

Economics and the theory of games

FERNANDO VEGA-REDONDO

Universidad de Alicante and Universitat
Pompeu Fabra, Spain



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa
<http://www.cambridge.org>

© Fernando Vega-Redondo 2003

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2003

Printed in the United States of America

Typeface Times 11/13 pt. *System* L^AT_EX 2_ε [TB]

A catalog record for this book is available from the British Library.

Library of Congress Cataloging in Publication Data available

ISBN 0 521 77251 6 hardback
ISBN 0 521 77590 6 paperback

Contents

<i>Preface</i>	<i>page xi</i>
1 Theoretical framework	1
1.1 Introduction and examples	1
1.2 Representation of a game in extensive form	4
1.3 Representation of a game in strategic form	12
1.4 Mixed extension of a game	16
Supplementary material	18
1.5 Mixed and behavioral strategies	18
1.6 Representation of a game in coalitional form	23
Summary	26
Exercises	26
2 Strategic-form analysis: theory	30
2.1 Dominance and iterative dominance	30
2.2 Nash equilibrium	35
2.3 Zero-sum bilateral games	45
Supplementary material	50
2.4 Nash equilibrium: formal existence results	50
2.5 Strong and coalition-proof equilibria	53
2.6 Correlated equilibrium	56
2.7 Rationalizability	61
Summary	68
Exercises	69
3 Strategic-form analysis: applications	72
3.1 Oligopoly (I): static models	72
3.2 Mechanism design (I): efficient allocation of public goods	83
3.3 Mechanism design (II): Nash implementation	90
3.4 Markets (I): macroeconomic coordination failures	99
Summary	104
Exercises	105
4 Refinements of Nash equilibrium: theory	110
4.1 Introduction	110
4.2 Refinements excluding “incredible threats”: examples	110
4.3 Subgame-perfect equilibrium	115
4.4 Weak perfect Bayesian equilibrium	117
Supplementary material	120

4.5	Refinements excluding “untenable beliefs”: examples	120
4.6	Sequential equilibrium	128
4.7	Perfect and proper equilibria	131
4.8	Strategic-form refinements	135
	Summary	143
	Exercises	144
5	Refinements of Nash equilibrium: applications	151
5.1	Oligopoly (II): sequential moves	151
5.2	Markets (II): decentralized price formation	159
5.3	Oligopoly (III): differentiated products	171
5.4	Mechanism design (III): efficient allocation of an indivisible object	176
	Summary	182
	Exercises	184
6	Incomplete information: theory	188
6.1	Introduction and examples	188
6.2	Bayesian games	191
6.3	Bayes-Nash equilibrium	196
6.4	Signaling games	204
	Supplementary material	217
6.5	Mixed strategies, revisited: a purification approach	217
6.6	Forward induction	221
	Summary	225
	Exercises	226
7	Incomplete information: applications	231
7.1	Markets (III): signaling in the labor market	231
7.2	Markets (IV): insurance markets and adverse selection	244
7.3	Mechanism design (IV): one-sided auctions	254
7.4	Mechanism design (V): buyer–seller trade	267
	Summary	275
	Exercises	276
8	Repeated interaction: theory	281
8.1	Introduction and examples	281
8.2	Repeated games: basic theoretical framework	283
8.3	Folk theorems: Nash equilibrium	286
8.4	Reputation and “irrationality”: informal discussion	294
	Supplementary material	300
8.5	Folk theorems: subgame-perfect equilibrium	300
8.6	Reputation and “irrationality”: formal analysis	311
	Summary	319
	Exercises	321
9	Repeated interaction: applications	324
9.1	Oligopoly (IV): intertemporal collusion in a Cournot scenario	324
9.2	Oligopoly (V): intertemporal collusion in a Bertrand scenario	334
9.3	Markets (V): efficiency wages and unemployment	341
	Summary	351
	Exercises	352

10	Evolution and rationality	355
10.1	Introduction	355
10.2	Static analysis	356
10.3	Basic dynamic analysis	363
10.4	Evolution in social environments	372
10.5	Evolution of cooperation: an example	387
	Summary	393
	Exercises	394
11	Learning to play	398
11.1	Introduction	398
11.2	Reinforcement learning	399
11.3	Static perceptions and myopic behavior	412
11.4	Memory, expectations, and foresight	420
	Summary	441
	Exercises	442
12	Social learning and equilibrium selection	446
12.1	Introduction	446
12.2	Evolutionary games: theoretical framework	447
12.3	Evolutionary games: alternative scenarios	449
12.4	Stochastic stability and equilibrium selection	453
12.5	Experimental evidence	470
	Supplementary material	474
12.6	Perturbed Markov processes: basic concepts and techniques	474
12.7	Reinforcement learning with flexible aspirations	482
	Summary	495
	Exercises	496
	<i>Bibliography</i>	501
	<i>Index</i>	507

Theoretical framework

1.1 Introduction and examples

In ordinary language, we speak of a “game” as a (generally amusing) process of interaction that involves a given population of individuals, is subject to some fixed rules, and has a prespecified collection of payoffs associated to every possible outcome. Here, the concept of a game mostly embodies the same idea. However, in contrast to the common use of this term, the kind of interaction to be studied may be far from amusing, as illustrated by the following example.

