GAME THEORY:
Analysis of Strategic Thinking

Pierpaolo Battigalli, Emiliano Catonini,
Nicodemo De Vito

December 2025

Abstract

This textbook introduces some concepts of the theory of games: rationality, dominance, rationalizability, and several notions of equilibrium (Nash, randomized, correlated, self-confirming, subgame perfect, Bayesian perfect equilibrium). For each of these concepts, the interpretative aspect is emphasized. Even though no advanced mathematical knowledge is required, the reader should nonetheless be familiar with the concepts of set, function, probability, and more generally be able to follow abstract reasoning and formal arguments.

Milano,
P. Battigalli, E. Catonini, and N. De Vito

Preface

This textbook provides an introduction to game theory, the formal analysis of strategic interaction. Game theory now pervades most non-elementary models in microeconomic theory and many models in the other branches of economics and in other social sciences. We introduce the necessary analytical tools to be able to understand these models, and illustrate them with some economic applications.

We also aim at developing an abstract analysis of strategic thinking, and a critical and open-minded attitude toward the standard game-theoretic concepts as well as new concepts.

Most of this textbook rely on relatively elementary mathematics. Yet, our approach is formal and rigorous. The reader should be familiar with mathematical notation about sets and functions, with elementary linear algebra and topology in Euclidean spaces, and with proofs by mathematical induction. Elementary calculus is sometimes used in examples.

Contents

1

Introduction

Game theory is the formal analysis of the behavior of interacting individuals. The crucial feature of an interactive situation is that the consequences of the actions of an individual depend (also) on the actions of other individuals. This is typical of many games people play for fun, such as chess or poker. Hence, interactive situations are called "games" and interactive individuals are called "players." If a player's behavior is intentional and he is aware of the interaction (which is not always the case), he should try and anticipate the behavior of other players. This is the essence of strategic thinking. In the rest of this chapter we provide a semi-formal introduction of some key concepts in game theory.

1.1 Decision Theory and Game Theory

Decision theory is a branch of applied mathematics that analyzes the decision problem of an individual (or a group of individuals acting as a single decision unit) in isolation. The external environment is a primitive of the decision problem. Decision theory provides simple decision criteria characterizing an individual's preferences over different courses of action, provided that these preferences satisfy some "rationality" properties, such as completeness (any two alternatives are comparable) and transitivity (if a is preferred to b and b is preferred to c, then a is preferred to c). These criteria are used to find optimal (or rational) decisions.

Game theory could be more appropriately called **interactive decision theory**. Indeed, game theory is a branch of applied mathematics that analyzes interactive decision problems: There are several individuals, called **players**, each facing a decision problem whereby the "external environment" (from the point of view of this particular player) is given by the other players' behavior (and possibly some random variables). In other words, the welfare (utility, payoff, final wealth) of each player is affected not only by his own behavior, but also by the behavior of other players. Therefore, in order to figure out the best course of action each player has to guess which course of action the other players are going to take.

1.2 Why Economists Should Use Game Theory

We are going to argue that game theory should be the main analytical tool used to build formal economic models. More generally, game theory should be used in all formal models in the social sciences that adhere to methodological individualism, i.e., try to explain social phenomena as the result of the actions of many agents, which in turn are freely chosen according to some consistent criterion.

The reader familiar with the pervasiveness of game theory in economics may wonder why we want to stress this point. Isn't it well known that game theory is used in countless applications to model imperfect competition, bargaining, contracting, political competition, and, in general, all social interactions where the action of each individual has a non-negligible effect on the social outcome? Yes, indeed! And yet it is often explicitly or implicitly suggested that game theory is not needed to model situations where each individual is negligible, such as perfectly competitive markets. We are going to explain why this is —in our view— incorrect: the bottom line will be that *every "complete" formal model of an economic (or social) interaction must be a game;* economic theory has analyzed perfect competition by taking shortcuts that have been very fruitful, but must be seen as such, just shortcuts.[1]

If we subscribe to methodological individualism, as mainstream economists claim to do, every social or economic observable phenomena we are interested in analyzing should be reduced to the actions of the

[1]Our view is very similar to the original motivations for the study of game theory offered by the "founding fathers" von Neumann and Morgenstern in their seminal book *The Theory of Games and Economic Behavior* [84].

individuals who form the social or economic system. For example, if we want to study prices and allocations, then we should specify which actions the individuals in the system can choose and how prices and allocations depend on such actions: if I is the set of agents, p is the vector of prices, y is the allocation, and $a = (a_i)_{i \in I}$ is the profile[2] of actions, one for each agent, then we should specify relations $p = f(a)$ and $y = g(a)$. This is done in all models of auctions. For example, in a sealed-bid, first-price, single-object auction, a is the profile of bids for the object on sale, $f(a) = \max\{a_i\}_{i \in I}$ (the object is sold for a price equal to the highest bid) and $g(a)$ is such that the object is allocated to the highest bidder,[3] who has to pay his bid. To be more general, we have to allow the variables of interest to depend also on some exogenous shocks x, as in the functional forms $p = f(a, x)$, $y = g(a, x)$. Furthermore, we should account for dynamics when choices and shocks take place over time, as in $y^t = g^t(a^1, x^1, ..., a^t, x^t)$. Of course, all the constraints on agents choices (such as those determined by technology) should also be explicitly specified. Finally, if we are to explain choices according to some rationality criterion, we should include in the model the preferences of each individual i over possible outcomes. This is what we call a "complete model" of the interactive situation.[4] We call variables that depend on actions and (possibly) exogenous shocks (such as variable y) "**endogenous**." (Actions themselves are "endogenous" in a trivial sense.) The rationale for this terminology is that we try to analyze/explain actions and variables that depend on actions.

At this point, you may think that this is just a trite repetition of some abstract methodology of economic modelling. Well, think twice! The most standard concept in the economist's toolkit, Walrasian equilibrium, is not based on a complete model and is able to explain prices and allocations only by taking a two-step shortcut: (1) The modeler "pretends" that prices are observed and taken as parametrically given by *all* agents (including all firms) before they act, hence before they can affect such prices; this is a

---

[2]Given a set I of individuals, we always call "profile" a list of elements from a given set (e.g., a set of actions), one for each individual i.

[3]Ties can be broken at random.

[4]The model is still in some sense incomplete: we have not even addressed the issue of what the individuals know about the situation and about each other. But the elements sketched in the main text are sufficient for the discussion. Let us stress that what we call "complete model" does not include the modeler's hypotheses on how the agents choose, which are key to provide explanations or predictions of economic and social phenomena.

kind of logical short-circuit, but it allows to determine demand and supply functions $D(p)$, $S(p)$. Next, (2) market-clearing conditions $D(p) = S(p)$ determine equilibrium prices. Well, this can only be seen as a (clever) reduced-form approach; absent an explicit model of price formation (such as an auction model), the modeler postulates that somehow the choices-prices-choices feedback process has reached a rest point and he describes this point as a market-clearing equilibrium. In many applications of economic theory to the study of competitive markets, this is a very reasonable and useful shortcut, but it remains just a shortcut, forced by the lack of what we call a complete model of the interactive situation.

So, what do we get when instead we have a complete model? As we are going to show a bit more formally in the next section, we get what game theorists call a "game." This is why game theory should be a basic tool in economic modelling, even if one wants to analyze perfectly competitive situations. To illustrate this point, we will present a purely game-theoretic analysis of a perfectly competitive market, showing not only how such an analysis is possible, but also that it adds to our understanding of how equilibrium can be reached.

1.3 Abstract Game Models

A completely formal definition of the mathematical object called "(non-cooperative) game" will be given in due time. We start with a semi-formal introduction to the key concepts, illustrated by a very simple example, a seller-buyer mini-game. Consider two individuals, S (Seller) and B (Buyer). Let S be the owner of an object and B a potential buyer. For simplicity, consider the following bargaining protocol: S can ask one Euro (1) or two Euros (2) to sell the object, B can only accept (a) or reject (r). The monetary value of the object for individual i ($i = S, B$) is denoted by V_i. This situation can be analyzed as a game which can be represented with a rooted tree with utility numbers attached to terminal nodes (leaves), player labels attached to nonterminal nodes, and action labels attached to arcs; see, for instance, Figure 1.1.

The game tree represents the formal elements of the analysis: the set of players (or roles in the game, such as seller and buyer), the actions, the rules of interaction, the consequences of complete sequences of actions and how players value such consequences.

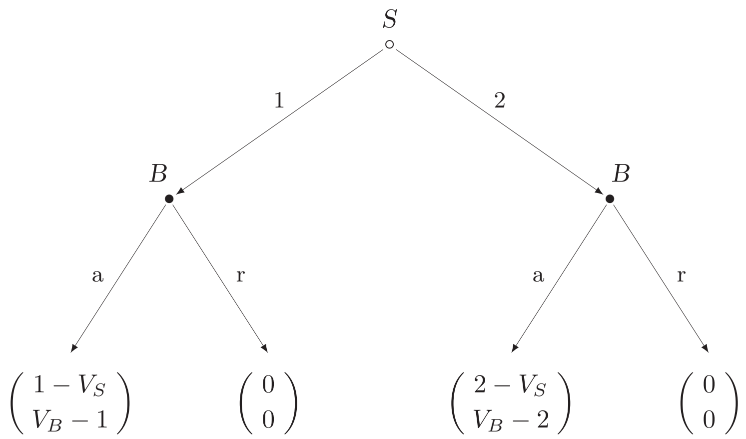

Figure 1.1: A seller-buyer mini-game.

Formally, a **game form** is given by the following elements:

- I, a set of **players**;
- For each player $i \in I$, a set A_i of **actions** which could conceivably be chosen by i at some point of the game as the play unfolds.
- Y, a set of **outcomes,** or **consequences**;
- $\mathcal{E}$, an **extensive-form** structure, that is, a mathematical representation of the rules saying whose turn it is to move, what a player knows, i.e., his or her information about past moves and random events, and what are the feasible actions at each point of the game; this determines a set Z of possible paths of play (sequences of feasible actions); a path of play $z \in Z$ may also contain some random events such as the outcomes of throwing dice;
- $g : Z \to Y$, an **outcome** (or consequence) **function** which assigns to each play z a consequence $g(z)$ in Y.

The above elements represent what the layperson would call the "rules of the game." To complete the description of the actual interactive situation (which may differ from how the players perceive it) we have to add players' preferences over consequences and—*via* expected utility calculations—lotteries over consequences:

- $(v_i)_{i \in I}$, where each $v_i : Y \to \mathbb{R}$ is a von Neumann-Morgenstern **utility function** representing player i's preferences over outcomes (consequences); preferences over lotteries are obtained *via* expected utility comparisons.

With this, we obtain the kind of mathematical structure called "game" in the technical sense of game theory. *This formal description does not say how such a game would (or should) be played.* The description of an interactive decision problem is only the first step to make a prediction (or a prescription) about players' behavior.

To illustrate, the mathematical description of the seller-buyer mini-game is as follows: $I = \{S, B\}$; $A_S = \{1, 2\}$, $A_B = \{a, r\}$; Y is a specification of the set of final allocations of the object and of monetary payments (for example, $Y = \{o_S, o_B\} \times \{0, 1, 2\}$, where (o_i, k) means object to i and k euros are given by B to S); $\mathcal{E}$ is represented by the game tree, without the utility numbers at endnodes; g is the rule stipulating that if B accepts the ask price p then B gets the object and gives p euros to S, if B rejects then S keeps the object and B keeps the money ($g(p, a) = (o_B, p)$, $g(p, r) = (o_S, 0)$); v_i is a risk-neutral (quasi-linear) utility function normalized so that the utility of no exchange is zero for both players ($v_S(o_B, p) = p - V_S$, $v_B(o_B, p) = V_B - p$, $v_i(o_S, 0) = 0$ for $i = S, B$). Game theory provides predictions on the behavior of S and B in this game based on hypotheses about players' knowledge of the rules of the game and of each other preferences, and on hypotheses about strategic thinking. For example, B is assumed to accept an ask price if and only if it is below his valuation. Whether the ask price is high or low depends on the valuation of S and what S knows about the valuation of B. If S knows the valuation of B, he can anticipate B's response to each offer and how much surplus he can extract.

A caveat on theoretical language It is common in game-theoretic work to refer to the *mathematical structure* describing an interactive situation (the "real game") as if it were the situation itself. Such abuse of language is usually innocuous, but sometimes the distinction is lost and this leads to unclear or confusing language. It is therefore useful to keep in mind the distinction between the real world objects and phenomena we are interested in and the mathematical structures used in the theoretical analysis. For example, the distinction is important in our analysis and discussion of games with incomplete or asymmetric information in Section 8.6 of Chapter 8.

1.4 Terminology and Classification of Games

Games come in different varieties and are analyzed with different methodologies. The same strategic interaction may be represented in a detailed way, using the mathematical objects described above, or in a much more parsimonious way. The amount of details in the formal representation constrains the methods used in the analysis. The terminology used to refer to different kinds of strategic situations and to different formal objects used to represent them may be confusing: Some terms that have almost the same meaning in the daily natural language, such as "perfect information" and "complete information," have very different meanings in game theory; other terms, such as "non-cooperative game," may be misleading. Furthermore, there is a tendency to confuse the substantive properties of the situations of strategic interaction that game theory aims at studying and the formal properties of the mathematical structures used in this endeavour. Here we briefly summarize the terminology and classification of game theory doing our best to dispel such confusion.

1.4.1 Cooperative vs Non-Cooperative Games

Suppose that the players (or at least some of them) could meet before the game is played and sign a binding agreement specifying their course of action (an external enforcing agency—e.g., the courts system—will force each player to follow the agreement). This could be mutually beneficial for them and if this possibility exists we should take it into account. But how? The theory of **cooperative games** does not model the process of bargaining (offers and counteroffers) which takes place before the game starts; this theory considers instead a simple and parsimonious representation of the situation, that is, how much of the total surplus each possible coalition of players can guarantee to itself by means of some binding agreement. For example, in the seller-buyer situation discussed above, the (normalized) surplus that each player can guarantee to himself is zero, while the surplus that the $\{S, B\}$ coalition can guarantee to itself (the "gains from trade") is the difference between the lowest and the highest valuation, that is, $V_B - V_S$ in the standard case where $V_B \geq V_S$. This simplified representation is called **coalitional game**. For every given coalitional game the theory tries to figure out which division of the surplus could result, or, at least, the set of allocations that are not excluded by strategic considerations; see, for example, Osborne and Rubinstein [65, Part IV].

The theory of **non-cooperative games** instead posits that either binding agreements are not feasible (e.g., in an oligopolistic market they could be forbidden by antitrust laws), or that the bargaining process which could lead to a binding agreement on how to play a game G is appropriately formalized as a sequence of moves in a "larger game" $\Gamma(G)$.[5] The players cannot agree on how to bargain in $\Gamma(G)$ and this is the game to be analyzed with the tools of non-cooperative game theory. It is therefore argued that non-cooperative game theory is more fundamental than cooperative game theory. For example, in the seller-buyer situation, $\Gamma(G)$ would be represented by the game-tree displayed in Figure 1.1 to illustrate the formal objects comprised by the mathematical representation of a game. The analysis of this game reveals that if S knows V_B he can use his first-mover advantage to ask for the highest price below V_B. Of course, the results of the analysis depend on details of the bargaining rules that may be unknown to an external observer and analyst. Neglecting such details, cooperative game theory typically gives weaker, but more robust predictions. Yet, in our view, the main advantage of non-cooperative game theory is not so much its sharper predictions, but rather the conceptual clarity of working with *explicit* assumptions and showing exactly how different assumptions lead to different results.

Cooperative game theory is an elegant and useful analytical toolkit. It is especially appropriate when the analyst has little understanding of the true rules of interaction and wishes to derive some robust results from parsimonious information about the outcomes that coalitions of players can

[5] Here we are using symbol G to represent the actual interactive situation being modeled. As we move on to the formal analysis of games we will use the same symbol to denote the *mathematical structure* representing such situations. The meaning of G and similar symbols will be either explicit or clear from the context.

implement. However, non-cooperative game theory is much more pervasive in modern economics, and more generally in modern formal social sciences. We will therefore focus on non-cooperative game theory.

One thing should be clear, focusing on this part of the theory does not mean that cooperation is neglected. *Non-cooperative game theory is **not** the study of un-cooperative behavior*, but rather a method of analysis. Indeed we find the name "non-cooperative game theory" misleading, but it is now entrenched and we will conform to it.

1.4.2 Static and Sequential Games

A game is **static** if each player moves only once and all players move *simultaneously* (or at least without any information about other players' moves). Examples of static games are: Matching Pennies, Stone-Scissor-Paper, and sealed-bid auctions.

A game has a **multistage** structure if some moves are *sequential*, players always know the number of previous moves, and some player may observe (at least partially) the previous behavior of his or her co-players. Examples of multistage games are: Chess, Poker, and open outcry auctions. Not all games with sequential moves have a multistage structure. We focus on this special case because it is simpler and it covers many applications.

Games with sequential moves are sometimes analyzed *as if* the players moved only once and simultaneously. The trick is to pretend that each player chooses in advance his **strategy**, i.e., a contingent plan of action specifying how to behave in every circumstance that may arise while playing the game. Each profile of strategies $s = (s_i)_{i \in I}$ determines a particular path, viz., $\zeta(s)$, hence an outcome $g(\zeta(s))$, and ultimately a profile of utilities, misleadingly called "**payoffs**," $(v_i(g(\zeta(s))))_{i \in I}$. This mapping $s \longmapsto (v_i(g(\zeta(s))))_{i \in I}$ is called the **normal form**, or **strategic form** of the game. The strategic form of a game can be seen as a static game where players simultaneously choose strategies in advance. For example, in the seller-buyer mini-game the strategies of S (seller) coincide with his possible ask prices; the set of strategies of B (buyer) instead contains four response rules: $S_B = \{a_1a_2, a_1r_2, r_1a_2, r_1r_2\}$ where a_p (respectively, r_p) is the instruction "accept (resp., reject) price p." The strategic form is as follows:

$S \backslash B$	a_1a_2	a_1r_2	r_1a_2	r_1r_2
$p = 1$	$1 - V_S, V_B - 1$	$1 - V_S, V_B - 1$	$0, 0$	$0, 0$
$p = 2$	$2 - V_S, V_B - 2$	$0, 0$	$2 - V_S, V_B - 2$	$0, 0$

Figure 1.2: Strategic form of the seller-buyer minigame.

A strategy in a static game is just the plan to take a specific action, with no possibility to make this choice contingent on the actions of others, as such actions are being simultaneously chosen. Therefore, the mathematical representation of a static game and its normal form must be the same. For this reason, static games are often called "normal-form games," or "strategic-form games." We are going to avoid this widespread abuse of language. In a correct language there should not be "normal-form games"; rather, looking at the normal form *of* a game is a *method of analysis.*

1.4.3 Assumptions about Information

Perfect Information and Observed Actions

A multistage game has **perfect information** if players move one at a time and each player—when it is his or her turn to move—is informed of all the previous moves (including the realizations of chance moves). If some moves are simultaneous, but each player observes all past moves, we say that the game has **observed actions (or perfect monitoring, or almost perfect information)**. Examples of games with perfect information are the seller-buyer mini-game, Chess, Backgammon, and Tic-Tac-Toe. An example of a game with observed actions is the repeated Cournot oligopoly, with simultaneous choice of outputs in every period and perfect monitoring of past outputs. Note, perfect information is an assumption about the *rules* of the game, i.e., an assumption about what information about past moves reaches players according to such rules.

Asymmetric Information

A game with imperfect information features **asymmetric information** if different players get different pieces of information about past moves, and in particular about the realizations of chance moves. Poker and Bridge are games with asymmetric information because players observe their cards, but not the cards received by others players. Like perfect information, also asymmetric information is entailed by the *rules* of the game.

Complete and Incomplete Information

An event E is **common knowledge** if everybody knows E, everybody knows that everybody knows E, and so on for all iterations of "everybody knows that." To use a suggestive semi-formal expression: for every $m = 1, 2, ...$ it is the case that [(everybody knows that)^m E is the case].

There is **complete information** in an interactive decision problem represented by $G = \langle I, A, Y, \mathcal{E}, g, (v_i)_{i \in I} \rangle$ if it is common knowledge that the actual interactive situation is the one represented by G. Conversely, there is **incomplete information** if, for some player i, either it is *not* the case that [i knows the actual interactive situation], or for some $m = 1, 2, ...$ it is *not* the case that [i knows that (everybody knows that)^m the actual interactive situation]. Most economic situations feature incomplete information because either the outcome function (represented by g), or players' preferences (represented by the utility functions $(v_i)_{i \in I}$) are not common knowledge. Note, complete (or incomplete) information is *not* an assumption *about* the *rules* of the game, it is an assumption on players' "interactive knowledge" concerning the rules *and* preferences. For example, in the seller-buyer mini-game it can be safely assumed that there is common knowledge of the outcome function (who gets the object and monetary transfers), but the valuations V_S and V_B need not be commonly known. For some types of objects on sale, V_i is known only to i; for other types of objects, the seller may know the quality of the object better then the buyer, so V_B could be given by an idiosyncratic component known to B plus a quality component known to S.

Although situations with *asymmetric* information about chance moves, such as Poker, are conceptually different from situations with *incomplete* information, we will see that there is a formal similarity which allows (at least to some extent) to use the same analytical tools to study both situations.

Examples: On the one hand, in games like chess and poker there is—presumably—complete information; indeed, the "rules of the game" are common knowledge, and it may be taken for granted that it is common knowledge that players like to win. On the other hand, auctions typically feature incomplete information because the competitors' valuations of the objects on sale are not commonly known.

1.5 Rational Behavior

The decision problem of a single individual i can be represented as follows: i can take an action in a set of feasible alternatives, or decisions, D_i; i's welfare (payoff, utility) is determined by his decision d_i and an external state $\omega_{-i} \in \Omega_{-i}$ (which represents the realizations of a vector of variables beyond i's control, e.g., the outcome of a random event and/or the decisions of other players) according to the outcome function

$$g : D_i \times \Omega_{-i} \to Y$$

and the utility function

$$v_i : Y \to \mathbb{R}.$$

We assume that i's choice cannot affect ω_{-i} and that i does not have any information about ω_{-i} when he chooses beyond the fact that $\omega_{-i} \in \Omega_{-i}$. The choice is made once and for all. (We will show in the part on sequential games how this model can be generalized.) Assume, just for simplicity, that D_i and Ω_{-i} are both finite and let, for notational convenience, $\Omega_{-i} = \{\omega_{-i}^1, \omega_{-i}^2, ..., \omega_{-i}^n\}$. In order to make a "rational" choice, i has to assess the probabilities of the different states in Ω_{-i}. Suppose that his beliefs about ω_{-i} are represented by a probability vector $\mu^i \in \Delta(\Omega_{-i})$ where

$$\Delta(\Omega_{-i}) = \left\{ \mu^i \in \mathbb{R}_+^n : \sum_{k=1}^{n} \mu^i(\omega_{-i}^k) = 1 \right\}.$$

Then i's **rational decision given** μ^i is to take any alternative that maximizes i's expected payoff, i.e., any d_i^* such that

$$d_i^* \in \arg\max_{d_i \in D_i} \mathbb{E}_{\mu^i} \left(v_i \left(g \left(d_i, \cdot \right) \right) \right),$$

where $\mathbb{E}_{\mu^i}$ denotes the expectation with respect to μ^i,

$$\mathbb{E}_{\mu^i} \left(v_i \left(g \left(d_i, \cdot \right) \right) \right) = \sum_{k=1}^{n} \mu^i(\omega_{-i}^k) v_i \left(g(d_i, \omega_{-i}^k) \right).$$

Decision d_i^* is also called a **best response**, or **best reply**, to μ^i. The probability vector could be exogenously given (roulette lottery) or simply represent i's **subjective beliefs** about ω_{-i} (horse lottery, game against an opponent). The composite function $v_i \circ g$ is denoted u_i, that is, $u_i(d_i, \omega_{-i}) = v_i\left(g(d_i, \omega_{-i})\right)$; u_i is called the **payoff function**, because in many interesting games the rules of the decision problem (or game) attach monetary payoffs to the possible outcomes of the game and it is taken for granted that the decision maker maximizes his expected monetary payoff. In the finite case, u_i is often represented as a matrix with (k, ℓ) entry $u_i^{k\ell} = u_i(d_i^k, \omega_{-i}^\ell)$: see Figure 1.3.

	ω_{-i}^1	ω_{-i}^2	$\ldots$
d_i^1	u_i^{11}	u_i^{12}	$\ldots$
d_i^2	u_i^{21}	u_i^{22}	$\ldots$
$\vdots$	$\vdots$	$\vdots$	$\ddots$

Figure 1.3: Matrix representation of the payoff function.

This representation of a decision maker's choice, of course, fits very well the case of static games. In this case $D_i = A_i$ is simply the set of feasible actions and Ω_{-i} may be the set A_{-i} of the opponents' feasible action profiles (or a larger space if there is incomplete information).

However, it has been argued that the representation is more general:[6] Consider a game with sequential moves, e.g., a multistage game. The rules of interaction specifies all the circumstances under which player i may have to choose an action and the information i would have in those circumstances. Then player i can formulate in advance a plan of action, or strategy, which specifies how he would behave in any such circumstance given the available information. Assuming that such plan is incentive compatible, that is, that player i has no incentive to deviate from it, then player i expects to follow his plan. His uncertainty concerns the information-dependent behavior of others, hence the strategies they are going to follow. Every profile of strategies, when carried out, induces a particular outcome. Therefore, the planning problem faced by a rational agent is the problem of selecting an incentive-compatible strategy under

[6] See von Neumann and Morgenstern [84].

uncertainty about the strategies used by the opponents (and possibly about other unknowns). In the part of this textbook devoted to multistage games we will discuss some subtleties related to the concepts of "strategy" and "plan."

1.6 Assumptions and Solution Concepts

Game theory provides a mathematical language to formulate assumptions about the "rules of the game" and players' preferences, that is, all the elements listed in Section 1.3. But such assumptions are not enough to derive conclusions about how the players behave in a given game. The central behavioral assumption in game theory is that players are rational (see Section 1.5). However, without further assumptions about what players believe about the variables which affect their payoffs, but are beyond their control (in particular their opponents' behavior), the rationality assumption has little behavioral content: we can only say that a rational player's choice must be **justifiable**, i.e., it it must be a best response to *some* belief. In most games, this is too little to derive interesting results.[7]

Then, how can we obtain interesting results from our assumptions about the rules of the game and preferences? The standard approach used in game theory is analogous to the one used in the undergraduate economics analysis of competitive markets, whereby we formulate assumptions about preferences and technology and then we assume that economic agents' plans are mutually consistent best responses to equilibrium prices.

As explained in Section 1.2, there is an important difference: the textbook analysis of competitive markets does not specify a price-formation mechanism and uses equilibrium market-clearing as a theoretical shortcut to overcome this problem. On the contrary, in a game-theoretic model all the observable variables we try to explain depend on players' actions (and exogenous shocks) according to an explicitly specified function, as in auction models.

Yet, there are also similarities with the analysis of competitive markets.

[7]There are important exceptions. For example, in many situations where each individual in a group decides how much to contribute to a public good, it is strictly dominant to give a minimal contribution (see Example 3 in Chapter 3).

One could say that the role of prices is played (sic) by players' beliefs since we assume that they are, in some sense, mutually consistent. The precise meaning of the statement "beliefs are mutually consistent" is captured by a **solution concept.** The simplest and most widely used solution concept in game theory, Nash equilibrium, assumes that players' beliefs about each other strategies are correct and each player best responds to her or his beliefs; as a result, each player uses a strategy that is a best response to the strategies used by other players.

Nash equilibrium is not the only solution concept used in game theory. Recent developments made clear that solution concepts implicitly capture expressible[8] *assumptions about players' rationality and beliefs*, and some assumptions that are appropriate in some contexts are too weak, or simply inappropriate, in different contexts. Therefore, it is very important to provide convincing motivations for the solution concept used in a specific application.

Let us somewhat vaguely define as "**strategic thinking**" what intelligent agents do when they are fully aware of participating in an interactive situation and form conjectures by putting themselves in the shoes of other intelligent agents. As the title of this book suggests, we mostly (though not exclusively) present game theory as an *analysis of strategic thinking.* We will often follow the traditional "textbook" game theory and present different solution concepts providing informal motivations for them, sometimes with a lot of hand-waving. However, the reader should always be aware that a more fundamental approach is being developed, where game theorists formulate explicit assumptions not only about the rules of the game and preferences, but also about players' rationality and initial beliefs, and also about how beliefs change as the play unfolds. Implications about choices and observables can be derived directly from such assumptions, without the mediation of solution concepts. Then, solution concepts become mere "shortcuts" to characterize the behavioral implications of assumptions about rationality and beliefs.

For example, a standard solution concept in game theory, subgame perfect equilibrium, requires players' strategies to form an equilibrium not only for the game itself, but also in every subgame.[9] In finite games

[8] Roughly, "expressible" means "something that can be expressed in a clear, precise, and not self-referential language."

[9] A precise definition of a subgame will be given in Part II. For the time being, you

with (complete and) perfect information, if there are no relevant ties, there is a unique subgame perfect equilibrium that can be computed with a backward induction algorithm. For two-stage games with perfect information, subgame perfect equilibrium can be derived from simple assumptions about rationality and beliefs: players are rational and the first mover believes that the second mover is rational. To illustrate, consider the seller-buyer mini-game and assume players have *complete information.* Specifically, suppose that it is common knowledge that $V_S < 1$and $V_B > 2$. The assumption of rationality of the buyer B, denoted with R_B, implies that B accepts every price below valuation V_B; since $V_B > 2$, the only strategy consistent with R_B is a_1a_2, that is, to accept both $p = 1$ and $p = 2$. Rationality of the seller S, denoted with R_S, implies that S attaches subjective probabilities to the four strategies of B and chooses the ask price that maximizes S's expected utility; if S has deterministic beliefs, S asks for the highest prices he expects to be accepted, e.g., S asks $p = 1$ if he is certain of strategy a_1r_2. But if S is certain that B is rational, given the complete information assumption, S must assign probability 1 to strategy a_1a_2. Then the rationality of S and S's certainty in the rationality of B imply that S asks for the highest price, i.e., S uses the bargaining power derived from the knowledge that $V_B > 2$ and the first-mover advantage. In symbols, write $\mathrm{B}_i(E)$ for "i believes (with probability 1) E," where E is some event; then $R_S \cap \mathrm{B}_S(R_B)$ implies that S asks for $p = 2$; R_B implies that $p = 2$ is accepted; $R_S \cap \mathrm{B}_S(R_B) \cap R_B$ implies that the action/strategy of S is $p = 2$ and the strategy of B is a_1a_2, which is the subgame perfect equilibrium obtained *via* backward induction.

Two-stage games with complete and perfect information are so simple that the analysis above may seem trivial. On the other hand, the *rigorous* analysis of more complicated multistage games based on higher levels of mutual belief in rationality presents conceptual difficulties that will be addressed only in the second part of this textbook. Here we further illustrate strategic thinking with a semi-formal analysis of a static game representing a perfectly competitive market. We hope this will also have the side effect to make the reader understand that *game theory is useful to analyze every "complete" model of economic interaction, including models of perfect competition.*

---

should think of a subgame as a "smaller"game that starts with a nonterminal node of the original game.

1.6.1 Analysis of a Perfectly Competitive Market

Let us analyze the decisions of a large number n of small, identical agricultural firms producing a crop (say, corn). For reasons that will be clear momentarily, it is convenient to index firms as equally spaced rational numbers in the interval $(0, 1]$, that is, $I = \{\frac{1}{n}, \frac{2}{n}, \ldots, \frac{n-1}{n}, 1\} \subset (0, 1]$ with generic element $i \in I$. Each firm i decides in advance how much to produce. This quantity is denoted $q(i)$. The output of corn $q(i)$ will be ready to be sold only in six months. Markets are incomplete: it is impossible to agree in advance on a price for corn to be delivered in six months. So, each firm i has to guess the price p for corn in six months. It is common knowledge that the demand for corn (not explicitly "micro-founded" in this example) is given by function $D(p) = n \max\{(\alpha - \beta p), 0\}$, that is, $D(p) = n(\alpha - \beta p)$ if $p \leq \frac{\alpha}{\beta}$ and $D(p) = 0$ otherwise. It is also common knowledge that each firm will fetch the highest price at which the market can absorb the total output for sale, $Q = \sum_{i=1}^{n} q(i)$. Thus, each firm will sell each unit of output at the uniform price $P\left(\frac{Q}{n}\right) = \frac{1}{\beta}\left(\alpha - \frac{Q}{n}\right)$ if $\frac{Q}{n} \leq \alpha$ and $p = 0$ if $\frac{Q}{n} > \alpha$ (for notational convenience, we express inverse demand as a function of the *average* output $\frac{Q}{n}$). Each firm i has total cost function $C(q) = \frac{1}{2m}q^2$, hence the marginal cost function is $MC(q) = \frac{q}{m}$. Each firm has obvious preferences, it wants to maximize the difference between revenues and costs. The technology (cost function) and preferences of each firm are common knowledge.[10]

Now, we want to represent mathematically the assumption that each firm is "negligible" with respect to the (maximum) size of the market α. Instead of assuming a large, but finite set of firms given by a fine grid I in the interval $(0, 1]$, we use a quite standard mathematical idealization and assume that there is a *continuum* of firms, normalized to the unit interval $I = (0, 1]$. Thus, the average output is $q = \int_0^1 q(i)\mathrm{d}i$ instead of $\frac{1}{n}\sum_{i=1}^{n} q(i)$. Each firm understands that its decision cannot affect the total output and the market price.[11]

Given that all the above is common knowledge, we want to derive the price implications of the following assumptions about rationality and

---

[10] The analysis can be generalized assuming heterogeneous firms, where each firm i is characterized by the marginal cost parameter $m(i)$. Then it is enough to assume that each i knows $m(i)$ and that the average $m = \frac{1}{n}\sum_{i=1}^{n} m(i)$ is common knowledge.

[11] This example is borrowed from Guesnerie [51].

beliefs:

R (*rationality*): each firm has a conjecture about (q and) p and chooses a best response to such conjecture, for brevity, each firm is rational;[12]

$\mathrm{B}(R)$ (*mutual belief in rationality*): all firms believe (with certainty) that all firms are rational;

$\mathrm{B}^2(R) = \mathrm{B}(\mathrm{B}(R))$ (*mutual belief of order 2 in rationality*): all firms believe that all firms believe that all firms are rational;

...

$\mathrm{B}^k(R) = \mathrm{B}(\mathrm{B}^{k-1}(R))$ (*mutual belief of order k in rationality*): [all firms believe that]^k all firms are rational;

...

Note, we are using symbols for these assumptions. In a more formal mathematical analysis, these symbols would correspond to *events* in a space of states of the world. Here they are just useful abbreviations. Conjunctions of assumptions will be denoted by the symbol $\cap$, as in $R \cap \mathrm{B}(R)$. It is a good thing that you get used to these symbols, but if you find them baffling, just ignore them and focus on the words.

What are the consequences of *rationality* (R) in this market? Let $p(i)$ denote the price *expected* by firm i.[13] Then i solves

$$\max_{q\geq 0}\left\{p(i)q - \frac{1}{2m}q^2\right\}$$

which yields

$$q(i) = m\ p(i).$$

Since firms know the inverse demand function $P(q) = \max\{(\alpha - q)/\beta, 0\}$, each firm i's expectation of the price is below the upper bound $\overline{p}^0 = \frac{\alpha}{\beta}$: $p(i) \leq \frac{\alpha}{\beta}$. It follows that

$$q = \int_0^1 mp(i)\mathrm{d}i = m\int_0^1 p(i)\mathrm{d}i \leq \overline{q}^1 := m\frac{\alpha}{\beta},$$

where $\overline{q}^1$ denotes the upper bound on average output when firms are rational. (We use ":=" to mean "equal by definition," when this is not obvious from the context.)

[12] The symbols R, $\mathrm{B}(R)$ etc. below should be read as abbreviations of sentences. They can also be formally defined using mathematics, but this goes beyond the scope of this textbook.

[13] This is, in general, the mathematical expectation of p according to the subjective probabilistic conjecture of firm i.

By *mutual belief in rationality*, each firm understands that $q \leq \overline{q}^1$ and $p \geq \max\{(\alpha - m\frac{\alpha}{\beta})/\beta, 0\}$. If $m \geq \beta$, this is not very helpful: assuming belief in rationality does not reduce the span of possible prices and does not yield additional implications for the endogenous variables we are interested in. It is not hard to understand that rationality and mutual belief in rationality of every order yields the same result $q \leq \frac{m\alpha}{\beta}$, where $\frac{m\alpha}{\beta} \geq \alpha$ because $m \geq \beta$.

So, let us assume from now on that $m < \beta$. Note, this is the so called "cobweb stability" condition. Given belief in rationality, $\mathrm{B}(R)$, each firm expects a price at least as high as the lower bound

$$\underline{p}^1 := \frac{1}{\beta}\left(\alpha - \overline{q}^1\right) = \frac{\alpha}{\beta}\left(1 - \frac{m}{\beta}\right) > 0.$$

Therefore $R \cap \mathrm{B}(R)$ implies that $q = m\int_0^1 p(i)di \geq m\underline{p}^1$, that is, average output is above the lower bound

$$\underline{q}^2 := m\underline{p}^1 = \alpha\frac{m}{\beta}\left(1 - \frac{m}{\beta}\right)$$

and hence the price is below the upper bound

$$\overline{p}^2 := \frac{1}{\beta}(\alpha - \underline{q}^2) = \frac{\alpha}{\beta}\left(1 - \frac{m}{\beta}\left(1 - \frac{m}{\beta}\right)\right) = \frac{\alpha}{\beta}\sum_{k=0}^{2}\left(-\frac{m}{\beta}\right)^k.$$

By $\mathrm{B}(R)$ and $\mathrm{B}^2(R)$ (*mutual belief of order 2 in rationality*), each firm predicts that $p \leq \overline{p}^2$. Rational firms with such expectations choose an output below the upper bound

$$\overline{q}^3 := m\overline{p}^2 = \alpha\frac{m}{\beta}\left(1 - \frac{m}{\beta}\left(1 - \frac{m}{\beta}\right)\right) = \alpha\frac{m}{\beta}\sum_{k=0}^{2}\left(-\frac{m}{\beta}\right)^k.$$

Thus, $R \cap \mathrm{B}(R) \cap \mathrm{B}^2(R)$ implies that $q \leq \overline{q}^3$ and the price must be above the lower bound

$$\underline{p}^3 := \frac{1}{\beta}(\alpha - \overline{q}^3) = \frac{\alpha}{\beta}\sum_{k=0}^{3}\left(-\frac{m}{\beta}\right)^k.$$

Can you guess the consequences for price and average output of assuming *rationality and common belief in rationality* (mutual belief in rationality of every order)? Draw the classical Marshallian cross of demand and supply, trace the upper and lower bounds found above and go on. You will find the answer.

More formally, define the following sequences of upper bounds and lower bounds on the price:

$$\overline{p}^L = \frac{\alpha}{\beta}\sum_{k=0}^{L}\left(-\frac{m}{\beta}\right)^k, \; L \text{ even}$$

$$\underline{p}^L = \frac{\alpha}{\beta}\sum_{k=0}^{L}\left(-\frac{m}{\beta}\right)^k, \; L \text{ odd}$$

It can be shown by mathematical induction that $R \cap \bigcap_{k=1}^{L-1} B^k(R)$ (rationality and mutual belief in rationality of order $k = 1, \ldots, L-1$, with $L \geq 2$) implies $p \leq \overline{p}^L$ if L is even, and $p \geq \underline{p}^L$ if L is odd. Since $\frac{m}{\beta} < 1$, the sequence $\overline{p}^{2\ell} = \frac{\alpha}{\beta}\sum_{k=0}^{2\ell}\left(-\frac{m}{\beta}\right)^k$ is decreasing in ℓ, the sequence $\underline{p}^{2\ell+1} = \frac{\alpha}{\beta}\sum_{k=0}^{2\ell+1}\left(-\frac{m}{\beta}\right)^k$ is increasing in ℓ, and they both converge to the competitive equilibrium price $p^* = \frac{\alpha}{\beta+m}$, see Figure 1.4.

A Comment on "Rational" *versus* "Rationalizable" Expectations

This price p^* is often called the **"rational expectations"** price, because it is the price that the competitive firms would expect if they performed the same equilibrium analysis as the modeler. *We refrain from using such terminology.* Game theory allows a much more rigorous and precise terminology, the one that we have been using above. In this textbook, "rationality" is a joint property of choices and beliefs and it means *only* that agents best respond to their beliefs. The phrase "rational beliefs," or "rational expectations" was coined at a time when theoretical economists did not have the formal tools to be precise and rigorous in the analysis of rationality and interactive beliefs. Now that these tools exist, such phrases as "rational expectations" should be avoided, at least in game theoretic analysis. On the other hand, expectations may be consistent with common belief in rationality (e.g., under common knowledge of the game), and it

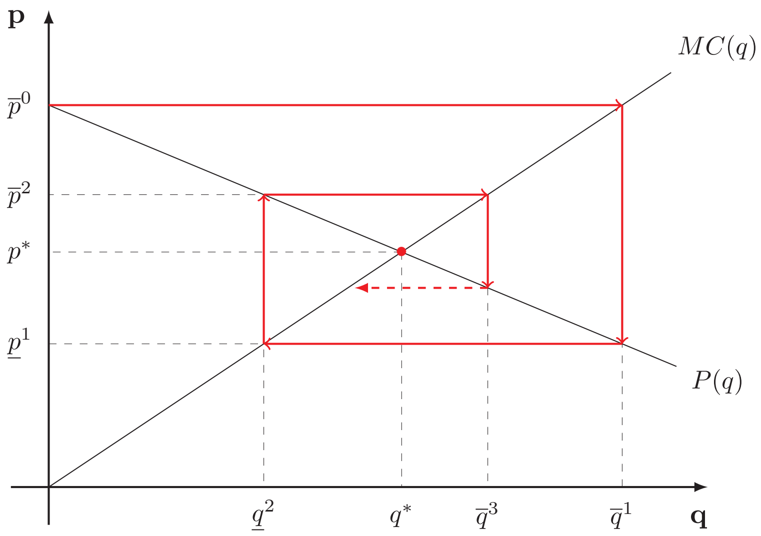

Figure 1.4: Strategic thinking in a perfectly competitive market.

makes sense to call such expectations **"rationalizable**." Note well, the fact that there is only one "rationalizable-expectations" price, p^*, is a *conclusion* of the analysis, it has not been assumed! Also, this conclusion holds only under the "cobweb stability" condition $m < \beta$. This contrasts with standard equilibrium analysis, whereby it is assumed at the outset that expectations are correct, whether or not this follows from assumptions like common belief in rationality.

What do we learn from this example? First, it shows that certain *assumptions about rationality and beliefs* (plus common knowledge of the game) may yield interesting implications or not according to the given model and parameter values. Here, rationality and common belief in rationality have sharp implications for the market price and average output if the "cobweb stability" condition holds. Otherwise they only yield a high upper bound on average output.

Second, recall that **strategic thinking** is what intelligent agents do when they are fully aware of participating in an interactive choice situation and form conjectures by putting themselves in the shoes of other intelligent agents; in this sense, the above is as good an analysis of strategic thinking as any, and yet it is the analysis of a perfectly competitive market!

Third, the example shows how solution concepts can be interpreted as shortcuts. Here we presented an iterated deletion of prices and conjectures, where one deletes prices that cannot occur when firms best respond to conjectures and delete conjectures that do not rule out deleted prices. Note, mutual beliefs of order k do not appear explicitly in this iterative deletion procedure. The procedure is a solution concept of a kind and it can be used as a shortcut to obtain the implications of rationality and common belief in rationality. Under "cobweb stability," the so-called "rational-expectations" equilibrium price p^* results. This shows that under appropriate conditions we can use the equilibrium concept to obtain the implications of rationality and common belief in rationality. This theme will be played time and again.

Part I

Static Games

2

Static Games: Description

A **static game**, or one-stage game, or simultaneous-moves game, is an interactive situation where all players move simultaneously and only once.[1] The key features of a static game are formally represented by a mathematical structure (a list of sets and functions)

$$G = \langle I, Y, (A_i)_{i\in I}, g, (v_i)_{i\in I} \rangle ,$$

where:

- I is the set of individuals or **players**, typically denoted by i, j;
- A_i is the *nonempty* set of possible **actions** for player i; a_i, a_i^*, a_i', $\hat{a}_i$ are alternative action labels we frequently use;
- $g : \times_{i\in I} A_i \to Y$ is the **outcome** (or consequence) **function** which captures the essence of the rules of the game, beyond the assumption of simultaneous moves;

[1]Sometimes static games are also called "normal-form games" or "strategic-form games." As we mentioned in Chapter 1, this terminology is somewhat misleading. The normal, or strategic form of a game Γ has the same structure of a static game, but the game Γ itself may have a sequential structure. The normal form of Γ shows the payoffs induced by any combination of plans of actions of the players, if such plans are implemented. Some game theorists, including the founding fathers von Neumann and Morgenstern [84], argue that from a theoretical point of view all strategically relevant aspects of a game are contained in its normal form. Be that as it may, here by "static game" we specifically mean a game where players move *simultaneously*.

- $v_i : Y \to \mathbb{R}$ is the Von Neumann-Morgenstern utility function of player i.

Structure $\langle I, Y, (A_i)_{i\in I}, g\rangle$, that is, the game without the utility functions, is called **game form**. The game form represents the essential features of the rules of the game. A game is obtained by adding to the game form a profile of utility functions $(v_i)_{i\in I}$, which represents players' preferences over lotteries of consequences, according to expected utility calculations.

From the consequence function g and the utility function v_i of player i, we obtain a function that assigns to each $a = (a_j)_{j\in I}$ the utility $v_i(g(a))$ for player i of consequence $g(a)$. This function

$$u_i = v_i \circ g : \bigtimes_{i\in I} A_i \to \mathbb{R}$$

is called **payoff function** of player i. The reason why u_i is called payoff function is that the early work on game theory assumed that consequences are distributions of monetary payments, or payoffs, and that players are risk neutral, so that it is sufficient to specify, for each player i, the monetary payoff implied by each action profile. But in modern game theory $u_i(a) = v_i(g(a))$ is interpreted as the von Neumann-Morgenstern utility of outcome $g(a)$ for player i. If there are monetary consequences, then

$$g = (g_i)_{i\in I} : \bigtimes_{i\in I} A_i \to \mathbb{R}^I,$$

where $m_i = g_i(a)$ is the net gain of player i when a is played. Assuming that player i is selfish, function v_i is strictly increasing in m_i and constant in each m_j with $j \neq i$ (note that selfishness is not an assumption of game theory, it is an economic assumption that may be adopted in game theoretic models). Thus, function v_i captures the risk attitudes of player i. For example, i is strictly risk averse if and only if v_i is strictly concave.

In some games the outcome function is *stochastic*, that is, each $a = (a_j)_{j\in I}$ corresponds to a lottery. For example, in a sealed-bid auction the object on sale is assigned at random to one of the high bidders if there is a tie. In the stochastic outcome case, u_i has to be interpreted as an expected payoff: if Y is finite,

$$u_i(a) = \sum_{y\in Y} g(a)(y)v_i(y),$$

where $g(a)(y)$ is the probability of outcome y according to lottery $g(a)$.

Games are typically represented in the reduced form

$$G = \langle I, (A_i, u_i)_{i \in I} \rangle$$

which shows only the payoff functions. We often do the same. However, it is conceptually important to keep in mind that payoff functions are derived from a consequence function and utility functions.

For ease of exposition, we often assume that all the sets A_i $(i \in I)$ are *finite*; when the sets of actions are infinite, we explicitly say so. This simplifies the analysis of probabilistic beliefs. A static game with finite action sets is called a **finite static game** (or just a finite game, if "static" is clear from the context). The following assumption is a generalization of finiteness:

Definition 1. *A (reduced form) static game $G = \langle I, (A_i, u_i)_{i \in I} \rangle$ is* ***compact-continuous*** *if, for each $i \in I$, A_i is a compact subset of a Euclidean space $\mathbb{R}^{k_i}$ and $u_i : \times_{j \in I} A_j \to \mathbb{R}$ is continuous.*[2]

We can always assume without loss of generality that, if A_i is finite, then it is a finite subset of some Euclidean space. Every such set is compact, and every function with a finite domain is continuous. Therefore, a finite game is compact-continuous. Most of our results hold for compact-continuous games. In some cases (for example, to prove the existence of an equilibrium) we may have to add other assumptions about the actions sets and the payoff functions.

We call **profile** a list of objects $(x_i)_{i \in I}$, one object x_i from some set X_i for each player $i \in I$. In particular, an **action profile** is a list

$$a = (a_i)_{i \in I} \in A = \times_{i \in I} A_i.$$

The payoff function u_i numerically represents the preferences of player i among the different action profiles $a, a', a'', \ldots \in A$. The strategic interdependence is due to the fact that the outcome function g depends on the entire action profile a and, consequently, the utility that a generic individual i can achieve depends not only on his choice, but also on those

[2] A subset of a Euclidean space is compact if and only if it is closed and bounded. For most of the results that rely on compactness and continuity, it is sufficient to assume that each action set A_i is a compact subset of a metrizable topological space.

of other players.[3] To stress how the payoff of i depends on a variable under i's control as well as on a profile of variables controlled by other individuals, we denote by $-i = I\backslash\{i\}$ the set of individuals different from i, we define

$$A_{-i} = \bigtimes_{j\in I\backslash\{i\}} A_j,$$

and we write the payoff of i as a function of the two arguments $a_i \in A_i$ and $a_{-i} \in A_{-i}$, that is, $u_i : A_i \times A_{-i} \to \mathbb{R}$.

In order to be able to reach some conclusions regarding players' behavior in a game G, we impose two *minimal hypotheses* (further hypotheses will be introduced later on as necessary):

(H1) Each player i knows A_i, A_{-i} and his own payoff function $u_i : A_i \times A_{-i} \to \mathbb{R}$;

(H2) Each player is rational (see Chapter 3 for a more formal definition of rationality).

Hypothesis (H1) (i knows u_i) may seem almost tautological: is it not obvious that every individual should know his preferences over A? Indeed not, because action profiles are not consequences (outcomes). We assume that each i knows his own preferences over (lotteries of) consequences, represented by utility function $v_i : Y \to \mathbb{R}$, but he may ignore the true outcome function $g : A \to Y$, hence he may ignore his payoff function $u_i = v_i \circ g$. The following example clarifies this point.

[3] Note, in our formulation, players' utility depends only on the consequence determined by action profile a through outcome function g. Thus, we are assuming a form of *consequentialism*. If you think that consequentialism is an obvious assumption, think twice. Substantial experimental evidence suggests agents also care about other players' *intentions*: agent i may evaluate the same action profile differently depending on his belief about opponents' intentions. To provide a formal analysis of players' intentions, we need to specify what an agent thinks other agents will do, what an agent thinks other agents think other agents will do, what an agent thinks other agents think other agents think that other agents will do and so on. Although the analysis of these issues is beyond the scope of this textbook, the interested reader is referred to the growing literature on *psychological games* (see Geanakoplos *et al* [47], Battigalli and Dufwenberg [12], Battigalli, Corrao and Dufwenberg [18], and Battigalli and Dufwenberg [13]).

Example 1. *(Knowledge of the utility and payoff functions)* Players 1 and 2 work in a team for the production of a public good. Their action is their effort $a_i \in [0,1]$. The output, Q, depends on efforts according to a Cobb-Douglas production function[4]

$$Q = K(a_1)^{\theta_1}(a_2)^{\theta_2}.$$

The cost of effort measured in terms of public good is $C_i = (a_i)^2$. Then, the outcome function is

$$g(a_1, a_2) = (K(a_1)^{\theta_1}(a_2)^{\theta_2}, (a_1)^2, (a_2)^2).$$

The utility function of player i is $v_i(Q, C_1, C_2) = Q - C_i$. The payoff function is

$$u_i(a_1, a_2) = v_i(g(a_1, a_2)) = K(a_1)^{\theta_1}(a_2)^{\theta_2} - (a_i)^2.$$

If i does not know all the parameters K, θ_1 and θ_2 (as well as the functional forms above), then he does not know the payoff function u_i, even if he knows the utility function v_i. ▲

Example 1 also clarifies the difference between "game" in the technical sense of game theory, and the "rules of the game" as understood by the layperson (see the Chapter 1). In the example above, the "rules of the game" specify the action set $[0,1]$ and the outcome function

$$(a_1, a_2) \longmapsto (Q, C_1, C_2) = \left(K(a_1)^{\theta_1}(a_2)^{\theta_2}, (a_1)^2, (a_2)^2\right).$$

In order to complete the description of the game in the technical sense, we also have to specify the utility functions $v_i(Q, C_1, C_2) = Q - C_i$ $(i = 1, 2)$.[5] Also, the example clarifies that the outcome map $a \mapsto g(a)$ may depend on personal features of the agents playing the game: in this case, the productivity θ_i of each agent affects how the pair of efforts (a_1, a_2) (say, the hours worked) is mapped to the output Q. Therefore, the outcome function g is not always fully determined by what we would call "rules of the game" in the natural language.

---

[4] In some applications players exert effort in a joint project and output q is interpreted as the probability that the project is successful.

[5] This specification assumes that players are risk neutral with respect to their consumption of public good.

Unlike the games people play for fun (and those that experimental subjects play in most game theoretic experiments), in many economic games the outcome function may not be fully known. In Example 1, for instance, it is possible that player i knows his own productivity parameter θ_i, but does not know K or θ_{-i}. Thus, assumption (H1) is substantive: i may not know g and hence he may not know $u_i = v_i \circ g$.

The **complete information** assumption that we will consider later on (e.g., in Chapter 4) is much stronger; recall from the Chapter 1 that there is complete information if the rules of the game and players' preferences over (lotteries over) consequences are common knowledge. Although, as we explained above, the rules of the game may be only partially known, there are still many interactive situations where it is reasonable to assume that they are not only known, but indeed commonly known. Yet, assuming common knowledge of players' *preferences* is often far-fetched. Thus, complete information should be thought of as an idealization that simplifies the analysis of strategic thinking. Chapter 8 will introduce the formal tools necessary to model the absence of complete information.

3

Rationality and Dominance

In this chapter we analyze static games from the perspective of *decision theory*. We take as given the conjecture[1] of a player about opponents' behavior. This conjecture is, in general, probabilistic, i.e., it can be expressed by assigning subjective probabilities to the different action profiles of the other players. A player is said to be rational if he maximizes his *expected payoff* given his conjecture. The concept of *dominance* allows to characterize which actions are consistent with the rationality assumption without knowing a player's conjecture.

3.1 Probabilities and Expected Payoff

Payoff function u_i, which is a derived utility function, implicitly represents the preferences of i over different probability measures on A, which induce corresponding **lotteries**, that is, probability measures over consequences. The preferred lotteries are those that yield a higher expected payoff. We explain this in detail below.

We first introduce some preliminary definitions about sets and functions. For any set X, we let 2^X denote the **power set** of X, i.e., the collection of all subsets of X, including the empty set $\emptyset$ and X itself. If X is finite, then $|2^X| = 2^{|X|}$, where $|\cdot|$ denotes the "cardinality of $\cdot$." For any pair of nonempty sets X and Y, we let Y^X denote the set of

[1] In general, we use the word "conjecture" to refer to a player's belief about variables that affect his payoff and are beyond his control, such as the actions of other players in a static game.

functions with domain X and codomain Y. If X and Y are both finite, then $|Y^X| = |Y|^{|X|}$.

The set of probability measures over a generic domain X is denoted by $\Delta(X)$. If X is finite[2]

$$\Delta(X) = \left\{ \mu \in \mathbb{R}_+^X : \sum_{x \in X} \mu(x) = 1 \right\}.$$

The definition above is the simplest one for finite domains. But there is an alternative, equivalent, definition, that can be more easily generalized to infinite domains. Since sometimes we will consider probability measures on infinite domains, we present here also the alternative. First consider the definition above for a finite X, and fix $\mu \in \Delta(X)$. Every subset of a finite uncertainty space X is called "event." For each event $E \subseteq X$, μ determines the probability of E as follows:

$$\mu(E) = \sum_{x \in E} \mu(x).$$

Thus, the map $E \mapsto \mu(E)$ satisfies the following properties:

- (*normalization*) $\mu(X) = 1$,
- (*additivity*) for all $E, F \in 2^X$, $E \cap F = \emptyset \Rightarrow \mu(E \cup F) = \mu(E) + \mu(F)$.

Note that additivity implies that $\mu(\emptyset) = 0$, because $\emptyset \cap E = \emptyset$, hence $\mu(E) = \mu(E \cup \emptyset) = \mu(E) + \mu(\emptyset)$, which implies $\mu(\emptyset) = 0$. These are the characterizing properties of a probability measure, defined as a function $\mu : 2^X \to [0, 1]$. Therefore, the alternative definition of $\Delta(X)$ is

$$\Delta(X) = \Big\{ \mu \in [0,1]^{2^X} : \mu(X) = 1, \\ \forall E, F \in 2^X, E \cap F = \emptyset \Rightarrow \mu(E \cup F) = \mu(E) + \mu(F) \Big\}.$$

When X is infinite, we typically restrict our attention to a sub-collection $\mathcal{X} \subseteq 2^X$ of subsets of X, called **events** (e.g., the Borel sets, when X is a subset of a Euclidean space, or—more generally—a metric space), and we strengthen the additivity property as follows:

---

[2] Recall that $\mathbb{R}_+^X$ is the set of nonnegative real-valued functions defined over the domain X. If X is the finite set $\{x_1, \ldots, x_n\}$, $\mathbb{R}_+^X$ is isomorphic to $\mathbb{R}_+^n$, the positive orthant of the Euclidean space $\mathbb{R}^n$.

- (*countable additivity*) if $(E_k)_{k=1}^{\infty} \in \mathcal{X}^{\mathbb{N}}$ is a sequence of pairwise disjoint events and $\bigcup_{k=1}^{\infty} E_k \in \mathcal{X}$, then

$$\mu\left(\bigcup_{k=1}^{\infty} E_k\right) = \sum_{k=1}^{\infty} \mu(E_k).$$

Whenever we consider infinite sets, we take for granted that there is a well-defined collection of events $\mathcal{X}$ (such as the Borel subsets of a metric space) and $\Delta(X)$ is the set of functions $\mu : \mathcal{X} \to [0,1]$ that satisfy normalization and countable additivity.

Note, *X can be regarded as a subset of $\Delta(X)$*: a point $x \in X$ corresponds to the degenerate probability measure δ_x, called **Dirac measure** (or deterministic, or dogmatic belief), that satisfies $\delta_x(E) = 1$ if and only if $x \in E$. Hence, with a slight abuse of notation, we can write $X \subseteq \Delta(X)$ (more abuses of notation will follow).[3]

Definition 2. *Consider a probability measure $\mu \in \Delta(X)$, where X is a* finite *set. The subset of the elements $x \in X$ to which μ assigns a positive probability is called **support** of μ and is denoted by*

$$\text{supp}\mu = \{x \in X : \mu(x) > 0\}.$$

If X is a closed subset of $\mathbb{R}^m$, we define the support of a probability measure $\mu \in \Delta(X)$ as the intersection of all closed sets C such that $\mu(C) = 1$.[4]

Suppose that there is a function $f : X \to Y$, where Y is a finite set of outcomes, or consequences,[5] and an individual has a preference relation $\succsim$ over the set of lotteries $\Delta(Y)$ represented by a Von Neumann-Morgenstern utility function $v : Y \to \mathbb{R}$, that is, for all lotteries $\lambda, \lambda' \in \Delta(Y)$,

$$\lambda \succsim \lambda' \Leftrightarrow \sum_{y\in Y} \lambda(y)v(y) \geq \sum_{y\in Y} \lambda'(y)v(y).$$

---

[3] If X has at least two elements, the inclusion is strict.

[4] If $X \subset \mathbb{R}^m$ is finite, this definition is equivalent to the previous one.

[5] If X is infinite, then it has to be endowed with a sigma-algebra of events $\mathcal{X} \subseteq 2^X$ (see Appendix 3.4) and f has to be $\mathcal{X}$-measurable, that is, $f^{-1}(y) \in \mathcal{X}$ for each outcome y in the finite set Y.

Then we can derive a corresponding preference relation on X as follows. First we define the **pushforward** function $\hat{f} : \Delta(X) \to \Delta(Y)$ that determines the lottery induced, through f, by a probability measure on X:

$$\hat{f}(\mu)(y) = \mu\left(f^{-1}(y)\right) = \sum_{x \in f^{-1}(y)} \mu(x),$$

where, of course, the second equality holds when X is finite. With this, we obtain a preference relation on $\Delta(X)$: for all $\mu, \mu' \in \Delta(X)$,

$$\mu \succsim \mu' \Leftrightarrow \sum_{y \in Y} \hat{f}(\mu)(y) v(y) \geq \sum_{y \in Y} \hat{f}(\mu')(y) v(y).$$

Consider the particular case where $X = A$ is the set of action profiles in a game, let $f = g : A \to Y$ be the consequence function, and fix any player i. Then, for all $\mu, \mu' \in \Delta(A)$,

$$\mu \succsim_i \mu' \Leftrightarrow \sum_{y \in Y} \hat{g}(\mu)(y) v_i(y) \geq \sum_{y \in Y} \hat{g}(\mu')(y) v_i(y).$$

Since $\hat{g}(\mu)(y) = \mu\left(g^{-1}(y)\right)$ and we have defined the payoff function as $u_i = v_i \circ g$, we have

$$\sum_{y \in Y} \hat{g}(\mu)(y) v_i(y) = \sum_{y \in Y} \sum_{a \in g^{-1}(y)} \mu(a) v_i(g(a)) = \sum_{a \in A} \mu(a) u_i(a)$$

for each $\mu \in \Delta(A)$ (to understand the second equality, note that $\{g^{-1}(y)\}_{y \in Y}$ is a partition of A, and recall that $u_i(a) = v_i(g(a))$). Therefore

$$\mu \succsim_i \mu' \Leftrightarrow \sum_{a \in A} \mu(a) u_i(a) \geq \sum_{a \in A} \mu'(a) u_i(a).$$

In other words, probability measures on A are ranked according to the corresponding expected payoffs.

Probabilities and expected payoffs in compact-continuous games
The foregoing analysis can be extended to infinite games.[6] In particular, we focus on compact-continuous games. For any compact subset X of a Euclidean space, we let $\mathcal{X} = \mathcal{B}(X)$ be the smallest sigma-algebra

[6] We elaborate on this topic in Section 3.4.1 of the Appendix.

containing all the closed subsets of X; $\Delta(X)$ is the set of all the probability measures $\mu : \mathcal{B}(X) \to [0,1]$. Let $A = \bigtimes_{i\in I} A_i$ and Y be compact subsets of Euclidean spaces. If outcome function $g : A \to Y$ is continuous and each utility function $v_i : Y \to \mathbb{R}$ $(i \in I)$ is also continuous, then

$$\left\langle I, (A_i, u_i)_{i\in I} \right\rangle = \left\langle I, (A_i, v_i \circ g)_{i\in I} \right\rangle$$

is a compact-continuous game. Since g is continuous, then it is also measurable, that is, $E \in \mathcal{B}(Y)$ implies $g^{-1}(E) \in \mathcal{B}(A)$ for each $E \subseteq Y$. With this, the pushforward function, $\hat{g} : \Delta(A) \to \Delta(Y)$, is defined as follows: for all $\mu \in \Delta(A)$ and $E \in \mathcal{B}(Y)$,

$$\hat{g}(\mu)(E) = \mu\left(g^{-1}(E)\right);$$

more compactly, the pushforward map is described as follows

$$\mu \mapsto \hat{g}(\mu) = \mu \circ g^{-1}.$$

When the probability measures $\mu, \mu' \in \Delta(A)$ have finite or countable supports, the expected utility formulas are almost as above, with $y \in Y$ replaced by $y \in g(\text{supp}\mu)$ and $y \in g(\text{supp}\mu')$ respectively, and $a \in A$ replaced by $a \in$suppμ and $a \in$suppμ' respectively. Otherwise, expected values are expressed as integrals:

$$\mathbb{E}_\mu(u_i) = \int_A u_i(a)\,\mu(\mathrm{d}a) = \int_Y v_i(y)\,\hat{g}(\mu)(\mathrm{d}y) = \mathbb{E}_{\hat{g}(\mu)}(v_i). \tag{3.1.1}$$

Example 2. To get some intuition about the generalized integrals in eq. (3.1.1), suppose that player i has only one opponent, whose action space is an interval $A_{-i} = \left[\underline{a}_{-i}, \overline{a}_{-i}\right] \subset \mathbb{R}$. Also suppose that player i plans to choose action $\hat{a}_i$, therefore his belief $\mu \in \Delta(A_i \times A_{-i})$ assigns probability one to the segment $\{\hat{a}_i\} \times \left[\underline{a}_{-i}, \overline{a}_{-i}\right]$ and μ corresponds to a probability measure μ^i on the interval $\left[\underline{a}_{-i}, \overline{a}_{-i}\right]$, the "conjecture" of i about $-i$. Such probability measures can equivalently be represented by cumulative distribution functions (cdf). So, let F_μ be the cdf representing μ^i; in particular,

$$F_\mu(a_{-i}) = \mu^i\left(\left[\underline{a}_{-i}, a_{-i}\right]\right) = \mu\left(\{\hat{a}_i\} \times \left[\underline{a}_{-i}, a_{-i}\right]\right).$$

Then $\mathbb{E}_\mu(u_i)$ can be expressed as a Riemann-Stieltjes integral:

$$\int_A u_i(a)\,\mu(\mathrm{d}a) = \int_{\underline{a}_{-i}}^{\overline{a}_{-i}} u_i(\hat{a}_i, a_{-i})\,\mathrm{d}F_\mu(a_{-i}).$$

If μ^i has a finite support, then F_μ is piecewise constant and suppμ^i is the set of points where F_μ is discontinuous (from the left). In this case,

$$\int_{\underline{a}_{-i}}^{\overline{a}_{-i}} u_i\left(\hat{a}_i, a_{-i}\right) \mathrm{d}F_\mu\left(a_{-i}\right) = \sum_{a_{-i}\in \mathrm{supp}\mu^i} u_i\left(\hat{a}_i, a_{-i}\right) \mu^i\left(a_{-i}\right).$$

If instead suppμ^i is an interval (or a union of intervals) and F_μ is differentiable with integrable derivative f_μ, then the differential formula $\mathrm{d}F_\mu\left(a_{-i}\right) = f_\mu\left(a_{-i}\right)\mathrm{d}a_{-i}$ applies, and

$$\int_{\underline{a}_{-i}}^{\overline{a}_{-i}} u_i\left(\hat{a}_i, a_{-i}\right) \mathrm{d}F_\mu\left(a_{-i}\right) = \int_{\underline{a}_{-i}}^{\overline{a}_{-i}} u_i\left(\hat{a}_i, a_{-i}\right) f_\mu\left(a_{-i}\right)\mathrm{d}a_{-i}.$$

▲

3.2 Conjectures

The reason why one needs to introduce preferences over lotteries and expected payoffs is that player i cannot observe other players' actions (a_{-i}) before making his own choice. Hence, he needs to form a conjecture about such actions. If i were certain of the other players' choices, then one could represent i's conjecture simply with an action profile $a_{-i} \in A_{-i}$. However, in general, i might be uncertain about other players' actions and assign a strictly positive (subjective) probability to several profiles a_{-i}, a'_{-i}, etc.

