

Bayesian game: A model for representing a game with incomplete information in which the players have beliefs about the state of the world which they update upon receiving their private information.
Bayesian Nash equilibrium: A strategy profile in a Bayesian game having the property that every player, given any of his possible types, chooses an action that maximizes his expected payoff in response to the other players' strategies.
Behavior strategy: A strategy in which the player's choices during the play of the game are allowed to be random but must be mutually independent.
Complete information: Qualifies a situation where the description of the game is common knowledge among the players.
Extensive form: A model for representing a game which includes a complete account of the decisions to be made by the players, the order in which they occur, the information available to the players at each stage, the distributions of the chance events, and the payoffs to the players for any possible play.
Game-tree: A rooted tree used in an extensive form representation of a game.
Information set: A set of decision nodes of a player which are indistinguishable to him at the stage when he has to make a choice at any one of them.
Maximin/minimax strategy: A strategy that is a player's best choice from the point of

view of the worst-case scenario.
Mixed extension: A game obtained from another game by allowing the players to use mixed strategies.
Mixed strategy: A strategy that consists of choosing at random, according to some distribution, one of the player's (pure) strategies.
Nash equilibrium: A strategy profile having the property that every player's strategy maximizes his payoff given the other players' strategies.
Perfect information: Qualifies a situation where every time a player has to make a decision he is fully aware of all the decisions and all the outcomes of chance moves that took place in the course of the play up to that stage.
Perfect recall: A weaker requirement than perfect information, where a player needs only to be fully aware of his own previous decisions.
State of the world: A full description of the parameters of the game and the beliefs of the players in a situation of incomplete information.
Strategic form: An abstract model for representing a game which specifies only the players' possible strategies and the payoffs they obtain for each strategy profile.
Strategy: A way for a player to play a game. In the case of a game in extensive form, it specifies the player's choices at each of his information sets.
Strategy profile: A choice of strategies, one for each player.
Subgame perfect equilibrium: A Nash equilibrium of a game in extensive form that remains a Nash equilibrium when restricted to any subgame.
Type: A full description of the attributes of a player, including his beliefs about the parameters of the game and the other players' beliefs, in a situation of incomplete information.
Value: In a zero-sum game, a number v such that player 1 can guarantee that he will receive at least v and player 2 can guarantee that he will pay at most v.
Winning strategy: In a chess-like game, a strategy that guarantees that the player who uses it will win.
Zero-sum game: A two-person game in which the two players' payoffs always add up to zero.

Summary

This article introduces the basic models and ideas of the theory of non-cooperative games. We begin by treating games in the usual sense of the word, such as chess. We show that for a certain class of games, the outcome is completely determined if the players play optimally. Then we indicate how the descriptive framework of game theory, including the extensive and strategic form representations, can serve to model interactions between agents which do not qualify as games in the usual sense. For zero-sum games, where one player's gain is the other's loss, we introduce the concept of value, which is the expected outcome when the game is played optimally. If players are allowed to use mixed (i.e., randomized) strategies, the minimax theorem asserts that the value exists. Non-zero-sum games are more complex, and we cannot hope to pinpoint their expected outcome as in the zero-sum case. The central concept for these games is that of a Nash equilibrium, which is a choice of strategies for the players having the property that every player does best by playing his strategy if the others do the same. Nash's theorem guarantees the existence of a Nash equilibrium in mixed strategies. Finally, we turn to the modeling of incomplete information, which occurs when the players lack information about the game they are facing. We present the concepts of a "state of the world" and the "type" of a player, and show how they are incorporated in the Bayesian game model.

1.Introduction

Non-cooperative game theory studies situations in which a number of agents are involved in an interactive process whose outcome is determined by the agents' individual decisions (sometimes in conjunction with chance) and affects the well-being of each agent in a possibly different way. The most obvious examples of such situations are parlor games, and the terminology we use has its roots in this area: the entire situation is called a game, the agents are called players, their acts are called moves, their overall plans of action are called strategies, and their evaluations of the outcome are called payoffs. But the range of situations that we have in mind is much wider, and includes interactions in areas such as economics, politics, biology, computing, etc. Thus, the significance of non-cooperative game theory to the understanding of social and natural phenomena is far bigger than its name may suggest.

The basic premise of our analysis is that players act rationally, meaning first and foremost that they strive to maximize their payoffs. However, since their payoffs are affected not only by their own decisions but also by the other players' decisions, they must reason about the other players' reasoning, and in doing so they take into account that the other players, too, act rationally.

The qualification "non-cooperative" refers to the assumption that players make their decisions individually, and are not allowed to forge binding agreements with other players that stipulate the actions to be taken by the parties to the agreement. The players may be allowed to communicate with each other prior to the play of the game and discuss joint plans of action. But during the play of the game they act as autonomous decision makers, and as such they will follow previously made joint plans only if doing so is rational for them.

The theory of non-cooperative games comprises three main ingredients. The first of these is the development of formal models of non-cooperative games that create unified frameworks for representing games in a manner that lends itself to formal mathematical analysis. The second ingredient is the formulation of concepts that capture the idea of rational behavior in those models. The main such concept is that of equilibrium. The third ingredient is the use of mathematical tools in order to prove meaningful statements about those concepts, such as existence and characterizations of equilibrium.

In any concrete application of the theory, the first step is to represent the situation at hand by one of the available formal models. Because real-life situations are typically very complex and not totally structured, it is often impossible and/or unhelpful to incorporate all the elements of the situation in the formal model. Therefore, this step requires judicious decisions identifying those important features that must be modeled. Once the model is constructed, its analysis is carried out based on the appropriate concept of equilibrium, drawing on general results of non-cooperative game theory or, as the case may be, exploiting attributes of the specific application. This analysis yields conclusions which may then be reformulated in terms of the real-life situation, providing insights into, or predictions about, the behavior of the agents and the long-term steady states of the system being investigated.

Another sort of application is sometimes called game theoretic engineering. It involves making recommendations to organizations on how to set up the "rules of the game" so that the rational behavior of the agents will lead to results that are desirable from the organization's point of view. Examples include the revision of electoral systems, the design of auctions, the creation of markets for emission permits as a means to efficiently control pollution, etc. (see *Mechanism Theory*). This sort of application seems to be gaining more and more recognition lately.

This article is organized according to several standard criteria for classifying non-cooperative games: how the payoffs of the players are related, the presence or the absence of chance, the nature of the information that the players have.

We do not attempt to present here a comprehensive survey of non-cooperative game theory. The omission of certain parts of the theory, even important ones, is unavoidable in such an article. Some of these areas are covered in other articles within this topic. What we try to do here is offer a gentle introduction to a small number of basic models, ideas and concepts.

2.Chess-Like Games

2.1.The Description of the Game

The game of chess is a prime example of a class of games that we call *chess-like games*. We first give a verbal description of what we mean by a chess-like game.

In such a game there are two players who take turns making moves. One of the players is designated as the one who starts, we call this player White and the other player Black. Whenever a player chooses a move, he is perfectly informed of all moves made prior to that stage. The play of the game is fully determined by the players' choices, that is, it does not involve any chance. For every initial sequence of moves made alternatingly by the two players, the rules of the game determine whether the player whose turn it is to play should choose a move—in which case they also determine what his legal moves are—or whether the play has ended—in which case they also determine the outcome: a win for White, a win for Black, or a draw. It may be the case that the rules allow for a play consisting of an infinite sequence of moves, but such infinite plays must also be classified as resulting in one of the three possible outcomes mentioned above.

Examples of chess-like games are chess, checkers, tic-tac-toe. Note that Kriegspiel (a version of chess in which a player does not observe his opponent's moves) is not a chess-like game, due to the lack of perfect information. Backgammon is not a chess-like game because it involves dice.