Consider the game usually known as the prisoner’s dilemma (PD). It involves two individuals, labeled 1 and 2, who have been arrested on the suspicion of having committed jointly a certain crime. They are placed in separate cells and each of them is given the option by the prosecutor of providing enough evidence to incriminate the other. If only *one* of them chooses this option (i.e., “defects” on his partner), he is rewarded with freedom while the other individual is condemned to a stiff sentence of twelve years in prison. On the other hand, if both defect on (i.e., incriminate) each other, the available evidence leads to a rather long sentence for both of, say, ten years in prison. Finally, let us assume that if neither of them collaborates with the prosecutor (i.e., they both “cooperate” with each other), there is just basis for a relatively light sentence of one year for each.

The payoff table corresponding to this situation (where payoffs are identified with the *negative* of prison years) is shown in Table 1.1.

Table 1.1: *Prisoner’s dilemma*

1 \ 2	<i>D</i>	<i>C</i>
<i>D</i>	−10, −10	0, −12
<i>C</i>	−12, 0	−1, −1

What would be your prediction on the most likely outcome of this situation? It seems clear that the prediction must be (D, D) because *D* is a *dominant strategy*, i.e., it is better than the alternative *C*, no matter what the other individual might choose to do; and this is so despite the fact that (C, C) would indisputably be a better “agreement” for both. However, unless the agents are somehow able to enforce such an agreement (e.g., through a credible threat of future revenge), they will not be able to achieve that preferred outcome. If both individuals are rational (in the sense

2 Theoretical framework

of aiming to maximize their *individual* payoffs), choosing D is the only course of action that makes sense under the circumstances described.

It is important to emphasize that the former line of argument continues to apply even if the individuals are *not* isolated in separate cells and may instead communicate with each other. As long as their decisions have to be taken *independently* (e.g., in the prosecutor's office, one by one), the same reasoning applies. No matter what they might have agreed beforehand, when the time comes to implement a decision, the fact that D is a dominant choice should lead both of them to adopt it.

The game just outlined is paradigmatic of many situations of interest. For example, the same qualitative dilemma arises when two firms are sharing a certain market and each one must decide whether to undertake an aggressive or conciliatory price policy (see Chapter 3). Now, we turn to another example with a very different flavor: the so-called battle of the sexes. It involves a certain young couple who have just decided to go out on a date but still have to choose where to meet and what to do on that occasion. They already anticipate the possibilities: they may either attend a basketball game or go shopping. If they decide on the first option, they should meet by the stadium at the time when the game starts. If they decide on the second possibility, they should meet at the entrance of a particular shopping mall at that same time.

Let us assume they have no phone (or e-mail), so a decision must be made at this time. The preferences displayed by each one of them over the different alternatives are as follows. The girl prefers attending the basketball game rather than going shopping, whereas the boy prefers the opposite. In any case, they always prefer doing something together rather than canceling the date. To fix ideas, suppose payoffs are quantified as in Table 1.2.

Table 1.2: *Battle of the sexes*

Girl \ Boy	B	S
	B	S
B	3, 2	1, 1
S	0, 0	2, 3

where B and S are mnemonic for “basketball” and “shopping,” respectively, and the pairs of numbers specified quantify the utilities obtained by each individual (first the girl's, second the boy's) for each choice combination. In principle, the couple could “agree” on implementing any pair of choices on the day in question. However, only (B, B) and (S, S) represent robust (or stable) agreements in the sense that if they settle on any of them and each believes that the other side is going to abide by it, both have incentives to follow suit. Each of these agreements will be labeled a *Nash equilibrium* and either of them may be viewed as a sensible prediction for the game. The problem, of course, is that there is an unavoidable multiplicity in the task of singling out *ex ante* which one of the two possible equilibria could (or should) be played. In contrast with the previous PD game, there is no natural basis to favor any one of those outcomes as more likely or robust than the alternative one.




Figure 1.1: Battle of the sexes, sequential version.

