

Lecture Notes on Game Theory and its Applications

Antonio Jiménez-Martínez¹

CIDE, Mexico City

¹CIDE, Division of Economics, Mexico City, MEXICO. Email: antonio.jimenez@cide.edu.

Contents

1. Preface	1
2. Introduction	3
2.1 Scope and (Most Common) Methodological Assumptions	3
2.2 (Brief) Historical Background	8
2.3 Non-Cooperative Games <i>versus</i> Cooperative Games	9
3. Extensive Form Games	11
3.1 Definitions and Examples	11
3.2 (Pure) Strategies	18
3.3 (Pure Strategy) Nash Equilibrium	21
3.4 “Backwards Induction” and Subgame Perfect Equilibrium	31
3.5 Behavior Strategies	37
4. Strategic Form Games	39
4.1 Strategic (or Normal) Form Representations	39
4.2 Mixed Strategies (and Their Relation to Behavior Strategies)	40
4.3 Beliefs and Best Replies	43
4.4 (Mixed Strategy) Nash Equilibrium	44
4.5 Dominated Strategies and <i>Rationalizability</i>	49
4.6 Correlated Equilibrium	53
5. Bayesian Games	61

6. Equilibrium Refinements	73
6.1 Perfect Bayes-Nash Equilibrium	73
6.2 Sequential Equilibrium	82
6.3 The “Global Games” Approach	85
7. Repeated Games	89
7.1 Finite <i>versus</i> Infinitely Repeated Games	97
7.2 “Folk” Theorems	98
8. Applications	105
8.1 Models of Communication: Sender-Receiver Games	105
8.2 Mechanism Design	130
8.3 Reputation	149
8.4 Bargaining Games	149
Bibliography	151

1. Preface

These notes were prepared for a course on game theory and its applications. They are preliminary and incomplete. The notes draw upon various published and unpublished sources. In particular, motivation of the main topics, basic definitions, notation, and examples follow closely the introductory textbooks by Martin Osborne ([Osborne \(2004\)](#)) and Joel Watson ([Watson \(2008\)](#)). The discussion of more advanced topics, and related examples, are built upon the textbook by Ariel Rubinstein and Martin Osborne ([Rubinstein and Osborne \(1994\)](#)). The exposition of some topics and applications is also based on a number of insightful notes that several authors have generously made available online, including Nageeb Ali, Dirk Bergemann, Navin Kartik, Stephen Morris, Ilya Segal, Joel Sobel, and Steve Tadelis. While these notes benefit substantially from all these materials, I am the only responsible for any errors contained in the current version. Suggestions and corrections are gratefully welcome.

2. Introduction

2.1. Scope and (Most Common) Methodological Assumptions

In many interactions, people regard their well-beings as interdependent: they consider that the actions of others affect them and, conversely, that their own actions influence the well-being of others. We label these situations as *strategic interaction* settings and, accordingly, say that individuals act *strategically* in such environments. These settings stand in sharp contrast with others where decision-makers perceive instead that the actions of a single person have no impact on the overall outcome and, therefore, neither on the well-being of other individuals. In economics, these situations where the well-being of the decision-makers is not regarded as interdependent are usually termed as *competitive* settings and decision-makers are said to act as price-takers.¹ Game theory is set of tools that allow us to analyze formally how individuals make decisions in situations of strategic interaction. Competitive environments are, therefore, not included in the scope of game theory. Competitive situations, though, are in fact present in many environments and the implications of competitive settings are of broad interest for economists and social scientists. A question of particular interest in competitive settings is that of how competitive outcomes reconcile with socially desired goals (i.e., the study of efficiency or of the social optimality of outcomes in competitive situations). The field of economics that systematically studies competitive settings is that of *competitive equilibrium theory*. In consonance with the requirement of strategic

¹This terminology comes from the most prominent case of such situations in economic relationships: a market for a good with many consumers that demand the good and many firms that offer it. In such environments, each individual is usually seen as very small compared to the size of the entire population in the market. As a consequence, it is natural to consider that the actions of a single consumer (or a single firm) do not impact the final price of the good. Consumers choose the quantities of the good they wish to demand (and firms choose the quantities they produce), but each decision-maker takes the final price as given. In other words, in consonance with their negligible sizes within the market, decision-makers believe that their individual actions alone are not able to influence the price of the product.

settings, game theory usually studies situations where a relatively small number of individuals interact, whereas competitive equilibrium theory explores environments where a large number of (relatively small) individuals interact.² We will use the term *game* to refer to a particular situation of strategic interactions that we are interested in studying, and the decision-makers engaged in such interactions will be referred to as *players*.

The main goal of game theory is to provide an analytical framework that helps us understand how people behave in a variety of strategic settings and predict, or explain, future, or hypothetical, behavior. To set up such a framework, we need to make particular assumptions about how individuals assess the best way to behave when they interact with others. First, we need to be specific about the goals of the individuals and about how they behave to accomplish such goals. The most common assumption in game theory—as well as in any field of modern economics or social sciences that follow a systematic scientific approach to study human behavior—in this respect is that of *rationality*. Rationality can be viewed as a very mild assumption that simply requires that individuals have a clearly formed idea of what makes them happy and that, given the information available to them in each situation, act so as to maximize their happiness. In other words, individuals have some well-specified preferences about the outcomes obtained by the combination of everyone’s actions and then choose the actions that lead to the outcomes they most prefer. Rationality says nothing about the form of the decision-makers’ preferences and any coherent preferences can be accommodated by this assumption.³ Rationality, in particular,

² Models of *market games* are usually seen as a bridge between the theoretical approaches taken by game theory and by competitive equilibrium theory. Market games explore the formation of prices in markets using techniques from game theory. Specifically, such models start with a relatively small number of decision-makers in the market which are assumed to behave strategically and then explore how the outcomes of the strategic interactions evolve as the number of individuals increase. Under certain conditions, the strategic outcomes converge to those predicted by the competitive equilibrium theory. In this way, market games can be seen as a game theoretical foundation of competitive equilibrium models.

³ Economists have traditionally debated on what basic premises must satisfy any coherent system of human preferences and have come up with a set of fundamental conditions, usually called axioms. Whereas the typical axioms in environments under certainty are fairly weak, the most commonly accepted axioms under uncertainty include one controversial requirement: the *independence (of irrelevant alternatives) axiom*. The independence axiom allows us to represent preferences under uncertainty in the quite convenient and tractable way of the expected utility form. The independence axiom have often been challenged by experimental evidence and its implications about the expected utility form representation are sometimes at odds with experimental and empirical findings. Motivated by the critiques raised from the experimental and empirical feedback, alternative representations of preferences under uncertainty, in which the independence axiom is relaxed or dropped, have been proposed over the last decades. However, the simple representation of preferences under the form of the expected utility, together with the lack of a unified theory that encompasses the alternative models, have lead to that accepting the independence axiom—and therefore resorting to

does not require that individuals maximize their monetary payoffs.⁴ To put it simply, rationality requires that individuals try to fulfill their happiness—whatever it is that makes them happy—or, equivalently, that they seek the course of action that suits best their preferences—regardless of the particular features of such preferences.

Secondly, to be more specific about the problems that the players face when trying to achieve the outcomes that they most prefer, we need to be precise about what they know of the situations where they have to decide. With regards to the information available about aspects relevant to their well-being, players may be uncertain about them, and, in particular, they can have perfect or *imperfect information*. Often, some players may have more precise information than others about such aspects. A player has *imperfect information* when, at some point during the play of the game, she lacks some information about the description of the game. Not having full information about some aspect of the game during the play generally implies that the player does not have all the strategically relevant information when she makes her choices. These features are present in most practical real-life situations and they are also captured by the set of tools provided by game theory. In this respect, game theory usually assumes that, *before playing the game*, players are fully aware about what they will know—and about what they will not—of the situations where they will interact with others. In other words, before they make their decisions, players are assumed to have full information about the entire description of the game they will be involved in, even of those parts of the game where, at the time of their decisions, they will have imperfect information about well-being-relevant aspects.⁵

Thirdly, to determine their best ways to behave, players need to draw inferences, or form *beliefs*, about how others make their own decisions. To deal with the inferences that each individual must

the expected utility form—constitutes in practice the most used approach to deal with preferences under uncertainty.

⁴For instance, if a decision-maker has altruistic interests and she prefers that others enjoy additional income even at the expense of sacrificing her own money, then rationality says that she will seek to donate part of her income to others. In this sense, there is nothing wrong in practice with that rational decision-makers may seek “non-material” or “spiritual” goals, may incur in monetary losses, or may even inflict themselves what others would regard as physical or emotional sufferings in the pursue of their happiness.

⁵To put it simply, players have a precise description of situations where they will face uncertainty and even where they will have imperfect information about some variable, or about the actions taken by others. In this sense, when a player does not know the value of the variable, we capture this by considering that the unknown value is the realization of a random variable. Then, the player is assumed to act as an statistician that knows the probability distribution of the random variable. Likewise, when a player does not know other player’s actions, she will also assign probabilities, usually termed as *beliefs*, to the set of possible actions.

make about how others will behave, game theory usually considers that players not only know fully the description of the game but that they also share a “common understanding” of such a description. The notion of *common knowledge* usually captures this idea of common understanding. By common knowledge of a feature we understand that each player knows the feature itself, that each player knows that each other player knows it, that each player knows that each other player knows that each other player knows it, and so on, *ad infinitum*.⁶ Game theory usually assumes that, prior to engaging in the game, players have common knowledge both about the fact that they are rational decision-makers and about the entire description of the game they will play. In dynamic interactions, nonetheless, sharing a common understanding of the description of the game does not imply that the players always commonly known the information they have, or what they are doing. When the play of the game takes place over time, players are assumed to share a common understanding of the description of the game—including these parts where they possibly will have imperfect information about some aspects—*before* they begin to play. However, during the play some players may have different information than others.

Finally, to study how players behave according to game theory, we need to use a solution, or equilibrium, notion that takes into account (1) how players think logically about how others will play in the game and (2) how they act in the light of such inferences in order to maximize their happiness. Furthermore, in order to be useful in explaining and predicting the behavior of decision-makers, this equilibrium notion must exist for a fairly broad set of games and must satisfy certain desirability properties. The *Nash equilibrium*—henceforth, NE—solution concept achieves all these goals in a brilliant way. In addition to the assumptions already mentioned, the NE concept requires that the players consider a (sometimes fictional) social institution (or mediator)

⁶That one knows the feature itself is viewed as a zero-order iterated belief about the feature. That one knows that other individual knows the feature is viewed as a first-order iterated belief. Knowing that other individual knows the feature, and knows that one knows it is viewed as second order iterated belief about the feature. In this way, we can have in general k th-order iterated beliefs about the feature. Thus, game theory usually assumes that players are able to compute up to k th-order iterated beliefs about any feature that forms part of the description of the situations in which they have to decide. Furthermore, that they are assumed to be able do so even when k tends to infinite. Calculating relatively high-order iterated beliefs may be seen as a too demanding computational requirement on the individuals. Some models of the strand of the literature that deal with *bounded rationality* considerations (see, e.g., Rubinstein (1998)) relax this assumption and place restrictions on the kind of high-order inferences that individuals may draw. For instance, some bounded rationality models consider that players are able to draw high-order iterated beliefs about unknown variables only up to a certain bounded level.

that helps them coordinate their actions in a way such that their beliefs and their actual actions be consistent. Not surprisingly, for the NE to work, players need to have common knowledge of such a social institution. We can rephrase this latter requirement as saying that players have common knowledge about the rationale behind the NE notion itself.

To summarize, the methodological assumptions most commonly imposed by game theory are:

1. Players are rational decision-makers.
2. Players are fully informed about the description of the game that they will play.
3. Players have common knowledge about:
 - The entire description of the game that they will play.⁷
 - The fact that they are rational decision-makers.
 - The requirements of the NE solution concept.

Although the list of assumptions above is present in most game theoretical analyses, game theory tools continue to be useful when some of these restrictions are modified. In particular, the following assumptions can sometimes be viewed as too demanding: (1) players represent their preferences using the expected utility form, (2) players can make all sorts of sophisticated computations when deciding about their most preferred course of action, and (3) players are able to keep track of arbitrary high-order beliefs about what they and others know. This, combined with the unsatisfactory experimental and empirical validation of some implications of the theory, have lead some economists to develop alternative models by using most of the standard game theoretical tools, yet by “relaxing,” “twisting,” or “enriching” some of the above discussed basic assumptions. For instance, models where players are assumed to have *cognitive biases* or make computations subject to *bounded rationality* constraints have flourished in the last decades. Despite the pertinence and usefulness of such alternative approaches, these notes will focus on the more conventional game theoretical approach that uses the all above listed assumptions.

⁷Formally, the most commonly imposed assumption 2. that the players know the description of the game is encompassed by the assumption 3. that the players have common knowledge about the description of the game. Assumption 2. is stated separately to emphasize that, even for particular approaches where the common knowledge assumption is relaxed, we need at least need to consider that the players know the description of the game (that is, even if they do not necessarily know that other players knows, that other players know, . . .).

2.2. (Brief) Historical Background

Earlier game theory formulations can be traced back to the works of french mathematicians Augustin Cournot (1801-1877) and Emile Borel (1871-1956). Cournot introduced (in 1838) the ideas underlying the concept of Nash equilibrium—in pure strategies—and reaction functions. Borel proposed (in 1921) the idea of mixed strategy. Hungarian native poli-mathematician John von Neumann (1903-1957) is perhaps the most important figure in the early development of game theory. In 1944, with the publication of their book *Theory of Games and Economic Behavior*, von Neumann and Oskar Morgenstern established the grounds for the subsequent development of the field as we know it today. In the early 1950's, American mathematician John Forbes Nash (1928-2015) made his seminal contributions to zero-sum games, non-cooperative game theory, and to bargaining theory. In 1951, Nash proposed formally the nowadays ubiquitous notion of Nash equilibrium and proved its existence for finite games. In 1952-1953, Lloyd Shapley made the earlier contributions to the theory of cooperative games and introduced the solution concept for cooperative games known as the *Shapley value*. In 1953, Harold Kuhn introduced formally the idea of extensive form games. In 1957, Robert Luce and Howard Raiffa popularized game theory, with the publication of *Games and Decisions: Introduction and Critical Survey*, among economists, psychologists, sociologists, and political scientists. In 1960, Thomas Schelling launched his vision of game theory as a unifying framework of strategic analysis for the social sciences with the publication *The Strategy of Conflict*. Schelling's contribution has proved crucial to the nowadays extensive use of game theory to address problems of military conflict and diplomacy issues between countries. In the late 1950's and early 1960's, Israeli mathematician Robert J. Aumann pioneered the studies of long-term cooperation, in a series of works which can be regarded as the seminal contributions on repeated games. Aumann's figure has been immense to the development of game theory in many respects. For instance, in 1976, he introduced formally the idea of *common knowledge*, which constitutes one of the most used assumptions by the theory. In 1974, Aumann introduced the solution concept of *correlated equilibrium*, from which the Nash equilibrium notion can emerge as a special case. Correlated equilibrium has subsequently proved to be a key solution notion to address contract

theory questions and within the mechanism design research agenda. In 1967-1968 John Harsanyi formalized his method for studying games with incomplete information, which has been crucial to broaden the scope of applications for the theory. In 1975, German economist Reinhard Selten introduced his most refined notion of subgame perfection for extensive form games.

Game theory has received immense appreciation and recognition within economics in the last decades. The unified structure of the theory, together with its broad applicability to key questions raised in the social sciences, have contributed to its fruitful development. In 1994, J. Harsanyi, J. Nash, and R. Selten were awarded the Nobel prize in economics for their seminal contributions to game theory. R. Aumann and T. Schelling were also awarded the Nobel prize in economics, in 2005, for providing key contributions to the subsequent development of the theory, which allowed its extension to other branches of the social sciences. Other Nobel laureates whose contributions to economics rest heavily on the set of analytical tools provided by game theory include: (1) George Akerlof, Michael Spence, and Joseph Stiglitz (2001), for their contributions to information economics; (2) Leonard Hurwicz, Eric Maskin, and Roger Myerson (2007), for their contributions to the study of incentives in markets and mechanism design; (3) Alvin Roth and Lloyd Shapley (2012), for their work on applied market design; (4) Jean Tirole (2014), for his industrial organization insights to market power and regulation, and (5) Oliver Hart and Bengt Holmström (2016) for their contributions to contract theory.

2.3. Non-Cooperative Games *versus* Cooperative Games

These notes restrict attention to *non-cooperative game theory*. Non-cooperative game theory explicitly models individual decisions as primitives, whereas *cooperative game theory* considers instead joint or team actions as primitives in their models. In this respect, cooperative game theory needs to assume the existence of binding agreements, which are previously made by the players, in order to develop the theory and to deliver its results. By taking an axiomatic approach, some strategic problems benefit from a simpler treatment using cooperative games, relative to using the non-cooperative tools. Nonetheless, given the crucial assumptions that the players need to be able to freely communicate among them and reach binding agreements, the scope of cooperative games

for applications is considerably reduced compared to the sort of situations that can be suitably addressed using non-cooperative games. For instance, cooperative games are useful for exploring decisions made by committees, juries, political parties, or groups where it is compelling to consider that their members must respect some binding agreements. In these notes, we will take the non-cooperative philosophical approach that players act individually, without any previous binding agreements. In this respect, non-cooperative game theory considers that the possibilities for communication or agreements must be explicitly modeled, as the result of individual decisions. Therefore, though certainly useful and quite relevant for some applications, cooperative game theory will not be discussed in these notes.

3. Extensive Form Games

3.1. Definitions and Examples

The richest and most structured way of describing a strategic interaction is the *extensive form representation* of the game. Extensive form games are quite convenient to describe situations where the players move sequentially over time. In particular, in an extensive form game we can keep track of when each player is due to decide, what each player knows when deciding, and what outcomes are derived from how the players move over time. In theory, any real-world game where players make their decisions sequentially can be represented as an extensive form game.¹

Definition 3.1. An extensive form game Γ consists of the following ingredients:²

1. A finite set $N = \{0, 1, \dots, n\}$ of players, which will be indexed by $i, j \in N$. More generally, the player indexed as player zero, $i = 0$, will in fact be present only in games with imperfect information. When present, we will refer to player zero as Nature.
2. For each player $i \in N$, a finite set of actions A_i , where an action taken by player i will be denoted as $a_i \in A$. Let $A \equiv \cup_{i \in N} A_i$, with generic element $a^k \in A$, be the set of actions available in the game, regardless of what player takes each particular action.
3. A set Σ of sequences of actions $\sigma = (\sigma^0, a^1, \dots, a^k, \dots, a^K)$, where K can be either finite or infinite, which we will call histories. Furthermore, the set of histories Σ must satisfy the properties:

¹In practice, however, because of the detailed description that it entails, the representation of some real-world games can become very complex and even intractable. Representing the chess game, for instance, turns out unfeasible. For this reason, in order to study particular dynamic strategic interactions, economic models aim at proposing very simple/parsimonious games that capture the relevant elements of such situations.

²The definition of extensive form game that is now standard in game theory is due to [Kuhn \(1953\)](#). The definition used in these notes borrows its particular details from [Rubinstein and Osborne \(1994\)](#).

- (a) $\sigma^0 \equiv \emptyset \in \Sigma$, where σ^0 will be called the initial history, or root, of the game.
- (b) If $\sigma^K = (\sigma^0, a^1, \dots, a^k, \dots, a^K) \in \Sigma$, then $\sigma^k = (\sigma^0, a^1, \dots, a^k) \in \Sigma$ as well, for each $k < K$.
In this case, the history σ^K is called a successor of each σ^k , each σ^k is called a predecessor of σ^K , and each σ^k is called an immediate predecessor of σ^{k+1} , for $k = 1, \dots, K - 1$.
- (c) If an infinite sequence $\sigma = (\sigma^0, a^1, \dots, a^k, \dots)$ satisfies $\sigma^k = (\sigma^0, a^1, \dots, a^k) \in \Sigma$ for each $1 \leq k < \infty$, then $\sigma \in \Sigma$.

A history $\sigma^K = (\sigma^0, a^1, \dots, a^k, \dots, a^K)$ is terminal if it is finite ($K < \infty$) and there is no action $a^{K+1} \in A$ such that $(a^1, \dots, a^k, \dots, a^K, a^{K+1}) \in \Sigma$. We will use $\widehat{\Sigma}$ to denote the set of terminal histories in the game and $A(\sigma) \equiv \{a \in A : (\sigma, a) \in \Sigma\}$ to denote the set of actions immediately available upon the non-terminal history σ .

4. A function $P : \Sigma \setminus \widehat{\Sigma} \rightarrow N$ that assigns a player $P(\sigma)$ to each non-terminal history σ . The function P is known as the player function of the game.
5. For each history σ where Nature plays (i.e., such that $P(\sigma) = 0$), a function $\mu(\cdot | \sigma) : A(\sigma) \rightarrow [0, 1]$ that assigns a probability $\mu(a | \sigma) \in [0, 1]$ to each action $a \in A(\sigma)$, that is, to each action immediately available upon history σ . Each function $\mu(\cdot | \sigma)$ will be called a prior at history σ and the set of functions $\mu \equiv \{\mu(\cdot | \sigma) : \sigma \in \Sigma \text{ with } P(\sigma) = 0\}$ will be a system of priors.
6. For each player $i \in N \setminus \{0\}$, a partition \mathcal{H}_i of the set of histories $\{\sigma \in \Sigma : P(\sigma) = i\}$ where player i plays, with the requirement that for each $\sigma, \sigma' \in h_i$, for each $h_i \in \mathcal{H}_i$, we have $A(\sigma) = A(\sigma')$. Accordingly, we will use $A(h_i)$ to denote $A(\sigma)$ for each $\sigma \in h_i$. Each element h of a partition \mathcal{H}_i will be referred to as an information set where i plays. In consonance with this, \mathcal{H}_i is the set of all information sets where i plays and $\mathcal{H} \equiv \{\mathcal{H}_i\}_{i \in N}$ is the collection of all the players' sets of information sets.
7. For each player $i \in N \setminus \{0\}$, a preference relation \succeq_i over the set of lotteries $\Delta(\widehat{\Sigma})$ on terminal histories.

We will use the list of ingredients $\Gamma \equiv \langle N, A, \Sigma, P, \mu, \mathcal{H}, (\succeq_i)_{i \in N} \rangle$ to denote an extensive form game.

The requirements (a)-(b) of 3. in Definition 3.1 above lead to that the set of histories Σ has the structure of a *tree* with a unique root at the initial history σ^0 . Since $\sigma^0 = \emptyset$ belongs to Σ , we have that σ^0 is a predecessor of each history $\sigma \in \Sigma$. Furthermore, since $\sigma^0 = \emptyset$ is unique by convention, then it is the unique predecessor of each history in the game. We can think of each history as a *node* and of two immediately adjacent histories σ and (σ, a) as two nodes connected by a *directed branch* or arrow that points from σ to (σ, a) . The tree given by the set of histories Σ is then described by the collection of nodes associated to the histories and by the directed branches that connect such nodes. Thus, we can start from a node, trace through the tree by following arrows, and, in this way, reach successors of the departing node.³

The conditional probabilities described in 5. of Definition 3.1 capture the probabilities that Nature chooses actions at each history where it plays. Such probabilities are exogenously given. Two types of uncertainty can be present in a strategic situation. The first one, called *extrinsic (or exogenous) uncertainty* is due to moves by Nature. Intuitively, here we deal with outcomes of chance, where there is no intentional choice by any player—of course, different from Nature. Formally, such random phenomena are captured by random variables. The conditional probabilities $\mu(\cdot | \sigma)$ give us the probabilities associated to such random variables. The second type of uncertainty is known as *intrinsic (or endogenous) uncertainty* and it is due to the fact that players are allowed to randomize between the actions available to them.⁴ Whereas extrinsic uncertainty is only present in situations where chance determines some strategically relevant moves, intrinsic uncertainty will always be present in any game since we always need to capture in the analysis the possibility of any player being uncertain about the decisions made by other players.

The elements of the partitions \mathcal{H}_i described in 6. of Definition 3.1 are known as information sets. They consist of the set of histories which the player called to move upon them cannot distinguish because she has imperfect information. In these cases, the player lacks full information about chance moves or about other players' choices. This will be particularly the case when players move simultaneously or when players simply do not have perfect information about what others have

³Formally, a tree is a directed graph that consists of a set of nodes and a set of directed branches (or arrows) that connects the nodes. In addition, each node must have a single immediate predecessor and there must be a unique node that precedes all nodes. All these requirements are ensured by the requirements in (a)-(b) of 3. of Definition 3.1.

⁴We will discuss this in detail when we introduce the notion of *mixed strategy*.

chosen previously. To summarize, a game will always include intrinsic uncertainty as we allow players to be uncertain about other players' moves when they make their own choices.⁵ Extrinsic uncertainty will be present if there are outcomes of chance—or chance moves. A player will have imperfect information if she does not have full information about previous moves when she is called to play. Formally, this will be captured by non-trivial information sets, where the player cannot distinguish between histories.

The stated condition that $A(\sigma) = A(\sigma')$ for each $\sigma, \sigma' \in h_i$, for each $h_i \in \mathcal{H}_i$, gives us the “informational consistency” requirement that the player is not able to deduce that $\sigma \neq \sigma'$ by observing $A(\sigma) \neq A(\sigma')$. In other words, if this requirement were not satisfied, then we would have a situation where the player, due to the immediately previous moves, does not know where she stands but, at the same time, can in fact learn where she is by observing different sets of actions available upon each history. Given this consistency condition, therefore, for an information set $h \in \mathcal{H}_i$, we can write $A(h) \equiv \{a \in A : (\sigma, a) \in \Sigma \text{ for each } \sigma \in h\}$ to denote the set of actions available upon any of the histories contained in the information set h .

Regarding preferences, as stated in 7. of Definition 3.1, the theory considers that players care only about the final history reached, regardless of the way in which such a history is reached. This “consequentialist” assumption is very common in decision theory under uncertainty.

The elements of the definition of an extensive form game can be illustrated by the game depicted in figure 3.1.

There are two players $i \in N = \{1, 2\}$ that interact in this game. The set of actions available to player 1 is $A_1 = \{A, B, C, F, G\}$ and the set of actions available to player 2 is $A_2 = \{D, E, H, I, J\}$.

⁵This can be due to that the player is not able to observe previous moves, or to that the players move simultaneously, or to that the player must decide before other players are due to make their own choices.




Figure 3.1 – Extensive Form Game with Imperfect Information

The entire set of actions in the game is $A = \{A, \dots, J\}$. The histories that constitute Σ are

$$\begin{aligned} \sigma^0 &= \emptyset, \sigma^1 = (\sigma^0, A), \sigma^2 = (\sigma^0, B), \sigma^3 = (\sigma^0, C), \\ \sigma^4 &= (\sigma^0, A, D), \sigma^5 = (\sigma^0, A, E), \\ \sigma^6 &= (\sigma^0, A, D, F), \sigma^7 = (\sigma^0, A, D, G), \sigma^8 = (\sigma^0, A, E, F), \sigma^9 = (\sigma^0, A, E, G), \\ \sigma^{10} &= (\sigma^0, B, H), \sigma^{11} = (\sigma^0, B, I), \sigma^{12} = (\sigma^0, B, J). \end{aligned}$$

Clearly, $\sigma^0, \dots, \sigma^5$ are non-terminal histories whereas $\sigma^6, \dots, \sigma^{12}$ are terminal histories. The player function is specified by $P(\sigma^0) = P(\sigma^4) = P(\sigma^5) = 1$ and $P(\sigma^1) = P(\sigma^2) = 2$. Nature does not move in this game so that there is no extrinsic uncertainty. The sets of information sets of the two players are given by

$$\begin{aligned} \mathcal{H}_1 &= \{\{\sigma^0\}, \{\sigma^4, \sigma^5\}\}, \\ \mathcal{H}_2 &= \{\{\sigma^1\}, \{\sigma^2\}\}. \end{aligned}$$

Therefore, player 1 cannot distinguish between histories σ^4 and σ^5 . Here $h = \{\sigma^4, \sigma^5\}$ is the only non-trivial information set: it captures a situation where player 1 has imperfect information about what is happening at some point during the game. Intuitively, she does not know whether player 2 have previously chosen action D or action E . Then, player 2 is not able to infer any new information by observing her available actions because $A(\{\sigma^4, \sigma^5\}) = \{F, G\}$, as required by our “informational

consistency” condition. Note that when drawing trees for extensive form games, we draw ellipses, or dotted lines, between nodes to indicate that the corresponding histories belong to the same information set. Both players are assumed to share common knowledge of the description of the game before they actually play it. In particular, player 1 knows that, at point where she is called to play upon histories σ^4 and σ^5 , she will not be able to know whether player 2 has chosen D or E . That is, before playing the game, she knows that she will have imperfect information in that moment. Also, player 2 knows this, and knows that player 1 knows this, and knows that player 1 knows that player 2 knows this, and so on *ad infinitum*. Finally, the pairs $(a, b), \dots, (p, q)$ specify the utility levels, or payoffs, that the players obtain, where the i -th entry of the vector indicates what player i gets. The letters a, \dots, q can be understood as descriptions of the outcomes or as monetary payoffs. Due to the fact that there is always intrinsic uncertainty in the game, players do not care only about the outcomes described by $(a, b), \dots, (p, q)$ but also about the set of lotteries or probability distributions over such outcomes. In this example, therefore, the letters $\{a, c, e, g, i, k, m, p\}$ indicate the payoffs to player 1 while $\{b, d, f, h, j, l, n, q\}$ indicate payoffs to player 2.

The extensive form game depicted in Figure 3.2 describes a situation where an incumbent firm in a market (player 1) interacts with a challenging firm (player 2) that considers whether or not to enter incumbent firm’s market and compete with it. First, the incumbent must choose whether to follow a high cost policy on advertising (H), a medium cost policy (M), or a low cost policy L . If the advertising policy chosen by the incumbent is aggressive, then the game ends and the final payoffs are described by (v_1^1, v_2^1) . In this case, the challenger is not given any option of choosing. However, if the advertising policy chosen by the incumbent is not high, then the challenger can choose whether to enter (e) the market or not enter (n). The challenger cannot distinguish whether the incumbent has chosen a medium or a low cost policy. After the challenger’s action, the incumbent must decide whether it chooses a sales price- bargain policy (x), to compete in the market with any other firm, or not, y . Importantly, here the incumbent does not know whether the challenger has entered the market or not. This might be due to the fact that it is not able to observe the presence of other firm or—perhaps more realistically—to that the incumbent chooses its policy on sales prices

at the same time when the challenger decides whether to enter or not the incumbent's market. Yet, notice that the incumbent does "recall" its previous decision on advertising.




Figure 3.2 – A Market Entry Game

Figure 3.3 represents a classical signaling game. Signaling games are of broad interest for economists and social scientists. The canonical signaling game was proposed by Spence (1973) and, since then, it is been applied to explore a plethora of situations where individuals strategically communicate their *private information*. The situation described here is one where an individual (the Sender) can be selected by chance as being either of high ability (H) or of low ability (L). Using the frequentist interpretation of probability, we can consider that 40% of the population are high-ability individuals. We observe that the Sender knows her own ability. Her ability is then seen as a characteristic of the Sender. When a characteristic due to a move of chance is known privately by some player, we say that such a characteristic is her *private information* and often refer to it as one possible realization of the *player's type*. Formally, a *player's private information or type* is a description of the information that the player has about the game and that it is not common knowledge among all the players. In extensive form games, the players' types are described by the information sets where they are called upon to move. After learning her characteristic, the Receiver chooses whether to invest in—verifiable—education (E) or not (NE). Then, a hiring institution (the Receiver) must choose whether to offer a managerial position (m) or a clerical position (c) to

the Sender. The Receiver does not have more information about the Sender beyond her education level. Therefore, the Receiver can distinguish whether the Sender has chosen E or NE , but cannot distinguish between the Sender's types. The vectors (x, y) that follow the terminal nodes indicate the monetary payoffs to the Sender and to the Receiver, respectively. Interestingly, such payoffs reflect situations where it is relatively more costly for the low ability Sender to invest in education, compared to the cost that the high ability Sender must incur. Such an intuitive feature in the payoff structure makes informative signaling possible in this sort of games.




Figure 3.3 – Signaling Game

3.2. (Pure) Strategies

A key concept in game theory is that of *strategy*. Players make their decisions throughout the play of the game by following a *complete contingent plan* which they elaborate *before interacting among them*. Since the entire description of the game is commonly known among the players, they can elaborate a plan that tells them how to subsequently act once they reach the histories—or, more generally for situations where they have imperfect information, the information sets—where they are due to play.

Definition 3.2. A pure strategy in an extensive form game Γ for a player i is a function $s_i : \mathcal{H}_i \rightarrow A_i$, with the requirement that $s_i(h) \in A(h)$ for each $h \in \mathcal{H}_i$, that specifies an action $s_i(h)$ for each information set

h at which player i is called to play. A pure strategy profile is a list, or profile, $s \equiv (s_i, s_{-i}) = (s_1, \dots, s_n)$ that includes pure strategies for all the players. We will use S_i to denote the set of all pure strategies of player i and $S \equiv S_1 \times \dots \times S_n$ to denote the set of all pure strategy profiles.

We can illustrate the set of pure strategies available to the players using the entry market game depicted in Figure 3.2. Player 1 chooses at the information sets $\{\sigma^0\}$, h_1 , and h'_1 . Using a vector that specifies a (pure) strategy for player 1 as $s_1 = (s_1(\{\sigma^0\}), s_1(h_1), s_1(h'_1))$, we have 12 different strategies, $S_1 = \{s_1^1, \dots, s_1^{12}\}$, specified as

$$\begin{aligned} s_1^1 &= (H, x, x), s_1^2 = (H, x, y), s_1^3 = (H, y, x), s_1^4 = (H, y, y), \\ s_1^5 &= (M, x, x), s_1^6 = (M, x, y), s_1^7 = (M, y, x), s_1^8 = (M, y, y), \\ s_1^9 &= (L, x, x), s_1^{10} = (L, x, y), s_1^{11} = (L, y, x), s_1^{12} = (L, y, y). \end{aligned}$$

As to player 2, we can consider that the single entry $s_2(h_2)$ gives us the two different strategies, $S_2 = \{s_2^1, s_2^2\}$, available to her as

$$s_2^1(h_2) = e, s_2^2(h_2) = n.$$

It is very important to make clear one particular aspect of the notion of strategy. Suppose for instance that you are playing a game all by yourself where you must decide between the existing railway routes in order to plan a train trip within Europe that begins in Paris. Then, starting from Paris, the notion of strategy requires that you choose between the available routes from each train station in Europe. Importantly, a strategy must specify a possible (direct) destination from each departing station in Europe even though your previously chosen route does in fact not allow you to reach such a departing station in the first place. In other words, even though the choices contained in a particular strategy do not let you reach Prague, such a strategy must specify a possible destination from Prague! This is the sense in which a strategy must be a *complete* contingent plan. Notice that planning for contingencies that would never arise if the player follows the strategy it is always theoretically feasible for the player—so long as one does not consider costs from designing strategies—and it does not affect the generality of the definition of strategy. More importantly, such choices at “never-reached” histories turn out crucial to study dynamic games. In particular,

the equilibrium notions typically used to study the players' behavior rely crucially on such choices.

Given an extensive form game representation, each (pure) strategy profile s leads to a terminal history which, in turn, gives each player a utility level or payoff. We will use $O(s)$ to denote the terminal history which is induced by the strategy profile s . In other words, O is a function $O : S \rightarrow \widehat{\Sigma}$ that associates a terminal history $\sigma = O(s)$ to each strategy profile s . Under the conditions we have imposed on the set of histories Σ , the function O turns out to be a *bijection* for any extensive form game. As a consequence, we can consider preferences either over terminal histories or over strategy profiles in a totally equivalent way.

Let $\pi \equiv (\pi(s))_{s \in S} \in \Delta(S)$ be a lottery over the set S of strategy profiles so that $\pi(s)$ indicates the probability that the strategy profile s is chosen by the players. Under some axioms usually assumed by the strand of theory that explores individual rational choice under uncertainty,⁶ a preference relation \succeq_i over lotteries π can be represented by a utility function that has the *expected utility form*. This utility representation under uncertainty was first proposed by [von Neumann and Morgenstern \(1944\)](#). Suppose that we use a (Bernoulli-type) utility function u_i to represent player i 's preferences over the set of pure strategy profiles S . Then, given that each strategy profile induces a terminal history and that each terminal history is achieved induced by a unique strategy profile (thanks to that the function O is a bijection), u_i can equivalently be used to represent player i 's preferences over the terminal histories σ : if $u_i(s) \geq u_i(s')$ whenever player i (weakly) prefers the terminal history $O(s)$ rather than $O(s')$. Then, a preference relation \succeq_i over the set of lotteries $\Delta(S)$ is represented by a utility function $U_i : \Delta(S) \rightarrow \mathbb{R}$ with the *expected utility form* if it has the form

$$U_i(\pi) \equiv \sum_{s \in S} \pi(s) u_i(s) = \sum_{O(s) \in \widehat{\Sigma}} \pi(O(s)) v_i(O(s)),$$

where the function v_i is used to represent player i 's preferences over (certain) terminal histories. Note that we simply express u_i as the composite function $u_i(s) \equiv v_i(O(s))$. Often, the consequences associated to the terminal histories of an extensive form game are represented by monetary payoffs.