We formally describe a chess-like game by means of a rooted tree (T, r), called the *game-tree*. Here T is a tree (that is, a connected acyclic graph which may be finite or infinite), and r, the root, is a designated node of T. We think of T as being oriented away from r. The edges that fan out from r correspond to the legal moves of White at his first turn. Each of these edges leads to a new node, and the edges that fan out from this new node correspond to the legal moves of Black after White chose the specified edge from r as his first move. It goes on like this: non-terminal nodes whose distance from r is even (respectively odd) are decision nodes of White (respectively Black), and the edges that fan out correspond to the legal moves of the respective player given the moves made so far. A maximal branch of (T, r) is a path that starts at the root r and either ends at a terminal node or is infinite. Every maximal branch corresponds to a play of the game. (Note the distinction we make between the usage of "game" and "play". A game is the totality of rules that define it, whereas a play of the game is a complete account of what happens a particular time when the game is played.) To complete the formal description of the game, we specify for each maximal branch whether it is a win for White, a win for Black, or a draw.

The realization that every chess-like game can in principle be represented by a game-tree as above, even though for most games constructing the actual game-tree is impractical, is an important conceptual step towards the analysis of such games.

2.2.The Determinacy of Chess-Like Games

The next important concept is that of a *strategy.* By a strategy we mean a complete set of instructions that tell a player what to do in every situation that may arise in the play of the game in which he is called upon to make a move. In terms of the formal description of a chess-like game by a game-tree, a strategy of a player is a function σ from the set of his decision nodes into the set of edges, such that $\sigma(x)$ is one of the edges that fan out of x.

It is clear from the above that a pair of strategies, σ for White and τ for Black, fully determines a play of the game, and in particular any such pair (σ, τ) results in a win for White, a win for Black, or a draw. In effect, the comprehensive definition of the concept of strategy renders the actual play of the game unnecessary, in principle: We may imagine the players announcing (simultaneously) their respective strategies to a referee or a machine, who can determine right away the outcome based on the announced strategies.

A strategy $\overline{\sigma}$ of White is a *winning* strategy, if for every strategy τ of Black, the pair $(\overline{\sigma}, \tau)$ results in a win for White. In other words, $\overline{\sigma}$ guarantees a win for White. A winning strategy for Black is defined similarly. A *drawing* strategy for a player is a strategy that guarantees at least a draw for that player, that is, using that strategy he will win or draw against any strategy of his opponent.

The following theorem was discovered by John von Neumann (but is often referred to, wrongly, as Zermelo's theorem).

Theorem 1. Let G be a chess-like game in which the length of a play is finitely bounded (i.e., there exists $M < \infty$ such that in every play of the game there are at most M moves). Then one of the following statements is true:
(a) White has a winning strategy in G.
(b) Black has a winning strategy in G.
(c) Both White and Black have drawing strategies in G.

It is important to understand that Theorem 1 does not merely state the tautological fact that every play of the game results in one of the three possible outcomes. Rather, the theorem asserts that every *game* (as opposed to every play of a game) satisfying the theorem's assumptions may be classified as a win for White, a win for Black, or a draw, in the sense that it will always end that way when played optimally. We refer to this property of a game G as *determinacy.*

The proof of Theorem 1 proceeds by "backward induction". Namely, one classifies every subgame starting at a terminal node of the game-tree (this is trivial), then every subgame starting at a node from which all moves lead to terminal nodes, and so on, working one's way to the root of the tree. For every subgame encountered in this process, it is classified as satisfying (a), (b), or (c), and at the end of the process G itself is classified. Thus, the proof of Theorem 1 is constructive, leading to an algorithm for classifying a game and finding a winning strategy for the player who has one (or drawing strategies for both players). Nevertheless, for most real-life games, the computational complexity of constructing the game-tree, let alone running this algorithm, is much too high. Thus, although Theorem 1 renders the games it applies to uninteresting, in principle, to players with unbounded computational abilities, in practice the games still hold interest to humans. The issues related to discovering winning strategies and the complexity of this task are studied within the field of "combinatorial games", which grew quite independently of the rest of game theory.

Incidentally, the question of whether or not the game of chess itself satisfies the assumptions of Theorem 1 depends on a careful scrutiny of the rules of chess regarding, e.g., the repetition of positions on the board. Assuming an interpretation of the rules that makes the length of a play of chess finitely bounded, we can conclude from Theorem 1 that chess is either a win for White, a win for Black, or a draw. But nobody knows which of these it is!

There are results extending Theorem 1 to certain classes of chess-like games with infinite plays. These results, starting from Gale and Stewart's theorem on "open games", depend on topological assumptions on the set of maximal branches of the game-tree which constitute a win for a given player. Using the Axiom of Choice, Gale and Stewart also showed that there exist infinite chess-like games which are undetermined, that is, for which the conclusion of Theorem 1 is false. This direction of research has revealed strong connections with the foundations of mathematics, and has grown quite independently of game theory, in the field called "descriptive set theory".

3.Representations of Non-Cooperative Games

3.1.An Informal Description of the Class of Games

The games that we consider here are more general than the chess-like games considered above in several respects:

- There is a finite number n of players, possibly $n > 2$.
- The order in which the players are called upon to make a move is arbitrary, and the identity of the player who has to move at a given stage may depend on the moves made up to that stage.
- Information may be imperfect, meaning that at the time when a player is called upon to make a move, he may have only partial information on the moves made prior to that time. This includes also the possibility of simultaneous moves, which may be represented as sequential moves with the provision that the player who moves later must do so without being informed of the choice of the player who moved earlier.
- There may be chance moves, that is, moves not controlled by any player but rather selected from a set of possible moves according to some probability distribution.
- The outcome associated with any play of the game, rather than being a win for some player or a draw, is represented by an n-tuple of real numbers $(u_1, \ldots, u_n)$, where u_i measures the utility that player i derives from the outcome. This permits to represent the outcomes of chess, for instance as $(1, 0)$, $(0, 1)$, or $(\frac{1}{2}, \frac{1}{2})$, and also allows for much more general situations, as we will see below.

This more flexible framework includes a variety of parlor games like Kriegspiel, backgammon, bridge, poker, monopoly. More importantly, many real-life situations which are not normally thought of as "games", may usefully be modeled as such. Examples include: competition between firms in an oligopolistic market, campaigns of opposing candidates running for election, struggle between genes as part of evolution, interaction between processors involved in a parallel computation, and more.

There is, however, one important implicit assumption about the games we consider here, the validity of which must be assessed for any real-life application. This is the assumption of *complete information*, meaning that a player knows the entire description of the game, and moreover he knows that every other player knows the entire description of the game, and furthermore he knows that every other player knows that every other player knows the entire description of the game, and so forth. This condition is expressed concisely by saying that the description of the game is common knowledge.

This concept should be distinguished from the concept of perfect information: while perfect information pertains to knowledge of what happened in the current play of the game, complete information pertains to knowledge of the game itself (its rules, the relevant probability distributions, who is informed of what and when, the utilities of the various outcomes to the various players). For example, bridge players who have mastered the rules of the game are engaged in a game of complete information which has, however, imperfect information: a player is not informed of the cards dealt to the other players. In modeling real-life situations, the assumption of complete information is more problematic. We deal with the modeling of incomplete information in Section 6.