Let us now explore a slight variation of the previous story that is not subject to the aforementioned multiplicity problem. On the day set for the date, rather than both individuals being out of reach, it turns out that the boy (only he) is at his home, where he can be contacted by phone. Suppose that the girl knows this and that, initially (i.e., when the plans were drawn), the boy managed to impose the “agreement” that they both would go shopping. The girl, angry at this state of affairs, may still resort to the following course of action: she can arrive at the stadium on the specified day and, shortly before the boy is ready to leave for the shopping mall, use the phone to let him know unambiguously where she is. Assume that it is no longer possible for the girl to reach the shopping mall on time. In this case, she has placed the boy in a difficult position. For, taking as given the fact that the girl is (and will continue to be) at the stadium waiting for him, the boy has no other reasonable option (if he is rational) than to “give in,” i.e., go to the stadium and meet the girl there. What has changed in this second scenario that, in contrast to the former one, has led to a single prediction? Simply, the time structure has been modified, turning from one where the decisions were independent and “simultaneous” to one where the decisions are sequential: first the girl, then the boy.

A useful way of representing such a sequential decision process diagrammatically is through what could be called a “multiagent decision tree,” as illustrated in Figure 1.1. In this tree, play unfolds from left to right, every intermediate (i.e., nonfinal) node standing for a decision point by one of the agents (the boy or the girl) and a particular history of previous decisions, e.g., what was the girl’s choice at the point when it is the boy’s turn to choose. On the other hand, every *final* node embodies a complete description of play (i.e., corresponds to one of the four possible outcomes of the game), and therefore has some payoff vector associated to it.

In the present sequential version of the game, it should be clear that the only intuitive outcome is (B, B) . It is true that, at the time when the plans for the date are discussed, the boy may threaten to go shopping (i.e., choose S) even if the girl phones him from the stadium on the specified day (i.e., even if she chooses B). However, as explained above, this is *not* a credible threat. Or, in the terminology

to be introduced in Chapter 4, such a threat does not belong to a (subgame) “perfect” equilibrium – only (B, B) defines a perfect equilibrium in the present case.

The representation of a game by means of a multiagent decision tree permits an *explicit* description of the order of movement of the different players as well as their information and possible actions at each point in the game. It is called its *extensive-form representation* and provides the most fundamental and complete way of defining any game. The next section formalizes this theoretical construct in a general and rigorous manner.

1.2 Representation of a game in extensive form

1.2.1 Formalization

The extensive form of a game requires the description of the following items.

1. *The set of players.* It will be denoted by $N = \{0, 1, 2, \dots, n\}$, where player 0 represents “Nature.” Nature performs every action that is exogenous to the game (whether it rains, some player wins a lottery, etc.). When it has no specific role to play, this fictitious player will be simply eliminated from the description of the game.
2. *The order of events.* It is given by a certain binary relation, R , defined on a set of nodes, K . More precisely, the set K is identified with the *collection of events* that can materialize along the game, whereas the relation R embodies a suitable criterion of *precedence* (not necessarily temporal, possibly only logical) applied to those events.¹ Here, the notion of event is the usual one, i.e., a description of “what is possible” at any given juncture in the game. Thus, in particular, an “elementary event”² is to be conceived simply as a sufficient description of a complete path of play, whereas the “sure event” refers to the situation that prevails at the beginning of the game (where still any path of play is attainable). As the players make their choices, the game advances along a decreasing (or nested) sequence of events, with a progressively narrower set of possibilities (i.e., paths of play) becoming attainable. Formally, this is captured through the relation R , which, for any pair of nodes $x, y \in K$, declares that $x R y$ whenever every path of play that is (possible) in y is (possible) as well in x . Thus, for example, if y stands for the event “both agents attend the basketball game” in the sequential battle of the sexes represented in Figure 1.1, the event x given by “the girl attends the basketball game” precedes y . Thus, by writing $x R y$ in this case,

¹ A binary relation R on K is defined as some subset of the Cartesian product $K \times K$. If $(x, y) \in R$, then we say that x is related to y and typically write $x R y$.

² In the language of traditional decision theory [see, e.g., Savage (1954)], an elementary event is the primitive specification of matters that would correspond to the notion of a “state,” i.e., a description of all relevant aspects of the situation at hand. For a formal elaboration of this approach, the reader is referred to the recent (and somewhat technical) book by Ritzberger (2002).

we mean that x logically precedes y in the set of occurrences that underlie the latter event – therefore, if y occurs, so does x as well.