⁶In particular, these axioms state that any preference order under uncertainty be: (I) complete, (II) rational, (III) continuous, and (IV) rank any two lotteries independently of any common third lottery. While Axioms (I)-(III) are the basic assumptions required by the theory of choice under certainty to represent preferences which impose relatively mild conditions, Axiom (IV) is the most controversial one when uncertainty is added to the analysis. Axiom (IV) is sometimes challenged by experimental and empirical evidence.

In such cases $u_i(s)$ is real number that indicates the “monetary value” that player i assigns to the terminal history $O(s)$.⁷ This common practice was already used in the signaling game depicted in figure 3.3.

3.3. (Pure Strategy) Nash Equilibrium

Nash (1951)’s solution concept is probably the most influential solution concept in game theory. The idea of Nash equilibrium is that if the players are figuring out which strategies they should play (in the pursue of their individual interests) and, in addition, if they share a common understanding of the logic of their optimal decisions, then they should either end up with a Nash equilibrium or assign irrational behavior to other players. This argument relies on the assumption that the players choose their strategies independently so that a player’s change of her selected strategy cannot change the strategies chosen by other players. Recall that the players choose their strategies before actually interacting with others so that this a natural assumption.

Definition 3.3 (NE in pure strategies). *A pure strategy Nash equilibrium (NE) of an extensive form game Γ is a pure strategy profile $s^* = (s_1^*, \dots, s_n^*)$ such that each player $i \in N$ weakly prefers the (certain) consequence $O(s^*)$ rather than $O(s_i, s_{-i}^*)$ for each $s_i \in S_i$.⁸*

Game theory commonly uses (Bernoulli-type) utility functions $u_i : S \rightarrow \mathbb{R}$ to describe how each player i evaluates each of the consequences induced by the strategy combinations chosen in the play of the game. Here, $u_i(s)$ indicates the utility level or payoff that accrues to player i from the (certain) terminal history $\sigma = O(s)$. As already mentioned, the number $u_i(s)$ can be interpreted in monetary terms but, in general, it can be taken as a much broader assessment. Then, using a typical utility representation under certainty, we can rewrite the definition of (pure- strategy) NE

⁷This is made by tractability reasons as expressing the utility that accrues to the players in monetary terms is appealing. As already discussed, however, rationality does not mean that the players pursue exclusively monetary goals.

⁸Given any profile $z = (z_1, \dots, z_n)$ for all players involved in a game, the notation $z = (z_1, \dots, z_n) = (z_i, z_{-i})$ is typical in game theory. Thus, z_{-i} indicates the profile obtained by excluding player i from the overall profile z , that is, $z_{-i} \equiv (z_1, \dots, z_{i-1}, z_{i+1}, \dots, z_n)$.

in 3.3 as s^* is a NE (in pure strategies) if for each player $i \in N$, we have:

$$u_i(s^*) \geq u_i(s_i, s_{-i}^*) \text{ for each } s_i \in S_i \text{ or, equivalently, if } s_i^* \in \arg \max_{s_i \in S_i} u_i(s_i, s_{-i}^*).$$

The rationale behind the NE notion is as follows. When choosing their most preferred strategies, players need to form beliefs about the strategies that the other players select. Consider an external (sometimes fictional) instrumental institution, or mediator, that fully informs each player about the choices made by the other players. Then, the NE notion requires that each player chooses the course of action that maximizes her utility given that she believes the information provided by the institution. Why would the players trust on the information provided by this mediator? The trick is that all players commonly know the description of the game, that other players are rational, and the logic underlying the NE concept. Then, all players can put themselves in the positions of the other players and compute how such players will behave given their own optimal choices. The NE concept then requires that the computations of how other players behave coincide with the beliefs provided by the external mediator. The players have every reason to believe the information provided by the external institution when they consider that the other players are also trusting such information and behaving accordingly to maximize their own expected utilities. If this story is commonly known among the players, then the requirement in Definition 3.3 gives us a reasonable and appealing description of how rational players will make their choices.⁹ In this sense, a NE strategy profile becomes a self-fulfilling prophecy: if all the players believe that the others are choosing their equilibrium strategies, then each player wishes to choose an equilibrium strategies so that the players' beliefs are in fact correct.

We can use the extensive form game depicted in Figure 3.4 to compute NEs. Player 1's (pure) strategies can be summarized as A and B , and Player 2's (pure) strategies can be described by

⁹Sometimes, the external institution or mediator is not fictional at all. Consider the case of a traffic light at a crossroad where we cannot see the cars passing by the intersecting road. When the light says red, the signal is informing us that the drivers in the intersections are passing. Our optimal strategy, if we believe such information, is then to stop and wait. We can put ourselves in the position of the drivers in the intersecting road and consider that they will anticipate that our most preferred choice is to stop and wait. Thus, they will believe on the information provided by the green light they are observing. Also, we will believe that their most preferred choice is to pass, which coincides with the information provided to them by the signal. Waiting when the signal says red and passing when it says green is a NE instrumented by the traffic light as an external mediator.

$CF, CG, DF, DG, EF,$ and EG . By combining the sets of pure strategies of both players, we can




Figure 3.4 – Computing NE—Example 1

write the matrix of payoffs in Figure 3.5. This matrix is very useful to compute how each player determines her most preferred choices, given the choice of the other player. Suppose that player 1 picks strategy A . Then, the most preferred strategies for player 2 are $\{CF, CG, EF, EG\}$, which give her a payoff of 2. Now, if player 2 chooses strategies EF or EG , then player 1's most preferred choice is B rather than A . Given that player 2 picks EF , by choosing B instead of A , player 1 moves from obtaining a payoff of 0 to getting 2. On the other hand, provided that player 2 picks EG , by moving from A to B player 1 obtains a payoff of 1 rather than 0. Therefore, such strategy combinations are not NEs. On the other hand, given that player 2 picks CF , player 1 would benefit by choosing B rather than A , which implies that (A, CF) is not a NE either. On the other hand, for the strategy profile (A, CG) , we observe that each of the two players is choosing a payoff maximizing strategy, given the strategy chosen by the other player. Thus, (A, CG) is a NE. When we obtain (strictly) higher payoffs by choosing a strategy different from a selected one, we will informally say that we have “(strict) incentives to deviate.” Clearly, no player can have (strict) incentives to deviate in a NE strategy profile. Contrary to the case of an equilibrium profile, if a player has (strict) incentives to deviate from an strategy candidate to equilibrium, then the suggested strategy combination is a self-denying prophecy. In practice, when computing NE, we will check for situations where no player has (strict) incentives to deviate, given the other players' strategies. We will rule out those situations where at least one player wishes to deviate. Analogously, we observe that (B, CG) , (B, DG) , and (B, EG) are NE strategy profiles as well. For an extensive form game Γ , we will use

the notation $NE(\Gamma)$ to indicate its set of NEs. In this example, we have

$$NE(\Gamma) = \{(A, CG), (B, CG), (B, DG), (B, EG)\}.$$

		②					
		<i>CF</i>	<i>CG</i>	<i>DF</i>	<i>DG</i>	<i>EF</i>	<i>EG</i>
①	<i>A</i>	1, 2	1, 2	1, -1	1, -1	0, 2	0, 2
	<i>B</i>	2, 2	1, 3	2, 2	1, 3	2, 2	1, 3

Figure 3.5 – Payoffs of Example 1

We can similarly compute the set of NEs of the extensive form game displayed in Figure 3.6. In this case, we can derive the payoff matrix represented in Figure 3.7. Here, we can look for the




Figure 3.6 – Computing NE—Example 2

strategies that yield each player maximum payoffs, taken as given the strategy chosen by the other player, so as to obtain

$$NE(\Gamma) = \{(L, b), (R, b)\}.$$

		②			
		<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
①	<i>L</i>	-1, 1	-1, 2	1, -1	1, -1
	<i>R</i>	1, -1	-1, 1	1, -1	-1, 1

Figure 3.7 – Payoffs of Example 2

As another example, consider the extensive form game depicted in Figure 3.8. For this game,




Figure 3.8 – Computing NE—Example 3

we can derive the payoff matrix represented in Figure 3.9. Here we can compute the set of NEs as

$$NE(\Gamma) = \{(AGIK, CE), (AGIK, CF), (AGIL, CE), (AGIL, CF), \\ (AGJK, CE), (AGJK, CF), (AGJL, CE), (AGJL, CF)\}.$$

Notice that this set of NEs can be described “in essence” as player 1 moves first A and then G, while player 2 moves C. Other possible moves are in fact captured by the complete description of the set $NE(\Gamma)$ above. However, we observe that some of the moves taken into account in the description are irrelevant with respect to the outcomes achieved in equilibrium.

When accounting for combinations of the players’ pure strategies sometimes one derives lengthy payoff matrices, such as the one obtained in Figure 3.9. This motivates the question of whether we can somehow “simplify” such payoff matrices by considering as one several pure strategies of a given player. More precisely, consider a given player and suppose that the combination of several of her strategies with the other players’ strategies yield exactly the same payoffs to each player. Then, such strategies of the given player are said to be *strategically equivalent* and merging them into just a single strategy does not affect neither the strategic interactions among the players nor the plausible outcomes that the game makes available to the players. In such instances, we can “simplify” the corresponding payoff matrices by merging all strategically equivalent strategies of a given player into a single strategy.

		②			
		CE	CF	DE	DF
	AGIK	6,3	6,3	1,2	1,2
	AGIL	6,3	6,3	1,2	1,2
	AGJK	6,3	6,3	1,2	1,2
	AGJL	6,3	6,3	1,2	1,2
	AHIK	1,4	1,4	3,2	3,2
	AHIL	1,4	1,4	3,2	3,2
	AHJK	1,4	1,4	3,2	3,2
①	AHJL	1,4	1,4	3,2	3,2
	BGIK	2,3	2,-3	2,3	2,-3
	BGIL	2,3	3,-2	2,3	3,-2
	BGJK	1,4	2,-3	1,4	2,-3
	BGJL	1,4	3,-2	1,4	3,-2
	BHIK	2,3	2,-3	2,3	2,-3
	BHIL	2,3	3,-2	2,3	3,-2
	BHJK	1,4	2,-3	1,4	2,-3
	BHJL	1,4	3,-2	1,4	3,-2

Figure 3.9 – Payoffs of Example 3

Definition 3.4. We say that two different strategies $s_i, s'_i \in S_i$ of a given player $i \in N$ are strategically equivalent if for each player $j \in N$, we have $u_j(s_i, s_{-i}) = u_j(s'_i, s_{-i})$ for each $s_{-i} \in S_{-i}$.

We can apply the logic of strategically equivalent strategies to the payoffs displayed in Figure 3.9. Notice that the four strategies $AGIK, AGIL, AGJK, AGJL$ of player 1 give exactly the same payoffs to each player 1 and 2 when combined with all the strategies available to player 2. These strategies could therefore be merged into a single one. In fact, the common actions to all these four strategies, which are also the relevant ones to determine the payoffs achieved by the players (when combined with player 2's strategies), are actions A and G . We can apply the idea behind strategically equivalent strategies to simplify the set of strategies $AHIK, AHIL, AHJK, AHJL$, where actions A and H are the relevant ones. In addition, following the same reasoning, we can merge the following pairs of strategies as well: $BGIK$ and $BHIK$ into BIK ; $BGIL$ and $BHIL$ into BIL ; $BGJK$ and $BHJK$ into BJK ; $BGJL$ and $BHJL$ into BJL .

The payoff matrix in Figure 3.9 could be simplified to the matrix displayed in Figure 3.10. We commonly use the term *semi-reduced payoff matrix* to refer to the matrix the results from merging all strategically equivalent strategies of the original payoff matrix.

		②			
		CE	CF	DE	DF
①	AG	6,3	6,3	1,2	1,2
	AH	1,4	1,4	3,2	3,2
	BIK	2,3	2,-3	2,3	2,-3
	BIL	2,3	3,-2	2,3	3,-2
	BJK	1,4	2,-3	1,4	2,-3
	BJL	1,4	3,-2	1,4	3,-2

Figure 3.10 – (Semi-Reduced) Payoff Matrix of Example 3

Figure 3.11 depicts the extensive form game that corresponds to the classical *prisoner's dilemma* situation. The prisoner's dilemma gives us a general framework which illustrates that NE outcomes do not need to coincide with socially desirable outcomes. The situation can be described by considering two players that have committed two crimes, one minor crime for which their guilt can be proved without any confession, and one major crime where conviction needs that at least one of the players confesses. The prosecutor establishes that if only one player confesses (*D* for "Default"), then she will go free (which gives her a payoff of 6) while the accused player will incur a long sentence (which gives her a payoff of 0). If none of the players confesses (*C* for "Cooperate"), then both will incur a short sentence for the minor offense (which gives each a payoff of 5). If both players confess, then they will be sentenced with a medium term period in jail that amounts to the long sentence for being found guilty of the major offense minus a reduction for confessing. This gives each a payoff of 1. The situation displayed in Figure 3.10 considers that player 1 decides first between confessing or not and then, without knowing the decision of player 1, player 2 makes the analogous choice.

For this game, we can derive the payoff matrix represented in Figure 3.12. We observe that the only NE is (D, D) which yields a payoff of 1 to each player. This is not a socially desirable outcome. In particular, an outcome is said to be (*weakly*) *Pareto efficient* if there is no other outcome that would make all players better off. In this example, the unique NE strategy profile entails an outcome that gives the players and outcome of 1 and yet there is another outcome that makes *all* of them (strictly) better off. For this outcome to be achieved both players have to not confess. This a situation socially more desirable than the situation that the NE concept predicts. This example illustrates the tension often present in economics between actual outcomes predicted by




Figure 3.11 – Prisoner’s Dilemma

the theory (*positive implications*) and socially desirable outcomes according to certain (very weak) philosophical criteria (*normative implications*). This tension is precisely what makes the message conveyed by the prisoner’s dilemma situation very interesting.

		②	
		D	C
①	D	1, 1	6, 0
	C	0, 6	5, 5

Figure 3.12 – Payoffs of the Prisoner’s Dilemma Game

The strategic situation described by the prisoner’s dilemma can be “rephrased” in terms more familiar to economists using typical motivations within the field of *industrial organization (IO)*. Consider a duopolistic market for some homogenous good where two firms compete by choosing the quantities they produce and offer. Suppose that the market is in equilibrium in the sense that supply equals demand. Firms choose quantities and the available demand for the good determines the final price of the good. The firms’ profits depend on the quantities that they produce (and sell) and on the final price of the good. Let $q_i \geq 0$ be the quantity of the good produced by firm $i \in \{1, 2\}$. Let $P = a - b(q_1 + q_2)$ be the (inverse) demand function for the good in this market, with $a, b > 0$. Let cq_i be the cost to each firm i from producing the quantity q_i , where $0 < c < a$. In this environment, the profits to firm i are given by the expression

$$\Pi_i(q_i, q_j) = [a - b(q_i + q_j)]q_i - cq_i,$$

where we use j to denote the player other than i , that is $\{j\} \equiv \{1,2\} \setminus \{i\}$. Suppose that firm i believes that firm j will choose quantity q_j^* . Then, firm i 's goal according to the NE requirements is to choose a quantity $q_i^* \geq 0$ that solves the problem

$$\max_{q_i \geq 0} [(a - c) - b(q_i + q_j^*)]q_i.$$

Assuming the nontrivial case $q_j^* > 0$, this problem is solved by considering the *first-order condition* $\partial \Pi_i(q_i^*, q_j^*) / \partial q_i = 0$, which gives us the linear expression¹⁰

$$q_i^* = \frac{a - c}{2b} - \frac{1}{2}q_j^*. \quad (3.1)$$

The expression of the so called *reaction function* $q_i^* = RF_i(q_j)$ in (3.1) is valid for both firms $i = 1, 2$ so that we obtained a linear system with two equations and two unknowns. Solving the system, we obtain a unique solution $q^* = (q_1^*, q_2^*)$ where

$$q_1^* = q_2^* = \frac{a - c}{3b},$$

which, in turn, yields the optimal profits

$$\Pi_i(q^*) = \frac{(a - c)^2}{9b} \quad \text{for both firms } i = 1, 2.$$

This solution gives us the unique NE of this game and the payoffs above describe the corresponding NE outcome. Figure 3.13 displays the two reaction functions of the competing firms. Here we can observe how quantities different from the NE choices would adjust as the firms use their reaction functions to pick their optimal quantities conditioned on the quantities chosen by their rivals.

Now, consider that instead of competing between them, the firms decide to collude in order

¹⁰The model proposed here of duopolistic competition in quantities, as well as the solution obtained, was first proposed in 1838 by the french economist Augustin Cournot (Cournot (1929)). Cournot termed the linear expression in (3.1) as *reaction function*. This equation tells us what quantity would a firm optimally choose conditioned on the quantity taken by its rival. By solving this model, Cournot indeed derived in fact a NE. Cournot did not propose a solution concept for games in general. Yet, the NE concept is also sometimes known as Cournot-Nash Equilibrium to acknowledge Cournot's seminal contribution to the ideas underlying the NE notion, which was not formally developed until more than a century later.




Figure 3.13 – Reaction Functions for Cournot Competition

to choose jointly the quantities that maximize the sum of their profits. This can be viewed as a situation where both firms merge into a single one. The profits of the merger would therefore be

$$\tilde{\Pi}(q) \equiv \Pi_1(q) + \Pi_2(q) = [(a - c) - b(q_1 + q_2)](q_1 + q_2).$$

Suppose that the collusion agreement considers that both firms must offer the same quantities. Then, the problem of the merge of firms would be to choose quantities $\tilde{q}_1, \tilde{q}_2 \geq 0$, with $\tilde{q}_1 = \tilde{q}_2$, that maximize the profit function derived above. Algebraically, we would use the *first-order conditions* $\partial \tilde{\Pi}(\tilde{q})/\partial q_1 = 0$, $\partial \tilde{\Pi}(\tilde{q})/\partial q_2 = 0$, and $\tilde{q}_1 = \tilde{q}_2$, to obtain

$$\tilde{q}_1 = \tilde{q}_2 = \frac{a - c}{4b}.$$

Now, these quantities yield the following profits

$$\Pi_i(\tilde{q}) = \frac{(a - c)^2}{8b} \quad \text{for both firms } i = 1, 2.$$

Thus, while the chosen quantities are larger in the NE outcome than in the collusive agreement, the profits from the NE to each of the firms are smaller than the ones under collusion. Now suppose

that firm j chooses indeed the collusive quantity instead of following the NE recommendation, that is, $\tilde{q}_j = (a - c)/4b$. Then, the expression obtained in (3.1), tell us that firm i would optimally choose the quantity $\bar{q}_i = 3(a - c)/8b$. This would give profits $\Pi_i(\bar{q}_i, \tilde{q}_j) = 9(a - c)^2/64b$ to firm i . We have that $\Pi_i(\bar{q}_i, \tilde{q}_j) > \Pi_i(\tilde{q})$. On the other hand, the profits to firm j would be $\Pi_j(\bar{q}_i, \tilde{q}_j) = 6(a - c)^2/64b$, so that $\Pi_j(\bar{q}_i, \tilde{q}_j) < \Pi_j(q^*)$. The analogous symmetric implications would follow if only firm i complies with the collusive agreement. Then, if we restrict attention to letting each of the firms to decide (simultaneously) only between offering the NE quantity or the collusive quantity, then we obtain a payoff matrix that corresponds exactly to a prisoner's dilemma situation. Suppose without loss of generality that $(a - c) = 1$. Then, such payoffs are displayed in Figure 3.14.

		②	
		C	D
①	C	$\frac{1}{8b}, \frac{1}{8b}$	$\frac{(3/4)}{8b}, \frac{(9/8)}{8b}$
	D	$\frac{(9/8)}{8b}, \frac{(3/4)}{8b}$	$\frac{1}{9b}, \frac{1}{9b}$

Figure 3.14 – To Collude or Not to Collude in Oligopolistic Cournot Competition

We observe that following the collusive agreement (C, C) yields higher payoffs to each of the firms. However, if the firms decide simultaneously and no additional considerations (such as penalties from not respecting the collusive agreement) are incorporated in the firms' final payoffs, then the firms would follow the NE recommendation which gives them lower profits.

3.4. "Backwards Induction" and Subgame Perfect Equilibrium

For some games where the players take decisions sequentially, their dynamic structures allow us to consider some NEs as "more reasonable or appealing" than others.¹¹ Intuitively, there are situations where players choose their optimal actions given the optimal actions of the others, as the NE notion requires, but where such situations would not be achieved in the first place if the players followed NE recommendations at earlier stages of the game. The additional criterion that we impose here is known as *sequential rationality*, which requires that each player chooses her most

¹¹ The process where we begin with several NEs and impose additional criteria to determine which of those NEs make more sense than others is known as *equilibrium refinement* or *equilibrium selection*.

desirable action conditional on being at each information set where she has to play. If this criterion is common knowledge, then each player will “look ahead” at each point in time to determine how all the players will behave optimally in the future. After incorporating such information, the player will determine her preferred option in the present. In this way, the players start by positioning themselves in the final stages of the game and then determine the optimal choices of everyone at each preceding stage by incorporating the information of how they behave after each stage. The players reason therefore in a “backwards” fashion. We refer as *backwards induction* to the process of analyzing an extensive form game from back to front and, in doing so, of determining how each player chooses her most desired actions at each information set where she is due to play, so that the sequential rationality condition is always satisfied.

Consider the extensive form game depicted in Figure 3.15, which is known as a *centipede game*. In a centipede game, each player is given, in alternating turns, the choice of stopping the game (S) or continuing it (C). The payoffs to both players increase as the game continues.




Figure 3.15 – Six Turn Centipede Game

In this game, each of the two players has eight possible (pure) strategies and the corresponding payoff matrix (after simplification of strategically equivalent strategies) is shown in Figure 3.16.

		②			
		S	CS	CCS	CCC
①	S	1,0	1,0	1,0	1,0
	CS	0,2	3,1	3,1	3,1
	CCS	0,2	2,4	5,3	4,6
	CCC	0,2	2,4	4,6	6,5

Figure 3.16 – (Semi-Reduced) Payoff Matrix of the Centipede Game

The unique NE of this game is (S, S) and its payoffs appear in the highlighted cell. Backwards

induction can be applied here as follows. Player 2 will choose S over C at the last information set where she is due to pay so as to collect a payoff of 6 instead of 5. Both players would commonly agree on that and discount it. Given this, player 1 would have previously chosen S rather than C so as to get a payoff of 5 instead of 4. Again, this would be commonly agreed upon and discounted by both players. As a consequence, player 2 would have previously chosen S rather than C so as to get a payoff of 4 instead of 3. Then, player 1 would have previously chosen S rather than C so as to get a payoff of 3 instead of 2. Given this, player 2 would have previously chosen S rather than C so as to get a payoff of 2 instead of 1. This takes us to the root of the tree where, *given the correctly anticipated future sequence of optimal choices*, player 1 wishes S rather than C so as to get a payoff of 1 instead of 0. In short, the only NE outcome compatible with the sequential rationality criterion here entails that player 1 moves S right at the beginning of the game, leading the play to an end from the very beginning. The NE strategies that satisfy the sequential rationality requirement are highlighted in red in Figure 3.14. We observe that the only NE of this game satisfies the sequential rationality criterion. Makes sense, right?

In some instances, sequential rationality (and the corresponding backwards induction reasoning process) does suggest us to draw attention away from some NEs because they are less plausible than others. Consider again the game depicted in Figure 3.4. Let us position ourselves on the last moves of the game and analyze how the players would optimally choose. In particular, upon history (σ^0, A) , player 2 will choose either action C or E , which give her a payoff of 2. Both players commonly agree on this. On the other hand, upon history (σ^0, B) , player 2 would choose G , which gives her a payoff of 3. Again, both players commonly agree on this. Then, upon history σ^0 player 1 would choose either action A or B because she knows that both players commonly agree on that, by doing so, player 1 would get a payoff of 1. Therefore, following the sequential rationality logic, we obtain the set of equilibria:

$$\{(A, CG), (B, CG), (B, EG)\}.$$

In this case, using the sequential rationality logic, we would rule out the previously obtained NE (B, DG) as less plausible than others.

As another example, consider the extensive form game depicted in Figure 3.17. The semi-reduced payoff matrix for this game is shown in Figure 3.18. We observe that both (UA, X) and (D, Y) are NEs for this game. However, it is not the case that both are subgame perfect. In particular, for the subgame that initiates at history (σ^0, U) , the strategy combination (UA, X) induces the NE (A, X) in the subgame. In addition, (UA, X) is a NE for the entire game. Therefore, it constitutes the only NE that is also subgame perfect. Here we can rule out the NE combination (D, Y) as less plausible than (UA, X) .




Figure 3.17 – Subgame Perfection Refines NEs

		②	
		X	Y
①	UA	3, 4	1, 4
	UB	2, 1	2, 0
	D	2, 6	2, 6

Figure 3.18 – (Semi-Reduced) Payoff Matrix

Sequential rationality can be applied in general to extensive form games. Selten (1965, 1975) developed a notion of equilibrium, known as *subgame perfect equilibrium*, that adds the sequential rationality requirement to the NE conditions. The concept of subgame perfection requires first to consider the description of what it is a subgame of an extensive form game.

Definition 3.5. Given an extensive form game Γ , a history $\sigma \in \Sigma$ in the tree representation of the game is said to be the root of a (proper) subgame if neither σ nor any of its successors are in an information set that contains histories that are not successors of σ . A (proper) subgame is the tree structure specified by σ and its successors.

To practice with identifying (proper) subgames, consider again the extensive form game depicted in Figure 3.1 and let us find the subgames available in this game. First, notice that the root of any extensive form game always initiates a subgame (which in fact coincides with the original game). Thus, here $\sigma^0 = \emptyset$ initiates the subgame that coincides with the entire game. Secondly, note that the history $\sigma^1 = (\sigma^0, A)$ also initiates a subgame. In particular, all the successors of σ^1 are in information sets that contain histories that are successors of σ^1 . We observe that the history $\sigma^2 = (\sigma^0, A, D)$ is not the root of a subgame. Specifically, σ^2 belongs itself to an information set that contains another history $\sigma^3 = (\sigma^0, A, E)$ which is not a successor of σ^2 . Analogously, σ^3 does not initiate a subgame either. Finally, we can verify that $\sigma^4 = (\sigma^0, B)$ is the root of a subgame since all its three successors are in (trivial) information sets that contain histories that are also successors of σ^4 . Therefore, this game has three (proper) subgames whose respective roots are σ^0 , σ^1 , and σ^4 .

With the notion of (proper) subgame at hand, we can now present the concept of subgame perfection.

Definition 3.6. Given an extensive form game Γ , a strategy profile s^* is a subgame perfect equilibrium if it specifies a NE in every (proper) subgame of the original game Γ .

Consider the extensive form game displayed in Figure 3.19. Here we have a subgame, highlighted in red, that corresponds to the prisoner's dilemma game. We know that (D, D) is the only NE that this subgame yields. In particular, player 1 optimally chooses D . This gives each player a payoff of 1. On the other hand, as highlighted in blue, we observe that player 1 optimally chooses D upon history (σ^0, B, C) but she chooses C upon history (σ^0, B, D) . Given this, player 2 would optimally choose D upon history (σ^0, B) since this gives her a payoff of 6, which she clearly prefers to the alternative achievable (under sequential rationality) payoff of 2. Given these NE choices in each of the three subgames already identifies, it remains to determine which will be the optimal

choice of player 1 at the root of the tree. We see that player 1 would optimally choose B so as to get a payoff of 2 rather than a payoff of 1.




Figure 3.19 – Computing Subgame Perfect Equilibria

As another example, consider the extensive form game displayed in Figure 3.20. Here, we consider that player 3 moves first and, after choosing between X and Y , players 1 and 2 choose simultaneously between X and Y . Instead of showing the corresponding subtree representations, Figure 3.19 displays directly the payoff matrices associated to the two subgames that initiate upon player 3's choice. Notice that, in order to compute NEs using the two displayed payoff matrices, we need to consider only the entries corresponding to players 1 and 2 because they are the only players due to move at such subgames. First, we observe that both (X, X) and (Y, Y) are NEs in subgame associated to history σ^1 . If asked, both players would agree on that they prefer the NE (Y, Y) . Indeed, if the society consisted only of players 1 and 2, then (Y, Y) Pareto would dominate (X, X) . Secondly, for the subgame associated to history σ^2 , we observe that (X, Y) is the only induced NE. Now, using backwards induction, player 3 would anticipate these three NEs and, therefore, her optimal choice would be Y . The only subgame perfect equilibrium of this game is (Y, X, Y) , which leaves players 1 and 2 with a lower payoff than their "preferred NE" for those subgames where they are the only players choosing.




Figure 3.20 – Three Player Game: Combining Extensive Form with Matrix Form

3.5. Behavior Strategies

To be able to fully analyze strategic settings in general, we need to account for the possibility that players choose randomly between their available options. Decisions where the players assign probabilities to their actions are naturally appealing in some environments where choosing randomly is interpreted as if the player “hesitates between her actions” and, when asked about her choice, she answers “with this given probability I would action such an action”. Also, using the frequentist interpretation of a probability, another way of viewing a random choice is that of considering that the player faces a (relatively large) number of times exactly the same decision (under the same circumstances). Then, she would choose in each instance an action without any hesitation whatsoever (thus, with probability one). The player, however, would be able to choose different actions across the identical situations, in accordance to the corresponding probabilities. For some environments, we indeed wish to use probabilities to describe best how the players interact.¹² One way to describe formally how the players choose randomly in situations represented as extensive form games is that given by the idea of *behavior strategy*.

Definition 3.7. *Given an extensive form game Γ , a behavior strategy for player $i \neq 0$ is a collection of*

¹²Other than the need of allowing for random decisions to capture more suitably outcomes in certain environments, allowing for random decisions is a crucial technical requirement for the set of tools of game theory to be useful in predicting/explaining behavior in strategic settings. Specifically, the NE solution can be proved to always exist only if we allow for random choices.

probability distributions

$$\rho_i = \{ \rho_i(\cdot | h) \in \Delta(A(h)) : h \in \mathcal{H}_i \text{ such that } P(\sigma) = i \text{ for each } \sigma \in h \}.$$

Here, $\rho_i(a | h)$ is the probability that, upon being called to play at each history σ in the information set h , player i chooses her immediately available action a . Thus, a behavior strategy ρ_i has conceptually the same meaning as the prior notion μ which is part of the description of an extensive form game. The only difference is that priors are exogenously given (as a choice by chance) whereas behavior strategies are endogenously chosen by players that pursue their goals. We will use $\rho \equiv (\rho_i, \rho_{-i}) = (\rho_1, \dots, \rho_n)$ to denote a *behavior profile*. Each behavior profile ρ , combined with the system of priors μ (if available) induces a probability distribution $\pi \in \Delta(\widehat{\Sigma})$ over the set of terminal histories of the game or, equivalently, over the final payoffs $v_i(\sigma)$ that accrue to each players i . In practice, given a system of priors μ , we can use the *total probability law* to compute how a profile ρ induces a probability distribution π over payoffs $v_i(\sigma)$.

4. Strategic Form Games

4.1. Strategic (or Normal) Form Representations

In some strategic situations, players make their choices once and for all and, furthermore, they pick their actions simultaneously. Such situations are best described by *strategic* or *normal* form games, which, by definition, are not explicit about any sequences of moves and consider that each player has not information whatsoever about the choices of others when making her own choices. Some environments adjust naturally to these features. Alternatively, game theorists sometimes find useful to “translate” a dynamic interaction, originally described by an extensive form game, into a game where all players move simultaneously. In this way, the analyst intentionally abstracts from the actual sequences of moves, as well as from the specific description of what each player knows at each point, during the play. Thus, by expressing an original extensive form game into its corresponding strategic form, we ignore pieces of information provided by the extensive form description. One advantage of working with the strategic form representation of a game is that it facilitates the computation of NEs. Then, we can always go back to the original extensive form game for further considerations as to how the players would optimally behave, or for selecting the most plausible equilibria from the obtained set of NEs.

Definition 4.1. A strategic form, or normal form, game Γ consists of $\Gamma = \langle N, S, (u_i)_{i \in N} \rangle$, where:

- (1) $N = \{1, \dots, n\}$ is a finite set of players;
- (2) $S = \times_{i \in N} S_i$ is a set of profiles $s = (s_1, \dots, s_n)$ of (pure) strategies for the players;
- (3) $u_i : S \rightarrow \mathbb{R}$ is utility function for player i such that $u_i(s)$ is the (certain) payoff that accrues to player i when the players choose the strategy profile s .

4.2. Mixed Strategies (and Their Relation to Behavior Strategies)

The key idea that players may randomize between their available choices can be modeled also when one consider games in strategic form.

Definition 4.2. *Given a game Γ , a mixed strategy for player i is a probability distribution β_i over the set of pure strategies S_i available to player i . That is, $\beta_i \in \Delta(S_i)$ and $\beta_i(s_i) \in [0, 1]$ indicates the probability that player i chooses (pure) strategy s_i . Let $\beta \equiv (\beta_i, \beta_{-i}) = (\beta_1, \dots, \beta_n)$ be a mixed strategy profile and use $\Delta(S) \equiv \Delta(S_1 \times \dots \times S_n) = \Delta(S_1) \times \dots \times \Delta(S_n)$ to indicate the set of all available mixed strategy profiles.*

Given an extensive form game Γ , let $\pi_\beta(\cdot)$ be the probability distribution over terminal histories induced by the mixed strategy profile β and, analogously, let $\pi_\rho(\cdot)$ be the probability distribution over terminal histories induced by the behavior strategy profile ρ . Then, we can use the expression of the expected utility to represent the preferences of each player i under the (endogenous) uncertainty associated to the random behavior of the players as

$$U_i(\beta) = \sum_{\sigma \in \tilde{\Sigma}} \pi_\beta(\sigma) v_i(\sigma) \quad \text{and} \quad U_i(\rho) = \sum_{\sigma \in \tilde{\Sigma}} \pi_\rho(\sigma) v_i(\sigma).$$

A key question that arises here is whether we can use mixed and behavior strategies in a fully equivalent way to represent the players' preferences under uncertainty and, therefore, to analyze their behavior. As it turns out, the answer to this question depends crucially on whether or not the players "forget" information that they previously had. To illustrate this point, consider the extensive form game depicted in Figure 4.1. The situation here described could be that of a single investor who must make her investment decisions along two different dimensions, in two sequential stages: first (u)p or (d)own, and then (l)eft or (r)ight. The terminal history that the investor reaches depends on an initial move by chance that determines whether the profitability from the investment is either (B)ad or (G)ood. The investor chooses in the first stage not knowing about the investment true profitability. After this, she chooses again in the second stage. In the second stage, she does not know about the investment profitability either but, more importantly,

she even “forgets” about her own investment choice at the first stage. Now consider a mixed strategy β for the investor such that $\beta(ul) = 1/3$ and $\beta(dr) = 2/3$. Such a strategy, together with the system of priors about the initial move of Nature, determines the distribution over terminal histories π_β with $\pi_\beta(\sigma^1) = p/3$, $\pi_\beta(\sigma^4) = 2p/3$, $\pi_\beta(\sigma^5) = (1-p)/3$, and $\pi_\beta(\sigma^8) = 2(1-p)/3$. Then,

$$U_1(\beta) = \frac{p}{3} [v(\sigma^1) + 2v(\sigma^4)] + \frac{(1-p)}{3} [v(\sigma^5) + 2v(\sigma^8)].$$

For this game, however, there is no behavior strategy capable of inducing such a distribution over terminal histories. To see this, consider any possible behavior strategy for the investor ρ , parameterized as $\rho(u|h) = x \in [0,1]$ and $\rho(l|h') = y \in [0,1]$. With this parameterization, any behavior strategy of the investor will assign probability pxy to the terminal history σ^1 and probability $p(1-x)y$ to the terminal history σ^3 . Therefore, in order to induce the distribution already achieved by the proposed mixed strategy β , we need that $pxy = 1/3$ and $p(1-x)y = 0$ be satisfied simultaneously. Clearly, these are two conflicting requirements, which leads to a contradiction.




Figure 4.1 – Imperfect Recall

The key feature that underlines the message conveyed by this example is that the player “forgets” a piece of information that she previously possessed, i.e., her own past move in the first investment stage. Intuitively, the use of mixed strategies allows the players to select their contingent course of action before playing the game. In this way, they can “insure” themselves

against forgetting information during the play of the game. This can no be done using behavior strategies where the players decide their course of action during the play of the game. Therefore, forgetting can lead to that we cannot use in an equivalent way mixed and behavior strategies to induce distributions over terminal histories. Equivalence between mixed and behavior strategies is, nonetheless, guaranteed when no player forgets any information that she previously had during the play of the game. We say that an extensive form game Γ features *perfect recall* if none of the players forgets any information she possessed in the past, including her previous moves in the game. Then, perfect recall allows us to obtain the following result.