3.2.The Extensive Form

The extensive form representation of a non-cooperative game uses a game-tree as in the case of chess-like games, but with additional structure. We present the ingredients of this representation one by one:

- A finite non-empty set of players N. Without loss of generality $N = \{1, \dots, n\}$.
- A rooted tree (T, r), where T is a tree with node-set V and edge-set E, and r is a node in V. The nodes represent positions that may arise in the course of the play, with r being the initial position. The edges represent moves in the game. We denote by E_x the set of edges that fan out of node x, that is, the set of possible moves at x.
- An ordered partition $(V_0, V_1, \dots, V_n)$ of the set of non-terminal nodes in V. For every $i \in N$, the set V_i consists of the decision nodes of player i. The set V_0 contains the nodes at which the move is selected randomly.
- For every $i \in N$, there is a partition Π_i of V_i with the following property: whenever $x, y \in I \in \Pi_i$, we have $|E_x| = |E_y|$. The sets $I \in \Pi_i$ are called *information sets* of player i. The interpretation of an information set is the following. Suppose we have $x \in I \in \Pi_i$. When the position x is reached, player i is called upon to choose a move without knowing that the actual position is x; he is only informed that the current position is an element of I. In other words, player i can distinguish between his different information sets, but not between positions in the same information set.
- For every $i \in N$ and every $I \in \Pi_i$, there is a set L_I of cardinality equal to that of E_x for every x in I, and a family of bijections $l_x : E_x \to L_I$, one for each x in I. The role of the set L_I and the bijections l_x is to attach a label to every move at every position in the information set I. This is necessary in order to enable player i to specify a chosen move at an unknown node within I.
- For every $x \in V_0$ there is a probability distribution P_x on the set E_x. It is understood that when the position x is reached, a move in E_x is selected at random according to the

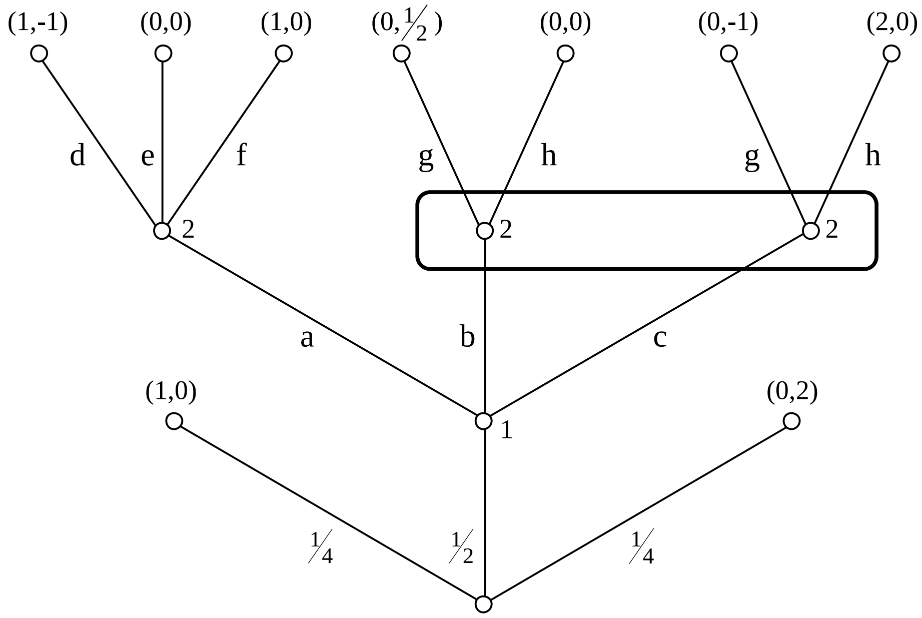

Figure 1: A two-person game in extensive form

distribution P_x, and that the choices at different nodes in V_0 are made independently of each other.

- To every maximal branch of (T, r) there is assigned an n-tuple $(u_1, \ldots, u_n) \in \Re^n$. Here u_i is the utility that player i derives from a play of the game along the given maximal branch. It is understood that these are von Neumann-Morgenstern utilities, meaning that the utility derived from a random play with a specified distribution equals the expected utility according to that distribution. The n-tuples $(u_1, \ldots, u_n)$ are called *payoff vectors.*

A game described as above is called an n-person *game in extensive form.* It is said to be *finite* if the tree T is finite. It is a game with *perfect information* if all the information sets are singletons. A weaker property is that of *perfect recall.* It requires that if two decision nodes x and y of player i are in the same information set, then the sequences of decisions of player i himself involved in reaching the two positions x and y must be identical.

It is customary to represent finite games graphically by drawing the game-tree with the following conventions:

- The players' names or numbers appear next to their decision nodes.
- Decision nodes belonging to the same information set are encircled together.
- Labels or probabilities attached to moves appear next to the corresponding edges.
- Payoff vectors appear next to the terminal nodes of the corresponding maximal branches.

We illustrate these conventions by depicting an example in Figure 1. The game is verbally described as follows. First, a lottery takes place. With probability $\frac{1}{4}$, the play ends with payoff

1 to player 1 and 0 to player 2. With probability $\frac{1}{4}$, the play ends with payoff 0 to player 1 and 2 to player 2. With probability $\frac{1}{2}$, player 1 has to choose between a, b, and c. If player 1 chose a, player 2 is informed of this and has to choose between d, e, and f, whereupon the play ends with the payoff vectors indicated in the figure. If player 1 chose b or c, player 2 is informed of this (but not which one of b, c was chosen) and has to choose between g and h, whereupon the play ends with the payoff vectors indicated in the figure.

Let G be a game in extensive form. A *strategy* of player i in G is a function σ defined on the set Π_i of information sets of i, such that for every $I \in \Pi_i$ we have $\sigma(I) \in L_I$. When a position x is reached, satisfying $x \in I \in \Pi_i$, the strategy σ dictates that the next move will be the edge $e \in E_x$ for which $l_x(e) = \sigma(I)$. Note, for example, that in the game given in Figure 1, player 1 has three strategies and player 2 has six strategies.

It turns out that it is sometimes beneficial to a player to choose his strategy in a randomized way. This is a key idea of non-cooperative game theory, which will become clearer in the next sections. Here we only introduce the necessary terminology. Let G be a game in extensive form. Let S_i be the set of strategies of player i in G. For the sake of clarity, the members of S_i will be called *pure* strategies.

A *mixed strategy* of player i in G is a probability distribution on the set S_i. If player i uses a mixed strategy, this means that he performs a lottery before the play of the game, the outcome of which determines which of his pure strategies he will play. The lotteries of different players are performed independently of each other.

A *behavior strategy* of player i in G is a function σ defined on the set Π_i of information sets of i, which assigns to every $I \in \Pi_i$ a probability distribution $\sigma(I)$ on the set L_I. If player i uses a behavior strategy, this means that whenever he is called upon to make a decision in the course of the play, he performs a lottery to determine his choice at that stage. The lotteries performed at different stages, as well as those performed by different players, are independent of each other.

Under the natural embeddings, the set of pure strategies of a player is contained in the set of his behavior strategies, which in turn is contained in the set of his mixed strategies. In general, the set of mixed strategies is larger than the set of behavior strategies, as it allows a player to correlate his choices at different information sets. However, in games with perfect recall, Harold Kuhn showed that there is no loss of generality in restricting players to use behavior strategies instead of mixed strategies.

An n-tuple of strategies, one for each player in G, is called a *strategy profile*. It is clear from the above that a strategy profile (pure, mixed, or behavior) determines a probability distribution on the set of plays of the game. (Note that even in the case of pure strategies, the resulting play is still in general random, due to the existence of chance moves.) This allows us to associate with every strategy profile a payoff vector which is the expectation, according to the probability distribution induced on the plays, of the payoff vectors assigned to the various plays of the game.

3.3.The Strategic Form

The observation we just made suggests a more compact representation of a game, which focuses on the mapping from strategy profiles to payoff vectors. This is called the strategic form representation of the game.

An n-person *game in strategic form* (the term "normal form" is also used in lieu of "strategic form") is specified by the following ingredients:

- A finite non-empty set of players N. Without loss of generality $N = \{1, \ldots, n\}$.
- For every $i \in N$, a non-empty set S_i of strategies of player i is given.
- For every $i \in N$, a real-valued payoff function $\pi_i : S_1 \times \cdots \times S_n \to \Re$ of player i is given.