Definition 3. *A* ***conjecture*** *of player i is a (subjective) probability measure $\mu^i \in \Delta(A_{-i})$. A* **deterministic conjecture** *is a probability measure $\mu^i \in \Delta(A_{-i})$ that assigns probability one to a particular action profile, that is, $\mu^i = \delta_{a_{-i}}$ for some $a_{-i} \in A_{-i}$.*

Note, we call "conjecture" a (probabilistic) belief about the behavior of other players, while we use the term "belief" to refer to a more general type of uncertainty.

Observation 1. *The set of deterministic conjectures of player i essentially coincides with the set of other players' action profiles, so that we can write $A_{-i} \subseteq \Delta(A_{-i})$.*

One of the most interesting aspect of game theory consists in determining players' conjectures, or, at least, in narrowing down the set of possible conjectures, combining some general assumptions about players' rationality and beliefs with specific assumptions about the given game G. However, in this chapter we will not try to "explain" players' conjectures. This is left for the following chapters.

For any given conjecture μ^i, the choice of a particular action $\bar{a}_i$ corresponds to the choice of the measure on A that assigns probability $\mu^i(a_{-i})$ to each action profile $(\bar{a}_i, a_{-i})$ $(a_{-i} \in A_{-i})$ and probability zero to the profiles (a_i, a_{-i}) with $a_i \neq \bar{a}_i$. Therefore, if a player i has conjecture μ^i and chooses (or plans to choose) action $\bar{a}_i$, the corresponding (subjective) expected payoff is

$$u_i(\bar{a}_i, \mu^i) = \mathbb{E}_{\mu^i}\left(u_i\left(\bar{a}_i.\cdot\right)\right),$$

where

$$\mathbb{E}_{\mu^i}\left(u_i\left(\bar{a}_i.\cdot\right)\right) = \sum_{a_{-i}\in \text{supp}\mu^i} \mu^i(a_{-i})u_i(\bar{a}_i, a_{-i})$$

if μ^i has a finite support.

There are many different ways to represent graphically actions, conjectures, and preferences when the number of available actions is small. Let us focus our attention on the instructive case where player i's *opponent* $(-i)$ has *only two actions*, denoted ℓ ("left") and r ("right"). Consider, for instance, the function u_i represented by the payoff matrix in Figure 3.1 for player i.

$i\backslash -i$	ℓ	r
a	4	1
b	1	4
c	2	2
e	4	0
f	1	1

Figure 3.1: Matrix 1.

Given that $-i$ has only two actions, we can represent the conjecture

μ^i of i about $-i$ with a single number: $\mu^i(r)$, the subjective probability that i assigns to r. Each action a_i corresponds to a function that assigns to each value of $\mu^i(r)$ the expected payoff of a_i. Since the expected payoff function $(1-\mu^i(r))u_i(a_i,\ell)+\mu^i(r)u_i(a_i,r)$ is linear in the probabilities,[7] every action is represented by a line, such as aa, bb, cc etc., in Figure 3.2.

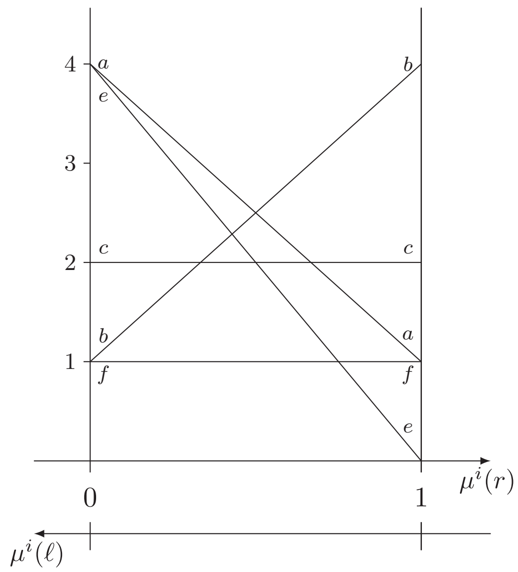

Figure 3.2: Expected payoff as a function of beliefs.

From Figure 3.2 it is apparent that only actions a, b, and e are "justifiable" by some conjecture. In particular, if $\mu^i(r)=0$, then a and e are both optimal; if $0<\mu^i(r)<\frac{1}{2}$, then a is the only optimal action (the line aa yields the maximum expected payoff). If $\mu^i(r)>\frac{1}{2}$, then b is the only optimal action (the line bb yields the maximum expected payoff). If $\mu^i(r)=\frac{1}{2}$, then there are two optimal actions: a and b.

[7]More precisely, the expected payoff function $u_i(a_i,\mu^i)$ is affine in μ^i, that is,

$$u_i(a_i,x\mu^i+(1-x)\nu^i)=xu_i(a_i,\mu^i)+(1-x)u_i(a_i,\nu^i)$$

for all $a_i\in A_i$, $\mu^i,\nu^i\in\Delta(A_{-i})$, and $x\in[0,1]$.

3.2.1 Mixed Actions

In principle, instead of choosing directly a certain action, a player could delegate his decision to a randomizing device, such as spinning a roulette wheel, or tossing a coin. In other words, a player could simply choose the *probability* of playing any given action.

Definition 4. *A random choice by player i, also called* ***mixed action****, is a probability measure $\alpha_i \in \Delta(A_i)$. An action $a_i \in A_i$ is also called* ***pure action****.*

Observation 2. *The set of pure actions can be regarded as a subset of the set of mixed actions, i.e., $A_i \subseteq \Delta(A_i)$ (cf. Observation 1).*

It is assumed that (according to i's beliefs) the random draw of an action of i is stochastically independent of the other players' actions. For example, the following situation is *excluded*: i chooses his action according to the (random) weather and he thinks his opponents are doing the same, so that there is correlation between a_i and a_{-i} even though there is no causal link between a_i and a_{-i} (this type of correlation will be discussed in section 6.2 of Chapter 6). More importantly, player i knows that moves are simultaneous and therefore by changing his actions he cannot cause any change in the probability distribution of opponents' actions. Hence, if player i has conjecture μ^i and chooses the mixed action α_i, where both have finite support, the subjective probability of each possible action profile (a_i, a_{-i}) is $\alpha_i(a_i)\mu^i(a_{-i})$ and i's expected payoff is

$$\begin{aligned} u_i(\alpha_i, \mu^i) &= \mathbb{E}_{\alpha_i \times \mu^i}(u_i) = \sum_{a_i \in \mathrm{supp}\alpha_i} \sum_{a_{-i} \in \mathrm{supp}\mu^i} \alpha_i(a_i)\mu^i(a_{-i})u_i(a_i, a_{-i}) \\ &= \sum_{a_i \in \mathrm{supp}\alpha_i} \alpha_i(a_i)u_i(a_i, \mu^i). \end{aligned}$$

If the opponent has only two feasible actions, it is possible to use a graph to represent the lotteries corresponding to pure and mixed actions. For each action a_i, we consider a corresponding point in the Cartesian plane with coordinates given by the utilities that i obtains for each of the two actions of the opponent. If the actions of the opponents are ℓ and r, we denote such coordinates $x = u_i(\cdot, \ell)$ and $y = u_i(\cdot, r)$. Any *pure* action a_i corresponds to the vector $(x, y) = (u_i(a_i, \ell), u_i(a_i, r))$ (a row in the payoff matrix of

i). The same holds for the mixed actions: α_i corresponds to the vector $(x, y) = (u_i(\alpha_i, \ell), u_i(\alpha_i, r))$. The set of points (vectors) corresponding to the mixed actions is simply the *convex hull* of the points corresponding to the pure actions.[8] Figure 3.3 represents such a set for the matrix in Figure 3.1.

How can conjectures be represented in such a figure? It is quite simple. Every conjecture induces a preference relation on the space of payoff pairs. Such preferences can be represented through a map of iso-expected payoff curves (or indifference curves). Let (x, y) be the generic vector of (expected) payoffs corresponding to ℓ and r respectively. The expected payoff induced by conjecture μ^i is $\mu^i(\ell)x + \mu^i(r)y$. Therefore, the indifference curves are straight lines defined by the equation $y = \frac{\overline{u}}{\mu^i(r)} - \frac{\mu^i(\ell)}{\mu^i(r)}x$, where $\overline{u}$ denotes the constant expected payoff. Every conjecture μ^i corresponds to a set of parallel lines with negative slope (or, in the extreme cases, horizontal or vertical slope) determined by the orthogonal vector $(\mu^i(\ell), \mu^i(r))$. The direction of increasing expected payoff is given precisely by such orthogonal vector. The optimal actions (pure or mixed) can be graphically determined considering, for any conjecture μ^i, the highest line of iso-expected payoff among those that touch the set of *feasible* payoff vectors (the shaded area in Figure 3.3).

Allowing for mixed actions, the set of feasible (expected) payoffs vectors is a convex polyhedron (as in Figure 3.3). To see this, note that if α_i and β_i are mixed actions, then also $p \cdot \alpha_i + (1-p) \cdot \beta_i$ (with $0 \leq p \leq 1$) is a mixed action, where $p \cdot \alpha_i + (1-p) \cdot \beta_i$ is the function that assigns to each pure action a_i the weight $p\alpha_i(a_i) + (1-p)\beta_i(a_i)$. Thus, all the convex combinations of feasible payoff vectors are feasible payoff vectors, once we allow for randomization. This geometrical intuition can be extended to all *finite* games. In general, each mixed action α_i corresponds to a vector of expected payoffs $(u_i(\alpha_i, a_{-i}))_{a_{-i} \in A_{-i}}$. Let

$$\mathbf{U} = \left\{ \mathbf{u} \in \mathbb{R}^{A_{-i}} : \exists \alpha_i \in \Delta(A_i), \mathbf{u} = (u_i(\alpha_i, a_{-i}))_{a_{-i} \in A_{-i}} \right\}$$

be the set of feasible payoff vectors when randomization is allowed; then $\mathbf{U}$ is a compact and convex polyhedron in $\mathbb{R}^{A_{-i}}$ with extreme points contained

[8] The **convex hull** of a set of points $X \subseteq \mathbb{R}^k$ is the smallest convex set containing X, that is, the intersection of all the convex sets containing X.

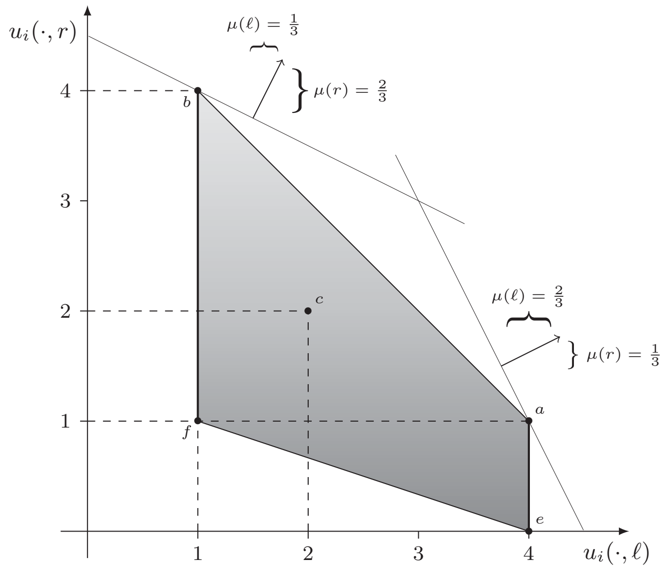

Figure 3.3: State-contingent representation of expected payoff.

in the set

$$\left\{\mathbf{u} \in \mathbb{R}^{A_{-i}} : \exists a_i \in A_i, \mathbf{u} = (u_i(a_i, a_{-i}))_{a_{-i} \in A_{-i}}\right\}$$

of payoff vectors induced by pure actions. For example, in Figure 3.3 $\mathbf{U}$ is the shaded area; the set of extreme points is

$$\{(u_i(a,\ell), u_i(a,r)), (u_i(b,\ell), u_i(b,r)), (u_i(e,\ell), u_i(e,r)), (u_i(f,\ell), u_i(f,r))\},$$

which is a subset of $\{(u_i(a_i,\ell), u_i(a_i,r))\}_{a_i \in A_i}$ because the vector

$$(u_i(c,\ell), u_i(c,r)) = (2,2) = \frac{1}{3}(1,1) + \frac{1}{3}(4,1) + \frac{1}{3}(1,4)$$
$$= \frac{1}{3}(u_i(f,\ell), u_i(f,r)) + \frac{1}{3}(u_i(a,\ell), u_i(a,r)) + \frac{1}{3}(u_i(b,\ell), u_i(b,r))$$

lies in the interior of $\mathbf{U}$.

However, the idea that players use coins or roulette wheels to randomize their choices may seem weird and unrealistic. Furthermore, as illustrated by Figure 3.3, for any conjecture μ^i and any mixed action α_i, there is always a pure action a_i that yields the same or a higher expected payoff than α_i (check all possible slopes of the iso-expected payoff curves and verify how the set of optimal points looks like).[9] Hence, a player cannot be *strictly* better off by choosing a mixed action rather than a pure action.

The point of view we adopt in this textbook is that expected-utility maximizing players never randomize (although their choices might depend on extrinsic, payoff-irrelevant signals, as in Chapter 6, Section 6.2). Nonetheless, it will be shown that in order to assess the justifiability of a given pure action it is analytically convenient to introduce mixed actions. (In Chapter 5, we discuss interpretations of mixed actions that do not involve randomization.)

3.3 Best Replies and Undominated Actions

Definition 5. *A (mixed) action α_i^* is a* ***best reply*** *to conjecture μ^i if*

$$\forall \alpha_i \in \Delta(A_i),\ u_i(\alpha_i^*, \mu^i) \geq u_i(\alpha_i, \mu^i),$$

that is

$$\alpha_i^* \in \arg \max_{\alpha_i \in \Delta(A_i)} u_i(\alpha_i, \mu^i).$$

The set of pure *actions that are best replies to conjecture μ^i is denoted by*[10]

$$r_i(\mu^i) = A_i \cap \arg \max_{\alpha_i \in \Delta(A_i)} u_i(\alpha_i, \mu^i).$$

The correspondence $r_i : \Delta(A_{-i}) \rightrightarrows A_i$ is called ***best reply correspondence.***[11] *An action a_i is called* ***justifiable*** *if there exists a conjecture $\mu^i \in \Delta(A_{-i})$ such that $a_i \in r_i(\mu^i)$.*

Note that, even if A_i is finite, one cannot conclude without proof that the set of pure best replies to a conjecture μ^i is nonempty, i.e., that $r_i(\mu^i) \neq \emptyset$. In principle, it could happen that in order to maximize

[9] A formal proof of this result will be provided in the next section.

[10] Recall: we regard A_i is a subset of $\Delta(A_i)$ (see Observation 2).

[11] A *correspondence* $\varphi : X \rightrightarrows Y$ is a multi-function that assigns to every element $x \in X$ a set of elements $\varphi(x) \subseteq Y$. A correspondence $\varphi : X \rightrightarrows Y$ can be equivalently expressed as a function with domain X and codomain 2^Y, the power set of Y. We allow $\varphi(x)$ to be the empty subset of Y.

expected payoff it is necessary to use mixed actions that assign positive probability to more than one pure action. However, we will show that $r_i(\mu^i) \neq \emptyset$ for every μ^i provided that A_i is finite (or, more generally, that A_i is compact and u_i continuous in a_i, two properties that trivially hold if A_i is finite).

The following result shows, as anticipated, that a rational player does not need to use mixed actions. Therefore, it can be assumed without loss of generality that his choice is restricted to the set of pure actions.

Lemma 1. *Consider a finite or compact-continuous game. Fix arbitrarily a player $i \in I$, a conjecture $\mu^i \in \Delta(A_{-i})$ and a mixed action α_i^*; then α_i^* is a best reply to μ^i if and only if every pure action in the support of α_i^* is a best reply to μ^i, that is*

$$\alpha_i^* \in \arg \max_{\alpha_i \in \Delta(A_i)} u_i(\alpha_i, \mu^i) \ \Leftrightarrow \ \mathrm{supp}\alpha_i^* \subseteq r_i(\mu^i).$$

Proof. We prove the result only for finite games. In the finite case the proof is quite simple. But since this is the first proof of this textbook, we will go over it rather slowly.

(Only if) We first show that if $\mathrm{supp}\alpha_i^*$ is not included in $r_i(\mu^i)$, then α_i^* is not a best reply to μ^i (this is the contrapositive of the "only if" implication).[12] Let a_i^* be a pure action such that $\alpha_i^*(a_i^*) > 0$ and assume that, for some α_i, $u_i(\alpha_i, \mu^i) > u_i(a_i^*, \mu^i)$, so that $a_i^* \in \mathrm{supp}\alpha_i^* \backslash r_i(\mu^i)$.[13] Since $u_i(\alpha_i, \mu^i)$ is a weighted average of the values $\left(u_i(a_i, \mu^i)\right)_{a_i \in A_i}$, there must be a pure action a_i' such that $u_i(a_i', \mu^i) > u_i(a_i^*, \mu^i)$. But then we can construct a mixed action α_i' that satisfies $u_i(\alpha_i', \mu^i) > u_i(\alpha_i^*, \mu^i)$ by "shifting probability weight" from a_i^* to a_i', which is possible because $\alpha_i^*(a_i^*) > 0$. Specifically, for every $a_i \in A_i$, let

$$\alpha_i'(a_i) = \begin{cases} 0, & \text{if } a_i = a_i^*, \\ \alpha_i^*(a_i') + \alpha_i^*(a_i^*), & \text{if } a_i = a_i', \\ \alpha_i^*(a_i), & \text{if } a_i \neq a_i^*, a_i'. \end{cases}$$

[12] The **contrapositive** of $p \Rightarrow q$ is $\neg q \Rightarrow \neg p$, where p and q are sentences and $\neg$ means "not." An implication $p \Rightarrow q$ holds if and only if its contrapositive $\neg q \Rightarrow \neg p$ holds. In this case, p says that $\alpha_i^* \in \arg\max_{\alpha_i \in \Delta(A_i)} u_i(\alpha_i, \mu^i)$, q says that $\mathrm{supp}\alpha_i^* \subseteq r_i(\mu^i)$, and $\neg q$ says that there is some $a_i^* \in \mathrm{supp}\alpha_i^*$ such that $a_i^* \notin r_i(\mu^i)$.

[13] Recall that $X \backslash Y$ denotes the set of elements of X that do not belong to Y: $X \backslash Y = \{x \in X : x \notin Y\}$.

α'_i is a mixed action since $\sum_{a_i \in A_i} \alpha'_i(a_i) = \sum_{a_i \in A_i} \alpha^*_i(a_i) = 1$. Moreover, it can be easily checked that

$$u_i(\alpha'_i, \mu^i) - u_i(\alpha^*_i, \mu^i) = \alpha^*_i(a^*_i)[u_i(a'_i, \mu^i) - u_i(a^*_i, \mu^i)] > 0,$$

where the inequality holds by assumption. Thus α^*_i is not a best reply to μ^i.

(If) Next we show that if each pure action in the support of α^*_i is a best reply, then α^*_i is also a best reply. It is convenient to introduce the following notation:

$$\hat{u}_i(\mu^i) = \max_{\alpha_i \in \Delta(A_i)} u_i(\alpha_i, \mu^i).$$

By definition of $r_i(\mu^i)$,

$$\forall a_i \in r_i(\mu^i), \ \hat{u}_i(\mu^i) = u_i(a_i, \mu^i) \tag{3.3.1}$$
$$\forall a_i \in A_i, \ \hat{u}_i(\mu^i) \geq u_i(a_i, \mu^i). \tag{3.3.2}$$

Assume that $\text{supp}\alpha^*_i \subseteq r_i(\mu^i)$. Then, for every $\alpha_i \in \Delta(A_i)$

$$\begin{aligned} u_i(\alpha^*_i, \mu^i) &= \sum_{a_i \in \text{supp}\alpha^*_i} \alpha^*_i(a_i) u_i(a_i, \mu^i) = \sum_{a_i \in r_i(\mu^i)} \alpha^*_i(a_i) u_i(a_i, \mu^i) = \hat{u}_i(\mu^i) \\ &= \hat{u}_i(\mu^i) \sum_{a_i \in A_i} \alpha_i(a_i) = \sum_{a_i \in A_i} \alpha_i(a_i) \hat{u}_i(\mu^i) \geq \sum_{a_i \in A_i} \alpha_i(a_i) u_i(a_i, \mu^i). \end{aligned}$$

The first equality holds by definition, the second follows from $\text{supp}\alpha^*_i \subseteq r_i(\mu^i)$, the third holds by eq. (3.3.1), the fourth and fifth are obvious and the inequality follows from (3.3.2). ■

In the matrix of Figure 3.1, for example, any mixed action that assigns positive probability only to a and/or b is a best reply to the uniform conjecture $\mu^i = \frac{1}{2}\delta_\ell + \frac{1}{2}\delta_r$ that assigns probability $\frac{1}{2}$ to ℓ and r.[14] Clearly, the set of pure best replies to the uniform conjecture is $r_i\left(\frac{1}{2}\delta_\ell + \frac{1}{2}\delta_r\right) = \{a, b\}$.

Note that if at least one pure action is a best reply among all pure and mixed actions, then the maximum that can be attained by constraining

[14] Recall that δ_x denotes the Dirac probability measure that assigns probability 1 to point x.

the choice to pure actions is necessarily equal to what could be attained choosing among (pure and) mixed actions, i.e.,

$$r_i(\mu^i) \neq \emptyset \;\Rightarrow \max_{\alpha_i \in \Delta(A_i)} u_i(\alpha_i, \mu^i) = \max_{a_i \in A_i} u_i(a_i, \mu^i).$$

This observation along with Lemma 1 yields the following:

Corollary 1. *Consider a finite or compact-continuous game. Fix arbitrarily a player $i \in I$ and a conjecture $\mu^i \in \Delta(A_{-i})$; then*

$$r_i(\mu^i) = \arg\max_{a_i \in A_i} u_i(a_i, \mu^i).$$

Hence, it is not necessary to use mixed actions to maximize expected payoff. Moreover, the best reply correspondence is nonempty-valued, that is, $r_i(\mu^i) \neq \emptyset$ for every $\mu^i \in \Delta(A_{-i})$.

Proof. For every conjecture μ^i, the expected payoff function $u_i(\cdot, \mu^i) : \Delta(A_i) \to \mathbb{R}$ is continuous in α_i and the domain $\Delta(A_i)$ is compact.[15] Hence, the function has at least one maximizer α_i^*. By Lemma 1, $\text{supp}\alpha_i^* \subseteq r_i(\mu^i)$, so that $r_i(\mu^i) \neq \emptyset$. As we have seen above, this implies that

$$\max_{\alpha_i \in \Delta(A_i)} u_i(\alpha_i, \mu^i) = \max_{a_i \in A_i} u_i(a_i, \mu^i)$$

and therefore $r_i(\mu^i) = \arg\max_{a_i \in A_i} u_i(a_i, \mu^i)$. ■

Recall that, for a given function $f : X \to Y$ and a given subset $C \subseteq X$,

$$f(C) = \{y \in Y : \exists x \in C, y = f(x)\}$$

denotes the set of images of elements of C. Analogously, for a given correspondence $\psi : X \rightrightarrows Y$,

$$\psi(C) = \{y \in Y : \exists x \in C, y \in \psi(x)\}$$

denotes the set of elements y that belong to the image $\psi(x)$ of some point $x \in C$. In particular, we use this notation for the best reply

[15] For the infinite case, see Appendix 3.4. If A_i is finite, $\Delta(A_i)$ is a compact subset of a Euclidean space.

correspondence. For example, $r_i\left(\Delta(A_{-i})\right)$ is the set of justifiable actions of player i (see Definition 5).

A question should spontaneously arise at this point: can we characterize the set of justifiable actions with no reference to conjectures and expected payoff maximization? In other words, can we conclude that an action will not be chosen by a rational player without checking directly that it is not a best reply to any conjecture? The answer comes from the concept of dominance. But, even if we are only concerned with the justifiability of pure actions, we will have to compare them with mixed actions. To anticipate: a (pure) action is justifiable if and only it is not dominated by any mixed action.

Definition 6. *A mixed action α_i **dominates** a (pure or) mixed action β_i if it yields a* strictly higher *expected payoff irrespective of the choices of the other players:*

$$\forall a_{-i} \in A_{-i},\ u_i(\alpha_i, a_{-i}) > u_i(\beta_i, a_{-i}).$$

The set of pure *actions of agent i that are* not dominated *by any mixed action is denoted by ND_i.*

The set of undominated actions can be formally written as follows:

$$\begin{aligned} ND_i &= \{a_i \in A_i : \forall \alpha_i \in \Delta(A_i), \exists a_{-i} \in A_{-i}, u_i(\alpha_i, a_{-i}) \leq u_i(a_i, a_{-i})\} \\ &= A_i \backslash \{a_i \in A_i : \exists \alpha_i \in \Delta(A_i), \forall a_{-i} \in A_{-i}, u_i(\alpha_i, a_{-i}) > u_i(a_i, a_{-i})\}. \end{aligned}$$

The proof of the following statement is left as an exercise for the reader.

Remark 1. *If a (pure) action a_i is dominated by a mixed action α_i then, for every conjecture $\mu^i \in \Delta(A_{-i})$, $u_i(a_i, \mu^i) < u_i(\alpha_i, \mu^i)$. Hence, a dominated action is not justifiable; equivalently, a justifiable action cannot be dominated.*

In the matrix of Figure 3.1, for instance, action f is dominated by c, which in turn is dominated by the mixed action $\alpha_i = \frac{1}{2}\delta_a + \frac{1}{2}\delta_b$ that assigns probability $\frac{1}{2}$ to a and b. Therefore, $ND_i \subseteq \{a, b, e\}$. Given that a, b, and e are best replies to at least one conjecture, we have that $\{a, b, e\} \subseteq ND_i$. Hence, $ND_i = \{a, b, e\}$.

The following lemma states that the converse of Remark 1 holds. Therefore, it provides a complete answer to our previous question about the characterization of the set of actions that a rational player would never choose.

Lemma 2. *(Wald [85], Pearce [66]) Fix arbitrarily a player $i \in I$ and an action $a_i^* \in A_i$ in a finite or compact-continuous game. There exists a conjecture $\mu^i \in \Delta(A_{-i})$ such that a_i^* is a best reply to μ^i* if and only if *a_i^* is not dominated by any (pure or) mixed action. In other words, the set of undominated (pure) actions and the set of justifiable actions coincide:*

$$ND_i = r_i\left(\Delta(A_{-i})\right).$$

Here we provide a proof for the *finite* case, which relies on an important result about linear algebra known as Farkas' lemma.[16] Intuitively, the payoff vector $\mathbf{u}^* = \left(u_i\left(a_i^*, a_{-i}\right)\right)_{a_{-i}\in A_{-i}}$ corresponding to an undominated action a_i^* must be on the "efficient frontier" of the convex set $\mathbf{U}$ of feasible payoff vectors (the North-East part of the boundary of the shaded area in Figure 3.3). Therefore we can find an hyperplane H (a line in Figure 3.3) going through $\mathbf{u}^*$ and separating space $\mathbb{R}^{A_{-i}}$ in two half-spaces (half-planes in Figure 3.3), so that one of them contains $\mathbf{U}$. Hyperplane H can be written as the set of vectors $\mathbf{y} \in \mathbb{R}^{A_{-i}}$ such that $\mathbf{y} \cdot \mathbf{x}^* = \mathbf{u}^* \cdot \mathbf{x}^*$, where $\mathbf{0} \neq \mathbf{x}^* \in \mathbb{R}^{A_{-i}}$. One can use Farkas' lemma to ensure that, since $\mathbf{y}^*$ is on the "efficient" part of the boundary, then $\mathbf{x}^*$ is a nonnegative vector, and therefore it can be normalized so that its elements sum to 1 and it becomes a conjecture $\mu^i \in \Delta\left(A_{-i}\right)$. By construction, this conjecture justifies action a_i^* as a best reply.

Lemma 3 (Farkas' Lemma)**.** *Let M be a $n \times m$ matrix and $\mathbf{c} \in \mathbb{R}^n$ be an n-dimensional vector. Then exactly one of the following statements is true:*
(1) there exists a vector $\mathbf{x} \in \mathbb{R}^m$ such that $M\mathbf{x} \geq \mathbf{c}$;
(2) there exists a vector $\mathbf{y} \in \mathbb{R}^n$ such that $\mathbf{y} \geq \mathbf{0}$, $\mathbf{y}^{\mathrm{T}}M = \mathbf{0}^{\mathrm{T}}$ and $\mathbf{y}^{\mathrm{T}}\mathbf{c} > 0$.

Proof of Lemma 2. We use Farkas' Lemma to prove the non-obvious part of Lemma 2: if an action is undominated then it is justifiable, or—by contraposition—if it is not justifiable then it is dominated. Fix i and a_i^* arbitrarily. Since the game is finite, let $k = |A_i|$ and $m = |A_{-i}|$, and label

[16] In the statement of Farkas' Lemma and in the related proofs, we use boldface symbols to denote vectors (i.e., functions from $\{1, ..., n\}$ to $\mathbb{R}$). When vectors are represented as matrices, the default intrepretation is that they are columns. Thus, the inner product $\mathbf{y} \cdot \mathbf{x} = \sum_s y_s x_s$ in matrix algebra notation is written as $\mathbf{y}^{\mathrm{T}}\mathbf{x}$, where $\mathbf{y}$ and $\mathbf{x}$ are column vectors and $\mathbf{y}^{\mathrm{T}}$ is the row vector obtained from $\mathbf{y}$ by transposition.

elements in A_i as $\{1, 2, ..., k\}$ and elements in A_{-i} as $\{1, 2, ..., m\}$.[17] Then, construct a $k \times m$ matrix U in which the (w, z)-th coordinate is given by

$$U_{w,z} = u_i(a_i^*, z) - u_i(w, z).$$

With this, $\Delta(A_{-i})$ is a subset of $\mathbb{R}^m$ and a_i^* is justifiable if and only if we can find $\mu^i \in \Delta(A_{-i})$ such that $U\mu^i \geq \mathbf{0}$. This condition can be rewritten as follows: a_i^* is justifiable if and only if we can find $\mathbf{x} \in \mathbb{R}^m$ satisfying inequality:

$$M\mathbf{x} \geq \mathbf{c}, \tag{3.3.3}$$

where M is a $(k + m + 1) \times m$ matrix and $\mathbf{c}$ is a $(k + m + 1)$-dimensional vector defined as follows:

$$M = \begin{bmatrix} U \\ I_m \\ \mathbf{1}_m^{\mathrm{T}} \end{bmatrix}, \ \mathbf{c} = \begin{bmatrix} \mathbf{0}_k \\ \mathbf{0}_m \\ 1 \end{bmatrix}$$

(I_m denotes the m-dimensional identity matrix, that is an $m \times m$ matrix having 1s along the main diagonal and 0s everywhere else; $\mathbf{1}_\ell$ and $\mathbf{0}_\ell$ ($\ell = k$, or $\ell = m$) denote the ℓ-dimensional vectors having respectively 1s and 0s everywhere; from now on, the dimensionality indexes will be omitted). To see this, note that matrix inequality (3.3.3) can be written as the system of inequalities

$$\begin{cases} U\mathbf{x} \geq \mathbf{0} \\ x_1 \geq 0, ..., x_m \geq 0 \\ \sum_{z=1}^m x_z \geq 1 \end{cases};$$

if $\sum_{z=1}^m x_z > 1$ instead of $\sum_{z=1}^m x_z = 1$, then $\mathbf{x} \notin \Delta(A_{-i})$; but this is immaterial because we can normalize by substituting $\mathbf{x} = (x_1, ..., x_m)$ with $\mathbf{x}' = \left(\frac{x_1}{\sum_{z=1}^m x_z}, ..., \frac{x_m}{\sum_{z=1}^m x_z}\right) \in \Delta(A_{-i})$ and satisfy inequality (3.3.3), since $U\mathbf{x} \geq 0$ if and only if $U\mathbf{x}' \geq 0$.
Thus, if a_i^* is *not* justifiable, inequality (3.3.3) does *not* hold and Farkas' lemma implies that we can find $\mathbf{y} \in \mathbb{R}^{k+m+1}$ such that

$$\begin{cases} \mathbf{y} \geq \mathbf{0} \\ \mathbf{y}^{\mathrm{T}}\mathbf{c} > 0 \\ \mathbf{y}^{\mathrm{T}}M = \mathbf{0}^{\mathrm{T}} \end{cases}.$$

[17] Recall that $|X|$ denotes the cardinality of set X.

By construction, $\mathbf{y}^{\mathrm{T}}\mathbf{c} = y_{k+m+1}$; thus the second condition can be rewritten as $y_{k+m+1} > 0$. Furthermore, $\mathbf{y}^{\mathrm{T}} M = \mathbf{0}^{\mathrm{T}}$ implies that

$$\sum_{w=1}^{k} y_w \left[u_i\left(a_i^*, z\right) - u_i\left(w, z\right) \right] = -\left(y_{k+z} + y_{k+m+1}\right) < 0$$

for every $z \in A_{-i} = \{1, 2, ..., m\}$, where the inequality holds because $y_{k+m+1} > 0$ and $y_{k+z} \geq 0$. Since the left hand side is non-zero and the vector $(y_1, ..., y_k)$ is nonnegative, we must have $y_w > 0$ for some $w \in A_i$. Then, we can construct a probability vector

$$\bar{\mathbf{y}} = \left(\frac{y_1}{\sum_{w=1}^{k} y_w}, ..., \frac{y_k}{\sum_{w=1}^{k} y_w} \right) \in \Delta(A_i) \subseteq \mathbb{R}^k$$

such that $u_i\left(a_i^*, z\right) < \sum_{w=1}^{k} \bar{y}_w u_i\left(w, z\right)$ for every $z \in A_{-i}$. We conclude that a_i^* is dominated by mixed action $\bar{\mathbf{y}}$. ■

We can thus conclude that *a rational player always chooses undominated actions* and that *each undominated action is justifiable as a best reply to some conjecture.*

In some interesting situations of strategic interaction an action dominates all others. In those cases, a rational player should choose that action. This consideration motivates the following:

Definition 7. *An action a_i^* is **dominant** if it dominates every other action, i.e., if*

$$\forall a_i \in A_i \backslash \left\{a_i^*\right\}, \forall a_{-i} \in A_{-i},\ u_i(a_i^*, a_{-i}) > u_i(a_i, a_{-i}).$$

You should try to prove the following statement as an exercise:

Remark 2. *Fix an action a_i^* arbitrarily. The following conditions are equivalent:*
(0) action a_i^* is dominant;
(1) action a_i^* dominates every mixed *action $\alpha_i \neq a_i^*$;*
(2) action a_i^* is the unique best reply to every *conjecture.*

The following example illustrates the notion of dominant action. The example shows that, assuming that players are motivated by their material self-interest, individual rationality may lead to socially undesirable outcomes.

Example 3. (Linear Public Good Game). In a community composed of n individuals it is possible to produce a quantity g of a public good using an input x according to the production function $y = kx$. Both y and x are measured in monetary units (say in Euros). To make the example interesting, assume that the productivity parameter k satisfies $1 > k > \frac{1}{n}$. A generic individual i can freely choose how many Euros to contribute to the community for the production of the public good. The community members cannot sign a binding agreement on such contributions because no authority with coercive power can enforce the agreement. Let W_i be player i's wealth. The game can be represented as follows: $A_i = [0, W_i]$, $a_i \in A_i$ is the contribution of i; consequences are allocations of the public and private good $(y, W_1 - a_1, \ldots, W_n - a_n)$; agents are selfish,[18] hence their utility function can be written as

$$v_i(y, W_1 - a_1, \ldots, W_n - a_n) = y - a_i,$$

which yields the payoff function

$$u_i(a_1, \ldots, a_n) = k \sum_{j=1}^{n} a_j - a_i.$$

It can be easily checked that $a_i^* = 0$ is dominant for any player i: just rewrite u_i as

$$u_i(a_i, a_{-i}) = k \sum_{j \neq i} a_j - (1 - k) a_i$$

and recall that $k < 1$. The profile of dominant actions $(0, \ldots, 0)$ is Pareto-dominated by any symmetric profile of positive contributions $(\varepsilon, \ldots, \varepsilon)$ (with $\varepsilon > 0$). Indeed, $u_i(0, \ldots, 0) = 0 < (nk - 1)\varepsilon = u_i(\varepsilon, \ldots, \varepsilon)$, where the inequality holds because $k > \frac{1}{n}$. Let $S(a) = \sum_{i=1}^{n} u_i(a)$ be the social surplus; the surplus maximizing profile is $\hat{a}_i = W_i$ for each i. ▲

An action could be a best reply only to conjectures that assign zero probability to some action profiles of the other players. For instance, action e in Figure 3.1 is justifiable as a best reply only if i is certain that $-i$ does not choose r. Let us say that a player i is **cautious** if his conjecture does not rule out any $a_{-i} \in A_{-i}$. Formally, let

$$\Delta^{\circ}(A_{-i}) = \left\{ \mu^i \in \Delta(A_{-i}) : \text{supp}\mu^i = A_{-i} \right\}$$

[18] And risk-neutral.

be the set of such *full-support* conjectures. In the *finite* case, we can write $\Delta^\circ(A_{-i})$ as follows:[19]

$$\Delta^\circ(A_{-i}) = \left\{\mu^i \in \Delta(A_{-i}) : \forall a_{-i} \in A_{-i}, \mu^i(a_{-i}) > 0\right\}.$$

A *rational and cautious* player i chooses actions in $r_i\left(\Delta^\circ(A_{-i})\right)$. These considerations motivate the following definition and results.

Definition 8. *A mixed action α_i* ***weakly dominates*** *another (pure or) mixed action β_i if it yields at least the same expected payoff for every action profile a_{-i} of the other players and strictly more for at least one $\bar{a}_{-i}$, that is,*

$$\begin{aligned} \forall a_{-i} &\in A_{-i},\ u_i(\alpha_i, a_{-i}) \geq u_i(\beta_i, a_{-i}), \\ \exists \bar{a}_{-i} &\in A_{-i},\ u_i(\alpha_i, \bar{a}_{-i}) > u_i(\beta_i, \bar{a}_{-i}). \end{aligned}$$

The set of pure *actions that are* not weakly dominated *by any mixed action is denoted by NWD_i. Such actions are also called* ***admissible.***

The set NWD_i of admissible actions can be formally written as follows:[20]

$$\begin{aligned} & NWD_i \\ = & \left\{ a_i \in A_i : \forall \alpha_i \in \Delta(A_i), \begin{array}{l} (\exists \bar{a}_{-i} \in A_{-i}, u_i(\alpha_i, \bar{a}_{-i}) < u_i(a_i, \bar{a}_{-i})) \\ \vee \\ (\forall a_{-i} \in A_{-i}, u_i(\alpha_i, a_{-i}) \leq u_i(a_i, a_{-i})) \end{array} \right\} \\ = & \left\{ a_i \in A_i : \forall \alpha_i \in \Delta(A_i), \begin{array}{l} (\exists \bar{a}_{-i} \in A_{-i}, u_i(\alpha_i, \bar{a}_{-i}) < u_i(a_i, \bar{a}_{-i})) \\ \vee \\ (\forall a_{-i} \in A_{-i}, u_i(\alpha_i, a_{-i}) = u_i(a_i, a_{-i})) \end{array} \right\}, \end{aligned}$$

which means that a_i is not weakly dominated if, for every α_i, either (1) a_i yields a strictly higher payoff than α_i for at least one a_{-i}, or (2) a_i yields a weakly higher expected payoff than α_i for every a_{-i}; if case (1) does not hold, then $u_i(\alpha_i, a_{-i}) \geq u_i(a_i, a_{-i})$ for every a_{-i}, and then a_i is not weakly dominated by α_i if and only if a_i and α_i are payoff-equivalent.

The reader should try to prove the following remark as an exercise:

---

[19] In the finite case, we can interpret $\Delta^\circ(A_{-i})$ as the relative interior of $\Delta(A_{-i})$, that is, the intersection of $\Delta(A_{-i})$ with the open set $\mathbb{R}^{A_{-i}}_{++}$.

[20] Symbol $\vee$ represents the non-exclusive "or." For example, if c is the proposition "the weather is cold" and r is the proposition "the weather is rainy," $c \vee r$ means that the weather is either cold, or rainy, or both; thus, $c \vee r$ is false if the weather is both warm (not cold) and dry (not rainy).

Remark 3. *If an action a_i^* in a finite or compact-continuous game is the* unique *best reply to some conjecture μ^i, then it is admissible.*

Next we analyze the relationship between admissible actions and the actions that might be chosen by a rational and cautious player. The following lemma (proved in the appendix) says that a rational and cautious player would never choose a weakly dominated action.

Lemma 4. *Fix arbitrarily a player $i \in I$ and an action $a_i^* \in A_i$ in a finite or compact-continuous game. If there exists a full-support conjecture μ^i that justifies a_i^*, then a_i^* is admissible:*

$$r_i \left(\Delta^\circ(A_{-i})\right) \subseteq NWD_i.$$

The following example shows that, if A_i is infinite, the inclusion can be strict, that is, there can be admissible actions that are not best replies to any full-support conjecture.

Example 4. Let $A_1 = [0,1]$, $A_2 = \{0,1\}$,

$$\begin{array}{rccl} u_1 : & A_1 \times A_2 & \to & \mathbb{R}, \\ & (a_1, a_2) & \mapsto & a_2 a_1 + (1-a_2)\sqrt{1-(a_1)^2}, \end{array}$$

and u_2 is any continuous function in $\mathbb{R}^{A_1 \times A_2}$. Clearly, this is a compact-continuous game. Let $\mu \in [0,1]$ denote the probability of $a_2 = 1$. Then

$$u_1 \left(a_1, \mu\right) = \mu a_1 + (1-\mu)\sqrt{1-(a_1)^2}.$$

This expected payoff function is strictly concave in a_1; hence, there is a unique best reply for each $\mu \in [0,1]$. The first-order condition yields the best reply function

$$r_1 \left(\mu\right) = \frac{\mu}{\sqrt{1-2\mu+2\mu^2}} \in [0,1].$$

Thus, $r_1(0) = 0$, $r_1(1) = 1$, and $0 < r_1(\mu) < 1$ if $0 < \mu < 1$, which implies that $0 \notin r_1\left(\Delta^\circ\left(A_2\right)\right)$ and $1 \notin r_1\left(\Delta^\circ\left(A_2\right)\right)$. Yet, neither 0 nor 1 is weakly dominated (see Remark 3). ▲

This example shows that the exact converse of Lemma 4 does not hold, because there are infinite compact-continuous games with admissible actions that cannot be justified as best replies to full-support conjectures. Yet, the following lemma (proved in the appendix) states that in all *finite* games admissibility is equivalent to justifiability by full-support conjectures.

Lemma 5. *(Cf. Wald [85], Pearce [66]) Let G be a finite game. Then, for every $i \in I$,*

$$NWD_i = r_i\left(\Delta^{\circ}(A_{-i})\right).$$

Observation 3. *The equivalence stated in Lemma 5 also holds for mixed actions: a mixed action in a finite game is not weakly dominated if and only if it is a best reply to a full-support conjecture.*

Next we consider an interesting special case:

Definition 9. *An action a_i^* is* ***weakly dominant*** *if it weakly dominates every other (pure) action, i.e., if for every other action $\widehat{a}_i \in A_i \backslash \{a_i^*\}$ the following conditions hold:*

$$\begin{aligned} \forall a_{-i} &\in A_{-i}, u_i(a_i^*, a_{-i}) \geq u_i(\widehat{a}_i, a_{-i}), \\ \exists \widehat{a}_{-i} &\in A_{-i}, u_i(a_i^*, \widehat{a}_{-i}) > u_i(\widehat{a}_i, \widehat{a}_{-i}). \end{aligned}$$

You should prove the following statement as an exercise:

Remark 4. *Fix an action a_i^* arbitrarily. The following conditions are equivalent:*
(0) a_i^* is weakly dominant;
(1) a_i^* weakly dominates every mixed *action $\alpha_i \neq a_i^*$;*
(2) a_i^* is the unique *best reply to every strictly positive conjecture $\mu^i \in \Delta^{\circ}(A_{-i})$;*
(3) a_i^* is the only *action with the property of being a best reply to* every *conjecture μ^i (note that if μ^i is not strictly positive, $\mu^i \in \Delta(A_{-i}) \backslash \Delta^{\circ}(A_{-i})$, then there may be other best replies to μ^i).*

If a rational and cautious player has a weakly dominant action, then he will choose such an action. There are interesting economic examples where individuals have weakly dominant actions.

Example 5. (Second Price Auction). An art merchant has to auction a work of art (e.g., a painting) at the highest possible price. However, he does not know how much such work of art is worth to the potential buyers. The buyers are collectors who buy the artwork with the only objective to keep it, i.e., they are not interested in what its future price might be. The authenticity of the work is not an issue. The potential buyer i is willing to spend at most $v_i > 0$ to buy it. Such valuation is completely subjective, meaning that if i were to know the other players' valuations, he would not change his own.[21] Following the advice of W. Vickrey, winner the of Nobel Prize for Economics,[22] the merchant decides to adopt the following auction rule: the artwork will go to the player who submits the highest offer, but the *price paid* will be equal to the *second highest offer* (in case of a tie between the maximum offers, the work will be assigned by a random draw). This auction rule induces a game among the buyers $i = 1, \dots, n$ where $A_i = [0, +\infty)$, and

$$u_i(a_i, a_{-i}) = \begin{cases} v_i - \max a_{-i}, & \text{if } a_i > \max a_{-i}, \\ 0, & \text{if } a_i < \max a_{-i}, \\ \frac{1}{1+|\arg\max a_{-i}|}(v_i - \max a_{-i}), & \text{if } a_i = \max a_{-i}, \end{cases}$$

where $\max a_{-i} = \max_{j \neq i}(a_1, ..., a_{i-1}, a_{i+1}, \dots, a_n)$. It turns out that offering $a_i^* = v_i$ is the weakly dominant action (can you prove it?). Hence, if the potential buyer i is rational and cautious, he will offer exactly v_i. Doing so, he will expect to make some profits. In fact, being cautious, he will assign a positive probability to event $[\max a_{-i} < v_i]$. Since by offering v_i he will obtain the object only in the event that the price paid is lower than his valuation v_i, the expected payoff from this offer is strictly positive. ▲

3.3.1 Comparative Risk Aversion and Justifiability

In this section we study the impact of risk aversion on the set of justifiable actions. In general, attitudes toward risk are captured by players' von

[21] If a buyer were to take into account a potential future resale, things might be rather different. In fact, other potential buyers could hold some relevant information that affects the estimate of how much the artwork could be worth in the future. Similarly, if there were doubts regarding the authenticity of the work, it would be relevant to know the other buyers' valuations.

[22] See Vickrey [82].

Neumann-Morgenstern utility functions $\upsilon_i : Y \to \mathbb{R}$ $(i \in I)$, where Y is the set of outcomes, or consequences. Assuming monetary consequences and selfish preferences, $Y = \mathbb{R}^I$ and $v_i\left((y_j)_{j\in I}\right) = \upsilon_i(y_i)$.[23] In this case, risk aversion is captured by the concavity of υ_i. In comparative terms, the preferences over monetary lotteries represented by $\hat{\upsilon}_i$ exhibit more risk aversion than the preferences represented by υ_i if $\hat{\upsilon}_i$ is a concave and strictly increasing transformation of υ_i, that is, there is a concave and strictly increasing function φ such that $\hat{\upsilon}_i = \varphi \circ \upsilon_i$.

The same comparative criterion can be directly applied to the payoff functions u_i. Consider again the case of monetary outcomes and selfish preferences and let $g = (g_i)_{i\in I} : A \to \mathbb{R}^I$, where $g_i : A \to \mathbb{R}$ is the monetary outcome for player $i \in I$. Then the payoff function of i is

$$\begin{array}{rccl} u_i = \upsilon_i \circ g_i : & A & \to & \mathbb{R}, \\ & a & \mapsto & \upsilon_i\left(g_i\left(a\right)\right). \end{array}$$

Thus, $\hat{u}_i = \hat{\upsilon}_i \circ g_i$ exhibits more risk aversion than $u_i = \upsilon_i \circ g_i$ if there is a concave and strictly increasing function φ such that

$$\hat{u}_i = \hat{\upsilon}_i \circ g_i = (\varphi \circ \upsilon_i) \circ g_i = \varphi \circ (\upsilon_i \circ g_i) = \varphi \circ u_i.$$

This motivates the following:

Definition 10. *Fix two games in reduced form $G = \langle I, (A_i, u_i)_{i\in I}\rangle$ and $\hat{G} = \langle I, (A_i, \hat{u}_i)_{i\in I}\rangle$ and a player $i \in I$, we say that i is **more risk averse in $\hat{G}$ than in G** if there exists a concave and strictly increasing function $\varphi : \mathbb{R} \to \mathbb{R}$ such that $\hat{u}_i = \varphi \circ u_i$.*

We show that, for players who rank actions according to subjective expected utility theory, higher risk aversion weakly expands, in the sense of set inclusion, the set of justifiable actions.[24] At first, the result may look counterintuitive: indeed, higher risk aversion of the decision maker, let us say player i, would increase the attractiveness of the "safer" actions, that is, actions $a_i \in A_i$ such that $u_i(a_i, a_{-i})$ is somewhat low for each profile $a_{-i} \in A_{-i}$, but does not change much with the choice of $a_{-i} \in A_{-i}$. On the other hand, "unsafe" actions a_i that are best replies to some

[23] We use the slightly different symbols v_i and υ_i because v_i is defined on $\mathbb{R}^I$ while υ_i is defined on $\mathbb{R}$.

[24] See Weinstein [89] and Battigalli et al. [21, Proposition 1].

deterministic conjectures may become instead *less* attractive. How can the set of justifiable actions change monotonically?

The key to understand the result intuitively is to note that the set of justifiable actions expands with risk aversion if and only if, for every action a_i which is justifiable with low risk aversion, the nonempty set $\left\{\mu^i \in \Delta\left(A_{-i}\right) : a_i \in r_i\left(\mu_i\right)\right\}$ of conjectures justifying a_i does *not* become *empty* when risk aversion is higher. In particular, if an "unsafe" action is justified by a deterministic conjecture, an increase in risk aversion cannot make such actions unjustifiable, because the set of best replies to deterministic conjectures is invariant to strictly increasing transformations of the payoff function u_i. The following example illustrates: as risk aversion increases, safe actions become easier to justify, in the sense that the set of justifying conjectures becomes larger; in particular, the set can be empty with low risk aversion and nonempty with higher risk aversion; on the other hand, if an unsafe action is—despite being unsafe—justifiable, then the set of justifying conjectures shrinks as risk aversion increases, but it cannot become empty.

Example 6. Consider the following game form with monetary outcomes, where $I = \{1, 2\}$, $A_1 = \{t, m, b\}$, $A_2 = \{\ell, r\}$, and outcome function $g : A \to \mathbb{R}^I$; player 1's component of the outcome function $g_1 : A \to \mathbb{R}$ is represented in Figure 3.4.

	ℓ	r
t	0	1
m	1/3	1/3
b	1	0

Figure 3.4: Function g_1 which represents monetary gains of player 1.

If player 1 is risk neutral, i.e., $v_1 : \mathbb{R}^I \to \mathbb{R}$ is defined by $v_1(x, y) = x$ and $u_1 = v_1 \circ g = g_1$, then action m is not justifiable. Indeed, for every belief $\mu^1 \in \Delta(A_2)$, we have

$$\mathbb{E}_{\mu^1}\left(u_1(m, \cdot)\right) = \frac{1}{3} < \frac{1}{2} \leq \max\left\{\mathbb{E}_{\mu^1}\left(u_1(t, \cdot)\right), \mathbb{E}_{\mu^1}\left(u_1(b, \cdot)\right)\right\}.$$

Then, let us suppose that the utility function v_1 is replaced by $v_{1,\theta}$:

$\mathbb{R}^I \to \mathbb{R}$ defined by

$$v_{1,\theta}(x,y) = x^{1/\theta},$$

where $\theta \geq 1$ parametrizes risk aversion. Then, the payoff function is defined by $u_{1,\theta}(a) = (g_1(a))^{1/\theta}$ for all $a \in A$.

It is not difficult to check that the set of justifiable actions of player 1 is

$$r_{1,\theta}\left(\Delta(A_{-1})\right) = \begin{cases} \{t,b\}, & \text{if } \theta \in [1, \log_2 3), \\ \{t,m,b\}, & \text{if } \theta \in [\log_2 3, \infty). \end{cases}$$

Therefore, the collection of justifiable actions expands as θ increases. On the other hand, notice that the set of beliefs justifying actions b and t *shrinks* as θ increases above the threshold $\overline{\theta} = \log_2 3$, see Figure 3.5. ▲

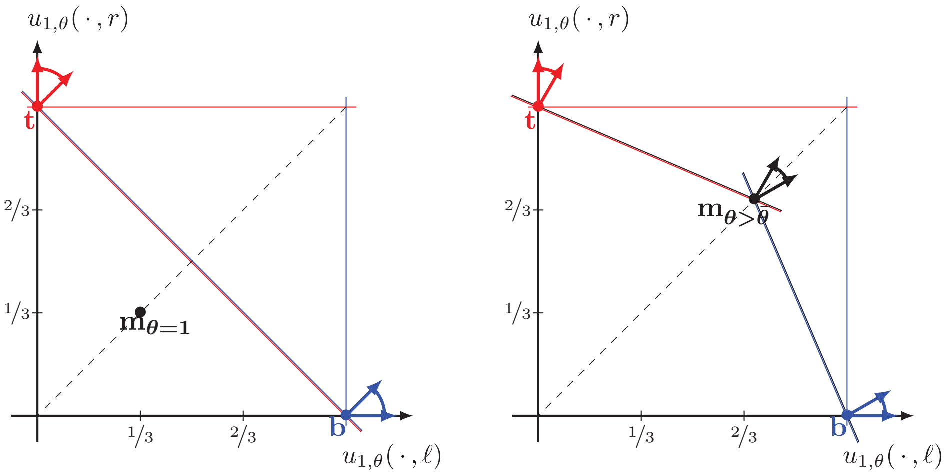

Figure 3.5: As θ increases, the set of beliefs justifying b and t shrinks.

Lemma 2 can be used to prove the following result:

Theorem 1. *Let $G = \langle I, (A_i, u_i)_{i\in I}\rangle$ and $\hat{G} = \langle I, (A_i, \hat{u}_i)_{i\in I}\rangle$ be two finite or compact-continuous games in reduced form, and suppose that player $i \in I$ is more risk averse in $\hat{G}$ than in G. Then*

$$r_i\left(\Delta(A_{-i})\right) \subseteq \hat{r}_i\left(\Delta(A_{-i})\right).$$

Proof. By assumption, there exists a concave and strictly increasing function $\varphi : \mathbb{R} \to \mathbb{R}$ such that $\hat{u}_i = \varphi \circ u_i$. Let us fix any justifiable action of i in G, that is, $a_i \in r_i\left(\Delta(A_{-i})\right)$. We must show that $a_i \in \hat{r}_i\left(\Delta(A_{-i})\right)$, where $\hat{r}_i$ is the best reply correspondence associated to payoff function $\hat{u}_i$. By Lemma 2, a_i is not dominated in G by any mixed action $\alpha_i \in \Delta(A_i)$:

$$\forall \alpha_i \in \Delta(A_i), \exists a_{-i} \in A_{-i}, u_i(a_i, a_{-i}) \geq u_i(\alpha_i, a_{-i}).$$

Therefore, for every mixed action $\alpha_i \in \Delta(A_i)$, there exists a profile $a_{-i} \in A_{-i}$ such that

$$\begin{aligned}\hat{u}_i(a_i, a_{-i}) &= \varphi\left(u_i(a_i, a_{-i})\right) \geq \varphi\left(u_i(\alpha_i, a_{-i})\right) = \varphi\left(\mathbb{E}_{\alpha_i}(u_i(\,\cdot\,, a_{-i}))\right) \\ &\geq \mathbb{E}_{\alpha_i}(\varphi(u_i(\,\cdot\,, a_{-i}))) = \mathbb{E}_{\alpha_i}(\hat{u}_i(\,\cdot\,, a_{-i})) = \hat{u}_i(\alpha_i, a_{-i}),\end{aligned}$$

that is, a_i is not dominated in $\hat{G}$. Here, the first inequality follows from the monotonicity of φ and the second follows from Jensen's inequality.[25]

Hence, again by Lemma 2, action a_i belongs to $\hat{r}_i\left(\Delta(A_{-i})\right)$. ■

3.3.2 Nice Games and Dominated Actions

Many models used in applications of game theory to economics and other disciplines have the following features: players choose actions in compact (closed and bounded) intervals of the real line, their payoff functions are continuous, and the payoff function of each player is strictly concave, or—at least—strictly quasi-concave,[26] in his own action. Such games are called "nice"; see Moulin [61].

Definition 11. *A static game $G = \langle I, (A_i, u_i)_{i\in I}\rangle$ is **nice** if, for every player $i \in I$, A_i is a compact interval $[\underline{a}_i, \bar{a}_i] \subset \mathbb{R}$, the payoff function $u_i : A \to \mathbb{R}$ is continuous (that is, jointly continuous in all its arguments), and the function $u_i(\cdot, a_{-i}) : [\underline{a}_i, \bar{a}_i] \to \mathbb{R}$ is strictly quasi-concave for each $a_{-i} \in A_{-i}$.*

[25] The concave transformation of the expectation of a random variable is larger than the expectation of the concave transformation of that random variable, see Aliprantis and Border [3, Theorem 11.24].

[26] A function $f : X \to \mathbb{R}$ with convex domain is **strictly quasi-concave** if for every pair of *distinct* points $x', x'' \in X$ and every $t \in (0, 1)$,

$$f(tx' + (1-t)x'') > \min\{f(x'), f(x'')\}.$$

Proposition 1. *Let* $A_i = [\underline{a}_i, \bar{a}_i] \subset \mathbb{R}$ *and fix* $a_{-i} \in A_{-i}$. *Function* $u_i(\cdot, a_{-i}) : [\underline{a}_i, \bar{a}_i] \to \mathbb{R}$ *is strictly quasi-concave if and only if the following two conditions hold: (1) there is a* unique *best reply* $r_i(a_{-i}) \in [\underline{a}_i, \bar{a}_i]$ *and (2)* $u_i(\cdot, a_{-i})$ *is strictly increasing on* $[\underline{a}_i, r_i(a_{-i})]$ *and strictly decreasing on* $[r_i(a_{-i}), \bar{a}_i]$.

Example 7. Consider Cournot's oligopoly game with capacity constraints: I is the set of firms, $a_i \in [0, \bar{a}_i]$ is the output of firm i and $\bar{a}_i$ is its capacity. The payoff function of firm i is

$$u_i(a) = a_i P \left(a_i + \sum_{j \neq i} a_j \right) - C_i(a_i),$$

where $P : [0, \sum_i \bar{a}_i] \to \mathbb{R}_+$ is the downward-sloping inverse demand function and $C_i : [0, \bar{a}_i] \to \mathbb{R}_+$ is i's cost function. The Cournot game is *nice* if $P(\cdot)$ and each cost function $C_i(\cdot)$ are continuous and if, for each i and a_{-i}, $P(a_i + \sum_{j \neq i} a_j)a_i - C_i(a_i)$ is strictly quasi-concave. The latter condition is easy to satisfy. For example, if each C_i is strictly increasing and convex, and $P(\sum_{j \in I} a_j) = \max\{0, \bar{p} - \beta \sum_{j \in I} a_j\}$, then strict quasi-concavity holds (prove this as an exercise).[27] ▲

Nice games allow an easy computation of the sets of justifiable actions and of dominated actions. We are going to compare best replies to deterministic conjectures, uncorrelated conjectures, correlated conjectures, pure actions not dominated by mixed actions, and pure actions not dominated by other pure actions. First note that, by definition, for every player i the following inclusions hold:[28]

$$\begin{aligned} r_i(A_{-i}) &\subseteq r_i(\mathrm{I}\Delta(A_{-i})) \subseteq r_i(\Delta(A_{-i})) \subseteq ND_i \\ &\subseteq \{a_i \in A_i : \forall b_i \in A_i, \exists a_{-i} \in A_{-i}, u_i(a_i, a_{-i}) \geq u_i(b_i, a_{-i})\}, \end{aligned}$$

[27] But strict concavity does *not* hold. Can you see why? Consider the neighborhood of a point where $P(\sum_{j \in I} a_j)$ becomes zero.

[28] If you do not know how to deal with probability measures on infinite sets, just consider the simple ones, i.e., those that assign positive probability to a finite set of points. The set of simple probability measures on an infinite domain X is a subset of $\Delta(X)$, but under the assumptions considered here it is possible to restrict one's attention to this smaller set.

where

$$\mathrm{I}\Delta(A_{-i}) = \left\{ \mu^i \in \Delta(A_{-i}) : \exists (\mu^i_j)_{j\in I\setminus\{i\}} \in \bigtimes_{j\in I\setminus\{i\}} \Delta(A_j), \mu^i = \prod_{j\neq i} \mu^i_j \right\}$$

is the set of *product* probability measures on A_{-i}, that is, the set of beliefs that satisfy independence across opponents. In words, the set of best replies to deterministic conjectures in A_{-i} is contained in the set of best replies to probabilistic independent conjecture, which is contained in the set of best replies to probabilistic (possibly correlated) conjectures, which is contained in the set of actions not dominated by mixed actions, which is contained in the set of (pure) actions not dominated by other pure actions. In each case the inclusion may hold as an equality. The following technical results imply that—in nice games—all these (weak) inclusions indeed hold as equalities. We start by providing conditions under which the best reply correspondence is actually a continuous *function*.

Lemma 6. *Consider a compact-continuous game such that, for each $i \in I$, each A_i is convex and $u_i(\cdot, a_{-i})$ is strictly quasi-concave for each $a_{-i} \in A_{-i}$. Then each player i has a well-defined and continuous best reply function on the domain A_{-i}.*

Proof. By convexity of A_i and strict quasi-concavity of u_i, the set $\arg\max_{a_i\in A_i} u_i(a_i, a_{-i})$ is a singleton for each a_{-i}. To prove this, it is sufficient to show that, if $a'_i \neq a''_i$ and $u_i(a'_i, a_{-i}) = u_i(a''_i, a_{-i})$, then $u_i(a'_i, a_{-i}) < \max_{a_i\in A_i} u_i(a_i, a_{-i})$. To see this, let a'_i and a''_i be as above. By convexity of A_i, for each $t \in (0, 1)$, $ta'_i + (1-t)a''_i \in A_i$; by strict quasi-concavity of u_i, $u_i(ta'_i+(1-t)a''_i, a_{-i}) > \min\{u_i(a'_i, a_{-i}), u_i(a''_i, a_{-i})\}$. Since $u_i(a'_i, a_{-i}) = u_i(a''_i, a_{-i})$ the result follows.