Given the interpretation of R as embodying some notion of precedence, it is natural to postulate that this binary relation is a (strict) partial ordering on K , i.e., it displays the following properties³:

$$\text{Irreflexivity: } \forall x \in K, \neg(xRx).$$

$$\text{Transitivity: } \forall x, x', x'' \in K, [xRx' \wedge x'R x''] \Rightarrow xRx''.$$

Associated to R , it is useful to define a binary relation, P , of *immediate precedence* in the following manner:

$$xPx' \Leftrightarrow [(xRx') \wedge (\nexists x'' : xRx'' \wedge x''Rx')].$$

Correspondingly, we may define the set of immediate predecessors of any given $x \in K$ as follows:

$$P(x) \equiv \{x' \in K : x'Px\}$$

and the set of its immediate successors by

$$P^{-1}(x) = \{x' \in K : xPx'\}.$$

Having interpreted (K, R) as the set of partially ordered events that reflect the unfolding of play in the game, it is useful to postulate that every $y \in K$ *uniquely* defines the set of its preceding events – or, expressing it somewhat differently, that y uniquely induces the chain (or history)⁴ of occurrences that give rise to it. In essence, this is equivalent to saying that (K, R) must have the structure of a *tree of events*, thus displaying the following two properties:

- (a) There exists a unique root (or initial node) x_0 that has no immediate predecessor ($P(x_0) = \emptyset$) and precedes all other nodes (i.e., $\forall x \neq x_0, x_0Rx$). This initial node is to be viewed as the beginning of the game.
- (b) For each $\hat{x} \in K, \hat{x} \neq x_0$, there exists a unique (finite) path of predecessors $\{x_1, x_2, \dots, x_r\}$ joining $\hat{x}$ to the root x_0 – i.e., $x_q \in P(x_{q+1})$, for all $q = 0, 1, \dots, r - 1$, and $x_r \in P(\hat{x})$.

As intended, (a) and (b) permit identifying each node in K with a (unique) particular history of the game – possibly partial and incomplete if it is an intermediate node, or even “empty” if it is the initial x_0 . Also note that, from (a) and (b), it follows that every $x \neq x_0$ has a *unique* immediate predecessor (i.e., $P(x)$ is a singleton). Indeed, this is precisely the key feature that allows one to associate to every node the set of its preceding events (i.e., the underlying history) in a univocal fashion. A possible such

³ As customary, we use the symbol $\neg(\cdot)$ to denote the negation of the statement in question, or $\wedge, \vee$ to join two statements by “and,” “or.” An alternative way of expressing negation is by superimposing $/$ on a certain symbol, e.g., $\nexists$ stands for the negation of existence.

⁴ Its temporal connotations notwithstanding, the term “history” is typically used in game theory to describe the unfolding of a path of play even when the implied irreversibility does not involve the passage of time. An illustration of this point may be obtained from some of our upcoming examples in Subsection 1.3.2.




Figure 1.2: Tree of events with $x_0 P x P x' P z$; $x_0 P x P z'$; $x_0 P x'' P z''$.

tree of events is graphically illustrated in Figure 1.2, where the play of the game unfolds from left to right and any two nodes linked by a line segment are taken to be immediately adjacent according to the relation P .

For simplicity, let us posit here that every path of the game reaches a definite end.⁵ Denote by $Z \equiv \{x \in K : P^{-1}(x) = \emptyset\}$ the set of final nodes, i.e., those nodes with *no* successors (for example, the nodes z , z' , and z'' in Figure 1.2). As explained, the interpretation of any such node is that of a primitive event, a complete history, or simply a *game play*. It is worth emphasizing that every final node includes not only information on the “characteristics” of the final outcome of the game but also describes in full detail its underlying history. To illustrate this point, consider for example the event “wearing the two gloves” resulting from the concatenation of the intermediate events “not wearing any glove” and “wearing just one glove.” Then, the two different ways in which one may end up wearing the two gloves (either the right or the left glove first) give rise to two different final nodes, even though they both display the same relevant features.

3. *Order of moves.* The set $K \setminus Z$ of intermediate nodes is partitioned into $n + 1$ subsets $K_0, K_1, \dots, K_n$. If $x \in K_i$, this simply means that when the event reflected by x materializes, it is player i ’s turn to take an action. For convenience, it is typically assumed that, if Nature moves in the game, it does so first, thus resolving once and for all any bit of exogenous uncertainty that may affect the course of play. In terms of our previous formalization, this amounts to making $K_0 \subseteq \{x_0\}$ – of course, K_0 is empty if Nature does not have any move in the game.
4. *Available actions.* Let $x \in K_i$ be any node at which some player $i \in N$ moves. The set of actions available to player i at that node is denoted by

⁵ Some of the game-theoretic models proposed at later points in this book (cf., for example, Subsections 5.2.1 and 8.2) admit the possibility that the game never ends, a case that requires a natural extension of the present formulation. Then, every infinite history must be interpreted as a different “end node,” which again embodies a full description of the whole turn of events that underlie it.

$A(x)$. Naturally, the cardinality of $A(x)$ must be identical to that of $P^{-1}(x)$, the set of immediate successors of x . This simply reflects the fact that it is player i who decides how the game proceeds after x along one of the possible ensuing directions. Formally, what is required is that the sets $A(x)$ and $P^{-1}(x)$ be isomorphic, i.e., each immediate successor of x must have a *unique* and *different* action a in the set $A(x)$ associated to it, and *vice versa*.