Proposition 4.1 (Kuhn's Theorem). *Consider an extensive form game Γ that features perfect recall. Then, for each mixed strategy profile β there exists a behavior strategy profile ρ that induces the same distribution over terminal histories. Conversely, for each mixed strategy profile β there exists a behavior strategy profile ρ that induces the same distribution over terminal histories.*

Thus, if we believe that perfect recall is a reasonable assumption, then we can exchangeably use behavior and mixed strategies, in a totally equivalent way, to represent the players' preferences under both sources of uncertainty—exogenous and endogenous—present in strategic settings. Given this, when we focus attention to NE—wherein strategic interactions are taken as if the players move simultaneously—, it turns out quite convenient to resort to mixed strategies, rather than to behavior strategies, and use the expression of the payoff $u_i(s)$ induced by an pure strategy profile s to compute player i 's expected utility as

$$U_i(\beta) = \sum_{s \in S} \prod_{i \in N} \beta_i(s_i) u_i(s) = \sum_{O(s) \in \hat{\Sigma}} \pi_\beta(s) v_i(O(s)).$$

Hence, when computing NEs where the players randomize between their actions, it is quite useful to specify the original game in its corresponding strategic form and then use mixed strategies to compute the players' expected payoffs. In the expression above, we consider that $\pi_\beta(\sigma) = \prod_{i \in N} \beta_i(s_i)$, where $\sigma = O(s)$. In other words, the players choose their strategies individually and, more importantly, *in a totally independent way*, as required by the theory. This, nonetheless, does

not preclude that the players' optimal decisions in equilibrium are usually correlated. Recall that we naturally consider that strategies are picked by the players before they play the game (i.e., they are chosen at an *ex-ante* stage), whereas an equilibrium is a final outcome of their interplay (i.e., it is realized at an *ex-post* stage).

4.3. Beliefs and Best Replies

When players are allowed to randomize between actions, they accordingly need to form probabilistic beliefs about other players' choices in order to determine their most preferred courses of action. Let us use $\psi_{-i}(s_{-i})$ to indicate *player i 's beliefs about the combination s_{-i} of pure strategies chosen by the rest of players*. In consonance with the fundamental premise that the players choose their pure strategies in an independent way, game theorists consider that players form beliefs about others' choices also in an independent manner. Therefore, if we let $\psi_{ij}(s_j)$ indicate *player i 's beliefs about the pure strategy s_j chosen by player $j \neq i$* , then $\psi_{-i}(s_{-i}) = \prod_{j \neq i} \psi_{ij}(s_j)$. Given player i 's beliefs ψ_{-i} about the strategies picked by the rest of players, we can then compute i 's expected utility when she chooses the mixed strategy β_i as

$$U_i(\beta_i, \psi_{-i}) = \sum_{s \in S} \beta_i(s_i) \psi_{-i}(s_{-i}) u_i(s_i, s_{-i}).$$

To determine her most preferred course of action upon each string of beliefs ψ_{-i} , player i wishes to choose a mixed strategy β_i so as to maximize $U_i(\beta_i, \psi_{-i})$. *Player i 's best-reply* in a game Γ is the mapping $BR_i : \Delta(S_{-i}) \rightarrow \Delta(S_i)$ specified as

$$\begin{aligned} BR_i(\psi_{-i}) &\equiv \{\beta_i \in \Delta(S_i) : U_i(\beta_i, \psi_{-i}) \geq U_i(\beta'_i, \psi_{-i}) \text{ for each } \beta'_i \in \Delta(S_i)\} \\ &= \arg \max_{\beta_i \in \Delta(S_i)} U_i(\beta_i, \psi_{-i}). \end{aligned} \tag{4.1}$$

Since players now optimally choose probability distributions β_i , the sets $BR_i(\psi_{-i})$ will typically include continuum subsets. In particular, this will be always the case when a player is indifferent between two pure strategies. Thus, the mappings BR_i will in general be *correspondences* rather than functions.

The rationale behind the idea of best-replies continues to work exactly the same way when one restricts attention to situations where players only pick pure strategies and, accordingly, their beliefs assign probability one to particular pure strategies chosen by their opponents.¹ Indeed, the reaction functions introduced in the earlier application to Cournot oligopolistic competition are examples of best-replies where the firms are restricted to choose quantities $q_i \geq 0$ as pure strategies.

In order to analyze behavior when the players make probabilistic choices and inferences, it turns out quite convenient to consider formally profiles of best-replies. Let us use $\psi = ((\psi_{ij})_{j \neq i})_{i \in N}$ to denote a *belief profile* of the players about the pure strategies picked by the rest of players. A *best-reply profile* in a game Γ is the mapping $BR : \Delta(S) \rightarrow \Delta(S)$ given by $BR(\psi) \equiv (BR_1(\psi_{-1}), \dots, BR_n(\psi_{-n}))$.

4.4. (Mixed Strategy) Nash Equilibrium

The logic behind the NE notion can be analogously applied to situations where the players choose probabilistic rules to determine their courses of action. In particular, at an equilibrium, each player i 's beliefs about the (probabilistic) choices of others coincide with their actual mixed strategies, provided that no player has strict incentives to deviate from her choices. Note that including randomized choices does not add any new considerations to the influential NE criterion already presented when introducing NE in pure strategies: all new needed ingredients are of form rather than of substance.

Definition 4.3 (NE in mixed strategies). *A mixed strategy Nash equilibrium (NE) of a game Γ is a mixed strategy profile $\beta^* = (\beta_1^*, \dots, \beta_n^*)$ such that for each player $i \in N$, $\beta_{-i}^* \in BR_i(\psi_{-i}^*)$ where $\psi_{-i}^* = \beta_{-i}^*$.*

As in the case of NE in pure strategies, the key consideration is that in equilibrium we impose the consistency requirement $\psi_{-i}^* = \beta_{-i}^*$, provided that each player i chooses her most preferred mixed strategies upon her beliefs ψ_{-i}^* . Using the concept of best-reply profile, we can then reformulate the

¹Formally, a *best-reply in pure strategies* can be defined as a mapping $BR_i : S_{-i} \rightarrow S_i$ specified by

$$BR_i(s_{-i}) \equiv \{s_i \in S_i : u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}) \text{ for each } s'_i \in S_i\}.$$

Here, $BR_i(s_{-i})$ gives us the set of pure strategies that player i optimally chooses when she believes (with probability one) that the rest of players are picking the strategy string s_{-i} .

NE definition in the following compact form: we say that β^* is a (mixed strategy) NE if $\beta^* \in BR(\beta^*)$.²

Consider the extensive form game depicted in Figure 4.2. The (semi-reduced) normal form representation of this game is given in Figure 4.3.




Figure 4.2 – Computing Mixed Strategy NE

We are using parameters $p, m, q \in [0, 1]$ (with $p + m + q = 1$) to account for the mixed strategies of player 1 and parameter $r \in [0, 1]$ to account for the mixed strategies of player 2. Using this simple notation, we already impose the condition that, in equilibrium, beliefs must coincide with actual (optimal) choices. We can compute the expected payoffs of the players as:

$$U_1(p, m, q; r) = p + (3 - 4r)m + (4 - 5r)q;$$

$$U_2(r; p, m, q) = (p + 2m - q) + (q - 2m)r.$$

Since the players' preferences are represented using expected utility forms, we naturally obtain linear expressions in the parameters p, m, q , for player, 1 and in the parameter r , for player 2. To obtain the best-replies of the players, note first that $(3 - 4r) - (4 - 5r) = r - 1$. Thus, *regardless of the value of $r \in [0, 1]$* , term $(4 - 5r)$ is always higher or equal than term $(3 - 4r)$ (in fact, they are equal only for $r = 1$). As a consequence, player 1 will optimally choose $m^* = 0$ for each $r \in [0, 1]$. This optimal choice will be part of the players' best-replies. Taking this optimal behavior into account, we can restrict attention to the following expression for the expected utilities of the players:

$$U_1(p, q; r) = p + (4 - 5r)q = 1 + (3 - 5r)q = U_1(q; r);$$

$$U_2(r; p, q) = (p - q) + qr = (1 - 2q) + qr = U_2(r; q),$$

since p and q must satisfy $p + q = 1$. Now, to compute the players' best replies, note that, for each

²Formally, β^* is required to be a *fixed point* of the correspondence BR .

given parameter $r \in [0, 1]$, player 1 wishes to pick parameter $q \in [0, 1]$ with to maximize $U_1(q; r)$. Analogously, for each given $q \in [0, 1]$, player 2 wishes to pick parameter r to maximize $U_2(r; q)$. It follows that:

$$BR_1(r) = \begin{cases} q = 0 & \text{if } r \in (3/5, 1]; \\ q \in [0, 1] & \text{if } r = 3/5; \\ q = 1 & \text{if } r \in [0, 3/5) \end{cases}$$

and

$$BR_2(q) = \begin{cases} r \in [0, 1] & \text{if } q = 0; \\ r = 1 & \text{if } q > 0. \end{cases}$$

To obtain the set of mixed strategy NEs, we now need to look for values of q^*, r^* of the parameters such that $q^* \in BR_1(r^*)$ and $r^* \in BR_2(q^*)$ simultaneously. By verifying both compatible and incompatible cases, we obtain:

1. $q^* = 0 \Rightarrow r^* \in [0, 1]$, which is only compatible with $r^* \in [3/5, 1]$.
2. $q^* \in (0, 1) \Rightarrow r^* = 1$, but this is incompatible with $r^* = 3/5$.
3. $q^* = 1 \Rightarrow r^* = 1$, but this is incompatible with $q^* = 0$.

		②	
		(r) A	(1-r) B
(p) a		1, 1	1, 1
① (m) bc		-1, 0	3, 2
(q) bd		-1, 0	4, -1

Figure 4.3 – Computing Mixed Strategy NE: (Semi-Reduced) Normal Form

To sum up, the set of mixed strategy NEs for the proposed game entails that player 1 chooses a with probability one, whereas player 2 chooses any randomized strategy between A and B that puts some weight in $[3/5, 1]$ to strategy A . Formally,

$$NE(\Gamma) = \{(p^*, m^*, q^*; r^*) \in [0, 1]^4 : (p^*, m^*, q^*) = (1, 0, 0), \text{ and } r^* \in [3/5, 1]\}.$$

For this example, we can depict the best-replies of the players, as shown in Figure 4.4.




Figure 4.4 – Best Replies

In particular, the strategy combination (a, A) is the only pure strategy NE of this game. This game has a continuum of mixed strategy NEs, in which player 2 randomizes between A and B . Nonetheless the pure strategy NE (a, A) is more plausible than any other NE according to the subgame perfection criterion. To see this, go back to the extensive form representation of the game. The, notice that at the subgame that initiates upon history (σ^0, b, B) , player 1 has strict incentives to choose d , which gives her a payoff of 4 rather than a payoff of 3. Knowing this, player 2 would never choose action B when call to play upon history (σ^0, b) so as to get a payoff of 0 rather than a payoff of -1. In other words, given that player 2 anticipates the optimal choice of player 1 in the subsequent subgame, she would play action B *with probability zero* (and action A with probability one), thus ruling out the continuum of NEs identified above. If we proceed by backwards induction, then at the root of the original game player 1 would choose action a with probability one, which would yield her a payoff of 1 rather than a payoff of -1.

As another example, consider the extensive form game depicted in Figure 4.10. The strategic form representation of this game can be described by the payoff matrix: If we use $p, q \in [0, 1]$ to parameterize all mixed strategies, we can derive the expected payoffs:

$$U_1(p; q) = (1 - 2q) + (4q - 2)p$$

$$U_2(q; p) = (2p - 1) + (2 - 4p)q.$$




Figure 4.5 – Computing mixed strategy NE (cont.)

		②	
		(q) L	(1 - q) R
①	(p) L	1, -1	-1, 1
	(1 - p) R	-1, 1	1, -1

We obtain the best-reply correspondences:


$$BR_1(q) = \begin{cases} p = 0 & \text{if } q \in [0, 1/2); \\ p \in [0, 1] & \text{if } q = 1/2; \\ p = 1 & \text{if } q \in (1/2, 1] \end{cases}$$

and

$$BR_2(p) = \begin{cases} q = 0 & \text{if } p \in (1/2, 1]; \\ q \in [0, 1] & \text{if } p = 1/2; \\ q = 1 & \text{if } p \in [0, 1/2). \end{cases}$$

By exploring possible compatible and incompatible cases from simultaneously requiring that $p^* \in BR_1(q^*)$ and $q^* \in BR_2(p^*)$, we see that $p^* = q^* = 1/2$ describes the unique NE of this game. There is no pure strategy NE and a unique (symmetric) mixed strategy NE. The best-replies of this game can be plotted as:

In addition to providing a compelling approach to many practical strategic settings, mixed strategy NE is a powerful analytical tool to explore the behavior of individuals since, as [Nash](#)



(1951) demonstrated for the case of finite games, mixed strategy NE always exists.³

Proposition 4.2 (Nash (1951)). *Each finite game Γ has (at least) one NE.*

4.5. Dominated Strategies and *Rationalizability*

Using the assumption that the players share a common understanding of the game they play (including the goals of each player and the logic behind the solution notion), NE is crucially based on the derived implication that each player *knows* the equilibrium behavior of all the other players. Game theorists, though, have developed alternative solution concepts which do not make use of this implication. One approach that has proved particularly appealing is known as *rationalizability*. Rationalizability requires that, rather than having common knowledge of the NE solution notion, the players just share a common understanding of their rationality. In this way, rationalizability asks for less from players, relative to NE, because they are not required to correctly guess others' choices at equilibrium. A caveat of this approach, however, is that rationalizability does not allow one to make precise predictions about the players' behavior for all games. Here is where the NE solution clearly improves over rationalizability. Importantly, rationalizability is not at odds with the NE logic. In particular, the NE requirements imply rationalizability but not the other way around.

³ A game Γ is said to be *finite* if the set of actions $A = \cup_{i \in N} A_i$ of the players is finite. Game theorists have been able to subsequently generalize the existence of NE to games with either countable infinite or continuum action sets. This was a challenge basically of a technical nature, rather than of philosophical substance.

To illustrate the ideas behind rationalizability, consider the strategic form game depicted in Figure 4.6. Suppose that, rather than making any inferences about player 1's plausible course of

		②		
		L	C	R
	U	2,6	0,4	4,3
①	M	3,3	0,0	1,2
	D	1,5	3,6	2,7

Figure 4.6 – Solving a Game with *Rationalizability*. I

action, player 2 considers only her own available strategies. In particular, suppose that player 2 computes her expected payoff under her strategy $\beta_2 = (x, 0, 1 - x)$, with $x \in [0, 1]$, for each given possible strategy β_1 followed by player 1. It follows that

$$U_2(x, U) = 3 + 3x > 4 \Leftrightarrow x > 1/3;$$

$$U_2(x, M) = 2 + x > 0 \text{ for each } x \in [0, 1];$$

$$U_2(x, D) = 7 - 2x > 6 \Leftrightarrow x < 1/2.$$

Therefore, for values of $x \in (1/3, 1/2)$, the mixed strategy β_2 gives player 2 a strictly higher payoff than the pure strategy C, regardless of the strategy β_1 chosen by player 1. If the players' rationality is assumed to be common knowledge, then both players would share a common understanding of such computation. Thus, both players would agree on that player 2 would never choose strategy C. In other words, they would rule out C as a plausible choice and pay attention only to the game in Figure 4.7. We can even apply the same logic a bit further. We observe that strategy U gives

		②	
		L	R
	U	2,6	4,3
①	M	3,3	1,2
	D	1,5	2,7

Figure 4.7 – Solving a Game with *Rationalizability*. II

player 1 a strictly higher payoff than strategy D , for each of the two remaining strategies of player 2. Now, both players would commonly agree on that player 1 will never play strategy D , which takes us to the relevant game shown in Figure 4.8. Now, both players would commonly agree on

		②	
		L	R
①	U	2, 6	4, 3
	M	3, 3	1, 2

Figure 4.8 – Solving a Game with *Rationalizability*. III

that strategy L always gives player 2 a strictly higher payoff than strategy R , which takes to the relevant game in Figure 4.9. Finally, taking this iterative process of reasoning one step further,

		②	
		L	
①	U	2, 6	
	M	3, 3	

Figure 4.9 – Solving a Game with *Rationalizability*. IV

we conclude that the *unique* plausible choice by the players is (M, L) , which yields a payoff of 3 to each players.

The example above illustrates the following notions, which we can use to solve games without making assumptions on how players form beliefs about others' choices.

Definition 4.4. A pure strategy $s_i \in S_i$ of a player i is strictly dominated if there exists a mixed strategy $\beta_i \in \Delta(S_i)$ such that

$$U_i(\beta_i, s_{-i}) > U_i(s_i, s_{-i}) \quad \text{for each } s_{-i} \in S_{-i}.$$

Common knowledge of rationality suggests that the players would commonly agree on neglecting strategies that are strictly dominated, thus ruling them out from consideration. Furthermore, as the example above illustrates, the fact the some strategies of a given player are ruled out can, in some cases, lead to that the strategies of other players become also strictly dominated. This iterative process could, in principle, go on for an arbitrary number of rounds. To be precise, use

$k = 1, 2, \dots$ to indicate the rounds of our iterative process of elimination of strictly dominated strategies. We then start by setting $\widehat{S}_{i,0} \equiv S_i$ for each player $i \in N$, and by rewriting the set S of pure strategies profiles as $S = \widehat{S}_0 \equiv \times_{i \in N} \widehat{S}_{i,0}$. Then, for a given player $i \in N$ and a given round $k \geq 1$, let

$$\widehat{S}_{i,k} \equiv \left\{ s_i \in \widehat{S}_{i,k-1} : \nexists \beta_i \in \Delta(\widehat{S}_{i,k-1}) \text{ s.t. } U_i(\beta_i, s_{-i}) > U_i(s_i, s_{-i}) \quad \forall s_{-i} \in \widehat{S}_{-i,k-1} \right\}$$

be the set of player i 's strategies that "survive" to the k -th round of elimination of strictly dominated strategies. Of course, we have $\widehat{S}_{i,k} \subseteq \widehat{S}_{i,k-1}$ for any $i \in N$. The particular description of the iterative process of elimination of strictly dominated strategies is then as follows. We start at round $k = 1$ by picking arbitrarily a player $i \in N$ and then checking for the set $\widehat{S}_{i,1} \subseteq \widehat{S}_{i,0}$. For the other players $j \neq i$, we then simply set $\widehat{S}_{j,1} \equiv \widehat{S}_{j,0} = S_j$ at this initial round $k = 1$. Then, we can consider the set $\widehat{S}_1 \equiv \times_{j \in N} \widehat{S}_{j,1}$ of pure strategies profiles that are not strictly dominated at round $k = 1$, conditional on the selected player i . Then, we proceed to round $k = 2$, where we pick another player $j \neq i$, and, in a totally analogous way as in the previous round $k = 1$, consider the set $\widehat{S}_{j,2} \subseteq \widehat{S}_{j,1}$. Again, we set $\widehat{S}_{l,2} \equiv \widehat{S}_{l,1}$ for each player $l \neq j$ other than j , and then consider the set $\widehat{S}_2 \equiv \times_{l \in N} \widehat{S}_{l,2}$. We proceed iteratively in this fashion for $k = 1, 2, 3, \dots$. We therefore obtain a sequence

$$\widehat{S}_1 \supseteq \widehat{S}_2 \supseteq \dots \supseteq \widehat{S}_k \supseteq \dots,$$

which is (weakly) decreasing with respect to set inclusion. Of course, if the sets of pure strategies are finite, so that $|S|$ is finite, then this naturally leads to that this iterative process reaches an end. In general, we say that $\widehat{S}(\Gamma)$ is a *rationalizable set of strategy profiles* for a game Γ if there is some $\bar{k} \in \{1, 2, \dots\}$ such that $\widehat{S}_{\bar{k}} = \widehat{S}_{\bar{k}+1} = \widehat{S}_{\bar{k}+2} = \dots = \widehat{S}(\Gamma)$. We can interpret these concepts using the example above. In particular, we started from the set of strategy profiles $\widehat{S}_0 = S = \{U, M, D\} \times \{L, C, R\}$. Then, we picked player $i = 2$ and then obtained $\widehat{S}_{2,1} = \{L, R\}$ so that $\widehat{S}_1 = \{U, M, D\} \times \{L, R\}$. Given this, at round $k = 2$, we were able to obtain $\widehat{S}_{1,2} = \{U, M\}$, for $j = 1$, so that $\widehat{S}_2 = \{U, M\} \times \{L, R\}$. These iterations proceeded in the example to $\widehat{S}_3 = \{U, M\} \times \{L\}$ and, finally, to $\widehat{S}_4 = \{M\} \times \{L\} = \widehat{S}(\Gamma)$.

Although the example above has shown that rationalizability can lead to a unique prediction of how the players behave, this is far from general. Rationalizability is not at odds to the NE criterion. Since rationalizability requires common knowledge of the players' rationality and the NE criterion requires this as well, and adds the assumption of common knowledge of the solution concept, it naturally follows that the set of NEs of any game is always a (weak) subset of the set of rationalizable strategy profiles. To see this, suppose that some strategy profile s^* is a NE of a certain game Γ but does not belong the set of its rationalizable strategy profiles, $s^* \notin \widehat{S}(\Gamma)$. Then, it must be the case that there is some player i , some round $k > 1$, and some strategy $\beta_i \in \Delta(S_{i,k-1})$ such that

$$U_i(\beta_i, s_{-i}) > U_i(s_i^*, s_{-i}) \quad \text{for each } s_{-i} \in \widehat{S}_{-i,k-1}.$$

In particular, it then must be the case that $U_i(\beta_i, s_{-i}^*) > U_i(s_i^*, s_{-i}^*)$ for some $\beta_i \in \Delta(\widehat{S}_{i,k-1})$. Since $\beta_i \in \Delta(\widehat{S}_{i,k-1})$ and $\widehat{S}_{i,k-1} \subseteq S_i$, we know that $\beta_i \in \Delta(S_i)$ as well. These implications, therefore, lead to a contradiction with the NE requirement that $s_i^* \in BR_i(s_{-i}^*)$.

Since the NE requirement considers the same conditions that a set of rationalizable profiles must satisfy (i.e., common knowledge of the game description and of the player's rationality) and then simply adds one extra condition (i.e, common knowledge of how players' beliefs about other decisions are made consistent with their actual choices) to such requirements, the set of NEs of any game is a (weak) subset of the set of rationalizable profiles. In particular, for any game Γ , the following relationships hold: $NE(\Gamma) \subseteq \widehat{S}(\Gamma) \subseteq S$. Although the set of rationalizable profiles generalizes the set of NEs, we cannot use this relationship at this point to make claims about existence of rationalizable sets since this is concept that refers to pure strategies.

4.6. Correlated Equilibrium

In the—perhaps, mainstream—view of game theory presented up to this point, players are assumed to choose their strategies individually and, more importantly, in a *fully independent way*. Modifying this view can, however, be useful to analyze strategic situations in some environments. Consider the strategic form game depicted in Figure 4.10. This game describes a typical environ-

ment commonly known in the social sciences as the *battle of the sexes*. One interpretation of this setting is that the two players are a couple and that they independently, and without any previous communication between them, decide one Friday evening between going to the movies (M) or to the university stadium to watch a basketball game (B). As a happy couple, each of them enjoys being with the other. Yet, player 1 would be (even) happier if they were able to coordinate to attending the basketball game, whereas player 2 (relatively) prefers that they both show up in the movie theatre.

		②	
		M	B
①	M	2,6	0,0
	B	0,0	6,2

Figure 4.10 – Battle of the Sexes. I

It can be verified that this game has 3 NEs, two of them in pure strategies and another one in mixed strategies: (M, M) (or, equivalently, $((1, 0), (1, 0))$), (B, B) (or, equivalently, $((0, 1), (0, 1))$), and $((0.25, 0.75), (0.75, 0.25))$. These NEs yield, respectively, the expected payoffs: $(2, 6)$, $(6, 2)$, and $(3/2, 3/2)$. Suppose, though, that the players can observe the realization of some *public* signal and are able to commit themselves to choose their actions based on the outcome of the signal. Suppose that the weather forecast predicts, with equal chance, that there can either rain (r) or be sunny (s) on that evening Friday. Suppose that the couple (ex-ante, before choosing their strategies) agrees on going to the movies if it is raining and on attending the basketball game if it is sunny. Notice that we in fact are enlarging the description of the original game depicted above. Here, conditional of the other player j following the prescribed strategy $s_j(\sigma^1) = M$ and $s_j(\sigma^2) = B$ it is also optimal for player i to follow such a strategy ($s_i(\sigma^1) = M$ and $s_i(\sigma^2) = B$). Furthermore, the expected payoff to each player of following this recommendation is the linear combination between 6 and 2, according to a 1/2 weight: an expected payoff of 4. More generally, if we were able to choose a random device that yields one public signal with probability p , then any linear combination $p \cdot (6, 2) + (1 - p) \cdot (2, 6)$ could be achieved as the players' expected payoffs. Importantly, notice that the expected payoff of 4 Pareto dominates the expected payoff derived from the only mixed




Figure 4.11 – The Battle of the Sexes with a Public Signal

strategy NE of this game and it is more “just” or “equitable” than the payoffs associated to any of the two pure strategy NEs.

It is interesting that players can be made better off when they are able to commit to public random devices. Furthermore, players could do even better if they were able to commit their actions to private, yet *correlated*, signals. To illustrate this point, consider the modification of the battle of the sexes setting described by figure 4.12. Now, we can verify that this game has again 3

		②	
		M	B
①	M	5,1	0,0
	B	4,4	1,5

Figure 4.12 – Battle of the Sexes. II

NEs, two of them in pure strategies and another one in mixed strategies: (M, M) (or, equivalently, $((1, 0), (1, 0))$), (B, B) (or, equivalently, $((0, 1), (0, 1))$), and $((0.5, 0.5), (0.5, 0.5))$. These NEs yield, respectively, the expected payoffs: $(5, 1)$, $(1, 5)$, and $(5/2, 5/2)$. As indicated earlier, by using a public randomization device with weight $p \in [0, 1]$ on a signal realization that recommends the two players to choose M (and the weight $(1 - p)$ on a signal that prescribes the two players to choose B), the players can achieve any expected payoff combination

$$p \cdot (5, 1) + (1 - p) \cdot (1, 5).$$

More generally, by considering a device that sends signals that give positive probability to recommendations leading to each of the three NEs, any expected payoff in the *convex hull* of the respective NEs payoff combinations⁴ can be achieved.

Now, suppose that the players were able to commit their actions to a device that sends correlated but *privately observed* signals. Suppose, for example, that they agree on hiring some external expert. The expert tosses a fair three-side dice and makes recommendations as follows. The expert (honestly) reveals: to player 1 whether the outcome is 1 or lies in the set $\{2, 3\}$, and to player 2 whether the outcome is 3 or lies in the set $\{1, 2\}$. Now, consider a strategy for player 1 that tells her to play M if she receives the outcome of 1 and to play B if she hears $\{2, 3\}$. Analogously, consider a strategy for player 2 that tells her to play B if she receives the outcome of 3 and to play M if she hears $\{1, 2\}$. Then, notice that if player 1 hears the outcome 1, then she believes that player 2 will choose M because she knows that player 2 will hear $\{1, 2\}$ (because $1 \in \{1, 2\}$). Then, it is optimal for player 1 to choose M . On the other hand, if player 1 hears the outcome $\{2, 3\}$, then she believes that player 2 will choose M with probability $1/2$ (and, thus, B with probability $1/2$) because she knows that player 2 will hear $\{1, 2\}$ with probability $1/2$ and 3 with probability $1/2$. Then, it is optimal for player 1 to randomize between M and B according to any weight (because player 1 is indifferent between the two actions). A totally analogous reasoning shows that it is optimal for player 2 to follow the recommendation prescribed for her, given that player 1 is following her own recommendation. The prescribed prescription profile is, therefore, self-enforcing and, more importantly allows the players to achieve a payoff profile $(10/3, 10/3)$, which is outside the linear combination of the two pure strategy NEs. More generally, it is outside the convex hull of the three possible payoffs achievable under the three original NEs.

These ideas were originally explored by [Aumann \(1974\)](#) and the previous examples motivate the following definition of *correlated equilibrium* proposed by R. J. Aumann.

Definition 4.5. *Given a game Γ , a correlated equilibrium (CE) is a probability distribution $p^* \in \Delta(S)$ over the set of pure strategy profiles such that for each $s_i \in S_i$ chosen with positive probability under p^* (that*

⁴Formally, the convex hull of a set of points $X \subset \mathbb{R}^n$ is the smallest convex set containing X and it is usually denoted as $co(X)$.

is, for each $s_i \in S_i$ such that $\sum_{s_{-i} \in S_{-i}} p^*(s_i, s_{-i}) > 0$, we have⁵

$$\sum_{s_{-i} \in S_{-i}} p^*(s_{-i} | s_i) u_i(s_i, s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} p^*(s_{-i} | s_i) u_i(s'_i, s_{-i})$$

for each $s'_i \in S_i$. We will use $CE(\Gamma)$ to denote the set of CE of game Γ .

The intuition behind the CE concept is as follows. All players are assumed to know (ex-ante, before picking their strategies) that some external “device” or “mediator” will recommend the players to choose some strategy profile s with probability $p^*(s)$. Each player, however, only gets to know (at an interim stage but still before picking her strategy) her own recommendation s_i . Contrary to the NE concept spirit and requirements, under the CE rationale players know that the strategies recommendations are correlated. Therefore, each player i can use the information s_i that she privately receives from the mediator to compute her beliefs—using the standard definition of conditional probability or, in other words, using Bayes’ rule—about the prescriptions s_{-i} received by the rest of the players. Then, given a CE p^* candidate, each player is required to choose the strategy s_i that maximizes her expected utility, given that all other players follow their own CE recommendations s_{-i} and, notably, using the conditional probability $p^*(s_{-i} | s_i)$ to compute her expected payoffs.

We can use the previous example to illustrate how CE can be computed. In the example, we dealt with a CE distribution proposal p^* such that

$$p^*(M | M) = 1, p^*(M | B) = p^*(B | B) = 1/2$$

for both player 1, whereas

$$p^*(B | B) = 1, p^*(M | M) = p^*(B | M) = 1/2$$

for player 2. Notice that conditional beliefs and, therefore, the induced expected payoffs have a

⁵The notation $p^*(s_{-i} | s_i)$ indicated the conditional probability of the string of strategies s_{-i} of all players other than i , given that player i picks strategy s_i .

symmetric structure. Then, we can compute the expected payoffs of player 1 as:

(1) If player 1 plays $s_1 = M$ based on hearing signal realization 1, then

$$\sum_{s_2 \in S_2} p^*(s_2 | s_1) u_1(s_1, s_2) = 1 \cdot [5] > 1 \cdot [4].$$

(2) If player 1 plays $s_1 = M$ based on hearing signal realization $\{2, 3\}$, then

$$\sum_{s_2 \in S_2} p^*(s_2 | s_1) u_1(s_1, s_2) = (1/2) \cdot [5] + (1/2) \cdot [0] = (1/2) \cdot [4] + (1/2) \cdot [1]$$

(3) If player 1 plays $s_1 = B$ based on hearing signal realization $\{2, 3\}$, then

$$\sum_{s_2 \in S_2} p^*(s_2 | s_1) u_1(s_1, s_2) = (1/2) \cdot [4] + (1/2) \cdot [1] = (1/2) \cdot [5] + (1/2) \cdot [0].$$

Therefore, conditional on the correlated recommendations p^* and given that player 2 follows her own prescription, player 1 has no (strict) incentives to deviate from what p^* prescribes for her. Given the symmetric structure of conditional beliefs and induced payoffs, we obtain the analogous implication for player 2. As indicated earlier, the expected payoff of any of the two players under the computed CE can be obtained as:

$$(1/3) \cdot [5] + (2/3) \cdot [5/2] = 10/3.$$

The expected payoffs that the players can achieve in this example are depicted in Figure 4.13.

Notice that the only difference with the NE criterion is that mixed strategies in a CE can be (ex-ante, before the players actually choose their strategies) correlated, whereas NE mixed strategies are assumed to be (ex-ante) independent—though, of course, NE strategies are in general ex-post (after the players pick them) correlated. This suggests that the set of CE generalizes the set of NEs, that is, $NE(\Gamma) \subseteq CE(\Gamma)$ for each game Γ . To see this, consider a NE $\beta^* = (\beta_1^*, \dots, \beta_n^*)$ of a game Γ . Then, propose a probability distribution $p^* \in \Delta(S)$ such that $p^*(s) = \prod_{i \in N} \beta_i^*(s_i)$ for each $s \in S$. Then, we have that $p^*(s_{-i} | s_i) = \prod_{j \neq i} \beta_j^*(s_j)$ so that the CE requirement that each player must




Figure 4.13 – Achievable Payoffs for the Battle of the Sexes II

choose $s_i \in S_i$ so as to maximize

$$\sum_{s_{-i} \in S_{-i}} p^*(s_{-i} | s_i) u_i(s_i, s_{-i})$$

translates into choosing $\beta_i^* \in \Delta(S_i)$ such that if $\beta^*(s_i) > 0$, then s_i must maximize

$$U_i(s_i, \beta_{-i}^*) = \sum_{s_{-i} \in S_{-i}} \prod_{j \neq i} \beta_j^*(s_j) u_i(s_i, s_{-i}).$$

In other words, the CE requirement translates into picking a mixed strategy $\beta_i^* \in BR_i(\beta_{-i}^*)$. Of course, this is ensured since we started by considering that β^* is a NE profile. As an interesting implication, note that since CE is a mixed strategy equilibrium notion, we can use the relationship that we have obtained to guarantee the existence of CE.

Proposition 4.3. *Each NE is a CE and, as a consequence, each finite game Γ has (at least) one CE.*

The notion of CE, and other solutions concepts that have been proposed using the general rationale behind the CE, have recently brought the attention of many game theorists, as well as applied social scientists, interested in exploring the incentives of individuals to behave according to some “social desirable” prescriptions. In this sense, NE perhaps gives us a tool more suitable to

study how strategic individuals actually behave, whereas CE enables us to explore what outcomes are achievable in environments where there are some socially desired outcomes. In short, the NE has been traditionally more linked to positive questions whereas the CE concept has drawn the game theorists' attention relatively more towards normative questions. In particular, the "no-incentive-to-deviate" conditions stated in Definition 4.5, and explored in the previous example, can be viewed as exactly as the type of the incentive compatibility conditions upon which rely the brach of the literature on mechanism design (or contract theory).

5. Bayesian Games

We have thus far assumed that players know other players' preferences over terminal histories or, equivalently, over strategy profiles. This might be seen as a very restrictive assumption in some environments. Formally, we say that a game has *incomplete information* when at least one player does not know (with certainty) the payoff that some other player receives from some strategy profile (or equivalently, terminal history).

Consider again the signaling situation described in Figure 3.3. Notice that we can represent such a situation in a strategic form fashion as follows. Here, after the move of Nature, player




Figure 5.1 – Signaling Game with Simultaneous Moves

S knows that she has either low ($t_S = L$) or high ability ($t_S = H$). Such possible outcomes of uncertainty, which are privately known by a player, are known as her *types*. As already introduced in Chapter 3, a player's type fully describes some payoff-relevant feature, privately known by the player, that results as an outcome of chance. For a game Γ where there is extrinsic uncertainty, let T_i , with typical element $t_i \in T_i$, be the *set of possible types of player i* . In this example, player S has

two possible types $t_S \in T_S = \{L, H\}$. On the other hand, player R does not know player S 's type and, furthermore, there is no feature which might be known privately only by R as a result from the move of Nature. Therefore, player R has here a trivial (singleton) set of types $T_R = \{R\}$, which can be described as stating that the only payoff relevant feature of player R is simply her name. Intuitively, when player S 's type is L , she knows that she is playing according to the right-hand payoff matrix. Also when her type is H , player S knows that she is playing according to the right-hand payoff matrix. Unlike this, player R does not know which payoff matrix describes her actual payoffs. Formally, this can be viewed as a situation of *incomplete information* rather than as a situation with imperfect information because player R 's uncertainty is relative to the players' preferences. Player R can, nonetheless, form beliefs about player S 's types according to the prior μ that describes the initial move of Nature. This allows us to re-interpret this game in fact as one of imperfect information.

Notice that dealing with incomplete information can get really messy because common knowledge of the description of the game would require a player form beliefs about others' preferences, beliefs about what others believe about what the player believes about others' preferences, and so on. Fortunately, as already mentioned, we can circumvent these difficulties by treating each incomplete information game as one of imperfect information, provided that each player does know both the possible payoffs that other players may have and can use a (consistent) probability distribution over such set of possible payoffs. J. C. Harsanyi ([Harsanyi \(1967-68\)](#)) provided an elegant way of making such a transformation from incomplete information games to imperfect information games. Specifically, we consider that the possible payoffs of the players are given by realizations of a random vector (or, equivalently, by a "historical" move of chance) such that each random variable in the vector gives the payoffs of each player. The random variables may be correlated. Then, each player observes the realization of his own random variable, though not necessarily those of the others. As it is the case with the rest of the description of the game, the (joint) probability distribution over the players' payoffs is (commonly) known by everyone. This extended game is known as a *Bayesian game (of incomplete information)* and a NE of such an extended game is known as a Bayes–Nash–Equilibrium. We have already referred to a possible realization

of each random variable as a possible type for the corresponding players. Notice that, using this extended game which asks us to include a move of chance, we have ended up with a game of imperfect information where a player (possibly) cannot distinguished between histories due to the move of Nature.