Next we show that the function $r_i(a_{-i}) = \arg\max_{a_i\in A_i} u_i(a_i, a_{-i})$ is continuous, that is, for any sequence $\left(a^k_{-i}\right)_{k\in\mathbb{N}}$ such that $\lim_{k\to\infty} a^k_{-i} = \bar{a}_{-i}$ the corresponding sequence of best replies $\left(r_i(a^k_{-i})\right)_{k\in\mathbb{N}}$ converges to $r_i(\bar{a}_{-i})$ $[\lim_{k\to\infty} r_i(a^k_{-i}) = r_i(\bar{a}_{-i})]$. Note first that, since the sequence $\left(r_i(a^k_{-i})\right)_{k\in\mathbb{N}}$ is contained in the compact set A_i, it must have at least one accumulation point $\bar{a}_i \in A_i$.[29] Therefore, there exists a subsequence $\left(a^{k_\ell}_{-i}\right)_{\ell\in\mathbb{N}}$ such that

[29]By the Bolzano–Weierstrass' theorem, every bounded sequence of real numbers has a convergent subsequence; see, e.g., Ok [64, Chapter A, p.52].

$\lim_{\ell\to\infty} r_i(a_{-i}^{k_\ell}) = \bar{a}_i$. The definition of best reply implies

$$\forall a_i \in A_i,\ \forall \ell \geq 1,\ u_i(r_i(a_{-i}^{k_\ell}), a_{-i}^{k_\ell}) \geq u_i(a_i, a_{-i}^{k_\ell}).$$

Taking the limit for $\ell \to \infty$ on both sides for any given a_i, the continuity of u_i yields

$$\forall a_i \in A_i,\ u_i(\bar{a}_i, \bar{a}_{-i}) \geq u_i(a_i, \bar{a}_{-i}).$$

Hence, $\bar{a}_i$ is a best reply to $\bar{a}_{-i}$. Since the best reply is unique, $\bar{a}_i = r_i(\bar{a}_{-i})$. This is true for every accumulation point. Therefore, the sequence $\left(r_i(a_{-i}^k)\right)_{k\in\mathbb{N}}$ has only one accumulation point, $r_i(\bar{a}_{-i})$, which is equivalent to say that $\lim_{k\to\infty} r_i(a_{-i}^k) = r_i(\bar{a}_{-i})$, as desired. ■

Lemma 6 implies that the set of best replies to deterministic conjectures is connected, which in turn yields the following convenient result:

Corollary 2. *In every* nice game, *the set of best replies of each player i to deterministic conjectures is a compact interval:*

$$r_i(A_{-i}) = \left[\min r_i(A_{-i}), \max r_i(A_{-i})\right].$$

Proof. Since $r_i(\cdot)$ is a continuous function and A_{-i} is convex (hence, connected) and compact, $r_i(A_{-i}) \subseteq \mathbb{R}$ is convex and compact,[30] hence it is a compact interval. ■

We can now state and prove the main result of this section.

Lemma 7. *In every* nice game, *the set of best replies of each player i to deterministic conjectures coincides with the set of actions not dominated by other* pure *actions, and therefore it also coincides with the set of best replies to independent or correlated probabilistic conjectures, and with the set of actions not dominated by mixed actions:*

$$\begin{aligned} r_i(A_{-i}) &= r_i(\mathrm{I}\Delta(A_{-i})) = r_i(\Delta(A_{-i})) = ND_i \\ &= \{a_i \in A_i : \forall b_i \in A_i, \exists a_{-i} \in A_{-i}, u_i(a_i, a_{-i}) \geq u_i(b_i, a_{-i})\}. \end{aligned}$$

Proof. Let $ND_{i,p} \subseteq A_i$ denote the set of player i's pure actions not dominated by other pure actions. We prove that $ND_{i,p} \subseteq r_i(A_{-i})$, that is, $A_i \backslash r_i(A_{-i}) \subseteq A_i \backslash ND_{i,p}$. Since we already noticed that

$$r_i(A_{-i}) \subseteq r_i(\mathrm{I}\Delta(A_{-i})) \subseteq r_i(\Delta(A_{-i})) \subseteq ND_{i,p},$$

[30] See Proposition 2 in Ok [64, D.2].

this implies the thesis. By Corollary 2,

$$r_i(A_{-i}) = [\min r_i(A_{-i}), \max r_i(A_{-i})].$$

Therefore, it is enough to show that all the actions below $\min r_i(A_{-i})$ or above $\max r_i(A_{-i})$ are dominated. Fix any $a_i < \min r_i(A_{-i})$. Thus, for each $a_{-i} \in A_{-i}$, $a_i < \min r_i(A_{-i}) \leq r_i(a_{-i})$. By definition, $u_i(\cdot, a_{-i})$ attains its maximum at $r_i(a_{-i})$; thus, by strict quasi-concavity, $u_i(\cdot, a_{-i})$ is strictly increasing on $[a_i, r_i(a_{-i})]$. It follows that

$$\forall a_{-i} \in A_{-i},\ u_i(a_i, a_{-i}) < u_i(\min r_i(A_{-i}), a_{-i}) \leq u_i(r_i(a_{-i}), a_{-i}).$$

Therefore, every $a_i < \min r_i(A_{-i})$ is strictly dominated by $\min r_i(A_{-i})$. A similar argument shows that every $a_i > \max r_i(A_{-i})$ is strictly dominated by $\max r_i(A_{-i})$. ■

3.3.3 Supermodular Nice Games

Best reply correspondences in nice games are continuous real-valued functions, which in general may not be monotone. We now consider a class of nice games where best reply correspondences are increasing functions, the so-called supermodular nice games, in which players have "strategic complementarities."[31] A game exhibits strategic complementarities if each player's incentive to increase his action becomes stronger for higher conjectured actions of the co-players. For example, if firms i and j are competing on prices and firm i thinks that j is going to increase its price, then firm i has a stronger incentive to increase (or, equivalently, a weaker incentive to decrease) its own price as well; see also Example 8 below. The notion of strategic complementarities is captured by the "supermodularity" of the payoff functions, which we define below for the special case of nice games.

We first introduce some notation. Let J be a finite index set, e.g., $J = \{1, \ldots, n\}$, or $J = I$, the player set in a game. Since $\mathbb{R}^J$ (the set of all functions from J to $\mathbb{R}$) is isomorphic to the Euclidean space $\mathbb{R}^{|J|}$, we call "vectors" the elements of $\mathbb{R}^J$. For each $x, y \in \mathbb{R}^J$, we say that $x \leq y$ if and only if $x_i \leq y_i$ for each $i \in J$; this is the usual coordinatewise incomplete

[31] Here we restrict our attention to nice games for the sake of simplicity, but the theory of supermodular games is much more general. See Topkis [81].

order on Euclidean spaces. Accordingly, for each $x, y \in \mathbb{R}^J$, we define the **order box** $[x, y] \subseteq \mathbb{R}^J$ as

$$[x, y] = \{z \in \mathbb{R}^J : x \leq z \leq y\}.$$

Note that $[x, y]$ is nonempty if and only if $x \leq y$, and $[x, y]$ is the usual compact interval if $x, y \in \mathbb{R}$.

Given a nonempty order box $[x, y] \subseteq \mathbb{R}^J$, a function $f : [x, y] \to \mathbb{R}$ is said to be **supermodular** if, for all $a, b \in [x, y]$,

$$f(a) + f(b) \leq f(\max(a, b)) + f(\min(a, b)), \tag{3.3.4}$$

where $\max(\cdot,\cdot)$ and $\min(\cdot,\cdot)$ denote the following binary operations on vectors:[32]

$$\begin{aligned}(a, b) &\mapsto (\max\{a_i, b_i\})_{i \in J},\\ (a, b) &\mapsto (\min\{a_i, b_i\})_{i \in J}.\end{aligned}$$

Definition 12. *A* nice *game $G = \langle I, (A_i, u_i)_{i \in I}\rangle$ is **supermodular** if, for each $i \in I$, the payoff function $u_i : A \to \mathbb{R}$ is supermodular.*

Note that the above definition is meaningful, because in nice games A is indeed an order box in $\mathbb{R}^I$. The idea behind the notion of supermodularity can be better understood by rewriting inequality (3.3.4) as

$$f(\max(a, b)) - f(a) \geq f(b) - f(\min(a, b)), \tag{3.3.5}$$

or also

$$f(\max(a, b)) - f(b) \geq f(a) - f(\min(a, b)), \tag{3.3.6}$$

which suggests a property of "increasing differences." To see this, consider the following example represented in Figure 3.6. Let $a = (2, 1)$ and $b = (1, 2)$ be points in the order box $X = [(1, 1), (2, 2)]$, and fix a supermodular function $f : X \to \mathbb{R}$. Then, horizontal arrows represent inequality (3.3.5), while vertical arrows represent inequality (3.3.6).

[32] When $a, b \in \mathbb{R}$ then $\max(a, b) = \max\{a, b\}$. But in the general case $a, b \in \mathbb{R}^J$, $\max(a, b)$ is the maximal element of the binary set $\{a, b\}$ if and only if, either $a_i \geq b_i$ for all $i \in J$, or $b_i \geq a_i$ for all $i \in J$; otherwise, $\max(a, b)$ does not even belong to $\{a, b\}$. A similar observation holds for $\min(a, b)$. In the more general framework of partially ordered sets, these binary operations are denoted by $a \vee b$ (maximum) and $a \wedge b$ (minimum).

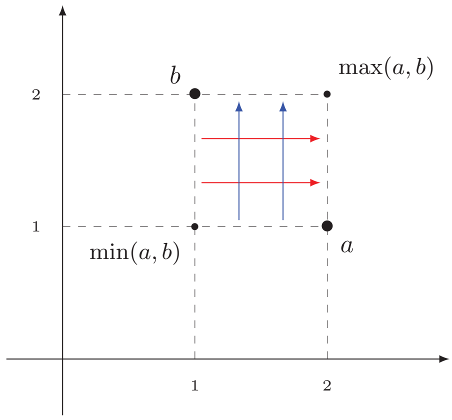

Figure 3.6: Increasing differences in a supermodular game.

Suppose that $X = A_1 \times A_2$ and let $f(a_1, a_2) = u_1(a_1, a_2)$ be the payoff function of player 1 in a two-person game. Then the "increasing differences" property says that player 1 gains more (or loses less) from increasing his action when the action of player 2 is higher.

We now make the connection with supermodularity more precise. Given nonempty order boxes $[x, y] \subseteq \mathbb{R}^{J_1}$ and $[a, b] \subseteq \mathbb{R}^{J_2}$, a function $f : [x, y] \times [a, b] \to \mathbb{R}$ is said to have **increasing differences in** $[a, b]$ if, for all $c \leq d$ in $[a, b]$, the function $\tilde{f} : [x, y] \to \mathbb{R}$ defined by

$$z \mapsto \tilde{f}(z) = f(z, d) - f(z, c)$$

is weakly increasing. Equivalently, $\tilde{f}(z) \leq \tilde{f}(w)$ for all $z, w \in [x, y]$ such that $z \leq w$.

The following result provides a useful criterion to check whether a function is supermodular:

Topkis' Characterization Theorem. (Topkis [81]) *Let* $[x, y] = \bigtimes_{i \in I}[x_i, y_i] \subseteq \mathbb{R}^I$ *be an order box. Then a function* $f : [x, y] \to \mathbb{R}$ *is supermodular if and only if* f *has increasing differences in* $\bigtimes_{j \in I \setminus \{i\}}[x_j, y_j]$ *for all* $i \in I$. *Additionally, if* f *is twice differentiable,*[33] *then* f *is*

[33] To avoid details involving boundary conditions, we stipulate that a function defined on an order box is differentiable if it has a differentiable extension defined on an open set containing the box.

supermodular if and only if the second-order partial derivatives $\partial^2 f/\partial x_i \partial x_j$ *are nonnegative for all* $1 \leq i < j \leq n$.

The second part of the Topkis' Characterization Theorem has an intuitive explanation. Indeed, if f is differentiable, then f has increasing differences if and only if $\partial f/\partial x_i$ is increasing in x_j for all $i \neq j$. It follows that, if f is also twice differentiable, then f has increasing differences if and only if $\partial^2 f/\partial x_i \partial x_j$ is a nonnegative function for all $i \neq j$.

Example 8. Consider a price-setting duopoly with differentiated products, so that $I = \{1, 2\}$ and $A_1 = A_2 = [0, \bar{p}]$, for some sufficiently large $\bar{p} > 0$. The joint demand function $D = (D_1, D_2) : A \to \mathbb{R}^2$ is defined by

$$(p_1, p_2) \mapsto (-d_1 p_1 + e_{12} p_2 + f_1, -d_2 p_2 + e_{21} p_1 + f_2)$$

for some collection of parameters $d_i, e_{ij}, f_i > 0$ $(i, j \in \{1, 2\},\ i \neq j)$. Such system of demand functions can be derived from quite standard preferences of consumers (see, for example, Motta [60]).

Each firm i has constant marginal cost $c_i < \bar{p}$. Since firms compete in prices, each firm i chooses p_i so as to maximize its payoff (profit)

$$u_i(p_i, p_j) = (p_i - c_i)\, D_i(p_i, p_j) = (p_i - c_i)\,(-d_i p_i + e_{ij} p_j + f_i)\,.$$

It follows from Topkis' Characterization Theorem that this is a supermodular nice game:

$$\frac{\partial u_i}{\partial p_i}(p_i, p_j) = -2d_i p_i + d_i c_i + e_{ij} p_j + f_i,$$

thus

$$\frac{\partial^2 u_i}{\partial p_i \partial p_j}(p_1, p_2) = e_{ij} > 0.$$

This means that each firm has the incentive to increase the price of its own product if the other firm is doing so. In particular, we verify that the best reply functions are weakly increasing, and indeed strictly increasing when they attain an interior solution. From the first-order condition $\frac{\partial u_i}{\partial p_i} = 0$ for interior solutions, and taking into account that $\frac{\partial^2 u_i}{\partial p_i^2} = -2d_i < 0$, we obtain the best reply function

$$r_i\,(p_j) = \min\left\{\frac{1}{2}\left(c_i + \frac{e_{ij} p_j + f_i}{d_i}\right), \bar{p}\right\},$$

that is, the best reply function is strictly increasing up to the point where it attains the upper bound $\bar{p}$. ▲

The following noteworthy lemma clarifies that the monotonicity result of Example 8 is a general property of supermodular nice games.

Lemma 8. *Let $G = \langle I, (A_i, u_i)_{i \in I} \rangle$ be a supermodular nice game. Then, for each $i \in I$, the best reply correspondences r_i are weakly increasing continuous functions. Moreover, if the payoff functions u_i are twice differentiable and $\frac{\partial^2 u_i}{\partial a_i^2} < 0$, $\frac{\partial^2 u_i}{\partial a_i \partial a_j} > 0$ for all $i, j \in I$, $j \neq i$, then the restriction of each r_i to the subset*

$$\{a_{-i} \in A_{-i} : \min A_i < r_i(a_{-i}) < \max A_i\}$$

where it attains interior values is a strictly increasing function.

Proof. We prove here the second part of the result, while the first part is proved in Appendix 3.4.2. Fix any $i \in I$ and note that, by Lemma 6, r_i is a continuous function. Pick any $a_{-i} \in A_{-i}$ such that $\min A_i < r_i(a_{-i}) < \max A_i$ (recall that $A_i = [\min A_i, \max A_i]$ because G is nice). Since $r_i(a_{-i})$ is an interior solution and u_i is twice differentiable, $r_i(a_{-i})$ satisfies the first and second-order conditions:

$$\begin{aligned} \frac{\partial u_i}{\partial a_i}(r_i(a_{-i}), a_{-i}) &= 0, \\ \frac{\partial^2 u_i}{\partial a_i^2}(r_i(a_{-i}), a_{-i}) &< 0, \end{aligned}$$

where the inequality is strict by assumption.[34] Hence, by the Implicit Function theorem, we obtain

$$\frac{\partial r_i}{\partial a_j}(a_{-i}) = -\frac{\frac{\partial}{\partial a_j}\left(\frac{\partial u_i}{\partial a_i}(r_i(a_{-i}), a_{-i})\right)}{\frac{\partial}{\partial a_i}\left(\frac{\partial u_i}{\partial a_i}(r_i(a_{-i}), a_{-i})\right)} = -\frac{\frac{\partial^2 u_i}{\partial a_i \partial a_j}(r_i(a_{-i}), a_{-i})}{\frac{\partial^2 u_i}{\partial a_i^2}(r_i(a_{-i}), a_{-i})} > 0$$

for each $j \neq i$. The last inequality follows from the hypothesis that $\frac{\partial^2 u_i}{\partial a_i \partial a_j} > 0$ for all $i \neq j$ and from the second-order condition. ■

[34] Strict quasi-concavity of u_i in a_i implies that $\frac{\partial^2 u_i}{\partial a_i^2} \leq 0$ at an interior maximum, but it does not imply that $\frac{\partial^2 u_i}{\partial a_i^2} < 0$. Consider, for example, a two-person game with $u_i(a_i, a_{-i}) = -(a_i - a_{-i})^4$. This payoff function is strictly concave (hence, strictly quasi-concave) in a_i, but $\frac{\partial^2 u_i}{\partial a_i^2} = 0$ at $a_i = a_{-i}$.

3.4 Appendix: Compact-Continuous Games

In this appendix, we add technical details about the analysis of compact-continuous games and we prove the general version of Lemma 2 (on best replies and dominance) and related results. The proof of Lemma 4 is similar. We first need some preliminaries about probability measures. We provide the necessary concepts to understand the analysis of mixed actions and probabilistic conjectures in the context of compact-continuous games. These preliminaries should be enough to allow the reader to understand the aforementioned proofs, but to fully master the mathematical concepts involved, the reader should consult appropriate primary sources, such as Aliprantis and Border [3]. In particular, the remarks stated in this Appendix are not proved; the reader should first try to prove them by herself or himself and then check on primary sources.

3.4.1 Probability Measures on Compact Sets

Sigma-Algebras and Probability Measures We expand on what we already explained in Section 3.1. For the reader's convenience, we repeat some concepts already introduced there. The following introductory discussion is heuristic.

To model uncertainty with an infinite state space X, we do not necessarily consider all the subsets of X as "events." We think of **events** as subsets of elements of X that verify some *statements*, which in turn can be obtained as negations, conjunctions, and disjunctions of simpler statements. For example, we may be uncertain about which point x is going to be selected, by an agent, or by a random device, in the compact interval $[0, 1]$. Then, the open sub-interval (a, b) of $[0, 1]$ corresponds to the statement "x is larger than a and smaller than b," which in turn is the conjunction of the two statements "x is larger than a" and "x is smaller than b." It is natural to assume that we can assign a probability to such elementary statements; with this, it is also natural to assume that we are able to assign a probability to their conjunction, i.e., to the open interval (a, b).

The negation of the conjunction "x is larger than a and smaller than b" is logically equivalent to the disjunction "either x is at most a, or x is at least b," which corresponds to the set $[0, a] \cup [b, 1] = [0, 1]\backslash(a, b)$. Of course, if we assign a probability p to a statement s, then we also assign probability $1-p$ to the negation of s. If statement s corresponds to set E, then the negation of s corresponds to the complement of E. Hence, if we can assign probability p to E, we also assign probability $1-p$ to E^c, the complement of E. In our example, we can assign probabilities to $(a,1]$, $[0,b)$, and (a,b), therefore we also assign probabilities to $[0,a]$, $[b,1]$, and $[0,a]\cup[b,1]$ respectively.

Finally, suppose we can assign a probability to each statement s_n, corresponding to event E_n. It is then quite natural to assume that we can also assign a probability to the statement "s_n holds for at least one $n\in\mathbb{N}$." For example, since we can assign probability to each statement "x is at most $(n-1)/(2n)$," then we can also assign probability to the statement "for at least one $n\in\mathbb{N}$, x is at most $\frac{1}{2}\frac{n-1}{n}$." Each elementary statement corresponds to the compact interval $\left[0,\frac{1}{2}\frac{n-1}{n}\right]$ and "for at least one $n\in\mathbb{N}$, x is at most $\frac{1}{2}\frac{n-1}{n}$" corresponds to

$$\bigcup_{n=1}^{\infty}\left[0,\frac{1}{2}\frac{n-1}{n}\right]=\left[0,\frac{1}{2}\right).$$

With this in mind, we abstract from the particular meaning of the statements and their representation as subsets of X and we identify the essential properties of the collection of subsets of X to which we are willing to assign a probability. A **sigma-algebra** on a set X is a collection of subsets $\mathcal{X}\subseteq 2^X$ called **events** such that

1. $X\in\mathcal{X}$,
2. for each $E\in\mathcal{X}$, also the complement of E belongs to $\mathcal{X}$, i.e., $E^c=X\backslash E\in\mathcal{X}$,
3. for each sequence of subsets E_n in $\mathcal{X}$, i.e., $(E_n)_{n=1}^{\infty}\in\mathcal{X}^{\mathbb{N}}$, also their union belongs to $\mathcal{X}$, that is, $\bigcup\limits_{n=1}^{\infty}E_n\in\mathcal{X}$.

Remark 5. *Let $\mathcal{X}$ be a sigma-algebra on X. Properties 1 and 2 imply that $\emptyset\in\mathcal{X}$. For every sequence of subsets $(E_n)_{n=1}^{\infty}$,*

$$\left(\bigcup_{n=1}^{\infty}E_n\right)^c=\bigcap_{n=1}^{\infty}(E_n)^c;$$

if follows that, by properties 2 and 3, if $(E_n)_{n=1}^{\infty}\in\mathcal{X}^{\mathbb{N}}$ then $\bigcap\limits_{n=1}^{\infty}E_n\in\mathcal{X}$.

Example 9. The collections $\{X, \emptyset\}$ and 2^X are sigma-algebras on X respectively called the **trivial** and the **discrete** sigma-algebra. ▲

Remark 6. *Let J be an arbitrary index set and consider the indexed family $\{\mathcal{X}_j\}_{j\in J}$ of sigma-algebras on X. Then also the intersection $\bigcap_{j\in J} \mathcal{X}_j$ is a sigma-algebra on X.*

A pair $(X, \mathcal{X})$ where $\mathcal{X}$ is a sigma-algebra on X is called a **measurable space**. A **probability measure** on a measurable space $(X, \mathcal{X})$ is a function $\mu : \mathcal{X} \to [0,1]$ such that

- $\mu(X) = 1$,
- for every sequence $(E_n)_{n=1}^{\infty} \in \mathcal{X}^{\mathbb{N}}$ of pairwise disjoint events

$$\mu\left(\bigcup_{n=1}^{\infty} E_n\right) = \sum_{n=1}^{\infty} \mu(E_n).$$

Remark 7. *For every probability measure μ on a measurable space $(X, \mathcal{X})$ the following holds:*
(1) $\mu(\emptyset) = 0$;
(2) if $(E_n)_{n=1}^{\infty} \in \mathcal{X}^{\mathbb{N}}$ is monotone increasing ($E_n \subseteq E_{n+1}$ for every $n \in \mathbb{N}$), then

$$\mu\left(\bigcup_{n=1}^{\infty} E_n\right) = \lim_{n\to\infty} \mu(E_n);$$

(3) if $(E_n)_{n=1}^{\infty} \in \mathcal{X}^{\mathbb{N}}$ is monotone decreasing ($E_{n+1} \subseteq E_n$ for every $n \in \mathbb{N}$), then

$$\mu\left(\bigcap_{n=1}^{\infty} E_n\right) = \lim_{n\to\infty} \mu(E_n).$$

Borel Sigma-Algebras and Probability Measures Let $X \subseteq \mathbb{R}^m$ be a nonempty compact subset.[35] In this case, the natural sigma-algebra of events to be considered is the so called **Borel sigma-algebra**, that is, the intersection of all the sigma-algebras containing all the closed subsets of

[35] All the definitions and results stated in this appendix for compact subsets of Euclidean spaces also apply for compact subsets of metric spaces.

X (or, equivalently, all the sets of the form $X \backslash C$ with C closed). Since 2^X is a sigma-algebra containing all the closed subsets, such intersection is well defined. We let $\mathcal{B}(X)$ denote the Borel sigma-algebra on any compact subset X of a Euclidean space. Whenever we consider a compact subset X of a Euclidean space, it is understood that the relevant measurable space is $(X, \mathcal{B}(X))$ and that the relevant probability measures are those with domain $\mathcal{B}(X)$, i.e., the Borel probability measures. The set of Borel probability measures on X is denoted $\Delta(X)$.

For every $\mu \in \Delta(X)$, with $X \subseteq \mathbb{R}^m$, we can define the corresponding cumulative distribution function (cdf) $F_\mu : \mathbb{R}^m \to [0,1]$ as follows:

$$\forall \hat{x} \in \mathbb{R}^m,\ F_\mu(\hat{x}) = \mu\left(\{x \in X : x_1 \leq \hat{x}_1, ..., x_m \leq \hat{x}_m\}\right).$$

Given F_μ, we can recover μ starting with the probability of events of the form $E = X \cap \times_{j=1}^m [\underline{x}_j, \overline{x}_j]$, that is

$$\mu(E) = \int_{\underline{x}_1}^{\overline{x}_1} \cdots \int_{\underline{x}_m}^{\overline{x}_m} \mathrm{d}F_\mu(x_1, \ldots, x_m);$$

then, intuitively, one can compute the probabilities of the unions, intersections, and complements of events whose probabilities have already been obtained.

Example 10. Let $X = [0,1]$. When the probability assigned to every interval is its length, so that—for example—$\mu((a,b]) = b - a$, then we obtain the uniform probability measure, or Lebesgue measure. If $\mu((a,b]) = F(b) - F(a)$ for some right-continuous non-decreasing function $F : \mathbb{R} \to [0,1]$, then $\mu = \mu_F$ is the probability measure generated by cdf F, and $F = F_\mu$ is the cdf generated by μ. If F is differentiable with integrable derivative f, then

$$\mu((a,b]) = \int_a^b f(x)\mathrm{d}x.$$

▲

Expected Values and Convergence of Sequences of Probability Measures For every continuous function $\varphi : X \to \mathbb{R}$ defined on a compact subset X of a Euclidean space and every measure $\mu \in \Delta(X)$, the **expected value** of φ given μ

$$\mathbb{E}_\mu(\varphi) = \int_X \varphi(x)\, \mu(\mathrm{d}x)$$

is the Riemann-Stieltjes integral

$$\mathbb{E}_{\mu}(\varphi) = \int \varphi(x)\,\mathrm{d}F_{\mu}.$$

Fix a sequence of measures $(\mu_n)_{n=1}^{\infty} \in [\Delta(X)]^{\mathbb{N}}$ and a measure $\bar{\mu}$; we say that μ_n **converges (weakly)** to $\bar{\mu}$, written $\mu_n \to \bar{\mu}$, if

$$\lim_{n\to\infty} \mathbb{E}_{\mu_n}(\varphi) = \mathbb{E}_{\bar{\mu}}(\varphi)$$

for every continuous (hence bounded) function $\varphi : X \to \mathbb{R}$.

Example 11. Let $X = [0,1]$. Then $\mu_n \to \bar{\mu}$ if and only if $\lim_{n\to\infty} F_{\mu_n}(x) = F_{\bar{\mu}}(x)$ for every x at which $F_{\bar{\mu}}$ is continuous. ▲

With this, we can define closed sets in $\Delta(X)$ and a notion of continuity of real-valued functions defined on $\Delta(X)$. First, we stipulate that a subset $C \subseteq \Delta(X)$ is **closed** if it contains the limits of all converging sequences of measures in C; that is, for all $(\mu_n)_{n\in\mathbb{N}} \in [\Delta(X)]^{\mathbb{N}}$ and $\bar{\mu} \in \Delta(X)$ such that $\mu_n \to \bar{\mu}$ and $\{\mu_n\}_{n\in\mathbb{N}} \subseteq C$, $\bar{\mu} \in C$. We say that $\psi : \Delta(X) \to \mathbb{R}$ is **continuous** if $\psi^{-1}(C)$ is closed for every closed subset $C \subseteq \mathbb{R}$. This is equivalent to requiring that $\lim_{n\to\infty} \psi(\mu_n) = \psi(\bar{\mu})$ for every converging sequence $\mu_n \to \bar{\mu}$. (See Aliprantis and Border [3, Theorem 15.11] and related material.) Note, when we consider the subset of all the Dirac measures $\{\mu \in \Delta(X) : \exists x \in X, \mu = \delta_x\}$, which can be identified with X, we obtain the usual notions of closed sets in X and continuity of real-valued functions defined on X.

A set $D \subseteq \Delta(X)$ is **open** if $X \backslash D$ is closed. We say that a collection of open sets $\mathcal{O}$ **covers** $\Delta(X)$ if the union of all such subsets contains $\Delta(X)$:

$$\Delta(X) \subseteq \bigcup \mathcal{O}.$$

$\Delta(X)$ endowed with its collection of open subsets is **compact** in the following sense: for every collection of open subsets $\mathcal{O}$ that covers $\Delta(X)$ there is a *finite* sub-collection $\mathcal{F} \subseteq \mathcal{O}$ that also covers $\Delta(X)$. Furthermore, $\Delta(X)$ is a **metrizable space**: this means that there exists a function $d : \Delta(X) \times \Delta(X) \to \mathbb{R}_+$, called **compatible metric** (or distance), such that

- for all $\mu', \mu'' \in \Delta(X)$, $d(\mu', \mu'') = 0$ if and only if $\mu' = \mu''$;

- for all $\mu', \mu'', \mu''' \in \Delta(X)$, $d(\mu', \mu''') \leq d(\mu', \mu'') + d(\mu'', \mu''')$;
- for all $(\mu_n)_{n=1}^{\infty} \in [\Delta(X)]^{\mathbb{N}}$ and $\bar{\mu} \in \Delta(X)$, $\mu_n \to \bar{\mu}$ if and only if for every $\varepsilon > 0$ there is some $n_\varepsilon \in \mathbb{N}$ such that $d(\mu_n, \bar{\mu}) < \varepsilon$ for all $n \geq n_\varepsilon$.[36]

Given nonempty compact subsets of Euclidean spaces X and Y, we can now define in a natural way notions of openness of subsets of $\Delta(X) \times \Delta(Y)$ and $X \times \Delta(Y)$, and of (joint) continuity for real-valued functions defined on $\Delta(X) \times \Delta(Y)$ and $X \times \Delta(Y)$. We omit the details.

3.4.2 Best Replies in Compact-Continuous Games

As explained in Section 3.4.1, for every compact subset X of a Euclidean space (more generally, for every compact subset of a metric space), the set of Borel probability measures $\Delta(X)$ is metrizable and compact; see [3, Theorem 15.11]. Every sequence in a compact metrizable space has a convergent subsequence. Also recall that, by definition of convergence in $\Delta(X)$, for every continuous function $f : X \to \mathbb{R}$, if $\mu_n \to \mu$ then $\mathbb{E}_{\mu_n}(f) \to \mathbb{E}_{\mu}(f)$. Therefore:

Lemma 9. *Let $X \subseteq \mathbb{R}^m$ be compact and let $f : X \to \mathbb{R}$ be continuous. For every sequence $(\mu_n)_{n=1}^{\infty} \in [\Delta(X)]^{\mathbb{N}}$ there is a convergent subsequence $(\mu_{n_k})_{k=1}^{\infty}$ and $\lim_{k\to\infty} \mathbb{E}_{\mu_{n_k}}(f) = \mathbb{E}_{\bar{\mu}}(f)$, where $\bar{\mu} = \lim_{k\to\infty} \mu_{n_k}$.*

It follows from Section 3.4.1 that if X and Y are compact subsets of Euclidean spaces, then $X \times \Delta(Y)$, or—equivalently—$\Delta(X) \times Y$, is a compact metrizable space. With this, we prove an important property of the best reply correspondence.

Lemma 10. *Fix a compact-continuous game G. For every player $i \in I$, and every closed subset $C_{-i} \subseteq A_{-i}$, the restriction $r_i|_{\Delta(C_{-i})}$ of the best reply correspondence to $\Delta(C_{-i})$ is nonempty-valued, has a closed graph, and a closed image, that is,*

$$\forall \mu^i \in \Delta(C_{-i}),\ r_i\left(\mu^i\right) \neq \emptyset,$$

[36] Compactness of the metrizable space $\Delta(X)$ implies that it is also complete (every Cauchy sequence converges) and separable ($\Delta(X)$ contains a countable subset D such that, for every $\mu \in \Delta(X)$ there is a sequence $(\mu_n)_{n=1}^{\infty} \in D^{\mathbb{N}}$ such that $\mu_n \to \mu$); see, e.g., Aliprantis and Border [3, Theorem 3.28].

the set

$$\mathrm{Gr}\left(r_i|_{\Delta(C_{-i})}\right) = \left\{\left(\mu^i, a_i\right) \in \Delta\left(C_{-i}\right) \times A_i : a_i \in r_i\left(\mu^i\right)\right\}$$

is closed in the compact metrizable space $\Delta(A_{-i}) \times A_i$*, and the set* $r_i(\Delta(C_{-i})) \subseteq A_i$ *is closed.*

Proof. Fix $\mu^i \in \Delta(C_{-i})$ arbitrarily. Since u_i is continuous on A, $u_i\left(\cdot, \mu^i\right) : A_i \to \mathbb{R}$ is continuous. By compactness of A_i, $\arg\max_{a_i \in A_i} u_i\left(a_i, \mu^i\right) \neq \emptyset$. Since $r_i\left(\mu^i\right) = \arg\max_{a_i \in A_i} u_i\left(a_i, \mu^i\right)$ (Corollary 1), $r_i\left(\mu^i\right) \neq \emptyset$.[37] Hence r_i is nonempty-valued.

To prove that $\mathrm{Gr}\left(r_i|_{\Delta(C_{-i})}\right)$ is closed, we show that, for every convergent sequence contained in $\mathrm{Gr}\left(r_i|_{\Delta(C_{-i})}\right)$, the limit of the sequence belongs to $\mathrm{Gr}\left(r_i|_{\Delta(C_{-i})}\right)$: Suppose that $\left(\mu_n^i, a_{i,n}\right) \to \left(\bar{\mu}^i, \bar{a}_i\right)$, where $\mu_n^i \in \Delta(C_{-i})$ and $a_{i,n} \in r_i\left(\mu_n^i\right)$ for every $n \in \mathbb{N}$, so that the range of the sequence is included in $\mathrm{Gr}\left(r_i|_{\Delta(C_{-i})}\right)$. Then,

$$\forall a_i \in A_i, \forall n \in \mathbb{N},\ u_i\left(a_{i,n}, \mu_n^i\right) \geq u_i\left(a_i, \mu_n^i\right).$$

By continuity of u_i, taking the limit as $n \to \infty$ for each fixed a_i, we obtain

$$\forall a_i \in A_i,\ u_i\left(\bar{a}_i, \bar{\mu}^i\right) \geq u_i\left(a_i, \bar{\mu}^i\right).$$

Therefore $\bar{a}_i \in r_i\left(\bar{\mu}^i\right)$ and $\bar{\mu}^i \in \Delta(C_{-i})$ (since $\Delta(C_{-i})$ is a closed subset of $\Delta(A_{-i})$), that is, $\left(\bar{\mu}^i, \bar{a}_i\right) \in \mathrm{Gr}\left(r_i|_{\Delta(C_{-i})}\right)$. To see that $r_i(\Delta(C_{-i}))$ is closed, fix arbitrarily a convergent sequence $(a_{i,n})_{n=1}^{\infty}$ such that $a_{i,n} \to \bar{a}_i$ and $\{a_{i,n}\}_{n=1}^{\infty} \subseteq r_i(\Delta(C_{-i}))$. We must show that $\bar{a}_i \in r_i(\Delta(C_{-i}))$. For every $n \in \mathbb{N}$ there is some conjecture $\mu_n^i \in \Delta(C_{-i})$ that justifies $a_{i,n}$ as a best reply: $a_{i,n} \in r_i\left(\mu_n^i\right)$. Since C_{-i} is a closed subset of a compact set, C_{-i} is compact. Hence also $\Delta(C_{-i})$ is compact and the sequence $\left(\mu_n^i\right)_{n=1}^{\infty}$ has a convergent subsequence $\left(\mu_{n_k}^i\right)_{k=1}^{\infty}$, with $\lim_{k\to\infty} \mu_{n_k}^i = \bar{\mu}^i \in \Delta(C_{-i})$ (Lemma 9). Since $\mathrm{Gr}(r_i)$ is closed, $\lim_{k\to\infty}\left(\mu_{n_k}^i, a_{n_k}\right) \in \mathrm{Gr}(r_i)$, that is,

$$\lim_{k\to\infty} a_{n_k} = \lim_{n\to\infty} a_n = \bar{a}_i \in r_i\left(\bar{\mu}^i\right).$$

Since $\bar{\mu}^i \in \Delta(C_{-i})$, we conclude that $\bar{a}_i \in r_i(\Delta(C_{-i}))$ as desired. ■

[37] Recall that we defined $r_i\left(\mu^i\right)$ as the set of mixed actions α_i that maxmize $u_i\left(\cdot, \mu^i\right) : \Delta(A_i) \to \mathbb{R}$ and are also Dirac measures, hence belong to A_i as well.

3.4.3 Proof of the Duality Lemma of Wald and Pearce

The structure of the argument is as follows: We first observe that a mixed action $\bar{\alpha}_i$ is justifiable if and only if

$$0 \geq \min_{\mu^i \in \Delta(A_{-i})} \max_{\alpha_i \in \Delta(A_i)} \left[u_i \left(\alpha_i, \mu^i \right) - u_i \left(\bar{\alpha}_i, \mu^i \right) \right].$$

Then we invoke the maxmin theorem, which states that we can invert the min and max operations in the expression above without changing the final value. Next we note that in the *max-min* formula the second operation can be replaced with a minimization with respect to deterministic conjectures, that is, co-players' action profiles. This finally yields the equivalence with the fact that $\bar{\alpha}_i$ is not dominated by any mixed action.

Since finite games can be regarded as compact-continuous, the proof applies to all finite games. Considering the case of finite games, the reader can understand the core of the argument without the distraction of measure-theoretic considerations.

For any given conjecture μ^i, mixed action $\bar{\alpha}_i$, and alternative mixed action α_i, $\left[u_i \left(\alpha_i, \mu^i \right) - u_i \left(\bar{\alpha}_i, \mu^i \right) \right]$ represents the strength of the "incentive to deviate" from $\bar{\alpha}_i$ to α_i. Thus, $\bar{\alpha}_i$ is justifiable if and only if the maximal incentive to deviate is non-positive for at least one μ^i, that is, the minimum of the maximal incentive to deviate is not positive. The following result and proof make the previous argument more formal.

Lemma 11. *Let $G = \langle I, (A_i, u_i)_{i \in I} \rangle$ be a compact-continuous game. Then, for every $i \in I$ and $\bar{\alpha}_i \in \Delta(A_i)$, the following are equivalent:*

(J) $\exists \bar{\mu}^i \in \Delta \left(A_{-i} \right), \forall \alpha_i \in \Delta \left(A_i \right), u_i \left(\bar{\alpha}_i, \bar{\mu}^i \right) \geq u_i \left(\alpha_i, \bar{\mu}^i \right)$,

(ND) $0 \geq \min_{\mu^i \in \Delta(A_{-i})} \max_{\alpha_i \in \Delta(A_i)} \left[u_i \left(\alpha_i, \mu^i \right) - u_i \left(\bar{\alpha}_i, \mu^i \right) \right]$.

Proof. *Preliminary observations:*[38] Let $X \subseteq \mathbb{R}^m$ be compact, and let the function $f : X \to \mathbb{R}$ be continuous. Then $\Delta \left(X \right)$ is compact metrizable and the map $\mu \mapsto \mathbb{E}_\mu \left(f \right)$ is continuous. Therefore

$$\sup_{x \in X} f(x) = \max_{x \in X} f(x), \quad \sup_{\mu \in \Delta(X)} \mathbb{E}_\mu \left(f \right) = \max_{\mu \in \Delta(X)} \mathbb{E}_\mu \left(f \right).$$

[38]To be skipped when considering finite games.

If X and Y are compact subsets of Euclidean spaces and $\varphi : \Delta(X) \times \Delta(Y) \to \mathbb{R}$ is continuous on $\Delta(X) \times \Delta(Y)$ (a product of compact metrizable spaces), then the maps $\mu \mapsto \max_{\nu \in \Delta(X)} \varphi(\nu, \mu)$ and $\nu \mapsto \min_{\mu \in \Delta(Y)} \varphi(\nu, \mu)$ are continuous.[39] Finally, recall that the payoff function u_i is extended to $\Delta(A_i) \times \Delta(A_{-i})$ as follows:

$$u_i(\alpha_i, \mu^i) = \mathbb{E}_{\alpha_i \times \mu^i}(u_i(\cdot, \cdot)).$$

Now, a probability measure $\mu^i \in \Delta(A_{-i})$ is such that $u_i(\alpha_i, \mu^i) \leq u_i(\bar{\alpha}_i, \mu^i)$ for all $\alpha_i \in \Delta(A_i)$, if and only if

$$0 \geq \sup_{\alpha_i \in \Delta(A_i)} \left[u_i(\alpha_i, \mu^i) - u_i(\bar{\alpha}_i, \mu^i)\right] = \max_{\alpha_i \in \Delta(A_i)} \left[u_i(\alpha_i, \mu^i) - u_i(\bar{\alpha}_i, \mu^i)\right].$$

Since the right-hand side is continuous in μ^i on the compact set $\Delta(A_{-i})$, there is a conjecture μ^i that makes it non-positive if and only if its minimum with respect to μ^i is non-positive. Hence, (J) holds if and only if

$$0 \geq \min_{\mu^i \in \Delta(A_{-i})} \max_{\alpha_i \in \Delta(A_i)} \left[u_i(\alpha_i, \mu^i) - u_i(\bar{\alpha}_i, \mu^i)\right],$$

that is, if and only if (ND) holds. ■

We wish to relate justifiability (condition J) to lack of dominance. To make this connection, we are going to use an important result that can be proved independently, the **maxmin theorem**:

Maxmin Theorem. (Sion [77]) *Let X and Y be compact subsets of Euclidean spaces, and let $f : X \times Y \to \mathbb{R}$ be continuous on $X \times Y$; then*

$$\min_{\mu \in \Delta(X)} \max_{\nu \in \Delta(Y)} \mathbb{E}_{\mu \times \nu}(f) = \max_{\nu \in \Delta(Y)} \min_{\mu \in \Delta(X)} \mathbb{E}_{\mu \times \nu}(f).$$

We can now prove the duality lemma of Wald and Pearce stating that an action is justifiable if and only if it is not dominated:

Lemma 12. (Wald [85], Pearce [66]) *Let G be a compact-continuous game. Then, for every $i \in I$ and $\bar{\alpha}_i \in \Delta(A_i)$, the following are equivalent:*

(J) $\exists \bar{\mu}^i \in \Delta(A_{-i}), \forall \alpha_i \in \Delta(A_i), u_i(\bar{\alpha}_i, \bar{\mu}^i) \geq u_i(\alpha_i, \bar{\mu}^i)$,

[39] See Aliprantis and Border [3, Theorem 17.31].

$$\text{(MM)}\ \forall \alpha_i \in \Delta(A_i), \exists a_{-i} \in A_{-i}, 0 \geq u_i(\alpha_i, a_{-i}) - u_i(\bar{\alpha}_i, a_{-i}).$$

Proof. By Lemma 11, (J) is equivalent to

$$0 \geq \min_{\mu^i \in \Delta(A_{-i})} \max_{\alpha_i \in \Delta(A_i)} \left[u_i\left(\alpha_i, \mu^i\right) - u_i\left(\bar{\alpha}_i, \mu^i\right)\right]. \tag{3.4.1}$$

By the Maxmin Theorem,

$$\min_{\mu^i \in \Delta(A_{-i})} \max_{\alpha_i \in \Delta(A_i)} \left[u_i\left(\alpha_i, \mu^i\right) - u_i\left(\bar{\alpha}_i, \mu^i\right)\right] = \max_{\alpha_i \in \Delta(A_i)} \min_{\mu^i \in \Delta(A_{-i})} \left[u_i\left(\alpha_i, \mu^i\right) - u_i\left(\bar{\alpha}_i, \mu^i\right)\right].$$

Therefore, (3.4.1) is equivalent to

$$\forall \alpha_i \in \Delta(A_i),\ 0 \geq \min_{\mu^i \in \Delta(A_{-i})} \left[u_i\left(\alpha_i, \mu^i\right) - u_i\left(\bar{\alpha}_i, \mu^i\right)\right]. \tag{3.4.2}$$

Since $\left[u_i\left(\alpha_i, \mu^i\right) - u_i\left(\bar{\alpha}_i, \mu^i\right)\right]$ is affine in μ^i,

$$\min_{\mu^i \in \Delta(A_{-i})} \left[u_i\left(\alpha_i, \mu^i\right) - u_i\left(\bar{\alpha}_i, \mu^i\right)\right] = \min_{a_{-i} \in A_{-i}} \left[u_i(\alpha_i, a_{-i}) - u_i(\bar{\alpha}_i, a_{-i})\right]$$

for every $\alpha_i \in \Delta(A_i)$. Therefore, (3.4.2) is equivalent to

$$\forall \alpha_i \in \Delta(A_i),\ 0 \geq \min_{a_{-i} \in A_{-i}} \left[u_i(\alpha_i, a_{-i}) - u_i(\bar{\alpha}_i, a_{-i})\right],$$

which is equivalent to condition (MM). The previous chain of equivalences proves the claim. ■

3.4.4 Cautiously Justifiable and Admissible Actions

Proof of Lemma 4. We must prove that, in a compact-continuous game, $r_i(\Delta^\circ(A_{-i})) \subseteq NWD_i$. By way of contraposition, we show that if a_i^* is weakly dominated, then it is not a best reply to any full-support conjecture, that is, $a_i^* \notin r_i(\Delta^\circ(A_{-i}))$. Thus, suppose that a_i^* is weakly dominated by some $\alpha_i \in \Delta(A_i)$, that is,

$$\begin{aligned} \forall a_{-i} &\in A_{-i},\ u_i(\alpha_i, a_{-i}) \geq u_i(a_i^*, a_{-i}), \\ \exists \bar{a}_{-i} &\in A_{-i},\ u_i(\alpha_i, \bar{a}_{-i}) > u_i(a_i^*, \bar{a}_{-i}). \end{aligned} \tag{3.4.3}$$

Let $\epsilon = u_i(\alpha_i, \bar{a}_{-i}) - u_i(a_i^*, \bar{a}_{-i}) > 0$. Continuity of u_i implies that the function

$$a_{-i} \mapsto u_i(\alpha_i, a_{-i}) - u_i(a_i^*, a_{-i}),$$

is continuous as well. Therefore, there exists $\eta > 0$ such that

$$\forall a_{-i} \in A_{-i},\, d\,(a_{-i}, \bar{a}_{-i}) < \eta \Rightarrow u_i(\alpha_i, a_{-i}) - u_i(a_i^*, \bar{a}_{-i}) > \frac{\epsilon}{2}$$

(d denotes the standard Euclidean distance). Note that the (relatively) open set $O_{-i} = \{a_{-i} \in A_{-i} : d\,(a_{-i}, \bar{a}_{-i}) < \eta\}$ is assigned strictly positive probability by every full-support probability measure μ^i.

So, fix any $\mu^i \in \Delta^\circ(A_{-i})$. Then, the definition of O_{-i} and eq.s (3.4.3) yield

$$\begin{aligned}
u_i(\alpha_i, \mu^i) - u_i(a_i^*, \mu^i) &= \int_{A_{-i}} [u_i(\alpha_i, a_{-i}) - u_i(a_i^*, a_{-i})]\, \mu^i\,(\mathrm{d}a_{-i}) \\
&= \int_{O_{-i}} [u_i(\alpha_i, a_{-i}) - u_i(a_i^*, a_{-i})]\, \mu^i\,(\mathrm{d}a_{-i}) \\
&\quad + \int_{A_{-i}\setminus O_{-i}} [u_i(\alpha_i, a_{-i}) - u_i(a_i^*, a_{-i})]\, \mu^i\,(\mathrm{d}a_{-i}) \\
&\geq \int_{O_{-i}} [u_i(\alpha_i, a_{-i}) - u_i(a_i^*, a_{-i})]\, \mu^i\,(\mathrm{d}a_{-i}) \\
&> \frac{\epsilon}{2}\mu^i(O_{-i}) \\
&> 0,
\end{aligned}$$

that is, $u_i(\alpha_i, \mu^i) > u_i(a_i^*, \mu^i)$. Since μ^i is arbitrary, we conclude that $a_i^* \notin r_i\,(\Delta^\circ(A_{-i}))$. ■

Proof of Lemma 5. The inclusion $r_i\,(\Delta^\circ(A_{-i})) \subseteq NWD_i$ holds by Lemma 4. We prove that if an action a_i^* of player i is not weakly dominated by any mixed action, then it is **cautiously justifiable**, i.e., justifiable by a full-support conjecture. With an argument similar to the proof of Lemma 2, we use Farkas' lemma to prove this statement by contraposition, that is, if a_i^* is not cautiously justifiable, then it is weakly dominated.

Fix i and a_i^* arbitrarily. Since the game is finite, we let $k = |A_i|$ and $m = |A_{-i}|$, and we label elements in A_i as $\{1, 2, ..., k\}$ and elements in

A_{-i} as $\{1, 2, ..., m\}$. Then, we construct a $k \times m$ matrix U in which the (w, z)-th coordinate is given by

$$U_{w,z} = u_i\left(a_i^*, z\right) - u_i\left(w, z\right).$$

With this, $\Delta^\circ\left(A_{-i}\right)$ is a subset of $\mathbb{R}^m_{++}$, the strictly positive orthant of $\mathbb{R}^m$, and a_i^* is cautiously justifiable if and only if there exists $\mu^i \in \Delta^\circ\left(A_{-i}\right)$ such that $U\mu^i \geq \mathbf{0}$. This condition can be rewritten as follows: a_i^* is cautiously justifiable if and only if, for some $p > 0$, we can find $\mathbf{x} \in \mathbb{R}^m$ satisfying

$$M\mathbf{x} \geq \mathbf{c}, \tag{3.4.4}$$

where M is a $(k + m + 1) \times m$ matrix and $\mathbf{c}$ is a $(k + m + 1)$-dimensional vector[40] defined as follows:

$$M = \begin{bmatrix} U \\ I_m \\ \mathbf{1}_m^{\mathrm{T}} \end{bmatrix}, \mathbf{c} = \begin{bmatrix} \mathbf{0}_k \\ \mathbf{p}_m \\ 1 \end{bmatrix}.$$

Here, I_m denotes the m-dimensional identity matrix, that is an $m \times m$ matrix having 1s along the main diagonal and 0s everywhere else; $\mathbf{1}_\ell$ and $\mathbf{0}_\ell$ ($\ell = k$, or $\ell = m$) denote the ℓ-dimensional vectors having respectively 1s and 0s everywhere; $\mathbf{p}_m$ denotes the m-dimensional vector having $p > 0$ everywhere; from now on, the dimensionality indexes will be omitted.

To see why a_i^* is cautiously justifiable if and only there exists $\mathbf{x} \in \mathbb{R}^m$ such that (3.4.4) holds, note that matrix inequality (3.4.4) can be written as the system of inequalities

$$\begin{cases} U\mathbf{x} \geq \mathbf{0} \\ x_1 \geq p, ..., x_m \geq p \\ \sum_{z=1}^m x_z \geq 1 \end{cases}.$$

Note that, if $\sum_{z=1}^m x_z > 1$ instead of $\sum_{z=1}^m x_z = 1$, then $\mathbf{x} \notin \Delta^\circ\left(A_{-i}\right)$; but this is immaterial because we can normalize by substituting $\mathbf{x} = (x_1, ..., x_m)$ with $\mathbf{x}' = \left(\frac{x_1}{\sum_{z=1}^m x_z}, ..., \frac{x_m}{\sum_{z=1}^m x_z}\right) \in \Delta^\circ\left(A_{-i}\right)$ and satisfy inequality (3.4.4), since $M\mathbf{x} \geq 0$ if and only if $M\mathbf{x}' \geq 0$.

[40] Recall that, in matrix calculus, an n-dimensional vector is represented as an $n \times 1$ matrix (column vector), and by transposition we obtain a $1 \times n$ matrix (row vector).

Thus, if a_i^* is *not* cautiously justifiable, inequality (3.4.4) does *not* hold and Farkas' lemma implies the existence of some $\mathbf{y} \in \mathbb{R}^{k+m+1}$ such that

$$\begin{cases} \mathbf{y} \geq \mathbf{0} \\ \mathbf{y}^{\mathrm{T}}\mathbf{c} > 0 \\ \mathbf{y}^{\mathrm{T}}M = \mathbf{0}^{\mathrm{T}} \end{cases} .$$

By construction,

$$\mathbf{y}^{\mathrm{T}}\mathbf{c} = p\sum_{z=1}^{m} y_{k+z} + y_{k+m+1} > 0. \tag{3.4.5}$$

Furthermore, $\mathbf{y}^{\mathrm{T}}M = \mathbf{0}^{\mathrm{T}}$ implies that

$$\sum_{w=1}^{k} y_w \left[u_i\left(a_i^*, z\right) - u_i\left(w, z\right)\right] = -\left(y_{k+z} + y_{k+m+1}\right) \leq 0 \tag{3.4.6}$$

for every $z \in A_{-i} = \{1, 2, ..., m\}$, where the inequality holds because $y_{k+m+1} \geq 0$ and $y_{k+z} \geq 0$. Since $\mathbf{y} \geq \mathbf{0}$ and $p > 0$, (3.4.5) implies that there exists $z' \in \{1, 2, ..., m+1\}$ such that $y_{k+z'} > 0$. If $z' = m+1$, then the left-hand side of (3.4.6) is non-zero for every $z \in \{1, 2, ..., m\}$. If $z' \in \{1, 2, ..., m\}$, then the left-hand side of (3.4.6) is non-zero for $z = z'$. Since the vector $(y_1, ..., y_k)$ is nonnegative, we must have $y_w > 0$ for some $w \in A_i$. Then, we can construct a probability vector

$$\bar{\mathbf{y}} = \left(\frac{y_1}{\sum_{w=1}^{k} y_w}, ..., \frac{y_k}{\sum_{w=1}^{k} y_w}\right) \in \Delta(A_i) \subseteq \mathbb{R}^k$$

such that $u_i\left(a_i^*, z\right) < \sum_{w=1}^{k} \bar{y}_w u_i\left(w, z\right)$ for at least one $z \in A_{-i}$, and $u_i\left(a_i^*, z\right) \leq \sum_{w=1}^{k} \bar{y}_w u_i\left(w, z\right)$ for every $z \in A_{-i}$. We conclude that a_i^* is weakly dominated by mixed action $\bar{\mathbf{y}}$. ■

3.4.5 Increasing Best Reply Functions

We conclude this appendix by providing the proof of the first part of Lemma 8. The proof follows from an important result on supermodular functions:

Topkis' Maximum Theorem. (Topkis [81]) *Let* $X \subseteq \mathbb{R}$ *and* $Y \subseteq \mathbb{R}^n$ *be nonempty order boxes, and fix a* continuous supermodular *function* $f : X \times Y \to \mathbb{R}$.

Then, for each $y \in Y$*, the set* $\arg\max_{x \in X} f(x, y)$ *has a greatest and a least element, which we denote by* $M_f(y)$ *and* $m_f(y)$*, respectively. In addition, the functions* M_f *and* m_f *are weakly increasing.*

Proof. By assumption, X is an order box of $\mathbb{R}$, i.e., it is a (nonempty) compact interval. Let $\hat{f} : Y \rightrightarrows \mathbb{R}$ be the correspondence defined by

$$y \mapsto \hat{f}(y) = \arg\max_{x \in X} f(x, y).$$

Since f is continuous, then standard arguments (see the proof of Lemma 10) show that, for each $y \in Y$, the set $\hat{f}(y)$ is a closed (hence compact) subset of $X \subseteq \mathbb{R}$. Therefore $M_f(y) = \max \hat{f}(y)$ and $m_f(y) = \min \hat{f}(y)$ are well defined.

It remains to show that $M_f(a) \leq M_f(b)$ for all $a, b \in Y$ such that $a \leq b$. Define $\bar{a} = M_f(a)$ and $\bar{b} = M_f(b)$ so that, in particular, we have $f(\bar{a}, a) \geq f(\min(\bar{a}, \bar{b}), a)$. Since f is supermodular on $X \times Y$, we obtain

$$\begin{aligned} f(\bar{a}, a) + f(\bar{b}, b) &\leq f(\max(\big((\bar{a}, a), (\bar{b}, b)\big) + f(\min\big((\bar{a}, a), (\bar{b}, b)\big)) \\ &= f(\max(\bar{a}, \bar{b}), b) + f(\min(\bar{a}, \bar{b}), a). \end{aligned}$$

By the previous observation, it follows that $f(\bar{b}, b) \leq f(\max(\bar{a}, \bar{b}), b)$, so that $\max(\bar{a}, \bar{b}) \in \hat{f}(b)$. Therefore the conclusion that $\bar{a} \leq \bar{b}$ follows from the fact that $\bar{b}$ is the greatest element of $\hat{f}(b)$. The proof for m_f is similar.
■

Proof of the First Part of Lemma 8. Recall that, according to Proposition 1, the best-reply correspondence r_i is actually a continuous function for each $i \in I$. The claim follows from Topkis' Maximum Theorem: it is enough to set $f = u_i$, $X = A_i$, and $Y = A_{-i}$, so that $r_i = M_f$, which is weakly increasing. ■

4

Rationalizability and Iterated Dominance

The analysis, so far, has been based on a set of minimal epistemic assumptions:[1] every player knows the sets of possible actions and his own payoff function. From this analysis we derived a basic, decision-theoretic principle of rationality: A rational player should not choose those actions that are dominated by mixed actions. This principle of rationality can be interpreted in a descriptive way (assuming that a rational player will not choose dominated actions), or in a normative way (a player should not choose dominated actions).

The concept of dominance is sufficient to obtain interesting results in some interactive situations: those where it is not necessary to guess the other players' actions in order to make a correct decision. Simple social dilemmas, like the Prisoners' Dilemma or contributing to a public good, have this feature. But when we analyze strategic reasoning, the decision-theoretic rationality principle is just a starting point: **thinking strategically** *means trying to anticipate the co-players' moves and plan a best response, taking into account that also the co-players are intelligent individuals trying to do just the same.*

Strategic thinking is based on knowledge of the rules of the game and preferences, knowledge of the other players' knowledge of rules and

[1] We call "epistemic" the assumptions about knowledge, conjectures, and beliefs of players.

preferences, and so on and so forth. As a starting point, we assume that players have *complete information*, i.e., common knowledge of the rules of the game and everybody's preferences. In Example 1 complete information implies that the functional forms and the parameters K, θ_1, θ_2 are common knowledge: both players know them, both players know that both players know them, and so on and so forth.

The complete-information assumption is sufficiently realistic for some economic situations in which the consequences of players' actions are purely monetary outcomes and players are selfish and risk neutral.[2] The complete information-assumption is also a useful simplification in that it allows to focus on other aspects of strategic thinking than players' knowledge of the game and of the knowledge of other players. Later on, we will introduce strategic thinking in games with *in*complete information and we will show that the techniques developed for games with complete information can be extended to the more general case.

4.1 Assumptions about Players' Beliefs

The analysis will proceed in a series of steps. After introducing the concept of rationality and its behavioral consequences, we will characterize which actions are consistent not only with rationality, but also with the belief that everybody is rational. Then we will characterize which actions are consistent not only with the previous assumptions, but also with the further assumption that each player believes that each other player believes that everybody is rational; and so on. Thus, at each further step we characterize more restrictive assumptions about players' rationality and beliefs.

Such assumptions can be formally represented as **events**, i.e., as subsets of a space Ω of "states of the world" where every state $\omega \in \Omega$ is a conceivable configuration of actions, conjectures, and beliefs concerning other players' beliefs. This formal representation provides an expressive and precise language that allows to make the analysis more rigorous

[2]Recall that what matters are players' preferences over *lotteries* of consequences, hence also their risk attitudes. But in some cases (for instance, some oligopolistic models) risk attitudes are irrelevant for the strategic analysis and risk neutrality can therefore be assumed with no loss of generality.

and transparent.[3] However, it calls for the preliminary introduction of some rather complex material, which is not strictly needed to understand the basics. Therefore we opt for a compromise. On the one hand, in order to make the presentation more transparent and to avoid verbose sentences, we use symbols to denote the assumptions concerning players' behavior and beliefs; we refer to these assumptions as "events" and denote a conjunction of assumptions by the symbol of intersection between events. On the other hand, rather than providing an explicit definition for such events, we only define mathematically the sets of actions (and conjectures) that are consistent with them. Of course, we can afford to do this precisely because we know how to represent assumptions about rationality and beliefs as events and how to derive characterizations of behavior based on increasingly sophisticated strategic thinking by means of rigorous mathematical methods. We ask you, the reader, to trust us on this derivation. Since the results are quite intuitive, we hope you will not object. (If you do object, perhaps you are ready for the complete mathematical analysis of strategic thinking: good news for us!)

To see what we mean, consider the rationality assumption. We denote by R_i the event "player i is rational," whereas $R = \bigcap_{i \in I} R_i$ denotes the event "everybody (in the game) is rational." In the previous chapter we saw that $\times_{i \in I} ND_i$ (where ND_i is the set of undominated actions of i) is the set of action profiles consistent with R (rationality).

The rationality assumption simply establishes a relation between conjectures (beliefs about the behavior of other players) and actions. This relation is represented by the best reply correspondence. Now we proceed introducing some assumptions concerning players' beliefs about each other beliefs, also called "**interactive beliefs**."

Denote by E a generic event about players' actions and beliefs (for example, $E = R$). We represent with the symbol $\mathrm{B}(E)$ the event "everybody believes E," meaning "everybody assigns probability one to E." Further, we write $\mathrm{B}(\mathrm{B}(E)) = \mathrm{B}^2(E)$ ("everybody believes that everybody believes E"), and more generally $\mathrm{B}^k(E) = \mathrm{B}(\mathrm{B}^{k-1}(E))$ for any integer $k > 1$.

Consider the conjunction of the following assumptions: R, $\mathrm{B}(R)$, $\mathrm{B}^2(R), \ldots, \mathrm{B}^k(R)$, etc. Is it possible to provide a characterization in terms

[3] See, for example, Battigalli and Bonanno [9], and Dekel and Siniscalchi [38].

of actions? In other words, what sets of action profiles are consistent with assumptions $R \cap \mathrm{B}(R)$, $R \cap \mathrm{B}(R) \cap \mathrm{B}^2(R), \ldots, R \cap \left(\bigcap_{k=1}^{K} \mathrm{B}^k(R)\right)$, etc.? In order to answer this question, it is useful to introduce a mapping, which is derived from the best reply correspondence and assigns to any "cross-product" subset C of A a subset of action profiles that are "rationalized," or "justified" by C.

4.2 The Justification Operator

Fix some *finite* set X and a collection $\mathcal{C}$ of subsets of X, i.e., $\mathcal{C} \subseteq 2^X$ (we postpone the analysis of the case of an infinite X). In this textbook we call "operator" a self-map from $\mathcal{C}$ to itself, that is, a map that associates each subset of X in $\mathcal{C}$ with another subset of X in $\mathcal{C}$.[4]

As a matter of notation, for a fixed set X_{-i} (for instance, $X_{-i} = A_{-i}$) and a subset $Y_{-i} \subseteq X_{-i}$, we denote by $\Delta(Y_{-i})$ the subset of probability measures on X_{-i} that assign probability one to Y_{-i}:

$$\Delta(Y_{-i}) = \{\mu^i \in \Delta(X_{-i}) : \mu^i(Y_{-i}) = 1\}$$

(another slight abuse of notation).

In particular, we are going to consider the set $X = A = \times_{i \in I} A_i$ of action profiles, and the collection $\mathcal{C}$ of all "Cartesian" subsets of A, i.e., subsets with the cross-product form $C = \times_{i \in I} C_i$, where $C_i \subseteq A_i$ for every i. Then, if the game is finite,

$$\mathcal{C} = \{C \in 2^A : \exists\, (C_i)_{i \in I} \in \bigtimes_{i \in I} 2^{A_i}, C = \bigtimes_{i \in I} C_i\}.$$

Let $C_{-i} = \times_{j \in I \setminus \{i\}} C_j$. We define the following sets:

$$\rho_i(C_{-i}) = \left\{a_i \in A_i : \exists \mu^i \in \Delta(C_{-i}), a_i \in r_i(\mu^i)\right\} = r_i(\Delta(C_{-i})),$$

$$\rho(C) = \bigtimes_{i \in I} \rho_i(C_{-i}).$$

[4]In mathematics, the term "operator" is mostly used for maps from a space of functions to a space of functions (possibly the same). But sets can always be represented by functions (e.g., indicator functions). Therefore the present use of the term "operator" is consistent with standard mathematical language.

The interpretation is as follows: $\rho(C)$ *is the set of action profiles that could be chosen by rational players if every* i *is certain that the co-players choose in* C_{-i}. Therefore, we say that $\rho(C)$ is the set of action profiles "rationalized" (or "justified") by C and we call the map $\rho : \mathcal{C} \to \mathcal{C}$ "**justification operator**."[5] Note that $\rho(\emptyset) = \emptyset$.

	ℓ	c	r
T	$3, 2$	$0, 1$	$0, 0$
M	$0, 2$	$3, 1$	$0, 0$
B	$1, 1$	$1, 2$	$3, 0$

Figure 4.1: A 3×3 game.