5. *Information sets.* For every player i , we postulate that her corresponding set of decision nodes K_i can be partitioned into a set H_i of disjoint sets, i.e., $K_i = \bigcup_{h \in H_i} h$ with $h \cap h' = \emptyset$ for all $h, h' \in H_i$ ($h \neq h'$). Each of these sets $h \in H_i$ is called an *information set* and has the following interpretation: player i is unable to discriminate among the nodes in h when choosing an action at any one of them. Intuitively, if player i cannot distinguish between two different nodes $x, x' \in h$, it must be that player i did not observe (or has forgotten – see Section 1.4) the preceding occurrences (choices) on which x and x' differ. Obviously, this interpretation requires that $A(x) = A(x')$ – that is, there must exist the same set of available actions at both x and x' . Otherwise, the inconsistency would arise that player i could in fact distinguish between x and x' on the basis of the different set of actions available at each node (an information that of course player i should have because she is the *decision maker* at both of those nodes).
6. *Payoffs.* Associated with every possible game play (i.e., final node or complete history of the game) there is a certain payoff for each of the different players. Thus, for every one of the final nodes $z \in Z$, we assign an n -dimensional real vector $\pi(z) = (\pi_i(z))_{i=1}^n$, each $\pi_i(z)$ identified as the payoff achieved by player $i = 1, 2, \dots, n$ if the final node z is reached. These real numbers embody how players evaluate any possible outcome of play and thus reflect every consideration they might deem relevant – pecuniary or not, selfish or altruistic. Payoffs for Nature are not specified since its behavior is postulated exogenously. (Fictitiously, one could simply posit constant payoffs for Nature over all final nodes.)

Payoff magnitudes are interpreted as von Neumann–Morgenstern utilities and, therefore, we may invoke the well-known theorem of expected utility⁶ when evaluating random outcomes. That is, the payoff or utility of a certain “lottery” over possible plays (or final nodes) is identified with its expected payoff, the weights associated with each one of those plays given by their respective *ex ante* probability. This implies that payoffs have a cardinal interpretation (i.e., payoff *differences* have meaning) and embody players’ attitude to risk. Formally, it amounts to saying that the specification of the payoffs in the game is unique *only* up to monotone *affine* transformations.⁷

Finally, note that even though “payoff accounting” is formally performed at the end of the game (i.e., payoffs are associated with final nodes alone),

⁶ See, e.g., Kreps (1990) for a classical textbook treatment of this topic.

⁷ A monotone affine transformation of a utility function $U(\cdot)$ is any function $\tilde{U}(\cdot)$ over the same domain, which may be written as follows: $\tilde{U}(\cdot) = \alpha + \beta U(\cdot)$ for any real numbers α, β , with $\beta > 0$.

this does not rule out that partial payoffs may materialize at intermediate stages. In those cases, the payoff associated with any final node is to be interpreted as the overall evaluation of the whole stream of payoffs earned along the unique history that leads to it.

The above six components define a game in extensive form. Often, we shall rely on a graphical description of matters where

- the unfolding events $x \in K$ induced by players' actions are represented through a tree structure of the sort illustrated in Figure 1.2;
- intermediate nodes are labeled with the index $i \in N$ of the player who takes a decision at that point;
- the edges departing from intermediate nodes $x \in K$ are labeled with the respective actions $a \in A(x)$ leading to each of its different successors in $P^{-1}(x)$;
- the intermediate nodes $\{x \in h\}$ that belong to the same information set h are joined by a dashed line;
- the final nodes $z \in Z$ have real vectors $\pi(z)$ associated with them, expressing the payoffs attained by each player in that game play.

A simple illustration of such a graphical way of describing a game in extensive form is displayed in Figure 1.3.

1.2.2 Examples

1.2.2.1 A simple entry game. Consider two firms, 1 and 2, involved in the following game. Firm 1 is considering whether to enter the market originally occupied by a single incumbent, firm 2. In deciding what to do (enter (E) or not (N)), firm 1 must anticipate what will be the reaction of the incumbent (fight (F) or concede (C)), a decision the latter will implement only after it learns that firm 1 has entered the market. Assume that the monopoly (or collusive) profits to be derived from the




Figure 1.3: A game in extensive form.




Figure 1.4: A simple entry game, extensive form.

market are given by two million dollars, which firm 2 either can enjoy alone if it remains the sole firm or must share with firm 1 if it concedes entry. On the other hand, if firm 2 fights entry, both firms are assumed to incur a *net* loss of one million dollars because of the reciprocal predatory policies then pursued.

The extensive-form representation of the entry game considered is described in Figure 1.4. In this simple extensive form, each firm has just one information set consisting of only one node. Thus, in both of these information sets, the corresponding firm is fully informed of what has happened at preceding points in the game. With this information at hand, each firm has two possible actions to choose from (N or E for firm 1; F or C for firm 2).