The game is said to be *finite* if the strategy sets S_i are finite. In the case of a finite two-person game, it is customary to represent it schematically by a table in which the rows correspond to the strategies in S_1, the columns correspond to the strategies in S_2, and the box of the table in position (σ, τ) contains the two numbers $\pi_1(\sigma, \tau)$, $\pi_2(\sigma, \tau)$ in this order.

We can pass from an extensive form representation to a strategic form representation of the same game by letting S_i be the set of pure strategies of player i, and defining π_i at a strategy profile to be the i-th coordinate of the expected payoff vector associated with that strategy profile. We illustrate this in Table 1 by representing the game given in Figure 1.

	d, g	d, h	e, g	e, h	f, g	f, h
a	$\frac{3}{4}, 0$	$\frac{3}{4}, 0$	$\frac{1}{4}, \frac{1}{2}$	$\frac{1}{4}, \frac{1}{2}$	$\frac{3}{4}, \frac{1}{2}$	$\frac{3}{4}, \frac{1}{2}$
b	$\frac{1}{4}, \frac{3}{4}$	$\frac{1}{4}, \frac{1}{2}$	$\frac{1}{4}, \frac{3}{4}$	$\frac{1}{4}, \frac{1}{2}$	$\frac{1}{4}, \frac{3}{4}$	$\frac{1}{4}, \frac{1}{2}$
c	$\frac{1}{2}, 0$	$\frac{5}{4}, \frac{1}{2}$	$\frac{1}{2}, 0$	$\frac{5}{4}, \frac{1}{2}$	$\frac{1}{2}, 0$	$\frac{5}{4}, \frac{1}{2}$

Table 1: The strategic form of the game in Figure 1

Passing from the extensive form to the strategic form entails a loss of information: we cannot reconstruct the extensive form from the strategic form. Nevertheless, for many purposes the strategic form is all we need to know about the game, and is more convenient to work with.

We observe, in particular, that while the concept of a behavior strategy cannot be captured by the strategic form, the concept of a mixed strategy can be defined directly for the strategic form: it is a probability distribution on the set S_i of pure strategies of player i. (In the case of a finite two-person game, we can visualize a mixed strategy of player 1 as a convex combination of rows, and a mixed strategy of player 2 as a convex combination of columns.) Thus, given a game G in strategic form with strategy sets S_i and payoff functions π_i, we can construct from it a new game G^* in strategic form, called the *mixed extension* of G. It has the same set of players as G, player i's strategy set S_i^* is the set of probability distributions on the set S_i, and player i's payoff function π_i^* is the multilinear extension of the function π_i (that is, we extend π_i by expectation).

It is often the case that a game is given only in strategic form. The extensive form of the game is either unknown or deemed irrelevant. In this case, the concept of a strategy becomes abstract, not related to the course of action in an interactive process. Rather, we envision the play of the game as consisting of a single stage in which the players simultaneously select strategies from their respective strategy sets. Once they have done so, their payoffs are determined by the payoff functions applied to the profile of selected strategies.

4.Two-Person Zero-Sum Games

4.1.The Concept of Value

A two-person *zero-sum game* is a game with two players in which the sum of the players' payoffs is identically zero. In other words, one player's gain is always the other player's loss; the interests of the two players are exactly opposed.

We may consider a zero-sum game in either of the two forms (extensive or strategic) presented in Section 3. In strategic form, it is customary to write a zero-sum game as $G = (S, T, \pi)$, with the understanding that S and T are player 1 and 2's strategy sets, respectively, and π, $-\pi$ are their respective payoff functions. Likewise, in the tabular representation (of a finite two-person zero-sum game) it suffices to write just one number, the payoff to player 1, in every box of the table. That is why these games are also called "matrix games". In such a game, player 1 chooses a row and tries to make the payoff as high as possible, while player 2 chooses a column and tries to make the payoff as low as possible (because his actual payoff is the negative of the entry in the table). Accordingly, player 1 is also referred to as the row player or the maximizer; player 2 is the column player or the minimizer.

8	3
1	5

Table 2: A two-person zero-sum game

How should players engaged in a zero-sum game choose their strategies? Consider for example the game given in Table 2. Player 1 may argue as follows: "If I choose the first row, the worst that can happen is that I get 3. If I choose the second row, the worst that can happen is that I get 1. Hence I should choose the first row and thereby guarantee that I get at least 3." Similarly, player 2 may argue as follows: "If I choose the first column, the worst that can happen is that I pay 8. If I choose the second column, the worst that can happen is that I pay 5. Hence I should choose the second column and thereby guarantee that I do not have to pay more than 5."

For a general zero-sum game (S, T, π), this logic suggests that a prudent player 1 should choose a strategy σ in S that attains the maximum in the expression

$$v_1 = \max_{\sigma \in S} \min_{\tau \in T} \pi(\sigma, \tau). \qquad (1)$$

Such a strategy of player 1 is called a *maximin* strategy. (We assume that the minima and the maximum are attained. This is clearly true if the game is finite.) Similarly, a prudent player 2 should choose a strategy τ in T that attains the minimum in the expression

$$v_2 = \min_{\tau \in T} \max_{\sigma \in S} \pi(\sigma, \tau). \quad (2)$$

Such a strategy of player 2 is called a *minimax* strategy.

The quantities v_1 and v_2 are the players' *optimal confidence levels*. Namely, v_1 is the highest number α such that player 1 is able to guarantee that the payoff will be at least α regardless of player 2's choice of strategy, and similarly v_2 is the lowest number β such that player 2 is able to guarantee that the payoff (to player 1) will be at most β regardless of player 1's choice of strategy.

In the above example (the game given in Table 2) we had $v_1 = 3$ and $v_2 = 5$. In general, it is easy to see that we always have $v_1 \leq v_2$. The case when $v_1 = v_2$ is of particular interest, and we proceed to analyze this case.

Let G be a zero-sum game, and let v be a real number. We say that v is a *value* of G, if player 1 has a strategy that guarantees that the payoff will be at least v, and player 2 has a strategy that guarantees that the payoff (to player 1) will be at most v. The following facts are straightforward to check:

- If G has a value then it is unique.
- G has a value if and only if $v_1 = v_2$; in this case $v = v_1 = v_2$.
- Assume G is finite, and consider its tabular representation. Then G has a value if and only if the table has a *saddle point*, that is, an entry which is the smallest in its row and the largest in its column (possibly with ties); in this case, that entry is the value of G. We illustrate this in Table 3.

1	3	−3	1	−2
4	2	3	**2**	5
2	0	1	−1	−1

Table 3: A two-person zero-sum game with a value

If G is a zero-sum game that has a value v, then the argument for playing the maximin/minimax strategies is very persuasive. Consider the situation from player 1's point of view. By using a maximin strategy, he guarantees a payoff of at least v, without making any assumptions about his opponent's behavior. If he does attempt to outguess his opponent's choice of strategy, a reasonable assumption is that the opponent, arguing along similar lines, will use a minimax strategy. Basing his choice on this assumption, player 1 sees that he cannot achieve anything higher than v, and choosing a maximin strategy gets him v. Thus, playing a maximin strategy is indicated both from a conservative, insurance seeking, perspective, and as a best response to the opponent's anticipated choice. A similar analysis supports the play of a minimax strategy by player 2.

It follows that when a game G with value v is played by two rational players, the expected outcome is v. This explains the use of the word "value" in this context. The number v is the answer to the question: How much is the privilege to play the game G worth to player 1?