Example 12. In a 3×3 game like the one in Figure 4.1, each player has $2^3 = 8$ subsets of actions (including the empty set). Hence, the collection of subsets $\mathcal{C}$ contains $8 \times 8 = 64$ elements, $7 \times 7 = 49$ of them are nonempty (since $\rho(\emptyset) = \emptyset$, we may ignore the empty subsets). It would be very tedious to determine $\rho(C)$ for each (nonempty) $C \in \mathcal{C}$. We give just a few examples:

$$\begin{aligned}
\rho(\{T, M, B\} \times \{\ell, c, r\}) &= \{T, M, B\} \times \{\ell, c\}, \\
\rho(\{T, M, B\} \times \{\ell, c\}) &= \{T, M\} \times \{\ell, c\}, \\
\rho(\{T, M\} \times \{\ell, c\}) &= \{T, M\} \times \{\ell\}, \\
\rho(\{M, B\} \times \{\ell, c\}) &= \{T, M\} \times \{\ell, c\}, \\
\rho(\{T, M\} \times \{\ell\}) &= \{T\} \times \{\ell\}, \\
\rho(\{T, M\} \times \{r\}) &= \{B\} \times \{\ell\}, \\
\rho(\{T\} \times \{\ell\}) &= \{T\} \times \{\ell\}, \\
\rho(\{B\} \times \{r\}) &= \{B\} \times \{c\}.
\end{aligned}$$

To see this, check the best replies to deterministic conjectures, note that r is strictly dominated by l and c, next note that B is strictly dominated by $\alpha_1 = \frac{1}{2}\delta_T + \frac{1}{2}\delta_M$ in the smaller game obtained by deleting the right column r. This shows that, in general, we may have $\rho(C) \subseteq C$, $C \subseteq \rho(C)$

[5] See Milgrom and Roberts [58].

(here this inclusion holds as an equality for $C = \{T\} \times \{l\}$), $\rho(C) \subsetneq C$ and $C \subsetneq \rho(C)$.[6] ▲

Remark 8. *Lemma 2 implies that $\rho_i(A_{-i})$ is the set of undominated actions for player i. Hence, $\rho(A) = \bigtimes_{i\in I} ND_i$.*

Note that, for every i, $E_{-i} \subseteq F_{-i}$ implies $\Delta(E_{-i}) \subseteq \Delta(F_{-i})$, which in turn implies $\rho_i(E_{-i}) = r_i(\Delta(E_{-i})) \subseteq r_i(\Delta(F_{-i})) = \rho_i(F_{-i})$. With this, we can conclude the following.

Remark 9. *The justification operator is* monotone*: for every pair of subsets $E, F \in \mathcal{C}$, if $E \subseteq F$ then $\rho(E) \subseteq \rho(F)$.*

Since the domain and codomain of ρ coincide, it makes sense to define the k-th iteration of ρ (the k-fold composition of ρ with itself) recursively as follows. For each $C \in \mathcal{C}$, define $\rho^0(C) = C$ for convenience; then for each $k \geq 1$,

$$\rho^k(C) = \rho(\rho^{k-1}(C)).$$

Note that we are using the standard definition of the iteration of a self-map $f : X \to X$, that is, $f^0 = \mathrm{Id}_X$ (the identity function on X) by convention, and $f^k = f \circ f^{k-1}$ for each $k \geq 1$. In this case $X = \mathcal{C}$ and $f = \rho$.

Infinite Games The analysis extends seamlessly from finite to compact-continuous games. To be completely general, we should define the justification operator for all the (Borel) measurable Cartesian subsets of A. But we can restrict our attention to the closed Cartesian subsets of A. In fact, compactness of A and continuity of the payoff functions imply that, for every $i \in I$ and every closed subset $C_{-i} \subseteq A_{-i}$, the set of conjectures $\Delta(C_{-i}) = \left\{\mu^i \in \Delta(A_{-i}) : \mu^i(C_{-i}) = 1\right\}$ is well defined[7] and that the set of best replies $r_i(\Delta(C_{-i}))$ is closed (see Lemma 15 in Appendix 4.6). Note that a Cartesian product $C = \bigtimes_{i\in I} C_i$ is closed if and only if each "factor" C_i is closed. Therefore, whenever $C = \bigtimes_{i\in I} C_i$ is closed, also $\rho(C) = \bigtimes_{i\in I} r_i(\Delta(C_{-i}))$ is closed, and the same holds for each $\rho^k(C)$ ($k \in \mathbb{N}$). For each $i \in I$, let $\mathcal{K}_{A_i} \subseteq 2^{A_i}$ denote the collection of *closed*

[6] Symbol "$\subsetneq$" means "not weakly included in."

[7] Because the closed C_{-i} is necessarily measurable, hence the probability $\mu^i(C_{-i})$ is well defined.

(hence compact) subsets of A_i, including A_i itself. With this, define the collection of closed Cartesian products

$$\mathcal{C} = \left\{ C \in 2^A : \exists \left(C_i\right)_{i\in I} \in \bigtimes_{i\in I} \mathcal{K}_{A_i}, C = \bigtimes_{i\in I} C_i \right\}.$$

Then, for each $C \in \mathcal{C}$, $\rho\left(C\right) = \bigtimes_{i\in I} r_i \left(\Delta\left(C_{-i}\right)\right)$ is well defined and $\rho\left(C\right) \in \mathcal{C}$. Therefore, $\rho : \mathcal{C} \to \mathcal{C}$ is a well-defined self-map.

4.3 Rationalizability: The Powers of ρ

The set of action profiles consistent with R (rationality of all the players) is $\rho(A)$. Therefore, if every player is rational and believes R, only those action profiles that are rationalized by $\rho(A)$ can be chosen. It follows that the set of action profiles consistent with $R \cap \mathrm{B}(R)$ is $\rho(\rho(A)) = \rho^2(A)$. Iterating the procedure one more step, it is relatively easy to see that the set of action profiles consistent with $R \cap \mathrm{B}(R) \cap \mathrm{B}^2(R)$ is $\rho(\rho^2(A)) = \rho^3(A)$. The general relationship between events about rationality and beliefs and corresponding sets of action profiles is given by the Table 4.1.

Table 4.1: Behavioral implications of assumptions on rationality and beliefs.

Assumptions about rationality and beliefs	Behavioral implications
R	$\rho(A)$
$R \cap \mathrm{B}(R)$	$\rho^2(A)$
$R \cap \mathrm{B}(R) \cap \mathrm{B}^2(R)$	$\rho^3(A)$
...	...
$R \cap \left(\bigcap_{k=1}^{K} \mathrm{B}^k(R)\right)$	$\rho^{K+1}(A)$
...	...

Note that, as one should expect, the sequence of subsets $(\rho^k(A))_{k=1}^{\infty}$ is weakly decreasing, i.e., $\rho^{k+1}(A) \subseteq \rho^k(A)$ $(k \in \mathbb{N})$. This fact can be easily derived from the monotonicity of the justification operator: by definition, $\rho^1(A) = \rho(A) \subseteq A = \rho^0(A)$; if $\rho^k(A) \subseteq \rho^{k-1}(A)$, the monotonicity of ρ implies $\rho^{k+1}(A) = \rho(\rho^k(A)) \subseteq \rho(\rho^{k-1}(A)) = \rho^k(A)$. By the induction principle $\rho^{k+1}(A) \subseteq \rho^k(A)$ for every k. Every monotonically decreasing sequence of subsets has a well-defined limit: the intersection of all the sets in the sequence. Therefore it makes sense to define:

$$\rho^\infty(A) = \bigcap_{k\geq 1} \rho^k(A).$$

Example 13. In the game of Example 12, iterating ρ starting from A we delete one action at each step and we stop after 4 steps with one action left for each player:

$$\begin{aligned}
\rho(A) &= \rho(\{T,M,B\} \times \{\ell,c,r\}) = \{T,M,B\} \times \{\ell,c\}, \\
\rho^2(A) &= \rho(\{T,M,B\} \times \{\ell,c\}) = \{T,M\} \times \{\ell,c\}, \\
\rho^3(A) &= \rho(\{T,M\} \times \{\ell,c\}) = \{T,M\} \times \{\ell\}, \\
\rho^k(A) &= \rho(\{T,M\} \times \{\ell\}) = \{T\} \times \{\ell\}, \ (\forall k \geq 4).
\end{aligned}$$

▲

It is easy to see that $\rho^\infty(A)$ may contain more than one element:

	b	s	c
B	$4,3$	$0,2$	$0,0$
S	$0,1$	$3,4$	$0,0$
C	$1,1$	$1,2$	$5,0$

Figure 4.2: A modified Battle of the Sexes.

Example 14. Here is a story for the payoff matrix in Figure 4.2: Rowena and Colin have to decide independently of each other whether to go to a Bach concert, a Stravinsky concert, or a Chopin concert. Rowena likes Bach more than Stravinsky, she loves Chopin, but she prefers to go to another concert with Colin rather than listening to Chopin alone. Colin *hates* Chopin, he likes Stravinsky more than Bach, but he prefers Bach with Rowena to Stravinsky alone. This is a variation on the "Battle of the Sexes." Iterating ρ from $A = \{B,S,C\} \times \{b,s,c\}$ we get

$$\begin{aligned}
\rho(A) &= \rho(\{B,S,C\} \times \{b,s,c\}) = \{B,S,C\} \times \{b,s\}, \\
\rho^2(A) &= \rho(\{B,S,C\} \times \{b,s\}) = \{B,S\} \times \{b,s\}, \\
\rho^k(A) &= \rho(\{B,S\} \times \{b,s\}) = \{B,S\} \times \{b,s\}, \ (\forall k \geq 3).
\end{aligned}$$

Strategic thinking leads Rowena to avoid the Chopin concert, but it does not allow to solve the Bach-Stravinsky coordination problem. ▲

We argued informally that an action profile a is consistent with rationality and common belief in rationality if and only if $a \in \rho^{\infty}(A)$. Such action profiles are called "rationalizable":

Definition 13. *(Bernheim [26], Pearce [66]) An action profile $a \in A$ is* ***rationalizable*** *if $a \in \rho^{\infty}(A)$.*

In finite games, the set of rationalizable action profiles is obtained in a finite number of steps (the number depends on the game). In infinite games with compact action spaces and continuous payoff functions the set of rationalizable profiles is obtained with countably many steps (i.e., finitely many, or a denumerable sequence of steps).

Theorem 2. *(a) If G is a finite game, then there exists a positive integer K such that $\rho^{K+1}(A) = \rho^{K}(A) = \rho^{\infty}(A) \neq \emptyset$. (b) If G is a compact-continuous game, then the set of rationalizable action profiles, $\rho^{\infty}(A)$, is nonempty, compact, and satisfies $\rho^{\infty}(A) = \rho\left(\rho^{\infty}(A)\right)$.*

Proof. For the proof of part (b), which contains measure-theoretic arguments, see Appendix 4.6. Here we prove only part (a). If A_i is finite, for every conjecture μ^i the set of best replies $r_i(\mu^i)$ is nonempty. Then for each nonempty Cartesian set $C \subseteq A$, $\rho(C)$ is nonempty (for every i, there exists at least a conjecture $\mu^i \in \Delta(C_{-i})$ with $\emptyset \neq r_i(\mu^i) \subseteq \rho_i(C_{-i})$). Thus, $\rho^k(A) \neq \emptyset$ implies $\rho^{k+1}(A) \neq \emptyset$. Since $A \neq \emptyset$, it follows by induction that $\rho^k(A) \neq \emptyset$ for every k.

The sequence of subsets $(\rho^k(A))_{k=1}^{\infty}$ is weakly decreasing. Also, if $\rho^k(A) = \rho^{k+1}(A)$, then $\rho^k(A) = \rho^{\ell}(A)$ for every $\ell \geq k$. Given that A is finite, the inclusion $\rho^{k+1}(A) \subseteq \rho^k(A)$ can be strict only for a finite number of steps K (in particular, $K \leq \sum_{i \in I}(|A_i| - 1)$, where $|A_i|$ denotes the number of elements of A_i: when $\rho^{k+1}(A) \subset \rho^k(A)$, at least one action for at least one player i is eliminated, but at least one action for each player is never eliminated). All the above implies that $\rho^{K+1}(A) = \rho^{K}(A) = \rho^{\infty}(A) \neq \emptyset$. ■

It is possible to give an alternative and useful characterization of rationalizable actions. Consider the following:

Definition 14. *A set $C \in \mathcal{C}$ has the* ***best reply property*** *if $C \subseteq \rho(C)$.*

Remark 10. *By Theorem 2, the set $\rho^\infty(A)$ of rationalizable action profiles of a finite or compact-continuous game has the best reply property.*

We leave the proof of the following as an exercise:

Remark 11. *For every pair of subsets $E, F \in \mathcal{C}$, if both E and F have the best reply property then also $C = \times_{i\in I}(E_i \cup F_i)$ has the best reply property.*

Theorem 3. *In every finite or compact-continuous game, an action profile $a \in A$ is rationalizable if and only if $a \in C$ for some subset $C \in \mathcal{C}$ with the best reply property.*

Proof. **(Only if)** If a is rationalizable then $a \in \rho^\infty(A) = \rho(\rho^\infty(A))$, where the equality follows from Theorem 2. Hence, a belongs to a set with the best reply property, viz., $\rho^\infty(A)$.

(If) Let $a \in C \subseteq \rho(C)$ for some $C \in \mathcal{C}$. We will prove by induction that $C \subseteq \rho^k(C) \subseteq \rho^k(A)$ for each k. Therefore $a \in C \subseteq \bigcap_{k\in\mathbb{N}} \rho^k(A) = \rho^\infty(A)$, i.e., a is rationalizable.

Basis step. Since $C \subseteq \rho(C)$, $C \subseteq A$ and ρ is monotone, it follows that $C \subseteq \rho^1(C) \subseteq \rho^1(A)$.

Inductive step. Suppose that $C \subseteq \rho^k(C) \subseteq \rho^k(A)$. By monotonicity

$$\rho(C) \subseteq \rho(\rho^k(C)) \subseteq \rho(\rho^k(A)).$$

Since $C \subseteq \rho(C)$ and $\rho(\rho^k(\cdot)) = \rho^{k+1}(\cdot)$, we obtain

$$C \subseteq \rho^{k+1}(C) \subseteq \rho^{k+1}(A).$$

■

Observation 4. *The proof of Theorem 3 clarifies that, in every finite or compact-continuous game, for every $C \in \mathcal{C}$ with the best reply property, $C \subseteq \rho^\infty(A)$. Since Theorem 2 implies that $\rho^\infty(A)$ has the best reply property, then $\rho^\infty(A)$ must be the largest set with the best reply property.*[8]

We gave a definition of rationalizability based on iterations of the justification operator ρ because it is the most intuitive. An alternative

[8]More precisely, given the partially ordered collection $(\mathcal{C}, \subset)$, $\rho^\infty(A)$ is the unique maximal set with the best reply property. In other words, $\rho^\infty(A)$ is the unique set $C \in \mathcal{C}$ such that $C \subseteq \rho(C)$ and, for every $D \in \mathcal{C}$, $C \subset D$ implies $D \nsubseteq \rho(D)$.

definition of rationalizability is suggested by Theorem 3: a is rationalizable if it belongs to some set $C \in \mathcal{C}$ with the best reply property. This second definition is equivalent to the first one for all games where $\rho^{\infty}(A) = \rho(\rho^{\infty}(A))$. But there are some "badly behaved" infinite games where $\rho(\rho^{\infty}(A)) \subset \rho^{\infty}(A)$ (where $\subset$ means *strict* inclusion). The second definition of rationalizability is valid for those games too, while the first one only gives a necessary condition. For the purposes of these lectures, this "technicality" can be neglected. We used the first definition because we find it more intuitive.

4.3.1 Comparative Risk Aversion and Rationalizability

We have shown in Theorem 1 of Chapter 3 that higher risk aversion implies a larger set of justifiable actions. An inductive argument extends such monotonicity result to the set of rationalizable action profiles: a weak increase in risk aversion for all players weakly expands the set of rationalizable actions of each player.[9] The interesting novelty is that the set of rationalizable actions of a player may expand just because the risk aversion of *another* player increases as shown by the following example:

Example 15. Consider the following two-person game form with monetary consequences:

$g_1, g_2:$	b'	b''
a'	0, 1	1, 0
a''	$\frac{1}{3}$,0	$\frac{1}{3}$,1
a'''	1, 1	0, 0

Let $u_{\theta,1} = g_1^{\frac{1}{\theta}}$, where $\theta \geq 1$ parametrizes risk aversion, and let $u_2 = v_2 \circ g_2$ for any continuous and strictly increasing v_2 (the risk attitudes of player 2 are immaterial). We consider the rationalizability correspondence $\theta \mapsto \rho_{\theta}^{\infty}(A)$. Note that every action of player 2 is a best response to some belief, and the set of rationalizable actions of player 2 is $\{b', b''\}$ if a'' is justifiable for player 1, and $\{b'\}$ if a'' is unjustifiable. In the latter case,

[9] See Weinstein [89], who also consider other solution concepts. For an analogous result about "ambiguity aversion" see the working paper version of Battigalli et al. [21, Proposition 1].

the only rationalizable action of player 1 is a''', the best reply to b'. With this, the calculations of Example 6 imply

$$\rho_\theta^\infty(A) = \rho_\theta^3(A) = \begin{cases} \{(a''', b')\}, & \text{if } 1 \le \theta < \log_2 3, \\ A_1 \times A_2, & \text{if } \theta \ge \log_2 3. \end{cases}$$

▲

Theorem 4. *Let $G = \langle I, (A_i, u_i)_{i\in I}\rangle$ and $\hat{G} = \langle I, (A_i, \hat{u}_i)_{i\in I}\rangle$ be two compact-continuous games, and suppose every player $i \in I$ is more risk averse in $\hat{G}$ than in G. Then the set of rationalizable profiles of G is included in the set of rationalizable profiles of $\hat{G}$.*

Proof. Let A_i^k and $\hat{A}_i^k$ respectively denote the sets of k-step rationalizable actions of i in G and $\hat{G}$. We prove by induction that $A_i^k \subseteq \hat{A}_i^k$ for every $i \in I$ and $k \in \mathbb{N}$.

Basis step. By Theorem 1, $A_i^1 = r_i(\Delta(A_{-i})) \subseteq \hat{r}_i(\Delta(A_{-i})) = \hat{A}_i^1$ for every $i \in I$, where $\hat{r}_i$ is i's best reply correspondence in game $\hat{G}$, that is, according to payoff function $\hat{u}_i$.

Inductive Step. Suppose, by way of induction, that $A_j^k \subseteq \hat{A}_j^k$ for every $j \in I$ and fix any $i \in I$, then $A_{-i}^k \subseteq \hat{A}_{-i}^k$. Since $A_i^{k+1} = r_i\left(\Delta\left(A_{-i}^k\right)\right)$ and $\hat{A}_i^{k+1} = \hat{r}_i\left(\Delta\left(\hat{A}_{-i}^k\right)\right)$, we must show that $r_i\left(\Delta\left(A_{-i}^k\right)\right) \subseteq \hat{r}_i\left(\Delta\left(\hat{A}_{-i}^k\right)\right)$. By Theorem 1, $r_i\left(\Delta\left(A_{-i}^k\right)\right) \subseteq \hat{r}_i\left(\Delta\left(A_{-i}^k\right)\right)$. Since $A_{-i}^k \subseteq \hat{A}_{-i}^k$ by the inductive hypothesis, then also $\Delta\left(A_{-i}^k\right) \subseteq \Delta\left(\hat{A}_{-i}^k\right)$ and $\hat{r}_i\left(\Delta\left(A_{-i}^k\right)\right) \subseteq \hat{r}_i\left(\Delta\left(\hat{A}_{-i}^k\right)\right)$. Therefore $r_i\left(\Delta\left(A_{-i}^k\right)\right) \subseteq \hat{r}_i\left(\Delta\left(A_{-i}^k\right)\right) \subseteq \hat{r}_i\left(\Delta\left(\hat{A}_{-i}^k\right)\right)$, which implies $r_i\left(\Delta\left(A_{-i}^k\right)\right) \subseteq \hat{r}_i\left(\Delta\left(\hat{A}_{-i}^k\right)\right)$. ■

4.4 Iterated Dominance

The equivalence between best replies and undominated actions allows us to characterize the actions that survive the iterated dominance procedure. In order to give the precise definition of this procedure, we first define the concept of dominance within a subset of action profiles.

Definition 15. *Fix a nonempty Cartesian subset C of A: $\emptyset \neq C \in \mathcal{C}$. An action a_i is **dominated in** C if $a_i \in C_i$ and there is a mixed action*

$\alpha_i \in \Delta(C_i)$ *such that*

$$\forall a_{-i} \in C_{-i},\ u_i(a_i, a_{-i}) < u_i(\alpha_i, a_{-i}).$$

We denote by $\mathrm{ND}_i(C) \subseteq A_i$ *the set of actions in* C_i *that are* not dominated *in* C *and let* $\mathrm{ND}(C) = \times_{i\in I} \mathrm{ND}_i(C) \subseteq A$ *denote the set of undominated action profiles in* C.

Observation 5. *Operator* ND *is not monotone.*

Using the standard notation for iterations (k-fold composition) of self-maps, we can represent the iterated dominance procedure through the following sequence of sets $\mathrm{ND}(A)$, $\mathrm{ND}(\mathrm{ND}(A)) = \mathrm{ND}^2(A)$, ..., $\mathrm{ND}^k(A)$, ... Essentially, the idea is to first eliminate all dominated actions, thus obtaining $\mathrm{ND}(A)$. Then one moves on to eliminate all dominated actions in the *restricted game* with set of action profiles $\mathrm{ND}(A)$, thus obtaining $\mathrm{ND}^2(A)$, then eliminate all the dominated actions from the restricted game with set of action profiles $\mathrm{ND}^2(A)$, thus obtaining $\mathrm{ND}^3(A)$, and so on.

Definition 16. $a \in A$ *is a profile of* ***iteratively undominated actions*** *if* $a \in \mathrm{ND}^\infty(A) = \bigcap_{k\geq 1} \mathrm{ND}^k(A)$.

Theorem 5. *(Pearce [66]) Fix a finite or compact-continuous game; for every* $k = 1, 2, ...,$ $\rho^k(A) = \mathrm{ND}^k(A)$. *Therefore, an action profile is rationalizable if and only if it is iteratively undominated.*

We first prove a quite simple preliminary result about the sequence of subsets $\left(\rho^k(A)\right)_{k=1}^\infty$. For any nonempty subset $C_i \subseteq A_i$ define the **constrained best reply** correspondence $r_i(\cdot|C_i) : \Delta(A_{-i}) \rightrightarrows C_i$ as

$$r_i(\mu^i|C_i) = \arg\max_{a_i\in C_i} u_i(a_i, \mu^i).$$

Recall that $r_i(\mu^i) = \arg\max_{a_i\in A_i} u_i(a_i, \mu^i)$. Therefore, for each $\mu^i \in \Delta(A_{-i})$,

$$\emptyset \neq r_i(\mu^i) \subseteq C_i \Rightarrow r_i(\mu^i) = r_i(\mu^i|C_i); \tag{4.4.1}$$

in words, if there is at least one best reply (as is the case in all finite, or compact-continuous games) and all the best replies satisfy the constraint of being in C_i, then constrained and unconstrained best replies must coincide (check you can prove this; more generally, prove that $r_i(\mu^i) \cap C_i \neq \emptyset \Rightarrow r_i(\mu^i|C_i) \subseteq r_i(\mu^i)$).[10] For each $C \in \mathcal{C}$ and $i \in I$, let

$$\overline{\rho}_i(C) = r_i(\Delta(C_{-i})|C_i) = \bigcup_{\mu^i \in \Delta(C_{-i})} \arg\max_{a_i \in C_i} u_i(a_i, \mu^i);$$

$$\overline{\rho}(C) = \bigtimes_{i \in I} \overline{\rho}_i(C).$$

Since $\overline{\rho}_i(C)$ is obtained by *constrained* maximization, it follows that $\overline{\rho}(C) \subseteq C$ for every C. The justification operator ρ, instead, does not satisfy this property (if C, for instance, is the set of dominated action profiles, then $\rho(C) \cap C = \emptyset$). Also, like ND, operator $\overline{\rho}$ is not monotone, whereas ρ is monotone. Yet, ρ and $\bar{\rho}$ coincide whenever the constraint "does not bite":

Lemma 13. *Suppose that $r_i(\mu^i) \neq \emptyset$ for each $i \in I$ and $\mu^i \in \Delta(A_{-i})$. Then,*

$$\forall C \in \mathcal{C},\ \rho(C) \subseteq C \Rightarrow \rho(C) = \bar{\rho}(C). \tag{4.4.2}$$

Proof. Suppose that $\rho(C) \subseteq C$. Since $\rho(C) = \bigtimes_{i \in I} r_i(\Delta(C_{-i}))$, $r_i(\mu^i) \subseteq r_i(\Delta(C_{-i})) \subseteq C_i$ for each i and $\mu^i \in \Delta(C_{-i})$ (the first inclusion holds by definition, the second by assumption). Under the stated assumptions $r_i(\mu^i) \neq \emptyset$ for every i and μ^i. Therefore, by (4.4.1), $r_i(\mu^i) = r_i(\mu^i|C_i)$ for every i and $\mu^i \in \Delta(C_{-i})$, which implies $\rho_i(C_{-i}) = r_i(\Delta(C_{-i})) = r_i(\Delta(C_{-i})|C_i) = \bar{\rho}_i(C)$ for each i, that is, $\rho(C) = \bar{\rho}(C)$. ■

This implies that if ρ and $\bar{\rho}$ are iterated starting from the set A of *all* action profiles, the same sequence of subsets obtains:

Lemma 14. *In a finite or compact-continuous game, for every $k = 1, 2, ...,$ $\rho^k(A) = \overline{\rho}^k(A)$.*[11]

Proof. By definition $\rho^0(A) = \overline{\rho}^0(A)$. Suppose by way of induction that $\rho^k(A) = \overline{\rho}^k(A)$. As already shown, the monotonicity of ρ implies $\rho(\rho^k(A)) = \rho^{k+1}(A) \subseteq \rho^k(A)$; hence $\rho(\rho^k(A)) = \rho(\overline{\rho}^k(A)) \subseteq \overline{\rho}^k(A)$. Under

[10] In some infinite games, for some i, μ^i, and C_i, $r_i(\mu^i) = \emptyset$ and $r_i(\mu^i|C_i) \neq \emptyset$. See Example 16.

[11] As should be clear from the proof, the lemma holds for all games where the best reply correspondences are non-empty-valued.

the compactness-continuity assumption, $r_i(\mu^i) \neq \emptyset$ for each $i \in I$ and $\mu^i \in \Delta(A_{-i})$. Therefore eq. (4.4.2) in Lemma 13 applies and yields $\rho(\bar{\rho}^k(A)) = \bar{\rho}(\bar{\rho}^k(A))$. Thus, $\rho^{k+1}(A) = \rho(\rho^k(A)) = \bar{\rho}(\bar{\rho}^k(A)) = \bar{\rho}^{k+1}(A)$. ■

The lemma can be reformulated as follows. Say that an action a_i is **iteratively justifiable** if (1) it is justifiable (2) it is justifiable in the reduced game G^1 obtained by elimination of all the non justifiable actions, (3) it is justifiable in the reduced game G^2 obtained by the elimination of all non justifiable actions in G^1, and so on. The actions that are iteratively justifiable are exactly those obtained by iterating the operator $\bar{\rho}$. Hence, Lemma 14 states that, for every player, the set of rationalizable actions coincides with the set of iteratively justifiable actions.

Example 16. Here is an example that violates the compactness-continuity hypothesis of Theorem 2 and where, as a consequence, $r_i(\mu^i)$ may be empty and the thesis of Lemma 14 does not hold: $I = \{1, 2\}$, $A_1 = \{0, 1\}$, $A_2 = \{1 - \frac{1}{k} : k = 1, 2, ...\} \cup \{1\}$, the payoff functions are given by the infinite bi-matrix in Figure 4.3.

$1\backslash 2$	0	...	$1-\frac{1}{k}$	...	1
0	$1, 0$	...	$1, 1-\frac{1}{k}$	...	$1, 0$
1	$0, 0$	...	$0, 0$	...	$0, 1$

Figure 4.3: A countably infinite game.

Thus, $u_2\left(1 - \frac{1}{k}, \mu^2\right) = \mu^2(0)\left(1 - \frac{1}{k}\right)$, $u_2(1, \mu^2) = \mu^2(1) = 1 - \mu^2(0)$. If $\mu^2(1) \geq \frac{1}{2}$ the best reply is $a_2 = 1$. If $\mu^2(1) < \frac{1}{2}$, then $u_2(1 - \frac{1}{k}, \mu^2) > u_2(1, \mu^2)$ for k sufficiently large, but there is no best reply because $u_2(1 - \frac{1}{k}, \mu^2)$ is strictly increasing in k. Note also that $a_1 = 0$ strictly dominates $a_1 = 1$. Clearly $\rho(A) = \bar{\rho}(A) = \{0\} \times \{1\}$ and $\rho_2(\{0\}) = \emptyset$, $\bar{\rho}_2(\{0\} \times \{1\}) = \{1\}$. Therefore $\rho^2(A) = \emptyset \neq \{0\} \times \{1\} = \bar{\rho}^2(A)$. How is compactness-continuity violated? A_1 is finite, hence compact; A_2 is a closed and bounded subset of the real line (A_2 contains the limit of the sequence $(1 - \frac{1}{k})_{k=1}^{\infty}$), hence it is also compact; but u_2 is discontinuous at

$(0,1)$: $\lim_{k\to\infty} u_2\left(0, 1-\frac{1}{k}\right) = 1 \neq 0 = u_2(0,1)$.[12] ▲

Theorem 5 follows quite easily from Lemma 14: indeed Lemma 2 implies that the iteratively undominated actions coincide with the iteratively justifiable ones, and by Lemma 14 the iteratively justifiable actions coincide with the rationalizable actions. The details are as follows:

Proof of Theorem 5. *Basis step.* Lemma 2 implies that $\rho(A) = \mathrm{ND}(A)$.

Inductive step. Assume that $\rho^{k-1}(A) = \mathrm{ND}^{k-1}(A)$ (inductive hypothesis) and consider the game G^{k-1} where the set of actions of each player i is $\rho_i^{k-1}(A)$, and the payoff functions are obtained from the original game by restricting their domain to $\rho^{k-1}(A)$. The inductive hypothesis implies that the set of undominated action profiles in G^{k-1} is $\mathrm{ND}(\rho^{k-1}(A)) = \mathrm{ND}(\mathrm{ND}^{k-1}(A)) = \mathrm{ND}^k(A)$. By Lemma 2, $\mathrm{ND}(\mathrm{ND}^{k-1}(A)) = \overline{\rho}(\mathrm{ND}^{k-1}(A))$. The inductive hypothesis and Lemma 14 yield $\overline{\rho}(\mathrm{ND}^{k-1}(A)) = \overline{\rho}(\rho^{k-1}(A)) = \overline{\rho}(\overline{\rho}^{k-1}(A)) = \rho^k(A)$. Hence $\mathrm{ND}^k(A) = \rho^k(A)$. ■

So far, we have considered a procedure of iterated elimination which is **maximal**, in the sense that at any step *all* the dominated actions of all players are eliminated (where dominance holds in the restricted game that resulted from previous iterated eliminations). However, one can show that to compute the set of action profiles that are iteratively undominated (and therefore rationalizable) it is sufficient to iteratively eliminate *some* actions which are dominated for *some* player until in the restricted game so obtained there are no dominated actions left. To simplify the exposition, we restrict the analysis of non-maximal iterated dominance procedures to *finite* games.[13]

Definition 17. *An* ***iterated dominance procedure*** *is a sequence* $(C^k)_{k=0}^K \in \mathcal{C}^{K+1}$ *of nonempty Cartesian subsets of A such that (i) $C^0 = A$, (ii) for each $k = 1, \ldots, K$, $\mathrm{ND}(C^{k-1}) \subseteq C^k \subset C^{k-1}$ ($\subset$ is the* strict *inclusion), and (iii) $\mathrm{ND}(C^K) = C^K$.*

[12] For the mathematically savy: it is easy to change the example so that A_2 is not compact and u_2 is trivially continuous, just endow A_2 with the discrete topology.

[13] For an extension to infinite, compact-continuous games see Dufwenberg and Stegeman [41], who also provide counterexamples when compactness-continuity fails.

In words, an iterated dominance procedure is a sequence of steps starting from the full set of action profiles (condition *i*) and such that at each step k, for at least one player, at least one of the actions that are dominated given the previous steps is eliminated (condition *ii*); the elimination procedure can stop only when no further elimination is possible (condition *iii*).

Theorem 6. *Fix a finite game. For every iterated dominance procedure* $(C^k)_{k=0}^K$, C^K *is the set of rationalizable action profiles.*

Proof. Fix an iterated dominance procedure, i.e., a sequence of subsets $(C^k)_{k=0}^K$ that satisfies *(i)*-*(ii)*-*(iii)* of Definition 17.

Claim. *For each* $k = 0, 1, \ldots, K$,

$$\rho^{k+1}(A) \subseteq \rho(C^k) = \mathrm{ND}(C^k).$$

The proof of the claim is by induction:

Basis step. By Lemma 2 $\rho(A) = \mathrm{ND}(A)$; by *(i)* $C^0 = A$. Therefore $\rho^1(A) \subseteq \rho(C^0) = \mathrm{ND}(C^0)$ (the weak inclusion actually holds as an equality).

Inductive step. Suppose that $\rho^{k+1}(A) \subseteq \rho(C^k) = \mathrm{ND}(C^k)$. By *(ii)*, $\mathrm{ND}(C^k) \subseteq C^{k+1} \subset C^k$. By monotonicity of ρ and the inductive hypothesis,

$$\rho^{k+2}(A) \subseteq \rho(\mathrm{ND}(C^k)) \subseteq \rho(C^{k+1}),$$

and

$$\rho(C^{k+1}) \subseteq \rho(C^k) = \mathrm{ND}(C^k) \subseteq C^{k+1}.$$

Since the game is finite, there is at least one best reply to every conjecture and eq. (4.4.2) holds; thus,

$$\rho(C^{k+1}) = \bar{\rho}(C^{k+1}) = \mathrm{ND}(C^{k+1}),$$

where the second equality follows from Lemma 2. Collecting inclusions and equalities, $\rho^{k+2}(A) \subseteq \rho(C^{k+1}) = \mathrm{ND}(C^{k+1})$. □

The claim implies that $\rho^{K+1}(A) \subseteq \rho(C^K) = \mathrm{ND}(C^K)$. By *(iii)*, $C^K = \mathrm{ND}(C^K)$. Therefore $\rho^{\infty}(A) \subseteq \rho(C^K) = C^K$, that is, every rationalizable profile belongs to C^K, and C^K has the best reply property. By Theorem 3, C^K must be the set of rationalizable profiles. ■

4.5 Rationalizability in Nice Games

Recall that in a **nice game** the action set of each player is a compact interval in the real line, payoff functions are continuous and each u_i is strictly quasi-concave in the own action a_i (see Definition 11). Therefore, nice games are compact-continuous. It turns out that the analysis of rationalizability and iterated dominance in nice games is nice indeed! The analysis, however, requires some technicalities.

Recall that, for the games under considerations, $\mathcal{C}$ denotes the collection of *closed* subsets of A with a cross-product form $C = \times_{i\in I} C_i$. Define the following operators on $\mathcal{C}$: for every $C \in \mathcal{C}$,

$$\begin{aligned} \mathrm{r}(C) &= \times_{i\in I} r_i(C_{-i}), \quad \iota\rho(C) = \times_{i\in I} r_i(\mathrm{I}\Delta(C_{-i})), \\ \mathrm{ND}_p(C) &= \times_{i\in I}\{a_i \in C_i : \forall b_i \in C_i, \exists a_{-i} \in C_{-i}, u_i(a_i, a_{-i}) \geq u_i(b_i, a_{-i})\}, \end{aligned}$$

where $\mathrm{I}\Delta(C_{-i})$ is the set of product probability measures on C_{-i}, or independent (uncorrelated) conjectures (see Section 3.3.2). Note that r and $\iota\rho$ are self-maps on $\mathcal{C}$, because $r_i(C_{-i})$ and $r_i(\mathrm{I}\Delta(C_{-i}))$ are closed for every closed set C_{-i} (see the Appendix 4.6). Similarly, for every $C \in \mathcal{C}$, $\mathrm{ND}_p(C)$ is closed; hence, also ND_p is a self-map on $\mathcal{C}$. In fact, for every $i \in I$ and $b_i \in C_i$, since u_i is continuous, the set

$$\{a_i \in C_i : \exists a_{-i} \in C_{-i}, u_i(a_i, a_{-i}) \geq u_i(b_i, a_{-i})\}$$

is closed. Therefore, also the intersection of such sets

$$\begin{aligned} &\bigcap_{b_i\in C_i}\{a_i \in C_i : \exists a_{-i} \in C_{-i}, u_i(a_i, a_{-i}) \geq u_i(b_i, a_{-i})\} \\ &= \{a_i \in C_i : \forall b_i \in C_i, \exists a_{-i} \in C_{-i}, u_i(a_i, a_{-i}) \geq u_i(b_i, a_{-i})\} \end{aligned}$$

is closed.

Iterating $\iota\rho$ from A yields a solution procedure called **independent rationalizability**,[14] which is what obtains if (a) players are rational, (b) they hold independent (i.e., uncorrelated) conjectures, and (c) there is common belief of (a) and (b).

[14] This was the solution originally called "rationalizability" by Bernheim [26] and Pearce [66] before the epistemic analysis proved that the cleaner and more basic concept is the "correlated" version $\rho^\infty(A)$.

Iterating r from A yields a solution procedure called **point-rationalizability**, see Bernheim [26]. This procedure has no autonomous conceptual interest,[15] because there is no good reason to assume at the outset that players hold deterministic conjectures. Point rationalizability is just an ancillary solution concept that turns out to be convenient in the analysis of special games, including nice games.

Iterating ND_p from A yields the **iterated deletion of (pure) actions strictly dominated by (pure) actions**. In general games with compact action sets and continuous payoff functions, $\mathrm{ND}^k(A) \subseteq (\mathrm{ND}_p)^k(A)$ for all k.[16] Clearly $((\mathrm{ND}_p)^k(A))_{k=1}^{\infty}$ is a much simpler procedure than $(\mathrm{ND}^k(A))_{k=1}^{\infty}$, but it is weaker and it does not have a satisfactory conceptual foundation. So, in general, it should be regarded as an ancillary algorithm that allows to simplify the analysis of compact-continuous games.

As an exercise, use monotonicity to prove the following statements:

Remark 12. *For each k, $\mathrm{r}^k(A) \subseteq \iota\rho^k(A) \subseteq \rho^k(A)$. Hence $\mathrm{r}^\infty(A) \subseteq \iota\rho^\infty(A) \subseteq \rho^\infty(A)$.*

Remark 13. *For each $C \in \mathcal{C}$, if $C \subseteq \mathrm{r}(C)$ ($C \subseteq \iota\rho(C)$), then $C \subseteq \mathrm{r}^\infty(A)$ ($C \subseteq \iota\rho^\infty(A)$).*

Thus, if $\mathrm{r}^\infty(A) \neq \emptyset$, then $\iota\rho^\infty(A) \neq \emptyset$ and $\rho^\infty(A) \neq \emptyset$. Also, suppose that there is some $C \in \mathcal{C}$ that contains at least two elements and has the "deterministic" best reply property $C \subseteq \mathrm{r}(C)$. Then, the results above imply that there is a multiplicity of rationalizable (and independent rationalizable) action profiles. These simple results show why point rationalizability is a useful solution concept even though it does not have an interesting conceptual foundation. One nice feature of nice games is that Lemma 7 yields

Theorem 7. *In every nice game, for every $k \in \mathbb{N}$,*

$$\mathrm{r}^k(A) = \iota\rho^k(A) = \rho^k(A) = \mathrm{ND}^k(A) = (\mathrm{ND}_p)^k(A)$$

[15] Point rationalizability obtains if (a) players are rational, (b) they hold deterministic conjectures, and (c) there is common belief of (a) and (b).

[16] The result holds more generally for games with compact actions sets, where the payoff functions are jointly continuous in the opponents' actions and upper semi-continuous in the own action.

and each set is a product of closed intervals; hence

$$\mathrm{r}^{\infty}(A) = \iota\rho^{\infty}(A) = \rho^{\infty}(A) = \mathrm{ND}^{\infty}(A) = (\mathrm{ND}_p)^{\infty}(A)$$

and, again, each set is a product of closed intervals.

Proof. The second statement follows easily from the first, which is trivially true for $k = 0$. Suppose, by way of induction, that the first statement holds for a given k. By the inductive hypothesis the set $C = \mathrm{r}^k(A)$ is a product of closed intervals. Then Lemma 7 and Corollary 2 (applied to the restricted action sets C_i, $i \in I$) yield

$$\mathrm{r}(\mathrm{r}^k(A)) = \iota\rho(\mathrm{r}^k(A)) = \rho(\mathrm{r}^k(A)) = \mathrm{ND}(\mathrm{r}^k(A)) = \mathrm{ND}_p(\mathrm{r}^k(A)).$$

By the inductive hypothesis

$$\mathrm{r}^{k+1}(A) = \iota\rho^{k+1}(A) = \rho^{k+1}(A) = \mathrm{ND}^{k+1}(A) = (\mathrm{ND}_p)^{k+1}(A).$$

■

Example 17. As an application, consider the market for a crop analyzed in Chapter 1 (Section 1.6.1), but this time suppose that there is a finite number $n \geq 2$ of firms who decide how much to produce of a crop to be sold on the market 6 months later at the highest price that allows demand to absorb total output; thus, they compete *à la* Cournot.[17] Each firm has cost function

$$C(q_i) = \frac{1}{2m}(q_i)^2,\ q_i \in \left[0, \frac{4\alpha}{\beta}\right]$$

($\frac{4\alpha}{\beta}$ is a capacity constraint given by the available land).[18] *Market demand* is given by the function $D(p; n) = \max\{0, n(\alpha - \beta p)\}$ with $0 < \beta < 2$; therefore price as a function of *average quantity* is

$$P\left(\frac{1}{n}\sum_{i=1}^{n} q_i\right) = \max\left\{0, \frac{1}{\beta}\left(\alpha - \frac{1}{n}\sum_{i=1}^{n} q_i\right)\right\}.$$

It is interesting to relate the analysis of rationalizable oligopolistic behavior to the standard competitive equilibrium for these market fundamentals.

[17] The following analysis is adapted from Boergers and Janssen [30].

[18] The capacity constraint is loose enough to simplify the calculations.

(1) First, verify that the game is nice. The payoff function of each firm i is the profit function

$$\pi_i(q_i, q_{-i}) = \begin{cases} \frac{q_i}{\beta}\left(\alpha - \frac{q_i}{n} - \frac{1}{n}\sum_{j\neq i} q_j\right) - \frac{1}{2m}(q_i)^2, & \text{if } n\alpha > q_i + \sum_{j\neq i} q_j, \\ -\frac{1}{2m}(q_i)^2, & \text{if } n\alpha \leq q_i + \sum_{j\neq i} q_j. \end{cases}$$

This function is clearly continuous.[19] Now, fix q_{-i} arbitrarily. If $\sum_{j\neq i} q_j \geq n\alpha$ then $\pi_i(q_i, q_{-i}) = -\frac{1}{2m}(q_i)^2$, which is strictly decreasing in q_i. If $\sum_{j\neq i} q_j < n\alpha$, then π_i is initially increasing in q_i up to the point $r_i(q_{-i})$ where marginal revenue equals marginal cost (see below), and then it becomes strictly decreasing.[20]

(2) The best reply function is easily obtained from the first-order conditions when $\alpha > \frac{1}{n}\sum_{j\neq i} q_j$, in the other case the best reply is the corner solution $q_i = 0$:

$$r_i(q_{-i}) = \max\left\{0, \frac{m(\alpha - \frac{1}{n}\sum_{j\neq i} q_j)}{\beta + \frac{2m}{n}}\right\}.$$

(3) The monopolistic output is

$$q^{n,M} := r_i(0, \ldots, 0) = \frac{m\alpha}{\beta + \frac{2m}{n}}.$$

[19]Note that $\frac{q_i}{\beta}\left(\alpha - \frac{q_i}{n} - \frac{1}{n}\sum_{j\neq i} q_j\right) - \frac{1}{2m}(q_i)^2 = -\frac{1}{2m}(q_i)^2$ if $n\alpha = q_i + \sum_{j\neq i} q_j$.

[20]There is a kink at $\hat{q}_i = n\alpha - \sum_{j\neq i} q_j$, but this is immaterial. Note also that π_i is *not* concave in q_i: in the left neighborhood of the kink point $\hat{q}_i$ the derivative is

$$\begin{aligned} \frac{\partial \pi_i}{\partial q_i}(q_i, q_{-i}) &= P\left(\frac{q_i}{n} + \frac{1}{n}\sum_{j\neq i} q_j\right) + \frac{q_i}{n}P'\left(\frac{q_i}{n} + \frac{1}{n}\sum_{j\neq i} q_j\right) - \frac{q_i}{m} \\ &\approx \frac{\hat{q}_i}{n}P'\left(\frac{\hat{q}_i}{n} + \frac{1}{n}\sum_{j\neq i} q_j\right) - \frac{\hat{q}_i}{m} < -\frac{\hat{q}_i}{m} \end{aligned}$$

where the approximation holds because q_i is close to $\hat{q}_i$ and $P\left(\frac{q_i}{n} + \frac{1}{n}\sum_{j\neq i} q_j\right) \approx 0$, and the inequality holds because $P' < 0$. In a right neighborhood of $\hat{q}_i$, $\frac{\partial \pi_i}{\partial q_i}(q_i, q_{-i}) = -\frac{q_i}{m} \approx -\frac{\hat{q}_i}{m}$. Therefore, for $\varepsilon > 0$ small enough,

$$\frac{\partial \pi_i}{\partial q_i}\pi_i(\hat{q}_i - \varepsilon, q_{-i}) < \frac{\partial \pi_i}{\partial q_i}\pi_i(\hat{q}_i + \varepsilon, q_{-i}),$$

which implies that π_i is *not* concave.

If every competitor produces at maximum capacity, $\frac{1}{n}\sum_{j\neq i} q_j > \alpha$ (because $\beta < 2$ and capacity is $4\alpha/\beta$), so the best reply is

$$r_i\left(\frac{4\alpha}{\beta}, \dots, \frac{4\alpha}{\beta}\right) = 0.$$

Since the game is nice, $\rho(A) = \left[0, q^{n,M}\right]^n$. If this set has the best reply property, this is also the set of rationalizable profiles. To check this, it is enough to verify whether the best reply to the most pessimistic conjecture consistent with rationality is zero, that is $r_i(q^{n,M}, ..., q^{n,M}) = 0$; in the affirmative case $\left[0, q^{n,M}\right]^n$ has the best reply property. The condition for this is $\frac{n-1}{n}\frac{m\alpha}{\beta+\frac{2m}{n}} \geq \alpha$, or

$$\beta \leq m\left(\frac{n-3}{n}\right).$$

Therefore, if $\beta < m$ there is $\bar{n}$ large enough so that, for each $n > \bar{n}$, $\left[0, q^{n,M}\right]^n$ has the best reply property, hence there is a huge multiplicity of rationalizable outcomes. Conversely, if $\beta > m$, then

$$\beta > m\left(\frac{n-3}{n}\right)$$

for every n; hence, in this case the set $\left[0, q^{n,M}\right]^n$ does *not* have the best reply property whatever the number n of firms. Furthermore, it can be shown that if $\beta > m$ there is a unique rationalizable outcome.[21] By symmetry, each firm has the same rationalizable output $q^{n,*}$ which must solve $q = \frac{m(\alpha - \frac{n-1}{n}q)}{\beta+\frac{2m}{n}}$, because the singleton $\{(q^{n,*}, \dots, q^{n,*})\}$ has the best reply property. Therefore

$$q^{n,*} = \frac{m\alpha}{\beta + m\frac{n+1}{n}}$$

with corresponding rationalizable price

$$p^{n,*} := P\left(\frac{m\alpha}{\beta + m\frac{n+1}{n}}\right) = \frac{\alpha + \frac{m}{n}\frac{\alpha}{\beta}}{\beta + m\frac{n+1}{n}}$$

[21] If $\beta > m$, the joint best reply function

$$(q_i)_{i=1}^n \mapsto (r_i(q_{-i}))_{i=1}^n$$

is a contraction, which implies that the sequence $\left(\rho^k(A)\right)_{k\in\mathbb{N}}$ shrinks to a point.

(of course, this is the Cournot-Nash equilibrium).
(4) Finally, the competitive equilibrium price with these market fundamentals is

$$p^* := \frac{\alpha}{\beta + m},$$

with average output

$$q^* := \frac{m\alpha}{\beta + m}.$$

As pointed out in Section 1.6.1 of the Chapter 1, $\beta > m$ is the "cobweb stability" condition. Under this condition, the unique rationalizable output is $q^{n,*} = \frac{m\alpha}{\beta + m\frac{n+1}{n}}$. Of course,

$$\lim_{n\to\infty} q^{n,*} = \lim_{n\to\infty} \frac{m\alpha}{\beta + m\frac{n+1}{n}} = \frac{m\alpha}{\beta + m} = q^*,$$
$$\lim_{n\to\infty} p^{n,*} = \lim_{n\to\infty} \frac{\alpha + \frac{m}{n}\frac{\alpha}{\beta}}{\beta + m\frac{n+1}{n}} = \frac{\alpha}{\beta + m} = p^*.$$

To sum up, under the "cobweb stability" condition $\beta > m$, the rationalizable outcome is unique and approximates the so called "rational expectations equilibrium." ▲

4.6 Appendix: Compact-Continuous Games

4.6.1 Preliminaries on the Justification Operator

Recall that, for the case of compact-continuous games, we defined $\mathcal{C}$ as the collection of *closed* Cartesian subsets of A. In light of the results presented in Section 3.4.2, we can show that ρ is a self map from $\mathcal{C}$ to $\mathcal{C}$.

Lemma 15. *For every $C \in \mathcal{C}$, $\rho(C) \in \mathcal{C}$; furthermore, if $C \neq \emptyset$ then $\rho(C) \neq \emptyset$.*

Proof. Let $C = \times_{i\in I} C_i \in \mathcal{C}$. Since C is closed and $\rho(C) = \times_{i\in I} r_i(\Delta(C_{-i}))$, Lemma 10 implies that the Cartesian set $\rho(C)$ is also closed; hence $\rho(C) \in \mathcal{C}$; furthermore, if $C \neq \emptyset$, Lemma 10 implies that $\rho(C) \neq \emptyset$. ■

For future reference, we record an important property of compact sets:

Lemma 16. (Finite-intersection property of compact sets) *Let X be compact and let $\{C_j\}_{j\in J}$ be an indexed family of* closed *subsets of X such that, for every finite $F \subseteq J$, the finite indexed subfamily $\{C_j\}_{j\in F}$ has nonempty intersection: $\bigcap_{j\in F} C_j \neq \emptyset$. Then also $\{C_j\}_{j\in J}$ has nonempty intersection: $\bigcap_{j\in J} C_j \neq \emptyset$.*

4.6.2 Proof of Theorem 2 (b)

We must prove that $\rho^{\infty}(A)$ is nonempty, compact, and satisfies $\rho^{\infty}(A) = \rho(\rho^{\infty}(A))$. Recall that we defined $\rho^0(A) = A$ for convenience. We first prove by induction on $k \in \mathbb{N}_0$ (where $\mathbb{N}_0$ is the set of nonnegative integers) that $\rho^k(A) \neq \emptyset$ for every $k \geq 0$. The basis step is trivial: $\rho^0(A) = A$ where $A \neq \emptyset$ by assumption. Now assume by way of induction that $\rho^k(A) \in \mathcal{C}$ is nonempty for some k. Then Lemma 15 implies that $\rho^{k+1}(A) = \rho\left(\rho^k(A)\right) \neq \emptyset$. Next we show that $\rho^{\infty}(A) \neq \emptyset$. Since the sequence $\left(\rho^k(A)\right)_{k=1}^{\infty}$ is weakly decreasing,

$$\forall \ell \in \mathbb{N}, \bigcap_{k=1}^{\ell} \rho^k(A) = \rho^{\ell}(A) \neq \emptyset.$$

Then, the finite-intersection property of compact sets applied to the indexed family $\{\rho^{\ell}(A)\}_{\ell\in\mathbb{N}}$ (Lemma 16) implies that

$$\rho^{\infty}(A) = \bigcap_{\ell=1}^{\infty} \rho^{\ell}(A) \neq \emptyset.$$

Since $\rho^{\infty}(A)$ is an intersection of closed set, it is closed. Since A is compact, the closed subset $\rho^{\infty}(A)$ is compact.

Finally, we prove that $\rho^{\infty}(A) = \rho(\rho^{\infty}(A))$. The inclusion $\rho(\rho^{\infty}(A)) \subseteq \rho^{\infty}(A)$ follows from the monotonicity of ρ: For every $k \in \mathbb{N}_0$, $\rho^{\infty}(A) \subseteq \rho^k(A)$. Since ρ is monotone,

$$\rho(\rho^{\infty}(A)) \subseteq \bigcap_{k=0}^{\infty} \rho\left(\rho^k(A)\right) = \bigcap_{\ell=1}^{\infty} \rho^{\ell}(A) = \rho^{\infty}(A).$$

Next we show that $\rho^{\infty}(A) \subseteq \rho(\rho^{\infty}(A))$. Pick any profile $(a_i)_{i\in I} \in \rho^{\infty}(A)$. For every k, let A_i^k denote the set of "k-rationalizable" actions of player

i, so that $\rho^k(A) = \times_{i\in I} A_i^k$ and $\rho^\infty(A) = \bigcap_{k=1}^{\infty} \times_{i\in I} A_i^k$. With this, for each $i \in I$, we can find a sequence of conjectures $\left(\mu_n^i\right)_{n=0}^{\infty}$ such that $\mu_n^i \in \Delta\left(A_{-i}^n\right) \subseteq \Delta(A_{-i})$ and $a_i \in r_i\left(\mu_n^i\right)$ for every $n \in \mathbb{N}_0$. By Lemma 9, $\left(\mu_n^i\right)_{n=0}^{\infty}$ has a convergent subsequence $\left(\mu_{n_k}^i\right)_{k=1}^{\infty}$ with limit $\mu_{n_k}^i \to \bar{\mu}^i \in \Delta(A_{-i})$. By Lemma 10, $a_i \in r_i\left(\bar{\mu}^i\right)$ for each $i \in I$. To prove that $(a_i)_{i\in I} \in \rho(\rho^\infty(A))$, we only have to show that $\bar{\mu}^i\left(\bigcap_{n=1}^{\infty} A_{-i}^n\right) = 1$ for each $i \in I$. Note that, for every n, there is some k such that $n_k \geq n$, so that $\mu_{n_k}^i\left(A_{-i}^n\right) = 1$ because $A_{-i}^{n_k} \subseteq A_{-i}^n$. Therefore, by the *portmanteau theorem* [3, Theorem 15.3], for every A_{-i}^n (a closed set)

$$\bar{\mu}^i\left(A_{-i}^n\right) \geq \limsup_{k\to\infty} \mu_{n_k}^i\left(A_{-i}^n\right) = 1,$$

that is, $\bar{\mu}^i\left(A_{-i}^n\right) = 1$. Since $\bar{\mu}^i$ is countably additive, it must be continuous.[22] Therefore,

$$\bar{\mu}^i\left(\bigcap_{n=1}^{\infty} A_{-i}^n\right) = \lim_{n\to\infty} \bar{\mu}^i\left(A_{-i}^n\right) = 1.$$

■

[22] This means that $\bar{\mu}^i\left(\bigcap_{k=1}^{\infty} E_k\right) = \lim_{k\to\infty} \bar{\mu}^i(E_k)$ for every weakly *decreasing* sequence of measurable sets $(E_k)_{k=1}^{\infty}$, and $\bar{\mu}^i\left(\bigcup_{k=1}^{\infty} F_k\right) = \lim_{k\to\infty} \bar{\mu}^i(F_k)$ for every weakly *increasing* sequence of measurable subsets $(F_k)_{k=1}^{\infty}$.

5

Pure Equilibrium

In Chapter 4, we analyzed the behavioral implications of rationality and common belief in rationality when there is complete information, that is, common knowledge of the rules of the game and of players' preferences. There are interesting strategic situations where these assumptions about rationality, belief, and knowledge imply that each player can correctly predict the opponents' behavior, e.g., many oligopoly games. But this is not the case in general: whenever a game has multiple rationalizable profiles there is at least one player i such that many conjectures μ^i are consistent with common belief in rationality, but at most one of them can be correct. In this chapter, we analyze situations where players' predictions are correct, and we discuss why predictions should (or should not) be correct. Let us give a preview.

We start with *Nash*'s classic definition of equilibrium (in pure actions): An action profile $a^* = (a_i^*)_{i \in I}$ is an **equilibrium** if each action a_i^* is a best reply to the actions of the other players a_{-i}^*.

Nash equilibrium play follows from the assumptions that players are rational and hold correct conjectures about the behavior of other players. But why should this be the case? Can we find scenarios that make the assumption of correct conjectures compelling, or at least plausible? We will discuss two types of scenarios: (1) those that make Nash equilibrium an "obvious way to play the game," and (2) those that make Nash equilibrium a steady state of an adaptive process. In each case, we will point out that Nash equilibrium is justified under rather restrictive assumptions, and that weakening such assumptions suggests interesting generalizations of the

Nash equilibrium concept whereby players form probabilistic conjectures about the behavior of co-players. We then analyze these generalizations in Chapter 6 on "probabilistic equilibria."

More specifically, Nash equilibrium may be the "obvious way to play the game" under complete information either as the outcome of sophisticated strategic reasoning, or as the result of a non-binding pre-play agreement. If it were obvious how the game should be played, then each player i could predict the obvious way to play of the co-players, viz. a^*_{-i}, and play the best reply a^*_i. Yet, if sophisticated strategic reasoning is given by rationality and common belief in rationality, then Nash equilibrium is obtained only in those games where rationalizable actions and equilibrium actions coincide. This coincidence is implied by some assumptions about action sets and payoff functions that are satisfied in several economic applications, but in general the relevant concept under this scenario is rationalizability, not Nash equilibrium.

What about pre-play agreements? True, a non-binding agreement to play a particular action profile a^* has to be a Nash equilibrium, otherwise the agreement would be self-defeating. But it will be shown in Chapter 6 that more general "probabilistic agreements" that make behavior depend on some extraneous random variables may be more efficient. Such probabilistic agreements are called "**correlated equilibria**" and generalize the Nash equilibrium concept.

Next we turn to adaptive processes. The general scenario is that a given game is played recurrently by agents that each time are drawn at random from large populations corresponding to different players/roles in the game (e.g., seller and buyer, or male and female). The agents drawn from each population are matched, play among themselves, and then are separated. If each player's conjecture about the opponents' behavior in the current play is based on his observations of the opponents' actions in the past and this player best responds, we obtain a dynamic process in which at each point in time some agent switches from one action to another. *If* the process converges so that each agent playing in a given role ends up choosing the *same* action over and over, the steady state must look like a Nash equilibrium: indeed the experience of those playing in role i is that the opponents keep playing some profile a^*_{-i} and thus they keep choosing a best reply, say, a^*_i.[1] Since this is true for each i, $(a^*_i)_{i\in I}$ must be a Nash

[1] Suppose for simplicity that there is only one best reply.

equilibrium.

But it is possible that the process converges to a heterogeneous situation: a fraction q_i^1 of the agents in the population i choose some action a_i^1, another fraction q_i^2 of agents choose action a_i^2 and so on. If the process has stabilized, $(q_j^1, q_j^2, \ldots)$ will (approximately) represent the observed frequencies of the actions of those playing in role j, and i will conjecture that a_j^1, $a_j^2, \ldots$ are played with probabilities q_j^1, $q_j^2, \ldots$. If each action a_i^1, a_i^2, ... is a best reply to such conjecture, given a bit of inertia, i will keep choosing whatever he was choosing before and the heterogeneous behavior within each population will persist. In this case, the steady state is described by *mixed* actions whose support is made of best responses to the mixed actions of the opponents. This is called "**mixed (Nash) equilibrium**." It turns out that a mixed equilibrium formally is a Nash equilibrium of the extended game where each player i chooses in the set $\Delta(A_i)$ of mixed actions and payoffs are computed by taking expected values (under the assumption of independence across players). Indeed, this is the standard definition of mixed equilibrium.

The analysis above relies on the assumption that, when a game is played recurrently, each agent can observe the frequency of play of different actions. But often the information feedback is poorer. For example, one may observe only a variable determined by the players' actions, such as the number of customers of a firm in a price-setting oligopoly. In such a context, it may happen that players' conjectures are wrong and yet they are confirmed by players' information feedback, so that they keep best responding to such wrong conjectures and the system is in a steady state which is not a Nash (or mixed Nash) equilibrium. A steady state whereby players best respond to confirmed (non contradicted) conjectures is called "conjectural equilibrium," or "**self-confirming equilibrium**." Since correct conjectures (conjectures that correspond to the actual behavior of the other players) are necessarily confirmed, every Nash equilibrium is necessarily self-confirming.

To sum up, a serious effort to provide reasonable justifications for the Nash equilibrium concept leads us to think about scenarios where strategic interaction would plausibly lead to players best responding to correct conjectures about the opponents. But for each such scenario we need additional assumptions (on the payoff functions, or on information feedback) to justify Nash play. Without such additional assumptions we are left with interesting generalizations of the Nash equilibrium concept.

The rest of the chapter is organized as follows. In Section 5.1, we define (pure) Nash equilibrium and provide sufficient conditions for its existence. In Section 5.2 we focus on symmetric equilibria of nice games and provide a simple existence proof. In Section 5.3 we analyze nice games with strategic complementarities and show that in these games there is a tight relationship between Nash equilibrium and rationalizability. Finally, in Section 5.4 we go back to the interpretation of the Nash equilibrium concept, paving the way for the generalizations analyzed in Chaper 6.

5.1 Nash Equilibrium

A Nash[2] equilibrium is a situation in which every player is rational and holds *correct* conjectures about the actions of the other players.

Definition 18. *An action profile* $a^* = (a_i^*)_{i \in I}$ *is a* ***Nash equilibrium*** *if, for every* $i \in I$, $a_i^* \in r_i(a_{-i}^*)$.

Observation 6. *An action profile* a^* *is a Nash equilibrium if and only if the singleton* $\{a^*\}$ *has the best reply property. Hence, by Theorem 3, every Nash equilibrium is rationalizable; if there is a unique rationalizable action profile, then it is necessarily the unique Nash equilibrium.*

By inspection of simple games, such as "Matching Pennies," it is obvious that not all games have Nash equilibria. The following classic theorem provides *sufficient* conditions for the existence of at least one Nash equilibrium.

Theorem 8. *Consider a* compact-continuous *game* $G = \langle I, (A_i, u_i)_{i \in I} \rangle$. *If, for each player* $i \in I$, A_i *is* convex *and function* $u_i(\cdot, a_{-i}) : A_i \to \mathbb{R}$ *is* quasi-concave *for all* $a_{-i} \in A_{-i}$, *then* G *has a Nash equilibrium.*

[2] John Nash, who has been awarded the Nobel Prize for Economics (joint with John Harsanyi and Reinhard Selten) in 1994, was the first to give a general definition of this equilibrium concept. Nash analyzed the equilibrium concept for the *mixed* extension of a game, and proved the existence of mixed equilibria for all games with a finite number of pure actions (see Definition 21 and Theorem 13).

Almost all others equilibrium concepts used in non-cooperative game theory can be considered as generalizations or "refinements" of the equilibrium proposed by Nash. Perhaps, this is the reason why the other equilibrium concepts that have been proposed do not take their name after one of the researchers who introduced them in the literature.

The proof of Theorem 8 follows from preliminary results of independent interest. The first one states that the graph of the best-reply correspondence is closed:

Lemma 17. ***(Closed-Graph Lemma)*** *If A is compact and $u_i : A \to R$ is continuous, then the restriction of the best reply correspondence to sub-domain A_{-i}, $r_i|_{A_{-i}} : A_{-i} \rightrightarrows A_i$, is nonempty valued and its graph*

$$\mathrm{Gr}\left(r_i|_{A_{-i}}\right) = \{(a_{-i}, a_i) \in A_{-i} \times A_i : a_i \in r_i\left(a_{-i}\right)\}$$

is closed.

The closed-graph property is a special case of Lemma 10 in Appendix 4.6, which considers best replies to probabilistic conjectures. Since the proof of Lemma 10 involves measure-theoretic arguments, it is useful to provide a separate proof for this simpler case (see also the proof of Lemma 6).

Proof of Lemma 17. Since u_i is continuous, for each $a_{-i} \in A_{-i}$, the section $u_i\left(\cdot, a_{-i}\right) : A_i \to \mathbb{R}$ is also continuous, therefore, it attains a maximum on the compact set A_i. Thus, $r_i(a_{-i}) \neq \emptyset$ for each $a_{-i} \in A_{-i}$.