1.2.2.2 A matching-pennies game. Consider the following game. Two players simultaneously choose “heads” or “tails.” If their choices coincide (i.e., both select heads, or both select tails) player 2 pays a dollar to player 1; in the opposite cases, player 1 pays this amount to player 2.

As explained above, the extensive form is to be conceived as the most basic and complete way of representing a game. However, since an extensive-form representation displays, by construction, a sequential decision structure (i.e., any decision node can belong to only a single agent), one might be tempted to think that it is inherently unsuited to model any simultaneity of choices such as the one proposed here. To resolve this puzzle, the key step is to grasp the appropriate interpretation of the notion of “simultaneity” in a strategic context. In any given game, the fact that certain actions are described as “simultaneous” does not necessarily reflect the idea that they are chosen at the same moment in *real time*. Rather, the only essential requirement in this respect is that at the time when one of the players takes her decision, she does not know any of the “simultaneous” decisions taken by the other players.

To formalize such a notion of simultaneity, we rely on the concept of information set, as formulated in Subsection 1.2.1. This allows us to model the matching-pennies game through *any* of the two extensive-form representations displayed in Figures 1.5 and 1.6 (recall the graphical conventions illustrated in Figure 1.3). In either of these alternative representations, each player has just one information set and two possible actions (heads (H) or tails (T)). However, while in the first representation it is player 1 who “fictitiously” starts the game and then player 2 follows, the second representation has the formal roles of the players reversed. Clearly, both of these




Figure 1.5: A matching-pennies game in extensive form, alternative 1.




Figure 1.6: A matching-pennies game in extensive form, alternative 2.

alternative extensive-form representations of the game should be viewed by the players as *strategically equivalent*. In both of them, no player is informed of the action played by the opponent, either because she moves first or because she is unable to distinguish between the possible “prior” moves of the other player.

1.2.2.3 Battle of the sexes. Along the lines pursued for the previous example, we may return to the battle of the sexes introduced in Section 1.1 and describe the extensive-form representation of its simultaneous version as displayed in Figure 1.7.

Again, the alternative representation of the simultaneous battle of the sexes where the formal roles of the boy and the girl are reversed is strategically equivalent to the one described in Figure 1.7. Of course, this is no longer the case if we consider instead the sequential version of the game where the girl moves first. Such a game has the extensive-form representation described in Figure 1.1. In it, the girl still has only one information set (she moves without knowing the decision her partner will make), but the boy has two information sets (he already knows the decision adopted by the girl at the time he makes his own decision). As explained in our informal discussion of Section 1.1, this sequential version of the game leads to a rather strong strategic position for the girl. It is obvious, however, that the relative strength of the strategic positions is reversed if the roles of the players (i.e., their order of move) is permuted. Thus, in contrast with the simultaneous version, such




Figure 1.7: Battle of the sexes, simultaneous version; extensive form.

a permutation does not produce strategically equivalent extensive-form games in this case.

1.2.2.4 The highest-card game. Two players use a “pack” of three distinct cards, $C \equiv \{h(\text{high}), m(\text{medium}), l(\text{low})\}$, to participate in the following game. First, player 1 picks a card, sees it, and then decides to either “bet” (B) or “pass” (P). If player 1 bets, then player 2 picks a card out of the two remaining ones, sees it, and chooses as well to either “bet” (B') or “pass” (P'). If both players bet, the player who has the highest card (no ties are possible) receives a hundred dollars from the opponent. On the other hand, if at least one of the players does *not* bet, no payments at all are made.

The extensive-form representation of this game is displayed in Figure 1.8. First, Nature moves at the root of the game (recall Subsection 1.2.1) and chooses one of the six possible card assignments for the two players in the set $D \equiv \{(c_1, c_2) \in C \times C : c_1 \neq c_2\}$. Next, there are three possible information sets for player 1, as she is informed of her own card but *not* of that of the opponent. (Again, we use the convention of joining the nodes included in the same information set by a discontinuous line.) In each of these information sets there are two nodes (those that correspond to the opponent receiving one of the two cards she herself has not received) and the same two choices available (B or P).⁸ In case player 1 decides to bet, three further information sets for player 2 follow, each of them reflecting analogous information considerations for this player. If both bet (i.e., choose B and B' , respectively), the induced final nodes have a payoff vector that assigns 100 to the player with the highest card and -100 to the opponent. On the other hand, if one of them does not bet (i.e., passes), the corresponding final node has a payoff vector $(0, 0)$ associated with it.

⁸ Note that, for notational simplicity, the same label (P or B) is attributed to pass or bet in different information sets of player 1. To be fully rigorous, however, we should have different labels in different information sets because their respective actions of passing and betting should be conceived as different in each of them. Of course, the same comment applies to the action labels of player 2 in the subsequent information sets.