The situation is more complicated when the game in question does not have a value, that is, when $v_1 < v_2$. Let us return to the example given in Table 2. We saw that the first row is the maximin strategy and the second column is the minimax strategy. Now, if player 1 anticipates the choice of the second column by his opponent, he concludes that it is better for him to choose the second row, since $5 > 3$. However, if player 2 anticipates that, he concludes that it is better for him to choose the first column, since $1 < 5$. This anticipatory logic leads to a vicious cycle, and it is not clear what the players should do.

It is interesting to note that Theorem 1 about chess-like games can be rephrased using the value terminology. We may view a chess-like game as a zero-sum game by assigning a payoff of 1 for a win, -1 for a loss, and 0 for a draw. Then it is easy to see that statements (a), (b), and (c) in the theorem correspond to the value being 1, -1, and 0, respectively. The theorem's assertion that one of (a), (b), and (c) must be true is equivalent to saying that the game has a value. By imitating the proof of Theorem 1, it can be shown that every finite two-person zero-sum game in extensive form with perfect information has a value.

This is an instance where the different roles of the two representations of a game—extensive and strategic—are reflected. The concept of a value was defined directly in terms of the strategic form. Yet, the perfect information feature, which can only be captured by the extensive form, is essential in proving that a game has a value.

4.2.The Minimax Theorem

Let us return once again to the game given in Table 2. We found that the optimal confidence levels of players 1 and 2 are 3 and 5, respectively. This was the case when we considered pure strategies only. Can the players do better by using mixed strategies?

Consider first player 1. Suppose that he flips a fair coin to determine if he will play the first or the second row. If player 2 chooses the first column, player 1 gets an expected payoff of $\frac{1}{2} \cdot 8 + \frac{1}{2} \cdot 1 = \frac{9}{2}$. If player 2 chooses the second column, player 1 gets an expected payoff of $\frac{1}{2} \cdot 3 + \frac{1}{2} \cdot 5 = 4$. In any case, player 1's expected payoff from his coin flipping strategy is at least 4. True, he will not get at least 4 every time the game is played, but on average he will get at least 4 regardless of what his opponent does (or whether his opponent is also randomizing). Even if the opponent finds out about player 1's coin flipping strategy, he cannot react to that in a way that drives the expected payoff below 4. So the coin flipping strategy achieves for player 1 a confidence level of 4, which is higher than his optimal confidence level of 3 when using a pure strategy. Actually, player 1 can do even better by a finer tuning of the probabilities. If he plays the first row with probability $\frac{4}{9}$ and the second with probability $\frac{5}{9}$, he guarantees an expected payoff of $\frac{37}{9}$.

Next, consider player 2. If he plays the first column with probability $\frac{2}{9}$ and the second with probability $\frac{7}{9}$, he guarantees an expected payoff of $\frac{37}{9}$, regardless of player 1's behavior. This is better than his optimal confidence level of 5 when using a pure strategy.

Let us observe what happened in this example. By introducing mixed strategies, both players improved on their optimal confidence levels in pure strategies. Moreover, we saw that in the mixed extension of the game, each player can guarantee a payoff of $\frac{37}{9}$ by using a suitable mixed strategy. So, whereas the original game did not have a value, the mixed extension does: its value is $\frac{37}{9}$. The corresponding mixed maximin strategy for player 1 is $(\frac{4}{9}, \frac{5}{9})$, and the corresponding mixed minimax strategy for player 2 is $(\frac{2}{9}, \frac{7}{9})$. The vicious cycle in the players reasoning that was pointed out in the discussion above disappears when mixed strategies are allowed. As in any game that possesses a value, the argument for using the mixed maximin/minimax strategies in the mixed extension is a compelling one. We can safely predict that if the game is played by rational players who are allowed to randomize, the expected payoff will be $\frac{37}{9}$ (or, put otherwise, if the game is played many times, the average payoff will approach $\frac{37}{9}$).

It was probably Émile Borel who first realized the effect of introducing mixed strategies. But John von Neumann succeeded to prove that this effect, as described in the analysis of the example above, holds in full generality. This is his celebrated minimax theorem.

Theorem 2. Let $G = (S, T, \pi)$ be a finite two-person zero-sum game in strategic form. Let $G^* = (S^*, T^*, \pi^*)$ be the mixed extension of G. Then G^* has a value, or equivalently,

$$\max_{\sigma \in S^*} \min_{\tau \in T^*} \pi^*(\sigma, \tau) = \min_{\tau \in T^*} \max_{\sigma \in S^*} \pi^*(\sigma, \tau). \qquad (3)$$

The minimax theorem is closely related to the duality theorem of linear programming. The two theorems are equivalent, in the sense that each one of them can be reduced to the other by an easy transformation. A direct proof of either theorem is based on the separation properties of convex polyhedra.

The connection to linear programming can be exploited for computational purposes. In order to compute the value of a given finite game in mixed strategies and to find mixed maximin/minimax strategies, one may use any algorithm for solving linear programs. In the special case when one of the players has only two pure strategies, an easy graphical method of solution is available.

Quite separately from its game theoretic meaning, the minimax theorem may be construed as asserting the equality of the max min and the min max for a bilinear form defined on the product of two simplices. Attempts to identify weaker conditions on the two spaces and/or on the function that suffice for the equality to hold have led to many generalizations and variants, with applications both in game theory and outside it.

5.Non-Zero-Sum Games

5.1.A Few Instructive Examples

In zero-sum games the interests of the two players are diametrically opposed. This is no longer the case for general non-cooperative games, and many interest structures are possible. We illustrate this with a few classical examples.

$1,1$	$0,0$
$0,0$	$1,1$

Table 4: A pure coordination game

Pure coordination games are the exact opposite of zero-sum games, in that the interests of the players coincide. A simple example is given in Table 4. A good way to think about this example is to imagine that the two players are drivers heading towards each other on a two-way road. Each of them has two available strategies: driving on the right or driving on the left. If both choose the right side or both choose the left side, all goes well and they get a payoff of 1 each. But if one of them chooses the right side and the other the left side, a collision results and the payoffs are 0 to both.

It is quite clear that the worst case analysis on which we based the reasoning of players in zero-sum games does not make sense here. Why should a player assume that the worst thing to him will happen, when the other player's interests are identical to his own? The only issue in this pure coordination game is to devise a mechanism for coordinating the two players' choices, so that they will end up getting the $(1,1)$ payoffs. In real life, this coordination is achieved through the common knowledge of the law that stipulates driving on the right side (in most countries) or on the left side (in some countries). Note that this law does not require enforcement: once it becomes common knowledge, it is in the players' best interest to obey it. This observation is important, because the unavailability of enforcement mechanisms is a basic premise of non-cooperative game theory.

$2,1$	$0,0$
$0,0$	$1,2$

Table 5: The battle of the sexes

Our next example is known as the "battle of the sexes". It is given in Table 5. The story that goes with this game is about a man and a woman who face a decision on where to spend the evening: the boxing arena or the concert hall. Quite stereotypically, the man prefers boxing while the woman prefers the concert, but both prefer spending the evening together to being apart. Each of them assigns a payoff of 2 to going together to the preferred event, 1 to going together to the other event, and 0 to being apart.

This game displays a combination of common interests (going to the same place) and opposing interests (concerning the choice of place). As in the previous example, it seems desirable to devise a procedure that will make sure that the two players achieve the $(2,1)$ or $(1,2)$ payoffs.

Each of the two possible agreements—going to the boxing arena or going to the concert hall—is self-enforcing in the sense that once it is reached, it is in the players' best interest to keep it. But in this example, in addition to the coordination issue, there is a bargaining issue: which of the two agreements will be reached?