To show that $\mathrm{Gr}\left(r_i|_{A_{-i}}\right)$ is closed, we prove that the limit of every convergent sequence in $\mathrm{Gr}\left(r_i|_{A_{-i}}\right)$ also belongs to $\mathrm{Gr}\left(r_i|_{A_{-i}}\right)$. Let $\left(a_{-i}^n, a_i^n\right)_{n=1}^{\infty} \in \mathrm{Gr}\left(r_i|_{A_{-i}}\right)^{\mathbb{N}}$ be such that $\lim_{n\to\infty}\left(a_{-i}^n, a_i^n\right) = (\bar{a}_{-i}, \bar{a}_i)$. By definition of $\mathrm{Gr}\left(r_i|_{A_{-i}}\right)$, a_i^n is a best reply to a_{-i}^n for each n, thus,

$$\forall a_i \in A_i, \forall n \in \mathbb{N},\ u_i\left(a_i^n, a_{-i}^n\right) \geq u_i\left(a_i, a_{-i}^n\right). \tag{5.1.1}$$

Since u_i is continuous, taking the limit of each side of (5.1.1) as $n \to \infty$ for each a_i, we obtain

$$\forall a_i \in A_i,\ u_i\left(\bar{a}_i, \bar{a}_{-i}\right) \geq u_i\left(a_i, \bar{a}_{-i}\right),$$

that is, $\bar{a}_i \in r_i\left(\bar{a}_{-i}\right)$. Therefore $(\bar{a}_{-i}, \bar{a}_i) \in \mathrm{Gr}\left(r_i|_{A_{-i}}\right)$. ■

Next we report without proof an important fixed-point theorem:[3]

[3] For a proof, see Ok [64], p. 331.

Theorem 9. ***(Kakutani)*** *Let X be a nonempty, convex, and compact subset of $\mathbb{R}^n$ and let $\varphi : X \rightrightarrows X$ be a correspondence such that the graph $Gr(\varphi) = \{(x, y) \in X \times X : y \in \varphi(x)\}$ is closed and, for each $x \in X$, $\varphi(x)$ is nonempty and convex; then φ has a fixed point, i.e., there is some $x^* \in X$ such that $x^* \in \varphi(x^*)$.*

Proof of Theorem 8. To ease notation, we use the symbol r_i also for the restriction of the best-reply correspondence to sub-domain A_{-i}. By the Closed-Graph Lemma, each r_i is nonempty valued and has a closed graph. Fix a_{-i} arbitrarily and let $u_i^*(a_{-i}) = \max_{a_i \in A_i} u_i(a_i, a_{-i})$ denote the maximum payoff of i given a_{-i}. By quasi-concavity of u_i, for every $a_{-i} \in A_{-i}$ and every real number $y \in \mathbb{R}$, the set $\{a_i \in A_i : u_i(a_i, a_{-i}) \geq y\}$ is convex. In particular, the set $\{a_i \in A_i : u_i(a_i, a_{-i}) \geq u_i^*(a_{-i})\}$ is convex. Since

$$r_i(a_{-i}) = \{a_i \in A_i : u_i(a_i, a_{-i}) \geq u_i^*(a_{-i})\}$$

(where the weak inequality holds as an equality), it follows that $r_i(a_{-i})$ is convex. Now define the joint best reply correspondence $r : A \rightrightarrows A$ as follows:

$$\forall a \in A,\ r(a) = \bigtimes_{i \in I} r_i(a_{-i}).$$

Since $r_i : A_{-i} \rightrightarrows A_i$ is nonempty and convex valued with a closed graph for each $i \in I$, the same holds also for the joint best reply correspondence $r : A \rightrightarrows A$. Then, by the Fixed-Point Theorem of Kakutani, there exists $a^* \in A$ such that $a^* \in r(a^*)$, that is, $a_i^* \in r_i(a_{-i}^*)$ for each $i \in I$. Such fixed point is a Nash equilibrium. ∎

5.2 Equilibrium in Symmetric Nice Games

As one can see from the proof of Theorem 8, rather advanced mathematical tools are necessary to prove a relatively general existence result. Here, we analyze the more specific case in which all players are in a symmetric position and have unidimensional action sets. This allows us to use more elementary tools to show that, under convenient simplifying assumptions, there exists an equilibrium in which all players choose the same action. The fact that all players choose the same action is not part of the thesis of Theorem 8, which did not assume symmetry. Therefore the next result is not a special case of Theorem 8.[4] Before we state the theorem, we must formally spell out the definition of symmetric game and symmetric equilibrium.

Definition 19. *A static* game $G = \langle I, (A_i, u_i)_{i\in I}\rangle$ *is* ***symmetric*** *if the players have the same set of actions, denoted by* $\hat{A}$ *(hence, for every* $i \in I$, $A_i = \hat{A}$*) and*

$$u_i(a) = u_{\pi(i)}\left(a \circ \pi^{-1}\right)$$

for every player $i \in I$*, every action profile* $a : I \to \hat{A}$*, and every bijection (permutation)* $\pi : I \to I$*. A Nash* equilibrium a^* *of a symmetric game* G *is* symmetric *if* $a^* : I \to \hat{A}$ *is constant, that is, if* $a_i^* = a_j^*$ *for all* $i, j \in I$*.*

Note that, if $a = (a_i)_{i\in I}$ then $a' = a \circ \pi^{-1} = \left(a_{\pi^{-1}(i)}\right)_{i\in I}$. In words, if $j = \pi(i)$, then the action a_j' played by j in the permuted profile a' is the action $a_i = a_{\pi^{-1}(j)}$ played by i in the original profile a. Hence, symmetry requires

$$u_i\left((a_j)_{j\in I}\right) = u_{\pi(i)}\left(\left(a_{\pi^{-1}(j)}\right)_{j\in I}\right)$$

so that player $\pi(i)$ plays the *same* action in profile $a' = a \circ \pi^{-1}$ as player i in profile a.

Example 18. The Cournot oligopoly game of Example 7 is *symmetric* if all firms have the same capacity constraint and total cost function: for some $k > 0$, and $C : [0, k] \to \mathbb{R}$, $\bar{a}_i = k$ and $C_i(\cdot) = C(\cdot)$ for all $i \in I$. To see this, fix $a = (a_j)_{j\in I} \in [0, k]^I = \hat{A}^I$ arbitrarily and let $\pi : I \to I$ be a bijection, or permutation, then

$$\begin{aligned} u_i(a) &= a_i P\left(\sum_{j\in I} a_j\right) - C_i(a_i) \\ &= a_{\pi^{-1}(\pi(i))} P\left(\sum_{j\in I} a_{\pi^{-1}(j)}\right) - C_{\pi(i)}(a_{\pi^{-1}(\pi(i))}) \\ &= u_{\pi(i)}\left(\left(a_{\pi^{-1}(j)}\right)_{j\in I}\right) = u_{\pi(i)}\left(a \circ \pi^{-1}\right), \end{aligned}$$

where the second equality holds because $\pi^{-1}(\pi(i)) = i$, summations are commutative, and $C_i = C = C_{\pi(i)}$ by assumption. ▲

[4] However, by adding symmetry to the hypotheses of Theorem 8, one can prove the existence of a symmetric Nash equilibrium.

Recall that a game is **nice** if the action set of each player is a compact interval in the real line, payoff functions are continuous, and each player's payoff function is strictly quasi-concave in his own action.

Theorem 10. *Every* symmetric nice *game has a symmetric Nash equilibrium.*

Proof. To ease notation, we assume without loss of generality that the common action set is the interval $\hat{A} = [0, 1]$, which is just a normalization. Since u_i is continuous and strictly quasi-concave in a_i, there exists one and only one best reply to every (deterministic) conjecture a_{-i}. Let $\hat{r}_i(x) \in [0, 1]$ denote the unique best reply to the symmetric (deterministic) conjecture a_{-i} such that $a_j = x$ for each $j \in I \backslash \{i\}$. Since the game is symmetric, $\hat{r}_i(x)$ must be independent of i. Thus, we let $\hat{r} : [0, 1] \to [0, 1]$ denote the common **best reply function for symmetric conjectures**. By Lemma 6, each $r_i(\cdot)$ is a continuous function; therefore, $\hat{r}(\cdot)$ must be *continuous.*

Now, let us introduce an auxiliary function $f : [0, 1] \to \mathbb{R}$ as follows: $f(x) = x - \hat{r}(x)$. Function f is continuous[5] with $f(0) \leq 0$ and $f(1) \geq 0$. It follows from the Intermediate Value Theorem[6] that there exists a point $x^* \in [0, 1]$ such that $f(x^*) = 0$, that is, $x^* = \hat{r}(x^*)$. For each player $i \in I$, action $a_i^* = x^*$ is a best reply to the symmetric conjecture a_{-i}^* such that $a_j^* = x^*$ for each $j \neq i$. This means that the profile $(a_i^*)_{i \in I}$ with $a_i^* = x^*$ for each $i \in I$ is a symmetric Nash equilibrium. ■

The reason why we considered symmetric nice games is that they are often used in applied theory and we could use elementary results in real analysis to prove the existence of a symmetric equilibrium. A key step in the proof is that every continuous function $\hat{r} : [0, 1] \to [0, 1]$ has a **fixed point** $x^* = \hat{r}(x^*)$. This is an elementary special case of the Fixed-Point Theorem of Kakutani, because $\hat{r} : [0, 1] \to [0, 1]$ is continuous if and only if the singleton-valued correspondence $x \mapsto \{\hat{r}(x)\}$ (which is trivially nonempty and convex-valued) has a closed graph.

[5] If $\varphi, \psi : [0, 1] \to \mathbb{R}$ are continuous functions, then, for every $k \in \mathbb{R}$, $(\varphi + k\psi) : [0, 1] \to \mathbb{R}$ $(x \mapsto \varphi(x) + k\psi(x))$ is also continuous.

[6] See, for example, Binmore [27], p. 88.

5.3 Equilibrium in Supermodular Nice Games

In Section 3.3.3, we have seen that best reply correspondences in supermodular nice games are weakly increasing functions (Lemma 8). This has an important consequence: the set of Nash equilibria has a largest and a smallest element, and they characterize the set of rationalizable action profiles. To articulate and prove this result, we first record in the following lemma a structural property of the set $\rho^{\infty}(A)$ of rationalizable action profiles in a supermodular nice game, namely, $\rho^{\infty}(A)$ is an order box in $\mathbb{R}^I$ such that the lowest (respectively, largest) rationalizable action of each player is the best reply to the lowest (respectively, largest) rationalizable actions of the co-players.

Lemma 18. *Let $G = \langle I, (A_i, u_i)_{i\in I}\rangle$ be a supermodular nice game, and—for each $i \in I$—let $\underline{a}_i^0$ and $\bar{a}_i^0$ respectively denote the smallest and largest action of player i, that is, $A_i = [\underline{a}_i^0, \bar{a}_i^0]$. Then, for each $n \in \mathbb{N}$,*

$$\rho^n(A) = \bigtimes_{i\in I}[\underline{a}_i^n, \bar{a}_i^n] = \bigtimes_{i\in I}\left[r_i\left(\underline{a}_{-i}^{n-1}\right), r_i\left(\bar{a}_{-i}^{n-1}\right)\right].$$

Therefore,

$$\rho^{\infty}(A) = \bigtimes_{i\in I}[\underline{a}_i^*, \bar{a}_i^*] = \bigtimes_{i\in I}\left[r_i\left(\underline{a}_{-i}^*\right), r_i\left(\bar{a}_{-i}^*\right)\right],$$

where, for each $i \in I$, $\underline{a}_i^* = \lim_{n\to\infty}\underline{a}_i^n$ and $\bar{a}_i^* = \lim_{n\to\infty}\bar{a}_i^n$.

Proof. Since G is a nice game, Theorem 7 implies that, for each $n \in \mathbb{N}$, $\rho^n(A)$ is an order box of best replies to deterministic conjectures, that is,

$$\rho^n(A) = \bigtimes_{i\in I}\left[\min r_i\left([\underline{a}_{-i}^{n-1}, \bar{a}_{-i}^{n-1}]\right), \max r_i\left([\underline{a}_{-i}^{n-1}, \bar{a}_{-i}^{n-1}]\right)\right],$$

where, for each $i \in I$, $[\underline{a}_{-i}^{n-1}, \bar{a}_{-i}^{n-1}]$ is the order box in $\mathbb{R}^{I\setminus\{i\}}$ of the $(n-1)$-rationalizable action profiles of the co-players. By Lemma 8, each r_i is a weakly increasing function; therefore

$$\begin{aligned}\min r_i\left([\underline{a}_{-i}^{n-1}, \bar{a}_{-i}^{n-1}]\right) &= r_i\left(\underline{a}_{-i}^{n-1}\right),\\ \max r_i\left([\underline{a}_{-i}^{n-1}, \bar{a}_{-i}^{n-1}]\right) &= r_i\left(\bar{a}_{-i}^{n-1}\right).\end{aligned}$$

Therefore, $\rho^n(A) = \bigtimes_{i\in I}[\underline{a}_i^n, \bar{a}_i^n] = \bigtimes_{i\in I}\left[r_i\left(\underline{a}_{-i}^{n-1}\right), r_i\left(\bar{a}_{-i}^{n-1}\right)\right]$.

Since $(\rho^n(A))_{n=1}^{\infty} = \left(\bigtimes_{i\in I}[\underline{a}_i^n, \bar{a}_i^n]\right)_{n=1}^{\infty}$ is a weakly decreasing sequence of Cartesian subsets, for each $i \in I$, $(\underline{a}_i^n)_{n=1}^{\infty}$ is weakly increasing and $(\bar{a}_i^n)_{n=1}^{\infty}$ is weakly decreasing, therefore these sequences have limits $\underline{a}_i^* = \lim_{n\to\infty} \underline{a}_i^n$ and $\bar{a}_i^* = \lim_{n\to\infty} \bar{a}_i^n$. Thus,

$$\rho^{\infty}(A) = \bigcap_{n=1}^{\infty} \rho^n(A) = \bigcap_{n=1}^{\infty} \bigtimes_{i\in I}[\underline{a}_i^n, \bar{a}_i^n] = \bigtimes_{i\in I}[\underline{a}_i^*, \bar{a}_i^*].$$

By Theorem 2 and 7,

$$\rho^{\infty}(A) = \rho(\rho^{\infty}(A)) = \mathrm{r}\left(\bigtimes_{i\in I}[\underline{a}_i^*, \bar{a}_i^*]\right) = \bigtimes_{i\in I} r_i\left(\left[\underline{a}_{-i}^*, \bar{a}_{-i}^*\right]\right) = \bigtimes_{i\in I}[\underline{a}_i^*, \bar{a}_i^*].$$

Again, Lemma 8 implies that

$$\begin{aligned} \min r_i\left(\left[\underline{a}_{-i}^*, \bar{a}_{-i}^*\right]\right) &= r_i\left(\underline{a}_{-i}^*\right), \\ \max r_i\left(\left[\underline{a}_{-i}^*, \bar{a}_{-i}^*\right]\right) &= r_i\left(\bar{a}_{-i}^*\right) \end{aligned}$$

for each $i \in I$. Therefore, $\rho^{\infty}(A) = \bigtimes_{i\in I}[\underline{a}_i^*, \bar{a}_i^*] = \bigtimes_{i\in I}\left[r_i\left(\underline{a}_{-i}^*\right), r_i\left(\bar{a}_{-i}^*\right)\right]$. ∎

Lemma 18 implies the main result of this section:

Theorem 11. *For every supermodular nice game G the set of rationalizable action profiles is an order box $\rho^{\infty}(A) = [\underline{a}^*, \bar{a}^*]$ where $\underline{a}^*$ is the smallest Nash equilibrium and $\bar{a}^*$ is the largest Nash equilibrium, that is, $\underline{a}^* \leq a^* \leq \bar{a}^*$ for every Nash equilibrium a^*.*

Proof. By Lemma 18,

$$[\underline{a}^*, \bar{a}^*] = \bigtimes_{i\in I}[\underline{a}_i^*, \bar{a}_i^*] = \bigtimes_{i\in I}\left[r_i\left(\underline{a}_{-i}^*\right), r_i\left(\bar{a}_{-i}^*\right)\right].$$

Therefore $\underline{a}_i^* = r_i\left(\underline{a}_{-i}^*\right)$ and $\bar{a}_i^* = r_i\left(\bar{a}_{-i}^*\right)$ for each $i \in I$, which implies that $\underline{a}^*$ and $\bar{a}^*$ are Nash equilibria. Fix any Nash equilibrium a^*. Since a^* is rationalizable (see Observation 6), it follows that $a^* \in [\underline{a}^*, \bar{a}^*]$, that is, $\underline{a}^* \leq a^* \leq \bar{a}^*$. ∎

Corollary 3. *Let $G = \langle I, (A_i, u_i)_{i \in I} \rangle$ be a supermodular nice game with a unique Nash equilibrium. Then there is a unique rationalizable action profile, the Nash equilibrium.*

Proof. By Theorem 11, the set of rationalizable action profiles is an order box $[\underline{a}^*, \bar{a}^*]$ where $\underline{a}^*$ and $\bar{a}^*$ are Nash equilibria. By uniqueness of the Nash equilibrium, $\underline{a}^* = \bar{a}^*$; therefore $[\underline{a}^*, \bar{a}^*]$ is a singleton. ■

For instance, suppose that the (increasing) best reply functions of a two-player supermodular nice game are represented in Figure 5.1. Then, there are three Nash equilibria x, y, and z. It follows from Theorem 11 that the set of rationalizable action profiles is the order box $[x, z]$, which is represented by the thick dashed line.

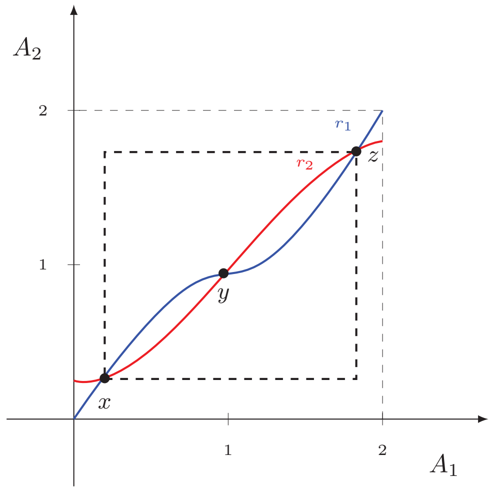

Figure 5.1: A supermodular nice game with $A_1 = A_2 = [0, 2]$.

Finally, we report an additional result which follows from Theorem 11.

Theorem 12. *In every* symmetric *supermodular nice game the set of rationalizable action profiles has a smallest element $\underline{a}^*$ and a largest element $\bar{a}^*$ and they are* symmetric *Nash equilibria.*

Part of the proof of Theorem 12 relies on the following lemma, which shows the symmetry of the equilibrium set in symmetric games.[7]

Lemma 19. *The set of Nash equilibria of a symmetric game is symmetric.*

Proof. If the set of Nash equilibria is empty, then the statement vacuously holds. So, in what follows, we assume that this set is nonempty. Let a^* be a Nash equilibrium of a symmetric game. We must show that, for any permutation $\pi : I \to I$, $a^* \circ \pi$ is also a Nash equilibrium. Let $\hat{A}$ denote the common action set. Since a^* is an equilibrium

$$\forall i \in I, \forall b \in \hat{A},\ u_i\left(a^*\right) \geq u_i\left(b, a^*_{-i}\right).$$

By symmetry

$$\begin{aligned}\forall i \in I, \forall b \in \hat{A},\ u_i\left(a^* \circ \pi\right) &= u_{\pi(i)}\left(\left(a^* \circ \pi\right) \circ \pi^{-1}\right) = u_{\pi(i)}\left(a^*\right) \\ &\geq u_{\pi(i)}\left(b, a^*_{-\pi(i)}\right) = u_i\left(b, \left(a^* \circ \pi\right)_{-i}\right).\end{aligned}$$

Therefore $a^* \circ \pi$ is an equilibrium. ■

Proof of Theorem 12. By Theorem 11, $\rho^\infty(A) = [\underline{a}^*, \bar{a}^*]$, where $\underline{a}^*$ and $\bar{a}^*$ are, respectively, the smallest and largest Nash equilibrium. Lemma 19 shows that the set of Nash equilibria of a symmetric game is symmetric, that is, for every Nash equilibrium $a^* \in \hat{A}^I$—where $\hat{A}$ is the common action set of each player—and every permutation (bijection) $\pi : I \to I$, $a^* \circ \pi$ is also a Nash equilibrium. Now, let $\hat{A} \subseteq \mathbb{R}$ (as is the case in a symmetric nice game), fix $a = (a_i)_{i \in I} \in \hat{A}^I$ arbitrarily, and note that $a \circ \pi \leq a$ for every permutation only if a is symmetric, that is, $a_i = a_j$ for all $i, j \in I$. To see this, pick any two players i and j so that $a_i \leq a_j$ and let π be any permutation such that $j = \pi(i)$; since $a \circ \pi \leq a$, then $a_j = a_{\pi(i)} \leq a_i$. Since $a_i \leq a_j$ and $a_j \leq a_i$, then $a_i = a_j$. By a similar argument, if $a \leq a \circ \pi$ for every permutation π, then a is symmetric. In particular, by symmetry of the set of Nash equilibria, $\underline{a}^* \circ \pi$ is a Nash equilibrium and $\underline{a}^* \leq \underline{a}^* \circ \pi$, $\bar{a}^* \circ \pi$ is a Nash equilibrium and $\bar{a}^* \circ \pi \leq \bar{a}^*$; thus, $\underline{a}^*$ and $\bar{a}^*$ are symmetric. ■

[7] Note that also the set of rationalizable profiles of a symmetric game is symmetric. Since it is a Cartesian product, it must have the form $\hat{C}^I$ for some subset $\hat{C}$ of the common action set $\hat{A}$. We could use the symmetry of the rationalizable set, rather than the symmetry of the equilibrium set, to prove Theorem 12.

5.4 Interpretations of Nash Equilibrium

Nash equilibrium is the most well known and applied equilibrium concept in economic theory, besides the competitive equilibrium. Indeed, we have argued in the Introduction that, in principle, any economic situation (and more generally any social interaction) can be represented as a non-cooperative game. The property according to which every action is a best reply to the other players' actions seems to be essential in order to have an equilibrium in a non-cooperative game.

Nonetheless, we should refrain from accepting this conclusion without further reflection. Why does the Nash equilibrium represent an interesting theoretical concept? When should we expect that the actions simultaneously chosen by the players form an equilibrium? Why should players hold correct conjectures regarding each other's behavior?

We propose a few different interpretations of the Nash equilibrium concept, each addressing the questions above. Such interpretations can be classified in two subgroups: (1) a Nash equilibrium represents "an obvious way to play," (2) a Nash equilibrium represents a stationary (stable) state of an adaptive process. In some cases, we will also introduce a corresponding generalization of the equilibrium concept which is appropriate under the given interpretation and will be analyzed in Chapter 6.

(1) Equilibrium as an "obvious way to play": Assume *complete information*, i.e., common knowledge of the interactive situation represented by G (recall from Chapter 1, G is a mathematical structure reperesenting the interactive situation; yet, to ease language, in this discussion we also let symbol G denote the situation itself).[8] It could be the case that from such common knowledge and shared assumptions about behavior and beliefs, or from some prior events that occurred before the game, the players could positively conclude that a specific action profile a^* represents an "obvious way to play" G. If a^* represents an obvious way to play, every player i expects that everybody else chooses his action in a^*_{-i}. Moreover, if a^* is to be played by rational players, it must be the case that no player i has incentives to choose an action different from a^*_i: if i had an incentive to choose $a_i \neq a^*_i$, not only a^*_i would not be chosen,

[8]Complete information is sufficient to justify the considerations that follow, but it is not strictly necessary.

but also the other players would have no reason to believe that a^* is the obvious way to play. Hence, a^* must be a Nash equilibrium. Indeed, in his PhD thesis, John Nash motivates his solution concept as follows (Nash [63], p. 23):

> "We proceed by investigating the question: what would be a 'rational' prediction of the behavior to be expected of rational playing the game in question? By using the principles that a rational prediction should be unique, that the players should be able to deduce and make use of it, and that such knowledge on the part of each player of what to expect the others to do should not lead him to act out of conformity with the prediction, one is led to the concept of a solution defined before."

What can make a^* an obvious way to play?

(1.a) Deductive interpretation: The players analyze G as an interactive decision problem and look for a rational solution. If such solution exists and is unique, then it corresponds to an obvious way to play. As an example of solution of a game, consider the case where G has a unique rationalizable profile a^* (this is the case for several models in economic theory, such as Example 17 in Section 4.5). Then, if all players are rational and their conjectures are derived from the assumption that there is common belief in rationality, they will choose exactly the actions in a^*. Moreover, a^* is necessarily the unique Nash equilibrium of the game (Observation 6).

Rationalizability captures deductive strategic thinking in a compelling way, and it justifies Nash equilibrium as an obvious way to play when there is a unique rationalizable outcome. But what about games where Nash equilibrium action profiles are a strict subset of rationalizable action profiles? If we insist in the deductive interpretation, perhaps we should just stick to rationalizability.

(1.b) Non-binding agreement: Suppose that the players are able to communicate before playing the game and that they reach an agreement to choose the actions specified in the action profile a^*. Suppose also that such agreement is not legally binding for the parties, it is simply a "gentlemen agreement" based on players honoring their words. Further, players attach little value to honoring their words: if i believes that a different action yields him a higher utility, he will not choose a_i^*. All players are perfectly aware of this. Therefore, the agreement is credible, or "self-enforcing,"

only if no player has an incentive to deviate from the agreement, i.e., only if a^* is a Nash equilibrium.

Is it really necessary that players agree on a specific action profile? Perhaps they could agree on making the action profile actually played (for instance, the way they coordinate in a "Battle of the Sexes" game) depend on some exogenous random variable, such as the weather. In some cases, this would allow to reach, in expectation, fairer outcomes. We will go back to this point in the next chapter.

(2) Equilibrium as a stationary state of an adaptive process.[9] When introductory economic textbooks explain why in a competitive market the price should reach the equilibrium level that equates demand and supply, they almost inevitably rely on informal dynamic justifications. They argue, for instance, that if there is an excess demand the sellers will realize that they are able to sell their goods for a higher price; conversely, if there is an excess supply the sellers will lower their prices to be able to sell the residual unsold goods. Essentially, such arguments rely more or less explicitly on the existence of a feedback effect that pushes the prices towards the equilibrium level. These arguments cannot be formalized in the standard competitive equilibrium model, where market prices are taken as parametrically given by the economic agents and then determined by the demand=supply conditions. They nonetheless provide an intuitive support to the theory.

Similar arguments can be used to explain why the actions played in a game should eventually reach the equilibrium position. In some sense, the conceptual framework provided by game theory is better suited to formalize this kind of arguments. As explained in the Introduction, in a game model every variable that the analyst tries to explain, i.e., any "endogenous" variable, is directly or indirectly determined by the players' actions, according to precise rules that are part of the game (unlike prices in a competitive market model, where the price-formation mechanism is not specified). Assuming that the given game represents an interactive situation that players face recurrently, one can formulate assumptions regarding how players modify their actions taking into account the

[9] To go deeper on this topic the reader can start with the survey by Battigalli *et al.* [19]. In the already mentioned paper by Milgrom and Roberts [58], the connections between the concept of rationalizability and adaptive processes are explored.

outcomes of previous interactions, thus representing formally the feedback process.

One can distinguish two different types of adaptive dynamics: learning and evolutionary dynamics. Here, we will present only a general and brief description.[10]

(2.a) Learning dynamics. Assume that a given game G is played recurrently and that players are interested in maximizing their current expected payoff (a reason for this may be that they do not value the future or that they believe that their current actions do not affect in any way future payoffs). Players have access to information on previous outcomes. Based on such information, they modify their conjectures about their opponents' behavior in the current period. Let a^* be a Nash equilibrium. If in period t every player i expects his opponents to choose a^*_{-i}, then a^*_i is one of his best replies. It is therefore possible that i chooses a^*_i. If this happens, what players observe at the end of period t will confirm their conjectures, which then will remain unchanged for the following period. So even a small inertia (that is a preference, *ceteris paribus*, to repeat the previous action) will induce players to repeat in period $t+1$ the previous actions a^*. Analogously, a^* will be played also in period $t+2$ and so on. Hence, the equilibrium a^* is a stationary state of the process.

We just argued that every Nash equilibrium is a stationary state of plausible learning processes (we presented the argument in an informal way, but a formalization is possible). (i) Do such processes always converge to a steady state? (ii) Is it true that, for any plausible learning process, every stationary state is a Nash equilibrium? It is not possible, in general, to give affirmative answers. First, one has to be more specific about the dynamics of the recurrent interaction. For instance, one has to specify exactly what players are able to observe about the outcomes of previous interactions: is it the action profile of the other players? Is it some variables that depend on such actions? Is it only their own payoffs? Another relevant issue is whether game G is played always by the same agents, or in every period agents are randomly matched with "strangers." An exact specification of these assumptions shows that the process does not always converge to a stationary state. Moreover, as we will see in the next section, there may be stationary states that do not satisfy the Nash equilibrium condition

[10] In Chapter 7, we present an elementary analysis of learning dynamics relating their long-run behavior to equilibrium concepts and rationalizability.

of Definition 18. There are two reasons for this: (a) Players' conjectures can be confirmed by observed outcomes even if they are not correct. (b) If players are randomly drawn from large populations, then the state variable of the dynamic process is given by the fractions of agents in the population that choose each action. But then it is possible that in a stationary state two different agents, playing in the same role, choose different actions. Even though no agent actually randomizes, such situations look like "mixed equilibria," in which players choose randomly among some of the best replies to their probabilistic conjectures, and such conjectures happen to be correct. We analyze notions of equilibrium corresponding to situations (a) and (b) in the next chapter.

(2.b) Evolutionary dynamics. In the analysis of adaptive processes in games, the analogy with evolutionary biology has often been exploited. Consider, for instance, a symmetric two-person game. In every period two agents are drawn from a large population. Then they meet, they interact, they obtain some payoff and finally they split. Individuals are compared to animals, or plants, whose behavior is determined by a genetic code transmitted to their offsprings. If action a is more successful than b, then the agents programmed to play a reproduce themselves faster than those programmed to play b, and therefore the ratio between the fraction of agents playing a and the fraction of agents playing b increases.

In evolutionary biology this theoretical approach based on game theory has been highly successful and has allowed to explain phenomena that appeared as paradoxical within a more naive evolutionary framework.[11] In the social sciences, evolutionary dynamics are used as a metaphor to represent, at an aggregate level, learning phenomena such as imitation, in which agents modify the way they play not on the basis of their direct personal experiences alone, but also observing the behavior of others in the same circumstances. Behaviors that turn out to be more successful are imitated and spread more rapidly.

As for learning dynamics, although Nash equilibria represent stationary states of evolutionary dynamics, the process need not always converge to a stationary state. Furthermore, there may be stationary states that do not satisfy the best reply property of Definition 18. Indeed, in this context a state is represented by the fractions in the population that use different actions. It may be the case that in a stationary state distinct agents of the

[11] See the monograph by Maynard Smith [57].

same population do different things.[12]

These considerations motivate the definition of more general equilibrium concepts whereby players may hold probabilistic, rather than deterministic conjectures about the opponents' behavior. For this reason we refer to these more general concepts—analyzed in Chapter 6—as "probabilistic equilibria."

[12] See the monographs by Weibull [87] and Sandholm [75].

6

Probabilistic Equilibria

So far, we considered mixed actions only with reference to dominance relations: if a (pure) action is dominated by a mixed action, then a rational player should not choose it. We already noticed (Lemma 1), though, that an expected utility maximizer has no strict incentive to choose a mixed action. If there were a small cost for tossing a coin or spinning a roulette wheel, then no expected utility maximizer would choose a mixed action.

We adopt the point of view that (rational) players do not actually choose mixed actions. Nevertheless, it is possible to reconcile this point of view with an interpretation of mixed actions that makes them relevant independently of the equivalence results stated in Lemma 2 and Theorem 5. Specifically, in Section 6.1 we adopt the so-called "mass-action interpretation" of mixed actions put forward by John Nash in his thesis (Nash [63], Weibull [88]): the individuals playing a given game G are drawn at random from large populations of agents and matched to play G; mixed actions represent *statistical distributions* of actions in populations; random matching implies that the share of population i playing a particular action a_j determines the probability that a_j is played in G; in a "mixed (Nash) equilibrium," any other player i correctly estimates this probability. In Section 6.2, we go back to the interpretation of an equilibrium as a non-binding self-enforcing agreement and notice that players may want to agree on probabilistic decision rules that link their behavior to the realization of extraneous random variables; such probabilistic agreements typically induce a spurious correlation between the actions of different players and are thus called "correlated equilibria." Finally, in Section 6.3

we interpret equilibria as the steady states of adaptive process whereby the game is played recurrently and players learn from their personal experience; this requires a specification of players' information feedback; in a steady state, called "self-confirming equilibrium," each agent holds a probabilistic conjecture confirmed by (i.e., consistent with) the feedback he receives and chooses a best reply; whether a steady state is a (mixed) Nash equilibrium depends on the properties of feedback.

6.1 Mixed Equilibrium

In 6.1.1, we elaborate on the mass-action interpretation of mixed equilibrium by means of an illustrative example and an informal analysis of population dynamics. In 6.1.2, we define the mixed eqilibrium concept and we provide existence and characterization results. In 6.1.3, we rely on the characterization of mixed equilibrium to offer a method of computation. In 6.1.4, we focus on an important special case: two-person zero-sum games. Finally, in 6.1.5, we analyze a more restrictive equilibrium concept, called "(trembling hand) perfect equilibrium," that—*inter alia*—eliminates mixed equilibria featuring weakly dominated (mixed) actions.

6.1.1 Motivation: population games

Consider the following game between home *owners* and *thieves*. Each owner has an apartment where he stores goods worth a total value of V. Owners have the option to keep an alarm system, which costs $c < V$.[1] The alarm system is not detectable by thieves. Each thief can decide whether to attempt a theft (burglary) or not. If there is no alarm system, the thieves successfully seize the goods and resell them to some dealer, making a profit of $V/2$. If there is an alarm system, the attempted burglary is detected and the police is automatically alerted. The thieves in such event need to leave all goods in place and try to escape. The probability they get caught is $\frac{1}{2}$ and in this case they are sanctioned with a monetary fine P and then released.[2]

[1] We assume for simplicity that c represents the cost of installing an alarm system as well as the cost of keeping it active. Furthermore, we neglect the possibility to insure against theft.

[2] Prisons do not exist. If thieves cannot pay the amount P, they are inflicted an equivalent punishment.

If the situation just described were a simple game between one owner and one thief, it could be represented using the matrix below:

O\T	*Burglary*	*No*
Alarm	$V - c, -\frac{P}{2}$	$V - c, 0$
No	$0, \frac{V}{2}$	$V, 0$

Figure 6.1: Matrix 2

It is easy to check that Matrix 2 has no equilibria according to Definition 18.

However, the game between owners and thieves is more complex. There are two large populations, and we assume for simplicity that they have the same size: the population of owners (each one with an apartment) and the population of thieves. Thieves randomly distribute themselves across apartments. For any given owner, the probability of an attempted burglary is equal to the fraction of thieves that decide to attempt a burglary. From the thieves' perspective, the probability that an apartment has an alarm system is given by the overall fraction of owners keeping an alarm system.

Assume that the game is played recurrently. The fractions of agents that choose the different actions evolve according to some adaptive process with the following features. At the end of each period it is possible to access (reading them on the newspapers) the statistics of the numbers of successful and unsuccessful burglaries. Players are fairly inert in that they tend to replicate the actions chosen in the previous period. However, they occasionally decide to revise their choices on the basis of the previous period statistics. Since in each period only a few agents revise their choices, such statistics change slowly.

An owner not equipped with an alarm system decides to install it if and only if the expected benefit is larger than the cost, i.e., if the proportion of attempted burglaries is larger than c/V. Conversely, an owner equipped with an alarm system will decide to get rid of it if and only if the proportion of attempted burglaries is lower than c/V. The proportion of attempted burglaries changes only slowly and the owners equate the probability of being robbed today with the one of the previous period. The probability that makes an owner indifferent between his two actions is c/V. When indifferent, an owner sticks to the previous period action.

Analogously, a thief that was active in the previous period decides not to attempt a burglary in the current period if and only if the fraction of apartments equipped with an alarm system (which is also the fraction of unsuccessful burglaries) is larger than $V/(V+P)$. A thief that was not active in the previous period attempts a burglary in the current one if and only if the fraction is lower than $V/(V+P)$.

The state variables of this process are the fraction α of installed alarm systems and the fraction β of attempted burglaries. It is not hard to see that α grows (resp., decreases) if and only if $\beta > \frac{c}{V}$ (resp. $\beta < c/V$). Similarly, β grows (resp., decreases) if and only if $\alpha < V/(V+P)$ (resp., $\alpha > V/(V+P)$). If $\alpha = V/(V+P)$ and $\beta = c/V$, then the state of the system does not change, i.e., $(\alpha, \beta) = (V/(V+P), c/V)$ is a stationary state, or rest point, of the dynamic process. Whether this rest point is stable depends on details that we have not specified. The mixed action pair $(\alpha_1, \alpha_2) = (V/(V+P), c/V)$ is said to be a *mixed* (Nash) *equilibrium* of the matrix game above.

In this example, we have interpreted the mixed action of j, say α_j, both as a *statistical distribution* of actions in population j and as a *conjecture* of the agents of population i about the action of the opponent. In a stable environment every conjecture about j, α_j, is *correct* in the sense that it corresponds to the statistical distribution of actions in population j. Furthermore, α_j is such that every agent in population i is *indifferent* among the actions that are chosen by a positive fraction of agents.

6.1.2 Existence and characterization of mixed equilibria

Next, we present a general definition of mixed equilibrium and then show that it is characterized by the aforementioned properties. To simplify the analysis, we restrict our attention to *finite* games, but the following concepts and results can be extended to compact-continuous games.

Definition 20. *The **mixed extension** of a finite game $G = \langle I, (A_i, u_i)_{i \in I} \rangle$ is a game $\overline{G} = \langle I, (\Delta(A_i), \overline{u}_i)_{i \in I} \rangle$ where*

$$\overline{u}_i(\alpha) = \sum_{a=(a_j)_{j \in I} \in A} u_i(a) \prod_{j \in I} \alpha_j(a_j),$$

for all i and $\alpha = (\alpha_i)_{i \in I} \in \times_{i \in I} \Delta(A_i)$.

Note, the payoff functions of the mixed extension are obtained calculating the expected payoff corresponding to a vector of mixed actions under the assumption that the actions of different players are *statistically independent.*

Definition 21. *(Nash [62]) Fix a finite game G. A mixed action profile α is a **mixed equilibrium** of game G if α is a Nash equilibrium of the mixed extension of G.*

Note that every (pure) Nash equilibrium $(a_i^*)_{i \in I}$ corresponds to a mixed equilibrium $(\alpha_i^*)_{i \in I}$ such that each α_i^* assigns probability one to a_i^*: mixed equilibrium is a generalization of (pure) Nash equilibrium.

Theorem 13. *Every finite game G has at least one mixed equilibrium.*[3]

Proof. It is easy to verify that the mixed extension $\overline{G}$ of a finite game G satisfies all the assumptions of Theorem 8: for each i, $\Delta(A_i)$ is a nonempty, convex and compact subset of $\mathbb{R}^{A_i}$, $\overline{u}_i$ is a multi-linear function and therefore it is both continuous in α and weakly concave, indeed linear,[4] in α_i. ■

When we introduced mixed actions in Chapter 3, we said that, even if we do not assume that players randomize, mixed actions still play an important technical role. The space of mixed actions is convex and expected payoff is multi-linear in the mixed actions of different players. In Chapter 3, such convexity-linearity structure was exploited to characterize the set of justifiable (pure) actions (Lemma 2). In Theorem 13, the convexification-linearization allowed by the mixed extension of a game yields the existence of an equilibrium.

The following result introduces an alternative characterization of mixed equilibria, which is not based on the mixed extension of the game. With a small abuse of notation, we let $r_i(\alpha_{-i})$ denote the set of best replies to the

---

[3] The same holds for any compact-continuous game.

[4] More precisely, for each α_{-i}, the section of $\bar{u}_i$ at α_{-i}, $\bar{u}_{i,\alpha_{-i}} : \Delta(A_i) \to \mathbb{R}$, is *affine*: for all $\alpha_i, \beta_i \in \Delta(A_i)$ and $\lambda \in [0,1]$,

$$\bar{u}_{i,\alpha_{-i}}(\lambda\alpha_i + (1-\lambda)\beta_i) = \lambda\bar{u}_{i,\alpha_{-i}}(\alpha_i) + (1-\lambda)\bar{u}_{i,\alpha_{-i}}(\beta_i).$$

conjecture $\mu^{\alpha_{-i}} \in \Delta(A_{-i})$ obtained as a product of the marginal measures α_j $(j \neq i)$, that is, $r_i(\alpha_{-i}) = r_i(\mu^{\alpha_{-i}})$ where

$$\forall \alpha_{-i} \in \bigtimes_{j \neq i} \Delta(A_j), \forall a_{-i} \in A_{-i},\ \mu^{\alpha_{-i}}(a_{-i}) = \prod_{j \neq i} \alpha_j(a_j).$$

Theorem 14. *A mixed action profile $(\alpha_i)_{i \in I}$ is a mixed equilibrium if and only if, for every $i \in I$, $\text{supp}\alpha_i \subseteq r_i(\alpha_{-i})$ (in the finite case, every action played with positive probability is a best reply to the mixed action profile played by the opponents).*

Proof. The statement follows directly from the definition of mixed equilibrium and from Lemma 1. ■

Corollary 4. *All actions played with positive probability in a mixed equilibrium are rationalizable.*

Proof. By Theorem 14, the set $C = \bigtimes_{i \in I} \text{supp}\alpha_i$ has the best reply property. Hence, Theorem 3 implies that every action played with positive probability is rationalizable. ■

The reader should verify by example that the converse of Corollary 4 does not hold: there are (finite) games with rationalizable actions that are not played with positive probability in any mixed equilibrium.

6.1.3 Computation of mixed equilibria

Theorem 13 and Corollary 4 provide a linear programming algorithm to compute all mixed equilibria of finite two-person games. In such games the payoff function of player i can be represented by a matrix with generic entry $u_i^{k\ell}$. To ease notation, let $I = \{1, 2\}$, and let player 1 choose the rows (indexed by k) and player 2 the columns (indexed by ℓ).

Step 1: Eliminate all iteratively dominated actions (by Theorem 5 and Corollary 4 such actions are played with zero probability in equilibrium). The order of elimination is irrelevant (Theorem 6).

Step 2: Let $G = \langle A_1, A_2, u_1, u_2 \rangle$ denote the residual game obtained after Step 1, i.e., the restriction to the set of rationalizable action pairs. For any pair of nonempty subsets $A_1^* \subseteq A_1$ and $A_2^* \subseteq A_2$, compute the set of mixed equilibria, (α_1, α_2), such that $\text{supp}\alpha_1 = A_1^*$ and $\text{supp}\alpha_2 = A_2^*$. This set, which could be empty, is computed as follows (we consider the non trivial case in which each set contains at least two actions). To ease notation, relable actions so that $A_1^* = \{1, ..., K\}$ and $A_2^* = \{1, ..., L\}$, and denote by α_1^k (respectively, α_2^ℓ) the probability of action k (resp. ℓ) of player 1 (resp. 2). Solve the following systems of linear equations and inequalities with unknowns $\alpha_1 = (\alpha_1^1, ..., \alpha_1^K) \in \Delta(A_1^*)$ and $\alpha_2 = (\alpha_2^1, ..., \alpha_2^L) \in \Delta(A_2^*)$:

$$\begin{aligned}\sum_{k=1}^{K} u_2^{k\ell}\alpha_1^k &= \sum_{k=1}^{K} u_2^{k1}\alpha_1^k,\ \ell = 2, ..., L, \\ \sum_{k=1}^{K} u_2^{k\ell}\alpha_1^k &\leq \sum_{k=1}^{K} u_2^{k1}\alpha_1^k,\ \ell = L+1, ..., |A_2|;\end{aligned} \tag{6.1.1}$$

$$\begin{aligned}\sum_{\ell=1}^{L} u_1^{k\ell}\alpha_2^\ell &= \sum_{\ell=1}^{L} u_1^{1\ell}\alpha_2^\ell,\ k = 2, ..., K, \\ \sum_{\ell=1}^{L} u_1^{k\ell}\alpha_2^\ell &\leq \sum_{\ell=1}^{L} u_1^{1\ell}\alpha_2^\ell,\ k = K+1, ..., |A_1|.\end{aligned} \tag{6.1.2}$$

The equations in (6.1.1) determine the set of mixed actions of player 1 that make player 2 indifferent between the actions in subset A_2^*. The inequalities determine the set of mixed actions of player 1 that make $a_2 = 1$ (and so *all the actions in* A_2^*) weakly preferred to the actions that do not belong to A_2^*. For any α_1 that satisfies (6.1.1), player 2 has no incentive to "deviate" from any mixed action with support A_2^*. Similar considerations hold for system (6.1.2). Such system determines the set of mixed actions of player 2 that make player 1 indifferent among all the actions in A_1^* and at the same time make such actions weakly preferred to all the others. Therefore, *the indifference conditions for player 1 determine the equilibrium randomization(s) of player 2, the indifference conditions for player 2 determine the equilibrium randomization(s) of player 1.*

In the intoductory example (Owners and Thieves) the equilibrium is determined as follows: First, note that the best reply to a deterministic conjecture is unique. However, no pair of pure actions is an equilibrium. Hence, the equilibrium is necessarily mixed (existence follows from Theorem 13). An equilibrium with support $A_1^* = \{Alarm, No\}$, $A_2^* =$

$\{Burglary, No\}$ is given by the following (we are writing $\alpha = \alpha_1(A)$ and $\beta = \alpha_2(B)$):

Indifference condition for $i = 1$ (Owner):

$$V - c = V(1 - \beta).$$

Indifference condition for $i = 2$ (Thief):

$$\frac{V}{2}(1 - \alpha) - \frac{P}{2}\alpha = 0.$$

Solving the system, $\alpha = \frac{V}{V+P}$, $\beta = \frac{c}{V}$, which is the stationary state identified by the (informal) analysis of a plausible dynamic process.

6.1.4 Zero-sum games and maxmin solution

Game theory was initially focused on games with monetary payoffs and two risk-neutral players such that the sum of their payoffs is *constant*. Given the assumption that players maximize their expected monetary payoffs, without loss of generality, such constant may be assumed to be 0, that is, $u_i = -u_{-i}$. Matching-Pennies and Rock-Scissors-Paper are well known examples of games that children play and can be assumed to have this zero-sum structure (this includes assumptions about personal preferences, which for such game forms seem quite plausible). Since most applications of game theory also feature a surplus-creation component, besides a surplus-distribution component, in modern game theory the analysis of zero-sum games is not as emphasized as at the birth of this discipline. Yet, two-person zero-sum games have properties that turn out to be important for both technical and conceptual reasons. Specifically, the set of mixed equilibria has special and noteworthy structure. To ease notation, given finite game $G = \langle A_1, A_2, u_1, u_2 \rangle$ with $u_1 = -u_2$, let $X := \Delta(A_1)$, $Y := \Delta(A_2)$, and let $V : X \times Y \to \mathbb{R}$ denote the payoff function of player 1 in the mixed extension of the game, that is,

$$\forall (x, y) \in X \times Y,\ V(x, y) := \sum_{a_1 \in A_1} \sum_{a_2 \in A_2} x(a_1) y(a_2) u_1(a_1, a_2).$$

Theorem 15. *(*Maxmin*) Fix a finite two-person, zero-sum game. Each player has a unique mixed equilibrium payoff; the equilibrium payoff of* *player 1, called "value" of the game and denoted V^*, satisfies the following equalities:*

$$\max_{x \in X} \min_{y \in Y} V(x, y) = V^* = \min_{y \in Y} \max_{x \in X} V(x, y).$$

*(*Exchangeability*) For any two equilibria (x', y') and (x'', y''), the mixed action pairs (x', y'') and (x'', y') are equilibria as well.*

This was one of the first important results in game theory and had its own independent proof.[5] However, one can take advantage of Theorem 13 to posit the existence of a mixed equilibrium and use the equilibrium inequalities and the zero-sum property to obtain an easier proof, which we leave as an exercise. An interpretation of this result is that, in any mixed equilibrium of a two-person, zero-sum game (where all equilibria are equivalent and exchangeable), each i plays a mixed action *"as if"* he were certain that such randomization (not its realization) were observed by the opponent $-i$, who would react by maximizing his own expected payoff, hence, minimizing i's expected payoff.[6] Thus, each player i solves a *maxmin* problem; since $u_{-i} = -u_i$, player $-i$, by solving his maxmin problem, also solves the problem of "minimaximizing" the payoff of i. The noteworthy feature of two-person, zero-sum (or constant-sum) games is that every pair of maxmin solutions, one for each player, necessarily forms an equilibrium pair. The maxmin criterion appears in several parts of game theory, such as the so-called "folk theorems" for repeated games[7] and the proof of the Wald-Pearce Lemma;[8] furthermore, it plays an independent and important role both in statistics and decision theory.[9]

6.1.5 Perfect mixed equilibrium

As elaborated throughout this chapter, we adopted the so-called "mass-action interpretation" of mixed equilibrium: for each role $i \in I$ in game $G = \langle I, (A_i, u_i)_{i \in I} \rangle$ there is a large population of agents; many games like G are played, and, in each such game, an agent playing in role i is drawn at random from population i. The state of the social system is described by a profile of statistical distributions $\alpha = (\alpha_j)_{j \in I} \in \times_{j \in I} \Delta(A_j)$. With

[5] See von Neumann and Morgenstern [84].
[6] Again, see von Neumann and Morgenstern [84].
[7] See, e.g., Chapter 8 in Osborne and Rubinstein [65].
[8] See Pearce [66] and the Appendix of Chapter 3.
[9] See, e.g., Cerreia-Vioglio *et al.* [35] and the relevant references therein.

this, if recent statistics α are observed, it is reasonable to assume that the conjecture of an agent in role i for the current period is the product measure $\mu^{\alpha_{-i}} = \prod_{j \neq i} \alpha_j \in \Delta(A_{-i})$, where independence follows from random matching. If rational, an agent in role i with such conjecture takes an action $a_i \in r_i(\alpha_{-i}) = r_i(\mu^{\alpha_{-i}})$.[10] If there is a positive fraction of agents in at least one population that do not play justifiable actions, then the state of the system is going to change and the fractions of agents playing such actions—as well as others that are not best replies to the state of the previous period—decrease, according to dynamics informally described in Section 6.1.1. The system is at rest only if[11]

$$\forall i \in I, \, \text{supp}\alpha_i \subseteq r_i(\alpha_{-i}),$$

which is—by Theorem 14—a necessary and sufficient condition for a mixed equilibrium. However, in the dynamic process, the fractions of agents playing actions that are not best replies to the limit frequencies vanish only in the limit.

Starting from these considerations we introduce a refined notion of mixed equilibrium. Suppose that the system starts from an *interior* point, that is, a situation where each action of each player/role i is used by a positive fraction of agents in population i; moreover, the system moves slowly, because only a small fraction of agents can change action in a short period of time. Then, the fractions of agents using "under-performing actions" vanishes only in the limit. Thus, at each time t, agents in each population i have strictly positive, statistically independent conjectures $\prod_{j \neq i} \alpha_j^t \in \Delta^\circ(A_{-i})$ (of course, statistical independence is relevant if there are at least two co-players). Suppose that $\alpha^t \to \alpha^*$; by finiteness of G and continuity of the (polynomial) expected payoff function $\bar{u}_i(\alpha)$, it is the case that $r_i(\alpha_{-i}^t) \subseteq r_i(\alpha_{-i}^*)$ for t large enough (that is, α_{-i}^t close enough to α_{-i}^*).[12] From then on, every action $a_i \in r_i(\alpha_{-i}^*) \backslash r_i(\alpha_{-i}^t)$ is suboptimal and does not survive in the limit. Therefore, $\text{supp}\alpha_i^* \subseteq r_i(\alpha_{-i}^t)$ for t large. This motivates the following requirement concerning mixed profile α^*.

[10]Recall that, with an abuse of notation, we let $r_i(\alpha_{-i})$ denote the set of best replies to the conjecture $\mu^{\alpha_{-i}} \in \Delta(A_{-i})$ obtained as a product of the marginal measures α_j $(j \neq i)$.

[11]Assuming a form of inertia, the condition is also sufficient to have a rest point. See Section (6.1.1).

[12]Formally, this follows from the closed-graph property of best reply correspondences. See Lemma 10.

Condition 1. *There is some sequence* $(\alpha^k)_{k=1}^{\infty}$ *of strictly positive mixed action profiles such that* $\lim_{k\to\infty}\alpha^k = \alpha^*$ *and* $\text{supp}\alpha_i^* \subseteq r_i(\alpha_{-i}^k)$ *for all* i *and* k*.*

Condition 1 yields a refinement of the mixed equilibrium concept. In particular, recall that, by Lemma 4, if an action is a best reply to a strictly positive conjecture,[13] then it is also admissible, i.e., it is not weakly dominated. Equivalently, weakly dominated actions cannot be best replies to strictly positive conjectures. One can show that this must hold also for the *mixed* equilibrium actions satisfying Condition 1, not only for the pure actions in their support (this follows from an application of Lemma 4 to the mixed extension of the given game). Furthermore, we will show below by example that this is only a necessary condition.

Reinhard Selten [76] arrived at the refined equilibrium condition—as per Condition 1—in a different way. He argued that the players can make mistakes, as if they were choosing with a "trembling hand." The "trembling-hand" idea is made formal by assuming that players cannot fully control their actions, because every pure action is selected in a partially random way with an exogenously given, strictly positive minimal probability. Specifically, for each $i \in I$ and $a_i \in A_i$, let $\varepsilon_i(a_i) > 0$ denote this minimal probability. Define a profile of minimal probabilities $\varepsilon = (\varepsilon_i(a_i))_{i\in I, a_i\in A_i} \in \times_{i\in I}\mathbb{R}_{++}^{A_i}$, where $\sum_{a_i\in A_i}\varepsilon_i(a_i) < 1$ for each i. With this, we can define an ε**-perturbed mixed extension** $\bar{G}^{\varepsilon} = \langle I, (\Delta^{\varepsilon}(A_i), \bar{u}_i)_{i\in I}\rangle$ of game G, that is, a game where the action set of each player i is

$$\Delta^{\varepsilon}(A_i) = \{\alpha_i \in \Delta(A_i) : \forall a_i \in A_i, \alpha_i(a_i) \geq \varepsilon_i(a_i)\}$$

(condition $\sum_{a_i\in A_i}\varepsilon_i(a_i) < 1$ ensures that the relative interior of $\Delta^{\varepsilon}(A_i)$ is not empty).

Definition 22. *A mixed action profile* α^* *is a* ***(trembling-hand) perfect equilibrium*** *of* G *if there exist a sequence* $(\varepsilon^k)_{k=1}^{\infty}$ *of profiles of minimal probabilities and a sequence* $(\alpha^k)_{k=1}^{\infty}$ *of equilibria of the corresponding* ε^k*-perturbed mixed extensions* $\bar{G}^{\varepsilon^k}$ *such that* $\varepsilon^k \to \mathbf{0}$ *(where* $\mathbf{0}$ *is the origin of* $\times_{i\in I}\mathbb{R}_{++}^{A_i}$*) and* $\alpha^k \to \alpha^*$*.*

[13] In infinite compact-continuous games, a full-support conjecture.

The “trembling-hand” qualification is useful to distinguish the appropriate extension of this equilibrium concept to sequential games from other notions of perfect equilibrium for such games. Since here we are only analyzing simultaneous-move games, we omit this qualification. The reader should be able to use the compactness of mixed action sets and the continuity of the payoff functions $\bar{u}_i$ $(i \in I)$ to prove the following statement.

Remark 14. *Every perfect equilibrium of G is also a mixed equilibrium of G.*

Theorem 16. *Every finite game G has a perfect equilibrium.*

Proof. Take any sequence of profiles of minimal probabilities $\left(\varepsilon^k\right)_{k=1}^{\infty}$ such that $\varepsilon^k \to \mathbf{0}$. Each game $\bar{G}^{\varepsilon^k}$ has a nonempty, compact and convex action set with continuous payoff functions; thus, it has a mixed equilibrium α^k (cf. Theorem 13). The sequence $\left(\alpha^k\right)_{k=1}^{\infty}$ is included in the compact set $\times_{i\in I} \Delta(A_i)$, hence, it has a converging subsequence $(\bar{\alpha}^m)_{m=1}^{\infty} = \left(\alpha^{k_m}\right)_{m=1}^{\infty}$. Let α^* be the limit, that is, $\lim_{m\to\infty} \bar{\alpha}^m = \alpha^*$. By definition, α^* is a perfect equilibrium. ■

It turns out that Condition 1 is indeed necessary and sufficient for a mixed equilibrium to be perfect.

Theorem 17. *A mixed action profile α^* is a perfect equilibrium if and only if it satisfies Condition 1.*

Proof. First, suppose that α^* is a perfect equilibrium. Then, by definition, there exist a sequence $\left(\varepsilon^k\right)_{k=1}^{\infty}$ of profiles of minimal probabilities and a sequence $\left(\alpha^k\right)_{k=1}^{\infty}$ of equilibria of the ε^k-perturbed mixed extensions $\bar{G}^{\varepsilon^k}$ such that $\varepsilon^k \to \mathbf{0}$ and $\alpha^k \to \alpha^*$. Consider any player i and action a_i such that, for some subsequence $\left(\alpha_i^{k_m}\right)_{m=1}^{\infty}$ of $\left(\alpha_i^k\right)_{k=1}^{\infty}$, $\alpha_i^{k_m}(a_i) = \varepsilon^{k_m}(a_i)$ for all m. Since $\lim_{m\to\infty} \alpha_i^{k_m} = \alpha_i^*$ and $\lim_{m\to\infty} \varepsilon^{k_m}(a_i) = 0$, it follows that $\alpha_i^*(a_i) = 0$, that is, $a_i \notin \operatorname{supp}\alpha_i^*$ for any such player and action. Thus, there must be K large enough such that

$$\forall k \geq K, \forall i \in I, \forall a_i \in \operatorname{supp}\alpha_i^*, \alpha_i^k(a_i) > \varepsilon^k(a_i).$$

Then, since α_i^k is a best reply to α_{-i}^k in the ε^k-perturbed mixed extension $\bar{G}^{\varepsilon^k}$, finiteness of G and continuity of the expected payoff functions imply that

$$\forall k \geq K, \forall i \in I, \text{supp}\alpha_i^* \subseteq r_i\left(\alpha_{-i}^k\right),$$

that is, Condition 1 holds.

Suppose now that Condition 1 holds for mixed action profile α^*. Let $(\alpha^k)_{k=1}^\infty$ be a sequence as per Condition 1. For each i and each $a_i' \in A_i \backslash \text{supp}\alpha_i^*$, let $\varepsilon^k(a_i') = \alpha_i^k(a_i')$. Moreover, for each $a_i \in \text{supp}\alpha_i^*$, let $\varepsilon^k(a_i) = \min_{a_i' \in A_i \backslash \text{supp}\alpha_i^*} \alpha_i^k(a_i')$. Since $\alpha^k \to \alpha^*$, it must be the case that $\varepsilon^k \to \mathbf{0}$. For k large enough, this choice of ε^k determines an ε^k-perturbed mixed extension $\bar{G}^{\varepsilon^k}$ of G. Hence, α^k is a mixed equilibrium of $\bar{G}^{\varepsilon^k}$. ■

Perfect equilibrium can also be characterized in terms of admissible (i.e., not weakly dominated) mixed actions. In particular:

Theorem 18. *Fix a finite game $G = \langle I, (A_i, u_i)_{i\in I}\rangle$ and a mixed action profile $\alpha^* = (\alpha_i^*)_{i\in I}$.*

(i) If α^* is a perfect equilibrium, then α^* is a mixed equilibrium such that α_i^* is admissible for every $i \in I$.

(ii) Let $|I| = 2$. If α^* is a mixed equilibrium such that α_i^* is admissible for every $i \in I$, then α^* is a perfect equilibrium.

Proof. Let $\alpha^* = (\alpha_i^*)_{i\in I}$ be a perfect equilibrium. By Remark 14, α^* is a mixed equilibrium. Fix a player $i \in I$. By Theorem 17, α_i^* is a mixed best reply to each element of a sequence $\left(\alpha_{-i}^k\right)_{k=1}^\infty$ of strictly positive mixed action profiles. Lemma 4 (which applies to the mixed extension of G, hence also to mixed actions) implies that α_i^* is admissible, that is, not weakly dominated.

To prove the second statement, suppose that G is a *two-person* game, and fix a mixed equilibrium $\alpha^* = (\alpha_i^*)_{i\in I}$ in admissible actions. By Lemma 5, for each $i \in I$ there exists $\hat{\alpha}_{-i} \in \Delta^\circ(A_{-i})$ such that α_i^* is a best reply to $\hat{\alpha}_{-i}$. Since $|I| = 2$, $\hat{\alpha}_{-i}$ trivially satisfies statistical independence, because $-i$ is the single co-player of i. Consider the sequence of strictly positive mixed action pairs $(\alpha^k)_{k=1}^\infty = \left(\frac{k-1}{k}\alpha_i^* + \frac{1}{k}\hat{\alpha}_i, \frac{k-1}{k}\alpha_{-i}^* + \frac{1}{k}\hat{\alpha}_{-i}\right)_{k=1}^\infty$. Clearly, $\alpha^k \to \alpha^*$. Furthermore, by affinity of u_i in the co-player's mixed action, since α_i^* is a best reply to both α_{-i}^* and $\hat{\alpha}_{-i}$, it is also a best reply to the convex combination $\alpha^k_{-i} = \frac{k-1}{k}\alpha^*_{-i} + \frac{1}{k}\hat{\alpha}_{-i}$ for each k. To see this, note that by assumption, for all $\alpha_i \in \Delta(A_i)$,

$$\begin{aligned} u_i(\alpha^*_i, \alpha^*_{-i}) &\geq u_i(\alpha_i, \alpha^*_{-i}), \\ u_i(\alpha^*_i, \hat{\alpha}_{-i}) &\geq u_i(\alpha_i, \hat{\alpha}_{-i}). \end{aligned}$$

Hence, for all $\alpha_i \in \Delta(A_i)$,

$$\begin{aligned} u_i(\alpha^*_i, \alpha^k_{-i}) &= \sum_{a_{-i}\in A_{-i}} u_i(\alpha^*_i, a_{-i})\alpha^k_{-i}(a_{-i}) \\ &= \frac{k-1}{k} \sum_{a_{-i}\in A_{-i}} u_i(\alpha^*_i, a_{-i})\alpha^*_{-i}(a_{-i}) + \frac{1}{k} \sum_{a_{-i}\in A_{-i}} u_i(\alpha_i, a_{-i})\hat{\alpha}_{-i}(a_{-i}) \\ &\geq \frac{k-1}{k} \sum_{a_{-i}\in A_{-i}} u_i(\alpha_i, a_{-i})\alpha^*_{-i}(a_{-i}) + \frac{1}{k} \sum_{a_{-i}\in A_{-i}} u_i(\alpha_i, a_{-i})\hat{\alpha}_{-i}(a_{-i}) \\ &= u_i(\alpha_i, \alpha^k_{-i}). \end{aligned}$$

Thus, by Theorem 17, α^* is a perfect equilibrium. ■

Notice that the admissibility condition in the statements of Theorem 18 refers to the equilibrium mixed actions α^*_i $(i \in I)$ and not to the pure actions in their supports. If a player has at least 3 pure actions, a mixed action can be weakly dominated even when none of the actions in its support are.[14] The following example shows that a mixed equilibrium of a two-person 3×3-game that assigns zero probability to all weakly dominated pure actions need not be perfect.

Example 19. Consider the following two-player 3×3-game:

	ℓ	c	r
t	0, 1	6, 1	0, 0
m	6, 1	0, 1	3, 0
b	3, 1	3, 1	2, 0

In this game, $\left(\frac{1}{2}\delta_t + \frac{1}{2}\delta_m, \frac{1}{2}\delta_\ell + \frac{1}{2}\delta_c\right)$ is a mixed equilibrium that assigns probability 0 to the only weakly (and strictly) dominated action, r; yet,

[14] Prove as an exercise that if a player has 2 pure actions, then a mixed action is weakly dominated if and only if its support contains a weakly dominated pure action. In this case, these mixed and pure actions are both dominated by the other pure action.

it is not perfect because the mixed action $\frac{1}{2}\delta_t + \frac{1}{2}\delta_m$ is weakly dominated by b. It can be checked that (t, c), (m, ℓ) and $\left(b, \frac{1}{2}\delta_\ell + \frac{1}{2}\delta_c\right)$ are all perfect equilibria. ▲

Assumption $|I| = 2$ in Theorem 18 is tight, that is, admissibility of mixed equilibrium actions is only a necessary, but not sufficient condition for perfectness in games with at least three players. There are two reasons for this that can be understood by focusing on 3-person games.

First, any two players i and j see the "trembles" of the third player k in the same way, that is, they attach the *same* minimal probabilities $\varepsilon_k(a_k)$ to the pure actions that k would rather not choose, as illustrated in the following:

Example 20. In the game below, Rowena (player 1) chooses the row, Colin (player 2) the column, and Mary (player 3) the matrix (L, M, or R).