Figure 1.8: The highest-card game, extensive form.

1.3 Representation of a game in strategic form

1.3.1 Formalization

Consider a game in extensive form:

$$\Gamma = \{N, \{K_i\}_{i=0}^n, R, \{H_i\}_{i=0}^n, \{A(x)\}_{x \in K \setminus Z}, \{[\pi_i(z)]_{i=1}^n\}_{z \in Z}\} \quad (1.1)$$

where each of its components has been formally defined in Subsection 1.2.1. All players involved in the game are assumed to be perfectly informed of its underlying structure, i.e., they know each of the components listed in Γ . Therefore, every one of them can precisely identify the different situations in which she might be called upon to play and consider, hypothetically, what her decision would be in every case. In modeling players who can perform *ex ante* such an exhaustive range of hypothetical considerations, we are led to the fundamental notion of *strategy*.

For each player, a strategy in Γ is a *complete* collection of *contingent* choices that prescribe what this player would do in each of the occasions in which she might have to act (i.e., make some decision). Thus, a strategy has to anticipate *every* possible situation in which the player could be asked to play and, for each of them, determine a particular choice among the alternative options available. Obviously, since it is impossible to demand from a player that she make a decision that depends on information she does *not* hold, a strategy must prescribe the *same* action for all the nodes included in any particular information set. Or, rephrasing it somewhat differently, a strategy can make players' decisions contingent only on their respective information sets, *not* on the particular decision nodes (among which they are not always able to discriminate).

A different and complementary perspective on the concept of strategy is to view it as a *sufficient* set of instructions that, if the player were to convey them to some intermediary, would allow the former to leave the game and have the latter act on her behalf. Only if this set of instructions can *never* leave the intermediary at a loss (i.e., not knowing how to proceed in some circumstances), can we say that it properly defines a strategy for the player in question.

To proceed now formally, recall that, for each player $i \in N$, H_i denotes the partition of her respective decision nodes K_i in disjoint informations sets. For any $h \in H_i$ and $x \in h$, consider the simplifying notation $A(h) \equiv A(x)$, and denote $A_i \equiv \bigcup_{h \in H_i} A(h)$. Then, as explained above, a strategy for player i is simply a function

$$s_i : H_i \rightarrow A_i, \quad (1.2)$$

with the requirement that

$$\forall h \in H_i, s_i(h) \in A(h), \quad (1.3)$$

i.e., any of the actions selected at given information set h must belong to the corresponding set of available actions $A(h)$.

Note that since any strategy s_i of player $i \in N$ embodies an *exhaustively* contingent plan of choice, every particular *strategy profile* $s \equiv (s_0, s_1, s_2, \dots, s_n)$ specifying the strategy followed by each one of the $n + 1$ players uniquely induces an associated path of play. Denote by $\zeta(s) \in Z$ the final node representing this path of play. Because all players are taken to be fully informed of the game Γ (i.e., know the different items specified in (1.1)), every player can be assumed to know as well the *mapping* $\zeta : S_0 \times S_1 \times \dots \times S_n \rightarrow Z$. Therefore, the decision problem faced by each player i can be suitably formulated in the following fashion: choose strategy $s_i \in S_i$ under some anticipation, conjecture, or guess concerning the strategies $s_j \in S_j$ to be chosen by the remaining players $j \neq i$. But then, if players may approach the strategic situation from such an *ex ante* viewpoint (i.e., by focusing on their own and others' plans of action), the same must apply to us, game theorists, who aim to model their behavior. Indeed, this is in essence the perspective adopted by the model of a game that is known as its *strategic (or normal) form representation*, which is denoted by $G(\Gamma)$.⁹ It consists of the following list of items:

$$G(\Gamma) = \{N, \{S_i\}_{i=0}^n, \{\pi_i\}_{i=1}^n\},$$

where

1. N is the set of players.
2. For each player $i = 1, 2, \dots, n$, S_i is her strategy space, i.e., the set of possible mappings of the form given by (1.2)–(1.3). Often, if we denote by $|H_i|$ the cardinality of the set H_i , it will be convenient to think of S_i

⁹ The notation $G(\Gamma)$ responds to the idea that Γ is taken to be the most fundamental representation of the game, whereas $G(\Gamma)$ is conceived as a “derived” representation. Nevertheless, we often find it convenient to formulate a game directly in strategic form, thus dispensing with the explicit detail of its extensive-form structure.