Definition 5.1. A Bayesian Game Γ consists of $\Gamma = \langle N, S, T, \pi, (u_i)_{i \in N} \rangle$, where:

- (1) $N = \{1, \dots, n\}$ is a finite set of players;
- (2) $S = \times_{i \in N} S_i$ is a set of profiles $s = (s_1, \dots, s_n)$ of (pure) strategies for the players;
- (3) $T = \times_{i \in N} T_i$ is a set of type profiles;
- (4) $\pi \in \Delta(T)$ is a probability distribution over the set of the players' type profiles;
- (5) $u_i : S \times T \rightarrow \mathbb{R}$ is utility function for player i such that $u_i(s, t)$ is the (certain) payoff that accrues to player i when the players choose the strategy profile s and their (actual) type profile is t .

Hence, in Bayesian Game each player is identified by a pair $(i | t_i)$, with $i \in N$ and $t_i \in T_i$, which, in addition to her name i , also tells us the particular payoff-relevant feature t_i that she has. The set S_i of pure strategies of each player i will remain constant when we consider a Bayesian game. However, since a player may adopt now different types which describe different pieces of information that she may have, she can now make her choices contingent of her type. Let $\beta_i(\cdot | t_i) \in \Delta(S_i)$ be a *mixed strategy* for player i , conditional on her type being t_i . The collection of all possible pure strategies for player i (across her set of possible types) is then $\beta_i \equiv \{\beta_i(\cdot | t_i)\}_{t_i \in T_i}$. In a Bayesian game, each player is interested in predicting the types of other players (which, in principle, may be unknown to her). Since $\pi(t_1, \dots, t_n)$ is assumed to be (commonly) known by the players, each player i can use the standard definition of conditional probability (or, equivalently, Bayes' rule) to obtain the conditional probability $\pi_i(t_{-i} | t_i)$. By doing so, player i assigns probabilities to the other players' types (and, therefore, to their payoff-relevant features) in a way *consistent* with the probability π with which Nature selects each type profile t . In general, the players derive the probability $\pi_i(t_{-i} | t_i)$ from the typical definition of conditional probability in modern probability or, equivalently, from Bayes' rule. Using Bayes' rule is indeed the the feature that gives its name to

this approach. In particular, we have

$$\pi_i(t_{-i} | t_i) = \frac{\pi(t_i, t_{-i})}{\sum_{t_{-i} \in T_{-i}} \pi(t_i, t_{-i})} \quad \text{whenever} \quad \sum_{t_{-i} \in T_{-i}} \pi(t_i, t_{-i}) > 0.$$

Then, given that a player identified as $(i | t_i)$ chooses a pure strategy s_i with positive probability (i.e., $\beta_i(s_i | t_i) > 0$), we can compute her expected payoffs, conditional on the rest of the players choosing mixed strategies, as

$$U_i(s_i, \beta_{-i} | t_i) \equiv \sum_{t_{-i} \in T_{-i}} \pi_i(t_{-i} | t_i) \sum_{s_{-i} \in S_{-i}} \prod_{j \neq i} \beta_j(s_j | t_j) u_i(s_i, s_{-i}, t_i, t_{-i}).$$

Notice that the expected payoffs derived above encompass both the intrinsic uncertainty that stems from the fact that other players may play according to mixed strategies and the extrinsic uncertainty due to the “historical” move of Nature that provide the players with their types (that is, their pieces of private information).

Definition 5.2. A Bayes–Nash–Equilibrium (BNE) of a Bayesian game Γ is a profile of collections of mixed strategies β^* such that for each player $i \in N$ and each type $t_i \in T_i$, we have that if $\beta_i^*(s_i | t_i) > 0$, then it must be the case that

$$U_i(s_i, \beta_{-i}^* | t_i) \geq U_i(s'_i, \beta_{-i}^* | t_i)$$

for each $s'_i \in S_i$.

In practice, computing BNE requires us to enlarge the original game and, in particular, to consider the best-reply correspondences of each player $(i | t_i)$. Each player $(i | t_i)$ will face other players $(j | t_j)$ according to the conditional probability $\pi_i(t_{-i} | t_i)$, yet player $(i | t_i)$ will not play against any player $(i | t'_i)$, for $t'_i \neq t_i$. Other than these considerations, finding BNEs is in practice not different from finding NEs. Indeed, a BNE is just NE of an enlarged game, where players are distinguished by their types (which describe their payoff-relevant private information), in addition to their names.

Consider the classical signaling example represented above as a Bayesian game, but suppose

that, instead of working only with the fixed value $\pi(H) = 0.4$, the probability $\pi(H) \in (0, 1)$ is allowed to vary to check for the generality of the messages conveyed by this example. Notice first that, for both type-players $(S | H)$ and $(S | L)$, the pure strategy E is strictly dominated by NE . The three type-players of the enlarged game, R , $(S | H)$, and $(S | L)$, would commonly agree on this, regardless of the particular value of the probability $\pi(H)$. Then, conditional on dropping out the pure strategy E , we can derive player R 's expected payoffs and, accordingly, obtain that R would (strictly) prefer to choose m rather than c if

$$[10]\pi(H) + [0](1 - \pi(H)) = 10\pi(H) > [4]\pi(H) + [4](1 - \pi(H)) = 4.$$

We then obtain that m is strictly preferred by R to c if $\pi(H) > 2/5 = 0.4$. Conversely, R strictly prefers c rather than m if $\pi(H) < 0.4$. Finally, R is indifferent between both options m and c if $\pi(H) = 0.4$, which is precisely the original proposed value for the probability of player S being of high productivity. To summarize, while (NE, NE, m) is the unique BNE of this game for $\pi(H) > 0.4$, (NE, NE, m) is the unique BNE for $\pi(H) < 0.4$. For the originally proposed priors, $\pi(H) = 0.4$, we have that both types of the Sender optimally choose to not get educated whereas the Receiver randomizes, according to any weight $\beta_R^*(m) \in [0, 1]$, between offering a managerial or a clerical position to the Sender.

The above described equilibria would seemingly appear as the most reasonable prediction of how the players behave in this signaling situation. This view, though, would be deceiving as it is fundamentally flawed. To see this, recall that the original description of this situation includes a round of sequential moves, where S moves first and then, given the information provided by S 's educational choice, R makes her own choice. In turn, R 's choice on the kind of position that it offers to S influences S 's educational choice given that S anticipates R 's optimal decisions. These crucial features are of substantial importance and, as it turns out, they are not incorporated into the analysis when we consider the Bayesian game representation of this setting. This is the case because (as in any strategic or normal form representation) the Bayesian representation always considers the situation as if the players move simultaneously instead, thus, ignoring key features of the game description such as the inherent dynamics and what each player knows at each point. More

importantly, in this particular environment, the Bayesian representation neglects what R can learn from S 's choice. Both information transmission and learning turn out to be the key phenomena that signaling games allow us to explore. Specifically, using the Bayesian representation of this game, we cannot resort to sequential rationality considerations. By computing the set of NEs of the enlarged signaling game, we are in fact missing some NEs from the original game. Intuitively, this is the case because the enlarged game *where players are considered to move simultaneously* does not coincide exactly with the original game. Incorporating sequential rationality considerations into the rationale behind the BNE solution concept for dynamic settings is the goal of the *perfect Bayes–Nash–Equilibrium* notion that we shall consider later on in Section 6.1.

The BNE solution concept is, nonetheless, fundamentally suitable to predict the players' behavior when the situation where they are involved is such that they decide simultaneously. In such settings, no relevant strategic considerations—neither on dynamic nor on informational grounds—, which were present in the original game, are left aside from the analysis. To illustrate the analytical power of the Bayesian approach to games with extrinsic uncertainty where players do move simultaneously, consider the following variant of the Cournot duopoly game explored earlier. As before, let $q_i \geq 0$ be the quantity of the good produced and offered by firm $i \in \{1, 2\}$, and let $P = a - b(q_1 + q_2)$ be the (inverse) demand function for the good in this market, with $a, b > 0$. However, suppose now that the cost to firm 1 from producing its quantity is zero and that the marginal cost of firm 2 can be either c_L (low) or c_H (high), with $0 < c_L < c_H < a$. The production cost of firm 2 is its private information, and Nature selects c_L with probability $p \in (0, 1)$. The set of types of firm 2 can therefore be established as $\{L, H\}$, indicating either low ($t = L$) or high ($t = H$) cost. Firm 2 knows its type whereas firm 1 can only assess that, with probability p , firm 2 has low production cost (and that, with probability $1 - p$, firm 2's production cost is high). Let q_2^L and q_2^H denote the quantities chosen by firm 2 when its cost is, respectively, c_L and c_H . The two firms choose their quantities simultaneously. If we consider this situation as a Bayesian game, then the profits to firm 2 ($2 | t$) (that is, conditional on its type being $t = L, H$) are given by the expression

$$\Pi_2(q_1, q_2^t) = [(a - c_t) - b(q_1 + q_2^t)]q_2^t \quad \text{for } t = L, H.$$

Considering the nontrivial case $q_2^{t*} > 0$, firm 2's problems are solved by considering the *first-order conditions* $\partial \Pi_2(q_1, q_2^t)/\partial q_2^t = 0$ (for $t = L, H$), which give us the linear expressions

$$q_2^t = \frac{a - c_t}{2b} - \frac{1}{2}q_1 \quad \text{for } t = L, H.$$

As for firm 1, its expected profits are given by

$$\Pi_1(q_1, q_2^L, q_2^H) = \left(p[a - b(q_1 + q_2^L)] + (1 - p)[a - b(q_1 + q_2^H)] \right) q_1.$$

The *first-order condition* $\partial \Pi_1(q_1, q_2^L, q_2^H)/\partial q_1 = 0$ yields

$$q_1 = \frac{a}{2b} - \frac{p}{2}q_2^L - \frac{1-p}{2}q_2^H$$

We have thus obtained a pair of linear systems with three equations and three unknowns. Yet, each of the systems involved only two equations and two unknowns, where one of the unknowns (q_1) is common to both systems. If we let $a = b = 1$, $p = 1/2$, $c_L = 0$, and $c_H = 1/4$, then we can write these two systems of equations as

$$\begin{cases} q_1 = \frac{1}{2} - \frac{1}{4}q_2^L - \frac{1}{4}q_2^H \\ q_2^L = \frac{1}{2} - \frac{1}{2}q_1 \end{cases}$$

and

$$\begin{cases} q_1 = \frac{1}{2} - \frac{1}{4}q_2^L - \frac{1}{4}q_2^H \\ q_2^H = \frac{3}{8} - \frac{1}{2}q_1 \end{cases}$$

Solving simultaneously the two systems, we obtain a unique solution

$$q_1^* = 3/8, \quad q_2^{L*} = 5/16, \quad \text{and} \quad q_2^{H*} = 3/16,$$

which constitutes the unique (pure strategy) BNE of this Cournot duopoly game with uncertain costs and asymmetric information.

Mixed strategy BNE can be computed in a totally analogous way. Consider a *public good* situation where the mayors of two nearby cities are considering whether or not to build a high-speed railway to improve the commute between their cities. Building the railway costs 20 billions (b) dollars but the railway is only valuable to a city if they build as well a convention center with a station for the train in it. If the convention center is built, then the value of the railway (net from any construction costs) is 30b to the city that uses it alongside with the convention center. The protocol to determine how the railway may be built is as follows. The two majors send independent messages to the state governor. Formally, major $i = 1, 2$ must send a message containing her city valuation $v_i \in \{0, 30\}$ for the railway project. The governor decides then that the project goes ahead and each city pays an equal share of the railway cost if $v_1 = v_2 = 30$. If only one city shows interest, then the project goes ahead and the interested city covers by itself the 20b of the railway project. If none of the majors reports interest for the project, then the railway is not built. In this situation the set of types v_i of each player $i = 1, 2$ consists of her private valuation for the project $T_i = \{0, 30\}$. Suppose that the probability of major i having no interest for the project, $v_i = 0$, is $p \in (0, 1)$. Suppose further that the valuations of the two majors are independent. In this case, we have

$$\pi(v_1, v_2) = \begin{cases} p^2 & \text{if } v_1 = v_2 = 0 \\ (1-p)^2 & \text{if } v_1 = v_2 = 30 \\ p(1-p) & \text{if } v_1 \neq v_2. \end{cases}$$

The set of pure strategies of each player i is simply $S_i = \{0, 30\}$ which describes the possible messages that the major can send to the governor. The utility to player i is then

$$u_i(s_1, s_2, v_1, v_2) = \begin{cases} v_i - 20 & \text{if } s_i > s_j \\ v_i & \text{if } s_i < s_j \\ v_i - 10 & \text{if } s_i = s_j = 30 \\ 0 & \text{if } s_i = s_j = 0. \end{cases}$$

A mixed strategy for player i is $\beta_i = \{\beta_i(\cdot | 0), \beta_i(\cdot | 30)\}$. Let us use $r_i = \beta_i(0 | 0) \in [0, 1]$ and

$q_i = \beta_i(0 | 30) \in [0, 1]$ to parameterize such an strategy. First, note that we must have $r_i^* = 1$ in equilibrium. This is intuitive: a major that does not value the railway will not report to the governor otherwise because she will have to contribute to build it. Algebraically, notice that

$$\begin{aligned} U_i(\beta_i, \beta_j | 0) &= \sum_{v_j \in T_j} \pi(v_j | 0) \sum_{s_i \in S_i} \sum_{s_j \in S_j} \beta_i(s_i | 0) \beta_j(s_j | v_j) u_i(s_i, s_j, v_i, v_j) \\ &= p \left[r_i r_j [0] + r_i (1 - r_j) [0] + (1 - r_i) r_j [-20] + (1 - r_i) (1 - r_j) [-10] \right] \\ &\quad + (1 - p) \left[r_i q_j [0] + r_i (1 - q_j) [0] + (1 - r_i) q_j [-20] + (1 - r_i) (1 - q_j) [-10] \right]. \end{aligned}$$

We observe that it must be the case that $r_i^* = 1$: reporting interest in the project is a (strictly) dominated strategy for any major that does not value the railway whatsoever. This is the case for both $i = 1, 2$. Exploring the incentives for type $v_i = 30$ is more tricky. Now, conditional on making use of the obtained implication that $r_i^* = 1$ for both $i = 1, 2$, we have

$$\begin{aligned} U_i(\beta_i, \beta_j^* | 30) &= p \left[q_i (1) [0] + q_i (1 - 1) [30] + (1 - q_i) (1) [10] + (1 - q_i) (1 - 1) [20] \right] \\ &\quad + (1 - p) \left[q_i q_j^* [0] + q_i (1 - q_j^*) [30] + (1 - q_i) q_j^* [10] + (1 - q_i) (1 - q_j^*) [20] \right]. \end{aligned}$$

This expected utility can be rearranged so as to obtain

$$\begin{aligned} U_i(\beta_i, \beta_j^* | 30) &= \left[10p + 10(1 - p)q_j^* + 20(1 - p)(1 - q_j^*) \right] \\ &\quad + \left[10 - 20p - 20(1 - p)q_j^* \right] q_i. \end{aligned}$$

To compute the set of BNEs is useful here to specify the function $\psi(p) \equiv (1 - 2p)/(2 - 2p)$. We observe that $\psi(p) \in (0, 1/2)$ is (strictly) decreasing in p , with $\psi(0) = 1/2$ and $\psi(1/2) = 0$. Then, the best-reply of player i , conditional on having valuation $v_i = 30$ for the project, is

$$BR_i(q_j | 30) = \begin{cases} q_i^* = 1 & \text{if } q_j < \psi(p); \\ q_i^* \in [0, 1] & \text{if } q_j = \psi(p); \\ q_i^* = 0 & \text{if } q_j > \psi(p). \end{cases}$$

By checking for possible compatibilities and inconsistencies we obtain the following BNEs:

1. For $p \in (0, 1/2)$, we have that $q_i^* = 1$ and $q_j^* = 0$ describes an equilibrium.
2. For $p \in (1/2, 1)$, we have that $q_i^* = q_j^* = 0$ describes an equilibrium.
3. For $p \in (0, 1/2)$, we have that $q_i^* = q_j^* = \psi(p) \in (0, 1/2)$ describes a set of equilibria.

Thus, we obtain multiple BNEs for the game describing this public good provision problem. We have both symmetric and asymmetric equilibria, as well as equilibria in pure strategies and in mixed strategies.

Another important field where BNE is widely applicable because the players' choices are made simultaneously is that of *auctions*. Consider a telecom public institution, or seller, that wishes to sell the rights to operate in a given spectrum to several mobile operators, or players. The public institution wishes to earn as much money as possible from selling the rights. The number of players is n and the valuation of player i for the spectrum is $v_i \in [0, M]$, for some $M > 0$. Furthermore, suppose that the valuation v_i of each operator $i \in N$ is drawn according to a *uniform* distribution from the interval $[0, M]$. In addition to being identically distributed, the random variables $\{v_i\}_{i \in N}$ are independent. The spectrum is offered to the mobile operators through an auction where each operator acts as a *bidder*. If operator i wins the spectrum, then it has to pay x for it so that $v_i - x$ is player i 's payoff. Each operator i knows its own valuation v_i before engaging in the auction but does not know other operators' valuations. The seller of the spectrum does not know the bidders' valuations but it knows that $v_i \in U[0, M]$ for each $i \in N$ and that the valuations v_i are independent. In the auction, the bidders i must simultaneously submit a bid $b_i \geq 0$ for the rights to operate in the spectrum.

Consider first a *second-price auction*. In this auction, the spectrum are awarded to highest bidder but it must pay a price equal to the second highest bid. If more than one bidder submits the highest bid, then the spectrum is randomly assigned (with equal probability) to one of them. Specifically, for a (pure strategy) profile $b = (b_i, b_{-i})$, let $\bar{B}(b) \equiv \max_{i \in N} \{b_i\}$ be the highest bid among the bids in the profile b . Then, under the second-price auction protocol, we have that $x(b) = \max_{j \in N} \{b_j \mid b_j \neq \bar{B}(b)\}$ so that $u_i(b) = v_i - x(b)$ if $b_i \in \arg \max_{i \in N} b_i$ and $u_i(b_i, b_{-i}) = 0$ otherwise. Finding the

BNE of a second-price auction is easy because submitting $b_i^* = v_i$, given that each other bidder is also submitting its actual valuation, is a strategy that gives (1) $u_i(b_i^*, b_{-i}^*) = v_i - x(b^*) > 0$ whenever $v_i = \bar{B}(b^*)$ (that is, whenever bidder i is the one that values most the spectrum rights), or (2) $u_i(b_i^*, b_{-i}^*) = 0$ $v_i \leq \bar{B}(b^*)$. Notice that no bidder has (strict) incentives to deviate from this equilibrium prescription. This BNE is efficient because the spectrum rights are awarded to the operator that values them most. However, the seller is not able to obtain all potential benefits from the trade because the winner bidder pays an amount lower than its actual valuation.

Consider now the case of a *first-price, sealed-bid auction* where the spectrum rights are awarded to highest bidder that, in addition, must pay a price equal to the its own bid. Unlike the case of a second-price auction, for which we have seen that bidders have incentives to reveal their actual valuations for the actioned good, the analysis of first-price auctions is more subtle. First, notice that no player has incentives to submit an amount higher than its valuation because it would obtain a negative payoff in the event that it wins the auction. Furthermore, intuition tells us that each player would wish to bid an amount lower than its actual valuation so that it can obtain positive payoffs in the event that it wins the auction. Assume that, in equilibrium, each player i bids a fraction $\alpha \in [0, 1]$ of its actual valuation so that b_i^* has the form $b_i^* = \alpha v_i$. Then, using the fact that each $v_i \sim U[0, M]$, with the v_i being independent random variables, we know that the probability that a given group of $n - 1$ bidders' valuations are simultaneously below x is $(x/M)^{n-1}$. So, if each other player $j \neq i$ bids according to $b_j^* = \alpha v_j$, then player i 's expected payoff of bidding and amount x is

$$U_i(x, b_i^* | v_i) = (v_i - x) \left(\frac{x}{\alpha M} \right)^{n-1}.$$

Notice that here the fraction $x/\alpha M$ gives us the probability some player j , other than player i , bids less than x , conditional on player j bidding according to αv_j and on $v_j \sim U[0, M]$. The *first-order condition* $\partial U_i(x, b_i^* | v_i)/\partial x = 0$ for player i 's best-reply is then

$$(n - 1)(v_i - x)x^{n-2} - x^{n-1} = 0,$$

which gives us $x^* = v_i(n - 1)/n$. Therefore, the optimal value α^* of parameter α must equal

$(n - 1)/n$. In other words, in the (pure strategy) BNE each operator uses the bidding strategy $b_i^* = (n - 1)v_i/n$. As in the case of a second-price auction, the BNE is efficient because the operator that values most the spectrum rights wins the auction. The winner pays here an amount lower than its actual valuation as well because $(n - 1)/n < 1$. Note that, if n becomes very large (so that n tends to infinity), then equilibrium bids approach their true valuations.

Finally, using the BNE solution concept, [Harsanyi \(1973\)](#) gave a very compelling interpretation to players choosing mixed strategies in a NE of a given game. His approach consists in starting with some original game under certainty and then (artificially) adding—a relatively small amount of—extrinsic uncertainty to such a game. If the original game has mixed strategies NEs, then each of these NEs can be obtained as a BNE of the slightly modified game by considering that the introduced uncertainty (asymptotically) vanishes. [Harsanyi \(1973\)](#)'s intuitive message is that, in general, we can reinterpret a NE that involves randomized strategies as the result of minor perturbations that have been omitted from the description of the original game. In intuitive terms, such perturbations make the players uncertain between different payoff specifications and, therefore, make them “hesitate” between some pure strategy optimal decisions and others.¹

¹[Myerson \(1991\)](#), Chapter 3.10, provides an insightful example to illustrate this point.

6. Equilibrium Refinements

The NE solution concept exists (at least in mixed strategies) in general for any game that we may consider. Many games, however, have multiple NEs and, depending on the particular circumstances described by the game, some of them can be argued to be more “appealing” or “reasonable” than others. Selecting among NEs is of interest because it enhances the power of game theory as a set of analytical tools for the social sciences. Thus, we wish to select, or *refine*, NEs in order to gain accuracy in our understanding and predictions about behavior in strategic settings.

There are typically several approaches that can be pursued to refine NEs, both within extensive and strategic form games. Some approaches are based on considering carefully the detailed description of how players move in dynamic situations (as it is the case, e. g., with the sequential rationality condition and the PBE solution concept) or of the information that the players have at each point. Other refinements are based either on perturbing the strategies optimally chosen by the players, on perturbing the payoffs of the game, or on introducing some amount of uncertainty to the original game.

6.1. Perfect Bayes-Nash Equilibrium

The first refinement of NE that we considered in these notes was that of PBE, where equilibria were selected by requiring that the players followed their optimal choices against others’ decisions *at each subgame* where they were due to play. When information is imperfect and this makes the players unable to distinguish in which history they are when making their decisions, applying sequential rationality can be more tricky. In fact, players do not choose at the root of a subtree when they must take their decisions at non trivial information sets. Perfect Bayes–Nash equilibrium is a

solution concept that, besides refining NE, naturally incorporates both (1) the sequential rationality criterion that we already considered for the PBE solution concept and (2) the way in which players make inferences regarding information they do actually not possess that we already considered for the BNE solution concept.

Given an extensive form game Γ , a *belief for player* $i \neq 0$ at an information set $h \in \mathcal{H}_i$ is a (conditional) probability distribution $\mu_i(\cdot | h) \in \Delta(h)$ so that $\mu_i(\sigma | h)$ is the probability that player i assigns to history $\sigma \in h$ being the true history where she must play. A *set of beliefs* μ_i for player $i \neq 0$ is then a set of probability distributions $\mu_i \equiv \{\mu_i(\cdot | h) | h \in \mathcal{H}_i\}$, one probability distribution for each information set where the player is due to play. A *system of beliefs* μ for the extensive form game Γ is then a profile of sets of beliefs for the players involved in the game, $\mu \equiv (\mu_i)_{i \in N}$. For a history $\sigma \in h \in \mathcal{H}_i$, let $U_i(\rho | \sigma)$ denote the expected payoff to player i when the players randomize their strategies according to the behavior strategy profile ρ , conditional on σ being the true history where player i must play. For a behavior strategy profile ρ , let $\mathbb{P}_\rho(\sigma)$ denote the probability of reaching history σ under the strategy profile ρ . At a practical level, the probabilities $\mathbb{P}_\rho(\sigma)$ are computed using the *total probability rule*, according to the probability weights given by the strategy profile ρ , across the tree representation of the extensive form game.

Definition 6.1. A perfect Bayes-Nash equilibrium (PBNE) of an extensive form game Γ is a profile of behavior strategies ρ^* and a system of beliefs μ^* such that

(1) each player $i \in N$ acts according to the sequential rationality criterion, conditional on the system of beliefs μ^* : for each $h \in \mathcal{H}_i$, we have

$$\sum_{\sigma \in h} \mu_i^*(\sigma | h) U_i(\rho_i^*, \rho_{-i}^* | \sigma) \geq \sum_{\sigma \in h} \mu_i^*(\sigma | h) U_i(\rho_i, \rho_{-i}^* | \sigma) \quad \forall \rho_i.$$

(2) the system of beliefs μ^* is consistent with the strategy profile ρ^* according to Bayes' rule: for each player $i \in N$ and each information set $h \in \mathcal{H}_i$, we have

$$\mu_i^*(\sigma | h) = \frac{\mathbb{P}_{\rho^*}(\sigma)}{\sum_{\sigma' \in h} \mathbb{P}_{\rho^*}(\sigma')},$$

whenever $\sum_{\sigma' \in h} \mathbb{P}_{\rho^*}(\sigma') > 0$.

The beliefs $\mu_i(\cdot | h)$ are commonly known as *posterior beliefs*, in contrast to the prior beliefs $\mu(\cdot | \sigma)$, with $P(\sigma) = 0$, introduced in the description of an extensive form game that describe how Nature moves. Intuitively, for sequential games where there is extrinsic uncertainty, the players begin with some prior beliefs about Also, in sequential games with no extrinsic uncertainty, but where some players are uncertain about previous moves of other players, the players are able to compute their posteriors $\mu_i(\cdot | h)$. The PBNE solution concept then requires that posteriors be obtained (based on the Bayesian machinery) by using “previous” behavior strategies.

Let us go back again to the classical signaling game. Consider the extensive form representation above. It coincides exactly with the game depicted earlier in Figure 3.3, except that now we also include some probabilities $p, q \in [0, 1]$ that describe player R 's posteriors. In this case, we have $p = \mu_R(H | NE)$ and $q = \mu_R(H | E)$. For this game, we will pay attention to two (extreme) categories,




Figure 6.1 – Signaling Game

or classes, of PBNEs that usually appear in general Sender-Receiver games which have either with the form of a signaling situation or a *cheap-talk* situation.¹

¹ Cheap-talk games are settings where one informed player sends a message to another player that must then take an action which is payoff-relevant to both players. Unlike signaling games, cheap-talk games do not incorporate explicit costs to the Sender from sending her messages to the Receiver. Message sending costs and other incentives are implicit in equilibrium. The canonical cheap-talk environment was proposed and analyzed by Crawford and Sobel (1982) who showed, among other relevant implications, that there are severe bounds to the information that can be transmitted in equilibrium in cheap-talk situations when the conflict of interests between the players is relatively high.

First, consider a situation where the Sender chooses her behavior strategy over the signals $\{E, NE\}$ in a way totally independent to the private information that she possesses. In other words, $\rho_S(\cdot | t_S) = \rho_S$ for both types $t_S = L, H$. Then, the posteriors that the Receiver would compute “in a consistent way” given such a behavior strategy of S are

$$p = \frac{0.4\rho_S(NE)}{0.4\rho_S(NE) + 0.6\rho_S(NE)} = 0.4 \quad \text{for } \rho_S(NE) > 0,$$

$$q = \frac{0.4\rho_S(E)}{0.4\rho_S(E) + 0.6\rho_S(E)} = 0.4 \quad \text{for } \rho_S(E) > 0.$$

The Receiver simply retains her priors and, therefore, learn nothing new from observing the level of education chosen by the Sender. The strategy chosen by the Sender is totally non-informative to the Receiver. Let us now explore the optimal behavior of the Sender given such posteriors. Let us use $x = \rho_R(m | NE) \in [0, 1]$ and $y = \rho_R(m | E) \in [0, 1]$ to parameterize R 's behavior strategies. The expected payoffs for the Receiver, conditional on each signal, NE and E , are then given by

$$U_R(\rho | NE) = 0.4 \left[x[10] + (1 - x)[4] \right] + 0.6 \left[x[0] + (1 - x)[4] \right] = 4,$$

$$U_R(\rho | E) = 0.4 \left[y[10] + (1 - y)[4] \right] + 0.6 \left[y[0] + (1 - y)[4] \right] = 4.$$

Interestingly, such payoffs do not depend neither on x nor on y , so that *any* behavior strategy is optimal for the Receiver. Given this, take $x = y = 0$ so that R chooses with probability one position c upon observing any of the two education levels. Finally, it remains to check whether the strategy initially proposed for the Sender is in fact optimal for her. To do this, let us use $w = \rho_S(NE | H) \in [0, 1]$ and $z = \rho_S(NE | L) \in [0, 1]$ to parameterize S 's behavior strategies. The expected payoffs for the Sender, conditional on each of her possible types, H and L , and taken as given the optimal strategy of the Receiver, are

$$U_S(\rho | H) = \left[w[4] + (1 - w)[0] \right] = 4w,$$

$$U_S(\rho | L) = \left[z[4] + (1 - z)[-3] \right] = -3 + 7z.$$

Therefore, R would optimally choose $w = z = 1$, which corresponds to $\rho_S(NE) = 1$. Notice that the

consistent beliefs described by q are left *undetermined*. This is the case because the information set that corresponds to receiving signal E is not reached under the path of the proposed equilibrium. This does not mean that such beliefs do not exist, neither that they are not well-defined, but simply that we need additional considerations in order to assess a particular value to such beliefs. In other words, using the Bayesian requirement present in the PBNE notion, any beliefs $q \in [0, 1]$ would work as consistent beliefs. Then, we need to re-think the optimal behavior of the Receiver in the hypothetical event that she observed E .² The expected payoffs to the Receiver, given any undetermined beliefs $q \in [0, 1]$ are given by

$$U_R(\rho | E) = q \left[y[10] + (1 - y)[4] \right] + (1 - q) \left[y[0] + (1 - y)[4] \right] = 4 + (10q - 4)y.$$

The proposed equilibrium was supported on considering $y = 0$. We see that this is indeed R 's optimal choice if $q \leq 0.4$. Therefore, for values $q \in [0, 0.4]$ of the (undetermined, under the equilibrium path) beliefs q , we have that the proposed behavior profile is a PBNE. Qualitatively, in this class of equilibria the Sender does nothing to provide additional information, or to *separate* her types. As a consequence, the Receiver does not learn anything new further than her priors. This class of equilibria is known as *pooling equilibria*.

Secondly, consider the totally opposed situation where the Sender chooses with probability one a different signal for each different piece of private information that she possesses. Furthermore, suppose that in particular we consider the (perhaps most in intuitive) behavior where a high-ability Sender chooses E , whereas a low-ability Sender chooses NE . In this case, we have $\rho_S(E | H) = \rho_S(NE | L) = 1$. The posteriors that the Receiver would obtain given such a behavior strategy of S are

$$p = \frac{0.4\rho_S(NE | H)}{0.4\rho_S(NE | H) + 0.6\rho_S(NE | L)} = 0,$$

$$q = \frac{0.4\rho_S(E | H)}{0.4\rho_S(E | H) + 0.6\rho_S(E | L)} = 1.$$

Thus, the Receiver does actually learn the true type of the Sender by observing the level of

²The Receiver is in fact never observing E under the proposed strategy of the Sender.

education chosen by the S. The strategy chosen by the Sender is now fully informative to the Receiver. Let us now explore the optimal behavior of the Sender given such posteriors. Again, we use $x = \rho_R(m | NE) \in [0, 1]$ and $y = \rho_R(m | E) \in [0, 1]$ to parameterize R's behavior strategies. The expected payoffs for the Receiver, conditional on each signal, NE and E , are then given by

$$U_R(\rho | NE) = \left[x[0] + (1 - x)[4] \right] = 4 - 4x,$$

$$U_R(\rho | E) = \left[y[10] + (1 - y)[4] \right] = 4 + 6y.$$

R's optimal behavior entails $x = 0$ and $y = 1$: with probability one, she offers the clerical position to an uneducated Sender and the managerial position to an education Sender. Finally, to do study the Sender's optimal behavior, use again $w = \rho_S(NE | H) \in [0, 1]$ and $z = \rho_S(NE | L) \in [0, 1]$ to parameterize S's behavior strategies. The expected payoffs for the Sender, conditional on each of her possible types, H and L , and taken as given the optimal strategy of the Receiver, are now

$$U_S(\rho | H) = \left[w[4] + (1 - w)[6] \right] = 6 - 2w,$$

$$U_S(\rho | L) = \left[z[4] + (1 - z)[3] \right] = 3 + z.$$

Therefore, the Receiver would optimally choose $w = 0$ and $z = 1$, which corresponds exactly with the proposed equilibrium behavior strategy. Qualitatively, in this class of equilibria the Sender does everything she is able to in order to provide additional information, or to *separate* her types. As a consequence, the Receiver does learn the Sender's private pieces of information. This class of equilibria is known as *(fully) separating equilibria*.

As a final insight from this signaling benchmark, recall that, by studying this dynamic situation as a Bayesian game (that is, as if both players moved simultaneously), we obtained that there was a unique (in qualitative terms) class of BNE. In this category of equilibria, the Sender always chose NE while the Receiver picked any randomization between m and c . We then argued that this approach was ill-suited to explore the players' behavior in this sequential interaction. Now, we can fully appreciate that, by transforming the dynamic situation into a Bayesian game, we did lose key NEs of the original extensive form game. The two—quite meaningful—classes of PBNEs that

we have now identified are some of these NEs and they could not be derived as BNE of the so transformed game.

There is a third—intermediate in terms of information transmission—class of equilibria that may arise in Sender-Receiver games, known as *semi-separating equilibria*. In a semi-separating equilibria the Sender typically randomizes between her available messages, the Receiver learns something about the unknown variable, yet does not end up being fully informed. The example represented by the game displayed in Figure 6.2 illustrates this class of equilibria. We can think of this situation as an investment project that requires two partners, players 1 and 2, and requires also two investment stages for partner 2. Think of partner 1 as an entrepreneur and of partner 2 as a funder. Thus, player 1 needs to raise funds from player 2 to fulfill the project, and both players benefit from the investment. The funder may have two different personalities, known privately only by her: benevolent (b) or ordinary (o). The entrepreneur begins with respective priors $1/4$ and $3/4$ about the true personality of the funder. In the first stage, the funder decides whether to provide the initial funds for the project (I) or not (NI). If the project is not initially funded, then the game ends with zero payoffs for both players. If the project is initially funded, then the entrepreneur takes into account the initial willingness of the funder to participate in the project and, accordingly, revises her priors (obtaining the posteriors p and $1 - p$). Then, given the initial investment, the entrepreneur chooses whether to invest her talent and efforts in the project (I) or not (NI). If the entrepreneur chooses not to invest, then the game ends and the (deceived) funder collects her payoffs, depending on her personality, as shown in the figure. If the entrepreneur chooses to invest, then the funder must decide in the final investment stage whether to ultimately comply with the project (C) or to act selfishly (S). If the investor does not comply with the project, her payoffs depend again on her true personality. If the project is fulfilled, then both players get a payoff of two. Notice that in this situation the entrepreneur is interested in assessing the true personality of the funder in order to decide whether or not to put her part in the project. Whether or not the funder is willing to provide the initial funds can be used here to obtain additional information about the true nature of the funder.

We begin by using backwards induction to solve this game and, then, observe that the decisions




Figure 6.2 – Trust Game

marked in red in the figure are optimal if the players are sequentially rational. The key to continue solving this game lies in determining how will player 2 act in information set h . Notice that, given her posteriors $p \in [0, 1]$, player 2 would optimally choose to invest if

$$p[2] + (1 - p)[-2] \geq 0 \Leftrightarrow p \geq 1/2.$$

In particular the entrepreneur would be indifferent between investing and not investing if it were the case that she finally assigns probability 1/2 to the funder being benevolent. As to the optimal behavior of the funder, conditional on each of her two possible personalities, let us use the parameterization $\rho_1(I | o) = x \in [0, 1]$.