L	r	s
r	3, 0, 0	2, 2, 0
s	2, 2, 0	2, 2, 0

M	r	s
r	2, 2	2, 2
s	2, 2	2, 2

R	r	s
r	0, 3, 0	2, 2, 0
s	2, 2, 0	2, 2, 0

Rowena and Colin choose between a risky action r and safe action s, while the middle action M is dominant for Mary. If Mary selects her dominant action with probability 1, Rowena and Colin are indifferent. Also, if Colin selects his safe action with probability 1, Rowena is indifferent, and the same holds with their roles reversed. Thus, their preferred action depends on how they think that Mary "trembles." Profile (r, r, M) is a (pure) equilibrium in admissible actions, but it is not perfect. To see this, let $\varepsilon_3(L) = \lambda$ and $\varepsilon_3(R) = \rho$. Similarly, let σ_1 (respectively, σ_2) denote the minimal probability of the safe action of Rowena (respectively, Colin). Perfectness holds (if and) only if the risky actions of Rowena and Colin are optimal when these minimal probabilities are very small, that is, if the following inequalities hold for small values of λ, ρ, σ_1 and σ_2:

$$\begin{aligned}(1-\sigma_2)(3\lambda + 2(1-\lambda-\rho)) &\geq 2,\\ (1-\sigma_1)(2(1-\lambda-\rho) + 3\rho) &\geq 2.\end{aligned}$$

It can be checked that this is impossible, because it requires $\lambda > 2\rho$ and $\rho > 2\lambda$. ▲

Second, perfect equilibrium actions must be best replies to strictly positive conjectures that satisfy *independence across co-players*, as illustrated in the following:

Example 21. In the game below, L and ℓ are, respectively, dominant for Colin (player 2, column) and Mary (player 3, matrix). Thus, α^* is a mixed equilibrium if and only if $\alpha^* = ((1-\beta)\delta_T + \beta\delta_B, L, \ell)$ with $\beta \in [0,1]$. Note T is the unique best reply to (R, ℓ) and (L, r), while B is the unique best reply to (R, r). Therefore, none of them is weakly dominated, and this implies that no mixture of T and B is weakly dominated. Thus, such α^* is an equilibrium in admissible (mixed) actions.

ℓ	L	R
T	1, 1, 1	1, 0, 1
B	1, 1, 1	0, 0, 1

r	L	R
T	1, 1, 0	0, 0, 0
B	0, 1, 0	1, 0, 0

Yet, $\alpha^* = ((1-\beta)\delta_T + \beta\delta_B, L, \ell)$ is perfect if and only if $\beta = 0$. To see this, consider any sequence of mixed action pairs

$$(\alpha_2^k, \alpha_3^k)_{k=1}^{\infty} = \left(\left(1-\varepsilon^k\right)\delta_L + \varepsilon^k\delta_R, \left(1-\eta^k\right)\delta_\ell + \eta^k\delta_r\right)_{k=1}^{\infty}$$

with $\varepsilon^k, \eta^k \to 0$ as $k \to \infty$. The expected payoffs of actions T and B are

$$\begin{aligned} u_1(T, \alpha_2^k, \alpha_3^k) &= \left(1-\eta^k\right) + \eta^k\left(1-\varepsilon^k\right), \\ u_1(B, \alpha_2^k, \alpha_3^k) &= \left(1-\eta^k\right)\left(1-\varepsilon^k\right) + \eta^k\varepsilon^k. \end{aligned}$$

Clearly, for k large enough,

$$\left(1-\eta^k\right) + \eta^k\left(1-\varepsilon^k\right) > \left(1-\eta^k\right)\left(1-\varepsilon^k\right) + \eta^k\varepsilon^k,$$

making T the unique best reply. ▲

6.2 Correlated Equilibrium

Mixed equilibria can be interpreted as statistical distributions over action profiles that are in some sense stationary. Such distributions are obtained as a product of the marginal distributions corresponding to the equilibrium mixed actions: if $(\alpha_i^*)_{i\in I}$ is a mixed equilibrium, then the probability of any profile $(a_i)_{i\in I}$ is $\prod_{i\in I}\alpha_i^*(a_i)$; this means that actions are statistically independent across players. Statistical independence is justified by the interpretation of mixed equilibrium as a stationary profile of statistical distributions of actions in $|I|$ populations whose agents are randomly matched with the agents of the other populations to play the game G.

Now we present a different concept of probabilistic equilibrium that, formally, generalizes the concept of mixed equilibrium by allowing *correlated* distributions over players' action profiles. The interpretation of this equilibrium concept given here is very different from the one proposed for the mixed equilibrium.[15]

6.2.1 Motivation and examples

Let us consider, as in Section 5.4 of Chapter 5, the case in which players are able to communicate and "sign" a *non-binding* agreement before the game is played. If this agreement simply prescribes to play a certain action profile a^*, then—as noted earlier—a^* must be a Nash equilibrium, otherwise the agreement would not be self-enforcing. However, players could reach more sophisticated self-enforcing agreements, in which the chosen actions depend on the realizations of extraneous random variables that do not directly affect the payoffs. We first illustrate this idea with a simple example.

	B	S
B	$3,1$	$0,0$
S	$0,0$	$1,3$

Figure 6.2: Matrix 3.

Matrix 3 represents the classic *Battle of the Sexes* (BoS): Rowena (row player) and Colin (column player) would like to coordinate and go to the same concert, either Bach or Stravinsky, but Rowena prefers Bach and Colin prefers Stravinsky. Suppose that Rowena and Colin have to agree on how to play the BoS the next day, when no further communication between them will be possible. There exists two simple self-enforcing agreements, (B, B) and (S, S). However, the first favors Rowena, the second favors

[15]Proposition 24 of Chapter 7 offers an adaptive interpretation.

Colin, and neither player wants to give up. How to sort this out? Colin can make the following proposal that would ensure in expected values a fair distributions of the gains from coordination: "If tomorrow's weather is $b$ad then both of us choose Bach, if instead it is $s$unny then we both choose Stravinsky." Notice that the weather forecasts are uncertain: there is a 50% probability of bad weather and 50% of sunny weather. The agreement generates an expected payoff of 2 for both players. Rowena understands that the idea is smart. Indeed, the agreement is self-enforcing as both players have an incentive to respect it if they expect the other to do the same. For instance, if Rowena expects that Colin sticks to the agreement and waking up she observes that the weather is $b$ad, then she expects that Colin will go to the Bach concert and she wants to play B. Similarly, if she observes that the weather is $s$unny, she expects that Colin will go to the Stravinsky concert and she wants to play S.

In other words, a sophisticated agreement can use exogenous and not directly relevant random variables to coordinate players' beliefs and behavior. In such an agreement the conjecture of a player about the behavior of others (and so their best replies) depends on the observed realization of such random variables and the actions of different players are correlated, albeit spuriously, i.e., without any direct or indirect causal influence from the choice of a player to the choice of some other player.

Clearly, it is not always possible to condition the choice on some commonly observable random variable. Furthermore, even if this were possible, players could still find it more convenient to relate their respective choices to random variables that are only partially correlated. Consider, for example, the following bi-matrix game:

	a	b
a	$6,6$	$2,7$
b	$7,2$	$0,0$

Figure 6.3: Matrix 4.

If Rowena and Colin rely on a commonly observed random variable to design a self-enforcing agreement, they can only achieve a probability distribution over pure Nash equilibria. Actually this is true in every game: suppose that the agreement says "if x is (commonly) observed, each $j \in I$

must take action $\sigma_j(x)$," then when i observes x he infers that each co-player j will take action $\sigma_j(x)$, and—in order to make the agreement self-enforcing—$\sigma_i(x)$ must be a best reply to $\sigma_{-i}(x)$ for each $i \in I$, thus $(\sigma_i(x))_{i\in I}$ must be a Nash equilibrium.[16] The two pure equilibria of the game in Matrix 4 yield the payoff pairs $(7, 2)$ and $(2, 7)$. So, the sum of the expected payoffs attainable with self-enforcing agreements that rely on a commonly observed random variable is $\mu \times (2+7) + (1-\mu) \times (7+2) = 9$, where μ is the probability of observing a realization x that yields (a, b) $(\mu = \mathbb{P}(\{x : \sigma_1(x) = a, \sigma_2(x) = b\}))$.

If instead Rowena and Colin rely on different (but correlated) random variables, they can do better. For example, suppose that Rowena observes whether $X = \mathbf{a}$ or $X = \mathbf{b}$, and Colin observes whether $Y = \mathbf{a}$ or $Y = \mathbf{b}$, where (X, Y) is a pair of random variables with the following joint distribution:

$X\backslash Y$	**a**	**b**
a	1/3	1/3
b	1/3	0

They agree that each one should play a when she or he observes **a**, and b when she or he observes **b**. It can be checked that this agreement is self-enforcing. Suppose that Rowena observes $X = \mathbf{a}$, then she assigns the same (conditional) probability, $\frac{1}{2} = \frac{\frac{1}{3}}{\frac{1}{3}+\frac{1}{3}}$, to actions a and b by Colin; thus, the expected payoff of choosing a is 4, while the expected payoff of choosing b is only 3.5. If she observes $X = \mathbf{b}$, then she is sure that Colin observes $Y = \mathbf{a}$ and plays a; so she best responds with b. A symmetric argument applies to Colin. Therefore no player wants to deviate, whatever she or he observes. The sum of the expected payoffs in this agreement is $\frac{1}{3}(6+6) + \frac{1}{3}(2+7) + \frac{1}{3}(7+2) = 10$.

[16]One can generalize and obtain distributions over mixed equilibria if players choose mixed actions according to the observed realization x.

6.2.2 Definition and characterization of correlated equilibrium

This rather informal discussion motivates the following definitions and results. As we did for mixed equilibria, for the sake of simplicity we restrict the analysis to *finite* games.

Definition 23. *(see Aumann [5] and [6]) A probabilistic self-enforcing agreement, or **correlated equilibrium**, is a structure* $\langle \Omega, p, (T_i, \tau_i, \sigma_i)_{i\in I}\rangle$, *where* (Ω, p) *is a finite probability space (with* $p \in \Delta(\Omega)$*), the sets* T_i *are finite*[17] *and the functions* $\tau_i : \Omega \to T_i$ *and* $\sigma_i : T_i \to A_i$ *are such that,*

$$\forall i \in I, \forall t_i \in T_i, \forall a_i \in A_i, \text{ if } p\left(\left\{\omega' : \tau_i(\omega') = t_i\right\}\right) > 0 \text{ then}$$

$$\sum_{\omega} p(\omega|t_i) u_i\left(\sigma_i(t_i), \sigma_{-i}(\tau_{-i}(\omega))\right) \geq \sum_{\omega} p(\omega|t_i) u_i\left(a_i, \sigma_{-i}(\tau_{-i}(\omega))\right).$$

In the above definition τ_i represents the random variable, also called **signal**, observed by player i, whereas σ_i represents the **strategy,** or **decision function,** imposed by the agreement on player i;

$$p(\omega|t_i) = \begin{cases} p(\omega)/p\left(\{\omega' : \tau_i(\omega') = t_i\}\right), & \text{if } \tau_i(\omega) = t_i, \\ 0, & \text{if } \tau_i(\omega) \neq t_i, \end{cases}$$

is the probability of state ω conditional on observing t_i (which is well defined if the probability of t_i is positive)[18] and

$$\sigma_{-i}(\tau_{-i}(\omega)) = \left(\sigma_j(\tau_j(\omega))\right)_{j\neq i}$$

is the action profile chosen by the other players (according to the agreement) if state ω occurs.

---

[17] Finiteness is just a simplification without substantial loss of generality when A is finite. If Ω and T_i are infinite, they must be endowed with sigma-algebras of events, probability measure p is defined on the sigma-algebra of Ω, and all the functions τ_i and σ_i must be measurable.

[18] This is the special case of the standard definition of conditional probability:

$$p(F) > 0 \Rightarrow p(E|F) = \frac{p(E \cap F)}{p(F)}.$$

If $E = \{\omega\}$ (a singleton) and $F = \{\omega' : \tau_i(\omega') = t_i\}$, then the formula in the main text obtains.

The definition requires that players have no incentive to deviate from the behavior prescribed by the agreement, whatever the possible realization of their signals. Thus, *conditional on every possible observation, the expected payoff of a deviation cannot be higher than the expected payoff obtained by sticking to the agreement.* The inequalities expressing these conditions are sometimes called *incentive compatibility* (or obedience) *constraints.*

Clearly, a correlated equilibrium induces a probability measure over players' action profiles, $\mu \in \Delta(A)$, where

$$\forall a \in A,\ \mu(a) = p\left(\{\omega : \forall i \in I, \sigma_i(\tau_i(\omega)) = a_i\}\right). \tag{6.2.1}$$

Formally, measure $\mu \in \Delta(A)$ is the image of measure $p \in \Delta(\Omega)$ through the pushforward given by the composite function $\sigma \circ \tau : \Omega \to A$, that is, $\mu = p \circ (\sigma \circ \tau)^{-1}$, where $\tau = (\tau_i)_{i \in I} : \Omega \to T$ and $\sigma = (\sigma_i)_{i \in I} : T \to A$; more explicitly,

$$\mu(E) = p((\sigma \circ \tau)^{-1}(E)) = \sum_{\omega : \sigma(\tau(\omega)) \in E} p(\omega)$$

for each subset of action profiles $E \subseteq A$. To analyze the properties of such induced probability measures we introduce the notion of "canonical" correlated equilibrium:

Definition 24. *A self-enforcing probabilistic agreement* $\langle \Omega', p', (\tau_i', \sigma_i')_{i \in I} \rangle$ *where*
(1) $\Omega' = A$,
(2) $p' \in \Delta(A)$,
(3) $\forall i \in I, \forall a = (a_i)_{i \in I} \in A$, $\tau_i'(a) = a_i$ *(*$\tau_i' = \mathrm{proj}_{A_i}$ *is the projection function from* A *onto* A_i*),*
(4) $\forall i \in I$, $\forall a_i \in A_i$, $\sigma_i'(a_i) = a_i$ *(*$\sigma_i' = \mathrm{Id}_{A_i}$ *is the identity function on* A_i*),*
is called ***canonical*** *correlated equilibrium.*

Although, formally, it is the whole structure $\langle \Omega', p', (\tau_i', \sigma_i')_{i \in I} \rangle$ that forms a canonical correlated equilibrium, it is standard to call "canonical correlated equilibrium" just the distribution $p' \in \Delta(A)$, as all the other elements of the structure are trivially determined by the given game G.

Theorem 19. *If* $\langle \Omega, p, (T_i, \tau_i, \sigma_i)_{i \in I} \rangle$ *is a correlated equilibrium, then* $p' = \mu$*, where* μ *satisfies (6.2.1), is a canonical correlated equilibrium* *distribution.*

A canonical correlated equilibrium with distribution μ can be interpreted as the following mechanism. An action profile is chosen at random according to the "agreed upon" probability measure $\mu \in \Delta(A)$. Then a "mediator" observes the realized profile $a = (a_i)_{i \in I}$ and *privately suggests* to each player i to play action a_i. Then player i chooses freely any action $a'_i \in A_i$, but he has no incentive to deviate from the suggested action a_i.

Proof of Theorem 19. Let $\langle \Omega, p, (T_i, \tau_i, \sigma_i)_{i \in I} \rangle$ be a correlated equilibrium; fix arbitrarily a player i and an action a_i with strictly positive marginal probability:

$$\mu(a_i) := (\text{marg}_{A_i}\mu)(a_i) = \sum_{a_{-i} \in A_{-i}} \mu(a_i, a_{-i}) = p\left(\{\omega : \sigma_i(\tau_i(\omega)) = a_i\}\right) > 0,$$

where the second equality follows from the definition of marginal probability, and the third one from eq. (6.2.1). We must show that i has no incentive to deviate from the "mediator's suggestion" a_i, that is

$$\forall a'_i, \sum_{a_{-i} \in A_{-i}} \mu(a_{-i}|a_i) \left[u_i(a_i, a_{-i}) - u_i(a'_i, a_{-i})\right] \geq 0,$$

where

$$\mu(a_{-i}|a_i) := (\text{marg}_{A_{-i}}\mu(\cdot|a_i))(a_{-i}) = \frac{p\left(\{\omega : \sigma(\tau(\omega)) = (a_i, a_{-i})\}\right)}{p\left(\{\omega : \sigma_i(\tau_i(\omega)) = a_i\}\right)}.$$

Since $\mu(a_i) > 0$, the previous system of inequalities is equivalent to

$$\forall a'_i, \ \mu(a_i) \sum_{a_{-i} \in A_{-i}} \mu(a_{-i}|a_i) \left[u_i(a_i, a_{-i}) - u_i(a'_i, a_{-i})\right] \geq 0.$$

Since $\mu(a_i, a_{-i}) = \mu(a_{-i}|a_i)\mu(a_i)$, the incentive compatibility constraints for μ can be expressed as

$$\forall a'_i, \quad \sum_{a_{-i} \in A_{-i}} \mu(a_i, a_{-i}) \left[u_i(a_i, a_{-i}) - u_i(a'_i, a_{-i})\right] \geq 0. \qquad (6.2.2)$$

Hence it is sufficient to show that (6.2.2) holds. This system of inequalities will be derived from the incentive compatibility constraints of the correlated equilibrium $\langle \Omega, p, (T_i, \tau_i, \sigma_i)_{i\in I}\rangle$ that induces μ, thus proving the result.

For every $t_i \in \sigma_i^{-1}(a_i)$ such that $p(\{\omega : \tau_i(\omega) = t_i\}) > 0$ (there must be at least one t_i like this), the incentive compatibility constraints yield

$$\forall a_i', \quad \sum_{\omega:\tau_i(\omega)=t_i} p(\omega|t_i)[u_i(a_i, \sigma_{-i}(\tau_{-i}(\omega))) - u_i(a_i', \sigma_{-i}(\tau_{-i}(\omega)))] \geq 0.$$

Multiplying each t_i-inequality by $p(\{\omega : \tau_i(\omega) = t_i\})$, taking into account that $p(\omega|t_i)p(\{\omega : \tau_i(\omega) = t_i\}) = p(\omega)$ for each $\omega \in \tau_i^{-1}(t_i)$, and taking the summation w.r.t. the t_i's in $\sigma_i^{-1}(a_i)$ yields

$$\forall a_i', \quad \sum_{\omega:\sigma_i(\tau_i(\omega))=a_i} p(\omega)[u_i(a_i, \sigma_{-i}(\tau_{-i}(\omega))) - u_i(a_i', \sigma_{-i}(\tau_{-i}(\omega)))] \geq 0. \tag{6.2.3}$$

Since u_i depends only on actions, the terms in (6.2.3) can be regrouped to obtain

$$\begin{aligned} \forall a_i', \quad & \sum_{a_{-i}} \sum_{\omega\in(\sigma\circ\tau)^{-1}(a_i,a_{-i})} p(\omega)[u_i(a_i, a_{-i}) - u_i(a_i', a_{-i})] \qquad (6.2.4) \\ = \; & \sum_{a_{-i}} [u_i(a_i, a_{-i}) - u_i(a_i', a_{-i})] \sum_{\omega\in(\sigma\circ\tau)^{-1}(a_i,a_{-i})} p(\omega) \geq 0, \end{aligned}$$

where $\sigma = (\sigma_j)_{j\in I} : T \to A$, $\tau = (\tau_j)_{j\in I} : \Omega \to T$, $\sigma \circ \tau : \Omega \to A$ is the composition of σ and τ, and $(\sigma \circ \tau)^{-1}(a_i, a_{-i})$ is the set of pre-images in Ω of action profile (a_i, a_{-i}).

By definition of μ, $\sum_{\omega\in(\sigma\circ\tau)^{-1}(a_i,a_{-i})} p(\omega) = \mu(a_i, a_{-i})$. Therefore (6.2.4) yields (6.2.2) as desired. ■

The first part of the proof yields the following:

Remark 15. *A distribution $\mu \in \Delta(A)$ is induced by a correlated equilibrium (hence, by Theorem 19, it is a canonical correlated equilibrium) if and only if it satisfies the following system of linear inequalities in the probabilities $(\mu(a))_{a\in A}$:*

$$\forall i \in I, \forall a_i, a_i' \in A_i, \quad \sum_{a_{-i}\in A_{-i}} \mu(a_i, a_{-i}) \left[u_i(a_i, a_{-i}) - u_i(a_i', a_{-i})\right] \geq 0.$$

Hence, the set of measures $\mu \in \Delta(A)$ induced by a correlated equilibrium (the set of canonical correlated equilibria) is a convex and compact polytope.[19] *The same holds for the set of expected payoff profiles induced by correlated equilibria.*

The following statement, the proof of which is elementary, clarifies that the correlated equilibrium concept generalizes the mixed equilibrium concept:

Remark 16. *Fix a mixed action profile α and let $\mu^\alpha \in \Delta(A)$ be the measure defined by $\mu^\alpha(a) = \prod_{i \in I} \alpha_i(a_i)$ for every $a \in A$. Then α is a mixed equilibrium if and only if μ^α is a canonical correlated equilibrium.*[20]

Like pure and mixed Nash equilibrium, also correlated equilibrium is a refinement of rationalizability:

Theorem 20. *If μ is a canonical correlated equilibrium, then every action to which μ assigns a positive marginal probability is rationalizable.*

Proof. For every i, let $C_i = \left\{a_i : \sum_{a_{-i}} \mu(a_i, a_{-i}) > 0\right\}$ be the set of i's actions to which μ assigns a positive marginal probability. It will be shown that $C = \times_{i \in I} C_i$ is a set with the best reply property. By Theorem 3, this implies that every action $a_i \in C_i$ is rationalizable, as desired. We write marginal and conditional probabilities using an obvious and natural notation: $\mu(a_i)$, $\mu(a_{-i}|a_i)$, $\mu(a_j|a_i)$.

For every i and a_i, if $a_i \in C_i$, then (by definition of C_i) $\mu(a_i) > 0$, and the conditional probability $\mu(\cdot|a_i) \in \Delta(A_{-i})$ is well defined. Also, for every $j \neq i$ and a_j, $\mu(a_j|a_i) > 0$ only if $\mu(a_j) > 0$. Hence, $\text{supp}\mu(\cdot|a_i) \subseteq C_{-i}$. Finally, given that μ is a canonical correlated equilibrium it must be the case that

$$\forall a_i' \in A_i, \sum_{a_{-i} \in A_{-i}} \mu(a_{-i}|a_i) u_i(a_i, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} \mu(a_{-i}|a_i) u_i(a_i', a_{-i}),$$

that is, $a_i \in r_i(\mu(\cdot|a_i))$. Since this is true for each player i, it follows that C has the best reply property. ■

[19] A polytope in $\mathbb{R}^2$ is just a polygon, a polytope in $\mathbb{R}^3$ is a polyhedron. "Polytope" is the generalization of these geometrical objects to $\mathbb{R}^n$.

[20] The word "correlated" here may be a bit misleading, as we are considering the case where players' actions are mutually independent. Nonetheless, this is a special case of the general definition of correlated equilibrium.

The concept of correlated equilibrium can be generalized assuming that different players can assign different probabilities to the states $\omega \in \Omega$. This generalization is called **subjective correlated equilibrium** (Brandenburger and Dekel [29]). It can be shown that an action a_i is rationalizable *if and only if* there exists a subjective correlated equilibrium in which a_i is chosen in at least one state ω (see Section 8.5.5 of Chapter 8).

6.3 Self-Confirming Equilibrium

We conclude this section on probabilistic equilibria introducing another generalization of the Nash equilibrium concept. In Section 5.4 of Chapter 5, we mentioned the possibility of interpreting an equilibrium as the stationary state of a learning process and we noted that such stationary states need not satisfy the Nash property. Consider the following simple example.

$1\backslash 2$	ℓ	r
t	$2,0$	$2,1$
b	$0,0$	$3,1$

Figure 6.4: Matrix 5.

The game in Matrix 5 has a unique rationalizable outcome, and so a unique Nash equilibrium (and a unique degenerate mixed equilibrium that coincides with the pure strategy equilibrium). If Rowena (player 1) knew the payoff function of Colin (player 2) and were to believe that Colin is rational, then she would expect him to play action r and would play the best reply b.

Instead, we assume the following: (i) information may be *incomplete*, each player knows his payoff function, but does not necessarily know the opponent's payoff function; (ii) the game is played recurrently and, at the end of each round, each player observes his realized payoff, but he does *not* observe directly the opponent's action; (iii) in every period players form probabilistic conjectures on the opponents' actions that are revised according to previous observations; consistently with Bayesian updating, when players observe something that was expected with probability one they do not revise their conjectures. Also, we assume for simplicity that (iv) players maximize the current expected payoff without worrying about the impact of current choices on future payoffs.

The situation of Colin is very simple: he will always choose his dominant action r. Conversely, the situation of Rowena is not so simple. If in any given period t she expects ℓ with probability larger than $\frac{1}{3}$, then she best responds with t. In this case Rowena's payoff is equal to 2 *independently* of the choice made by Colin, hence she is not able to infer Colin's choice from observing the realization $u_1 = 2$, rather she observes something that she was expecting with certainty (to obtain a payoff of 2) and thus she does not revise her probabilistic conjecture. Hence also in period $t + 1$ she repeats the same choice t and according to the same reasoning her conjecture does not change. This shows that if the starting conjecture of Rowena assigns a sufficiently high probability to ℓ the process is stuck in (t, r), which is therefore a stationary state, even if it is not an equilibrium in the sense of Definitions 18 (Nash equilibrium) or 21 (mixed equilibrium).[21]

The example just described illustrates a situation in which players best respond to their conjectures and the information obtained *ex post* does not induce them to change their conjectures, even if they are incorrect. Situations of this kind are known as "**self-confirming equilibria**," or "conjectural equilibria."[22]

6.3.1 Pure Self-Confirming Equilibrium

As it is apparent from the previous example, in order to verify whether a certain situation is a self-confirming equilibrium, it is necessary to

[21] One could object to that if Rowena is patient, then she will want to *experiment* with action b so as to be able to observe (indirectly) Colin's behavior, even if that implies an expected loss in the current period. The objection is only partially valid. It can be shown that for any "degree of patience" (discount factor) there exists a set of initial conjectures that induce Rowena to choose always the safe action a rather than experimenting with b. It is true, however, that the more Rowena values future payoffs, the less is it plausible that the process will be stuck in (t, r).

[22] Both terms ("conjectural" and "self-confirming") have been used in the literature with reference to the same idea (see, for example, the literature review in [20]). For an analysis of how this concept has originated and its relevance for the analysis of adaptive processes in a repeated-iteration context see the survey by Battigalli *et al.* [19] and Chapter 7.

specify what players are able to observe *ex post*. We represent this information with a **feedback** function, $f_i : A \to M_i$, where M_i is a set of "messages" that i could receive at the end of each period. Let $f_{i,a_i} : A_{-i} \to M_i$ denote the section of f_i at a_i. Assuming that i remembers his choice, if i receives message m_i after he has chosen action a_i, he infers that the opponents must have played an action profile from the set $f_{i,a_i}^{-1}(m_i) = \{a_{-i} : f_i(a_i, a_{-i}) = m_i\}$ (In the previous example, $f_1(\cdot) = u_1(\cdot)$, $M_1 = \{0, 2, 3\}$, m_i means "You got m_i euros," $\{a_2 : f_1(t, a_2) = 2\} = \{\ell, r\}$, $\{a_2 : f_1(b, a_2) = 0\} = \{\ell\}$, $\{a_2 : f_1(b, a_2) = 3\} = \{r\}$). Suppose that i's conjecture is μ^i, i chooses the best reply $a_i^* \in r_i(\mu^i)$ and then observes m_i; suppose also that μ^i assigns probability one to the set $f_{i,a_i^*}^{-1}(m_i) = \{a_{-i} : f_i(a_i^*, a_{-i}) = m_i\}$; then the conjecture of i is confirmed (he observes what he expected with probability 1), he sticks to it and keeps choosing a_i^*.

This preliminary discussion shows that in order to ascertain the stability of behavior under learning from personal experience, we must enrich the standard definition of "game" with an *ex post* information structure that specifies what players can learn after any round of a recurrent game has been played. We do so by adding to the static game (in reduced form) $G = \langle I, (A_i, u_i)_{i \in I} \rangle$ a profile of feedback functions $f = (f_i : A \to M_i)_{i \in I}$. We call the pair (G, f) **game with feedback**. Differently from the analysis of mixed and correlated equilibria, here we consider both finite and infinite games.

Definition 25. *A game with feedback (G, f) is **compact-continuous** if G is compact-continuous and, for each player i, the message space M_i is a compact subset of a Euclidean space and the feedback function f_i is continuous.*

With this, we can define the self-confirming equilibrium concept,[23] which captures with a static definition the possible steady states of learning dynamics. We start with equilibria in pure actions, which capture stability in a game recurrently played by the same set of individuals. Then we will generalize to distributions of pure actions with the backdrop of a population game.

Definition 26. *Fix a game with feedback (G, f). A profile of actions and conjectures $(a_i^*, \mu^i)_{i \in I} \in \times_{i \in I} (A_i \times \Delta(A_{-i}))$ is a **self-confirming***

[23] Also called "conjectural equilibrium" as mentioned above.

equilibrium *of* (G, f) *if, for every player* $i \in I$*, the following conditions hold:*
*(1) (*rationality*)* $a_i^* \in r_i(\mu^i)$*,*
*(2) (*confirmed conjectures*)* $\mu^i \left(f_{i,a_i^*}^{-1} \left(f_i(a^*) \right) \right) = 1$.
Profile $(a_i^*)_{i \in I}$ *is a* ***self-confirming equilibrium action profile*** *of* (G, f) *if there is some profile of conjectures* $(\mu^i)_{i \in I}$ *such that* $(a_i^*, \mu^i)_{i \in I}$ *is a self-confirming equilibrium of* (G, f)*.*

Note that, on the one hand, the feedback obtained with non-justifiable actions is irrelevant for the set of self-confirming equilibria. Indeed, the rationality condition (1) implies that self-confirming equilibrium actions are necessarily justifiable; hence, the confirmed conjecture condition (2) actually applies only to justifiable actions. On the other hand, unlike Nash equilibria, in the analysis of self-confirming equilibria one cannot ignore the unjustifiable actions, because a self-confirming equilibrium action can be justified by a confirmed conjecture that (wrongly) assigns positive probability to an unjustifiable action of a co-player.

For example, in the game of Matrix 5 the action profile (t, r) is a self-confirming equilibrium pair if each player observes only his own payoff ($f_i = u_i$, $i = 1, 2$): Pick any (μ^1, μ^2) with $\mu^1(\ell) \geq \frac{1}{3}$ and note that $t \in r_1(\mu^1)$, $\{a_2 : u_1(t, a_2) = u_1(t, r)\} = A_2$, $\{a_1 : u_2(a_1, r) = u_2(t, r)\} = A_1$. If instead $\mu^1(\ell) < \frac{1}{3}$, then $r_1(\mu^1) = \{b\}$. Therefore, t is justified by a confirmed conjecture μ^1 only if μ^1 assigns positive probability to the unjustifiable action ℓ of player 2.

Remark 17. *Every Nash equilibrium* a^* *is a self-confirming equilibrium action profile for every profile of feedback functions* f *(let* $\mu^i(a_{-i}^*) = 1$ *for each* i*). If, for every player* $i \in I$ *and every justifiable action* $a_i \in r_i(\Delta(A_{-i}))$*, the section* $f_{i,a_i} : A_{-i} \to M_i$ *is one-to-one (injective), then the sets of Nash equilibrium and self-confirming equilibrium profiles of actions coincide.*

6.3.2 Anonymous Interaction

The Nash equilibrium concept has been generalized to allow for randomization, thus obtaining the mixed equilibrium concept. We motivated mixed equilibrium as a characterization of stationary patterns of behavior in situations of recurrent anonymous interaction, i.e., situations where a game is played recurrently and every role (e.g., the role of the row player) is played by agents drawn at random from large populations. A similar generalization can be applied to the self-confirming equilibrium concept. Again, we will interpret a mixed action $\alpha_i \in \Delta(A_i)$ as a representation of the fractions of agents in population i playing the different actions $a_i \in A_i$. Since different agents in population i may hold different conjectures, the extension should allow for some heterogeneity: in a stationary state different agents of the same population may choose different actions that are justified by different conjectures. This stands in stark contrast with the mixed Nash equilibrium concept, which relies on the assumption that conjectures are correct, hence common to all agents in any given population i. We will assume for the sake of simplicity that all the agents playing the same action hold the same conjecture. This restriction is without loss of generality. As we did for mixed equilibria, here we simplify the analysis by assuming that *action and message sets are finite*.

Consider an agent from population i who holds conjecture μ^i and keeps playing a best reply $a_i \in r_i(\mu^i)$. In the case of anonymous interaction, what an agent playing in role i observes at the end of the game depends on the behavior of the agents playing in the other roles, but these agents are drawn at random; therefore, each message m_i will be observed, in the long run, with a frequency determined by the fraction of agents in population(s) $-i$ choosing the actions that (together with a_i) yield message m_i. Conjecture μ^i is confirmed (that is, consistent with the evidence) if the subjective probabilities that μ^i assigns to each message given a_i, that is $(\mu^i(f_{i,a_i}^{-1}(m_i))_{m_i \in M_i}$, coincide with the observed long-run frequencies of these messages.

The following example clarifies this point:

$1\backslash 2$	ℓ	r
t	$2,3$	$2,1$
b	$0,0$	$3,1$

Figure 6.5: Matrix 6.

Example 22. Consider the game in Matrix 6 and assume that feedback and payoff functions coincide. Suppose that 25% of the column agents choose ℓ and 75% choose r. Then the probability that a row agent receives the message "You got zero euros" when he chooses b is $\frac{1}{4}$. Row agents do not know this fraction and could, for instance, believe that ℓ occurs with probability $\frac{1}{2}$, which would induce them to choose t; alternatively, a probability of $\frac{1}{5}$ would induce them to play b. It may happen, for instance, that half of the row agents believe that $\mathbb{P}(\ell) = \frac{1}{2}$ and the other half believe $\mathbb{P}(\ell) = \frac{1}{5}$. The former ones choose t, the latter ones b. If the fractions of column agents playing ℓ and r remain 25% and 75% respectively, then those who play b will observe "Zero euros" 25% of the times and "Three euros" 75% of the times. They will then realize that their beliefs were not correct and will keep on revising them until eventually they will coincide with the objective proportions, 25%:75%. These row agents will continue to play b. The other half of the row agents, those that believe that $\mathbb{P}(\ell) = \frac{1}{2}$ and choose t, do not observe anything new and therefore keep on believing and doing the same things. But is it possible that the fractions of column agents playing ℓ and r stay constant? Indeed it is. Suppose that 25% of them believe that $\mathbb{P}(t) = \frac{2}{5}$ and the rest believe that $\mathbb{P}(t) = \frac{1}{5}$. The former ones will choose ℓ (expected payoff: $\frac{6}{5} > 1$), the latter ones will choose r (as $1 > \frac{3}{5}$). Those choosing r do not receive any new information and keep on doing the same thing. Those choosing ℓ half of the times observe "Three euros," and the other half "Zero Euros." Their conjectures are not confirmed, but they will keep revising upward the probability of t and best replying with ℓ, until their conjectures converge to the long-run frequencies 50%-50%, i.e., the actual proportions of row agents choosing t and b. Then there is a stable situation characterized by the following fractions, or mixed actions: $\alpha_1(t) = \frac{1}{2}$, $\alpha_2(\ell) = \frac{1}{4}$. These fractions do not form a mixed Nash equilibrium, because the indifference conditions for a mixed Nash equilibrium yield $\alpha_1^*(t) = \frac{1}{3}$ and $\alpha_2^*(\ell) = \frac{1}{3}$. ▲

It should be clear from the discussion that here, as in our interpretation of the mixed equilibrium concept, mixed actions are interpreted as statistical distributions of pure actions in populations of agents playing in the same role.

Let us move to a general definition. Denote by $\mathbb{P}^{f_i}_{a_i,\mu^i}(m_i)$ the probability of receiving message m_i determined by action a_i and conjecture μ^i, given the feedback function f_i: in the finite case,

$$\mathbb{P}^{f_i}_{a_i,\mu^i}(m_i) = \sum_{a_{-i}:f_i(a_i,a_{-i})=m_i} \mu^i(a_{-i}).$$

Similarly, $\mathbb{P}^{f_i}_{a_i,\alpha_{-i}}(m_i)$ is the probability of m_i determined by a_i and the mixed action profile α_{-i}, given f_i:

$$\mathbb{P}^{f_i}_{a_i,\alpha_{-i}}(m_i) = \sum_{a_{-i}:f_i(a_i,a_{-i})=m_i} \prod_{j\neq i} \alpha_j(a_j).$$

In general, $\mathbb{P}^{f_i}_{a_i,\mu^i}(\cdot) \in \Delta(M_i)$ and $\mathbb{P}^{f_i}_{a_i,\alpha_{-i}}(\cdot) \in \Delta(M_i)$ are, respectively, the probability measures on M_i obtained as pushforward of μ^i and $\mu^{\alpha_{-i}}$ through $f_{i,a_i} : A_{-i} \to M_i$, the section of i's feedback function at a_i:[24]

$$\mathbb{P}^{f_i}_{a_i,\mu^i}(\cdot) = \mu^i \circ f^{-1}_{i,a_i},\ \mathbb{P}^{f_i}_{a_i,\alpha_{-i}}(\cdot) = \mu^{\alpha_{-i}} \circ f^{-1}_{i,a_i}.$$

Definition 27. *Fix a game with feedback* (G, f)*. A profile of mixed actions and conjectures* $(\alpha_i, (\mu^i_{a_i})_{a_i\in\text{supp}\alpha_i})_{i\in I}$ *is an* ***anonymous self-confirming equilibrium*** *of* (G, f) *if for every role* $i \in I$ *and all actions* $a_i \in \text{supp}\alpha_i$ *the following conditions hold:*
*(1) (*rationality*)* $a_i \in r_i(\mu^i_{a_i})$,
*(2) (*confirmed conjectures*)* $\mathbb{P}^{f_i}_{a_i,\mu^i_{a_i}}(\cdot) = \mathbb{P}^{f_i}_{a_i,\alpha_{-i}}(\cdot)$.
$\alpha = (\alpha_i)_{i\in I}$ *is an* ***anonymous self-confirming equilibrium mixed action profile*** *if there is a profile of conjectures* $((\mu^i_{a_i})_{a_i\in\text{supp}\alpha_i})_{i\in I}$ *such that* $(\alpha_i, (\mu^i_{a_i})_{a_i\in\text{supp}\alpha_i})_{i\in I}$ *is an anonymous self-confirming equilibrium.*

Remark 18. *Every mixed Nash equilibrium of* G *is an anonymous self-confirming equilibrium profile for every* f*.*

Remark 19. *Every pure self-confirming equilibrium* $(a^*_i)_{i\in I}$ *is a degenerate anonymous self-confirming equilibrium profile, which can be interpreted as a* symmetric *equilibrium where all agents of the same population* i *play the same action* a^*_i*.*

Since the anonymous version of the self-confirming equilibrium concept is more general, from now on we refer to it simply as "self-confirming equilibrium" without further qualifications.

[24]If A_{-i} is infinite, it is endowed with a sigma-algebra of subsets $\mathcal{A}_{-i} \subseteq 2^{A_{-i}}$ containing all the singletons $\{a_{-i}\}$ $(a_{-i} \in A_{-i})$, conjectures are sigma-additive functions $\mu^i : \mathcal{A}_{-i} \to [0, 1]$, and each section f_{i,a_i} of the feedback function is assumed to be $\mathcal{A}_{-i}$-measurable, so that—in particular— $f^{-1}_{i,a_i}(m_i) \in \mathcal{A}_{-i}$ for each message m_i. These assumptions are satisfied in compact-continuous games with feedback.

Example 23. It can be checked as an exercise that the game in Matrix 6, assuming that the players observe ex post only their own payoff ($f_i = u_i$), admits three types of anonymous self-confirming equilibrium:

1. t is chosen by all the agents in population 1: α_2 can take any value since the conjecture of 1 is necessarily "confirmed" (1, choosing t, does not receive new information); $\alpha_1(t) = 1$, $\mu^1_t(\ell) \geq \frac{1}{3}$, $0 \leq \alpha_2(\ell) \leq 1$, $\mu^2_\ell(t) = 1$, $\mu^2_r(t) \leq \frac{1}{3}$;

2. r is chosen by all agents in population 2: α_1 can take any value since the conjecture of 2 is necessarily "confirmed" (2, choosing r, does not receive new information); $0 \leq \alpha_1(t) \leq 1$, $\mu^1_t(\ell) \geq \frac{1}{3}$, $\mu^1_b(\ell) = 0$, $\alpha_2(\ell) = 0$, $\mu^2_r(t) \leq \frac{1}{3}$;

3. all actions are chosen by a positive fractions of agents: the beliefs of those who choose "informative" actions (b for $i = 1$ and ℓ for $i = 2$) are correct; $\frac{1}{3} \leq \alpha_1(t) \leq 1$, $0 < \alpha_2(\ell) \leq \frac{1}{3}$, $\mu^1_t(\ell) \geq \frac{1}{3}$, $\mu^1_b(\ell) = \alpha_2(\ell)$, $\mu^2_\ell(t) = \alpha_1(t)$, $\mu^2_r(t) \leq \frac{1}{3}$.

Notice that the equilibria of type 1 include the Nash equilibrium (t, ℓ), those of type 2 include the Nash equilibrium (b, r), and those of type 3 include the mixed equilibrium $\alpha^*_1(t) = \frac{1}{3}$, $\alpha^*_2(\ell) = \frac{1}{3}$. ▲

Beliefs about Distributions

In the definition of anonymous self-confirming equilibrium we check whether actions played by a positive fraction of agents are justified by confirmed conjectures. This definition is adequate because a subjective expected utility maximizer ultimately cares only about the probabilities of the actions of his co-players. Yet, what an agent is uncertain about in a population game is the actual *distribution* of actions of the co-players, that is—in a two-population game—, the true $\alpha_{-i} \in \Delta(A_{-i})$.[25] Thus, an agent in population i has a belief $\nu^i \in \Delta(\Delta(A_{-i}))$. Suppose that an agent with such belief ν^i keeps playing a_i and that the true distribution of actions in the other population is α_{-i}, then his belief is confirmed if ν^i assigns probability one to the set of α'_{-i} that yield the same frequency distribution

[25] With more than one co-player drawn at random from different populations, one has to take into account that draws are independent across populations, which makes the analysis slightly more complex.

of messages as α_{-i} given a_i. Using the "pushforward" notation, the confirmation condition is

$$\nu^i\left(\left\{\alpha'_{-i}\in\Delta\left(A_{-i}\right):\alpha'_{-i}\circ f^{-1}_{i,a_i}=\alpha_{-i}\circ f^{-1}_{i,a_i}\right\}\right)=1.$$

Belief $\nu^i\in\Delta\left(\Delta\left(A_{-i}\right)\right)$ yields a corresponding conjecture $\mu^i\in\Delta\left(A_{-i}\right)$ by taking the subjective averages, according to ν^i, of the objective probabilities $\alpha_{-i}(a_{-i})$:[26]

$$\mu^i(a_{-i})=\int_{\Delta(A_{-i})}\alpha_{-i}(a_{-i})\nu^i\left(\mathrm{d}\alpha_{-i}\right).$$

With this, the definition of anonymous self-confirming equilibrium given above is equivalent to requiring that each action played by a positive fraction of agents be justified by the conjecture derived from a confirmed belief over distributions of opponents' actions. However, such equivalence holds only under the assumption of subjective expected utility maximization. With more general decision criteria one has to use the conceptually more precise definition involving beliefs about distributions.[27]

6.3.3 Properties of Feedback and Equilibrium

In simultaneous-moves games, if it is possible to perfectly observe ex post the opponents' actions, then a self-confirming equilibrium (mixed) action profile is necessarily a Nash equilibrium. Indeed, in this case, conjectures are confirmed if and only if they are correct (Remark 17).[28] However, simultaneous-moves games are also used as reduced forms to analyze the equilibria of corresponding games with sequential moves. Specifically, it will be shown that each sequential game admits a "strategic-form" representation whereby it is assumed that players choose in advance and simultaneously the contingent plans (strategies) that will determine their

[26] These subjective averages are also called the "predictive probabilities" implied by belief ν^i.

[27] See Battigalli *et al.* [20]. Furthermore, if we model explicitly the learning process of agents who experiments rationally, the relevant beliefs are those about distributions (see, e.g., Fudenberg and Levine [45] and Battigalli *et al.* [22]).

[28] In the case of anonymous interaction, we can consider at least two scenarios: (a) the individuals observe the statistical distribution of the actions in previous periods, (b) the individuals observe the long-run frequencies of the opponents' actions. In both cases we can say that a conjecture is correct if it corresponds to the observed frequencies.

actions in every possible situation that may arise as the play unfolds (see Section 1.4.2 of the Introduction, and Sections 9.3 and 9.3.1 of Chapter 9). In sequential games, even if, at the end of the game, it is possible to observe all the *actions* that have *actually* been *taken*, it is not possible to observe *how the opponents would have played in circumstances that did not occur*. Hence, it is impossible, at least for some player, to observe the strategies of other players. For this reason, in sequential games it is easier to find self-confirming equilibria that do not correspond to Nash equilibria of the strategic form.

Next we provide conditions on feedback that imply the equivalence of self-confirming equilibria and mixed equilibria. First we define the ex post information partition of A_{-i} corresponding to a given action a_i. Recall that $f_{i,a_i} = f_i(a_i, \cdot) : A_{-i} \to M_i$ is the section at a_i of the feedback function f_i. Section f_{i,a_i} shows how the feedback received by player i depends on the opponents' actions given action a_i. For all $i \in I$ and $a_i \in A_i$, we let

$$\mathcal{F}_{-i}(a_i) = \{C_{-i} \in 2^{A_{-i}} : \exists m_i \in M_i, C_{-i} = f^{-1}_{i,a_i}(m_i)\}$$

denote the partition of A_{-i} given by the pre-images of function f_{i,a_i}. For example, the game in Matrix 6 with $f_i = u_i$ $(i = 1, 2)$ has the following ex post partitions, for each player and action:

$$\begin{aligned} \mathcal{F}_{-1}(t) &= \{A_2\}, \mathcal{F}_{-1}(b) = \{\{\ell\}, \{r\}\}, \\ \mathcal{F}_{-2}(\ell) &= \{\{t\}, \{b\}\}, \mathcal{F}_{-2}(r) = \{A_1\}. \end{aligned}$$

Definition 28. *Feedback function f_i satisfies* ***own-action independence of feedback about others*** *relative to payoff function u_i if all the* justifiable *actions induce the same partition of A_{-i}, that is,*

$$\forall a'_i, a''_i \in r_i\left(\Delta\left(A_{-i}\right)\right), \mathcal{F}_{-i}(a'_i) = \mathcal{F}_{-i}(a''_i). \qquad (6.3.1)$$

Game with feedback (G, f) satisfies ***own-action independence*** *if condition (6.3.1) holds for each player $i \in I$.*

Note that condition (6.3.1) applies only to justifiable actions because, as explained above, only the feedback about such actions matters for the determination of self-confirming equilibria. It is easily checked that the game in Matrix 6 with feedback $f_i = u_i$ $(i = 1, 2)$ does *not* satisfy own-action independence of feedback about others.

Definition 29. *Feedback function f_i satisfies* ***observed payoffs*** *relative to payoff function u_i if for every action of i the feedback received by i determines the payoff of i, that is,*

$$\forall a_i \in A_i, \forall a'_{-i}, a''_{-i} \in A_{-i},\ f_i(a_i, a'_{-i}) = f_i(a_i, a''_{-i}) \Rightarrow u_i(a_i, a'_{-i}) = u_i(a_i, a''_{-i}). \tag{6.3.2}$$

Game with feedback (G, f) satisfies ***observed payoffs*** *if condition (6.3.2) holds for each player $i \in I$.*

Note that the same condition can be expressed as

$$u_i(a_i, a'_{-i}) \neq u_i(a_i, a''_{-i}) \Rightarrow f_i(a_i, a'_{-i}) \neq f_i(a_i, a''_{-i}).$$

In other words, the observed payoff condition says that for each player $i \in I$ there is a function $\pi_i : A_i \times M_i \to \mathbb{R}$ such that $u_i(a_i, a_{-i}) = \pi_i(a_i, f_i(a_i, a_{-i}))$ for every profile (a_i, a_{-i}).[29] The following example illustrates both properties of feedback.

Example 24. Let G be a Cournot oligopoly. Suppose each firm $i \in I$ knows its own total cost function $C_i(\cdot)$ and the inverse market demand function $P(\cdot)$, and that it observes ex post only the market price:

$$f_i(a_i, a_{-i}) = P\left(a_i + \sum_{j \neq i} a_j\right),$$

where a_j is the output of firm j. The utility of firm i is its profit:[30]

$$u_i(a_i, a_{-i}) = \pi_i(a_i, f_i(a_i, a_{-i})) = a_i P\left(a_i + \sum_{j \neq i} a_j\right) - C(a_i).$$

Of course, the observed payoffs condition holds for this game with feedback. Furthermore, if $P(\cdot)$ is invertible, also own-action independence holds. Indeed, each firm i observes ex post the total output of the competitors independently of its own output choice, because

$$\sum_{j \neq i} a_j = P^{-1}(p) - a_i$$

[29] In a more compact and abstract form, $u_{i,a_i} = \pi_{i,a_i} \circ f_{i,a_i}$ for every a_i.

[30] If firm i is owned by a possibly risk averse agent, then $u_i = v_i \circ \pi_i$, where v_i is concave and strictly increasing.

for each a_i and observed price p. Next suppose instead that each firm i only observes the revenue

$$f_i(a_i, a_{-i}) = R_i(a_i, a_{-i}) = a_i P\left(\sum_{j \in I} a_j\right)$$

fetched for its output a_i, but it does not observe the unit price. In this case, the observed payoffs property still holds, because

$$u_i(a_i, a_{-i}) = \pi_i(a_i, f_i(a_i, a_{-i})) = R_i(a_i, a_{-i}) - C(a_i).$$

But own-action independence holds *if and only if* producing no output ($a_i = 0$) is *not justifiable.* Indeed, if $a_i = 0$, then i's revenue is zero for every output profile of the competitors and i cannot recover $\sum_{j \neq i} a_j$. If instead $a_i > 0$, then i can back out from its revenue the unit price

$$p = P\left(\sum_{j \in I} a_j\right) = R_i(a_i, a_{-i}) / a_i$$

and the total output of the competitors:

$$\sum_{j \neq i} a_j = P^{-1}\left(\frac{R_i(a_i, a_{-i})}{a_i}\right) - a_i.$$

Whether $a_i = 0$ is justifiable or not depends on details of the oligopoly model that we did not specify here. For example, if $C_i(a_i) = ca_i$ and there is some competitors' output profile $\hat{a}_{-i}$ such that $P\left(\sum_{j \neq i} \hat{a}_j\right) < c$, then $a_i = 0$ is justifiable as a best response to $\hat{a}_{-i}$ and own-action independence does *not* hold. ▲

Comment on Own-Action Independence The example clarifies that own-action independence of feedback *about others* may hold even if, for some or all i, feedback function f_i depends on a_i: market price in a quantity-setting oligopoly depends on own output, but the observation of market price gives feedback about the output of others that does not depend on own output.

Comment on Payoff Observability As the example suggests, observability of payoffs holds in many applications: what we call "payoff" in game theory is a player's "utility" expressed as a function of all players' actions, and it seems obvious that a player should observe how much utility he gets. Yet, payoff observability is not a tautological assumption. Suppose, for example, that players have other-regarding preferences. This means that they do not only care about the consequences of interaction for themselves, but also for (some) other players. In particular, they may care about the consumption of other players. If player i cares about his own income y_i and the incomes of other players y_{-i}, then his utility is a function of the form $v_i(y_i, y_{-i})$. The outcome function is $y = (y_j)_{j\in I} = (g_j(a))_{j\in I} = g(a)$. Assume that i observes only his own income; then $f_i(a) = g_i(a)$ and $u_i(a) = v_i(f_i(a), g_{-i}(a))$. Hence, we may have $f_i(a') = f_i(a'')$ and yet $u_i(a') \neq u_i(a'')$. Intuitively, $u_i(a)$ is not necessarily a utility "experienced"—hence observed—by i, it is a way of ranking action profiles according to the preferences over the induced income allocations. But player i observes only his own component of the income allocation; hence, he cannot know if the realized allocation is one he ranks highly or not.

The following lemma says that under own-action independence and payoff observability, an action can be justified by a confirmed conjecture if and only if it is a best reply to the true distribution of actions of the co-players.

Lemma 20. *Let (G, f) be a finite or compact-continuous game with feedback; fix any $i \in I$ and suppose that f_i satisfies own-action independence of feedback about others and observed payoffs relative to u_i. Then, for all $a_i^* \in A_i$ and $\alpha_{-i} \in \times_{j\neq i} \Delta(A_j)$ the following are equivalent:*
(1) there is a conjecture $\mu^i \in \Delta(A_{-i})$ such that $a_i^* \in r_i(\mu^i)$ and $\mathbb{P}^{f_i}_{a_i^*,\mu^i}(\cdot) = \mathbb{P}^{f_i}_{a_i^*,\alpha_{-i}}(\cdot)$,
(2) $a_i^* \in r_i(\alpha_{-i})$.

Proof. We prove the result for *finite* games.[31] Fix a_i^* and α_{-i}. It is

[31] The proof for compact-continuous games is left as an exercise for the mathematically savvy reader. Consider that, in this (more general) case, feedback functions are measurable because they are continuous, and best reply correspondences are non-empty valued. The proof shows that the result can be further generalized (cf. Battigalli *et al.* [23]).

obvious that (2) implies (1): let $\mu^i(a_{-i}) = \prod_{j\neq i} \alpha_j(a_j)$ for all $a_{-i} \in A_{-i}$, then (1) holds.

To prove that (1) implies (2), first note that *observation of payoffs* implies that, for each action $a_i \in A_i$, the section of the payoff function of i at a_i, $u_{i,a_i} : A_{-i} \to \mathbb{R}$, is constant on each cell of the ex post information partition $\mathcal{F}_{-i}(a_i)$. Indeed, fix any $C_{-i} \in \mathcal{F}_{-i}(a_i)$ and $a'_{-i}, a''_{-i} \in C_{-i}$. By definition of $\mathcal{F}_{-i}(a_i)$, a'_{-i} and a''_{-i} yield the same message given a_i: $f_i(a_i, a'_{-i}) = f_i(a_i, a''_{-i})$; hence, observability of payoffs implies $u_i(a_i, a'_{-i}) = u_i(a_i, a''_{-i})$. With this, we write $u_i(a_i, C_{-i})$ for the common value of $u_i(a_i, \cdot)$ on cell $C_{-i} \in \mathcal{F}_{-i}(a_i)$. Similarly, we write $f_i(a_i, C_{-i})$ for the common feedback message on cell $C_{-i} \in \mathcal{F}_{-i}(a_i)$ given a_i.[32] Thus,

$$\forall a_i \in A_i,\ \bar{u}_i(a_i, \alpha_{-i}) = \sum_{C_{-i}\in\mathcal{F}_{-i}(a_i)} u_i(a_i, C_{-i}) \sum_{a_{-i}\in C_{-i}} \prod_{j\neq i} \alpha_j(a_j). \quad (6.3.3)$$

Let (1) hold for $\mu^i \in \Delta(A_{-i})$. Then, in particular, a_i^* is justifiable, that is, $a_i^* \in r_i(\Delta(A_{-i}))$. The set of "objective best replies" $r_i(\alpha_{-i})$ contains at least one action a_i^o, which is necessarily justifiable. We are going to prove that $\bar{u}_i(a_i^*, \alpha_{-i}) \geq \bar{u}_i(a_i^o, \alpha_{-i})$. Since $a_i^o \in r_i(\alpha_{-i})$, this implies (2): $a_i^* \in r_i(\alpha_{-i})$.

Since both a_i^* and a_i^o are justifiable, own-action independence implies the common-feedback condition $\mathcal{F}_{-i}(a_i^o) = \mathcal{F}_{-i}(a_i^*)$. Common feedback and eq. (6.3.3) yield

$$\bar{u}_i(a_i^o, \alpha_{-i}) = \sum_{C_{-i}\in\mathcal{F}_{-i}(a_i^*)} u_i(a_i^o, C_{-i}) \sum_{a_{-i}\in C_{-i}} \prod_{j\neq i} \alpha_j(a_j). \quad (6.3.4)$$

The confirmed-conjecture condition stated in (1) implies that the objective and subjective probabilities of each message $f_i(a_i^*, C_{-i})$ coincide:

$$\forall C_{-i} \in \mathcal{F}_{-i}(a_i^*),\ \sum_{a_{-i}\in C_{-i}} \prod_{j\neq i} \alpha_j(a_j) = \sum_{a_{-i}\in C_{-i}} \mu^i(a_{-i}). \quad (6.3.5)$$

Eqs. (6.3.4)-(6.3.5) yield

$$\bar{u}_i(a_i^o, \alpha_{-i}) = \sum_{C_{-i}\in\mathcal{F}_{-i}(a_i^*)} u_i(a_i^o, C_{-i}) \sum_{a_{-i}\in C_{-i}} \mu^i(a_{-i}) = u_i(a_i^o, \mu^i).$$

[32] The message is constant on $C_{-i} \in \mathcal{F}_{-i}(a_i)$ by definition of $\mathcal{F}_{-i}(a_i)$, whether or not realized payoffs are observed.

Eqs. (6.3.3)-(6.3.5) yield

$$\bar{u}_i(a_i^*, \alpha_{-i}) = \sum_{C_{-i} \in \mathcal{F}_{-i}(a_i^*)} u_i(a_i^*, C_{-i}) \sum_{a_{-i} \in C_{-i}} \mu^i(a_{-i}) = u_i(a_i^*, \mu^i).$$

Since $a_i^* \in r_i\left(\mu^i\right)$, the foregoing equalities imply

$$\bar{u}_i(a_i^*, \alpha_{-i}) = u_i(a_i^*, \mu^i) \geq u_i(a_i^o, \mu^i) = \bar{u}_i(a_i^o, \alpha_{-i}).$$

Hence, $a_i^* \in r_i\left(\alpha_{-i}\right)$. ■

Theorem 21. *Let (G, f) be a finite or compact-continuous game with feedback that satisfies own-action independence and observed payoffs. Then the sets of mixed Nash equilibria and self-confirming equilibrium profiles of mixed actions coincide.*

Proof. Every mixed Nash equilibrium is also self-confirming (see Remark 18). To prove the converse, fix any selfconfirming equilibrium $(\alpha_i, (\mu^i_{a_i})_{a_i \in \text{supp}\alpha_i})_{i \in I}$. By definition, for each player $i \in I$ and action $a_i \in \text{supp}\alpha_i$, we have $a_i \in r_i\left(\mu^i_{a_i}\right)$ and $\mathbb{P}^{f_i}_{a_i, \mu^i_{a_i}}(\cdot) = \mathbb{P}^{f_i}_{a_i, \alpha_{-i}}(\cdot)$. Thus, Lemma 20 implies that $\text{supp}\alpha_i \subseteq r_i\left(\alpha_{-i}\right)$ for each $i \in I$. Hence, by Theorem 14, $\left(\alpha_i\right)_{i \in I}$ is a mixed Nash equilibrium. ■

Although observability of payoffs is not tautological, it is common in economic applications. Therefore the above theorem says that the key property for the comparison between (anonymous) self-confirming and (mixed) Nash equilibrium is own-action dependence/independence of feedback.

6.3.4 Comparative Statics for Self-Confirming Equilibrium

Here we analyze how the set of self-confirming equilibria is affected by changes in the informativeness of feedback and in risk aversion. We show that coarser feedback and higher risk aversion expand the set of self-confirming equilibria.[33] This means that the long-run outcome of learning is less predictable and less likely to coincide with Nash equilibrium when feedback is less informative or players are more risk averse.

[33] In the case of risk aversion, the result only refers to the set of pure, or non-anonymous self-confirming equilibria of games with obvervable payoffs.

Coarseness of Feedback By inspection of the definition of self-confirming equilibrium it is easily checked that if the feedback functions f_{i,a_i} are constant for all i and a_i, that is, if there is no feedback about co-players' actions, then every profile $(\alpha_i)_{i\in I}$ that assigns positive weight only to justifiable actions is a self-confirming equilibrium profile of mixed actions. This suggests that poor feedback yields a large set of self-confirming equilibria. We can make this intuition precise. First recall that a partition $\bar{\mathcal{F}}$ is coarser than a partition $\mathcal{F}$ if each cell of $\bar{\mathcal{F}}$ is a union of cells of $\mathcal{F}$.

Definition 30. *We say that game with feedback $\left(G,\bar{f}\right)$ has **coarser feedback** than (G,f) if, for all $i\in I$ and $a_i\in A_i$, $\bar{\mathcal{F}}_{-i}(a_i)$ is coarser than $\mathcal{F}_i(a_i)$.*

The following result says that making feedback coarser expands the set of self-confirming equilibria.

Theorem 22. *Fix two compact-continuous games with feedback (G,f) and $\left(G,\bar{f}\right)$. If $\left(G,\bar{f}\right)$ has **coarser feedback** than (G,f), then every self-confirming equilibrium of (G,f) is also a selfconfirming equilibrium of $\left(G,\bar{f}\right)$.*

Proof. We prove the result for *finite* games.[34] Fix a self-confirming equilibrium $(\alpha_i,(\mu^i_{a_i})_{a_i\in\text{supp}\alpha_i})_{i\in I}$ of (G,f). Since (G,f) and $\left(G,\bar{f}\right)$ differ only with respect to the feedback functions, the given profile satisfies the best response property (1) of Definition 27 in both games with feedback; therefore, we only have to show that the confirmed conjectures property (2) is satisfied also for $\left(G,\bar{f}\right)$. Fix $i\in I$ and $a_i\in$suppα_i arbitrarily. Since the justifying conjecture $\mu^i_{a_i}$ is confirmed in (G,f), then

$$\forall C_{-i}\in\mathcal{F}_{-i}(a_i)\,,\ \mu^i_{a_i}(C_{-i})=\sum_{a_{-i}\in C_{-i}}\prod_{j\neq i}\alpha_j(a_{-i})\,.$$

Since $\bar{\mathcal{F}}_{-i}(a_i)$ is coarser than $\mathcal{F}_{-i}(a_i)$,

$$\forall\bar{C}_{-i}\in\bar{\mathcal{F}}_{-i}(a_i)\,,\ \bar{C}_{-i}=\bigcup_{C_{-i}\in\mathcal{F}_{-i}(a_i):C_{-i}\subseteq\bar{C}_{-i}}C_{-i}.$$

[34] Again, the proof for compact-continuous games is left as an exercise for the mathematically savvy reader. Furthermore, the result holds for more general games (cf. Battigalli *et al.* [23]).

By additivity of $\mu^i_{a_i}$,

$$\forall \bar{C}_{-i} \in \bar{\mathcal{F}}_{-i}(a_i), \; \mu^i_{a_i}\left(\bar{C}_{-i}\right) = \sum_{C_{-i}\in\mathcal{F}_{-i}(a_i):C_{-i}\subseteq\bar{C}_{-i}} \mu^i_{a_i}(C_{-i}) = \sum_{a_{-i}\in\bar{C}_{-i}} \prod_{j\neq i} \alpha_j(a_{-i}).$$

Thus, $\mu^i_{a_i}$ is confirmed in $\left(G, \bar{f}\right)$ as well. ■

Risk Aversion Using Jensen's inequality and Lemma 2 (the duality result of Wald and Pearce), we have shown that an increase in risk aversion expands the set of rationalizable actions (Theorems 1 and 4). A simpler argument relying only on Jensen's inequality shows that an increase in risk aversion expands the set of non-anonymous (i.e., pure) self-confirming equilibria of games with observed payoffs.[35] The intuition is quite simple: under payoff observability, in a pure self-confirming equilibrium each player is certain about the payoff of his equilibrium action, but he may be uncertain about the payoff of deviations; keeping conjectures fixed as we increase risk aversion, deviations with uncertain payoffs become less attractive. Therefore, the same profile of actions and conjectures is a self-confirming equilibrium also with higher risk aversion. On the other hand, higher risk aversion may make "safe" actions more attractive and thus yield new self-confirming equilibria. The following example illustrates.

Example 25. Consider again the game form with monetary payoffs of Example 15 and assume that feedback functions coincide with monetary payoff functions:

$g_1 = f_1, g_2 = f_2:$	b'	b''
a'	0, 1	1, 0
a''	$\frac{1}{3}$,0	$\frac{1}{3}$,1
a'''	1, 1	0, 0

The ex post information partitions are

$$\begin{aligned} \mathcal{F}_{-1}\left(a'\right) &= \mathcal{F}_{-1}\left(a'''\right) = \left\{\{b'\},\{b''\}\right\}, \; \mathcal{F}_{-1}\left(a''\right) = \{A_2\}, \\ \mathcal{F}_{-2}(b') &= \mathcal{F}_{-2}(b'') = \left\{\{a',a'''\},\{a''\}\right\}. \end{aligned}$$

[35] See Battigalli *et al.* [20, Theorem 1] for a similar result about mixed self-confirming equilibria and "ambiguity aversion." See Weinstein [89] for results about Nash equilibrium and risk aversion.