Table 1.3: *A simple entry game, strategic form*

1 \ 2	F	C
N	0, 2	0, 2
E	-1, -1	1, 1

as contained in the Cartesian product $A_i^{|H_i|}$, a set that is isomorphic to the family of functions of the form (1.2).¹⁰

3. For each player $i = 1, 2, \dots, n$,

$$\pi_i : S_0 \times S_1 \times \dots \times S_n \rightarrow \mathbb{R} \quad (1.4)$$

is her payoff function where, abusing former notation, the payoff associated to every strategy profile $s \equiv (s_0, s_1, s_2, \dots, s_n)$ is identified with the payoff $\pi_i(\zeta(s))$ earned by player i in the final node $z = \zeta(s)$ uniquely induced by those strategies.

The apparent simplicity displayed by the strategic-form representation of a game is somewhat misleading. For if the underlying game is complex (e.g., displays an involved sequential structure), a complete specification of the strategy spaces may become a quite heavy task. Then, the full richness of detail (order of movement, dispersion of information, player asymmetries, etc.), which is *explicitly* described by the representation of the game in extensive form, becomes implicitly “encoded” by a large set of quite complex strategies. To illustrate matters, we now focus on the collection of leading examples introduced in Subsection 1.2.2 and describe for each of them in turn their corresponding strategic form.

1.3.2 Examples

1.3.2.1 A simple entry game (continued). Consider the entry game whose extensive form is described in Figure 1.4. In this game, both players have only one information set. Therefore, their respective strategy sets can be simply identified with the set of possible actions for each of them. That is, $S_1 = \{N, E\}$ and $S_2 = \{F, C\}$. To complete the specification of the strategic form, one still has to define the players’ payoff functions. These may be characterized by a list of payoff pairs $[(\pi_i(s_1, s_2))_{i=1,2}]_{(s_1, s_2) \in S_1 \times S_2}$, as displayed in Table 1.3, where each row and column, respectively, is associated with one of the strategies of individuals 1 and 2.

1.3.2.2 A matching-pennies game (continued). Consider now the matching-pennies game whose two equivalent extensive-form representations are described in Figures 1.5 and 1.6. Again in this case, because each player has only one

¹⁰ Let the elements of H_i be indexed as $h_1, h_2, \dots, h_r$. Then, any $y = (y_1, y_2, \dots, y_r) \in A_i^{|H_i|}$ can be identified with the mapping $s_i(\cdot)$ such that $s_i(h_k) = y_k$. Of course, for such a mapping to qualify as a proper strategy, it has to satisfy (1.3).

Table 1.4: *A matching-pennies game, strategic form*

1 \ 2	H	T
H	1, -1	-1, 1
T	-1, 1	1, -1

Table 1.5: *Battle of the sexes, sequential version; strategic form*

Girl \ Boy	(B, B)	(B, S)	(S, B)	(S, S)
B	3, 2	3, 2	1, 1	1, 1
S	0, 0	2, 3	0, 0	2, 3

information set, her strategies can be identified with the actions available in that single information set. That is, $S_1 = \{H, T\}$, $S_2 = \{H, T\}$ where, for simplicity, we do not distinguish notationally between each player's strategies. Finally, to define the payoff functions, the induced payoff pairs are arranged in Table 1.4 with the aforementioned conventions.

1.3.2.3 Battle of the sexes (continued). The simultaneous version of the battle of the sexes (cf. Figure 1.7) is formally analogous to the previous example, its payoffs as given by Table 1.2. On the other hand, its sequential version, whose extensive-form representation is given by Figure 1.1, has the girl displaying one information set and the boy displaying two of them. Thus, for the girl, her strategy set is simply $S_g = \{B, S\}$, whereas for the boy we have $S_b = \{(B, B), (B, S), (S, B), (S, S)\}$. Here (recall Subsection 1.3.1), we view each of the boy's strategies as an element of $\{B, S\}^{|H_b|} = \{B, S\}^2$, with the information sets indexed downward (i.e., the upper one first, the lower one second). With this notational convention, the payoff functions are as indicated in Table 1.5.

An interesting point to note in this case is that the payoff table displays a number of payoff-vector equalities across pairs of different cells. This simply reflects the fact that, given any particular girl's strategy, only that part of the boy's strategy that pertains to the information set induced by the girl's decision is payoff relevant. Therefore, the two different boy's strategies that differ only in the information set *not* reached (given the girl's chosen strategy) lead to the same payoff for both players.

1.3.2.4 The highest-card game (continued). Consider the game whose extensive-form representation is described in Figure 1.8. In this game, the players' strategy spaces (including Nature in this case) are as follows:

- *Nature*: $S_0 = \{(c_1, c_2) \in C \times C : c_1 \neq c_2\}$.
- *Player 1*: $S_1 = \{s_1 : C \longrightarrow \{\text{Bet}(B), \text{Pass}(P)\}\} = \{B, P\}^{|C|}$.
- *Player 2*: $S_2 = \{s_2 : C \longrightarrow \{\text{Bet}(B'), \text{Pass}(P')\}\} = \{B', P'\}^{|C|}$.