Suppose first that the funder chooses “always”—following the frequentist or Laplacian interpretation of probability—to invest when she is an ordinary person: $x = 1$. Then, we obtain the entrepreneur posteriors $p = 3/4$. This is a pooling strategy by the funder—i.e., choosing the same strategy regardless of her private information—that leaves the entrepreneurs with posteriors that coincide with her priors. Since $1/4 < 1/2$, the entrepreneur would then optimally choose not to invest. Given this, notice that, conditional on being ordinary, the funder would then receive a

payoff of -2 , which is lower than the zero payoff that she would obtain instead by switching her initial choice to not invest. Therefore, picking $x = 1$ is not sequentially rational for the funder.

Secondly, suppose that the funder chooses “always” to not invest when she is an ordinary person: $x = 0$. Then, we obtain the entrepreneur posteriors $p = 1$. This is a fully separating strategy by the funder—i.e., choosing different strategies for each possible different piece of her private information—that leaves the entrepreneurs with posteriors that allow her to learn the true personality of the funder. Now, since $1 > 1/2$, the entrepreneur would then optimally choose to invest. Given this, notice though that, conditional on being ordinary, the funder would then receive a zero payoff, which is less than the payoff of 6 that she would obtain instead by switching her initial choice to invest. Therefore, picking $x = 0$ is not sequentially rational for the funder either. The first message that we obtain is that there is no pure strategy PBNE in this situation.

Then, we turn to explore plausible equilibria where the funder (purely) randomizes between investing and not investing when she is ordinary. Let us now use $\rho_2(I|h) = y \in [0, 1]$ to parameterize the entrepreneur decision at information set h . Notice that if the entrepreneur chooses $y = 0$, then it would be optimal for the ordinary funder to pick I with probability one. On the other hand, if the entrepreneur chooses $y = 1$, then it would be optimal for the ordinary funder to pick NI with probability one. In other words, the funder will purely randomize between I and NI only if the entrepreneur herself purely randomizes between I and NI . We must therefore look for equilibria where the entrepreneur chooses $y \in (0, 1)$. As already pointed out, this would be the case if and only if $p = 1/2$. Notice that the ordinary funder’s strategy, as parameterized x , that leads to such posteriors must then satisfy

$$\frac{(1/4) \cdot 1}{(1/4) \cdot 1 + (3/4) \cdot x} = \frac{1}{2} \Rightarrow x = \frac{1}{3}.$$

Now, given this strategy for the funder, and a strategy parameterized by $y \in (0, 1)$ for the entrepreneur, the expected payoffs to the funder, conditional on being ordinary, would be given by

$$U_1(\rho|o) = (2/3)[0] + (1/3)\left[y[6] + (1-y)[-2]\right].$$

Therefore, the ordinary funder would be in the first place indifferent between choosing I and NI if and only if $y = 1/4$. We have thus obtained a (unique) PBNE where $x^* = 1/3$, $y^* = 1/4$, and $p^* = 1/2$. Prior to involving in her interaction with the funder, the entrepreneur believed that only one fourth of the funders are benevolent and, upon observing that the funder invested initially, concluded instead that half of the funders are benevolent. The entrepreneur learned something about the funder's nature, yet she did not end up being fully informed.

6.2. Sequential Equilibrium

While dealing with the signaling environment described by Figure 3.3, we came across a situation where the consistent beliefs of a player were left undetermined. The PBNE notion does not provide any criterion to select among such beliefs, therefore, neither to select among the corresponding optimal behaviors and equilibria. The general issue that we identified by means of that example is that some PBNE proposals may imply that some information sets are not reached along the suggested equilibrium paths. In such cases, using the Bayesian approach to obtain beliefs is not very useful. Motivated by this problem of undetermined posteriors, [Kreps and Wilson \(1982\)](#) proposed a brilliant way of refining further among the set of PBNEs. In intuitive terms, their approach consists of perturbing slightly the proposed equilibrium in a way such that each immediate successor action be achieved with positive probability under the corresponding behavior strategies. We will deal then with a sequence of behavior profiles that approaches to the suggested PBNE behavior strategy. Importantly, every element of the sequence is such that no beliefs are left undetermined by using Bayes' rule (because all information sets are reached with some positive probability in each of the perturbed profiles in the sequence). Then, we work with the corresponding sequence of consistent beliefs. In particular, we compute the limiting beliefs of such a sequence. Finally, for such limiting beliefs and the corresponding limiting behavior profile (which was proposed at the beginning of the process), we need to check for the sequential rationality of the players. An important feature of [Kreps and Wilson \(1982\)](#)'s approach is that we need to check for the optimal behavior of the players only at the limiting behavior profile and its associated beliefs, yet *not at any of the behavior profile—belief pairs along the sequence*. Specifically, a *fully mixed behavior strategy profile* is a profile ρ

such that $\rho_i(a | h_i) > 0$ for each immediate successor $a \in A(h_i)$, for each information set $h_i \in \mathcal{H}_i$, and for each player $i \in N$. The equilibrium refinement proposed by [Kreps and Wilson \(1982\)](#) is known as *sequential equilibrium*.

Definition 6.2. A sequential equilibrium (SE) of an extensive form game Γ is a profile of behavior strategies ρ^* and a system of beliefs μ^* that constitutes a PBNE of such a game and, in addition, such that: (1) there is a sequence of fully mixed behavior strategy profiles $\{\rho_k\}$ which converges to ρ^* , and (2) the sequence of beliefs $\{\mu_k\}$, where each μ_k is obtained from each ρ_k by applying Bayes' rule, converges to μ^* .

To see how the computation of SE works in practice consider the extensive form game depicted in Figure 6.3. Consider first the (pure) strategy profile ρ where the players pick (B, U, L) . Because




Figure 6.3 – Computing SE

the information set h is not reached under such a strategy combination, any belief $p \in [0, 1]$ satisfies the consistency requirement imposed by the PBNE solution concept. To check whether player 3 is being sequentially rational at h , we compute her expected payoffs at h by using $\rho_3(L | h) = x \in [0, 1]$ to parameterize an arbitrary behavior strategy for her. We obtain

$$\begin{aligned} U_3(\rho | h) &= p \left[x[1] + (1 - x)[3] \right] + (1 - p) \left[x[2] + (1 - x)[1] \right] \\ &= (1 - 2p) + (1 - 3p)x. \end{aligned}$$

Therefore, choosing L with probability one ($x = 1$) is optimal for player 3 if she assigns sufficiently low probability, $p \in [0, 1/3]$, to player 2 having chosen action U . Conditional on player 3 playing L

is then optimal for player 2 to pick U since this gives her a payoff of 2 which exceeds the payoff of 1 that she would obtain by deviating to D . As for player 1, conditional on players 2 and 3 choosing (U, L) , then it is optimal for her to choose B . By doing so, player 1 obtains a payoff of 2 which exceeds the payoff of 1 that she would obtain by deviating to T . Therefore, for any induced belief $p \in [0, 1]$ of player 3, (B, U, L) is a PBNE of this game.

Nonetheless, there exists another (pure) strategy profile ρ' that constitutes a PBNE of this game as well. Suppose now that the players choose (T, U, R) instead, as indicated in blue in the figure. Now $p = 1$ gives us the unique belief that satisfies the consistency requirement of the PBNE notion. For such a posterior, it follows from the above derived expression for player 3's expected payoff that choosing R with probability one ($x = 0$) is her unique sequentially rational decision. Conditional on this, it is now easy to check that players 2 and 3, respectively, have strict incentives to play U and T , as prescribed by the proposed strategy profile. Furthermore, we can verify that this strategy profile (T, U, R) is indeed the unique SE of this extensive form game. To see this, let us parameterize the behavior strategies of players 1 and 2, respectively, as $\rho_1(T | \sigma^0) = y \in [0, 1]$ and $\rho_2(U | (\sigma^0, T)) = z \in [0, 1]$. Then, the consistent posteriors of player 3 are obtained as

$$p = \frac{yz}{yz + y(1 - z)} = z$$

as long as $y > 0$. As required by the definition of SE, we now need to consider sequences $\{y_k\}$ and $\{z_k\}$ of fully mixed behavior strategies that converge, respectively, to $y = 1$ and $z = 1$. Note that to work with fully mixed behavior strategies, we need to consider $y_k > 0$ and $z_k > 0$ for each term k along the sequence. Then, it is clear that the corresponding sequence of posteriors $\{p_k\}$ for player 3 satisfies $p_k = z_k$. Notice that any sequences $\{y_k\} \rightarrow 1$ and $\{z_k\} \rightarrow 1$, that satisfy $y_k, z_k > 0$ for each k , would be suitable to obtain the desired sequence of posteriors $\{p_k\} \rightarrow 1$.

On the other hand, we observe that there is no way of proposing sequences of behavior strategies $\{y_k\} \rightarrow 0$ and $\{z_k\} \rightarrow 1$, with $y_k, z_k > 0$ for each k , such that $\{p_k\} \equiv \{z_k\}$ converges to a belief lower than $1/3$. Therefore, while (B, U, L) is a PBNE of this extensive form game, it does not constitute a SE. In this example the SE criterion has allowed us to narrow down the set $\{(B, U, L), (T, U, R)\}$ of plausible NEs, and thereby to conclude that the most appealing prediction here is given by

(T, U, R) .

6.3. The “Global Games” Approach

The *global games approach* can be seen as a particular equilibrium selection criterion through “perturbations.” In particular, the proposed perturbations here affect what the players privately know about some unknown “historical” move by Nature. As already described, [Harsanyi \(1967-68\)](#) provided a unified framework to analyze situations where the players’ payoffs depend on the realization of some move of Nature, typically termed also as *unknown fundamental* or *state of the world*. In such classical Bayesian Games environments, the players’ optimal choices not only depend on their beliefs of the state of the world but also on their beliefs of others’ beliefs about the state of world, of their beliefs about others’ beliefs about everyone else’s beliefs about the state of the world, and so on. Such higher-order beliefs are, however, very complex to compute by the players and by the analyst herself that explores such strategic settings with extrinsic uncertainty. Therefore, it would be interesting to identify environments where arbitrarily higher-order beliefs about some unknown parameter play an important role in the individuals’ decisions, yet that allow for a systematic and tractable analysis. The global games approach, first proposed by [Carlsson and van Damme \(1993\)](#) enables such a desirable treatment. Consider a game with extrinsic uncertainty where the move of Nature is captured by a state of the world θ . For this Bayesian game, the global games approach consists of assuming that the private information, signal, or type, that each player receives about the state θ has some “small” amount of noise. The joint distribution of the noises that affect the players’ signals—or noise technology—is assumed to be commonly known among the players. The trick that allows for a tractable approach to how players compute arbitrarily higher-order beliefs about θ is considering that they are smart Laplacian statisticians that correctly interpreted probabilities following their frequentist interpretation. The global games approach is appealing when one considers that there is a very large number of (ex-ante identical) players involved in the game of interest. In such cases, when hypothesizing the higher-order beliefs of the players about θ , each player only needs to draw inferences about the proportion of other players that are choosing each particular choice among the available choices. This approach can be seen,

therefore, as a “short-cut” that allows the game theorist for a systematic approach to situations that may otherwise appear to be intractable. At a more fundamental level, the classical assumptions underlying the NE notion that (1) the players have common knowledge about the entire situation in which they interact and (2) that beliefs about other players’ decisions are consistent with their actual optimal choices invite for the presence of multiple equilibria. Through the heuristic device incorporated into the analysis by the Laplacian view of the the players, the global games approach allows to identify which set of self-fulfilling beliefs will prevail in equilibrium and, therefore, to select among the set of NEs in a practical way.

To illustrate the rationale behind the global games approach, consider the game depicted in Figure 6.4, with $\theta \in \mathbb{R}$, which was originally proposed by Carlsson and van Damme (1993). In

		②	
		<i>I</i>	<i>NI</i>
①	<i>I</i>	θ, θ	$\theta - 1, 0$
	<i>NI</i>	$0, \theta - 1$	$0, 0$

Figure 6.4 – Investment Game under Extrinsic Uncertainty

this situation the two players must decide whether to Invest (*I*) or Not Invest (*NI*) in a risky project. We observe that the investment activity entails complementarities, that is, the return to each player from investing is higher if the other player also invests.

Suppose first that there is in fact no extrinsic uncertainty and the players have complete information about the state of the world $\theta \in \mathbb{R}$. In this case, we can distinguish three possible outcomes:

1. If $\theta > 1$, then *I* is a (strictly) dominant strategy for each player.
2. If $\theta \in [0, 1]$, then there are two pure strategy NEs: either both players choose *I* or both players choose *NI*.
3. If $\theta < 0$, then *NI* is a (strictly) dominant strategy for each player.

The most interesting situation, however, is the Bayesian game that we obtain when the players are instead uncertain about the value of θ . Consider a situation where θ is randomly drawn

from the real line. That is, $\theta \sim U[\underline{\theta}, \bar{\theta}]$, where $\underline{\theta} \rightarrow -\infty$ and $\bar{\theta} \rightarrow +\infty$.³ Suppose that each player i observes a private signal (or type) $t_i = \theta + \varepsilon_i$. Suppose further that each noise term ε_i is independently normally distributed with mean zero and standard deviation σ : $\varepsilon_i \sim N(0, \sigma^2)$ for each $i = 1, 2$. Given such distributional assumptions on θ and each ε_i , we know that the (conditional on one's own signal or type) beliefs about the state and about the other player's signal realization (or type) are given, respectively, as $(\theta | t_i) \sim N(t_i | \sigma^2)$ and $(t_j | t_i) \sim N(t_i | 2\sigma^2)$. This conditional distributions describe the beliefs of any player i about, respectively, the state of the world and the private information that the other player has.

Let consider now the following (pure) strategy for each player i : for some real number, or "cutoff" point k , take $\beta_i(I | t_i) = 1$ if $t_i > k$ and $\beta_i(NI | t_i) = 1$ if $t_i \leq k$. Thus, this strategy prescribes the player to Invest (with probability one) if the value of her type exceeds a certain value k and Not Invest (again, with probability one) if the value of her type does not exceed the value k . This kind of "switching" strategy seems quite natural in this context: the player chooses to invest whenever she receives a signal pointing out that the value of the state is sufficiently high, which turns out very intuitive given the payoffs of the players. Suppose now that each player thinks that her opponent is in fact following this "switching" strategy around the cutoff point k . Let $F(\cdot)$ denote the (cumulative) distribution function of the standard normal distribution. Then, notice that each player i will assign probability $F\left(1/\sqrt{2}\sigma(k - t_j)\right)$ to her opponent j observing a signal t_j less than k . In particular, given the cutoff point $k = t_j$ for the proposed strategies, each player i will assign probability $1/2$ to her opponent investing. As a consequence, there will be an equilibrium where both players follow the suggested switching strategies with cutoff value $k^* = 1/2$. Furthermore, some analytical arguments—perhaps beyond the scope of these notes—can be used to show that this is in fact the only strategy that survives the iterated elimination of strictly dominated strategies, at the stage where is player knows her own type.

Consider now the following many-player analog of this game. Suppose that a continuum of players is deciding now whether to invest. The payoff from not investing is zero. The payoff from investing to each player is $\theta - 1 + \rho$, where ρ is the proportion of other players that choose

³This is a non-standard distributional assumption that nevertheless poses no technical difficulties.

to invest. The information structure is exactly the same as in the two-player version. For this version of the investment game, switching strategies with cutoff value $k^* = 1/2$ are also the only equilibrium that survives iteration elimination of strictly dominated strategies. Importantly, to achieve this equilibrium, each player i needs to keep track of an infinitely iterated layer of beliefs about whether each of the other players j is observing a signal t_j higher than $k^* = 1/2$. Nonetheless, the beautiful trick that underlies the global games approach allows us to circumvent, in a fully consistent and systematic way, such a complex higher-order belief computation process. In particular, the suggested short-cut requires instead each player i to make the following, far less demanding, computations:

1. Estimate θ by using her signal realization t_i .
2. Postulate that ρ is uniformly distributed on the interval $[0, 1]$.
3. Take her optimal action.

Recall that the estimation that the player makes about θ given her type is just the type itself: $(\theta | t_i) \sim N(t_i | \sigma^2)$ so that $E[\theta | t_i] = t_i$. Therefore, if we consider that $\rho \sim U[0, 1]$, then the expected payoff from investing is $t_i - 1/2$. On the other hand, the expected payoff from not investing is zero. Thus, following the procedure above, player i would invest if $t_i \geq 1/2$, which coincides with the equilibrium obtained earlier. Postulating that ρ is uniformly distributed captures the idea that the player has not any precise information that leads her to hold some specific beliefs about the proportion of other players that invest. Then, the player considers that the such a proportion is completely random, which represent a type of “diffuse” or “agnostic” beliefs on others’ actions.⁴ For such “uninformed inferences,” the player then is assumed to act as a Laplacian statistician that considers that at least a proportion ρ of other players is investing whenever she receives a signal $t_i \geq \rho$. The messages conveyed by the example presented here can be extended generally to a wide class of binary choice games.⁵

⁴More formally, the fact that taking a uniform distribution for ρ is an appropriate choice is fundamentally linked to the implication that the transformation of any random variable according to its (cumulative) distribution function is uniformly distributed.

⁵Morris and Shin (2003) provide the required arguments in their comprehensive view of the theory and applications of the global games approach.

7. Repeated Games

In many environments, some interactions do not occur just once and forever but players face them repeatedly: firms, couples, professionals, institutions, or countries are involved among them in the same strategic settings over time, where they receive the same possible outcomes and have the same preferences over such outcomes in each period. The strand of the literature on *repeated games* aims at exploring situations where individuals interact repeatedly, by making crucial assumptions on who interacts with whom and on the information that permeates—either privately or publicly—from each round of interaction to the subsequent ones. By analyzing formally repeated interactions, important issues of dynamic behavior can be explored, such as the relationship between reputation and long-term stable outcomes, career building concerns, the design of institutions, or the rise of stable social norms. The term “reputation” is commonly used to describe how a player’s past behavior affects what others believe about her and, therefore, how others behave in the future with respect to such a player. The theory allows us to interpret long-term stable interactions as “social norms,” which capture how individuals reward or punish others that do not follow the behavior prescribed by the norms. The literature on repeated games typically takes on two approaches: (1) exploring the structure of the equilibrium strategies, and (2) describing the set of payoffs that can be sustained in equilibria. The first approach allows us to analyze explicitly how incentives affect reputation, shapes social norms, and determines long-term behavior. The second approach tells us about the welfare levels that can be ensured to the players in repeated interactions. At an instrumental level, SPE and PBNE/SE are the solution concepts typically used to predict behavior in repeated games.

A *repeated game* is a strategic situation where a particular—usually, simultaneous or “one-shot”—game is played several times under the assumption that the primitives of such a game

do not change over time. Formally, a repeated game is a sequence of repetitions of some *stage game*. For simplicity, we will consider that a *stage game* is a strategic form game with no extrinsic uncertainty $\Gamma = \langle N, A, (v_i)_{i \in N} \rangle$, where $A = \times_{i \in N} A_i$, and $v_i : A \rightarrow \mathbb{R}$ for each $i \in N$.¹ We consider that the stage game Γ is played repeatedly over a discrete sequence of periods $t = 0, 1, 2, \dots, T$, where T could tend to infinite. Given this, we can then capture this repeated situation by means of a “multiple-shot” extensive form game Γ_R , which gives us the repeated game associated with the stage game Γ .

For each period $t = 0, 1, 2, \dots, T$, we will use $a^t = (a_1^t, \dots, a_n^t) \in A$ to denote an *action profile chosen by the players at period t* . In these notes, we restrict attention to repeated games with *perfect monitoring*: at the end of each period t , all players observe (or “learn”) the action profile chosen a^t . In other words, the actions of each player are perfectly monitored by all other players. The set of histories Σ of the repeated game Γ_R can then be specified as follows.

1. Use σ^0 to denote the initial history or root of the tree that gives the extensive form game Γ_R . Then, at $t=0$, the players choose an action profile $a^0 \in A$.
2. At $t = 1$, for each history $\sigma^1 = (\sigma^0, a^0) \in A$, the players choose an action profile $a^1 \in A$.
3. In general, at each $t = 0, 1, 2, \dots, T$, upon each history $\sigma^t = (\sigma^0, a^0, \dots, a^{t-1}) \in A^t$, the players choose an action profile $a^t \in A$.

Hence, for $t > 0$, a t -*period history* of the extensive form game Γ_R that describes the repeated game (associated to our stage game Γ) is a history with the form $\sigma^t = (\sigma^0, a^0, a^1, \dots, a^{t-1}) \in \Sigma^t \equiv A^t$, where Σ^t is the *set of possible histories at t* . For $t = 0$, we simply set $\Sigma^0 \equiv \{\sigma^0\}$. The tree associated to the extensive form game is therefore formally given as $\Sigma \equiv \bigcup_{t=0}^T \Sigma^t$. A repeated game Γ_R is *finite* if T is finite and *infinite* if $T \rightarrow \infty$.

A *pure strategy for player i in the repeated game Γ_R* is then a mapping $s_i : \Sigma \rightarrow A_i$. Also, a *behavior strategy for player i in the repeated game Γ_R* is a mapping $\rho_i : \Sigma \rightarrow \Delta(A_i)$. As usual, we will use $s = (s_1, \dots, s_n)$ and $\rho = (\rho_1, \dots, \rho_n)$ to denote *strategy profiles* for the repeated game Γ_R .

¹ When considering repetitions of a stage game, it is notationally quite convenient to identify each pure strategy of the stage game with its corresponding action $a_i \in A_i$. Other than the change of labels from “ s ”s to “ a ”s, this does not actually imply neither a conceptual change nor even a change in notation since actions do coincide with pure strategies in strategic form games.

In dynamic situations, it is natural to assume that individuals do not regard equally a certain payoff obtained in the present than the same payoff received in a future period. In particular, the theory usually considers that players have some degree of “impatience” and that they discount future payoffs following an exponential pattern. In particular, we will consider a discount factor $\delta \in [0, 1]$ and then

$$u_i(s) = \sum_{t=0}^T \delta^t v_i(a^t)$$

will account for the aggregate payoffs that the player receives when $T < \infty$. For $T \rightarrow \infty$, the aggregate payoffs are instead typically normalized as the average discounted payoff²

$$u_i(s) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t v_i(a^t).$$

To fix ideas about the concepts introduced thus far, consider that two players interact according to a prisoner’s dilemma stage game whose payoffs are depicted in Figure 7.1. Suppose that the

		②	
		C	D
①	C	3, 3	0, 4
	D	4, 0	1, 1

Figure 7.1 – Prisoner’s Dilemma Stage Game

players repeat their interaction twice, $t = 0, 1$, and that they are not impatient, $\delta = 1$. For instance, if they play (C, C) in the first period and (C, D) in the second period, then they obtain a payoff profile $(3 + 0, 3 + 4) = (3, 7)$. Then, we can derive the extensive form game Γ_R is this two-period repeated game with payoffs as shown in Figure 7.2. The red diagram includes the players’ moves at $t = 0$ whereas the blue diagram includes the players’ moves at $t = 1$.

Furthermore, we can plot the set of attainable payoffs for this two-period repeated situation where $\delta = 1$ as shown in Figure 7.3.

Some additional—more technical—elements are sometimes useful to explore equilibria in repeated games. Given a history σ^t , with $t < T$, the *continuation game from σ^t* is the repeated game

²The specific reason for using the average discounted payoff for infinite horizon repeated games is that if the player obtains a constant payoff across periods, that is, $v_i(a^t) = \bar{v}$ for each $t = 0, 1, \dots$, then $u_i(s) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \bar{v} = \bar{v}$.




Figure 7.2 – Two-Period Repeated Prisoner's Dilemma Game

$\Gamma_R|_{\sigma^t}$ that begins following history σ^t . Given a strategy profile $s \in S \equiv A^T$ of the repeated game, the continuation strategy induced by history σ^t , for $t < T$, is the strategy $s_i|_{\sigma^t}(\sigma^\tau)$ specified as

$$s_i|_{\sigma^t}(\sigma^\tau) \equiv s_i(\sigma^t \sigma^\tau),$$

where $\sigma^t \sigma^\tau$ stands as shorthand notation for the string (σ^t, σ^τ) , which is simply the concatenation of history σ^t followed by the sequence of action profiles $\sigma^\tau = (a^{t+1}, a^{t+2}, \dots, a^T)$. Then, let us use $s|_{\sigma^t} \equiv (s_1|_{\sigma^t}, \dots, s_n|_{\sigma^t})$ to denote a continuation strategy profile induced by history σ^t . For $t < T$, we will analogously use the notation

$$\rho_i|_{\sigma^t}(\sigma^\tau) \equiv \rho_i(\sigma^t \sigma^\tau),$$

to indicate the continuation behavior strategy induced by history σ^t and $\rho|_{\sigma^t} \equiv (\rho_1|_{\sigma^t}, \dots, \rho_n|_{\sigma^t})$ to denote a behavior continuation strategy profile induced by history σ^t . Intuitively, $s_i|_{\sigma^t}$ and $\rho_i|_{\sigma^t}$ are just the restrictions of strategies s_i and ρ_i , respectively, to the subgames that follow history σ^t . To




Figure 7.3 – Attainable Payoffs for the Two-Period Prisoner’s Dilemma Repeated Game

illustrate how continuation strategies work, consider again the two-period repeated game depicted in Figure 7.3. Suppose that player 1 chooses a (pure) strategy s_1 specified as:

$$s_1(\sigma^0) = C; s_1(\sigma^0, CC) = C; s_1(\sigma^0, CD) = D; s_1(\sigma^0, DC) = D; s_1(\sigma^0, DD) = C.$$

Such a strategy is indicated in blue in Figure 7.4. Also, suppose that player 2 chooses a (pure) strategy s_2 specified as:

$$s_2(\sigma^0) = D; s_2(\sigma^0, CC) = D; s_2(\sigma^0, CD) = D; s_2(\sigma^0, DC) = C; s_2(\sigma^0, DD) = C.$$

Such a strategy is indicated in red in Figure 7.4. Notice that the proposed strategy combination $s = (s_1, s_2)$ gives us an “action path,” which is depicted by arrows in the figure. The final payoffs achieved under such a strategy profile are (1, 5). Let us now fix the particular history $\bar{\sigma}^1 = (\sigma^0, CD) \in \Sigma^1$. Then, for the suggested strategy profile we have the continuation strategies³ $s_1|_{\bar{\sigma}^1}(\bar{\sigma}^1) = D$ and $s_2|_{\bar{\sigma}^1}(\bar{\sigma}^1) = D$, so that $s|_{\bar{\sigma}^1}(\bar{\sigma}^1) = DD$. On the other hand, if we take another particular history such as $\tilde{\sigma}^1 = (\sigma^0, DD) \in \Sigma^1$, then we have the continuation strategies $s_1|_{\tilde{\sigma}^1}(\tilde{\sigma}^1) = C$ and $s_2|_{\tilde{\sigma}^1}(\tilde{\sigma}^1) = C$, so that $s|_{\tilde{\sigma}^1}(\tilde{\sigma}^1) = CC$.

Because of the description of how players move sequentially in a repeated game with perfect

³ Notice that, in this case, any concatenation (σ^t, σ^τ) specified in the definition of a continuation strategy must have the form σ^t since any $\sigma^1 \in \Sigma^1$ is the last history at which the players make decisions in this example where $T = 1$.




Figure 7.4 – Two-Period Repeated Prisoner’s Dilemma Game: Checking for Equilibria

monitoring, rather than using NE, the most appealing solution notion here is SPE.

Definition 7.1. A behavior strategy profile ρ^* is a SPE of a repeated game Γ_R , with perfect monitoring, if for each history $\sigma^t \in \Sigma$, the continuation strategy profile $\rho^*|_{\sigma^t}$ is a NE of the repeated game.

The definition above simply says that the a SPE is a strategy profile that induces a NE in each subgame of the extensive form game associated to the repeated game. This coincides with our previous general definition of SPE.

Although it gives us a very appealing solution notion, computing SPE by following directly its definition raises potentially formidable technical difficulties. Notice that checking for subgame perfection entails in principle verifying whether a very large number—in fact, an infinite number for infinite repeated games—of strategy profiles are NEs. Added to this, checking whether a profile $\rho|_{\sigma^t}$ is a NE involves checking that each player i ’s strategy $\rho_i|_{\sigma^t}$ makes her no worse off than under a very large—even infinite—number of potential deviations. To illustrate these difficulties, let us

go back to the example in Figure 7.4 and to the particular strategy profile suggested in the picture. If we were to follow the definition of SPE to check whether or not the proposed strategy profile is a SPE, then we would have to verify separately whether there are profitable deviations for both player 1 and player 2 at the five histories $\sigma^0, \sigma_a^1, \sigma_b^1, \sigma_c^1,$ and σ_d^1 . This total of 10 verifications is what it would take us to check for equilibria in a two-player, two-period, two-action repeated game, which in fact leave us with the simplest example we might consider.

This potentially unmanageable task in general repeated games can be simplified enormously by applying a principle borrowed from dynamic programming which is commonly known as the *one-shot deviation principle*. In particular, the *one-shot deviation principle* works as follows. First, fix a strategy profile ρ , which we wish to check whether it constitutes an equilibrium. Then, we define a *one-shot profitable deviation for player i* —with respect to ρ_i —as a strategy $\rho'_i \neq \rho_i$ such that the following two conditions are satisfied:

1. There is a *single* history $\sigma^{\tilde{t}} \in \Sigma$ (that is, for some $\sigma^{\tilde{t}} \in \Sigma^{\tilde{t}}$, for some $\tilde{t} \in \{0, 1, \dots, T\}$) such that for each $\sigma^t \neq \sigma^{\tilde{t}}$, we have $\rho'_i(\sigma^t) = \rho_i(\sigma^t)$.
2. Given the continuation game $\Gamma_R|_{\sigma^{\tilde{t}}}$, we have $u_i(\rho'_i|_{\sigma^{\tilde{t}}}, \rho_{-i}|_{\sigma^{\tilde{t}}}) > u_i(\rho|_{\sigma^{\tilde{t}}})$.

Condition 1 above simply says that there is exactly only one history at which ρ'_i differs from ρ_i . Given this, condition 2 then requires that, conditional on reaching history $\sigma^{\tilde{t}}$ by means of the proposed strategy profile, player i is strictly better by deviating than by following the prescribed strategy ρ_i . We thus need to consider alternatives that deviate from the suggested profile once and then return to the prescriptions of the suggested profile. This does not imply that the path of generated actions will differ from the equilibrium strategies in only one period. The one-shot profitable deviation (if it exists) leads to a different history than it does in the suggested profile, and then the equilibrium strategies may respond at this history by making different prescriptions.

Proposition 7.1 (One-Shot Deviation Principle). *A strategy profile ρ^* is a SPE of the repeated game Γ_R if and only if there are no one-shot profitable deviations.*

The relevance of the one-shot deviation principle lies in the allowed reduction of the space of deviations that we need to check for. In particular, one does not need to worry about alternative

strategies that might deviate from the equilibrium in period t , and then again in period $t' > t$, and again in period $t'' > t'$. To illustrate this point, let us go back to the example displayed in Figure 7.4. Here we can conclude that the suggested strategy profile does not constitute a SPE simply because if player 2 deviates in history σ^0 from defaulting to cooperating, then she gets a payoff of 7, which exceeds the induced payoff of 5. Of course, we can find more profitable one-shot deviations. Now, suppose that the number of periods $T < \infty$ is instead arbitrary and that the players follow a (pure) strategy combination such that $s_1(\sigma^t) = s_2(\sigma^t) = D$ for each $\sigma^t \in \Sigma^t$, for each $t = 0, 1, \dots, T$. Then, for $T < \infty$, we have that each player is obtaining a payoff of

$$1 \cdot \sum_{t=0}^T \delta^t = \frac{1 - \delta^{T+1}}{1 - \delta}$$

under the prescribed strategy. Now suppose that one of the two players deviates only at some history $\sigma^{\tilde{t}}$ for some period \tilde{t} by choosing C instead and continues to choose D in each other history. Then, such a player would receive instead a payoff of

$$1 \cdot \sum_{t=0}^{\tilde{t}-1} \delta^t + 0 \cdot \delta^{\tilde{t}} + 1 \cdot \sum_{t=\tilde{t}+1}^T \delta^t = \frac{1 - \delta^{\tilde{t}}}{1 - \delta} + \frac{\delta^{\tilde{t}+1} - \delta^{T+1}}{1 - \delta} = \frac{1 - \delta^{T+1}}{1 - \delta} - \left(\frac{\delta^{\tilde{t}}}{1 - \delta} \right).$$

Since, for $\delta \in (0, 1]$, we have $\delta^{\tilde{t}}/(1 - \delta) > 0$, it follows that there are not one-shot profitable deviations and, therefore, that always defaulting is a SPE.

Finally, notice that this logic behind the non-existence of one-shot profitable deviations from always defaulting does not go through when the repeated interaction is infinite instead. Intuitively, this is the case because the loss from deviating to C in a single history vanishes as $T \rightarrow \infty$. In general, the loss from deviating to cooperate in a finite number of histories becomes negligible in infinitely repeated games. This argument will be crucial to obtain equilibria where the players may cooperate for the infinitely repeated prisoner's dilemma stage game. Furthermore, the one-deviation principle can be carefully used to obtain that always defaulting is the unique SPE of the finitely repeated prisoner's dilemma game.

Importantly, the one-shot deviation principle leads directly to the following useful result for

the interesting class of strategy profiles that prescribe the same actions at all the histories available at the same period. Formally, we say that a strategy profile is *history-independent* if $\rho(\sigma^t) = \rho(\tilde{\sigma}^t)$ for each $\sigma^t, \tilde{\sigma}^t \in \Sigma^t$, for each $t \geq 0$.

Proposition 7.2. *A history-independent strategy profile ρ^* is a SPE of the repeated game Γ_R if and only for each t and each $\sigma^t \in \Sigma^t$, $\rho^*(\sigma^t)$ is a NE of the stage game Γ .*

To prove the “if” part note that since each $\rho^*(\sigma^t)$ is a NE of the stage game, then there is no profitable one-shot deviation. Thus, by the one-shot deviation principle, there is no profitable deviation in the extensive form game associated to the repeated interaction. As to the “only if” part, suppose that ρ^* is a SPE of the repeated game and that some $\rho^*(\sigma^t)$ is not a NE of the stage game. Then, there must exist a profitable deviation for some player i in which she deviates from ρ_i^* at σ^t but otherwise plays just as prescribed by ρ_i^* . Because the rest of the players are playing strategies that do not vary with the particular histories available in each period, this must be a profitable deviation in the subgame that initiates at history σ^t . This, in turn, contradicts that ρ^* is a SPE of the repeated game.

Therefore, NEs of the stage game are never ruled out when the players play it repeatedly. This result shows that equilibria are not refined by considering repeated interactions. On the contrary, the set equilibria is in general enlarged. In particular, for the case of infinitely repeated games, this enlargement takes the extreme form—not desirable when the analysis seeks to predict behavior—of “everything goes” implications showed by “folk” theorems.

7.1. Finite *versus* Infinitely Repeated Games

What repeated interactions can offer, both in terms of equilibria structure and of the set of achievable payoffs, depends crucially on whether the interactions are finite in time or go on forever with no particular end date. Although real-world players do not actually live for ever, one appealing justification for assuming that a repeated game is infinite is that, in many cases, players may in practice not envision any ending date whatsoever when they make their choices. For finitely repeated games, if the stage game has a unique NE, then such a NE gives us the only equilibrium

behavior that persists in the long run. The argument behind this result makes use of generalized backwards induction.

Proposition 7.3. *Suppose that $T < \infty$ and that the stage game Γ has a unique (possibly in mixed strategies) NE β^* . Then, the unique SPE of the repeated game Γ_R is the strategy profile ρ^* such that for each t and each $\sigma^t \in \Sigma^t$, we have $\rho^*(\sigma^t) = \beta^*$.*

As already indicated earlier when going through the prisoner's dilemma example, infinite repetitions can allow for much broader sets of equilibria, relative to the case of finitely repeated games. The interesting implication of the theory here is therefore that it allows to rationalize long-term behavior features, strongly based on social norms and on the role of reputation, which cannot be otherwise sustained (neither in one-shot interactions nor in finitely repeated interactions). The drawback of the analysis of infinitely repeated games, though, is that the set of equilibria—and, more importantly, of attainable equilibrium payoffs as well—becomes potentially so large that the predictive/explanatory power of the theory gets critically compromised. These messages, both with their positive and negative implications, can be summarized by the so called “folk” theorems of repeated game theory.

7.2. “Folk” Theorems

To fix ideas, let us go back again to the prisoner's dilemma game depicted in Figure 7.1. Suppose that $\delta \in (0, 1)$ and that this stage game is played during an infinite number of periods, $T \rightarrow \infty$. We turn to explore plausible long-term equilibrium behavior for this situation where, in intuitive terms, the players do not envision an ending for their sequence of future interactions.