In particular, action a'' does not allow to observe ex post whether player 2 chose b' or b''. Under risk neutrality, a'' is not justifiable, hence it cannot be part of a self-confirming equilibrium. If instead risk aversion is sufficiently high, the uniform conjecture on A_2 justifies a'' and is confirmed, which implies that (a'', b'') is a pure self-confirming equilibrium action pair. Assume $v_1(m_1) = (m_1)^{1/\theta}$ (the risk attitudes of player 2 are immaterial). The pure (non-anonymous) equilibrium correspondence is

$$pSCE_\theta = \begin{cases} \{(a''', b')\}, & \text{if } 1 \leq \theta < \log_2 3, \\ \{(a''', b'), (a'', b'')\}, & \text{if } \theta \geq \log_2 3. \end{cases}$$

(Compare with Examples 6 and 15.) ▲

Theorem 23. *Suppose game with feedback (G, f) satisfies observed payoffs and let players in $\hat{G}$ be more risk averse than in G, then every pure self-confirming equilibrium of (G, f) is also a pure self-confirming equilibrium of $\left(\hat{G}, f\right)$.*

Proof. Fix arbitrarily a self-confirming equilibrium $\left(a_i^*, \mu^i\right)_{i \in I}$ of (G, f) and a player $i \in I$. Then $a_i^* \in r_i\left(\mu^i\right)$ and μ^i is confirmed (given profile a^*); since realized payoffs are observed, $u_i\left(a_i^*, \mu^i\right) = u_i\left(a_i^*, a_{-i}^*\right)$; therefore,

$$\forall a_i \in A_i, u_i\left(a_i^*, \mu^i\right) = u_i\left(a_i^*, a_{-i}^*\right) \geq u_i\left(a_i, \mu^i\right) = \mathbb{E}_{\mu^i}\left(u_{i,a_i}\right), \qquad (6.3.6)$$

where $\mathbb{E}_{\mu^i}(u_{i,a_i})$ is just the explicit expression of $u_i\left(a_i, \mu^i\right)$ as an expectation. We must show that $\left(a_i^*, \mu^i\right)_{i \in I}$ is a self-confirming equilibrium of $\left(\hat{G}, f\right)$. Since we keep the feedback functions fixed, conjecture μ^i is confirmed in $\left(\hat{G}, f\right)$ as well; therefore, we only have to show that $a_i^* \in \hat{r}_i\left(\mu^i\right)$, where $\hat{r}_i$ is the best reply correspondence of i in $\hat{G}$, that is, given payoff function $\hat{u}_i$. Payoff observability holds for $\left(\hat{G}, f\right)$ as well: Let φ be the concave and strictly increasing function such that $\hat{u}_i = \varphi \circ u_i$, then

$$\begin{aligned} f_i(a_i, a'_{-i}) &= f_i(a_i, a''_{-i}) \Rightarrow u_i(a_i, a'_{-i}) = u_i(a_i, a''_{-i}) \\ &\Rightarrow \varphi\left(u_i(a_i, a'_{-i})\right) = \varphi\left(u_i(a_i, a''_{-i})\right) \Rightarrow \hat{u}_i(a_i, a'_{-i}) = \hat{u}_i(a_i, a''_{-i}). \end{aligned}$$

Therefore,

$$\hat{u}_i\left(a_i^*, \mu^i\right) = \hat{u}_i\left(a_i^*, a_{-i}^*\right).$$

With this,

$$\forall a_i \in A_i,\ \hat{u}_i\left(a_i^*, \mu^i\right) = \hat{u}_i\left(a_i^*, a_{-i}^*\right) = \varphi\left(u_{i,a_i^*}\left(a_{-i}^*\right)\right)$$

$$\geq \varphi\left(\mathbb{E}_{\mu^i}\left(u_{i,a_i}\right)\right) \geq \mathbb{E}_{\mu^i}\left(\varphi \circ u_{i,a_i}\right) = \mathbb{E}_{\mu^i}\left(\hat{u}_{i,a_i}\right) = \hat{u}_i\left(a_i, \mu^i\right)$$

where the first inequality follows from eq. (6.3.6) and the monotonicity of φ, and the second is Jensen's inequality, which follows from the concavity of φ. Since i was fixed arbitrarily, this proves that the given profile satisfies the best-reply and confirmed-conjectures properties in $\left(\hat{G}, f\right)$ and is therefore a self-confirming equilibrium of the latter. ■

As is clear from the proof, the result can be generalized as follows: whenever a pure or mixed action profile yields sure payoffs to the players and is a self-confirming equilibrium, then it still is a self-confirming equilibrium with more risk averse players. The following example shows that when equilibrium payoffs are uncertain, or risky, an equilibrium may disappear when risk aversion increases.

Example 26. Consider the game form with monetary outcomes in the following matrix:[36]

g_1, g_2	ℓ	r
t	6, 0	0, 1
m	4, 1	2, 1
b	3, 1	3, 0

(I) First assume that each player *observes his monetary payoff*, that is, $f_i = g_i$ $(i = 1, 2)$. It can be checked that if player 1 is *risk neutral* then $\left(m, \frac{1}{2}\delta_\ell + \frac{1}{2}\delta_r\right)$ is a mixed (Nash and self-confirming) equilibrium profile where player 1 gets 3 dollars on average. Note that player 1 is indifferent among his actions in this equilibrium, despite the fact that m is chosen with probability 1. If player 1 is risk averse this profile cannot be an equilibrium, because the safe action b becomes more attractive given the former equilibrium conjecture. (II) Next assume that player 2 observes his monetary payoff, but *player 1 observes nothing.* Under

[36]The risk attitudes of player 2 are immaterial.

risk neutrality, $\left(\left(m, \mu^1(\ell) = \frac{1}{2}\right), \left(r, \mu^2(m) + \mu^2(t) = 1\right)\right)$ gives a set of self-confirming equilibria (parameterized by $\mu^2(m) \in [0,1]$), but these equilibria disappear under risk aversion, which makes b more attractive given conjecture $\mu^1(\ell) = \frac{1}{2}$. ▲

7

Learning and Solution Concepts

So far we avoided the details of learning dynamics. The analysis of such dynamics requires the use of mathematical tools whose knowledge we do not take for granted in this textbook (differential equations, difference equations, stochastic processes). Nonetheless, it is possible to state some elementary results about learning dynamics by addressing the following question: When is it the case that a *trajectory*, that is, an infinite sequence of action profiles $(a^t)_{t=1}^{\infty}$ (with $a^t = (a_i^t)_{i \in I} \in A$), is *consistent with adaptive learning*? We start with a qualitative answer.

Consider a *finite game* G that is played recurrently. Assume that all actions are observed *ex post* and consider the point of view of a player i that observes that a given profile a_{-i} has *not* been played for a very long time, say for at least T periods. Then it is reasonable to assume that i assigns to a_{-i} a very small probability. If T is sufficiently large, and a_{-i} was not played in the periods $\hat{t}, \hat{t}+1, \ldots, \hat{t}+T$, then the probability of a_{-i} in $t > \hat{t}+T$ will be so small that the best reply to i's conjecture in t will *also* be the best reply to a conjecture that assigns probability *zero* to a_{-i}. In other words, i will choose in $t > \hat{t}+T$ only those actions that are best replies to conjectures μ^i such that $\operatorname{supp}\mu^i \subseteq \left\{a_{-i}^{\tau} : \hat{t} \leq \tau < t\right\}$, that is, only actions in the set $r_i\left(\Delta\left(\left\{a_{-i}^{\tau} : \hat{t} \leq \tau < t\right\}\right)\right)$.[1] Notice that this argument

[1] A_i is finite and $u_i(a_i, \mu^i)$ is continuous (in fact affine) with respect to probabilities $(\mu^i(a_{-i}))_{a_{-i} \in A_{-i}}$. It follows that if a_i^* is a best reply to a conjecture that assigns a "sufficiently small" probability to a_{-i}, then a_i^* is a best reply also to a slightly different

assumes only that i is able to compute best replies to conjectures. The argument is therefore consistent with a high degree of incompleteness of information.

7.1 Perfectly Observed Actions

Building on this intuition, we can use the justification operator ρ to define the consistency of a trajectory $(a^t)_{t=1}^{\infty}$ with adaptive learning under the assumption that the opponents' past actions are perfectly observed.[2] Recall that, for any set $C = \times_{i \in I} C_i \subseteq A$ we let $\rho_i(C_{-i}) = r_i(\Delta(C_{-i}))$, which gives the set of profiles "justified" or "rationalized" by C: $\rho(C) = \times_{i \in I} \rho_i(C_{-i})$. Now, the set of action profiles chosen in a given interval of time is not, in general, a Cartesian product. For this reason, it is useful to generalize the definition of ρ as follows. For $C_{-i} \subseteq A_{-i}$ (not necessarily a product set) let $\rho_i(C_{-i}) = r_i(\Delta(C_{-i}))$. Then, let $C \subseteq A$, and define

$$\rho(C) = \times_{i \in I} \rho_i(\text{proj}_{A_{-i}} C),$$

where $\text{proj}_{A_{-i}} C = \{a_{-i} : \exists a_i \in A_i, (a_i, a_{-i}) \in C\}$ is the set of the other players' action profiles that are consistent with C. (If C is a product set, the usual definition obtains.) Note that, even with this more general definition, the ρ operator is *monotone*: $E \subseteq F$ implies $\rho(E) \subseteq \rho(F)$.

According to the reasoning developed above, suppose that players base their conjectures on past observations and they maximize their expected payoffs in each period. Then, if for a "very long" time only action profiles in C are observed, the current profile of best replies must be in $\rho(C)$. This explains the following definition.

Definition 31. *A trajectory $(a^t)_{t=1}^{\infty}$ is* ***consistent with adaptive learning*** *(when opponents' past actions are perfectly observed) if for every $\hat{t}$ there exists a T such that, for all $t > \hat{t} + T$, $a^t \in \rho\left(\left\{a^\tau : \hat{t} \leq \tau < t\right\}\right)$.*

Next we define when a trajectory $(a^t)_{t=1}^{\infty}$ generates a "limit distribution" $\mu \in \Delta(A)$.

---

conjecture that assigns probability *zero* to a_{-i}.

[2] The following results are adapted from Milgrom and Roberts [58].

Definition 32. *$\mu \in \Delta(A)$ is the* ***limit distribution*** *of $(a^t)_{t=1}^{\infty}$ if for every $a \in A$*

$$\text{(i)} \lim_{t \to \infty} \frac{|\{\tau : 1 \leq \tau < t, a^\tau = a\}|}{t} = \mu(a)$$

$$\text{(ii) if } \mu(a) = 0, \text{ then } \exists \hat{t} : \forall t > \hat{t},\ a^t \neq a.$$

This definition requires that the "long-run frequency" of every a is $\mu(a)$, and furthermore that if $\mu(a) = 0$ there exists a time starting from which the profile a is no longer chosen. Convergence of a trajectory to an action profile a^* is a special case: let $\mu(a^*) = 1$ in the definition, then $(a^t)_{t=1}^{\infty}$ converges to a^*, written $a^t \to a^*$, if there exists a time $\hat{t}$ such that, for every $t > \hat{t}$, $a^t = a^*$ (this is the standard notion of convergence for discrete spaces, that are formed by isolated points). Also, notice that not all trajectories admit a limit distribution.

The following lemma points out some properties necessarily satisfied by limit distributions.

Lemma 21. *If μ is the limit distribution of $(a^t)_{t=1}^{\infty}$, then for every $\hat{t}$, there exists T such that, for every $t > \hat{t} + T$,*

$$a^t \in \text{supp}\mu \subseteq \{a^\tau : \hat{t} \leq \tau < t\}.$$

Proof. Let μ be the limit distribution of $(a^t)_{t=1}^{\infty}$ and fix any given time $\hat{t}$ and profile a. If a is never chosen from $\hat{t}$ onwards, that is if $a \notin \{a^\tau : \hat{t} \leq \tau < t\}$ for every $t > \hat{t}$, then, by Definition 32 (i),

$$\mu(a) = \lim_{t \to \infty} \frac{|\{\tau : 1 \leq \tau < t, a^\tau = a\}|}{t} \leq \lim_{t \to \infty} \frac{\hat{t}}{t} = 0,$$

i.e., $a \notin \text{supp}\mu$. Therefore, for every $a \in \text{supp}\mu$ there exists a time T'_a such that, for every $t > \hat{t} + T'_a$, $a \in \{a^\tau : \hat{t} \leq \tau < t\}$. Let $T' = \max_{a \in \text{supp}\mu} T'_a$ (this is well defined because A is finite). Then, for every $t > \hat{t} + T'$, $\text{supp}\mu \subseteq \{a^\tau : \hat{t} \leq \tau < t\}$. Moreover, if $\mu(a) = 0$ there exists a T''_a such that $\forall t > \hat{t} + T''_a$, $a^t \neq a$ (Definition 32 (ii)). Let $T'' = \max_{a \in A \setminus \text{supp}\mu} T''_a$, then $\forall t > \hat{t} + T''$, $a^t \in \text{supp}\mu$. Let $T = \max\{T', T''\}$, then

$$\forall t > \hat{t} + T,\ a^t \in \text{supp}\mu \subseteq \{a^\tau : \hat{t} \leq \tau < t\}.$$

■

Note that the membership relation means that, from a certain t onwards, only action profiles in the support of the limit distribution occur. The inclusion states that *all* action profiles in the support occur infinitely often.

It is now possible to present a few elementary results that relate adaptive learning with the solution and equilibrium concepts introduced earlier. The first result identifies a sufficient condition for consistency with adaptive learning.[3]

Theorem 24. *Let $(a^t)_{t=1}^{\infty}$ be a trajectory. If the limit distribution of $(a^t)_{t=1}^{\infty}$ (provided that it exists) is a canonical correlated equilibrium, then $(a^t)_{t=1}^{\infty}$ is consistent with adaptive learning.*

Proof. Let μ be the limit distribution of $(a^t)_{t=1}^{\infty}$ and fix an arbitrary $\hat{t}$. It follows from Lemma 21 and the monotonicity of ρ that there exists T such that, for every $t > \hat{t} + T$, $a^t \in \text{supp}\mu$ and

$$\rho(\text{supp}\mu) \subseteq \rho(\{a^\tau : \hat{t} \leq \tau < t\}).$$

Suppose that μ is a canonical correlated equilibrium. Then $\text{supp}\mu \subseteq \rho(\text{supp}\mu)$ (the proof is almost identical to the one of Theorem 20). The claim follows. ■

The next two results show that, in the long run, adaptive learning induces behavior consistent with the "complete-information" solution concepts of earlier chapters, even if consistency with adaptive learning does not require complete information.

Theorem 25. *Let $(a^t)_{t=1}^{\infty}$ be a trajectory consistent with adaptive learning. Then (1)*

$$\forall k \geq 0, \exists t_k, \forall t \geq t_k,\ a^t \in \rho^k(A), \tag{7.1.1}$$

so that only rationalizable actions are chosen in the long run; (2) if $a^t \to a^*$, then a^* is a Nash equilibrium.

Proof. (1) First recall that since A is finite there exists a K such that $\rho^k(A) = \rho^\infty(A)$ for every $k \geq K$. Then, (7.1.1) implies that from some time t_K onwards only rationalizable actions are chosen. The proof of (7.1.1) is by induction. The statement trivially holds for $k = 0$. Suppose

[3] See Definition 23 and Remark 19.

by way of induction that the statement is true for a given k. By consistency with adaptive learning, there exists a T_k such that

$$\forall t > t_k + T_k,\ a^t \in \rho(\{a^\tau : t_k \le \tau < t\}).$$

By the inductive assumption, $\tau \ge t_k$ implies $a^\tau \in \rho^k(A)$. Hence,

$$\{a^\tau : t_k \le \tau < t\} \subseteq \rho^k(A).$$

By monotonicity of ρ,

$$\rho(\{a^\tau : t_k \le \tau < t\}) \subseteq \rho(\rho^k(A)) = \rho^{k+1}(A).$$

Taking $t_{k+1} = t_k + T_k + 1$, the above inclusions yield

$$\forall t \ge t_{k+1},\ a^t \in \rho(\{a^\tau : t_k \le \tau < t\}) \subseteq \rho^{k+1}(A)$$

as desired.

(2) If $a^t \to a^*$, then there exists a $\hat{t}$ such that $a^t = a^*$ for every $t > \hat{t}$. Consistency of $(a^t)_{t=1}^\infty$ with adaptive learning implies that there exists T such that

$$\forall t > \hat{t} + T,\ a^* = a^t \in \rho(\{a^\tau : \hat{t} \le \tau < t\}) = \rho(\{a^*\}).$$

Since $a^* \in \rho(\{a^*\})$, a^* is a Nash equilibrium. ■

7.2 Imperfectly Observed Actions

The definition of consistency with adaptive learning needs to be modified when players do not perfectly observe the actions previously chosen by their opponents. Assume that if an action profile $a = (a_j)_{j\in I}$ is chosen player i only observes a signal (or message) $m_i = f_i(a)$. As in the section on self-confirming equilibria (Section 6.3), the profile of functions $f = (f_i : A \to M_i)_{i\in I}$ is a primitive element of the analysis.

Take a trajectory $(a^t)_{t=1}^\infty$. The set of signals observed by player i from time $\hat{t}$ (included) to time t (excluded) is $\{m_i : \exists\tau, \hat{t} \le \tau < t, m_i = f_i(a^\tau)\}$. From i's perspective, and assuming that i recall the actions he chose, the set of other players' action profiles that could have been chosen in the same time interval is

$$\{a_{-i} : \exists\tau, \hat{t} \le \tau < t, f_i(a^\tau) = f_i(a_i^\tau, a_{-i})\}.$$

This motivates the following definition:

Definition 33. *A trajectory* $(a^t)_{t=1}^{\infty}$ *is* f-consistent with adaptive learning *if for every* $\hat{t}$ *there exists a* T *such that, for every* $t > \hat{t} + T$ *and* $i \in I$, $a_i^t \in \rho_i\left(\left\{a_{-i} : \exists \tau, \hat{t} \leq \tau < t, f_i(a^\tau) = f_i(a_i^\tau, a_{-i})\right\}\right)$.

Observation 7. *Imperfect observability implies that, in general, two players* i *and* j *obtain different information about the past actions of a third player* k*. This is the reason why the justification operator* ρ *could not be used to define the property of* f*-consistency with adaptive learning. If each section* f_{i,a_i} *of feedback function* f_i *is injective (*$i \in I$*,* $a_i \in A_i$*), then there is perfect observability of other players' previous actions and the definition of the previous section (Definition 31) obtaines.*

It is now possible to generalize the results that relate consistency with adaptive learning and equilibrium. Given the discussion of Section 6.3, it should not be surprising that the relevant equilibrium concept for this scenario is self-confirming equilibrium.[4]

Theorem 26. *Fix a game with feedback* (G, f) *and a trajectory* $(a^t)_{t=1}^{\infty}$*. If* $(a^t)_{t=1}^{\infty}$ *has a limit distribution which is part of an anonymous self-confirming equilibrium of* (G, f)*, then* $(a^t)_{t=1}^{\infty}$ *is* f*-consistent with adaptive learning.*

Proof. Fix an anonymous self-confirming equilibrium $(\alpha_i, (\nu_{a_i}^i)_{a_i \in \mathrm{supp}\alpha_i})_{i \in I}$. By the confirmed conjectures condition, for every $i \in I$ and $a_i \in \mathrm{supp}\alpha_i$, $\hat{f}_{i,a_i}(\nu_{a_i}^i) = \hat{f}_{i,a_i}(\alpha_{-i})$, where $\hat{f}_{i,a_i} : \Delta(A_{-i}) \to \Delta(M_i)$ is the pushforward induced by the section $f_{i,a_i} : A_{-i} \to M_i$ of i's feedback function at a_i. This implies $f_{i,a_i}(\mathrm{supp}\nu_{a_i}^i) = f_{i,a_i}(\mathrm{supp}\alpha_{-i})$ and therefore $\mathrm{supp}\nu_{a_i}^i \subseteq f_{i,a_i}^{-1}(f_{i,a_i}(\mathrm{supp}\alpha_{-i}))$, which can be rewritten more explicitly as

$$\mathrm{supp}\nu_{a_i}^i \subseteq \{a_{-i} \in A_{-i} : \exists \hat{a}_{-i} \in \mathrm{supp}\alpha_{-i},\ f_i(a_i, a_{-i}) = f_i(a_i, \hat{a}_{-i})\}. \tag{7.2.1}$$

In words, the conjecture that justifies a given action in a self-confirming equilibrium deems possible only (but possibly not all) the action profiles of the opponents that are consistent with the feedback generated by some of the profiles that actually occur with positive probability. This fact will be used in the proof of the statement, which we are about to undertake.

[4]The following results are adapted from Gilli [50].

Before we delve into the mathematical details, we can provide the intuition behind the result. Recall that, starting from a certain period, say t, the only profiles of actions that occur are those in the support of the limit distribution—provided that it exists—and all such profiles occur infinitely often. If such limit distribution is an anonymous self-confirming equilibrium, then every agent, from t onwards, plays only actions in the support of the equilibrium distribution and keeps playing each of them infinitely often. Thus, the set of messages that every player receives starting from t coincides with the set of messages that are possible according to a justifying conjecture, and all such messages occur infinitely often. This implies that the given trajectory is f-consistent with adaptive learning.

Assume that product measure $\prod_{i\in I}\alpha_i$ of the self-confirming equilibrium is the limit distribution of trajectory $(a^t)_{t=1}^{\infty}$, which implies that α_i is the limit distribution of $(a_i^t)_{t=1}^{\infty}$ for each $i \in I$. Throughout, let $\hat{t}$ be fixed but arbitrary.

Fix $i \in I$. For every $a_i \in \text{supp}\alpha_i$, by the rationality condition, we have $a_i \in r_i(\nu^i_{a_i})$. If we denote by $C^{a_i}_{-i}$ the set on the right hand side of (7.2.1), we have $\nu^i_{a_i} \in \Delta(C^{a_i}_{-i})$, and thus

$$a_i \in r_i(\Delta(C^{a_i}_{-i})) = \rho_i(C^{a_i}_{-i}).$$

Next, we show that there exists T'_i such that, for every $a_i \in \text{supp}\alpha_i$ and $t > \hat{t} + T'_i$,

$$\rho_i(C^{a_i}_{-i}) \subseteq \rho_i\big(\{\hat{a}_{-i} \in A_{-i} : \exists\tau,\, \hat{t} \le \tau < t,\, f_i(a^\tau) = f_i(a_i^\tau, \hat{a}_{-i})\}\big). \quad (7.2.2)$$

By monotonicity of ρ_i, it is sufficient to show

$$C^{a_i}_{-i} \subseteq \{\hat{a}_{-i} \in A_{-i} : \exists\tau,\, \hat{t} \le \tau < t,\, f_i(a^\tau) = f_i(a_i^\tau, \hat{a}_{-i})\}. \quad (7.2.3)$$

Assume $\hat{a}_{-i} \in C^{a_i}_{-i}$, that is, there exists $a_{-i} \in \text{supp}\alpha_{-i}$ such that

$$f_i(a_i, \hat{a}_{-i}) = f_i(a_i, a_{-i}) \quad (7.2.4)$$

(i.e., $\hat{a}_{-i}$ and a_{-i} generate the same feedback given a_i). Since both a_i and a_{-i} are in the support of the *actual* mixed strategies in the limit (product) distribution, the profile $a = (a_i, a_{-i})$ is also played with positive probability, that is, $a \in \text{supp}\alpha$. Thus, a occurs infinitely often in the trajectory $(a^t)_{t=1}^{\infty}$ and there exists $\tau^* > \hat{t}$ such that $a^{\tau^*} = a$, which also implies $f_i(a^{\tau^*}) = f_i(a) = f_i(a_i, \hat{a}_{-i})$. Taking $T'_{i,a_i} > \tau^* - \hat{t}$, it is immediately verified that $\hat{a}_{-i}$ belongs to the right hand side of (7.2.3) for every $t > \hat{t} + T'_{i,a_i}$ (τ^* is the integer τ whose existence is required in the specification of the set). Thus, $\max\{T'_{i,a_i}\}_{a_i \in \mathrm{supp}\alpha_i}$ is the desired T'_i. It follows that, for every $t > \hat{t} + T'_i$

$$\mathrm{supp}\alpha_i \subseteq \bigcup_{a_i \in \mathrm{supp}\alpha_i} \rho_i(C_{-i}^{a_i}) \subseteq \rho_i\big(\{\hat{a}_{-i} \in A_{-i} : \exists \tau,\, \hat{t} \leq \tau < t,\, f_i(a^\tau) = f_i(a_i^\tau, \hat{a}_{-i})\}\big).$$

By Lemma 21, there exists T''_i such that, for every $t > \hat{t} + T''_i$, $a_i^t \in \mathrm{supp}\alpha_i$. Choosing $T_i > \max\{T'_i, T''_i\}$ and $T > \max\{T_i\}_{i \in I}$, we get:

$$\forall t > \hat{t} + T, \forall i \in I,\, a_i^t \in \rho_i\big(\{\hat{a}_{-i} \in A_{-i} : \exists \tau,\, \hat{t} \leq \tau < t,\, f_i(a^\tau) = f_i(a_i^\tau, \hat{a}_{-i})\}\big).$$

Since $\hat{t}$ is arbitrary, this complete the proof. ■

The following result is a partial converse:

Theorem 27. *Fix a game with feedback (G, f) and a trajectory $(a^t)_{t=1}^{\infty}$. If $(a^t)_{t=1}^{\infty}$ is f-consistent with adaptive learning and $a^t \to a^*$ then a^* is a self-confirming equilibrium action profile.*

Proof. If $a^t \to a^*$, then there exists $\hat{t}$ such that $a^t = a^*$ for every $t > \hat{t}$. By f-consistency of $(a^t)_{t=1}^{\infty}$ with adaptive learning, it follows that there exists T such that, for every $t > \hat{t} + T$ and for every $i \in I$,

$$\begin{aligned} a_i^* &= a_i^t \in \rho_i\left(\left\{a_{-i} : \exists \tau, \hat{t} \leq \tau < t, f_i(a^\tau) = f_i(a_i^\tau, a_{-i})\right\}\right) \\ &= \rho_i\left(\left\{a_{-i} : f_i(a_i^*, a_{-i}^*) = f_i(a_i^*, a_{-i})\right\}\right). \end{aligned}$$

Hence, for every player i there exists a belief μ^i such that: (1) $a_i^* \in r_i(\mu^i)$, (2) $\mu^i\left(\left\{a_{-i} : f_i(a_i^*, a_{-i}^*) = f_i(a_i^*, a_{-i})\right\}\right) = 1$, which means that $(a_i^*, \mu^i)_{i \in I}$ is a self-confirming equilibrium. ■

8

Incomplete Information

Let $G = \langle I, (A_i, u_i)_{i\in I}\rangle$ be a static game where the payoff functions $u_i : A \to \mathbb{R}$ are derived from an outcome function $g : A \to Y$ and utility functions $v_i : Y \to \mathbb{R}$, that is, $u_i = v_i \circ g$, for each $i \in I$. To make sense of the strategic analysis of the game, it is typically assumed that G is common knowledge, or—in other words—that there is *complete information.* The reason is apparent for "deductive" solutions concepts, such as rationalizability, that characterize the set of action profiles consistent with common knowledge of the game, rationality and common belief in rationality. Also the interpretation of the Nash equilibrium concept (or more generally of correlated equilibrium) as a "self-enforcing agreement" makes more sense under the complete information assumption. To see this, suppose that before the game is played the players agreed to play some Nash equilibrium profile a^*. Then, nobody has an incentive to deviate *if he expects that the others comply with the agreement.* But why should the players expect that the agreement is complied with? If the payoff functions of others are unknown, a player could suspect that some other player has an incentive to deviate and hence he may want to deviate too.

In some games the incentive to deviate is strong even if the chance that the opponent might not honor the agreement is small. For example, consider the Pareto efficient agreement (a, c) of the following payoff matrix:

If Rowena does not know Colin's payoff function, she may suspect—for example—that d is dominant for Colin, and hence switch to her safe action b provided she assigns to d a probability above 1%. A similar argument

	c	d
a	$100, 100$	$0, 99$
b	$99, 0$	$99, 99$

applies to Colin. Now suppose that payoff functions are mutually known, but Rowena is not sure that Colin knows her payoff function. Then she assigns a positive probability to Colin playing his safe action d instead of the agreed upon action c, and again she would go for the safe action, b, provided she assigns to d a probability above 1%. Continuing like this, it can be shown that even a high degree of mutual knowledge of the game does not ensure that the agreement will be honored. On the other hand, common knowledge of the game makes the agreement self-enforcing.[1]

8.1 Games with Payoff Uncertainty

How can we represent a game with incomplete information? The problem can be viewed as follows: the outcome function $g : A \to Y$ and the utility functions $v_i : Y \to \mathbb{R}$ depend on a vector of parameters θ which is not commonly known. Hence, the payoff functions $u_i = v_i \circ g$ depend on a parameters vector which is not commonly known.[2] To represent this uncertainty, payoff functions can be written in a parameterized form

$$u_i : \Theta \times A \to \mathbb{R},$$

where Θ is the set of values of θ that are not excluded by what is commonly known among the players.

In order to represent what a player knows about θ, we assume for simplicity that θ is decomposable into subvectors θ_j, with $j \in \{0\} \cup I$, and that each player $i \in I$ knows the true value of θ_i, whereas θ_0 represents residual uncertainty that would persist even if the players could credibly share their private information. To further simplify the analysis, we will sometimes assume that there is no such residual uncertainty (in many

[1] Vice versa, recall that the interpretation of a Nash equilibrium (pure or mixed) as a stationary state of an adaptive process does *not* require common knowledge, nor mutual knowledge of the game.

[2] The set of the opponents' actions could be unknown as well. We omit this source of uncertainty for the sake of simplicity.

cases, this assumption is rather innocuous). In this case Θ_0 is a singleton and it makes sense to omit it from the notation, writing $\theta = (\theta_i)_{i\in I}$.

Example 27. Consider again the team of two players producing a public good (Example 1). Now we express the production function and cost-of-effort functions in a slightly different parametric form. The parameterized production function is

$$Q = \theta_0 (a_1)^{\theta_1^q} (a_2)^{\theta_2^q}.$$

The parameterized cost of the effort of player i measured in terms of output is

$$C_i = \frac{1}{2\theta_i^c}(a_i)^2.$$

The parameterized consequence function specifies a triple (Q, C_1, C_2) as a function of (θ, a_1, a_2):

$$g(\theta, a_1, a_2) = \left(\theta_0 (a_1)^{\theta_1^q} (a_2)^{\theta_2^q}, \frac{1}{2\theta_1^c}(a_1)^2, \frac{1}{2\theta_2^c}(a_2)^2\right).$$

The *commonly known utility function* of player i is $v_i(Q, C_1, C_2) = Q - C_i$. The *parameterized payoff function* is then

$$u_i(\theta, a_1, a_2) = v_i(g(\theta, a_1, a_2)) = \theta_0 (a_1)^{\theta_1^q} (a_2)^{\theta_2^q} - \frac{1}{2\theta_i^c}(a_i)^2.$$

It is assumed to be common knowledge that $\theta_0 \in \Theta_0 \subseteq \mathbb{R}_+$ and $\theta_i = (\theta_i^q, \theta_i^c) \in \Theta_i \subseteq \mathbb{R}_+^2$ where the sets Θ_0, Θ_1 and Θ_2 are specified by the given model. Each player i knows his efficiency parameter $\theta_i = (\theta_i^q, \theta_i^c)$ and nobody knows θ_0, unless Θ_0 is a singleton. ▲

In Example 27 the set $\Theta = \Theta_0 \times \Theta_1 \times \Theta_2$ only includes parameters that directly affect the payoff functions of players and, for this reason, we call them **payoff-relevant** (θ_1^q and θ_2^q are output elasticities with respect to agents' actions, θ_1^c and θ_2^c are coefficients of the cost functions, and θ_0 is a "total factors productivity" parameter). However, parameter vector θ_i can also include **payoff-irrelevant** information, namely information that does not *directly* affect payoffs. As we are going to show in this chapter, even payoff-irrelevant parameters may affect the strategic analysis. For instance, a player may believe that some payoff-irrelevant information that he or other players observe is correlated with parameters affecting payoffs.

Unlike θ_i $(i \in I)$, the parameter θ_0 that captures residual uncertainty may be assumed to affect the payoff of at least one player: since it is common knowledge that players ignore θ_0 and cannot condition their behavior on θ_0, there is no reason to include payoff-irrelevant components in θ_0. Formally, it may be assumed without any loss of generality that, whenever $\theta_0' \neq \theta_0''$, there must be some player j, information profile $(\theta_i)_{i\in I}$, and action profile $a = (a_i)_{i\in I}$ such that $u_j(\theta_0', (\theta_i)_{i\in I}, a) \neq u_j(\theta_0'', (\theta_i)_{i\in I}, a)$.

Example 28. Consider the following Cournot oligopoly game with incomplete information: The market inverse demand function is

$$P(\theta_0, Q) = \max\{0, \theta_0 - Q\}, \text{ with } Q = \sum_{i=1}^{n} q_i.$$

The cost function of each firm i is given by

$$C_i(\theta_i, q_i) = \frac{1}{2\theta_i^c}(q_i)^2.$$

Then

$$u_i(\theta, q_1, ..., q_n) = q_i \max\left\{0, \theta_0 - \sum_{i=1}^{n} q_i\right\} - \frac{1}{2\theta_i^c}(q_i)^2. \tag{8.1.1}$$

Thus, θ_0 is an (unknown) shift parameter that captures market conditions, while θ_i^c parametrizes the cost function of firm i. Obviously, θ_0 and $(\theta_i^c)_{i\in I}$ are payoff relevant. Further, assume that each firm has a marketing department which performs a market analysis and issues a score to summarize its report; in particular, for every $i \in I$, let $\theta_i^m = \theta_0 + \varepsilon_i$, where ε_i is normally distributed with mean 0 and precision τ_i and $\varepsilon_i, \varepsilon_j$ are independent for every i, j with $i \neq j$, conditional on θ_0.[3] The payoff function of each firm i, given by eq. (8.1.1), is independent of $(\theta_j^m)_{j\in I}$; thus, each θ_i^m $(i \in I)$ is payoff-irrelevant. Nevertheless, θ_i^m may play a role in the strategic analysis: (1) firm i can condition its behavior on θ_i^m, (2) firm $j \neq i$ may believe that i will condition its behavior on θ_i^m and try to exploit the correlation between θ_j^m, θ_0 and θ_i^m to make infererence about the behavior of i, and so on. ▲

[3] Assume, for simplicity, that the marketing department operational cost is equal to 0.

In the previous narrative, θ is mostly interpreted as a fixed parameter vector imperfectly and asymmetrically known by the players. To stress this interpretation we sometimes call θ **state of nature**. However, in many models, either the whole θ, or some of its components, are intepreted as the realization of an exogenous random variable. To illustrate, in Example 28 the shift parameter θ_0 may represent the realization of a demand shock. The specific interpretation of θ (or some of its components) as a fixed parameter or the realization of an exogenous random variable is more important for extensions of the Nash and selfconfirming equilibrium concepts to enviroments with incomplete and asymmetric information. Thus, we will consider it again as we discuss these extensions in Sections 8.6 and 8.7.

8.1.1 Private and Interdependent Values

In general, player i's payoff function, u_i, may depend on the whole profile $\theta = (\theta_j)_{j \in I}$, and not just on θ_i (see Example 27). However, the analysis is simpler when the consequence function g is commonly known and uncertainty only concerns the utility functions $v_i : Y \to \mathbb{R}$. In this case, each player i knows his own payoff function $u_i = v_i \circ g$ (given that g is commonly known and v_i represents i's preferences over lotteries of consequences). To represent other players' ignorance of i's utility (and payoff) function, such a function can be written in parameterized form $v_i(\theta_i, y)$, where θ_i is known only to player i. This yields parameterized payoff functions $u_i(\theta_i, a) = v_i(\theta_i, g(a))$, where i knows θ_i. Whenever u_i varies only with θ_i (and a) the game is said to have **private values**.

Suppose now that the consequence function g is *not* common knowledge. Uncertainty about the consequence function can be expressed by representing g in the parameterized form $g(\theta, a)$. Then, θ_i identifies not only i's preferences, but also i's private information about the consequence function. Payoff functions have the general parameterized form $u_i(\theta, a) = v_i(\theta_i, g(\theta, a))$. Since u_i depends on the whole parameter vector θ, the game is said to have **interdependent values**, meaning that u_i may vary with θ_j $(j \neq i)$.

Finally, we say that there is distributed knowledge of θ if parameter θ can be identified by pooling the information of all the players. More formally, let $\theta = (\theta_0, (\theta_i)_{i \in I})$ and let Θ_j denote the set of possible values of θ_j $(j \in \{0\} \cup I)$. Then there is **distributed knowledge of** θ when Θ_0

is a *singleton*.

Example 29. Consider again Example 28. Then, there is distributed knowledge of θ if the (inverse) demand function is commonly known, i.e., $\Theta_0 = \{\bar{\theta}_0\}$. In this case, the model also features private values. ▲

Example 30. In the game of Example 27 there is distributed knowledge of θ if and only if there is only one possible value, say $\bar{\theta}_0$, of the total productivity parameter, i.e., $\Theta_0 = \{\bar{\theta}_0\}$. There are private values if and only if there is distributed knowledge of θ, hence common knowledge of total productivity $\bar{\theta}_0$, and moreover there is common knowledge of the output elasticities with respect to players' actions, $\bar{\theta}_1^q$ and $\bar{\theta}_2^q$, that is, $\Theta_i = \{\bar{\theta}_i^q\} \times \Theta_i^c$ $(i = 1, 2)$. Otherwise there are interdependent values.

▲

To summarize, a strategic environment with incomplete information can be simply described by the following mathematical structure

$$\hat{G} = \langle I, \Theta_0, (\Theta_i, A_i, u_i : \Theta \times A \to \mathbb{R})_{i \in I} \rangle ,$$

where all the sets Θ_j $(j \in I \cup \{0\})$ and A_i $(i \in I)$ are nonempty, Θ_0 is the space of residual uncertainty, Θ_i the set of parameter values that i's opponents consider possible, and—as usual—$\Theta = \times_{i \in I \cup \{0\}} \Theta_i$, $A = \times_{i \in I} A_i$. We call θ_i, i's private information about parameter θ, the **information-type** of i (the term "type" will be used later for a more complex object). We call $\hat{G}$ "**game with payoff uncertainty**" rather than "game with incomplete information;" the latter terminology, although perfectly legitimate, is often used to indicate a game endowed with a description of the possible beliefs of players about θ—see Section 8.4. We assume that all the sets and functions specified by structure $\hat{G}$ are commonly known. Otherwise, it would mean that some aspects of players' uncertainty have been omitted from the formal representation.

Note, *the "size" of Θ_j measures the ignorance of each player $i \neq j$ about θ_j*. It is important to specify how ignorant each player is. Even if the modeler knew the true value of θ, say $\bar{\theta}$, he should still put the sets Θ_j in the formal model. This is true even in the private value case. Why? Because in the strategic analysis of the game each player i puts himself in the shoes of each other player j and considers how j would behave for each information-type θ_j that i deems possible.

We can extend to games with payoff uncertainty some properties defined earlier.

Definition 34. *A game with payoff uncertainty $\hat{G}$ is **compact-continuous** if all the sets Θ_j $(j \in I \cup \{0\})$ and A_i $(i \in I)$ are compact subsets of Euclidean spaces, and every payoff function u_i is continuous $(i \in I)$.*

Definition 35. *A game with payoff uncertainty $\hat{G}$ is **nice** if (1) it is compact-continuous, (2) all the sets Θ_j $(j \in I \cup \{0\})$ are convex subsets of Euclidean spaces,*[4] *and, (3) for every player $i \in I$, A_i is a compact interval in $\mathbb{R}$, and the section of his payoff function $u_{i,\theta,a_{-i}} : A_i \to \mathbb{R}$ is strictly quasi-concave for all $(\theta, a_{-i}) \in \Theta \times A_{-i}$.*

The games of Examples 27 and 28 are nice if all the parameter sets and action sets are convex and compact.

8.2 Rationalizability and Payoff Uncertainty

The specification of a game with payoff uncertainty is sufficient to obtain some results about the outcomes of strategic interaction that are consistent with rationality and standard assumptions about players' beliefs (for instance, some degrees of mutual belief in rationality). Consider the following examples:

Example 31. There are two possible payoff functions for each player i, $u_i(\theta^a, \cdot)$ and $u_i(\theta^b, \cdot)$. Player 1 (Rowena) knows the true payoff functions while player 2 (Colin) does not know them. The situation can be represented with two matrices, corresponding to the two possible values of the parameter θ, assuming that Rowena knows the true payoff matrix.[5] Action a dominates b in payoff matrix θ^a, while b is justifiable (not dominated) in matrix θ^b.

---

[4] It is enough to assume that each Θ_j is connected, i.e., it is not contained in a union of disjoint open sets. Every convex set is connected. A subset of $\mathbb{R}$ is connected if and only if it is convex, but a connected subset of $\mathbb{R}^k$ may be nonconvex. What matters for our analysis is that, for every continuous function $f : X \to \mathbb{R}$ $(X \subseteq \mathbb{R}^k)$, if X is convex, or connected, then $f(X)$ is an interval.

[5] In this case, we can drop indexes from θ as only player 1 has private information and therefore there is a one-to-one correspondence between Θ and Θ_1. Thus, $\Theta_1 = \{\theta_1^a, \theta_1^b\}$, $\Theta_2 = \{\bar{\theta}_2\}$, $\Theta = \{\theta^a, \theta^b\}$ with $\theta^a = (\theta_1^a, \bar{\theta}_2)$ and $\theta^b = (\theta_1^b, \bar{\theta}_2)$.

$\hat{G}^1$:

θ^a	c	d
a	$4, 0$	$2, 1$
b	$3, 1$	$1, 0$

θ^b	c	d
a	$2, 0$	$0, 1$
b	$0, 1$	$1, 2$

Consider the following epistemic assumptions: R_1, R_2, $\mathrm{B}_1(R_2)$, $\mathrm{B}_2(R_1)$, $\mathrm{B}_1(\mathrm{B}_2(R_1))$.[6] Obviously the implications of such hypotheses for Rowena's choice depend on the true value of θ (which is known only to her). As theorists, we do not make any assumption about θ, but rather we want to explain what may happen (consistently with the above mentioned assumptions) for each possible value $\theta \in \Theta$.
R_1 implies that Rowena chooses a if $\theta = \theta^a$ (a is dominant in matrix θ^a).
$R_2 \cap \mathrm{B}_2(R_1)$ implies that Colin chooses d. Indeed, $\mathrm{B}_2(R_1)$ implies that Colin is certain that Rowena would choose a if $\theta = \theta^a$. Hence, Colin assigns probability zero to the pair (θ^a, b). Also, notice that d is "dominant" when the set of possible pairs is restricted to $\{(\theta^a, a), (\theta^b, a), (\theta^b, b)\}$.
$R_1 \cap \mathrm{B}_1(R_2) \cap \mathrm{B}_1(\mathrm{B}_2(R_1))$ implies that Rowena expects d and as a consequence chooses b if $\theta = \theta^b$.
To sum up, the aforementioned assumptions imply that the outcome is (a, d) if $\theta = \theta^a$ and (b, d) if $\theta = \theta^b$. Note, the set of i's actions consistent with rationality and these epistemic assumptions can depend only on what i knows. Then, in this case, the set of possible actions of Colin (the singleton $\{d\}$) is independent of θ. ▲

Example 32. Players 1 and 2 receive an envelope with a money prize of k thousands Euros, with $k = 1, ..., K$. It is possible that both players receive the same prize. Every player knows the content of her own envelope and can either offer to exchange her envelope with that of the other player (action OE=*Offer to Exchange*), or not (action N=*do Not offer to exchange*). The OE/N actions are taken simultaneously and the exchange takes place only if it is offered by both. In order to offer an exchange a player has to pay a small transaction cost ε. The players' payoff is given by the amount of money they end up with at the end of the game. For this example, $\Theta_i = \{1, ..., K\}$ and $u_i(\theta, a)$ is given by the following table:

Therefore, a necessary condition for i to offer to exchange is that he assigns a positive probability to event $[\theta_j > \theta_i] \cap [a_j = \mathrm{OE}]$. It can be

[6] The meaning of these symbols is explained in Chapter 4.

$a_i \backslash a_j$	OE	N
OE	$\theta_j - \varepsilon$	$\theta_i - \varepsilon$
N	θ_i	θ_i

shown that the assumptions of rationality and common belief in rationality imply that i keeps his envelope, whatever the content. Let us consider, for simplicity, only the case $K = 3$ (the general result can be obtained by induction):
R_i implies that i keeps the envelope if $\theta_i = 3$, since by offering to exchange he could obtain at most $3 - \varepsilon$.
$R_i \cap \mathrm{B}_i(R_j)$ implies that i keeps the envelope if $\theta_i \geq 2$. Indeed, i is certain that j would not offer to exchange if $\theta_j = 3$; hence i is certain that by offering to exchange he would obtain at most $2 - \varepsilon$.
$R_i \cap \mathrm{B}_i(R_j) \cap \mathrm{B}_i(\mathrm{B}_j(R_i))$ implies that i keeps the envelope whatever the value of θ_i. Indeed i is certain that j could offer to exchange only if $\theta_j = 1$; hence i is certain that by offering to exchange he could obtain at most $1 - \varepsilon$. ▲

Intuitive, step-by-step solutions such as those of the previous examples can be obtained by a simple extension of the rationalizability solution concept from games with complete information to games with payoff uncertainty. In general, we deem "rationalizable" any outcome consistent with the assumptions of *rationality and common belief in rationality*, which we denote using the symbol $R \cap \mathrm{CB}(R)$,[7] given the *common background knowledge*. Rationalizability in games of complete information, for instance, identifies the choices that are consistent with assumptions $R \cap \mathrm{CB}(R)$ given that the game is common knowledge.

In the analysis of strategic interaction with incomplete information, it is interesting to address the following question: What behavior is consistent with $R \cap \mathrm{CB}(R)$ given that the features of strategic interaction encoded by structure $\hat{G} = \langle I, \Theta_0, (\Theta_i, A_i, u_i)_{i \in I} \rangle$ are common knowledge? The intuitive solutions of the previous examples suggest how to obtain the answer: eliminate, step by step, the pairs (θ_i, a_i) such that a_i is not a best reply to any probabilistic conjecture $\mu^i \in \Delta(\Theta_0 \times \Theta_{-i} \times A_{-i})$ that assigns

[7]Recall that $\mathrm{CB}(R) = \bigcap_{k \geq 1} \mathrm{B}^k(R)$ denotes common probability-one belief in rationality.

probability zero to every (θ_{-i}, a_{-i}) eliminated in the previous steps. By Lemma 2, in a given step of the elimination procedure one has to delete the pairs (θ_i, a_i) such that, for θ_i, a_i is dominated by some mixed action α_i given the previous steps of elimination.

More formally, this solution procedure can be defined *via* a generalization of the justification operator. Consider a finite or compact-continuous $\hat{G}$. First we extend the definition of best reply correspondence to the payoff-uncertainty setup: for every $\mu^i \in \Delta(\Theta_0 \times \Theta_{-i} \times A_{-i})$ and $\theta_i \in \Theta_i$, let

$$r_i(\mu^i, \theta_i) = \arg\max_{a_i \in A_i} \mathbb{E}_{\mu^i}\left(u_{i,\theta_i,a_i}\right), \tag{8.2.1}$$

where $u_{i,\theta_i,a_i} : \Theta_0 \times \Theta_{-i} \times A_{-i} \to \mathbb{R}$ is the section of i's parameterized payoff function at (θ_i, a_i). When conjecture μ^i has a finite support, the expected utility formula is

$$\mathbb{E}_{\mu^i}\left(u_{i,\theta_i,a_i}\right) = \sum_{(\theta_0,\theta_{-i},a_{-i}) \in \operatorname{supp}\mu^i} u_i\left(\theta_0, \theta_i, \theta_{-i}, a_i, a_{-i}\right) \mu^i\left(\theta_0, \theta_{-i}, a_{-i}\right).$$

As usual, we slightly abuse notation and write $\Delta(\Theta_0 \times C_{-i})$ for the set of probability measures on $\Theta_0 \times (\bigtimes_{j \neq i} \Theta_j \times A_j)$ that assign probability one to $\Theta_0 \times C_{-i}$. The following lemma says that the best reply correspondence $(\mu^i, \theta_i) \mapsto r_i(\mu^i, \theta_i)$ defined above is "well behaved," that is, it is nonempty valued and it has a closed graph:[8]

Lemma 22. *If $\hat{G}$ is a compact-continuous game with payoff uncertainty, for every $i \in I$, and every nonempty closed subset $C_{0,-i} \subseteq \Theta_0 \times \Theta_{-i} \times A_{-i}$,*

$$\operatorname{Gr}\left(r_i|_{\Delta(C_{0,-i}) \times \Theta_i}\right) = \left\{\left(\mu^i, \theta_i, a_i\right) \in \Delta\left(C_{0,-i}\right) \times \Theta_i \times A_i : a_i \in r_i\left(\mu^i, \theta_i\right)\right\}$$

is closed, and $r_i\left(\mu^i, \theta_i\right) \neq \emptyset$ for all $\left(\mu^i, \theta_i\right) \in \Delta\left(C_{0,-i}\right) \times \Theta_i$.

Let $\mathcal{C}$ denote the collection of closed Cartesian subsets of the form

$$C = \Theta_0 \times \left(\bigtimes_{i \in I} C_i\right) \subseteq \Theta \times A,$$

[8] The mathematical meaning of "closed" in the lemma is clear when $\Theta_0 \times \Theta_{-i} \times A_{-i}$ is finite, which implies that $\Delta\left(\Theta_0 \times \Theta_{-i} \times A_{-i}\right)$ is a compact (and convex) subset of $\mathbb{R}^{\Theta_0 \times \Theta_{-i} \times A_{-i}}$. For the more general compact-continuous case see Section 4.6 and, in particular, Lemma 10.

where $C_i \subseteq \Theta_i \times A_i$ for each $i \in I$, and let

$$C_{0,-i} = \Theta_0 \times \left(\bigtimes_{j \neq i} C_j \right).$$

With this, define the **justification operator** $\rho : \mathcal{C} \to \mathcal{C}$ as follows: for each $C \in \mathcal{C}$,

$$\begin{aligned} \rho_i(C_{0,-i}) &= \left\{ (\theta_i, a_i) \in \Theta_i \times A_i : \exists \mu^i \in \Delta(C_{0,-i}), a_i \in r_i(\mu^i, \theta_i) \right\}, \\ \rho(C) &= \Theta_0 \times \left(\bigtimes_{i \in I} \rho_i(C_{0,-i}) \right). \end{aligned}$$

Intuitively, we keep the whole residual uncertainty space Θ_0 as is, because strategic thinking can rule out some pairs (θ_i, a_i) for some player i, but cannot rule out any possible θ_0. Formally, one can reinterpret θ_0 as the information-type of an *inactive* player 0, so that A_0 is a singleton.

Remark 20. *ρ is* monotone*, that is, $E \subseteq F$ implies $\rho(E) \subseteq \rho(F)$ for all $E, F \in \mathcal{C}$; therefore, the sequence $\left(\rho^k(\Theta \times A)\right)_{k=1}^{\infty}$ is (weakly) decreasing.*

With this, we can extend the definition of rationalizability from games with complete information to games with payoff uncertainty.[9]

Definition 36. *A profile $\left(\theta_0, (\theta_i, a_i)_{i \in I}\right)$ is **rationalizable** if*

$$\left(\theta_0, (\theta_i, a_i)_{i \in I}\right) \in \rho^{\infty}(\Theta \times A) = \bigcap_{k \in \mathbb{N}} \rho^k(\Theta \times A);$$

*an action a_i is **rationalizable for information-type** θ_i if pair (θ_i, a_i) is part of a rationalizable profile, that is,*

$$(\theta_i, a_i) \in \operatorname{proj}_{\Theta_i \times A_i} \rho^{\infty}(\Theta \times A).$$

A proper formalization of events (like rationality, R) and belief operators (like $\mathrm{B}_i(\cdot)$ and $\mathrm{B}(\cdot)$) would allow to derive the following table, where it is implicitly assumed that there is common knowledge of the game with payoff uncertainty $\hat{G}$.

[9] See Battigalli [8].

Assumptions about rationality and beliefs	Behavioral implications
R	$\rho(\Theta \times A)$
$R \cap \mathrm{B}(R)$	$\rho^2(\Theta \times A)$
$R \cap \mathrm{B}(R) \cap \mathrm{B}^2(R)$	$\rho^3(\Theta \times A)$
...	...
$R \cap \left(\bigcap_{k=1}^{K} \mathrm{B}^k(R)\right)$	$\rho^{K+1}(\Theta \times A)$
...	...

In Example 31 $\rho^3(\Theta \times A) = \{(\theta^a, a), (\theta^b, b)\} \times \{d\}$. In Example 32 $\rho^K(\Theta \times A) = (\Theta_1 \times \{N\}) \times (\Theta_2 \times \{N\})$ (after K steps, where K is the number of possible prizes, only the pairs (θ_i, N) survive).

The results about rationalizability in games with complete information can be extended to games with payoff uncertainty:

Theorem 28. *If $\hat{G}$ is a finite or compact-continuous game with payoff uncertainty, then*

$$\mathrm{proj}_{\Theta}\rho^k(\Theta \times A) = \Theta$$

for all $k \in \mathbb{N} \cup \{\infty\}$, and

$$\rho(\rho^\infty(\Theta \times A)) = \rho^\infty(\Theta \times A);$$

furthermore, for every $C \in \mathcal{C}$,

$$C \subseteq \rho(C) \Rightarrow C \subseteq \rho^\infty(\Theta \times A).$$

The first part follows from Lemma 22 and says that we cannot eliminate any information-type θ_i, we can only eliminate pairs (θ_i, a_i) such that a_i is not rationalizable given θ_i; this implies that *the set of rationalizable actions is nonempty for every information-type of every player.* The second part says that restarting the iterations from the set of rationalizable profiles cannot further reduce it. The third part says that every set $C \in \mathcal{C}$ with the best reply property is contained in the set of rationalizable profiles. The proof is a rather straightforward extension of the proofs of Theorems 2 and 3.

Furthermore, also the equivalence between rationalizability and iterated dominance can be extended to the payoff-uncertainty setup. For any $C \in \mathcal{C}$, let $C_{i,\theta_i} = \{a_i \in A_i : (\theta_i, a_i) \in C_i\}$ denote the θ_i-section of C_i.

Definition 37. *Fix $C \in \mathcal{C}$, $i \in I$, $a_i \in A_i$, $\alpha_i \in \Delta(A_i)$ in a game with payoff uncertainty $\hat{G}$; mixed action α_i* ***dominates*** *a_i* ***given*** *θ_i* ***within*** *C, written $\alpha_i \gg_{(\theta_i,C)} a_i$, if* $\text{supp}\alpha_i \subseteq C_{i,\theta_i}$ *and*

$$\forall(\theta_0, \theta_{-i}, a_{-i}) \in \Theta_0 \times C_{-i}, \, \mathbb{E}_{\alpha_i}(u_{i,\theta_0,\theta_i,\theta_{-i},a_{-i}}) > u_i(\theta_0, \theta_i, \theta_{-i}, a_i, a_{-i}),$$

where $u_{i,\theta,a_{-i}} : A_i \to \mathbb{R}$ is the section of i's parameterized payoff function at $(\theta, a_{-i}) = (\theta_0, \theta_i, \theta_{-i}, a_{-i})$.

When mixed action α_i has a finite support, the expected utility formula is

$$\mathbb{E}_{\alpha_i}(u_{i,\theta_0,\theta_i,\theta_{-i},a_{-i}}) = \sum_{a'_i \in \text{supp}\alpha_i} u_i\left(\theta_0, \theta_i, \theta_{-i}, a'_i, a_{-i}\right) \alpha_i\left(a'_i\right).$$

This is like the standard notion of dominance by a mixed action (within a restricted set of possible profiles), but it is extended to take into account i's private information and i's uncertainty about (θ_0, θ_{-i}). Next define the set of undominated profiles given $C \in \mathcal{C}$, that is $\text{ND}(C)$, as follows:

$$\begin{aligned} \text{ND}_i(C) &= C_i \backslash \left\{(\theta_i, a_i) \in C_i : \exists \alpha_i \in \Delta(C_{i,\theta_i}), \alpha_i \gg_{(\theta_i,C)} a_i\right\}, \\ \text{ND}(C) &= \Theta_0 \times \left(\bigtimes_{i \in I} \text{ND}_i(C) \right). \end{aligned}$$

As in the complete-information case, $\text{ND} : \mathcal{C} \to \mathcal{C}$ is *not monotone.* Lemma 2 can be restated as follows:

Lemma 23. *Let $\hat{G}$ be a finite or compact-continuous game with payoff uncertainty. For each $C \in \mathcal{C}$, $i \in I$ and $(\theta_i, a_i^*) \in C_i$, the following statements are equivalent:*
(1) there is no *mixed action α_i such that $\alpha_i \gg_{(\theta_i,C)} a_i^*$,*
(2) there is a conjecture $\mu^i \in \Delta(\Theta_0 \times C_{-i})$ such that $a_i^* \in \arg\max_{a_i \in C_{i,\theta_i}} \mathbb{E}_{\mu^i}(u_{i,\theta_i,a_i})$.

By a straightforward extension of the methods of Chapter 4, one can use Lemma 23 to prove the following result:

Theorem 29. *If $\hat{G}$ is a finite or compact-continuous game with payoff uncertainty, then $\rho^k(\Theta \times A) = \text{ND}^k(\Theta \times A)$ for all $k \in \mathbb{N} \cup \{\infty\}$.*

Therefore, a profile of information-types and actions $((\theta_i, a_i))_{i\in I}$ is rationalizable if and only if it survives the iterated elimination of those pairs (θ_i, a_i) such that a_i is dominated given θ_i within the set of profiles that survived the previous rounds of elimination.

Finally, in the case of *nice* games with payoff uncertainty, rationalizability is characterized by the iterated deletion of pairs (θ_i, a_i) such that either a_i is not a best reply to any deterministic conjecture, or a_i is dominated by another pure action. We can formally state this result by extending the definition of operators r and ND_p from games with complete information (cf. Section 4.5) to games with payoff uncertainty. For each $C \in \mathcal{C}$, let $\mathrm{r}_i(C_{0,-i})$ denote the set of pairs (θ_i, a_i) such that a_i is a best reply for θ_i to a deterministic conjecture in $C_{0,-i}$, that is,

$$\mathrm{r}_i(C_{0,-i}) = \bigcup_{\theta_i\in\Theta_i} r_i\left(C_{0,-i}\times\{\theta_i\}\right)$$

$$= \left\{(\theta_i, a_i)\in\Theta_i\times A_i : \exists\left(\theta_0, (\theta_j, a_j)_{j\neq i}\right)\in C_{0,-i}, a_i\in r_i\left(\left(\theta_0, (\theta_j, a_j)_{j\neq i}\right), \theta_i\right)\right\},$$

where (with the usual abuse of notation) we identify $C_{0,-i}\subseteq\Theta_0\times\Theta_{-i}\times A_{-i}$ with the set of Dirac measures on $\Theta_0\times\Theta_{-i}\times A_{-i}$ supported by points in $C_{0,-i}$, and let

$$\mathrm{r}\left(C\right) = \Theta_0\times\left(\bigtimes_{i\in I}\mathrm{r}_i\left(C_{0,-i}\right)\right).$$

The iterations of operator (self-map) r give rise to an extension of the **point rationalizability** concept to games with payoff uncertainty. For each $C\in\mathcal{C}$, let

$$\mathrm{ND}_{p,i}(C) = C_i\backslash\left\{(\theta_i, a_i)\in C_i : \exists a_i'\in C_{i,\theta_i}, a_i'\gg_{(\theta_i,C)} a_i\right\}$$

and

$$\mathrm{ND}_p(C) = \Theta_0\times\left(\bigtimes_{i\in I}\mathrm{ND}_{p,i}(C)\right).$$

The iterations of operator ND_p (self-map) yield a simple notion of iterated dominance by pure actions. With these definitions of the operators r and ND_p, we have:

Theorem 30. *If $\hat{G}$ is a nice game with payoff uncertainty, then*

$$\mathrm{r}^k(\Theta \times A) = \rho^k(\Theta \times A) = \mathrm{ND}_p^k(\Theta \times A)$$

for all $k \in \mathbb{N} \cup \{\infty\}$.

The proof is based on a straightforward extension of Lemma 7.

8.3 Directed Rationalizability

According to the previous analysis of rationalizability, restrictions on players' conjectures are only derived from strategic thinking, while restrictions on players' beliefs about the parameter profile θ do not play any role. To see this more formally, let $\rho_j^\infty(\Theta \times A) = \mathrm{proj}_{\Theta_j \times A_j} \rho^\infty(\Theta \times A)$ denote the set of rationalizable pairs (θ_j, a_j) of player j. By Theorem 28, an action a_i is rationalizable for information-type θ_i if and only if it is a best reply to some conjecture μ^i that assigns probability 1 to $\Theta_0 \times \left(\times_{j \neq i} \rho_j^\infty(\Theta \times A)\right)$, where $\mathrm{proj}_{\Theta_j} \rho_j^\infty(\Theta \times A) = \Theta_j$ for each co-player j (for each type, the set of rationalizable actions is nonempty). Therefore, the marginal beliefs about residual uncertainty and co-players' information-types are not in any way restricted by rationalizability, the restrictions concern only the relationship between information-types and actions implied by rationality and common belief in rationality.

Yet, when we analyze strategic interaction in environments with incomplete information, it may be natural to assume that some contextual restrictions on players' beliefs about θ hold, and that it is common belief that they hold. For example, when two sellers compete they may have statistical information about the distribution of valuations among buyers. If this statistical information is common, we obtain a common restriction on sellers' beliefs about buyers' valuations.[10]

To simplify our informal terminology, let us say that some event E is **transparent** if E is true and there is common belief of E, in symbols, if $E \cap \mathrm{CB}(E)$ is the case. We are interested in characterizing the behavioral implications of rationality and common belief in rationality under the assumption that some features of players' conjectures are transparent.

[10] If buyers are modeled as non-strategic agents rather than players, this is a common restriction on beliefs about θ_0.

Specifically, we present an extension of the notion of rationalizability for games with payoff uncertainty whereby some restrictions on players' conjectures are posited and assumed to be transparent, thereby "directing" the result of the solution procedure toward a subset of outcomes. For this reason, we call this more flexible approach "directed rationalizability." To gain intuition on what we mean and how this procedure works, we first look at an example.

Example 33. Consider a modification of the game in Example 31. There are two possible payoff functions for each player i, $u_i(\theta^a, \cdot)$ and $u_i(\theta^b, \cdot)$. Player 1 (Rowena) knows the true payoff functions, while player 2 (Colin) does not know them. We represent this situation with two matrices, corresponding to the two possible values of the parameter θ.

θ^a	c	d
a	4,0	2,1
b	3,1	1,0

θ^b	c	d
a	2,1	0,0
b	0,0	1,1

In this game,

$$\rho^{\infty}(\Theta \times A) = \rho^{1}(\Theta \times A) = \left\{(\theta^a, a), (\theta^b, a), (\theta^b, b)\right\} \times \{c, d\}.$$

Indeed, action a is dominant for information-type θ^a, while both a and b are justifiable for θ^b. Moreover, both actions for Colin are justifiable. The procedure ends at the first step because, as it can be verified, action c (resp. d) is optimal for Colin if he thinks that θ^b (resp. θ^a) is more likely than θ^a (resp. θ^b).

Let $p = \mu^2(\{\theta^a\} \times A_1)$ denote the probability that Colin assigns to information-type θ^a. Assume now that it is transparent that $p > \frac{1}{2}$. Given such restriction, consider the following epistemic assumptions (we are not making the transparency of $p > \frac{1}{2}$ explicit in the notation): R, $R \cap \mathrm{B}(R)$ and $R \cap \mathrm{B}(R) \cap \mathrm{B}^2(R)$.

R implies that Rowena chooses a if $\theta = \theta^a$ (a is dominant in matrix θ^a).

$R \cap \mathrm{B}(R)$ implies that Colin chooses d. Indeed, $\mathrm{B}_2(R_1)$ implies that Colin is certain that Rowena would choose a if $\theta = \theta^a$. Hence, Colin assigns probability zero to (θ^a, b), which implies $\mu^2(\theta^a, a) = p$. Since $p > \frac{1}{2}$, action d yields a strictly higher expected payoff than c. To see this, let $q = \mu^2\left(\theta^b, b\right)$ denote the probability that Colin assigns to the pair (θ^b, b).