Equity considerations seem to suggest a compromise that would give each player a payoff of $\frac{3}{2}$. This may be achieved in a number of ways: choosing an agreement by a coin toss, alternating between agreements each time the game is played, or making the more satisfied player compensate the other one. However, none of these arrangements is an option in our framework, where the game is non-cooperative and is played just once. Observe, in particular, that mixed strategies cannot achieve the $(\frac{3}{2}, \frac{3}{2})$ payoff vector, because no pair of mixed strategies yields a probability of $\frac{1}{2}$ each to the $(2, 1)$ and $(1, 2)$ outcomes. The point here is that mixed strategies allow each player to perform an independent lottery in order to determine his choice of strategy, but they do not allow the players to correlate their random choices.

$2, 2$	$0, 3$
$3, 0$	$1, 1$

Table 6: The prisoner's dilemma

Our third example is the "prisoner's dilemma", perhaps the most famous example in game theory. It is given in Table 6. Here, the story is about two suspects who are accused by the police of having committed a crime together. Each of them has the option to deny the charges or to confess and incriminate the other. If both deny the charges then for lack of evidence they will not be brought to justice for this crime, but they will be tried and convicted for a minor offense, which results in a payoff of 2 to each of them. If both confess, they will be convicted for the crime, but their sentence will be lenient in view of the confession—this results in a payoff of 1 to each of them. If one of them confesses and incriminates the other (who denies) then, in exchange for giving state's evidence against the other suspect, he will not be indicted. His payoff in this case is 3, while the other gets 0 (corresponding to the maximum prison sentence).

When a player in this game weighs his options, he realizes that regardless of whether the other player denies or confesses, he is better off confessing. Indeed, $3 > 2$ and $1 > 0$. (We express this fact by saying that confessing is a *dominant* strategy, that is, for any given strategy of the opponent, this strategy yields a higher payoff than the alternative strategy.) Thus, rational considerations lead both players to confess, thereby obtaining a payoff of 1 each. This result seems paradoxical, since they could both get 2 by denying the charges. If they could sign a binding agreement committing them to deny the charges, they would be well advised to do so. However, such an agreement is not self-enforcing: the players have incentives to breach it. Therefore, in the absence of enforcement possibilities, this agreement cannot be sustained and the players will end up with a payoff of 1 each. Note that the police does not have to isolate the two suspects in order to induce confessions. Even if they are allowed to communicate and coordinate their actions, they may promise each other to stick together, but at the end of the day their interests will lead them to confess.

Attempts to understand this paradoxical conclusion and reconcile it with empirical evidence have generated a significant body of research in game theory. The most important idea is that of repeated games: if the players envision the game being played repeatedly an infinite number of times, it may become rational for them to play cooperatively (i.e., to deny the charges every time), because a player who deviates from this behavior can be punished by the other player in a tit-for-tat fashion (see *Stochastic and Repeated Games*).

Finally, it should be emphasized that the prisoner's dilemma plays an important role in game theory not because of the peculiar story that gave it its name, but because it reflects a theme which is common to many real-life conflict situations. Examples include price wars, arms races, pollution effects, and more. In all these examples, everybody would be happier if everybody behaved cooperatively, but the individual incentives lead the players away from that behavior.

5.2.The Concept of Nash Equilibrium

Let $G = (S_1, S_2, \ldots, S_n, \pi_1, \pi_2, \ldots, \pi_n)$ be a game in strategic form. Here n is the number of players, S_i is player i's set of strategies, and π_i is player i's payoff function. A strategy profile $(\overline{\sigma_1}, \overline{\sigma_2}, \ldots, \overline{\sigma_n})$ is a *Nash equilibrium point* (or simply an *equilibrium*) of G, if for every player i and every strategy $\sigma_i \in S_i$, we have

$$\pi_i(\overline{\sigma_1}, \ldots, \overline{\sigma_{i-1}}, \sigma_i, \overline{\sigma_{i+1}}, \ldots, \overline{\sigma_n}) \leq \pi_i(\overline{\sigma_1}, \ldots, \overline{\sigma_{i-1}}, \overline{\sigma_i}, \overline{\sigma_{i+1}}, \ldots, \overline{\sigma_n}). \quad (4)$$

In words, an equilibrium has the property that no player can gain by a unilateral deviation from it. Put otherwise, if a player assumes that the other players behave as prescribed in the equilibrium, it is in his best interest to do so himself.

The notion of equilibrium, introduced by John Nash, is a fundamental concept in non-cooperative game theory. It embodies the idea of a self-enforcing agreement between players. In a non-cooperative framework, where players act according to their individual interests and binding agreements are not available, only Nash equilibria can be sustained.

Let us identify the equilibrium points in the three examples of the previous subsection.

- In the pure coordination game of Table 4, there are two equilibrium points. In one of them, both players drive on the right; in the other, both drive on the left. The mixed extension of this game has one additional equilibrium, in which both players toss fair coins to choose a side. The payoff vector associated with this (rather silly) equilibrium is $(\frac{1}{2}, \frac{1}{2})$.
- In the battle of the sexes (Table 5) there are two equilibrium points. In one of them, both players go to the boxing arena; in the other, both go to the concert hall. The mixed extension of this game has one additional equilibrium, in which each player goes to the preferred event with probability $\frac{2}{3}$ and to the other event with probability $\frac{1}{3}$. The payoff vector associated with this equilibrium is $(\frac{2}{3}, \frac{2}{3})$.
- In the prisoner's dilemma (Table 6) there is a unique equilibrium point in which both players confess. The same is true for the mixed extension.

The concept of Nash equilibrium is a static one. It tells us if a given strategy profile represents a viable agreement, but it does not address the question of dynamics: how will that agreement be reached, and what will determine which one of several equilibria of the same game will be reached? These issues are addressed by the theory of learning in games (see *Evolution and Learning in Games*).

When applied to the special class of two-person zero-sum games, the concept of Nash equilibrium agrees with the maximin/minimax analysis. More precisely, if $G = (S_1, S_2, \pi_1, \pi_2)$ is a zero-sum game then a strategy profile $(\overline{\sigma_1}, \overline{\sigma_2})$ is an equilibrium of G if and only if $\overline{\sigma_1}$ is a maximin strategy of player 1, $\overline{\sigma_2}$ is a minimax strategy of player 2, and $\pi_1(\overline{\sigma_1}, \overline{\sigma_2})$ is the value of G. Note that, of the two properties of maximin/minimax strategies in zero-sum games with a value—that of achieving the optimal confidence levels and that of being best responses to each other—the Nash equilibrium concept retains only the latter for non-zero-sum games.

The question of existence of a Nash equilibrium is related to the question of existence of a value in zero-sum games. An example of a zero-sum game without a value (e.g., that of Table 2) is also an example of a game with no Nash equilibrium. The positive results on the existence of a value also have analogues concerning the existence of a Nash equilibrium in non-zero-sum games. We present them in the following subsections.

5.3.Existence of a Pure Strategy Equilibrium in Games with Perfect Information

The concept of equilibrium was defined for games in strategic form. It applies of course also to games in extensive form, and in this case the existence of an equilibrium may be inferred from properties of the representation. The following theorem, due to Harold Kuhn, is the counterpart of Theorem 1 for non-zero-sum games.

Theorem 3. Let G be a finite n-person game in extensive form with perfect information. Then G has a Nash equilibrium in pure strategies.

Similar to Theorem 1, this theorem is proved by backward induction. As a by-product of this method of proof, we obtain that every game satisfying the assumptions of Theorem 3 has a pure strategy equilibrium with a stronger property, called subgame perfect equilibrium. We proceed to explain this stronger property, which was introduced by Reinhard Selten.

Let G be an n-person game in extensive form. Let (T, r) be the rooted game-tree in the representation of G. For every node x of T, we denote by G_x the game represented by the subtree of T rooted at x. The nodes of this subtree are all nodes y of T such that the path from r to y includes x. All the other ingredients in the representation of G_x are naturally inherited from the corresponding ones for G. Every strategy profile in G induces in a natural way a strategy profile in G_x. A strategy profile in G is a *subgame perfect equilibrium* of G if for every node x of T, the induced strategy profile in G_x is a Nash equilibrium.