Consider first the (pure) strategy profile where the recommendation for each player is verbally described as: (1) play C in $t = 0$, and (2) in any subsequent period $t \geq 1$, play C if the corresponding history is such that no player has played D in the past, and play D otherwise. This sort of strategy is commonly known as a *grim-trigger* strategy. To explore whether this grim-trigger strategy corresponds to a SPE of the infinitely repeated game, we only need to check whether there exists one-shot profitable deviations from it. The simple structure of the proposed profile turns out

helpful to perform this verification task. Note first that there is no profitable one-shot deviation at any history where D has played by some player in the past (i.e., either CD , DC , or DD has been played at least once in the past). This is the case simply because such a one-shot deviation only reduces the player's current period payoff while it does not affect her future payoffs. Therefore, we only need to pay attention to one-shot deviations at histories where CC has always been played in the past. Without loss of generality, we can consider a one-shot deviation for player 1 to playing D in period $t = 0$. By doing so, player 1 "triggers" a perpetual (mutual) punishment from CC to (DC, DD, DD, \dots) . As a consequence, this one-shot deviation is profitable if and only if

$$3 \geq (1 - \delta) \left[4 + 1 \cdot \frac{\delta}{1 - \delta} \right] \Leftrightarrow \delta \geq \frac{1}{3}.$$

Quite interestingly, if the players are sufficiently patient, then we obtain SPE of the repeated game where mutual cooperation is sustained in *all* periods. The (average discounted) payoffs attained under this equilibrium are $(3, 3)$.

Consider now another (pure) strategy profile where the prescription for each player is verbally described as: (1) play C in $t = 0$, and (2) in any subsequent period $t \geq 1$, play whatever the opponent has played in the previous period. This sort of strategy is commonly known as a *tit-for-tat* strategy. Given the symmetry of the proposed profile, we can focus, without loss of generality, on possible one-shot profitable deviations for player 1. Notice that whether or not a one-shot deviation is profitable at some history σ^t depends only on the action profile played in the previous period $t - 1$. In short, what happens in periods $0, 1, \dots, t - 3, t - 2$ is irrelevant. Let us then study the four possibilities $a^{t-1} = CC$, $a^{t-1} = CD$, $a^{t-1} = DC$, and $a^{t-1} = DD$.

For $a^{t-1} = CC$, we have that not deviating from C at σ^t leads to CC in each period along the path of play. On the other hand, deviation to D by player 1 leads to DC, CD, DC, CD, \dots , so that the required condition for the one-shot deviation to be not profitable is

$$3 \geq (1 - \delta)[4 + \delta(0) + \delta^2(4) + \delta^3(0) + \delta^4(4) + \dots] = (1 - \delta) \frac{4}{1 - \delta^2} \Leftrightarrow \delta \geq \frac{1}{3}.$$

For $a^{t-1} = CD$, not deviating from C at σ^t leads to DC, CD, DC, CD, \dots in each period along the

path of play. On the other hand, deviation to D by player 1 leads to CC, CC, CC, CC, \dots , so that the required condition for the one-shot deviation to be not profitable is

$$(1 - \delta) \frac{4}{1 - \delta^2} \geq 3 \Leftrightarrow \delta \leq \frac{1}{3}.$$

For $a^{t-1} = DC$, not deviating from D at σ^t leads to CD, DC, CD, DC, \dots in each period along the path of play. On the other hand, deviation to D by player 1 leads to DD, DD, DD, DD, \dots , so that the required condition for the one-shot deviation to be not profitable is

$$(1 - \delta) \frac{4\delta}{1 - \delta^2} \geq 1 \Leftrightarrow \delta \geq \frac{1}{3}.$$

Finally, for $a^{t-1} = DD$, not deviating from D at σ^t leads to DD, DD, DD, DD, \dots in each period along the path of play. On the other hand, deviation to C by player 1 leads to CD, DC, CD, DC, \dots , so that the required condition for the one-shot deviation to be not profitable is

$$1 \geq (1 - \delta) \frac{4\delta}{1 - \delta^2} \Leftrightarrow \delta \leq \frac{1}{3}.$$

Obviously, there is one (and only one) value for δ that satisfies simultaneously the two requirements $\delta \geq 1/3$ and $\delta \leq 1/3$: for $\delta = 1/3$, the mutual tit-for-tat profile is a SPE of the infinitely repeated prisoner's dilemma game. Interestingly, any payoff combination in the convex hull of the set of payoffs $(1, 1)$, $(1, 3)$, $(3, 1)$, and $(3, 3)$ can be achieved. Geometrically, this tit-for-tat SPE allows the players to attain the payoffs in the area depicted in the graph below.

The examples above raise the natural question of what are the payoff profiles that can be attained by a SPE of an infinitely repeated game. As already mentioned, there is a number of "folk" theorems in the repeated games literature that provide answers. Suppose that $\tilde{Y} = \{(v_1(a), \dots, v_n(a)) : a \in A\}$ is the set of all payoff combinations that are attained by some action profile $a \in A$ of the stage game. The *set of feasible payoffs* is then given by the convex hull of the set \tilde{Y} , which is commonly denoted as $\Upsilon \equiv co(\tilde{Y})$. For instance, the convex hull of the set of payoffs attainable in our prisoner's dilemma stage game is as depicted below.




Figure 7.5 – Attainable Payoffs for the Infinite Prisoner’s Dilemma Repeated Game under Tit-for-Tat

It is intuitive that no payoff combination outside the set of attainable payoffs $\tilde{\Upsilon}$ can be achieved as an average discounted payoff of a SPE of the repeated game. However, the “folk” theorems show that as $\delta \rightarrow 1$ any payoff in the convex hull $co(\tilde{\Upsilon})$ can be obtained as an average discounted payoff by some SPE of the repeated game. In addition to the players being sufficiently patient, these “every goes” results require that the players optimally choose appropriate sequences of action profiles which might change as time evolves.

A couple of definitions are required to state two important “folk” theorems.

Definition 7.2 (Nash-Threat Payoff). *The Nash-threat payoff for player i is the payoff \underline{v}_i specified as*

$$\underline{v}_i \equiv \inf \{ v_i : \text{there exists a stage game NE strategy } \beta^* \text{ (possibly mixed) such that } v_i(\beta^*) = v_i \} .$$

The Nash-threat payoff can be interpreted as a punishment payoff that the players would impose on a given player if she deviates from some prescribed long-run strategy. This gives us formally a threat to those that do not comply to some (equilibrium) pre-established “social norm.” The reputations of the players that deviate from the “agreed” social norm will be damaged and, therefore, they will receive a punishment by the society. It is intuitive that the players can always inflict a punishment to deviators where they allow them at most the payoff of the less beneficial




Figure 7.6 – Convex Hull of Attainable Payoffs for the Infinite Prisoner’s Dilemma Repeated Game

NE in the stage game. All they need to do is specify in the long-run strategy a punishment where they resort forever to their best replies of the NE that yields the lowest payoff to a deviator. Since it is a NE, any deviator i will have to best reply (to the punishment) as well and, therefore, will end up with a payoff of \underline{v}_i from the period of the deviation onwards. This kind of equilibrium punishment is commonly known as *Nash reversion*.

Definition 7.3 (Individually Rational Payoff). A minmax payoff for player i is a payoff \tilde{v}_i specified as

$$\tilde{v}_i \equiv \min_{\beta_{-i} \in \Delta(S_{-i})} \max_{\beta_i \in \Delta(S_i)} v_i(\beta_i, \beta_{-i}).$$

Then, a payoff profile $v = (v_1, \dots, v_n)$ is strictly individual rational if $v_i > \tilde{v}_i$ for each player $i \in N$.

Suppose that a given player plays at each period t of a repeated game the optimal reply to her opponents’ choices. Then, a minmax payoff is the lowest payoff that the players can impose on such a player. In other words, by appropriately choosing her long-run strategy, any player can always guarantee herself her minmax payoff in the stage game. Therefore, in this sense, a payoff lower to a player’s minmax payoff is not individually rational for such a player. It is intuitive that a player’s minmax payoff gives us a credible punishment for such a player in the event that she deviates from some prescribed long-run strategy.

Our two main “folk” theorems are then stated as follows.

Theorem 7.1 (Nash-Threat Folk Theorem). *Take a stage game Γ and pick any $v \in \Upsilon$ such that $v_i > \underline{v}_i$ for each player $i \in N$. Then, there is some $\underline{\delta} \in [0, 1)$ such that for each $\delta > \underline{\delta}$, there is a SPE of the infinitely repeated game Γ_R with average discounted payoff profile v .*

Theorem 7.2 (Minmax Folk Theorem). *Take a (finite) stage game Γ and let $\Upsilon^* \subseteq \Upsilon$ be the set of feasible and strictly individual rational payoffs. Then, if v belongs to the interior of Υ^* , there is some $\underline{\delta} \in [0, 1)$ such that for each $\delta > \underline{\delta}$, there is a SPE of the infinitely repeated game Γ_R with average discounted payoff profile v .*

As already indicated, standard equilibrium requirements do little in repeated interactions to narrow down the set of equilibria. This makes difficult to use the theory to predict behavior and to conduct “comparative statics” analyses. In applications, though, game theorists usually introduce some refinements to narrow down achievable equilibria and payoffs. Some examples of the requirements typically used to refine equilibria in repeated games are efficiency constraints, stationarity in equilibrium proposals, the use of Markov strategies, or of symmetric strategies.

Finally, although “folk” theorems deliver undesirable “anything can happen” messages, they are useful because the constructions of their proofs tell us exactly how payoffs can be achieved and—perhaps more importantly—what are the qualitative features of rewards or punishments implicit in equilibrium social norms.

8. Applications

8.1. Models of Communication: Sender-Receiver Games

Dynamic games where players hold different initial information are particularly rich for applications of economic, political, and social interest. The prominent field of *information economics* is widely recognized for having explored two basic market failures, known as *adverse selection* and *moral hazard*, which are caused by information asymmetries between contracting parties. Such models have also proposed some plausible (partial) solutions to the market failures, such as *signaling* or *screening*. Overall, this literature has made it clear that asymmetries of information between contracting parties lead to inefficiencies.¹ In practice, this kind of failures translate into that (1) job candidates invest too much in their education just to signal their abilities, (2) high quality goods are driven out of the market, and even markets for some goods collapse, when sellers have more information than buyers about the intrinsic characteristics of the goods, or (3) insurance companies fail to offer full insurance to their clients when they cannot observe the effort that the insured agents exert on avoiding accidents.

To explore more specifically these applications, let us consider first a general model where there are two players, an informed Sender (S) and an uninformed Receiver (R). These players' preferences depend on an (endogenous) action $a \in A$ and a (exogenous) state of the world $\theta \in \Theta$. Use $u_i(a, \theta)$ to indicate the utility that accrues to player $i = S, R$ under the action/state pair (a, θ) . The state θ is selected by Nature according to a probability distribution $F(\theta)$. For most parts of the analysis, consider that S learns the true value of θ at no cost.

¹The seminal contributions to information economics are [Akerlof \(1970\)](#), [Spence \(1973\)](#), and [Rothschild and Stiglitz \(1976\)](#). In 2001, George Akerlof, Michael Spence, and Joseph Stiglitz were awarded the Nobel prize in Economics for their work on these topics.

8.1.1. Adverse Selection

Our first application of information economics focuses on the labor market and it is based on [Akerlof \(1970\)](#)'s influential work. To introduce it, we need to add more structure to our general model. In particular, consider that player S summarizes a continuum set of job candidates. Each candidate has a productivity level $\theta \in \Theta \equiv [\underline{\theta}, \bar{\theta}] \subseteq \mathbb{R}_+$ that is drawn independently from the common distribution $F(\theta)$. Therefore, player S can be regarded as a representative job candidate that receives initially her productivity θ according to the distribution $F(\theta)$. In addition, note that, for a given value θ , $F(\theta)$ represents as well the proportion of candidates with a productivity level no higher than θ . Assume further that $F(\theta)$ is a smooth function with a density $f(\theta) > 0$. Player S learns the true value of θ at no cost but R knows only the distribution $F(\theta)$ of productivities. In other words, θ is S 's *private information* or *type*. If S has productivity level θ , then she has a *reservation wage* $r(\theta)$ that describes what she would earn outside of this labor market. Assume that $r(\theta)$ is a continuous function that increases (strictly) in θ ; the candidate obtains higher earnings outside the job market if she has higher levels of productivity. Player R summarizes a continuum of firms that wish to hire job candidates and the representative firm offers a salary $a \in [\underline{a}, \bar{a}] \subseteq \mathbb{R}_+$ to the representative candidate; a salary here corresponds to the action a of the general framework.

Firms act competitively so that (the representative) firm R cannot offer a salary a lower than θ if she knew that the true productivity of the (representative) candidate is θ .² As to the other side of the market, notice that S accepts a job offer a only if $r(\theta) \leq a$.

The specific preferences of these players are given by

$$u_S(a, \theta) = \begin{cases} a & \text{if } r(\theta) \leq a \\ r(\theta) & \text{if } r(\theta) > a, \end{cases}$$

and

$$u_R(a, \theta) = \theta - a.$$

²Suppose that θ were known to firms. Then, if a particular firm offered a salary $w < \theta$, competing firms would benefit by offering a salary $a + \varepsilon$, for $\varepsilon > 0$. Competition for candidates would lead firms to act in such a fashion until they offered a salary no less than θ .

Suppose for a moment that the productivity level θ were known by firms as well so that there were no informational asymmetries between the players. Then, an *efficient equilibrium* would imply that player R offers a salary $\widehat{a}(\theta) = \theta$ and that player S accepts whenever $r(\theta) \leq \theta = \widehat{a}(\theta)$. The efficient fraction of hired candidates would be $\widehat{\Theta} = \{\theta \in \Theta : r(\theta) \leq \theta\}$. The utility levels at this efficient situation for the players would be $u_S(\widehat{a}(\theta), \theta) = \theta$ and $u_R(\widehat{a}(\theta), \theta) = 0$. Importantly, in this case firms would be able to make wage offers $\widehat{a}(\theta)$ which are contingent on the (observed) productivity θ of the candidates.

However, in the considered job market, where θ is in fact S 's private information, a *competitive equilibrium* is a salary a^* —which cannot be conditional on θ —and a fraction of hired candidates Θ^* that satisfy:

$$\Theta^* = \{\theta \in \Theta : r(\theta) \leq a^*\} \quad (8.1)$$

and

$$a^* = \mathbb{E}[\theta \mid \theta \in \Theta^*]. \quad (8.2)$$

Thus, to find equilibria, we need to solve both equations (8.2) and (8.1) simultaneously. Substitution of (8.1) into (8.2), leads to the following equilibrium condition for the salary

$$a^* = \mathbb{E}[\theta \mid r(\theta) \leq a^*].$$

Then, if we consider a function $\phi : [\underline{\theta}, \bar{\theta}] \rightarrow [\underline{\theta}, \bar{\theta}]$ specified as $\phi(a) \equiv \mathbb{E}[\theta \mid r(\theta) \leq a]$, we observe that the required equilibrium condition tells us that an equilibrium salary a^* is a *fixed point* for the function ϕ .³ We can guarantee the existence of (at least) a fixed point for a real value function such as $\phi : [\underline{\theta}, \bar{\theta}] \rightarrow [\underline{\theta}, \bar{\theta}]$ if $\phi(\theta)$ is a continuous function.⁴ Notice that continuity of the function $\phi(\theta)$ holds because we are assuming that $r(\theta)$ is a continuous function of θ and that the density $f(\theta)$ is positive.

With these elements at hand, we can now compare the efficient situation associated to having no informational asymmetries with the competitive equilibrium under private information. To make

³ A *fixed point* of a function $\phi : B \rightarrow B$ is a value $x \in B$ such that $x = \phi(x)$.

⁴ This claim follows directly from *Brower's fixed point theorem*.

our comparison as clear as possible, assume further that the function $r(\theta)$ satisfies $r(\theta) \equiv \theta - \varepsilon$ for each $\theta \in \Theta$, for some $\varepsilon > 0$. Then, the efficient situation would imply that all candidates are hired, $\widehat{\Theta} = \Theta$. As to the competitive equilibrium we have already argued that there is at least one equilibrium wage $a^* \in [\underline{\theta}, \bar{\theta}]$. We cannot rule out the possibility that there are multiple equilibria but, for the sake of simplicity, let us restrict attention to situations with a unique competitive equilibrium. One possibility then is that $a^* \in (\underline{\theta}, \bar{\theta})$. In this case, we observe that there must be some $\tilde{\theta}(\varepsilon) \in (a^*, \bar{\theta}]$ such that $r(\theta) > a^*$ for each $\theta \in (\tilde{\theta}(\varepsilon), \bar{\theta}]$. Then, as $\varepsilon > 0$ vanishes, it follows from the expression in (8.1) that the fraction of hired candidates tends to $\Theta^* = [\underline{\theta}, a^*]$, with $a^* < \bar{\theta}$. In other words, as the (outside) reservation wage approaches the actual productivity of candidates, the most efficient candidates will be faced with an average salary offer a^* lower than what they could earn outside the job market. This phenomenon is known as *market unravelling* and describes a situation where the most productive candidates are driven out of the market. Intuitively, when the average offered salary $a^* = \mathbb{E}[\theta \mid r(\theta) \leq a^*] < \bar{\theta}$ is less than their outside option, the most productive candidates reject the offer (and turn to the outside option). An extreme case of this market unravelling would be that $a^* = \underline{\theta}$, where the market completely breaks down. The (ex-post) utility levels at this inefficient equilibrium for the players would be $u_S(a^*, \theta) = \mathbb{E}[\theta \mid r(\theta) \leq a^*]$ and $u_R(a^*, \theta) = \theta - \mathbb{E}[\theta \mid r(\theta) \leq a^*]$.

8.1.2. Signaling

To mitigate the market inefficiencies caused by adverse selection, in practice individuals make certain decisions that, under some conditions, allow them to transmit credibly their private information to the uninformed parties. Following our job market story, one possibility here is that candidates make *costly* education choices to transmit information to employers about their true productivity. Notably, although education might actually improve one's productivity, it is crucial for the mechanism that we now present to work that education does not affect the candidate's productivity. Considering that education does improve one's productivity is conceivably a plausible assumption in some applications but *signaling* does not consider this assumption as part of its mechanism. In our story, firms are aware of this and they do not expect to find gains in a candi-

date's productivity due to education. Instead, firms know how the cost of investing in education relates to the candidates' actual productivities—that is, as usually, they know the entire description of the game—which allows them to make inferences on the productivity of the candidates.

The *signaling mechanism* considered here is based on Spence (1973)'s work. The intuitive message is that, in some circumstances, high-productivity candidates may have incentives to invest in costly education simply to distinguish themselves from low-productivity candidates. Then, firms may infer that candidates that invest in education have in fact high productivity levels. We have already analyzed, in Section 6.1, two extensive form games that describe this sort of signaling situations: the canonical job market signaling game with binary action choices and the investment trust game depicted in Figure 6.2. As considered then, PBNE and sequential equilibrium are the suitable solution concepts that allow us to study behavior in these environments.

The general setting of *signaling* follows by enlarging the domain of preferences in the adverse selection model described earlier in Subsection 8.1.1. In addition, we drop out for simplicity the outside activity and, therefore, also the reservation wage $r(\theta)$ from the preference specification. Neither the formal structure of the model nor its intuitive messages rely on this element.⁵ Specifically, preferences depend now also on the choice of an education level $m \in [\underline{m}, \bar{m}] \subseteq \mathbb{R}_+$ by player S . In general, a choice such as m is considered in signaling games as a *signal* or message. *Signals* are actions, observable by R , that can be chosen conditional on S 's unobservable private information. This correlation may allow R to learn something about S 's private information. The players' preferences are described now by utility functions $u_S(m, a, \theta)$ and $u_R(m, a, \theta)$. A particularly natural specification for these utilities in many applications would be: $u_S(m, a, \theta) = a - c(m, \theta)$ and $u_R(m, a, \theta) = \theta - a$, where $c : [\underline{m}, \bar{m}] \times [\underline{\theta}, \bar{\theta}] \rightarrow \mathbb{R}_+$ is a function that describes the cost to a candidate of productivity θ from investing in an education level m .

In this setting, behavior strategies are given by a collection of probability distributions⁶ $\sigma = \{\sigma(\cdot | \theta) \in \Delta([\underline{m}, \bar{m}]) : \theta \in [\underline{\theta}, \bar{\theta}]\}$, for job candidates, and a collection $\alpha = \{\alpha(\cdot | m) \in \Delta([\underline{a}, \bar{a}]) : m \in [\underline{m}, \bar{m}]\}$, for firms. In particular, $\alpha(m | \theta)$ indicates the probability that the representative candidate chooses

⁵This reservation wage $r(\theta)$ was a crucial element behind the rationale of the adverse selection phenomenon, but it is no longer needed to explore how players can credibly signal their pieces of private information.

⁶All notation is presented for the case where the sets of actions are uncountably infinite. The model continues to work through for the case where actions are finite by suitably modifying notation.

education level m when her productivity is θ and $\alpha(a | m)$ is the probability that the representative firm offers salary a when it observes a candidate with education level m . Recall that $F(\theta), f(\theta)$ describe the prior beliefs in this model. Then, let $\mu(\theta | m)$ be the posterior probability that the representative firm assigns to θ being the true productivity of the candidate when it observes a résumé with education level m . That is, posteriors are described by a collection of probability distributions $\mu = \{\mu(\cdot | m) \in \Delta(\Theta) : m \in [\underline{m}, \bar{m}]\}$. The application of the general definition of *PBNE* to this setting is:

Definition 8.1. *A signaling PBNE is a pair of behavior strategies (σ^*, α^*) and a system of (posterior) beliefs μ^* such that*

1. *Optimal response by S: for each $\theta \in [\underline{\theta}, \bar{\theta}]$, we have that $\sigma^*(m | \theta) > 0$ implies that*

$$\int_{\underline{a}}^{\bar{a}} \alpha^*(a | m) u_S(m, a, \theta) da \geq \int_{\underline{a}}^{\bar{a}} \alpha^*(a | m') u_S(m', a, \theta) da \quad \forall m' \in [\underline{m}, \bar{m}].$$

2. *Optimal response by R: for each $m \in [\underline{m}, \bar{m}]$ such that $\sigma^*(m | \theta) > 0$ for some $\theta \in [\underline{\theta}, \bar{\theta}]$, we have that $\alpha^*(a | m) > 0$ implies that*

$$\int_{\Theta} \mu^*(\theta | m) u_R(m, a, \theta) d\theta \geq \int_{\Theta} \mu^*(\theta | m) u_R(m, a', \theta) d\theta \quad \forall a' \in [\underline{a}, \bar{a}].$$

3. *Consistent posteriors: posteriors follow from priors and from S's signaling choice according to Bayes' rule, that is, for each $m \in [\underline{m}, \bar{m}]$ such that $\sigma^*(m | \theta') > 0$ for some $\theta' \in [\underline{\theta}, \bar{\theta}]$, we have that*

$$\mu^*(\theta | m) = \frac{\sigma^*(m | \theta) f(\theta)}{\int_{\Theta} \sigma^*(m | \theta') f(\theta') d\theta'}.$$

Condition 1 above says that a candidate with productivity θ places positive probability on an education level m only if that choice maximizes her expected utility given the optimal choice α^* of the firm. Condition 2 states that, upon observing education level m , the firm places positive probability on a salary a only if it maximizes its expected utility, conditional on the induced posteriors μ^* . Lastly, condition 3 requires that the posteriors μ^* be consistent with the candidate's

optimal choice σ^* and the priors $f(\theta)$, according to Bayes' rule. Importantly, the required condition for consistent posteriors only states how posteriors follow from the standard notion of conditional probability for information sets $m \in [\underline{m}, \bar{m}]$ that are in fact reached under the optimal choice of the candidate. To determine how posteriors would be obtained at education levels chosen with zero probability under S 's proposed behavior strategy, we would need to add the requirement in the definition of the sequential equilibrium refinement stated in Section 6.1.

As already explored in Section 6.1, the extent to which S 's strategy credibly communicates can be categorized using three "qualitative" classes of equilibrium proposals: pooling, completely separating, and semi-separating equilibria.

1. We say that a signaling equilibrium $(\sigma^*, \alpha^*, \mu^*)$ is *pooling* if S 's strategy is independent of her type: for each $\theta \in [\underline{\theta}, \bar{\theta}]$, we have $\sigma^*(m | \theta) = \sigma^*(m)$ for each $m \in [\underline{m}, \bar{m}]$. As a consequence, notice that the requirement in 3 of the equilibrium definition leads to that posteriors coincide with priors in this class of equilibria: for each m such that $\sigma^*(m) > 0$, we have $\mu^*(\theta | m) = f(\theta)$ for each $\theta \in [\underline{\theta}, \bar{\theta}]$.
2. On the opposite extreme of the spectrum, we say that a signaling equilibrium $(\sigma^*, \alpha^*, \mu^*)$ is *completely separating* if each type θ of S sends a different signal with probability one: for each $\theta \in [\underline{\theta}, \bar{\theta}]$, we have $\sigma^*(m | \theta) = 1$ if $m = m^*(\theta)$, where $m^*(\theta)$ is a (strictly) monotone function in $\theta \in [\underline{\theta}, \bar{\theta}]$. Then, the requirement in 3 of the equilibrium definition leads to that, according to its posteriors, the firm learns always the true productivity of the candidate: for each m such that $\sigma^*(m | \theta) > 0$ for some $\theta \in [\underline{\theta}, \bar{\theta}]$, we have $\mu^*(\theta | m) = 1$ whenever $m = m^*(\theta)$.

In addition to these two extreme cases, we can also obtain very interesting *semi-separating equilibria* that are partially revealing of S 's private information.

It is not difficult to construct a pooling—completely uninformative—equilibrium for the presented model. Suppose that all candidates, regardless of their true productivities, pick with probability one some given education level $\tilde{m} \in [\underline{m}, \bar{m}]$, that is, $\sigma^*(\tilde{m} | \theta) = 1$ for each $\theta \in [\underline{\theta}, \bar{\theta}]$. Then, consider a best-reply by the firms such that, upon observing \tilde{m} , they offer with probability one a salary \tilde{a} that maximizes their expected utility when the posterior belief coincide with the

prior, that is,

$$\tilde{a} \in \arg \max_{a \in [\underline{a}, \bar{a}]} \int_{\Theta} f(\theta) u_R(\tilde{m}, a, \theta) d\theta.$$

On the other hand, suppose that for any other education level $m \neq \tilde{m}$, firms offer with probability one a salary that maximizes their expected utility when the posterior assigns probability one to $\underline{\theta}$ being the true productivity, that is,

$$\hat{a} \in \arg \max_{a \in [\underline{a}, \bar{a}]} u_R(m, a, \underline{\theta}).$$

For instance, for the specification $u_R(m, a, \theta) = \theta - a$, we have $\tilde{a} = \mathbb{E}[\theta]$ and $\hat{a} = \underline{\theta}$ —of course, provided that firms are in a competitive environment where they cannot offer a salary lower than θ . Then, provided that the least productive candidate does not strictly prefer salary \hat{a} over \tilde{a} , all types θ of S prefer salary \tilde{a} over \hat{a} so that everyone will in equilibrium pick education level \tilde{m} rather than any other education level. Intuitively, starting from a situation where all candidates pick a given education level \tilde{m} and receive in return the salary \tilde{m} firms would offer when they have no other information beyond the prior $F(\theta), f(\theta)$, no type θ of player S has incentives to choose an education level different from \tilde{m} . This is the case because deviating from \tilde{m} would lead firms to offer the salary \hat{a} they would optimally offer to the least skilled candidate. By construction of the equilibrium, \hat{a} acts as the highest possible penalty, which would leave any deviating candidate strictly worse off.

It is not surprising that signaling situations have equilibria where Senders do not transmit any information whatsoever because, owing to the conflict captured by preferences, they have some interest in not letting Receivers know the true value of the unknown state. It is the existence of separating equilibria what constitutes the truly interesting case for signaling situations. When separating equilibria exists, this means that it is possible for Senders to transmit credibly their private information even though there is a conflict of interest between them and the Receivers. As it turns out, the existence of separating equilibria requires a systematic relationship between signals and types. A particularly appropriate specification for this relationship is commonly known as the *single-crossing condition* and it plays an important role in games of information transmission and,

more generally, in situations of asymmetric information.

Definition 8.2 (Single-Crossing Condition). *The preferences of the Sender satisfy the single-crossing condition if $u_S(m', a', \theta) > u_S(m, a, \theta)$ for two signals $m' > m$ implies $u_S(m', a', \theta') > u_S(m, a, \theta')$ for each $\theta' > \theta$.*

In words, the single-crossing condition requires that a candidate of a certain productivity θ' that picks an education level m' , higher than another m and, in return, receives a salary a' be (strictly) better off relative to having a pair (m, a) whenever any candidate with lower productivity θ ranks (weakly) the pairs (m', a') and (m, a) in the same way. In most applications, $u_S(m, a, \theta)$ is decreasing in m and increasing in a so that indifference curves over the space of commodities $(m, a) \in [\underline{m}, \bar{m}] \times [\underline{a}, \bar{a}]$ are well-defined for each $\theta \in [\underline{\theta}, \bar{\theta}]$ (think of m as a bad and of a as a good in typical applications). Then, this condition can geometrically be seen as any two indifference curves over the space of commodities $(m, a) \in [\underline{m}, \bar{m}] \times [\underline{a}, \bar{a}]$ of two different types of S , θ and θ' , crossing exactly once—hence, the name. Notice that if a given type θ is indifferent between two pairs (m', a') and (m, a) , then the single-crossing condition states that a higher type θ' must prefer to pick m' rather than m . Therefore the single-crossing condition relates signals to types in a way such that higher productivity candidates prefer (weakly) to pick higher education levels in equilibrium. Since firms know the entire description of the candidates' preferences, they can accordingly infer in equilibrium that higher education levels correspond to higher productivity candidates.

Suppose that $u_S(m, a, \theta)$ is a smooth function, strictly decreasing in m and strictly increasing in a . Then, we can study the expression $a(m; \theta)$ of an indifference curve for type θ of S by considering $u_S(m, a(m; \theta), \theta) = \bar{u}_S(\theta)$ so that

$$\frac{\partial u_S}{\partial m} + \frac{\partial u_S}{\partial a} \cdot \frac{da(m; \theta)}{dm} \Big|_{\bar{u}_S(\theta)} = 0 \quad \Rightarrow \quad \frac{da(m; \theta)}{dm} \Big|_{\bar{u}_S(\theta)} = - \frac{\partial u_S(\cdot, \theta) / \partial m}{\partial u_S(\cdot, \theta) / \partial a}.$$

Then, if the ratio $-\frac{\partial u_S(\cdot, \theta) / \partial m}{\partial u_S(\cdot, \theta) / \partial a}$ decreases in θ , we obtain the single-crossing condition. The differentiable version of the single-crossing condition is often known as the *Spence-Mirrlees condition*.

Figure 8.1 illustrates the Spence-Mirrlees condition with two indifference curves that yield a

common level of utility $\bar{u}_S = \bar{u}_S(\theta') = \bar{u}_S(\theta)$ to two different types θ', θ of the Sender, with $\theta' > \theta$. We observe that the feature that the slope $da(m; \theta)/dm|_{\bar{u}_S(\theta)}$ of the indifference curve decreases with θ leads to that the indifference curves cross exactly once. Figure 8.1 also depicts another indifference curve for the high type θ' of the Sender that yields such a type a higher utility level, that is, $\tilde{u}_S(\theta') > \bar{u}_S(\theta')$. Notice that starting from an education/salary pair (m_1, a_1) , there is a region, bounded from below by the blue indifference curve and from above by the red indifference curve, where type θ' increases her utility by raising the education level but where type θ utility decreases instead. This gives us the geometric illustration of how the single-crossing condition may lead higher types benefit by raising their education levels while, at the same time, it is not in the interest of lower types to do raise their education levels. In particular, the low productivity candidate is indifferent between the pairs (m_1, a_1) and (m_2, a_2) whereas the high productivity candidate strictly prefers (m_2, a_2) over (m_1, a_1) . In this way, candidates with different productivity levels may have incentives to “separate” themselves from each other. Consider the preference specification




Figure 8.1 – Spence-Mirrlees condition ($\theta' > \theta$)

$u_S(m, a, \theta) = a - c(m, \theta)$ and assume further that $c(m, \theta)$ is twice-continuously differentiable with $\partial c(m, \theta)/\partial m > 0$ for each $\theta \in [\underline{\theta}, \bar{\theta}]$; obtaining higher levels of education is more costly. Then, the Spence-Mirrlees condition—which, in turn, implies the single-crossing condition—requires that the derivative $\partial c(m, \theta)/\partial m > 0$ decreases in θ . Intuitively, the marginal cost of investing in

education must be higher for low productivity candidates. Does it make sense?

The single-crossing condition is useful to obtain the following version of the existence of a completely separating equilibrium for our signaling game. A few additional (and more technical conditions) conditions are required in this version of the existence result.

Proposition 8.1. *Consider our signaling game and, furthermore, suppose that*

1. $\Theta = \{\theta_1, \dots, \theta_k, \dots, \theta_K\}$ is a finite set of states of the world.
2. $u_R(m, a, \theta)$ and $u_S(m, a, \theta)$ are continuous in (m, a) .
3. $u_S(m, a, \theta)$ is strictly decreasing in m and strictly increasing in a , for each $\theta \in \Theta$.
4. R 's best-reply when it leans fully S 's type,

$$BR_R(\theta) \equiv \arg \max_{a \in [\underline{a}, \bar{a}]} u_R(m, a, \theta),$$

is uniquely defined, independent of the value of m .

5. *The single-crossing condition holds.*
6. *There exist some education levels $m_1^*, m^H \in [\underline{m}, \bar{m}]$ such that*

$$u_S(m^H, BR_R(\theta_K), \theta_K) < u_S(m_1^*, BR_R(\theta_1), \theta_1).$$

Then, this signaling game has a completely separating equilibrium where the set of education levels $[\underline{m}, \bar{m}]$ is partitioned into a finite set

$$\{[\underline{m}, m_1^*], (m_1^*, m_2^*], \dots, (m_{k-1}^*, m_k^*], \dots, (m_{K-1}^*, m_K^*]\}, \quad \text{with } m_K^* = \bar{m},$$

of equilibrium education subsets, each type $\theta_k \in \Theta$ sends with probability one the education level m_k^ , and R 's posterior puts probability one on θ_k being the true type of S when it observes education level m_k^* .*

The requirement in 6. of the proposition is a boundary condition that makes sending high levels of education, such as m^H , unattractive to S . In particular, it requires that the most skilled

candidate prefers to be treated by the job market as the least skilled candidate rather than choosing a very high education level m^H .

To show how the key result in the proposition works, let us construct a separating equilibrium using the sufficient conditions stated above. Consider the following equilibrium proposal:

- Type θ_1 of S chooses an education level m_1^* that solves her utility maximization problem, provided that R learns the true type of S , that is

$$m_1^* \in \arg \max_{m \in [\underline{m}, \bar{m}]} u_S(m, BR_R(\theta_1), \theta_1)$$

- Suppose that $m_\kappa^* \in \arg \max_{m \in [\underline{m}, \bar{m}]} u_S(m, BR_R(\theta_\kappa), \theta_\kappa)$ have been specified for $\kappa = 1, \dots, k-1$ and set the value functions $v_S^*(\theta_\kappa) \equiv u_S(m_\kappa^*, BR_R(\theta_\kappa), \theta_\kappa)$. Then, let m_k^* be an education level that maximizes type θ_k utility (again, conditional on that R learns S 's true productivity), subject to the condition that type θ_k does not (strictly) wish to mimic the immediately lower productivity candidate, that is,

$$m_k^* \in \arg \max_{m \in [\underline{m}, \bar{m}]} u_S(m, BR_R(\theta_k), \theta_k)$$

$$\text{s.t.: } u_S(m, BR_R(\theta_k), \theta_k) \leq v_S^*(\theta_{k-1}).$$

Then, let this process go on iteratively to determine the remaining education levels m_κ^* for $\kappa = k+1, \dots, K$. Clearly, with this process, each player is best responding to the strategy chosen by her opponent, conditional on the information that she possesses according to Bayesian posteriors.

- As always, to check whether there are incentives to deviate, we also need to specify the players' strategies at information sets that fall out of the path described by equilibrium behavior. Then, suppose that R chooses (a) $BR_R(\theta_k)$ upon education levels in the interval $[m_k^*, m_{k+1}^*)$, (b) $BR_R(\theta_1)$ upon education levels lower than m_1^* , and (c) $BR_R(\theta_K)$ upon education levels higher than m_K^* . That is, firms are assumed to regard (a) any education level that would correspond to θ_k , but which indeed exceeds m_k^* , as if it actually corresponds to type θ_k , (b)

any education level lower than the lowest level in equilibrium as if it corresponds to the least skilled candidate, and (c) any education level higher than the highest level in equilibrium as if it corresponds to the most skilled candidate.

In this equilibrium proposal, the single-crossing condition begins playing its role for types no lower than θ_k . Suppose that the best-reply of the firms $BR_R(\theta_\kappa)$ is strictly increasing in θ_κ . Then, the single-crossing condition implies that the candidates' choices m_κ^* are strictly increasing in their true productivity θ_κ .

This construction of equilibria involves an “unraveling argument” where high-productivity types find it profitable to send signals different from lower productivity types. Furthermore, the equilibria that we have constructed involve inefficient levels of signaling. Under the condition that $u_S(m, a, \theta)$ is decreasing in the education level m , it follows that all types θ_κ , for $\kappa = 2, \dots, K$, invest in education more than they would do in the absence of informational asymmetries in order avoid being treated by firms like less skilled candidates.

As we can observe from the earlier constructions of pooling and separating equilibria, signaling games exhibit a profound multiplicity of equilibria. The reason behind this multiplicity is that the PBNE notion does not place restrictions on the strategies that Receivers may choose at information sets that are not reached under the Senders' strategies in equilibrium proposals. To select among the multiple equilibria, the literature offers some refinements, which rely on [Kohlberg and Mertens \(1986\)](#)'s notion of *strategy stability* and are appealing to signaling situations.

8.1.3. Cheap Talk

Although the systematic relationship between pieces of private information and signals captured by the single-crossing condition makes sense in many environments, other situations do not fit into this assumption. An extreme instance of environments where the single-crossing condition is absent are those where, because of their preference specification, sending signals is costless for Senders. Given the general notation $u_i(m, a, \theta)$ to describe the preferences of each player $i = S, R$, here $a \in A$ is an action taken by the uninformed Receiver, $m \in M$ is a message chosen by the

informed Sender,⁷ and $\theta \in \Theta$ is a payoff-relevant state of the world. Then, a *cheap-talk* model is a signaling game where $u_i(m, a, \theta)$ is independent of m : $u_i(m, a, \theta) = u_i(a, \theta)$ for each (a, θ) . Intuitively, player S can talk as she pleases about her private information and such talking bears her no direct cost. These situations abound, right? Notably, unlike what we studied in repeated games, we are not considering here a repeated interaction between S and R so that we might think of costs to S in terms of reputation. Cheap-talk settings are static, with a one-shot interaction between the players. Thus, since there is no cost to S from not revealing her private information, neither in her preference specification nor in terms of reputation, one may wonder why would she have incentives to transmit any information at all. The key is that, by not transmitting (at least some of) her private information, S may induce R to take an action that, in turns, hurts S . In this way, cheap-talk situations may capture endogenous costs to Senders that are realized in equilibrium. The interesting point of cheap-talk models is that such endogenous costs may lead Senders to reveal some information, despite of their conflicting interests with Receivers and the absence of direct costs from hiding information.

The model that we know present follows closely Crawford and Sobel (1982). Assume that $\Theta \equiv [0, 1]$ and that the prior on θ is given by a continuous distribution $F(\theta)$ with density $f(\theta)$. Assume also that the utility functions $u_i(a, \theta)$ are twice-continuously differentiable for $i = S, R$ and they have a unique maximum $a_i(\theta)$ for each θ . A typical example of utility functions that satisfy these assumptions are those corresponding to *single-peaked preferences* where one player is biased relative to the other: $u_S = -(a - (\theta + b))^2$ and $u_R = -(a - \theta)^2$, with $b > 0$. In this case, $a_S(\theta) = \theta + b$ while $a_R(\theta) = \theta$, so that the conflict of interest between S and R is described by the *bias* b . Most applications of the cheap-talk theory consider such preference specification and, furthermore, assume that $F(\theta)$ corresponds to the uniform distribution on the interval $[0, 1]$.⁸

The *cheap-talk game* proceeds as follows. First, Nature select a value θ for the state and S learns such a value. Then, S sends to R a (costless) message $m \in M$, where M is any infinite set. Finally,

⁷We rather use the terminology “message” instead of “signal” because this is the usual convention in the related literature. Other than the assumption that they do not enter the players’ utility functions, messages m play a totally analogous role to signals in signaling models. Their purpose in both classes of models is for the informed player to transmit information to the uninformed player.

⁸We usually refer to this case as the “uniform-quadratic” leading example.

R uses the message received from S to update her priors and to choose (optimally, based on her posteriors) an action $a \in A$. The selected pair (a, θ) determines the payoffs $u_i(a, \theta)$ to the players. In this environment, a (pure) strategy⁹ for S is a function $a_R : \Theta \rightarrow M$ and a (pure) strategy for R is a function $m_R : M \rightarrow A$. Then, $m_S(\theta)$ indicates the message sent by S when she learns that the true value of the state is θ and a_R indicates the action chosen by R when she hears message m .

A PBNE for this game is defined in the usual way.

Definition 8.3. A cheap-talk (pure-strategy) PBNE is a pair of strategies (m_S^*, a_R^*) and a system of beliefs μ^* such that

1. Optimal response by S : for each $\theta \in \Theta$, we have that $m_S^*(\theta) = m$ implies that

$$u_S(a_R^*(m), \theta) \geq u_S(a_R^*(m'), \theta) \quad \forall m' \in M.$$

2. Optimal response by R : for each $m \in M$ such that $m_S^*(\theta) = m$ for some $\theta \in [0, 1]$, we have that $a_R^*(m) = a$ implies that

$$\int_0^1 \mu^*(\theta | m) u_R(a, \theta) d\theta \geq \int_0^1 \mu^*(\theta | m) u_R(a', \theta) d\theta \quad \forall a' \in A.$$

3. Consistent posteriors: posteriors follow from priors and from S 's signaling choice according to Bayes' rule, that is, for each $m \in M$ such that $m_S^*(\theta') = m$ for some $\theta' \in [0, 1]$, we have that

$$\mu^*(\theta | m) = \begin{cases} f(\theta) / \int_{[\theta' : m_S^*(\theta')=m]} f(\theta') d\theta' & \text{for } m_S^*(\theta) = m; \\ 0 & \text{for } m_S^*(\theta) \neq m. \end{cases}$$

We say that a pair of equilibrium strategies (m_S^*, a_R^*) induces action a if the probability (according to priors) that some states lead R to choose action a under (m_S^*, a_R^*) is positive. Then, the set of

⁹Unlike the exposition of signaling games, we consider now only strategies without randomizations because this suffices to explore the most important features of cheap-talk models. In particular, the stated assumptions guarantee that R 's best-reply is a unique and that R will not randomize in equilibrium. Also, Crawford and Sobel (1982) demonstrate that S will neither randomize in equilibrium.

actions induced in equilibrium is

$$\left\{ a \in A : \int_{[\theta \in \Theta : a_R^*(m_S^*(\theta))=a]} f(\theta) d\theta > 0 \right\}.$$

Notice that a pooling equilibrium would induce a single action because the Receiver would maximize her expected utility using her priors regardless of the messages that she receives. On the other extreme, under the stated assumptions, if $a_R(\theta)$ is a monotone function, then a completely separating equilibrium necessarily induces an infinite uncountable set of actions. Specifically, in this case, the cardinality of the set of induced actions in a completely separating equilibrium must coincide with the cardinality of the unit interval $[0, 1]$. This is the case because the Receiver learns the true value of the state. Then, if $a_R(\theta)$ is monotone, R chooses a different optimal action for each $\theta \in [0, 1]$.¹⁰ Crawford and Sobel (1982) showed that the set of induced actions in any cheap-talk PBNE lies between 1 and a finite upper bound. More specifically, Crawford and Sobel (1982) demonstrated that there exists a positive integer N^* such that (1) there is no equilibrium than induces more than N^* actions and (2) for each integer N with $1 \leq N \leq N^*$ there exists at least one PBNE where the set of induced actions has cardinality N . In other words, for cheap-talk situations, (1) there is no completely separating equilibrium and (2) there is always at least one pooling equilibrium and, perhaps, other semi-separating equilibrium where S transmits partially her private information. Full communication is not possible in cheap-talk situations and it is always possible for the Sender to “babble” when she talks cheap. These implications are intuitive. It is perhaps less intuitive that there may exist other equilibria where the Sender does communicate partially despite of the absence of costs from hiding information and the presence of conflicting interests with the Receiver. Under uniform priors, and the preference specification $u_S = -(a - (\theta + b))^2$ and $u_R = -(a - \theta)^2$, it follows that the upper bound $N^*(b)$ depends continuously on the bias b . Furthermore, $\lim_{b \rightarrow 0} N^*(b) = +\infty$ and $\lim_{b \rightarrow +\infty} N^*(b) = 1$. That is, for the “uniform-quadratic” example, equilibrium can be fully revealing if the bias vanishes and only “babbling” equilibria exist when the bias is sufficiently high.

As it was the case in signaling games, cheap-talk games have multiple equilibria. In cheap-talk

¹⁰ Clearly, this is necessarily the case for the specification $u_R(a, \theta) = -(a - \theta)^2$.

situations, though, equilibria multiplicity is even more profound. Specifically, cheap-talk games have three types of equilibrium indeterminacies:

1. Multiple best-replies to messages that are not sent with positive probability in the equilibrium proposal. This type of indeterminacy is typical in dynamic games with imperfect information and it is not a particular feature of cheap-talk models.
2. Multiple meaning of messages. This type of indeterminacy is a particular feature of cheap-talk situations. Notice that any given system of equilibrium beliefs where S does not babble may be induced by multiple different associations between S 's type θ and the message m . In other words, when there is some communication, messages have no intrinsic meaning and—unless we impose further assumptions about the intrinsic meaning of messages—the set of messages does not give us a “language” where each vocable could be interpreted literally. One can determine the amount of information transmitted for each equilibrium but the actual meaning of messages cannot be determined in equilibrium. In other words, language is endogenously determined simultaneously by both S and R .
3. Multiple associations between states of the world θ and induced actions in equilibrium. This type of indeterminacy follows from the lack of intrinsic meaning of messages.

Many equilibrium refinements have been proposed in the literature to reasonably select among equilibria. Interestingly, using the [Kohlberg and Mertens \(1986\)](#)'s strategic stability approach, [Chen et al. \(2008\)](#) propose a selection criterion that selects the most informative equilibria in cheap-talk situations, that is, equilibria that induce exactly a number N^* of actions. Their criterion requires that a condition, which they label as “no incentive to separate (NITS),” be satisfied. In words, the NITS condition requires that the lowest type $\theta = 0$ of S prefers her equilibrium payoff to the payoff that she would obtain if R knew her type and chose accordingly her optimal action. Selection of the most informative equilibria requires also an additional condition regarding monotonicity of R 's best-reply. Such conditions required for the selection of the most informative equilibria under [Chen et al. \(2008\)](#)'s approach are in particular satisfied in the “uniform-quadratic” case, which is frequently considered in applications.

While fully revealing equilibria exist in signaling games that satisfy the single-crossing condition, standard cheap-talk games do not have fully communicative equilibria. This leads to the question of whether one can propose interesting variations of the standard cheap-talk game in order to obtain fully communicative equilibria. Unfortunately, the mechanism behind the standard cheap-talk framework that leads to severe limits in communication is robust. This is why variations of the model where communication is enhanced are necessarily based on introducing some common interests between the players. One way of doing this is by considering that both the set of states and the set of actions are multi-dimensional. In a nutshell, notice that effective communication requires that different types of S have different preferences over their final payoffs. In signaling games, the required difference in preferences appear because there is a direct cost from sending messages. Then, we can reasonably assume that different types bear different costs from sending their messages. In the standard cheap-talk setting, the difference appears when different types of S prefer different optimal actions picked by R . Interestingly, when states and actions have multiple dimensions, the heterogeneity may appear simply from the fact that different types may have different preferences about the relative importance of how the action matches the state across their different dimensions. With this logic, [Chakraborty and Harbaugh \(2010\)](#) show that fully informative equilibria exist in multi-dimensional settings when the preferences of S do not depend on the state.¹¹ Following this reasoning, [Battaglini \(2002\)](#) shows that in n -dimensional models with quadratic preferences, S and R have common interests along $n - 1$ dimensions which, under some conditions, enables full communication.

Another interesting variation is to consider more than one Sender and ask whether competition between Senders can enhance communication. For one-dimensional cheap-talk models, [Krishna and Morgan \(2001\)](#) show that full revelation is possible in equilibrium when Senders speak out simultaneously. Their argument is as follows. Consider two Senders $j = 1, 2$, with positive biases $b_j > 0$, and a Receiver that wishes to choose an action that coincides with the true value of the state, $a_R(\theta) \in \theta$ for each $\theta \in [0, 1]$. Let us restrict attention to equilibria where messages have

¹¹ Also, [Chakraborty and Harbaugh \(2007\)](#) show that full communication exists in multi-dimensional settings when (1) the different dimensions of the state are independently drawn from a common distribution, (2) preferences are additively separable across the different dimensions, and (3) messages are constrained to stating whether one coordinate is larger than another.

intrinsic meaning so that $m_j = \theta \in [0, 1]$ is a statement that θ is the true value of the state. Let $a_R(m_1, m_2)$ be the action chosen by the Receiver when she hears message m_j from Sender j . Then, consider an equilibrium proposal where the Receiver responds by choosing the announced state $a_R(m, m) = m$ if both Senders report some common message $m = \theta$. Now suppose that one of the Senders, say $j = 2$, considers whether or not to deviate from telling the truth, provided that the opponent is telling the truth, $m_1 = \theta$. Then, for each given state θ , Sender 2 has no incentives to deviate if the Receiver responds by harming her more under any $m_2 \neq m_1 = \theta$ than under $m_2 = \theta$. A simple way of completing this equilibrium proposal then is by making the Receiver to choose the worse possible action for any Sender $a_R(m_1, m_2) = a_R(0)$ whenever $m_1 \neq m_2$. Since the Senders' biases are positive, any Sender prefers the Receiver to pick the action that coincides with the true value of the state rather than choosing $a = 0$.

Krishna and Morgan (2001)'s analysis shows the logic of how competition can improve communication when competing Senders report simultaneously.¹² Their equilibrium, though, is not robust to perturbations of the game. For instance, if Senders do not obtain identical information or if the Receiver hears the messages with some noise, then the proposed equilibrium breaks down. Less fragile revealing equilibria can be constructed with several Senders when cheap-talk is multi-dimensional. Battaglini (2002) and (Ambrus and Takahashi (2008)) provide conditions under which fully revealing equilibria exists for multi-dimensional settings with several conflicting Senders. Specifically, Battaglini (2002) considers a setting where $a, \theta \in \mathbb{R}^n$, each Sender j wishes to minimize the distance $\|a - (\theta + b_j)\|$, and the Receiver wishes to minimize the distance $\|a - \theta\|$. In this case, the Receiver and each Sender j share common interests on the hyperplane specified as

$$\left\{ x \in \mathbb{R}^n : \frac{1}{\|b_j\|} b_j \cdot x = c \right\} \quad \text{where } c \in \mathbb{R}.$$

Battaglini (2002) then shows that if two Senders $j = 1, 2$ have different biases, $b_1 \neq b_2$, then it is possible to construct an equilibrium where the Receiver becomes fully informed. Again, let $a_R(m_1, m_2)$ be the action chosen by the Receiver when she hears message m_j from Sender j . Then,

¹² Note that the mechanism the underlies the proposed revealing equilibrium does not work if Senders report sequentially instead.

in a fully revealing equilibrium, each Sender j optimally announces $m_j = \theta_j$ (with probability one) and the Receiver has an optimal best-reply

$$a_R^*(m_1, m_2) = \begin{cases} \theta_1 & \text{if } \theta_1 = \theta_2; \\ z & \text{if } \theta_1 \neq \theta_2, \end{cases}$$

where $z \in \mathbb{R}^n$ is any point

$$z \in \left\{ x \in \mathbb{R}^n : b_1 \cdot x = \frac{\|b_2\|}{\|b_1\|} b_1 \cdot \theta_1 \right\} \cap \left\{ x \in \mathbb{R}^n : b_2 \cdot x = \frac{\|b_1\|}{\|b_2\|} b_2 \cdot \theta_2 \right\}.$$

Recall that each Sender j wishes to minimize the distance

$$\begin{aligned} \|a - (\theta + b_j)\| &= +\sqrt{((a - \theta) - b_j) \cdot ((a - \theta) - b_j)} \\ &= +\sqrt{b_j \cdot b_j + (a - \theta) \cdot (a - \theta) + 2[b_j \cdot \theta - b_j \cdot a]}. \end{aligned} \tag{8.3}$$

Then, in the situation where both Senders truthfully reveal the state $m_j = \theta$, each Sender j suffers a loss $\|b_j\| = +\sqrt{b_j \cdot b_j}$. Given this, if one of them, say $j = 2$, deviates to choosing any message $m_2 \neq m_1 = \theta$, we observe from the specification of R 's optimal strategy and from the expression in (8.3) that she receives even a higher loss because it must be the case that either $b_1 > b_2$ or $b_2 > b_1$. Specifically, $[b_2 \cdot \theta - b_2 \cdot a] > 0$ so that

$$\|a_R^*(\theta_1, m_2) - (\theta + b_2)\| > +\sqrt{b_2 \cdot b_2} = \|a_R^*(\theta_1, \theta_2) - (\theta + b_2)\|.$$

Figure 8.2 illustrates how fully revealing equilibrium can be constructed for multi-dimensional cheap-talk situations with two Senders. The circumferences in red and blue represent indifference curves for the Senders. Suppose that the true state is $\bar{\theta}$. Then, the Receiver wishes to pick action $a_R(\bar{\theta}) = \bar{\theta}$, and each Sender j 's most preferred action is $\bar{\theta} + b_j$. The biases are depicted in blue and red arrows. Then, a fully revealing equilibrium is constructed by considering that both Senders announce $\bar{\theta}$ as the true state and the Receiver responds by choosing action $a = \bar{\theta}$. Both Senders would then receive payoffs lower than their optimal payoffs. However, things can get even worse

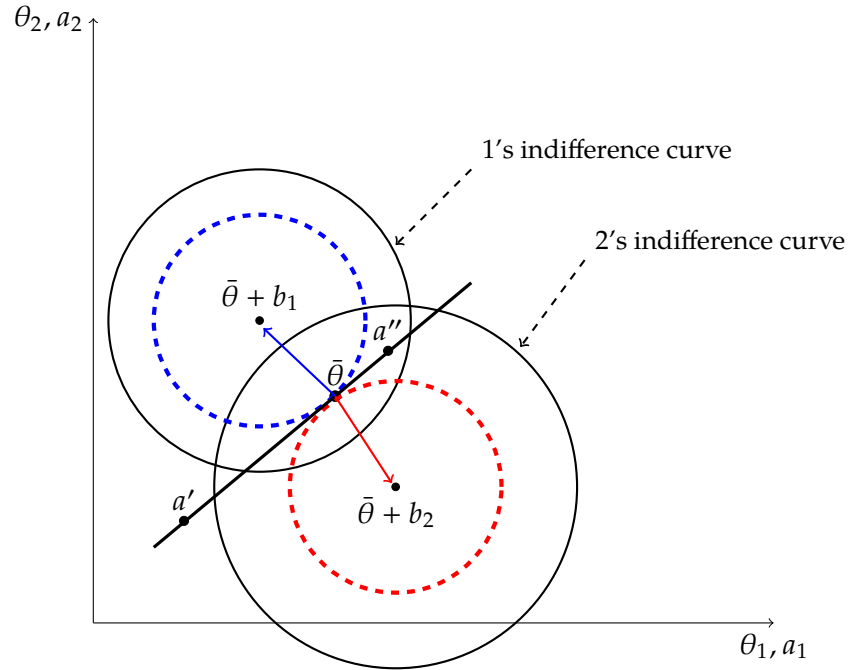


Figure 8.2 – Revealing Equilibria: Two-Dimensional Cheap-Talk with Two Senders

for a Sender that deviates unilaterally, as depicted in the figure. In particular, we can make the Receiver to respond, upon receiving any message $m_j \neq \bar{\theta}$, with actions such as a' or a'' . Clearly, the any deviating Sender $j = 1, 2$ would be worse off when the Receiver picks action a' , relative to the case where the Receiver chooses action $\bar{\theta}$.

Despite these interesting attempts to rationalizing communicative behaviors, the general insight provided by the standard cheap-talk setting, where effective communication is severely restricted, continues to hold even for the above mentioned variations. In particular, even for multi-dimensional settings, cheap-talk communication is restricted if either the state space is bounded (Ambrus and Takahashi (2008)) or if the dimensions of uncertainty are strongly correlated between them according to the prior (Levy and Razin (2007)).

8.1.4. (Bayesian) Persuasion

A feature of dynamic asymmetric information games where an informed speaker communicates with an uninformed decision maker is that the Sender is ex-ante uninformed, exactly as it is the Receiver. The only information that both players (commonly) know at the ex-ante stage is the prior

distribution of the payoff-relevant state. It is at the interim stage when the Sender fully learns the true value of the state, and yet at this point the Receiver still remains with the priors. Given this, a crucial assumption of cheap-talk models is that, at the interim stage and after learning the true value of θ , the Sender decides which message m she sends to the Receiver. In other words, the Sender announces publicly a relationship between θ and m that she has privately chosen only after learning the value of θ . The Receiver then uses the relation that the Sender establishes between θ and m to update the priors (following Bayes's rule) and picks her optimal action a . Since the Sender establishes her optimal relation between θ and m conditional on having learned the true value of the state, she has an "informational advantage" over R at that moment.

Consider now a different plausible environment where the Sender instead selects the relation between θ and m at the ex-ante stage, before knowing the true value of θ , and then commits to such a relationship despite of what she might learn subsequently at the interim stage. This assumption can be interpreted as if S is potentially informed and then chooses to commit to public information acquisition and disclosure. Since the Sender chooses a relation between θ and m for each possible value of θ , this information choice is best viewed as acquiring an *information structure*. A typical example would be that of designing an experiment that releases subsequently information about θ that cannot then be concealed or distorted. In this way, when she makes her information decision, S intentionally chooses to "tie her hands" and commits to disclose publicly every bit of information that the experiment provides. It is the commitment assumption what makes this problem essentially different from cheap-talk models. Intuitively, we see that now S has no "informational advantage" over R at the moment when she picks her optimal strategy. One may then wonder whether this informational leverage may mitigate the conflict of interests between the players and, therefore, allow for more effective communication relative to the cheap-talk benchmark. It turns out that, under certain conditions, this is indeed the case.

The benchmark described is known as *Bayesian persuasion*.¹³ An important advantage of Bayesian persuasion over cheap-talk models is that, by choosing the information structure, the Sender (1) decides exactly how posteriors are obtained and (2) implicitly explains the meaning of

¹³Kamenica and Gentzkow (2011) receive most of the credit for the study of these models. Rayo and Segal (2010)'s paper explored similar ideas and appeared at approximately the same time.

each message to R . In this way, the sort of indeterminacy typically present in cheap-talk games due to the endogenous meaning of messages is eliminated. Note that language continues to be determined endogenously in equilibrium but now this is done only by S , not by both players simultaneously. Intuitively, when designing the experiment, the Sender tells the Receiver how to interpret the messages that the experiment releases. A drawback of the Bayesian persuasion approach relative to the cheap-talk framework, though, is the suitability of the strong commitment assumption to certain real-world environments. Also, assuming that an information structure, or experiment, can reveal any information about the state implies making strong assumptions about the technology behind the experiment/communication device.

The details that we now present are based on [Kamenica and Gentzkow \(2011\)](#)'s exposition. We use the notation introduced earlier for the benchmark cheap-talk model, except that we now also allow for the Sender to randomize between messages. Given a set of states of the world $\theta \in \Theta$ and a set of messages $m \in M$, player S acquires an information structure by picking a collection $\sigma \equiv \{\sigma(\cdot | \theta) \in \Delta(M) : \theta \in \Theta\}$ of *signals*. A signal $\sigma(\cdot | \theta)$ is a probability distribution where $\sigma(m | \theta)$ indicates the probability that the experiment delivers message m when the true value of the state is θ . Then, each given information structure σ chosen by the Sender induces on the Receiver a (Bayesian) posterior distribution $\mu(\cdot | m) \in \Delta(\Theta)$ for each $m \in M$. It turns out quite convenient here to denote posteriors as $\mu_m \equiv \mu(\cdot | m)$ just for expositional reasons. Since each message m is heard with a certain probability, according to the prior and the picked information structure, it follows that each information structure induces a distribution over posterior distributions μ_m . Specifically the probability that message m be delivered by the experiment—and, therefore, that the posterior distribution μ_m be obtained—is described by $g(m) = \int_{\Theta} \sigma(m | \theta) f(\theta) d\theta$. Then, clearly any distribution $g(m)$ over posteriors induced by an information structure must satisfy the following “Bayesian plausibility” condition: $f(\theta) = \int_M \mu_m(\theta) g(m) dm$ for each $\theta \in \Theta$. In other words, induced posteriors must coincide “in average” with the prior.¹⁴ The Receiver chooses, for each particular posterior distribution μ_m , an action $\hat{a}_R(\mu_m)$ that maximizes her expected utility,

¹⁴This “Bayesian plausibility” condition is just an adding up constraint. In fact, any distribution of posteriors such that the induced posteriors are derived from the priors using Bayes’ rule satisfy by construction such a condition.

that is,

$$\widehat{a}_R(\mu_m) \in \arg \max_{a \in A} \int_{\Theta} \mu_m(\theta) u_R(a, \theta) d\theta.$$

Then, S 's decision problem consists of choosing an information structure σ —or, equivalently, of inducing a system of beliefs $\mu \equiv \{\mu_m \in \Delta(\Theta) : m \in M\}$ —that maximizes her ex-ante expected utility, under the restriction of the “Bayesian plausibility” condition. That is, the Sender wishes to solve the problem

$$\begin{aligned} \max_{\sigma} \int_M g(m) \left[\int_{\Theta} \mu_m(\theta) u_S(\widehat{a}_R(\mu_m), \theta) d\theta \right] dm \\ \text{s.t.: } f(\theta) = \int_M \mu_m(\theta) g(m) dm. \end{aligned}$$

To simplify notation, for a system of beliefs μ , let us denote

$$U_S(\mu_m) \equiv \int_{\Theta} \mu_m(\theta) u_S(\widehat{a}_R(\mu_m), \theta) d\theta.$$

In words, $U_S(\mu_m)$ is the interim expected utility to the Sender, conditional on message m , when S picks an information structure that induces the posteriors μ_m on the Receiver, and the Receiver responds optimally given such posteriors. An information structure σ delivers information that S would like to share if it induces (with positive probability $g(m) > 0$) a posterior distribution μ_m that is better (at the interim stage) for S than keeping the prior f . Algebraically, S wishes to publicly share the information provided by an experiment σ at the interim stage for some $m \in M$ with $g(m) > 0$ if

$$U_S(\mu_m) > \int_{\Theta} \mu_m(\theta) u_S(\widehat{a}_R(f), \theta) d\theta.$$

As to the Receiver, any information structure that delivers information different from the priors is valuable to the Receiver. Notice that any posterior $\mu_m \neq f$ allows the Receiver to optimally choose her action more suitably than f .

Given these details, the interesting result provided by [Kamenica and Gentzkow \(2011\)](#) works

as follows. Recall that S 's objective function is

$$V(\mu_m) \equiv \int_M g(m)U_S(\mu_m)dm,$$

so that the sender wishes to publicly share the information provided by an experiment σ at the ex-ante stage if

$$V(\mu_m) \geq V(f) = \int_M g(m)U_S(f)dm = U_S(f).$$

On the other hand, we know that the function $U_S(\mu_m)$ is convex in μ_m if

$$\int_M g(m)U_S(\mu_m)dm \geq U_S\left(\int_M g(m)\mu_m dm\right) = U_S(f),$$

where the last equality above follows from the Bayesian plausibility condition. Almost by definition, therefore, convexity of $U_S(\mu_m)$ is the key feature that characterizes when the Sender wishes to commit to publicly disclose an information structure more informative than the priors. Building on this insight, [Kamenica and Gentzkow \(2011\)](#) then provide a mathematical characterization of the Sender's optimal information structure choice that relates the solution of its problem to the smallest concave function greater than or equal to $U_S(\mu_m)$. Such a function of $U_S(\mu_m)$ is by definition known as the *concave closure* of $U_S(\mu_m)$.¹⁵

8.1.5. Verifiable Information

Bayesian persuasion enhances communication relative to cheap-talk because of the key commitment assumption. Without the commitment assumption, Senders may have incentives to truthfully reveal their private information when there are exogenous restrictions to the messages available to them. In particular, [Grossman \(1981\)](#), [Milgrom \(1981\)](#), and [Milgrom and Roberts \(1986\)](#) demonstrate that fully revealing equilibrium exist if Senders are restricted to provide information that is verifiable. The typical game explored to study communication when information is verifiable considers that the set of states $\Theta = \{\theta_1, \dots, \theta_K\}$ is finite and ordered, and that the set of available

¹⁵ Formally, the concave closure of $U_S(\mu_m)$ is the supremum $\sup\{y : (\mu, y) \in co(U_S(\mu_m))\}$, where $co(U_S(\mu_m))$ is the convex hull of the graph of $U_S(\mu_m)$.

messages M is the set of all subsets of Θ , that is, $M = 2^\Theta$. Suppose that S 's preferences do not depend on θ and that she always prefers R to take higher actions. In most applications, states correspond to qualities of a good and actions correspond to the price the Receiver is willing to pay for such a good. While the Receiver is assumed to wish to match the price with the quality of the good, the Sender wishes that the Receiver chooses the highest price. The game proceeds as follows. Player S learns the true value θ of the state and then reports a message m under the constraint that $\theta \in m$. Given this, player R uses S 's message to update her priors (following Bayes' rule) and takes an action.

In this game, equilibrium behavior features "skepticism" by R and S 's optimal strategy allows R to learn the true value of the state. To see how the mechanism behind this kind of equilibria works, assume that R regards a message $m \in M = 2^\Theta$ as indicating that θ is the lowest element in m (skepticism). Then, notice that the highest type θ_K of S wishes to separate herself by choosing $m = \{\theta_K\}$. Since R takes action $a = \min_\theta m$ in this equilibrium proposal, S 's type θ_{K-1} wishes to send the message $m' = \{\theta_{K-1}\}$. The Sender would not receive a higher payoff by deviating from this equilibrium proposal. Notice that the Sender would obtain a lower payoff if she deviates to sending a message m'' that includes types lower than θ_{K-1} as, in this case, $\min_\theta m'' < \theta_{K-1}$. Also, the Sender would not obtain a higher payoff by deviating to a message that contains types higher than θ_{K-1} . By iterating this argument from higher to lower types, it follows that truthful communication "unravels" from the top of the set of states in this kind of equilibria.

8.2. Mechanism Design

In many interactions under uncertainty with conflicting players, it is usually the case that some players have no control over how their opponents align their actions with the state of the world but, at the same time, they enjoy some power to choose or affect (at least some of the) rules of the game in which the players are involved. Prominent examples where a player cannot control how her opponents' actions are aligned with the state happen because the player cannot observe directly either the state or her opponents' actions. We have already discussed some of these examples: buyers cannot observe the quality of a good, firms cannot observe the productivity of

job candidates, patients cannot observe the medical biases of doctors, insurance companies cannot observe the initial health of the insured, or their disease prevention habits, shareholders cannot observe the productivity of the company's executives, or their financial decisions, firm's owners cannot observe the decisions of their managers, governments cannot observe the actual production costs in a natural monopoly, or the monopolist's efforts to reduce such costs. These situations are particularly relevant because, as we have already seen for the case of adverse selection, they typically lead to socially inefficient outcomes. In the above mentioned examples, we also notice that the uninformed party can, to some extent, affect the rules of the game. Customer organizations can promote rules to regulate markets, firms can decide how they evaluate their candidates, insurance companies can decide over their insurance policies and contract terms, shareholders can choose how their executives are evaluated and remunerated, governments can regulate natural monopolies. In short, it is usually the case that one party has an "information advantage" over the other while, at the same time, the latter holds some "legal rights" over the former. There is a broad, and systematically structured, approach that captures this sort of situations and seeks to deal with the socially undesirable outcomes caused by informational asymmetries. This approach is known under the synonymous terminologies of *mechanism design*, *contract theory*, or *institutional design*.¹⁶ The basic rationale behind the *mechanism design* approach is that the party with the power to "design the rules of the game" does so in a way such that makes it beneficial for the informed party not to deviate much from designer's objectives. In this way, by designing the rules of the game (or, synonymously, mechanism or contract), the designer imposes additional incentives over the informed party that make the final incentives of both parties somehow more compatible than they were initially. Not surprisingly, these rules are generally known as *incentive compatibility conditions*. Importantly, by reducing the final discrepancies in incentives, incentive compatibility conditions help alleviate the initial inefficiencies of the interactions.

¹⁶ Because of their historical developments, there are some subtle differences between the focus of mechanism design and contract theory models. Their substances, or more philosophical reasonings, though, are indistinguishable similar. These notes considers the three terms as synonymous.

8.2.1. Mechanism Design with Adverse Selection (or Hidden Types)

When conflicting players interact and some of them hold more information than others, the better informed ones naturally have incentives to hide such information, or to lie about it. By doing so, they typically seek to benefit from their informational advantage. We have already explored this issue when players have some private information about an exogenous payoff-relevant characteristic, modeled through types or states of the world, in the context of adverse selection. We have also seen how communication, in the form of Sender-Receiver games, can somehow mitigate the inefficiencies that arise in these situations. In fact, the sort of instruments that we already introduced through Sender-Receiver games can be regarded as special cases of the broad mechanism design approach. When considering Sender-Receiver games, it is usually the case that one pays more attention to positive features, such as the description of equilibrium and its properties. When pursuing the classical mechanism approach, one is typically more concerned about normative features, addressing directly inefficiency issues. Rather than the essence of their settings and functioning, it is a certain distinction in focus/perspective what makes strategic communication models somewhat different from mechanism design models.

There is nonetheless a plethora of (seemingly) similar models that deal with how a party may affect the incentives of others who held private information, through mechanisms or contracts, to mitigate the initial conflict of interests and achieve socially desirable goals. The distinctions between the models are mainly based on who controls what, who communicates with whom, or what can be communicated. For instance, Bayesian persuasion models and mechanism design models are tightly connected. Let us summarize some of these interrelated literatures and present their differences.

Bayesian Games with Communication

A strand of the mechanism design literature that deals with the incentives of informed players to reveal publicly their information considers that the “designer” is a *social planner* that has no preferences over the actions chosen by the players and that it is only interested in achieving socially desirable outcomes. Unlike the Bayesian persuasion models, the designer cannot provide

the players with additional information about the relevant state but it can affect the “mechanism” under which the players choose their actions. These models are known as Bayesian Games with Communication.

Let us go back to the formulation of Bayesian Games seen in Section 5. We will enlarge a little bit the model explored then. Also, to work with a notation similar to that developed earlier in Sender-Receiver games, consider that each player $i \in N = \{1, \dots, n\}$ must choose a payoff-relevant action $a_i \in A$ and that the players care about an unknown state of the world $\theta \in \Theta$ as well. The payoff to each player i is given by a utility $u_i(a, \theta)$, where $a = (a_i)_{i \in N}$ is an action profile. The players commonly know a prior distribution $p \in \Delta(\Theta)$ about θ . Suppose that for each state realization $\theta \in \Theta$, each player i privately learns a type $t_i \in T_i$ that describes the information that the player initially has about θ .

One way of making the “states of the world” and “players’ types” formulations compatible is that of considering a mapping $\tau : \Theta \rightarrow T$, with $T = \times_{i \in N} T_i$, that specifies a type profile $t = (t_1, \dots, t_n) = (\tau_1(\theta), \dots, \tau_n(\theta)) = \tau(\theta)$ for each state realization θ . Assume that: (1) $\tau_i(\theta) \neq \tau_i(\theta')$ implies $\theta \neq \theta'$ for each $i \in N$, so that no different types of any player can follow from a common state of the world, and (2) $\theta \neq \theta'$ implies $\tau(\theta) \neq \tau(\theta')$.¹⁷ In this context, a standard Bayesian game that captures how the players choose their actions consists of $\Gamma = \langle N, A, T, \pi, (u_i)_{i \in N} \rangle$, where the joint probability distribution over types π is related to the prior over states p according to the function τ as follows:

$$\pi(\tau(\theta)) \equiv p(\theta) \quad \text{for each } \theta \in \Theta.$$

The prior p , the set of type profiles T , and the function τ define an *information structure* in this setting. Intuitively, an information structure describes what each player knows about the state, what each player knows about what each other player knows about the state, and so on.

Now, following Myerson (1982) and Forbes (1986),¹⁸ consider a modification of this game where, prior to choosing her action, each player i is allowed to report confidentially a message $m_i \in M_i \equiv T_i$ about her type to an external mediator. The players report their messages simultaneously to the

¹⁷In other words, τ is assumed to be a one-to-one function. It is not required, though, that each particular τ_i be one-to-one, that is, different states of the world can provide the the same type to any given player.

¹⁸See also Myerson (1991), Chapter 6.3.

mediator who then submits a private recommendation $r_i(m)$ to each player i , based on the profile of messages $m = (m_1, \dots, m_n)$ received from the players. Finally, each player i chooses her action based on the recommendation $r_i(m)$ that she hears from the mediator. Let $r = (r_i)_{i \in N} \in A$ be a profile of recommendations by the mediator and let $\sigma(r | m)$ be the probability that the mediator recommends action profile r when it receives a message profile m from the players. The conditional probability $\sigma(r | m)$ specifies a *mechanism* in this environment. Importantly, the players' types are assumed to be unverifiable by the mediator so that any player can lie about her type. The mediator is not modeled as a player with incentives in this game.

Notice that, rather than a model of communication between the different players, as we explored in Sender-Receiver games, we are considering here a model of (private bilateral) communication between each player and some external (quite metaphorical) mediator. Also, only the player can provide the mediator with additional information as the mediator has no information of its own to share with the players.