We illustrate the distinction between a mere Nash equilibrium and a subgame perfect equilibrium by means of an example (Figure 2). In this game, the backward induction proceeds as follows. At node x, player 2 prefers c to d, since $1 > 0$. Taking this into account, player 1, at his decision node, prefers b to a, since b yields him a payoff of 2 (given the choice of c over d) while a gives him a payoff of 1. We obtain the strategy profile where 1 chooses b and 2 chooses c, which is a subgame perfect equilibrium.

But the game has an additional Nash equilibrium, in which player 1 chooses a and player 2 chooses d. This is a Nash equilibrium, because player 1's choice of a is best for him assuming that player 2 chooses d at x, and player 2's choice of d at x is immaterial assuming that player 1 chooses a. However, this equilibrium is not subgame perfect, because in the subgame G_x

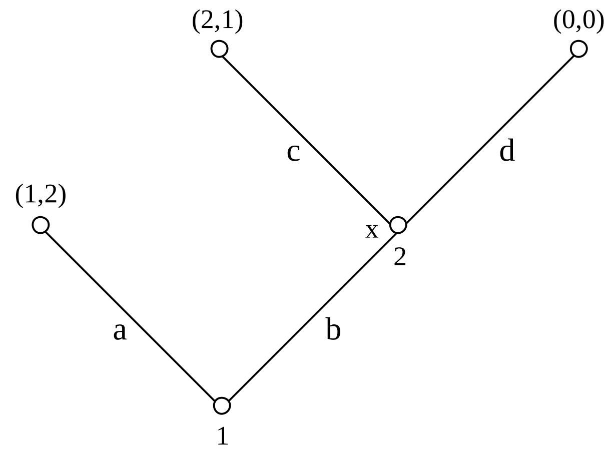

Figure 2: A game with a Nash equilibrium which is not subgame perfect

it induces a non-equilibrium. In other words, the "irrational" choice of d at node x does not violate the conditions for a Nash equilibrium, because when this strategy profile is played node x is not reached at all. But the conditions for a subgame perfect equilibrium require it to be in equilibrium even at nodes which are off the equilibrium path.

Clearly, a subgame perfect equilibrium represents more stability than a Nash equilibrium which is not subgame perfect. But in real life, one does observe sometimes non-subgame perfect equilibria. The game given in Figure 2 corresponds to the following situation. Player 1 ("the authorities") has to make a decision that affects player 2. He considers b to be the correct decision, but player 2 wants him to choose a, which is better for player 2. In order to make him do that, player 2 threatens that if player 1 chooses b, he will do d, something terrible to himself (e.g., jump off the roof) which is also undesirable to player 1. It can be argued that this threat is incredible, since it would be irrational for player 2 to carry it out, and therefore player 1 should ignore it. But if player 1 accepts the threat and chooses a to avoid the consequences of b, the players are actually at the non-subgame perfect equilibrium found above. The interpretation of such an equilibrium is closely related to the role of "incredible" threats.

5.4.Existence of a Mixed Strategy Equilibrium

In the absence of perfect information, or when the game is given in strategic form, the existence of a Nash equilibrium in pure strategies is not guaranteed. But as in the zero-sum case, the introduction of mixed strategies restores existence. The following theorem is due to John Nash.

Theorem 4. Let G be a finite n-person game in strategic form. Let G^* be the mixed extension of G. Then G^* has a Nash equilibrium.

This theorem includes the minimax theorem (Theorem 2) as a special case. But unlike the minimax theorem, whose proof employs linear methods only, the proof of Nash's theorem requires the use of topological methods. One defines a continuous mapping of the space of mixed strategy profiles to itself in such a way that any fixed point must be a Nash equilibrium, and one invokes Brouwer's fixed point theorem to prove the existence of such a point.

The difference in the proofs is also reflected in the computational aspect. Computing the value and the maximin/minimax strategies for the mixed extension of a finite two-person zero-sum game is easy. For the two-person, non-zero-sum case, it is still possible to find a mixed strategy equilibrium by the Lemke-Howson algorithm. But in the general n-person case the situation is worse, and one has to resort to the use of algorithms that approximate fixed points of continuous maps.

While the existence issue is solved by Nash's theorem, the non-uniqueness of Nash equilibrium remains an important concern. As we have seen in examples, a game may have multiple equilibria which differ significantly from each other. This undermines our ability to predict the outcome of a game based on the Nash equilibrium concept. It is natural to try to introduce additional requirements from an equilibrium, which make it more stable, more desirable, or more likely to be achieved, hopefully without destroying existence.

This direction of research has generated a variety of concepts, called "refinements" of Nash equilibrium. One example is the concept of a subgame perfect equilibrium defined for games in extensive form, that we discussed in the previous subsection. There are many other refinements, both for games in extensive form and for games in strategic form. A detailed presentation of these concepts is beyond the scope of this article.

6.Games with Incomplete Information

6.1.The Modeling of Incomplete Information

In Subsection 3.2 we showed how the notion of information sets allows us to handle imperfect information, which occurs when a player has to make a decision with only partial knowledge about the course of the play up to that moment. Here we will address the issue of incomplete information, which occurs when a player has only partial knowledge about the description of the game itself.

To illustrate this distinction, consider the example of an English auction, in which the participants make public offers to buy a certain item at increasingly higher prices. Assuming that the auction is conducted in a manner that prevents different participants from making offers simultaneously, we are looking at a game with perfect information: at every stage each participant knows the offers made up to that point. What a participant is typically uncertain about is how much the item is worth to each of the other participants. These valuations determine the utility levels that the participants associate with the outcomes of the auction. Thus, the players lack information about the payoff functions, which are part of the description of the game. This is a situation of incomplete information.

In such a situation, a player has beliefs about the values of the unknown parameters, which may be expressed in the form of a (subjective) probability distribution over some set of possible values. For example, in an auction with three participants, participant A may believe that with probability $\frac{2}{3}$ the item is worth \$200 to participant B and \$300 to participant C, and with probability $\frac{1}{3}$ it is worth \$250 to participant B and \$400 to participant C. Similarly, participant B has beliefs about the valuations of A and C, and C has beliefs about those of A and B. We assume that each participant knows his own valuation.

Clearly, a player's beliefs about the values of the unknown parameters of the game are relevant to his choice of action. But it is important to note that a player's beliefs about the other players' beliefs about the values of the unknown parameters are also relevant to his choice of action. And so are his beliefs about the other players' beliefs about his own and each other's beliefs about the values of the unknown parameters, and so on. This leads to an infinite hierarchy of beliefs, that needs to be specified in order to allow an analysis of the situation.

The modeling of these infinite hierarchies is facilitated by the introduction of two abstract concepts: the set of all possible states of the world, and for every player, the set of all possible types of that player. A state of the world ω encodes a full description specifying the parameters of the game and each player's type. A type t_i of player i encodes a full description of his beliefs, represented as a probability distribution on the set of those possible states of the world in which he is of type t_i.