Formally, we are enlarging the original Bayesian game Γ into another Bayesian game $\widehat{\Gamma}$ with communication. The new Bayesian game $\widehat{\Gamma}$ is specified as

$$\widehat{\Gamma} = \langle N, \widehat{S}, T, \pi, (\widehat{u}_i)_{i \in N} \rangle .$$

Here, $\widehat{S} = \times_{i \in N} \widehat{S}_i$ is the set of the players' pure strategy profiles and $(\widehat{u}_i)_{i \in N}$ is the list of their utility functions. More specifically, the set \widehat{S}_i of pure strategies for each player i in this enlarged Bayesian game is $\widehat{S}_i = \{(m_i, \delta_i) : m_i \in T_i, \delta_i : A_i \rightarrow A_i\}$, where m_i is a message that player i reports to the external mediator (with the interpretation that $m_i = t_i$ means "my type is t_i ") and $\delta_i(r_i)$ is an action chosen by player i upon hearing recommendation r_i from the mediator. Each function $\delta_i : A_i \rightarrow A_i$ is known as a *decision rule*. Let $\delta \equiv (\delta_i)_{i \in N}$ be a profile of decision rules. In addition, each utility function $\widehat{u}_i : \widehat{S} \rightarrow \mathbb{R}$ for the modified Bayesian Game is specified as

$$\widehat{u}_i(m, \delta | t) \equiv \sum_{r \in A} \sigma(r | m) u_i(\delta(r), \tau^{-1}(t)).$$

This is the expression for player i 's (ex-post) expected utility when the players report message

profile m to the mediator and choose a profile of decision rules δ , conditional on the the players having a type profile t . Since τ is one-to-one, we have that $\theta = \tau^{-1}(t)$ gives us the unique state of the world that corresponds to the type profile t . However, players do not actually know the types of their opponents. Therefore, to explore the BNE of this game, we need to consider the (interim) expected utility of each player:

$$\begin{aligned}\widehat{U}_i(m, \delta | t_i) &= \sum_{t_{-i} \in T_{-i}} \pi_i(t_{-i} | t_i) \widehat{u}_i(m_i, m_{-i}, \delta_i, \delta_{-i} | t_i, t_{-i}) \\ &= \sum_{t_{-i} \in T_{-i}} \sum_{r \in A} \pi_i(t_{-i} | t_i) \sigma(r | m_i, m_{-i}) u_i(\delta_i(r_i), \delta_{-i}(r_{-i}), \tau^{-1}(t_i, t_{-i})).\end{aligned}$$

Of course, the appropriate equilibrium notion to solve this enlarged Bayesian game is that of BNE. We then assume that the players act according to the BNE requirements that, conditional on the private information described by each type t_i , each player i chooses her message m_i and obeys the mediator according to the decision rule $\delta_i(r_i)$ in a way such that, given the (optimal) choices of the other players, i has no (strict) incentives neither to report any other message nor to disobey the mediator's recommendation.

In general, these equilibrium conditions give us many BNEs. As mentioned earlier, the mechanism design approach is interested mainly in normative goals. To this end, let us now pay attention to a subset of BNEs for the proposed modified game where each player in fact reports truthfully, $m_i = t_i$ for each $t_i \in T_i$, and follows the mediator's recommendation, $\delta_i(r_i) = r_i$ for each $r_i \in A_i$. By doing so, each player would reveal her private information to the others in equilibrium, thus eliminating the initial informational asymmetries and allowing for a socially efficient outcome. Then, taking this sort of BNE as the attainable goal of a social planner, the mechanism design approach asks what mechanisms $\sigma(r | m)$ would induce the players to play according to such BNEs. We then say the the decision rule or mechanism σ is *incentive compatible* if it induces a BNE of the Bayesian game $\widehat{\Gamma}$ with communication where each player reports honestly her type and follows the mediator's recommendation. Formally, σ^* is *incentive compatible* if for each type $t_i \in T_i$ of each

player $i \in N$, we have:¹⁹

$$\widehat{U}_i(t_i, t_{-i}, r_i, r_{-i} | t_i) \geq \widehat{U}_i(m'_i, t_{-i}, \delta_i(r_i), r_{-i} | t_i) \quad \forall m'_i \in T_i \quad \forall \delta_i(r_i) \in A_i,$$

or, equivalently,

$$\begin{aligned} \sum_{t_{-i} \in T_{-i}} \sum_{r \in A} \pi_i(t_{-i} | t_i) \sigma^*(r | t_i, t_{-i}) u_i(r_i, r_{-i}, \tau^{-1}(t_i, t_{-i})) \geq \\ \sum_{t_{-i} \in T_{-i}} \sum_{r \in A} \pi_i(t_{-i} | t_i) \sigma^*(r | m'_i, t_{-i}) u_i(\delta_i(r_i), r_{-i}, \tau^{-1}(t_i, t_{-i})) \quad \forall m'_i \in T_i \quad \forall \delta_i(r_i) \in A_i. \end{aligned}$$

If σ^* is incentive compatible according to the condition above, we also say that σ^* *implements* (pure strategy) BNEs where the players reveal their private information and follow the mediator's recommendations.²⁰

Notably, even when one restricts attention to such (socially desirable) BNEs where $m_i = t_i$ and $\delta_i(r_i) = r_i$ for each player i , there is a profound multiplicity of BNEs that can be implementable. Pretty much like in the case of cheap-talk games, this is due to the fact that messages have no intrinsic meaning here and, in particular, each player can reveal completely her type to the mediator by picking a different message for each different type realization. Fortunately, unlike cheap-talk situations, for a given selected mechanism, here the players alone are the ones who determine the amount of information transmitted. Since the designer receives no payoffs in this setting, the amount of information revealed is not determined simultaneously by Sender and Receiver, only by the Sender. This simplifies the equilibrium indeterminacy problem, relative to cheap-talk models, and allows for a set of convenient results commonly known under the term *revelation principle* arguments. The Revelation principle was first proposed by [Aumann \(1974\)](#), in the context of his CE notion proposal, and subsequently by [Myerson \(1991\)](#) (Chapter 6.3), for the case of Bayesian games with communication we are presenting now. For Bayesian games with communication, the *revelation principle* works as follows. Suppose that we consider a more general communication system between each player i and the mediator where the player can send messages $m_i \in M_i \neq T_i$

¹⁹ Attention is restricted to pure strategy BNE in this definition of incentive compatibility.

²⁰ This strand of the mechanism design literature is also known as *Bayesian implementation*.

and the mediator can make recommendations $r_i \in R_i \neq A_i$. Consider a mechanism σ^* that is incentive compatible under such a broader communication system. Then, the revelation principle shows that there exists another mechanism σ^{**} which is also incentive compatible for the narrower communication system where $m_i \in T_i$ and $r_i \in R_i \neq A_i$, and that provides the player with the same expected utility as σ^* . In this way, there is no loss of generality in considering that the players use a language where they just communicate their types and the mediator directly communicates actions.

Finally, determining which mechanisms σ are incentive compatible depends crucially on the information structure described by (p, T, τ) . In practice, this is problematic since information structures are in general hard to observe. Interestingly, following the global games approach presented in Section 6.3, [Bergemann and Morris \(2013\)](#) propose a benchmark that allows one to provide conditions under which the players wish to follow the external mediator recommendations for *regardless of the information structure*.

Information Design

To fix ideas, recall that Bayesian persuasion considers situations where a Sender chooses information structures to provide a Receiver with additional information about the state of the world. To “homologate terms” with the approach that we now present, let us refer here to a Sender as an *information designer*.²¹ In the Bayesian persuasion literature, the designer has incentives of her own regarding the Receiver’s choice of actions and the state of the world. Nonetheless, this strand of the literature does not consider cases where several (possibly conflicting) Receivers interact among them in their actions. A crucial feature of this literature is that the designer can (ex-ante) commit to provide information about the state—in spite of the discrepancy between her ex-ante and interim objectives.

On the other hand, recall that Bayesian games with communication consider that the designer does not have incentives of her own regarding the choice of actions of the Receivers and the state of the world. In this literature, the designer cannot either provide Receivers with additional

²¹ In mechanism design models, it is also common to use the respective terms *Principal* for the Sender, or information designer, and *Agent* for the Receiver.

information about the state. However, this literature does cover situations where the Receivers interact among them in their choices of actions. Another key feature of this literature is that the designer can (ex-ante) commit to make recommendations based on the information disclosed by the Receivers but not on the realization of the state.

Intuitively, Bayesian persuasion explores situations where the designer chooses an experiment, or trial, that discloses publicly information about an unknown issue whereas Bayesian implementation studies situations where the designer chooses the rules of an interaction in ways such that the final outcome crucially depends on how private information is revealed. A typical example of the latter class of situations is that of a public auction where its rules induce the bidders to reveal their true valuations for the auctioned good.

There is also a recent, broader approach, that covers both situations where the designer can provide Receivers with additional information about the state and the Receivers strategically interact among them in their action choices. [Bergemann and Morris \(2016\)](#) coin this class of problems as *information design* problems. As in the Bayesian persuasion approach, a crucial assumption of the information design approach is that the designer can (ex-ante) commit to providing the Receivers with additional information about the state of the world. Interestingly, rather than focusing on induced posteriors, the information design approach relies on action recommendations and, specifically, it makes use of an extension of the CE notion used in games with complete information to incomplete information settings, which is termed as Bayesian Correlated Equilibrium.

The following model of information design follows [Bergemann and Morris \(2017\)](#)'s exposition. There is finite set Θ of states θ . There is a set $N = \{1, \dots, n\}$ of players and each player i must choose an action a_i from a finite set of actions A_i . The payoff to each player i is given by a utility function $u_i(a, \theta)$, where $a = (a_i)_{i \in N}$. The state θ is picked according to a prior probability distribution $p(\theta) > 0$. There is also an information designer, whose payoffs are described by a utility function $v(a, \theta)$, that can (ex-ante) commit to provide the players with additional information about θ . Notice that, other than the fact that now there are several players that interact among them in their choice of actions, the setting is similar to the one presented earlier to describe Bayesian persuasion.

An *information structure* is described the prior p , by a finite set T_i of types t_i for each player i , and by a collection of (conditional) probability distributions $\{\pi(\cdot | \theta) \in \Delta(T)\}_{\theta \in \Theta}$ over type profiles $t = (t_i)_{i \in N} \in T = \times_{i \in N} T_i$. Thus, $\pi(t | \theta)$ is the probability that the players have type profile t when the realization of the state is θ . As in the case of Bayesian games with communication, an information structure (p, T, π) describes what each player knows about the state, what each player knows about what other players know about the state, and so on.

The game that describes the information design problem proceeds as follows. First, the information designer chooses and commits to a *decision rule*

$$\sigma \equiv \{\sigma(\cdot | t, \theta) \in \Delta(A)\}_{(t, \theta) \in T \times \Theta},$$

where each $\sigma(a | t, \theta)$ is the probability that the chosen decision rule delivers action recommendation a to the players when they have private information described by t and the true realization of the state is θ . By making a decision rule conditional on the state of the world, the designer can choose the amount of additional information that it provides to the players. Then, the true value θ of the state is realized and the players learn their types according to the conditional distribution $\pi(t | \theta)$. After this, the players receive the action recommendations that the decision rule $\sigma(a | t, \theta)$ discloses. Finally, the players pick their optimal actions, based on the prior and the action recommendations, and obtain accordingly their payoffs. Given this game, we say that the decision rule σ is *incentive compatible*²² if $\sigma^*(a^* | t, \theta) > 0$ for some $a^* = (a_i^*, a_{-i}^*)$ implies that, for each type $t_i \in T_i$ and for each player $i \in N$, we have

$$a_i^* \in \arg \max_{a_i \in A_i} \sum_{\theta \in \Theta} p(\theta) \sum_{t_{-i} \in T_{-i}} \pi(t | \theta) \sum_{a_{-i} \in A_{-i}} \sigma^*((a_i^*, a_{-i}^*) | t, \theta) u_i(a_i, a_{-i}, \theta). \quad (8.4)$$

In other words, if an incentive compatible decision rule σ^* recommends an action profile a^* with positive probability, then it must be the case that no player i has (strict) incentives to choose any other action $a_i \neq a_i^*$, given her type t_i . Using a revelation principle argument to ignore more complex communication systems and to solve equilibria indeterminacy issues, [Bergemann and](#)

²² [Bergemann and Morris \(2016\)](#) term this incentive compatibility condition as *obedience condition*.

Morris (2016) state that, without loss of generality, σ^* is a *Bayes Correlated Equilibrium* (BCE) if and only if it is incentive compatible. The condition (8.4) gives us a restriction that the decision rule chosen by the designer must satisfy in order for the players to actually follow her recommendation. To complete the description of the model, we need an objective function for the designer. The designer wishes to choose decision rules that maximize her (ex-ante) expected utility given that the players follow her recommendations. Specifically, the designer's problem is

$$\begin{aligned} \max_{\sigma} \quad & \sum_{\theta \in \Theta} p(\theta) \sum_{t \in T} \pi(t | \theta) \sum_{a \in A} \sigma(a | t, \theta) v(a, \theta) \\ \text{s.t.:} \quad & a_i^* \in \arg \max_{a_i \in A_i} \sum_{\theta \in \Theta} p(\theta) \sum_{t_{-i} \in T_{-i}} \pi(t | \theta) \sum_{a_{-i} \in A_{-i}} \sigma^*(a^* | t, \theta) u_i(a_i, a_{-i}^*, \theta). \end{aligned} \quad (8.5)$$

Let us now explore the similarities between the key incentive compatible condition in (8.4) and (1) the CE's key condition and (2) the optimality condition imposed on the Receiver in the Bayesian persuasion approach.

First, consider the complete information version of the game presented above, so that $\Theta = \{\bar{\theta}\}$, for some state value $\bar{\theta}$. Then, the set of types of each player i collapses into a trivial set $T_i = \{\bar{t}_i\}$, for some \bar{t}_i , as well. The information structure (p, T, π) in this particular case satisfies $p(\bar{\theta}) = 1$ and $\pi(\bar{t} | \bar{\theta}) = 1$. As a consequence, the incentive compatibility condition in (8.4) requires now

$$\sum_{a_{-i} \in A_{-i}} \sigma^*(a_i^*, a_{-i}^*) u_i(a_i^*, a_{-i}^*, \theta) \geq \sum_{a_{-i} \in A_{-i}} \sigma^*(a_i^*, a_{-i}^*) u_i(a_i, a_{-i}^*, \theta) \quad \forall a_i \in A_i$$

where $\sigma \in \Delta(A)$ is now a joint probability distribution over action profiles a . Then, by dividing the inequality above over the (marginal) probability $\sum_{a_i \in A_i} \sigma^*(a_i, a_{-i}) > 0$, the required condition can be rewritten as

$$\sum_{a_{-i} \in A_{-i}} \sigma^*(a_i^* | a_{-i}^*) u_i(a_i^*, a_{-i}^*, \theta) \geq \sum_{a_{-i} \in A_{-i}} \sigma^*(a_i^* | a_{-i}^*) u_i(a_i, a_{-i}^*, \theta) \quad \forall a_i \in A_i,$$

a condition formally identical to the one derived earlier in the definition of CE.

Secondly, consider the one-player version of the game presented above ($n = 1$ so that $i = R$) and, in consonance with the tools typically developed to study Bayesian persuasion models,

suppose that the information designer chooses a collection $\sigma \equiv \{\sigma(\cdot | \theta) \in \Delta(M) : \theta \in \Theta\}$ of *signals*. Again, a signal $\sigma(\cdot | \theta)$ is a probability distribution where $\sigma(m | \theta)$ indicates the probability that the experiment delivers message m when the true value of the state is θ . Then, each given information structure σ chosen by the Sender induces on the Receiver a posterior distribution $\mu_m \in \Delta(\Theta)$ over states θ , for each $m \in M$. Then, the (unique) Receiver chooses for each $\mu_m \in \Delta(\Theta)$ an action $\widehat{a}_R(\mu_m)$ that maximizes her expected utility $\sum_{a \in A} \mu_m(\theta) u_R(a, \theta)$. Then, following the benchmark model of Bayesian persuasion, the information designer must choose an information structure σ that maximizes her ex-ante expected utility, under the restriction of the “Bayesian plausibility” condition:

$$\begin{aligned} \max_{\sigma} \quad & \sum_M g(m) \sum_{\Theta} \mu_m(\theta) v(\widehat{a}_R(\mu_m), \theta) \\ \text{s.t.:} \quad & f(\theta) = \sum_M \mu_m(\theta) g(m). \end{aligned}$$

Since the “Bayesian plausibility” condition is satisfied for any information structure σ that leads to a collection of posteriors $\{\mu_m\}$ through Bayes’ rule, let us simply focus on the objective function of the Sender’s problem. Using Bayes’ rule, such an (ex-ante) expected utility can be rewritten as

$$\begin{aligned} \sum_M g(m) \sum_{\Theta} \frac{\sigma(m | \theta) p(\theta)}{g(m)} v(\widehat{a}_R(\mu_m), \theta) \\ = \sum_{\Theta} p(\theta) \sum_M \sigma(m | \theta) v(\widehat{a}_R(\mu_m), \theta). \end{aligned}$$

This expected utility is analytically equivalent to the objective function of the designer’s problem specified in (8.5) for the particular case where players are uncertain only about the stage. As to the incentive compatibility condition, note that, by picking an action $\widehat{a}_R(\mu_m)$ that maximizes her expected utility $\sum_{a \in A} \mu_m(\theta) u_R(a, \theta)$, the Receiver is, equivalently, choosing an action to maximize

$$\sum_{a \in A} \frac{\sigma(m | \theta) p(\theta)}{g(m)} u_R(a, \theta).$$

By dividing such an objective function by the positive term $g(m)$, it follows that the Receiver must

choose an action $\widehat{a}_R(\mu_m)$ that maximizes the objective $\sum_{a \in A} \sigma(m | \theta) p(\theta) u_R(a, \theta)$ for each $\theta \in \Theta$ and each $m \in M$ such that $g(m) > 0$. Therefore, by averaging over θ according to priors, we obtain that, for each given message m heard with positive probability, action $\widehat{a}_R(\mu_m)$ maximizes the objective function $\sum_{a \in A} \sigma(m | \theta) p(\theta) u_R(a, \theta)$ if and only if it maximizes the ex-ante expected utility

$$\sum_{\theta \in \Theta} p(\theta) \sum_{a \in A} \sigma(m | \theta) u_R(a, \theta).$$

Solving this problem is analytically equivalent to satisfying the incentive compatibility condition stated in (8.4) for the case where the players are initially uncertain only about the state θ .

We observe that Bayesian persuasion is a particular case of the information design approach when there is a single Receiver and, in addition, she is only initially uncertain about the state. Also, we can notice that the required incentive compatibility condition for the Receiver to be willing to choose her action after learning the information disclosed by the Sender is closely related to the CE solution concept presented in Subsection 4.6.

8.2.2. Mechanism Design with Moral Hazard (or Hidden Actions)

Another usual type of asymmetric, or hidden information, problems concern endogenously chosen actions rather than exogenous characteristics. In these cases, the less informed players are restricted in ways such that they cannot monitor perfectly the actions taken by the better informed ones. For example, insurance companies cannot monitor whether their insured are careful enough so as to prevent accidents, shareholders cannot monitor whether the firm's executive board makes the most profitable decisions, or land owners cannot monitor whether their hired farmers work diligently to obtain good harvests. When players interact and some of them hold more information than others about *the actions* that they choose, we refer to this kind of situations as *moral hazard*. Pretty much like the situations of adverse selection, moral hazard leads to socially undesirable, or inefficient, outcomes. Moral hazard can be viewed as another failure of markets caused by asymmetries of information between contracting parties. Together with adverse selection, moral hazard and its implications are central topics in the field of information economics.

The mechanism design approach is also useful to deal with the inefficiencies that arise with

moral hazard. One way of thinking about mechanism design in moral hazard environments is that of considering certain relationships between parties that are actually contractible in real world interactions. Such relationships are contractible because they can be naturally related to observable (or verifiable) outcomes. For instance, although shareholders cannot observe the decisions taken by the members of the executive board (or, e.g., whether such decisions follow the prescriptions of economic theory), they can observe the actual profits of the company at the end of the term. Likewise, although the insurance company cannot observe how careful a certain driver is, it can observe her record of accidents. A university cannot observe the effort exerted by its researchers but it can observe the number and qualities of their publications. Then, contracts can be made contingent on such type of observable (or verifiable) features. Once a particular set of relationships (which must be based on observables) is established, we can then model formally such a contract as a game and explore it using our game theoretical tools. In this way, a mechanism (or contract/institution) is defined as the rules of a game. Rather than solving games, though, the approach taken by mechanism design proceeds as follows. Given the constraints imposed to an interaction by the physical environment, or by the informational/cognitive endowment of some players, mechanism design considers a set of outcomes that are socially desirable and then ask what games will deliver such outcomes as solutions when the players interact. In this sense, mechanism design seeks to propose rules for games such that, once the involved parties play them, some socially desirable outcomes will be achieved. This logic behind mechanism design can be traced back historically to market socialism concerns where social philosophers realized that the existing markets were just one of some plausible mechanisms. Their proposal was then to consider other mechanisms that enabled more efficient outcomes for the sort of market interactions captured by the original mechanism. In practice, this is achieved through market regulation, or certain interventions, or modifications, to the rules under which such markets operate. The tricky exercise that mechanism design pursues then is how to propose the precise interventions that can lead to more efficient outcomes. The question is so good, and the taken approaches have proved so insightful, that this area of research has received immense recognition in economics.²³ Alongside

²³For instance, recognition in the form of Nobel awards has been paid to researchers that have contributed to mechanism design, either from more theoretical or more applied perspectives: Leonard Hurwicz, Eric Maskin, and

with their closely related literatures of information design and Sender-Receiver games, mechanism design is currently one of the most prolific fields of research in economics.

The classical model of moral hazard²⁴ considers two players, a *Principal* (or owner of a firm) and an *Agent* (or manager of the firm) that can potentially be involved in a certain project. The Principal hires the Agent to choose a certain action $a \in A = [\underline{a}, \bar{a}] \subset \mathbb{R}_+$, usually interpreted as an *effort level*. Each action a is a (pure) strategy for the Agent and it yields a probability distribution over the payoffs, or revenue, that the Principal receives from the project. Specifically, the revenue of the Principal is given by $q = q(a, \varepsilon) \in Q = [\underline{q}, \bar{q}] \subseteq \mathbb{R}_+$, where ε is a random variable with (commonly) known distribution. The Principal “legally owns the rights” to the revenue and can decide how to assign a part of it to the Agent. In this sense, the Principal enjoys a certain bargaining power to set up the project relationship (or contract) with the Agent. In most interesting cases, the action a chosen by the Agent is not observable by the Principal but the Principal can observe the revenue q . Therefore, any enforceable contract can be made contingent only on the revenue. Conditional on the observed revenue the Principal gives the Agent a payment $w = w(q)$, so that a (pure) strategy for the Principal is specified by a function $w : Q \rightarrow \mathbb{R}$. A strategy for the Principal $w(\cdot)$ is also known as a *payment scheme*. Use \mathcal{W} to denote the set of all feasible payment schemes.

The Principal and the Agent have a conflict of interests respect the Agent’s choice $a \in A$. The payoffs of the Agent are $u(w) - c(a)$, where $u(w)$ is the utility from payment w and $c(a)$ is the cost from exerting an effort level a . It is usually assumed that the functions u and c are twice-continuously differentiable with $u' > 0$ and $u'' < 0$, whereas $c' > 0$ and $c'' > 0$. That, is the Agent is assumed to be risk-averse and to face convex costs from her effort. The Principal’s utility is the revenue less the payment to the Agent, $v(a, w) = q(a, \varepsilon) - w$. Thus, the Principal is assumed to be risk-neutral.²⁵ Assume that the random variable ε is such that q is a continuous random variable

Roger Myerson (2007); Alvin Roth and Lloyd Shapley (2012); Jean Tirole (2014); and Oliver Hart and Bengt Holmström (2016).

²⁴ These notes present a continuous choice version of the classical model. Beyond technical differences, versions where decisions are taken over finite sets work similarly to the setting presented here.

²⁵ Risk-neutrality for the Principal is assumed for tractability but, more importantly, it is the most innocuous assumption regarding risk attitude we can impose in this environment. Since contracting on the effort level a will be unfeasible to the Principal, she will be harmed by the informational advantage that the Agent possesses. This constitutes the central efficiency message of the model. Considering risk-aversion would only exacerbate the damage to the Principal in this environment under uncertainty and, more crucially, asymmetric information.

with (conditional) density $f(q|a) > 0$ on Q . Assume also that $f(q|a)$ is twice-continuously differentiable with respect to a .

First Best Benchmark: Observable Actions

Consider first the (hypothetical) case where the Principal can in fact observe (or verify) the effort choice a . In this case, the contracting parties have symmetric information and one expects that the Agent be remunerated according to the marginal return from her effort. Recall that, under the logic that she holds the legal rights about what the project delivers, the Principal has the power of setting a contract and submitting it to the Agent, who must then decide whether to accept it or not. The contract is assumed to be legally enforceable under the terms that it states. Specifically, the problem that the Principal solves, to set up the contract, is to pick an effort level $a \in A$ and a payment scheme $w(\cdot) \in \mathcal{W}$ so as to maximize her expected utility under one condition. This condition is known as *Participation Constraint* (PC) and its goal is to ensure that the Agent is in the first place interested in signing up the contract. The PC condition requires that the expected utility of the Agent be no less than a certain exogenous minimum utility level \underline{U} . The utility level \underline{U} can be interpreted as the utility that the Agent can secure just by choosing an alternative outside option, rather than participating in the Principal's project. The Principal's goal is to solve

$$\begin{aligned} \max_{\{a, w(\cdot)\}} \int_{\underline{q}}^{\bar{q}} f(q|a)[q - w(q)]dq \\ \text{s.t. : } \int_{\underline{q}}^{\bar{q}} f(q|a)u(w(q))dq - c(a) \geq \underline{U}. \end{aligned} \tag{8.6}$$

The *Lagrangian* associated to the problem in (8.6) is

$$\mathcal{L}(a, w(\cdot); \lambda) = \int_{\underline{q}}^{\bar{q}} f(q|a) [q - w(q) + \lambda u(w(q))] dq - \lambda c(a) - \lambda \underline{U}.$$

Considering interior solutions, the (pointwise) F.O.C. with respect to $w(\cdot)$ gives us

$$\lambda = 1/u'(w(q)) \quad \text{for each } q \in Q.$$

This condition is known as the *Borch rule* and describes optimal risk sharing. In particular, it follows that the optimal payment scheme $w(q)$ must be constant, independently of the obtained revenue q . Also such a condition implies that the PC condition must be binding at the optimal payment scheme so that the Agent is indifferent between signing up the contract or choosing the outside option. Let us use the short-hand notation $f_a(q | a) = \partial f(q | a) / \partial a$ for the partial derivative of the conditional density of the revenue with respect to the action choice. Then the F.O.C. with respect to a yields

$$\int_{\underline{q}}^{\bar{q}} f_a(q | a) [q - w(q) + \lambda u(w(q))] dq = \lambda c'(a).$$

Let \tilde{w} be the constant payment that solves the equation $\lambda = 1/u(w)$. Then, by combining both F.O.C.s, we obtain the condition that characterizes optimal payment \tilde{w} and effort \tilde{a} in the First Best contract:

$$\tilde{w} = \int_{\underline{q}}^{\bar{q}} f_a(q | \tilde{a}) q dq + \frac{u(\tilde{w}) - c'(\tilde{a})}{u'(\tilde{w})}.$$

Intuitively, the optimal constant payment equals the expected revenue, conditional on the optimal effort, plus the (normalized) excess of the Agent' utility over her marginal cost.

Second Best Benchmark: Unobservable Actions

Consider now the (more realistic) case where the Principal cannot observe (or verify) the effort choice a and can only observe the revenue q that the project yields. The Principal “designs” the contract and offers it to the Agent who then decides whether or not to accept it. Again, the contract is assumed to be legally enforceable under the terms stated in it. The problem that the Principal now faces, when writing down the contract, is to pick an effort level $a \in A$ and a payment scheme $w(\cdot) \in \mathcal{W}$ so as to maximize her expected utility under the PC and an another, more interesting, restriction, known as *Incentive Compatibility* (IC). Following the same rationale that we explored for mechanism design with hidden types, the IC condition seeks that the Agent be interested in choosing her effort level conditional on the payment scheme that she receives from the Principal. Notice that we are considering an approach where one player, the Principal, optimally chooses the

strategies of both players. Algebraically, the Principal wishes to solve

$$\begin{aligned}
& \max_{\{a, w(\cdot)\}} \int_{\underline{q}}^{\bar{q}} f(q | a) [q - w(q)] dq \\
& \text{s.t. : } \int_{\underline{q}}^{\bar{q}} f(q | a) u(w(q)) dq - c(a) \geq \underline{U}; \quad (\text{PC}) \\
& a \in \arg \max_{a' \in [\underline{a}, \bar{a}]} \int_{\underline{q}}^{\bar{q}} f(q | a') u(w(q)) dq - c(a'). \quad (\text{IC})
\end{aligned} \tag{8.7}$$

As we already discussed with respect to the case of hidden types, the logic behind the mechanism design approach is that the designer puts herself in the Agent's position and discounts the Agent's optimal behavior under the mechanism that the designer establishes. In this way, starting from an ex-ante stage with conflicting interests, a mechanism seeks to align (to a certain extent) the player's interests at the interim stage through the (designed) incentives that it imposes on them.

Let us solve the problem in (8.7) by following an intuitive method known in the literature as the *first order approach* to the Second Best design problem.²⁶ This method considers that we solve first the Agent's decision problem stated in the IC condition. The F.O.C. to the Agent's problem is

$$\int_{\underline{q}}^{\bar{q}} f_a(q | a) u(w(q)) dq - c'(a) = 0. \tag{8.8}$$

We then can solve the Principal's problem subject to the PC constraint and to the condition derived in (8.8) above. The *Lagrangian* associated to such a problem is

$$\begin{aligned}
\mathcal{L}(a; \lambda, \mu) = & \int_{\underline{q}}^{\bar{q}} f(q | a) [q - w(q)] dq + \lambda \left[\int_{\underline{q}}^{\bar{q}} f(q | a) u(w(q)) dq - \lambda c(a) - \lambda \underline{U} \right] \\
& + \mu \left[\int_{\underline{q}}^{\bar{q}} f_a(q | a) u(w(q)) dq - c'(a) \right].
\end{aligned}$$

Then, by considering interior solutions, the (pointwise) F.O.C. with respect to $w(\cdot)$ gives us the

²⁶ Though intuitive, formal application of this method yet requires that we impose further conditions on the conditional density $f(q | a)$.

condition that characterizes the optimal contract for the Second Best scenario:

$$\frac{1}{u'(w(q))} = \lambda + \mu \frac{f_a(q|a)}{f(q|a)}. \quad (8.9)$$

To deliver its main insights, the classical model of moral hazard makes use of an intuitive and appealing assumption, known as the *monotone likelihood ratio condition (MLRC)*. Formally, this assumption requires that, for each $a \in A$, the likelihood ratio $L(q; a) \equiv f_a(q|a)/f(q|a)$ be non-decreasing in $q \in Q$. The intuition behind this condition is simply that higher effort levels are associated to higher (conditional) probabilities of achieving higher revenues, relative to lower effort levels. To put it bluntly, even though the presence of uncertainty due to the random component ε can make it possible for poor revenue realizations even when the Agent is doing her best, both players commonly agree on that higher effort levels are more likely to allow higher revenues than lower effort choices. Provided that $\lambda, \mu > 0$ in the Second Best scenario,²⁷ the condition obtained in (8.9), together with the MLRC assumption, leads to that the optimal payment $w(q)$ must vary with q and, in particular, it must be non-decreasing in q . Intuitively, the optimal contract agreed by both parties leads to that remunerations will be based on the observed revenue in a quite natural way. It is not only the case that the Agent will receive a higher payment because the “pie” is higher. More crucially, by knowing that the payment is positively related to the observed outcome of her effort, the Agent will be suitably incentivized (at the interim stage) so as to mitigate the profoundly misaligned incentives which are caused in this environment by having both conflicting preferences and different information about actions.

²⁷ This is not evident and requires formal proof.

8.3. Reputation

TO BE WRITTEN

8.4. Bargaining Games

TO BE WRITTEN

Bibliography

- AKERLOF, G. A. (1970): "The Market for "Lemons": Quality Uncertainty and the Market Mechanism," *The Quarterly Journal of Economics*, 84, 488–500.
- AMBRUS, A. AND S. TAKAHASHI (2008): "Multi-Sender Cheap Talk with Restricted State Spaces," *Theoretical Economics*, 3, 1–27.
- AUMANN, R. J. (1974): "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, 1, 67–96.
- BATTAGLINI, M. (2002): "Multiple Referrals and Multidimensional Cheap Talk," *Econometrica*, 70, 1379–1401.
- BERGEMANN, D. AND S. MORRIS (2013): "Robust Predictions in Games with Incomplete Information," *Econometrica*, 81, 1251–1308.
- (2016): "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, 11.
- (2017): "Information Design: A Unified Perspective," Tech. rep., Yale.
- CARLSSON, H. AND E. VAN DAMME (1993): "Global Games and Equilibrium Selection," *Econometrica*, 61, 989–1018.
- CHAKRABORTY, A. AND R. HARBAUGH (2007): "Comparative Cheap Talk," *Journal of Economic Theory*, 132, 70–94.
- (2010): "Persuasion by Cheap Talk," *American Economic Review*, 100, 2361–2382.
- CHEN, Y., N. KARTIK, AND J. SOBEL (2008): "Selecting Cheap-Talk Equilibria," *Econometrica*, 76, 117–136.
- COURNOT, A. (1929): *Researches into the Mathematical Principles of the Theory of Wealth*, New York: MacMillan.
- CRAWFORD, V. P. AND J. SOBEL (1982): "Strategic Information Transmission," *Econometrica*, 5, 1431–1451.
- FORBES, F. (1986): "An Approach to Communication Equilibrium," *Econometrica*, 54, 1375–1385.
- GROSSMAN, S. J. (1981): "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 24, 461–483.

- HARSANYI, J. C. (1967-68): "Games with Incomplete Information Played by 'Bayesian' Players," *Management Science*, 14, 159–182, 320–334, 486–502.
- (1973): "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points," *International Journal of Game Theory*, 2, 1–23.
- KAMENICA, E. AND M. GENTZKOW (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.
- KOHLBERG, E. AND J.-F. MERTENS (1986): "On the Strategic Stability of Equilibria," *Econometrica*, 54, 1003–1037.
- KREPS, D. M. AND R. WILSON (1982): "Sequential Equilibria," *Econometrica*, 50, 863–894.
- KRISHNA, V. AND J. MORGAN (2001): "Asymmetric Information and Legislative Rules: Some Amendments," *American Political Science Review*, 95, 435–452.
- KUHN, H. W. (1953): "Extensive Games and the Problem of Information," in *Contributions to the Theory of Games*, ed. by H. W. Kuhn and A. W. Tucker, Princeton University Press, vol. I.
- LEVY, G. AND R. RAZIN (2007): "On the Limits of Communication in Multidimensional Cheap Talk: A Comment," *Econometrica*, 75, 885–803.
- MILGROM, P. AND J. ROBERTS (1986): "Relying on the Information of Interested Parties," *RAND Journal of Economics*, 17, 18–32.
- MILGROM, P. R. (1981): "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, 12, 380–391.
- MORRIS, S. AND H. S. SHIN (2003): *Global Games: Theory and Applications*, Cambridge University Press, vol. 1 of *Econometric Society Monographs*, 56–114.
- MYERSON, R. B. (1982): "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems," *Journal of Mathematical Economics*, 10, 67–71.
- (1991): *Game Theory: Analysis of Conflict*, Cambridge, MA, London, England: Harvard University Press.
- NASH, J. F. (1951): "NonCooperative Games," *Annals of Mathematics*, 54, 289–295.
- OSBORNE, M. J. (2004): *An Introduction to Game Theory*, Oxford University Press.
- RAYO, L. AND I. R. SEGAL (2010): "Optimal Information Disclosure," *Journal of Political Economy*, 118, 949–987.
- ROTHSCHILD, M. AND J. STIGLITZ (1976): "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 80, 629–649.
- RUBINSTEIN, A. (1998): *Modeling Bounded Rationality*, The MIT Press.
- RUBINSTEIN, A. AND M. OSBORNE (1994): *A Course in Game Theory*, Cambridge, MA, London, England: MIT Press.

- SELTEN, R. (1965): "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit," *Zeitschrift für die Gesamte Staatswissenschaft*, 121, 301–324 and 667–689.
- (1975): "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, 25–55.
- SPENCE, A. M. (1973): "Job Market Signalling," *Quarterly Journal of Economics*, 90, 225–243.
- VON NEUMANN, J. AND O. MORGENSTERN (1944): *Theory of Games and Economic Behavior*, New York: John Wiley and Sons.
- WATSON, J. (2008): *Strategy: An Introduction to Game Theory*, Norton and Company, second ed.