The above "definitions" are admittedly circuitous, and they should be regarded as declarations of intent rather than definitions. The seminal work of John Harsanyi and subsequent work by other researchers, in particular Mertens and Zamir, has shown how these concepts can be made rigorous and at the same time tractable. The general model used to represent a game with incomplete information is that of a *Bayesian game*. We describe its ingredients one by one:

- A finite non-empty set of players N. Without loss of generality $N = \{1, \ldots, n\}$.
- A non-empty set Ω, whose elements are thought of as the possible states of the world (although formally they are just elements of an abstract set). We allow Ω to be infinite, in which case it also comes with a measurability structure.
- For every $i \in N$, there is a probability distribution P_i on Ω. It is called the prior of player i, and it represents his beliefs about the state of the world before he receives any information.
- For every $i \in N$, there is a (measurable) partition Π_i of Ω. For $\omega \in \Omega$ we denote by $\Pi_i(\omega)$ the part in Π_i which contains ω. It is understood that if the actual state of the world is ω, player i is informed that the state of the world belongs to $\Pi_i(\omega)$. Upon receiving this information, he updates his prior to a posterior $P_i|\Pi_i(\omega)$ obtained from P_i by conditioning on the event that the actual state of the world is in $\Pi_i(\omega)$. The parts of the partition Π_i can be thought of as the possible types of player i.
- For every $i \in N$, there is a non-empty set A_i of possible actions of player i. The term "action" is used here in the same sense as "strategy" was used for games with complete information, that is, an action of player i fully determines his conduct in the entire play of the game. However, we reserve the term "strategy" in the current context to mean a type-dependent choice of action (see below). The set A_i is assumed to be the same in all states of the world, and in particular independent of the type of player i. This entails no loss of generality, because we can embed the actual action sets in some universal set.
- For every $\omega \in \Omega$ and every $i \in N$, there is a function $\pi_i^\omega : A_1 \times \cdots \times A_n \to \Re$ which represents the payoff to player i when the state of the world is ω. The dependence of the payoffs on the state of the world ω reflects the idea that ω encodes, among other things, data about the parameters of the game.

This completes the description of a Bayesian game. A *strategy* of player i is a function $\sigma : \Pi_i \to A_i$. By choosing a strategy, a player is in effect announcing before he knows his type what his action will be for any of his possible types.

For a given strategy profile $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_n)$ and a given type t_i of player i (that is, part of his partition Π_i) which intersects the support of P_i, player i can compute his expected payoff, which we denote by $\mathbf{E}_{t_i}\pi_i(\sigma)$. By this we mean the expectation of $\pi_i^\omega(\sigma_1(\Pi_1(\omega)), \ldots, \sigma_n(\Pi_n(\omega)))$ when ω is chosen at random from the posterior distribution $P_i|t_i$. (The assumption that the type t_i intersects the support of the prior P_i means, in the finite case, that player i assigns a positive probability to his being of type t_i. This is necessary for the posterior $P_i|t_i$ to be defined. In the case that Ω is infinite, some further, technical conditions are needed for the above expectation to be well-defined.)

We arrive now to the concept of equilibrium for a Bayesian game G. For a strategy profile $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_n)$, a player i, a type t_i, and an action $a \in A_i$, we denote by $(\sigma|t_i \mapsto a)$ the profile obtained from σ by changing σ_i so that at t_i it chooses a. A strategy profile $\overline{\sigma} = (\overline{\sigma_1}, \overline{\sigma_2}, \ldots, \overline{\sigma_n})$ is a *Bayesian Nash equilibrium* of G if for every player i, for every type t_i of player i which intersects the support of P_i, and for every action $a \in A_i$, we have

$$\mathbf{E}_{t_i}\pi_i(\overline{\sigma}|t_i \mapsto a) \leq \mathbf{E}_{t_i}\pi_i(\overline{\sigma}). \tag{5}$$

The condition for equilibrium means that each player, when he is informed of his type and he evaluates his payoff by computing the expectation with respect to his posterior, finds it optimal to stick to the action indicated for his type in his equilibrium strategy, if he assumes that the other players follow their own equilibrium strategies.

6.2.Consistency and the Extensive Form Representation

In describing the Bayesian game model, we assumed that every player has a prior probability distribution P_i on the set Ω of possible states of the world. We did not discuss how these priors come into being. It may be argued that differences in beliefs between players can only be explained by differences in their information. Since a prior is supposed to represent beliefs as they are before receiving any information, such an argument may lead to the conclusion that in a proper sense the players must all have the same prior.

We refer to the case when all players have a common prior, that is, there exists a probability distribution P on Ω such that $P_i = P$ for all $i \in N$, as the *consistent* case. There are good conceptual arguments in favor of consistency, but there are also dissenting views. We do not go into the debate here. In any case, consistency has the advantage of allowing a simpler representation of the Bayesian game as a game in extensive form.

Let G be a consistent Bayesian game, with player set N, a set Ω of states of the world, a common prior P on Ω, partitions Π_i of Ω for $i \in N$, action sets A_i for $i \in N$, and payoff functions π_i^ω for $\omega \in \Omega$ and $i \in N$. We construct from G the following game in extensive form. First, a chance move selects a state of the world ω from Ω according to the probability distribution P. Each player $i \in N$ is informed of his type, that is, of the part in Π_i to which the selected ω belongs. Then, the players simultaneously choose actions from their respective action sets A_i, $i \in N$. (The simultaneity is captured in the extensive form by grouping together in an information set of player i all of his decision nodes which follow a selection of any ω within one of his types, regardless of the actions taken "earlier" by other players.) The payoff to player i corresponding to a play in which ω was selected by chance and the players chose the actions $a_1, a_2, \ldots, a_n$ is $\pi_i^\omega(a_1, a_2, \ldots, a_n)$.

In passing to the extensive form representation, we have eliminated the incompleteness of information from our modeling. In effect, the lack of information about the game was converted into a lack of information about the outcome of the chance move, yielding a game with complete but imperfect information. We may analyze this game as we would any such game, and in particular assume that the description of the game itself in extensive form is common knowledge.

It can be checked that the transition to the extensive form representation preserves the equilibrium concept: the Bayesian Nash equilibria of the consistent Bayesian game are precisely the Nash equilibria of the extensive form game that we constructed from it.

Bibliography

Aumann R.J. and Hart S.(Eds.)(1992, 1994, 2002). *Handbook of Game Theory with Economic Applications*, Volumes I, II, and III, Amsterdam: North-Holland. [This handbook contains surveys of various areas in game theory. Chapter 2, Games in Extensive and Strategic Forms, by S. Hart (pp. 19-40) introduces basic concepts of non-cooperative game theory. Chapter 3, Games with Perfect Information, by J. Mycielski (pp. 41-70) surveys work on infinite games with perfect information. Chapter 20, Zero-Sum Two-Person Games, by T.E.S. Raghavan (pp. 735-759) presents zero-sum games, the minimax theorem and some of its extensions. Chapter 24, Communication, Correlated Equilibria and Incentive Compatibility, by R.B. Myerson (pp. 827-847) includes, among other things, a discussion of correlated equilibrium, an important concept that is not covered in this article.]

Berlekamp E.R., Conway J.H. and Guy R.K.(1982). *Winning Ways for Your Mathematical Plays*, London: Academic Press. [This is the main reference to the area of combinatorial games.]

Fudenberg D. and Tirole J.(1991). *Game Theory*, Cambridge, MA: MIT Press. [This textbook presents non-cooperative game theory with special attention to its economic applications.]

Luce R.D. and Raiffa H.(1957). *Games and Decisions*, New York: John Wiley & Sons. [This classical book presents the main ideas and concepts of game theory with a minimum of technicalities.]

Osborne M.J. and Rubinstein A.(1994). *A Course in Game Theory*, Cambridge, MA: MIT Press. [This textbook presents game theory and in particular treats the concept of rationality and its bounded versions.]

Owen G.(1995). *Game Theory*, 3rd Edition, San Diego: Academic Press. [This textbook is oriented mainly to the mathematics student. It includes chapters on linear programming and its connections to zero-sum games.]

van Damme E.(1991). *Stability and Perfection of Nash Equilibria*, 2nd Edition, Berlin: Springer-Verlag. [This is a comprehensive reference on refinements of Nash equilibrium.]

von Neumann J. and Morgenstern O.(1944). *Theory of Games and Economic Behavior*, Princeton: Princeton University Press. [This classical book laid the foundations of game theory. The second and third editions include an appendix on utility theory.]