The expected payoff of choosing d is

$$\mu^2(\theta^a, a) + \mu^2\left(\theta^b, b\right) = p + q \geq p > \frac{1}{2},$$

and the expected payoff of choosing c is

$$\mu^2\left(\theta^b, a\right) \leq 1 - p < \frac{1}{2}.$$

Since it is transparent that $p > \frac{1}{2}$, $R \cap \mathrm{B}(R) \cap \mathrm{B}^2(R)$ implies that Rowena expects d and she chooses b if $\theta = \theta^b$.
Summing up, the aforementioned assumptions imply that the outcome is (a, d) if $\theta = \theta^a$ and (b, d) if $\theta = \theta^b$. ▲

The example has two features that make it very simple. First, there is no residual uncertainty and only one player is uninformed. Second, we only posited a restriction on beliefs about the co-player's information-type. It follows that there is no room for a possible dependence of the posited belief restriction on the own information-type of a player: in Example 33 there is only one possible information-type of Colin. Yet, there are reasons to make the analysis more flexible. For example, in many economic applications a player's information-type represents private information about an unknown statistical distribution. To consider a more specific example, when a couple splits and they have to agree on who is going to keep their house and pay a compensation to the other, it may be common belief who values the house more among the two of them, but not exactly by how much. This can be modeled as a type-dependent restriction on beliefs about θ: let i be the high-valuation player, if i's type/valuation is θ_i, then player i is certain that $\theta_{-i} < \theta_i$.

It may also be reasonable to posit joint restrictions concerning the subjective probabilities assigned to information-types *and* actions of co-players, such as (without residual uncertainty) the independence restriction

$$\mu^i(\theta_{-i}, a_{-i}) = \prod_{j \neq i} \mu^i(\theta_j, a_j).$$

Next we present an example of a different kind. It may be plausible to assume that the following is transparent in a two-person game with no residual uncertanity about θ: player i assigns strictly positive probability to every type of j, and if a certain action a_j^* is *weakly* dominant for a type θ_j^*, then player i assigns strictly positive probability to a_j^* given θ_j^*. Thus, it is also transparent that $\mu^i\left(\theta_j^*, a_j^*\right) = \mu^i\left(a_j^*|\theta_j^*\right)\mu^i\left(\theta_j^*\right) > 0$. This seemingly innocuous assumption may have a significant impact.

Example 34. Consider the Envelope Game of Example 32 *without transaction costs*, that is, with $\varepsilon = 0$. Suppose that the only transparent features of players' conjectures are as follows: (a) if i's prize (information-type) is θ_i, then player i assigns strictly positive probability to every prize $\theta_{-i} \leq \theta_i$ of the co-player, and (b) if the co-player had the lowest prize, absent transaction costs, he would offer to exchange (weakly dominant action) with strictly positive probability. Then, an "unraveling" argument quite similar to the one offered in Example 32 shows that no type above the lowest offers to exchange. Note that the assumed restriction is plausible and consistent with strategic reasoning, but it does not *derive* from strategic reasoning. Indeed, if the posited restrictions are transparent, rationality and common belief in rationality imply the following: if a player has the lowest prize, then he is certain that—whatever he does—he ends up with the same prize; hence, he is indifferent. As a consequence, offering to exchange would be rationalizable for every type, because a player could assign probability 1 to the event that the co-player does not exchange. ▲

With this in mind, for each information-type θ_i of each player i, we posit a subset of conjectures in $\Delta(\Theta_0 \times \Theta_{-i} \times A_{-i})$. To simplify the analysis, we assume that all the sets Θ_j ($j \in I \cup \{0\}$) and A_i ($i \in I$) are *finite*.[11] For each $i \in I$ and $\theta_i \in \Theta_i$, we let $\Delta_{i,\theta_i} \subseteq \Delta(\Theta_0 \times \Theta_{-i} \times A_{-i})$ denote the restricted set of conjectures for information-type θ_i of player i. We maintain throughout this section that $\Delta_{i,\theta_i} \neq \emptyset$ for each i and θ_i. To illustrate, in Example 33 player 2 has no private information and

$$\Delta_2 = \left\{\mu^2 \in \Delta\left(\Theta_1 \times A_1\right) : \mu^2\left(\{\theta^a\} \times A_1\right) > \frac{1}{2}\right\},$$

whereas $\Delta_{1,\theta^a} = \Delta_{1,\theta^b} = \Delta\left(\{c,d\}\right)$ (no restrictions for any information-

[11] This simplification is relatively innocuous. The analysis can be extended to compact-continuous games with a modicum of additional technicalities.

type of player 1); in Example 34

$$\Delta_{i,\theta_i} = \left\{ \mu^i : \begin{array}{cc} \forall \theta_{-i} \in \{1, ..., \theta_i\}, & \mu^i(\{\theta_{-i}\} \times A_{-i}) > 0, \\ & \mu^i(1, \text{OE}) > 0 \end{array} \right\}.$$

For a given profile $\Delta = (\Delta_{i,\theta_i})_{i \in I, \theta_i \in \Theta_i}$, define the **$\Delta$-justification operator** $\rho_\Delta : \mathcal{C} \to \mathcal{C}$ as follows: for each $C \in \mathcal{C}$,

$$\begin{aligned} \rho_{i,\Delta}(C_{0,-i}) &= \left\{ (\theta_i, a_i) \in \Theta_i \times A_i : \exists \mu^i \in \Delta(C_{0,-i}) \cap \Delta_{i,\theta_i}, a_i \in r_i(\mu^i, \theta_i) \right\}, \\ \rho_\Delta(C) &= \Theta_0 \times \left(\bigtimes_{i \in I} \rho_{i,\Delta}(C_{0,-i}) \right). \end{aligned}$$

It can be verified that, for any fixed Δ, the operator ρ_Δ is *monotone*, that is, $E \subseteq F$ implies $\rho_\Delta(E) \subseteq \rho_\Delta(F)$ for all $E, F \in \mathcal{C}$. There is also another noteworthy form of monotonicity. Consider two profiles $\Delta = (\Delta_{i,\theta_i})_{i \in I, \theta_i \in \Theta_i}$ and $\Delta' = \left(\Delta'_{i,\theta_i}\right)_{i \in I, \theta_i \in \Theta_i}$, and write $\Delta \subseteq \Delta'$ if $\Delta_{i,\theta_i} \subseteq \Delta'_{i,\theta_i}$ for each $i \in I$ and $\theta_i \in \Theta_i$. The proof of the following statement is left as an exercise for the reader.

Remark 21. *Fix two profiles of restrictions, Δ and Δ'. If $\Delta \subseteq \Delta'$, then $\rho_\Delta(C) \subseteq \rho_{\Delta'}(C)$ for each $C \in \mathcal{C}$.*

For a fixed profile Δ, we define the k-th iteration of ρ_Δ recursively. For each $C \in \mathcal{C}$, let $\rho^0_\Delta(C) = C$, and, for each $k \geq 1$, let $\rho^k_\Delta(C) = \rho_\Delta(\rho^{k-1}_\Delta(C))$.

Definition 38. *Fix a profile $\Delta = (\Delta_{i,\theta_i})_{i \in I, \theta_i \in \Theta_i}$ of restricted sets of conjectures. A profile $\left(\theta_0, (\theta_i, a_i)_{i \in I}\right)$ is **Δ-rationalizable** if*

$$\left(\theta_0, (\theta_i, a_i)_{i \in I}\right) \in \rho^\infty_\Delta(\Theta \times A) = \bigcap_{k \in \mathbb{N}} \rho^k_\Delta(\Theta \times A);$$

*an action a_i is **Δ-rationalizable for information-type** θ_i if the pair (θ_i, a_i) is part of a Δ-rationalizable profile, that is,*

$$(\theta_i, a_i) \in \text{proj}_{\Theta_i \times A_i} \rho^\infty_\Delta(\Theta \times A).$$

First note that we obtain the previous definition of rationalizability for games with payoff uncertainty when we let $\Delta_{i,\theta_i} = \Delta(\Theta_0 \times \Theta_{-i} \times A_{-i})$ for every i and θ_i, that is, when no restrictions are assumed. Second, the set

$\rho_{\Delta}^{\infty}(\Theta\times A)$ of Δ-rationalizable profiles depends on the specific profile Δ of restricted sets; different specifications of Δ may yield different solution sets. In Example 33, if it is transparent that Colin deems θ^a strictly more likely than θ^b, the solution is

$$\rho_{\Delta}^{\infty}(\Theta\times A)=\rho_{\Delta}^{3}(\Theta\times A)=\left\{(\theta^a,\mathrm{a}),\left(\theta^b,\mathrm{b}\right)\right\}\times\{\mathrm{d}\};$$

if instead Colin may *also* deem θ^b at least as likely as θ^a (that is, if there are no restrictions for Colin), the solution is

$$\rho_{\Delta}^{\infty}(\Theta\times A)=\rho_{\Delta}^{1}(\Theta\times A)=\left\{(\theta^a,\mathrm{a}),\left(\theta^b,\mathrm{a}\right),\left(\theta^b,\mathrm{b}\right)\right\}\times\{\mathrm{c},\mathrm{d}\}.$$

This inclusion result (with respect to the restrictions represented by Δ) is not a coincidence: the reader can use Remark 21 to prove by induction the following statement.

Remark 22. *Fix two profiles of restrictions,* Δ *and* Δ'*. If* $\Delta\subseteq\Delta'$*, then* $\rho_{\Delta}^{\infty}(\Theta\times A)\subseteq\rho_{\Delta'}^{\infty}(\Theta\times A)$.

We call "directed rationalizability" the map $\Delta\mapsto\rho_{\Delta}^{\infty}(\Theta\times A)$ when we refer to the final result, and the map $\Delta\mapsto\left(\rho_{\Delta}^{k}(\Theta\times A)\right)_{k=1}^{\infty}$ when we refer to the iterated elimination procedure.

It can be shown that, for each profile Δ, $\rho_{\Delta}^{\infty}(\Theta\times A)$ charaterizes the behavioral implications of rationality and common belief in rationality when it is transparent that players' conjectures satisfy the restrictions represented by Δ. Of course, since ρ_{Δ} is monotone, the sequence $\left(\rho_{\Delta}^{k}(\Theta\times A)\right)_{k=1}^{\infty}$ is (weakly) decreasing. Thus, the abstract analysis of directed rationalizability is similar to the analysis of the simpler "undirected" version. But there is a conceptual difference: if the (transparent) restrictions represented by Δ also concern conjectures about co-players' *behavior*, then they may turn out to be inconsistent with implications of strategic thinking, that is, of the assumptions R, $R\cap\mathrm{B}(R)$, $R\cap\mathrm{B}(R)\cap\mathrm{B}^2(R)$,...; then, the set of Δ-rationalizable actions is empty. This is the case in the Envelope game of Example 32 (i.e., with positive transaction costs) if it is transparent that the lowest type offers to exchange with strictly positive probability, because rationality and common belief in rationality imply that *no* type offers to exchange. We now show that if, instead, the Δ-restrictions only concern beliefs about θ, then the solution is nonempty. This requires some preliminary notations and definitions.

Recall that, given a probability measure μ on a space $X \times Y$, we let $\mathrm{marg}_X \mu$ denote its marginal on X, that is, $\mathrm{marg}_X \mu(E) = \mu(E \times Y)$ for every $E \subseteq X$. We say that $\Delta = (\Delta_{i,\theta_i})_{i \in I, \theta_i \in \Theta_i}$ is a profile of **restrictions on exogenous beliefs**[12] if, for each $i \in I$ and each $\theta_i \in \Theta_i$, there is a nonempty set $\bar{\Delta}_{i,\theta_i} \subseteq \Delta(\Theta_0 \times \Theta_{-i})$ such that

$$\Delta_{i,\theta_i} = \left\{ \mu^i \in \Delta(\Theta_0 \times \Theta_{-i} \times A_{-i}) : \mathrm{marg}_{\Theta_0 \times \Theta_{-i}} \mu^i \in \bar{\Delta}_{i,\theta_i} \right\}.$$

With this, we obtain the following extension of Theorem 28.

Theorem 31. *Let $\hat{G}$ be a finite game with payoff uncertainty. Consider a profile $\Delta = (\Delta_{i,\theta_i})_{i \in I, \theta_i \in \Theta_i}$ of restrictions on* exogenous *beliefs. Then*

$$\mathrm{proj}_\Theta \rho^k_\Delta (\Theta \times A) = \Theta$$

for all $k \in \mathbb{N} \cup \{\infty\}$, and

$$\rho_\Delta (\rho^\infty_\Delta (\Theta \times A)) = \rho^\infty_\Delta (\Theta \times A);$$

furthermore, for every $C \in \mathcal{C}$,

$$C \subseteq \rho_\Delta (C) \Rightarrow C \subseteq \rho^\infty_\Delta (\Theta \times A).$$

Like the proof of Theorem 28, also the proof of Theorem 31 is left to the reader as an exercise. The key additional step is to show (by induction) that, for each step k and for all i and θ_i there is at least one conjecture $\mu^i \in \Delta_{i,\theta_i}$ such that $\mu^i \left(\Theta_0 \times \left(\bigtimes_{j \neq i} \mathrm{proj}_{\Theta_j \times A_j} \rho^{k-1}_\Delta (\Theta \times A) \right) \right) = 1$. This is so because the latter condition only concerns conjectures about the behavior of co-players' *given* their information-types, while the former only restricts the probabilities of their information-types.

8.4 Equilibrium and Beliefs: The Simplest Case

The directed rationalizability approach *allows for* the introduction of some restrictions on exogenous beliefs, i.e., beliefs about parameters or

[12] Recall that we call "exogenous" something that we are not trying to explain as game theorists: we are not trying to explain (θ_0, θ_{-i}) nor—more generally—any beliefs about exogenous things.

exogenous variables. In this section and the next we argue that an extension of the traditional Nash equilibrium approach *requires* the precise specification of exogenous beliefs.

In examples 31 and 32 the repeated elimination of dominated actions selects a unique outcome for each θ. In the complete information case we proved that if there is a unique rationalizable profile, this must be the unique equilibrium. Also in the present incomplete-information context we can interpret the rationalizable mapping from information-types to actions, *if* it is *unique*, as a unique equilibrium of the following kind. For every player i and every information-type θ_i there is a corresponding choice, which in general can be denoted by $\sigma_i(\theta_i)$—in Example 31 $\sigma_1(\theta^a) = a$, $\sigma_1(\theta^b) = b$, and $\sigma_2 = d$; in Example 32 $\sigma_i(\theta_i) = \text{N}$. The unique profile of rationalizable decision functions $\sigma = (\sigma_i : \Theta_i \to A_i)_{i \in I}$ has the following property: there exists no player i, no information-type θ_i and no belief over $\Theta_0 \times \Theta_{-i}$, such that i has an incentive to deviate from $\sigma_i(\theta_i)$ *provided that i expects every other player j to act according to the decision function* $\sigma_j(\cdot)$; furthermore, σ is the unique profile of decision functions with this property.

A decision function of player j can be interpreted as the conjecture of player i about how j would behave as a function of his information-type θ_j. Then, one can tentatively define an equilibrium in this context as a profile of decision functions such that every player, given his information-type, chooses a best reply to his conjecture about other players, and furthermore this conjecture is correct, i.e., it corresponds to how each co-player chooses as a function of his own information-type.

The problem of this tentative definition is that it is silent on the belief of each player i over $\Theta_0 \times \Theta_{-i}$. But in many games it is not possible to ascertain whether a profile of decision functions $(\sigma_i)_{i \in I}$ has the property that no player i has incentives to unilaterally deviate from σ_i (i.e., what the others expect him to choose) without further assumptions about players' beliefs about θ. This is shown by the following example.

Example 35. Consider the following variation of game $\hat{G}^1$ (Rowena knows the true payoff matrix, Colin does not):

How can we determine whether d is a best reply to the decision function $(\sigma_1(\theta^a), \sigma_1(\theta^b)) = (a,b)$? It all depends on the probability that player 2 (Colin) assigns to θ^a and θ^b. If Colin thinks that θ^a is more likely than

$\hat{G}^2$:

θ^a	c	d
a	$4,0$	$2,1$
b	$3,1$	$1,0$

θ^b	c	d
a	$1,1$	$0,0$
b	$0,1$	$2,0$

θ^b, $\mathbb{P}_2(\theta^a) \geq \frac{1}{2}$, then d is a best reply, otherwise it is not. Suppose that $\mathbb{P}_2(\theta^a) < \frac{1}{2}$. In an "equilibrium" (σ_1, σ_2), we must have $\sigma_1(\theta^a) = a$, because a is dominant for information-type θ^a of Rowena.[13] In matrix θ^b Colin's payoff is independent of Rowena's action; thus, if $\mathbb{P}_2(\theta^a) < \frac{1}{2}$, whatever the value of $\sigma_1(\theta^b)$, Colin's best reply to the decision function σ_1 is c (indeed, c yields expected payoff $1-\mathbb{P}_2(\theta^a) > \frac{1}{2}$, while d yields expected payoff $\mathbb{P}_2(\theta^a) < \frac{1}{2}$). Rowena's best response to c in matrix θ^b is a. Then, if $\mathbb{P}_2(\theta^a) < \frac{1}{2}$, the equilibrium profile should be $\sigma_1(\theta^a) = \sigma_1(\theta^b) = a$, $\sigma_2 = c$. ▲

The example shows that equilibrium analysis requires a specification of beliefs about other players' information-types. In general, *in order to define equilibrium behavior, it is necessary to enrich the mathematical structure $\hat{G}$ with the probabilistic beliefs of every player i about the residual uncertainty parameter θ_0 and his co-players' information-types* $\theta_{-i} = (\theta_j)_{j\in I\setminus\{i\}}$. We denote such beliefs by $p^i \in \Delta(\Theta_0 \times \Theta_{-i})$. It is then natural to ask: How is p^i determined? What does player j $(j \neq i)$ know or believe about p^i? If p^i is just a subjective probability, then it does not seem very plausible to assume that j "knows" p^i, and it is even less plausible to assume that the belief profile $(p^i)_{i\in I}$ is "common knowledge." Then, even though by fixing a belief profile $(p^i)_{i\in I}$ we can determine whether each action $\sigma_i(\theta_i)$ $(\theta_i \in \Theta_i)$ is a best reply to the profile of decision functions $\sigma_{-i} = (\sigma_j)_{j\neq i}$, it is not clear whether this best reply property is sufficient to ensure that there is no incentive to deviate. How can i be confident that j will follow his prescribed choice $\sigma_j(\theta_j)$ if i does not know p^j and therefore cannot know whether each $\sigma_j(\theta_j)$ $(\theta_j \in \Theta_j)$ is in turn a best reply to σ_{-j}?

Actually, if we are to use a rigorous language, the word **"knowledge"** should only be used for a true belief determined by direct observation or logical deduction. As long as we assume that players cannot peer into other players' minds, we have to exclude that they can *know* anything

[13] Of course, "equilibrium" has yet to be defined in this context. The formal definition will follow.

about the beliefs of others. However, a player may know facts that, according to some psychological hypotheses, imply that the co-players' beliefs have some features, or maybe completely pin down the co-players' beliefs about unknown parameters, beliefs about such beliefs, and so on. If the psychological hypotheses are correct, this player's beliefs about the beliefs of others are also correct. In general, as we did in Section 8.3, we say that some event E is **transparent** if E is true and there is common belief of E. Then, the correctly posed question is whether the belief profile $(p^i)_{i\in I}$ is transparent. In the affirmative case, we can carry out a relatively simple equilibrium analysis with incomplete information.

Next, we consider a scenario where it is plausible to assume that $(p^i)_{i\in I}$ is transparent.[14] For each player/role $i \in I$ there is a large population of agents that can play in that role. Each one of them is characterized by an information-type θ_i (for instance, his preferences, his ability, his strength). Let $q_i \in \Delta(\Theta_i)$ denote the statistical distribution of θ_i within population i: $q_i(\theta_i)$ is the fraction of agents in the population i with information-type θ_i. Furthermore, assume that θ_0 is determined through a random experiment whose probabilities are fixed and transparent (for instance, assume that θ_0 depends on the color of a ball extracted from an urn whose composition has been publicly announced in front of all players at the same time); let the probability of θ_0 according to this experiment be $q_0(\theta_0)$. If players are randomly chosen from the corresponding populations, then the probability of meeting opponents characterized by the profile of information-types $\theta_{-i} = (\theta_j)_{j\neq i}$ is $\prod_{j\in I\setminus\{i\}} q_j(\theta_j)$.[15] If these statistical distributions are commonly known, then it is reasonable to assume that, for all i and θ_{-i}, $p^i(\theta_0, \theta_{-i}) = \prod_{j\in\{0\}\cup I\setminus\{i\}} q_j(\theta_j)$ and that the profile of "objective" beliefs $(p^i)_{i\in I}$ is transparent. Thus, we obtain a structure

$$BG^s = \left\langle I, \Theta_0, (\Theta_i, A_i, u_i, p^i)_{i\in I}\right\rangle$$

called **simple Bayesian game with type-independent beliefs** (in the following section we consider a more general definition of Bayesian game). Assuming that all the features of the interactive situation described by

[14] This is what Harsanyi (1967-68) calls the "prior lottery model."

[15] As usual, we are assuming that every Θ_j is finite. The generalization to the infinite case, though conceptually trivial, requires the use of measure theoretic concepts.

BG^s are transparent, we can meaningfully define as **equilibrium** a profile of decision functions $(\sigma_i)_{i\in I} \in \times_{i\in I} A_i^{\Theta_i}$ such that, for each i and θ_i,

$$\sigma_i(\theta_i) \in \arg\max_{a_i\in A_i} \sum_{\theta_0,\theta_{-i}} p^i(\theta_0,\theta_{-i})u_i(\theta_0,\theta_i,\theta_{-i},a_i,\sigma_{-i}(\theta_{-i})), \qquad (8.4.1)$$

where $\sigma_{-i}(\theta_{-i}) = (\sigma_j(\theta_j))_{j\in I\setminus\{i\}}$.[16] Note, for any fixed conjecture σ_{-i} about the decision functions of the co-players, the expected payoff being maximized in eq. (8.4.1) is well defined because we have specified an exogenous belief p^i. Without such specification (for each player), we could not ascertain whether a profile of decision functions $(\sigma_i)_{i\in I}$ is an equilibrium.

Example 36. (*First Price Auctions with independent private values*) A given object is put on sale by means of an auction. The set of participant is $I = \{1, ..., n\}$, the monetary value of the object for participant i is θ_i. Therefore the model has private values. The set of possible valuations is $\Theta_i = [0, 1]$ and the profile $(\theta_1, ..., \theta_n)$ is uniformly distributed on $[0, 1]^n$. (This means that each bidder i believes that his competitors' valuations are mutually independent, uniform random variables.) The object is assigned to the player that makes the highest offer. In case of a draw, it is randomly assigned among the highest bidders. Whoever is given the object has to pay his bid (First Price rule). The resulting payoff function is:

$$u_i(\theta, a) = \begin{cases} (\theta_i - a_i)\frac{1}{|\arg\max_j a_j|}, & \text{if } a_i = \max_j a_j, \\ 0, & \text{if } a_i < \max_j a_j. \end{cases}$$

It turns out that this game has a symmetric equilibrium in which each player bids $\sigma_i(\theta_i) = \frac{n-1}{n}\theta_i$.

How can we guess that these functions form an equilibrium? Consider the following heuristic derivation. If $\theta_i = 0$ it does not make sense (indeed, it is weakly dominated) to offer more than zero. Now assume $\theta_i > 0$. If bidder i conjectures that each competitor bids according to the linear rule $\sigma_j(\theta_j) = k\theta_j$, where $k \in (0, 1)$ *is a coefficient that we have to determine*,

[16]Furthermore, even if the structure BG^s were not transparent (perhaps because the statistical distributions are not commonly known), the definition of equilibrium given in the text could perhaps be motivated as a steady state of an adaptive process. This is certainly the case under private values.

then i believes that, if he bids a_i, the probability of winning the object is

$$\mathbb{P}\left([\forall j \neq i, \sigma_j(\widetilde{\theta}_j) < a_i]\right) = \mathbb{P}\left(\left[\forall j \neq i, \widetilde{\theta}_j < \frac{a_i}{k}\right]\right)$$

($\widetilde{\theta}_j$ denotes the random valuation of j from i's viewpoint; we can neglect ties because, for each competitor j, the probability of event $[\widetilde{a}_j = a_i]$ is zero). Since we assumed independent and uniform distributions of values, we have

$$\mathbb{P}\left(\left[\forall j \neq i, \widetilde{\theta}_j < \frac{a_i}{k}\right]\right) = \begin{cases} \left(\frac{a_i}{k}\right)^{n-1}, & \text{if } \frac{a_i}{k} < 1, \\ 1, & \text{if } \frac{a_i}{k} \geq 1. \end{cases}$$

Thus, if bidder i is rational he offers the smallest of the following numbers: $\arg\max_{0 \leq a_i < k} \left(\frac{a_i}{k}\right)^{n-1} (\theta_i - a_i) = \frac{n-1}{n}\theta_i$ and k,[17] that is

$$a_i = \min\left\{k, \frac{n-1}{n}\theta_i\right\}.$$

In a symmetric equilibrium each player has a correct conjecture about the bidding functions of the competitors and each player has the same (best reply) bidding function, therefore $k = \frac{n-1}{n}$. This implies

$$\min\left\{k, \frac{n-1}{n}\theta_i\right\} = \min\left\{\frac{n-1}{n}, \frac{n-1}{n}\theta_i\right\} = \frac{n-1}{n}\theta_i.$$

Therefore, the optimal bid when the symmetric conjecture of i about each j is $\sigma_j(\theta_j) = \frac{n-1}{n}\theta_j$ turns out to be precisely $a_i = \sigma_i(\theta_i) = \frac{n-1}{n}\theta_i$. ▲

The best method to compute the equilibria of a simple Bayesian game depends on the specific game. In Example 36 we followed a "conjecture and verify" method to compute the symmetric equilibrium. For finite games, one can adopt the method of constructing the so called "strategic form" of the game and compute its Nash equilibria. We describe this method formally in our analysis of general Bayesian games.

[17] To compute the maximizer in $[0, 1)$ use the first order condition:

$$\frac{(n-1)}{k}\left(\frac{a_i}{k}\right)^{n-2}(\theta_i - a_i) - \left(\frac{a_i}{k}\right)^{n-1} = 0.$$

8.5 The General Case: Bayesian Games

In order to provide a more general definition of equilibrium in games with incomplete information it is necessary to consider the case in which the beliefs of a generic player i on $\Theta_0 \times \Theta_{-i}$ are not known to his co-players $-i$. This means that i is not characterized by his private information θ_i alone, and the belief p^i has to be separately specified. From j's point of view the pair $(\theta_i, p^i) \in \Theta_i \times \Delta(\Theta_0 \times \Theta_{-i})$ is unknown. Why should j care about p^i? The problem, from j's point of view, is that j's payoff depends on i's choice a_i which in turn depends on θ_i *and* p^i. Thus, in order to extend the equilibrium concept to this more general case it is necessary to specify j's beliefs about (θ_i, p^i).

According to the subjective formulation, known in economic theory as the **Bayesian approach**, a *decision maker forms probabilistic beliefs over all the relevant variables that are unknown to him.* In our example j has probabilistic beliefs about (θ_i, p^i).

To simplify the exposition, we first focus on the two-person case with distributed knowledge of parameter θ (Θ_0 is a singleton): the only opponent of i is j and *viceversa.* The beliefs of j about θ_i and p^i are therefore a joint probability measure $q^j \in \Delta(\Theta_i \times \Delta(\Theta_j))$ from which belief $p^j \in \Delta(\Theta_i)$ can be recovered by computing the marginal distribution on Θ_i. Beliefs $\left(q^j\right)_{j\in I}$ are called **second-order beliefs**, whereas beliefs $\left(p^j\right)_{j\in I}$ are called **first-order beliefs**: second-order beliefs are beliefs about the primitive uncertainty and the opponent's first-order beliefs. Hence, first-order beliefs can be derived from second-order beliefs.

For some applications there is no need to go beyond second-order beliefs. Suppose, for example that $\Theta_i = \{\theta_i^0, \theta_i^1\}$, $i = 1, 2$. Then we may think of p^i as i's subjective probability of θ_j^1: $\Delta(\Theta_j)$ is isomorphic to $[0, 1]$. Suppose that it is transparent that the second-order beliefs of each player j have the form $q^j = p^j \times \overline{q}^j$ where $\overline{q}^j$ is given by the uniform distribution on $[0, 1]$, which is the set of possible values of the subjective probability $p^i(\theta_j^1)$ (recall: by "transparent" we mean that second-order beliefs do indeed have this form and furthermore it is common belief that they have this form). Then each first-order belief p^j also pins down the second-order belief $q^j = p^j \times \overline{q}^j$ and there is no need to explicitly consider beliefs of higher order.

However, since beliefs are subjective, *we have to allow for more* *general situations where second-order beliefs are not pinned down by first-order beliefs.* Since behavior may depend on second-order beliefs, it is necessary to introduce **third-order beliefs** $r^j \in \Delta(\Theta_i \times \Delta(\Theta_j) \times \Delta(\Theta_j \times \Delta(\Theta_i)))$,[18] from which the first and second-order beliefs can be recovered by marginalization. At this point, the formalism is already quite complex. Furthermore, there is no compelling reason to stop at third-order beliefs. How should we proceed? Is it possible to use a formal and compact representation of strategic interactions with incomplete information which is not too complex, but at the same time allows a representation of beliefs of arbitrarily high order and a meaningful definition of equilibrium? The answer is Yes.

8.5.1 Bayesian Games

The solution relies on the adoption of a more abstract and self-referential approach proposed by John Harsanyi [52].[19] Starting from the game with payoff uncertainty $\langle I, \Theta_0, (\Theta_i, A_i, u_i)_{i \in I} \rangle$, let us consider a richer structure with a set of "states of the world" Ω (assumed *finite* for the sake of simplicity). A **state of the world** ω characterizes each player's knowledge and "interactive beliefs." This can be formalized mathematically introducing functions $\tau_i : \Omega \to T_i$ and $\vartheta_i : T_i \to \Theta_i$ for each player i, a function $\vartheta_0 : \Omega \to \Theta_0$ and a probability measure $p_i \in \Delta(\Omega)$. It is also implicitly assumed that the situation of strategic interaction represented by these elements is transparent.

To grasp the meaning of the above functions we follow Harsanyi and present a *metaphor*.[20] Imagine an *ex ante* state in which all players are equally ignorant. Each player i is endowed with a *prior* (subjective) probability measure $p_i \in \Delta(\Omega)$. Before choosing an action, each player i receives a "signal" t_i about the state of the world. From signal t_i player i infers that $\theta_i = \vartheta_i(t_i)$ and $\omega \in \tau_i^{-1}(t_i) = \{\omega' : \tau_i(\omega') = t_i\} \subseteq \Omega$. Assume

[18] Clearly, the symbol r^j used here is not to be confused with the symbol denoting the best reply correspondence.

[19] This fundamental contribution led to the award to John Harsanyi of the 1994 Nobel Prize for Economics, jointly with John Nash and Reinhard Selten.

[20] This is what Harsanyi calls the "random vector model" of the Bayesian game. The only difference is that such model (not his general analysis of incomplete information) posits a common prior ($p_i = p \in \Delta(\Omega)$ for every $i \in I$), which represents an objective probability distribution.

for simplicity that p_i assigns positive probability to all signals that i could conceivably receive,[21] i.e.,

$$p_i\left(\tau_i^{-1}(t_i)\right) = \sum_{\omega:\tau_i(\omega)=t_i} p_i(\omega) > 0$$

for all $t_i \in T_i$.[22] Then, for each $t_i \in T_i$, the corresponding conditional probability measure $p_i(\cdot|t_i)$ is well defined:

$$\forall E \subseteq \Omega,\ p_i(E|t_i) = \frac{p_i\left(E \cap \tau_i^{-1}(t_i)\right)}{p_i\left(\tau_i^{-1}(t_i)\right)}.$$

With this, the beliefs of player i at state of the world ω are given by probability measure $p_i(\cdot|\tau_i(\omega)) \in \Delta(\Omega)$. Since τ_i and p_i are common knowledge, the function $\omega \mapsto p_i(\cdot|\tau_i(\omega))$ is also common knowledge. It follows that a state of the world determines a player's *beliefs about the parameter θ and about other players' beliefs about it.* To verify this claim, let us focus on the case of *two players and distributed knowledge of* Θ (this is done only to simplify the notation) and let us derive the beliefs of i about t_j given signal t_i:

$$\forall t_j \in T_j,\ p_i(t_j|t_i) = p_i\left(\tau_j^{-1}(t_j)|t_i\right),$$

where $\tau_j^{-1}(t_j) = \{\omega : \tau_j(\omega) = t_j\}$. Next we derive the **first-order beliefs** about θ (since i knows θ_i, his beliefs about θ are determined by his beliefs about Θ_j):

$$\forall \theta_j \in \Theta_j,\ p_i^1(\theta_j|t_i) = p_i\left(\vartheta_j^{-1}(\theta_j)|t_i\right),$$

where $\vartheta_j^{-1}(\theta_j) = \{t_j : \vartheta_j(t_j) = \theta_j\}$. The functions $t_j \mapsto p_j^1(\cdot|t_j) \in \Delta(\Theta_i)$ ($j \in \{1,2\}$) are also common knowledge. Hence, we can derive the **second-order beliefs** about θ, that is, the joint belief of a player about θ and about the opponent's first-order beliefs about θ:

$$\forall(\overline{\theta}_j, \overline{p}_j^1) \in \Theta_j \times \Delta(\Theta_i),\ p_i^2\left(\overline{\theta}_j, \overline{p}_j^1|t_i\right) = \sum_{t_j:\vartheta_j(t_j)=\overline{\theta}_j, p_j^1(\cdot|t_j)=\overline{p}_j^1} p_i(t_j|t_i).$$

---

[21] It can be shown that this comes without substantial loss of generality provided that we allow the subjective priors of different players to be different.

[22] We denote by $p_i(\cdot)$ and $p_i(\cdot|\cdot)$, respectively, the prior and conditional probabilities of events. For instance, $p_i(t_{-i}|t_i)$ is the probability that event $\{\omega : \tau_{-i}(\omega) = t_{-i}\}$ occurs conditional on the event $\{\omega : \tau_i(\omega) = t_i\}$.

[It can be verified that

$$\forall \overline{\theta}_j \in \Theta_j,\, p_i^1\left(\overline{\theta}_j|t_i\right) = \sum_{\overline{p}_j^1 : \exists t_j \in T_j, p_j^1(\cdot|t_j)=\overline{p}_j^1} p_i^2\left(\overline{\theta}_j, \overline{p}_j^1|t_i\right),$$

in words, $p_i^1(\cdot|t_i)$ is the marginal on Θ_j of the joint distribution $p_i^2(\cdot|t_i) \in \Delta(\Theta_j \times \Delta(\Theta_i))$.]

It should be clear by now that it is possible to iterate the argument and compute the functions that assign to each type the corresponding third-order beliefs about θ, fourth-order beliefs about θ, and so on. To sum up, we can conclude that the information and beliefs of all orders of player i about θ are determined by t_i according to the function $t_i \mapsto (\vartheta_i(t_i), p_i^1(\cdot|t_i), p_i^2(\cdot|t_i), ...)$. Signal t_i is called the **type**[23] of player i. Sequence $(p_i^1(\cdot|t_i), p_i^2(\cdot|t_i), ...)$ is called the **hierarchy of beliefs** of type t_i. The information and beliefs of each player i in a given state of the world ω are those corresponding to the type $t_i = \tau_i(\omega)$. The **information-type** of player i, $\theta_i = \vartheta_i(t_i)$, is just one component of his overall type t_i, which also specifies i's beliefs about all the relevant exogenous variables/parameters, θ_{-i}, p_{-i}^1, p_{-i}^2, etc.

Definition 39. *A (finite)* Bayesian Game *is a structure*

$$BG = \langle I, \Omega, \Theta_0, \vartheta_0, (\Theta_i, T_i, A_i, \tau_i, \vartheta_i, p_i, u_i)_{i\in I}\rangle$$

where each set in BG is finite; for every player $i \in I$, $p_i \in \Delta(\Omega)$, $\vartheta_0 : \Omega \to \Theta_0$, $\tau_i : \Omega \to T_i$, $\vartheta_i : T_i \to \Theta_i$, $p_i(t_i) = p_i\left(\tau_i^{-1}(t_i)\right) > 0$ *for all* $t_i \in T_i$,[24] *and* $u_i : \Theta \times A \to \mathbb{R}$.

The analysis of solution concepts for Bayesian games will clarify that *only the beliefs* $p_i(\cdot|t_i)$ $(i \in I,\, t_i \in T_i)$ *really matter.* The priors p_i are just a convenient mathematical tool to represent the beliefs of each type.

In what follows, we will often use the phrase "type t_i chooses a_i" to mean that if player i *were* of type t_i *he would choose* action a_i.

Transparent Restrictions on First-Order Exogenous Beliefs

It is worth noting that Harsanyi's self-referential and implicit representation of hierarchical exogenous beliefs allows—in principle—to model the transparent restrictions on exogenous (first-order) beliefs

[23]The expression "type *à la* Harsanyi" is also used.
[24]This condition implies that $\tau_i : \Omega \to T_i$ is *onto*.

considered in Section 8.3. Fix a game with payoff uncertainty $\hat{G} = \langle I, \Theta_0, (\Theta_i, A_i, u_i)_{i\in I}\rangle$. To obtain a Bayesian game, we append to $\hat{G}$ a belief structure $\langle \Omega, \vartheta_0, (T_i, \tau_i, \vartheta_i, p_i)_{i\in I}\rangle$ as per Definition 39. Fix a profile $\bar{\Delta} = \left(\bar{\Delta}_{i,\theta_i}\right)_{i\in I,\theta_i\in\Theta_i}$, where $\bar{\Delta}_{i,\theta_i} \subseteq \Delta\left(\Theta_0 \times \Theta_{-i}\right)$ for each i and θ_i. With a slight abuse of language, we say that $\bar{\Delta}$ is transparent to mean that it is both true and commonly believed that, if the information-type of player i is θ_i, then his first-order exogenous beliefs belong to set $\bar{\Delta}_{i,\theta_i}$.[25] Belief structure $\langle \Omega, \vartheta_0, (T_i, \tau_i, \vartheta_i, p_i)_{i\in I}\rangle$ yields the transparency of $\bar{\Delta}$ if the implied first-order beliefs satisfy these restrictions, that is, for every $i \in I$, $\theta_i \in \Theta_i$ and $t_i \in \vartheta_i^{-1}(\theta_i)$, we have $p^1(\cdot|t_i) \in \bar{\Delta}_{\theta_i}$. If, additionally, there is some $t_i \in \vartheta_i^{-1}(\theta_i)$ such that $p^1(\cdot|t_i) \in \bar{\Delta}_{\theta_i}$ for every $i \in I$ and $\theta_i \in \Theta_i$, then the given belief structure *exactly* captures the transparency of $\bar{\Delta}$.[26] Our analysis of the connection between Harsanyi's approach and the Directed Rationalizability approach in Section 8.5.5 relies on this observation.

8.5.2 Bayesian Equilibria

In general, players' choices depend not only on their private information on θ, but also more generally on their *types*. In equilibrium, each player has a correct conjecture about the dependence of each opponent's choice on his type, and maximizes his expected payoff given his type:

Definition 40. *A **Bayesian equilibrium** is a profile of decision functions $(\sigma_i : T_i \to A_i)_{i\in I}$ such that*

$$\forall i \in I, \forall t_i \in T_i,\ \sigma_i(t_i) \in \arg\max_{a_i\in A_i} \mathbb{E}_{p_i,\sigma_{-i}}(u_{i,\vartheta_i(t_i),a_i}|t_i),$$

where

$$\begin{aligned} & \mathbb{E}_{p_i,\sigma_{-i}}(u_{i,\vartheta_i(t_i),a_i}|t_i) \\ = & \sum_{\omega\in\Omega} p(\omega|t_i) u_i(\vartheta_0(\omega), \vartheta_i(t_i), \vartheta_{-i}(\tau_{-i}(\omega)), a_i, \sigma_{-i}(\tau_{-i}(\omega))) \end{aligned}$$

[25] We call this a "slight abuse of language" because we defined transparency as a property of *events*, while profile $\bar{\Delta}$ is not in itself an event. However, $\bar{\Delta}$ corresponds to an event concerning players' exogenous first-order beliefs, as explained in the text.

[26] If at least one set $\bar{\Delta}_{\theta_i}$ is infinite, then the beliefs structure has to be infinite as well, which requires some technical conditions.

(for every ω *and* t_{-i}, $\tau_{-i}(\omega) = (\tau_j(\omega))_{j\neq i}$, $\vartheta_{-i}(t_{-i}) = (\vartheta_j(t_j))_{j\neq i}$, $\sigma_{-i}(t_{-i}) = (\sigma_j(t_j))_{j\neq i}$*).*

Note that the expected payoff of type t_i of player i choosing action a_i can be re-written as

$$\mathbb{E}_{p_i,\sigma_{-i}}(u_{i,\vartheta_i(t_i),a_i}|t_i) = \sum_{(\theta_0,t_{-i})\in\Theta_0\times T_{-i}} p(\theta_0,t_{-i}|t_i)u_i(\theta_0,\vartheta_i(t_i),\vartheta_{-i}(t_{-i}),a_i,\sigma_{-i}(t_{-i})),$$

where $p(\theta_0, t_{-i}|t_i)$ is the probability that event

$$\{\omega : \vartheta_0(\omega) = \theta_0,\ \tau_{-i}(\omega) = t_{-i}\}$$

occurs conditional on the event $\{\omega : \tau_i(\omega) = t_i\}$. Again, as we did in Section 8.4, we stress that without the specification of the belief structure (that is, of a specific exogenous belief for each Harsanyi type t_i of each player i), we cannot determine the expected payoff function maximized for each type t_i for some given conjecture σ_{-i} about co-players' decision functions. Thus, such specification is necessary to ascertain whether a profile of decision functions $(\sigma_i : T_i \to A_i)_{i\in I}$ is an equilibrium.

Example 37. We illustrate the concepts of Bayesian game and equilibrium elaborating on game $\hat{G}^2$ of Example 35. Suppose that Rowena does not know Colin's first-order beliefs. From her point of view such beliefs can assign either probability $\frac{1}{3}$ or $\frac{3}{4}$ to θ^a. The two possibilities are regarded as equally likely by Rowena and all this is common knowledge. This situation can be represented by the following Bayesian game:

$$\Omega = \{\alpha, \beta, \gamma, \delta\},$$

$$\Theta_1 \cong \Theta = \left\{\theta^a, \theta^b\right\},\ T_1 = \left\{t_1^a, t_1^b\right\},$$

$$\tau_1(\alpha) = \tau_1(\beta) = t_1^a,\ \tau_1(\gamma) = \tau_1(\delta) = t_1^b,\ \vartheta_1(t_1^a) = \theta^a,\ \vartheta_1(t_1^b) = \theta^b,$$

$$\forall\omega \in \Omega,\ p_1(\omega) = \frac{1}{4},$$

$$\Theta_2 = \left\{\hat{\theta}_2\right\},\ T_2 = \{t_2', t_2''\},$$

$$\tau_2(\alpha) = \tau_2(\gamma) = t_2',\ \tau_2(\beta) = \tau_2(\delta) = t_2'',$$

$$p_2(\alpha) = \frac{3}{8},\ p_2(\beta) = \frac{1}{6},\ p_2(\gamma) = \frac{1}{8},\ p_2(\delta) = \frac{1}{3},$$

	t_2'	t_2''
θ^a, t_1^a	$\alpha, \frac{1}{4}, \frac{3}{8}$	$\beta, \frac{1}{4}, \frac{1}{6}$
θ^b, t_1^b	$\gamma, \frac{1}{4}, \frac{1}{8}$	$\delta, \frac{1}{4}, \frac{1}{3}$

and where the functions u_i are as in $\hat{G}^2$. The probabilistic structure is described in the table below:

To verify that this Bayesian game represents the situation outlined above we compute the following:

$$\begin{aligned} p_2^1\left(\theta^a|t_2'\right) &= p_2\left(t_1^a|t_2'\right) = \frac{p_2(\alpha)}{p_2(\alpha)+p_2(\gamma)} = \frac{3/8}{3/8+1/8} = \frac{3}{4}, \\ p_2^1\left(\theta^a|t_2''\right) &= p_2\left(t_1^a|t_2''\right) = \frac{p_2(\beta)}{p_2(\beta)+p_2(\delta)} = \frac{1/6}{1/6+1/3} = \frac{1}{3}, \\ p_1\left(t_2'|t_1^a\right) &= p_1\left(t_2'|t_1^b\right) = \frac{1}{2}. \end{aligned}$$

This means that in every state of the world Rowena believes that the two events *[Colin assigns probability* $\frac{3}{4}$ *to* θ^a*]* and *[Colin assigns probability* $\frac{1}{3}$ *to* θ^a*]* are equally likely. We now derive the Bayesian equilibria. Let $\sigma = (\sigma_1, \sigma_2)$ be an equilibrium. Since a is dominant when Rowena knows that $\theta = \theta^a$, $\sigma_1(t_1^a) = a$. Hence, the equilibrium expected payoff accruing to type t_2' if he chooses c is

$$p_2\left(t_1^a|t_2'\right) u_2(\theta^a, a, c) + p_2\left(t_1^b|t_2'\right) u_2(\theta^b, \sigma_1(t_1^b), c) = \frac{3}{4}\times 0 + \frac{1}{4}\times 1 = \frac{1}{4};$$

the equilibrium expected payoff for t_2' if he chooses d is

$$p_2\left(t_1^a|t_2'\right) u_2(\theta^a, a, d) + p_2\left(t_1^b|t_2'\right) u_2(\theta^b, \sigma_1(t_1^b), d) = \frac{3}{4}\times 1 + \frac{1}{4}\times 0 = \frac{3}{4}.$$

It follows that in equilibrium $\sigma_2(t_2') = d$. The expected payoffs for the two actions of type t_2'' are

$$\begin{aligned} p_2\left(t_1^a|t_2''\right) u_2(\theta^a, a, c) + p_2\left(t_1^b|t_2''\right) u_2(\theta^b, \sigma_1(t_1^b), c) &= \frac{1}{3}\times 0 + \frac{2}{3}\times 1 = \frac{2}{3}, \\ p_2\left(t_1^a|t_2''\right) u_2(\theta^a, a, d) + p_2\left(t_1^b|t_2''\right) u_2(\theta^b, \sigma_1(t_1^b), d) &= \frac{1}{3}\times 1 + \frac{2}{3}\times 0 = \frac{1}{3}. \end{aligned}$$

Since the maximizing choice for type t_2'' is c, $\sigma_2(t_2'') = c$. We can now determine the equilibrium choice for type t_1^b:

$$\begin{aligned} p_1\left(t_2'|t_1^b\right) u_1(\theta^b, a, d) + p_1\left(t_2''|t_1^b\right) u_1(\theta^b, a, c) &= \frac{1}{2}\times 0 + \frac{1}{2}\times 1 = \frac{1}{2}, \\ p_1\left(t_2'|t_1^b\right) u_1(\theta^b, b, d) + p_1\left(t_2''|t_1^b\right) u_1(\theta^b, b, c) &= \frac{1}{2}\times 2 + \frac{1}{2}\times 0 = 1. \end{aligned}$$

Therefore, $\sigma_1(t_1^b) = b$. ▲

8.5.3 Bayesian Equilibrium and Nash Equilibrium

The concept of Bayesian equilibrium for a game of incomplete information BG can be restated in an equivalent way as a Nash equilibrium of two games with complete information that are associated with the original game: the *ex ante* strategic form and the *interim* strategic form.

The *ex ante* strategic form refers to the metaphor that was previously introduced to explain the elements of the Bayesian game: if there is an *ex ante* stage in which all players are "ignorant," then at such stage each player i can make a contingent plan of action, or strategy, $\sigma_i : T_i \to A_i$, that specifies the action to take for every possible signal t_i that player i might receive. If i believes that the other players follow the strategy profile σ_{-i}, the expected payoff of strategy σ_i is

$$U_i(\sigma_i, \sigma_{-i}) = \mathbb{E}_{p_i,\sigma_i,\sigma_{-i}}(u_i), \tag{8.5.1}$$

where (in a finite game)

$$\begin{aligned} & \mathbb{E}_{p_i,\sigma_i,\sigma_{-i}}(u_i) \\ = & \sum_{\omega\in\Omega} p_i(\omega) u_i(\vartheta_0(\omega), \vartheta_i(\tau_i(\omega)), \vartheta_{-i}(\tau_{-i}(\omega)), \sigma_i(\tau_i(\omega)), \sigma_{-i}(\tau_{-i}(\omega))) \\ = & \sum_{t_i\in T_i} p_i(t_i) \sum_{(\theta_0,t_{-i})\in\Theta_0\times T_{-i}} p_i(\theta_0, t_{-i}|t_i)\, u_i(\theta_0, \vartheta_i(t_i), \vartheta_{-i}(t_{-i}), \sigma_i(t_i), \sigma_{-i}(t_{-i})). \end{aligned}$$

Let $\Sigma_i = A_i^{T_i}$ (recall, this is the set of functions with domain T_i and codomain A_i). The **ex ante strategic form** of BG is the static game

$$\mathcal{AS}(BG) = \langle I, (\Sigma_i, U_i)_{i\in I}\rangle,$$

where U_i is defined by (8.5.1). The reader should be able to prove the following result by inspection of the definitions and using the assumption that each type t_i is assigned positive probability by the prior p_i.

Remark 23. *A profile* $(\sigma_i)_{i\in I}$ *is a Bayesian equilibrium of BG if and only if it is a Nash equilibrium of the game* $\mathcal{AS}(BG)$.

The *interim* strategic form is based on a different metaphor. Assume that for each role i in the game there is a set of potential players T_i. Assume for notational simplicity that $T_i \cap T_j = \emptyset$ for each $i, j \in I$, $i \neq j$ (this is just a matter of labelling). A potential player t_i is characterized by the payoff function $u_i(\vartheta_i(t_i), \cdot, \cdot, \cdot) : \Theta_0 \times \Theta_{-i} \times A \to \mathbb{R}$ and the beliefs $p_i(\cdot, \cdot|t_i) \in \Delta(\Theta_0 \times T_{-i})$. In the event that t_i is selected to play the game in the role of agent i, he will assign probability $p_i(\theta_0, t_{-i}|t_i)$ to the event that the residual uncertainty is θ_0 and that he is facing exactly the profile of "opponents" $t_{-i} = (t_j)_{j\neq i}$. The set of actions available to t_i is A_i, that is, $A_{t_i} = A_i$ ($i \in I$, $t_i \in T_i$). If each potential player $t_j \in T_j$ chooses the action $a_{t_j} \in A_{t_j}$, t_i's expected payoff is computed as follows:

$$\overline{u}_{t_i}((a_{t_j})_{j\in I, t_j\in T_j}) = \mathbb{E}_{p_i, \overrightarrow{a}_{-i}}\left(u_{i,a_{t_i}}|t_i\right), \tag{8.5.2}$$

where $\overrightarrow{a}_{-i} = \left(a_{t_j}\right)_{t_j\in T_j, j\neq i}$ and (in a finite game)

$$\mathbb{E}_{p_i, \overrightarrow{a}_{-i}}\left(u_{i,a_{t_i}}|t_i\right) = \sum_{(\theta_0, t_{-i})\in\Theta_0\times T_{-i}} p_i(\theta_0, t_{-i}|t_i)\, u_i(\theta_0, \vartheta_i(t_i), \vartheta_{-i}(t_{-i}), a_{t_i}, (a_{t_j})_{j\neq i}).$$

The **interim strategic form** of BG is the following static game with $\sum_{i\in I}|T_i|$ players:

$$\mathcal{IS}(BG) = \left\langle \bigcup_{i\in I} T_i, (A_{t_i}, \overline{u}_{t_i})_{i\in I, t_i\in T_i} \right\rangle,$$

where $\overline{u}_{t_i}$ is defined by (8.5.2). The reader should be able to prove the following result by inspection of the definitions:

Remark 24. *A profile* $(\sigma_i^*)_{i\in I}$ *is a Bayesian equilibrium of BG if and only if the corresponding profile* $(a_{t_i}^*)_{i\in I, t_i\in T_i}$ *such that* $a_{t_i}^* = \sigma_i^*(t_i)$ ($i \in I$, $t_i \in T_i$) *is a Nash equilibrium of the game* $\mathcal{IS}(BG)$.

Example 38. Consider the simple Bayesian game of Example 35, where $T_1 = \Theta_1 \cong \Theta = \{\theta^a, \theta^b\} = \Omega$, $\tau_1(\theta) = \theta$ for each θ (player 1 knows θ), $\tau_2(\theta^a) = \tau_2(\theta^b)$ (player 2 does not know θ). To ease notation, assume $p_1 = p_2$ (this does not affect the strategic analysis of the game), and let

$p = p_i(\theta^a)$. The ex ante strategic form of the game is a 4×2 matrix game ($a.a$ means "a for each type," $a.b$ means "a if θ^a and b otherwise," $b.a$ means the opposite, etc.).

$\sigma_1 \backslash \sigma_2$	c	d
$a.a$	$3p+1, 1-p$	$2p, p$
$a.b$	$4p, 1-p$	$2, p$
$b.a$	$2p+1, 1$	$p, 0$
$b.b$	$3p, 1$	$2-p, 0$

If $p < 1/2$, the unique equilibrium of the matrix game is $(\sigma_1, \sigma_2) = (a.a, c)$, which can be obtained by iterated dominance: first note that $a.a$ dominates $b.a$, and $a.b$ dominates $b.b$, then the computation unravels. The interim strategic form is formally a three-person game, where the prior probability $p = p_2$ matters only to compute the payoffs of (the unique type of) player 2. By Remark 24, if $p < 1/2$, this three-person game has a unique equilibrium where both players/types θ^a and θ^b choose a, and player 2 chooses c. Also in this case, the equilibrium can be obtained by iterated dominance: one starts noticing that b is dominated by a for type θ^a, then the computation unravels (see the analysis in Example 35). ▲

In the previous example, iterated dominance in the ex ante strategic form gives the same result as iterated dominance in the interim strategic form. This is not always true. In the Appendix of this chapter, we analyze in detail the relationships between iterated dominance in the ex ante and interim strategic form and different notions of rationalizability for Bayesian games.

8.5.4 Bayesian and Correlated Equilibrium

Even though it may seem counter-intuitive, the definition of Bayesian game admits as a special case that there is common knowledge of the payoff functions and yet there are multiple types for some players. This is the case when there are functions $\bar{u}_i \in \mathbb{R}^A$ $(i \in I)$ such that $u_i(\theta, \cdot) = \bar{u}_i$ for each θ and i.[27] In other words, *we can define a Bayesian game with many*

[27] When the sets Θ_0 and Θ_i are all singletons we obtain a special case of this special case!

types even starting from a game with complete information! An even more special case can be analyzed: the one in which, besides having complete information, all players are characterized by the same prior belief ($p_i = p$ for each $i \in I$), which we refer to as **"common prior."** What can we say about Bayesian equilibria of a game with complete information?

To make the exposition more precise, we have to introduce some terminology and notation: given a complete-information game $G = \langle I, (A_i, \bar{u}_i)_{i\in I}\rangle$ we call **Bayesian elaboration** of G any Bayesian game

$$BG = \langle I, \Omega, \Theta_0, \vartheta_0, (\Theta_i, T_i, A_i, \tau_i, \vartheta_i, p_i, u_i)_{i\in I}\rangle$$

such that, Θ_0 and every Θ_i ($i \in I$) contain a single element, say $\bar{\theta}_0$ and $\bar{\theta}_i$ ($i \in I$). Thus, for every profile $a \in A$, we can write $u_i(\bar{\theta}, a) = \bar{u}_i(a)$. The following remark follows directly from the definitions:

Remark 25. *Let BG be a Bayesian elaboration of a finite game with complete information G and suppose that there is a common prior. Then, every Bayesian equilibrium of BG corresponds to a correlated equilibrium of G.*[28]

8.5.5 Bayesian Equilibrium and Rationalizability

The equivalence between the correlated equilibria of a game with complete information G and the Bayesian equilibria of its Bayesian elaboration BG holds under the assumption that players share a common prior. Removing this common prior assumption, one obtains the notion of subjective correlated equilibrium: a **subjective correlated equilibrium** of a complete-information game G is a Bayesian equilibrium of a Bayesian elaboration of G.

Theorem 32. *An action profile $(a_i)_{i\in I}$ of the finite complete-information game G is rationalizable if and only if it is selected in a subjective correlated equilibrium, that is, if and only if there are a Bayesian elaboration BG of G, an equilibrium $(\sigma_i)_{i\in I}$ of BG, and a state of the world ω in BG such that $a_i = \sigma_i(\tau_i(\omega))$ for every $i \in I$.*

Proof. (**If**) Let σ be an equilibrium of a Bayesian elaboration of $G = \langle I, (A_i, u_i)_{i\in I}\rangle$. To show that every image of σ is a profile of

[28] In a way, one could say that Harsanyi [52] had implicitly defined the correlated equilibrium concept before Aumann [5], but without being aware of it!

rationalizable actions, define the sets $C_i = \sigma_i(T_i) = \sigma_i(\tau_i(\Omega))$ $(i \in I)$. We are going to verify that the cross-product $C = \times_{i\in I} C_i$ has the best reply property $C \subseteq \rho(C)$, which implies that each action profile of $\sigma(\tau(\Omega)) \subseteq C$ is rationalizable in G (Theorem 3).[29] For any i, $a_i^* \in C_i = \sigma_i(T_i)$ and $t_i \in \sigma_i^{-1}(a_i^*)$, let $\mu^{t_i} = p_i(\cdot|t_i) \circ (\tau_{-i} \circ \sigma_{-i})^{-1} \in \Delta(A_{-i})$ denote the conjecture induced by the beliefs on Ω of type t_i given σ_{-i}. More explicitly,

$$\forall a_{-i} \in A_{-i},\ \mu^{t_i}(a_{-i}) = p_i\left((\sigma_{-i} \circ \tau_{-i})^{-1}(a_{-i})\,|t_i\right) = \sum_{\omega:\sigma_{-i}(\tau_{-i}(\omega))=a_{-i}} p_i(\omega|t_i).$$

By construction, $u_i(a_i, \mu^{t_i}) = \mathbb{E}_{p_i,\sigma_{-i}}(u_{i,a_i}|t_i)$ for all $a_i \in A_i$ (recall that u_i is independent of θ). Since σ is an equilibrium,

$$a_i^* = \sigma_i(t_i) \in \arg\max_{a_i\in A_i} \mathbb{E}_{p_i,\sigma_{-i}}(u_{i,a_i}|t_i) = \arg\max_{a_i\in A_i} u_i(a_i, \mu^{t_i}).$$

Furthermore, again by construction, $\text{supp}\mu^{t_i} \subseteq \sigma_{-i}(T_{-i}) = C_{-i}$. Thus, $C_i \subseteq r_i(\Delta(C_{-i}))$ for each $i \in I$, that is, $C \subseteq \rho(C)$.

(**Only if**) It is sufficient to show that there exists a subjective canonical correlated equilibrium in which every rationalizable profile is played in some state of the world. Let $C = \rho^\infty(A)$ be the set of rationalizable profiles. Then $C = \rho(C)$ (Theorem 2), and for every $i \in I$ and $a_i \in C_i$ there exists a conjecture $\mu^{a_i} \in \Delta(C_{-i})$ such that $a_i \in r_i(\mu^{a_i})$. We construct the Bayesian elaboration of G as follows. Types are "copies" of rationalizable actions (they can be interpreted as "suggested actions"); if the type of i is a copy of action a_i, then the belief of i is derived from the justifying conjecture μ^{a_i}. Formally, $\Omega = C = T$; for every $i \in I$, $T_i = C_i$ and τ_i is the identity function, $\tau_i = \text{Id}_{C_i}$; for every rationalizable action/type $t_i \in C_i = T_i$, let the interim belief be $p_i(\cdot|t_i) = \mu^{t_i}$;[30] define the prior belief of i as follows: for every $t = (t_j)_{j\in I} \in C = T = \Omega$, $p_i(t) = \frac{1}{|C_i|} p_i(t_{-i}|t_i)$. Then, the profile of identity functions (for every $i \in I$ and for every $a_i \in T_i = A_i$, $\sigma_i(a_i) = a_i$) is an equilibrium of this Bayesian elaboration.[31]

■

[29] Note that, for each i, $C_i = \text{proj}_{A_i}\sigma(\tau(\Omega))$ because τ_i is onto; but $\sigma(\tau(\Omega))$ may be not Cartesian. Hence, in general, $\sigma(\tau(\Omega)) \subsetneq C$.

[30] Recall that $\mu^{t_i} \in \Delta(C_{-i}) = \Delta(T_{-i})$.

[31] The "type" $t_i = a_i$ is not to be interpreted as the action that i necessarily has to play, but rather as a specification of i's belief about the opponents and of the best reply to such belief, a best reply that i *freely chooses* in equilibrium.

This result can be generalized. Fix a game with payoff uncertainty

$$\hat{G} = \langle I, \Theta_0, (\Theta_i, A_i, u_i : \Theta \times A \to \mathbb{R})_{i \in I} \rangle ;$$

a Bayesian game based on $\hat{G}$ is a Bayesian game obtained by appending a state space, signal functions etc. to $\hat{G}$, that is,

$$BG = \langle I, \Omega, \Theta_0, \vartheta_0, (\Theta_i, T_i, A_i, \tau_i, \vartheta_i, p_i, u_i)_{i \in I} \rangle .$$

Using arguments that generalize the proof of Theorem 32, one can prove the following result.

Theorem 33. *A profile of information-types and actions $(\theta_i, a_i)_{i \in I}$ is rationalizable in the game with payoff uncertainty $\hat{G}$ if and only if there is a Bayesian game BG based on $\hat{G}$, an equilibrium σ of BG and a state of the world ω in BG such that $(\theta_i, a_i)_{i \in I} = (\vartheta_i(\tau_i(\omega)), \sigma_i(\tau_i(\omega)))_{i \in I}$.*

These results show that, without restrictions on players' exogenous beliefs, the Bayesian equilibrium assumption that players best respond to correct conjectures about their opponents' decision functions has no behavioral implications beyond those that can be derived from rationality and common belief in rationality.[32] In a sense, *rationalizability gives the "robust" implications of Bayesian equilibrium analysis.*

Suppose that the modeler is confident that contextual considerations make some restrictions on exogenous first-order beliefs transparent. Then the robust implications of Bayesian equilibrium are captured by Directed Rationalizability. More formally, consider a Bayesian game BG obtained by appending to game with payoff uncertainty $\hat{G}$ a belief structure $\langle \Omega, \vartheta_0, (T_i, \tau_i, \vartheta_i, p_i)_{i \in I} \rangle$. Say that BG yields the restrictions on exogenous beliefs Δ if, for all $i \in I$, $\theta_i \in \Theta_i$, and $t_i \in \vartheta_i^{-1}(\theta_i)$, there is some $\mu^i \in \Delta_{i,\theta_i}$ such that $p_i^1(\cdot|t_i) = \mathrm{marg}_{\Theta_0 \times \Theta_{-i}} \mu^i$ (see Section 8.5.1). With this, we obtain a variation of the previous result.

Theorem 34. *Fix a profile $\Delta = (\Delta_{i,\theta_i})_{i \in I, \theta_i \in \Theta_i}$ of restrictions on* exogenous *beliefs. A profile of information-types and actions $(\theta_i, a_i)_{i \in I}$ is Δ-rationalizable in the game with payoff uncertainty $\hat{G}$ if and only if there is a Bayesian game BG based on $\hat{G}$ that yields the restrictions on exogenous beliefs Δ, an equilibrium σ of BG, and a state of the world ω in BG such that $(\theta_i, a_i)_{i \in I} = (\vartheta_i(\tau_i(\omega)), \sigma_i(\tau_i(\omega)))_{i \in I}$.*

[32] See Brandenburger and Dekel [29], and Battigalli and Siniscalchi [17].

8.6 Incomplete and Asymmetric Information

Recall from Chapter 1 that it is important to distinguish between the mathematical structure used to represent a real world interactive decision situation[33] and the situation itself. With this in mind, let us go back to Harsanyi's metaphor, presented in Section 8.5.1: the information players have, the residual uncertainty and players' payoffs depend on a random variable, or chance move.[34] Let $\omega \in \Omega$ denote a typical realization of this random variable and assume that this realization occurs before players make their choices. For example, ω may be the order of cards in a deck before some of these cards are distributed to the players. The payoff of player i is determined by a function $\hat{u}_i : \Omega \times A \to \mathbb{R}$, that is, how payoffs depend on actions is determined by the initial chance move. Let $p_i \in \Delta(\Omega)$ denote the subjective probability measure of i over the possible realizations. Each player first receives a signal regarding the initial chance move, then chooses an action $a_i \in A_i$ simultaneously to the other players. Let $\tau_i : \Omega \to T_i$ denote player i's signal function. The prior belief and signal function of each player i are such that $p_i\left(\tau_i^{-1}(t_i)\right) > 0$ for all $t_i \in T_i$. We call such interactive situation a **game with asymmetric information** about an initial move by chance.

In this context, it may make sense to suppose the subjective priors p_i coincide with an objective probability measure $p \in \Delta(\Omega)$ and that this is commonly believed. More generally, even if the priors do not coincide, it is assumed that the profile of subjective measures $(p_i)_{i\in I}$ is transparent.[35]

An **equilibrium** of such game with asymmetric information is a strategy profile $(\sigma_i : T_i \to A_i)_{i\in I}$ such that the strategy of each player is a best reply to the strategy profile of the other players. More explicitly, define the payoff function in strategic form

$$U_i((\sigma_i)_{i\in I}) = \sum_{\omega\in\Omega} p_i(\omega)\hat{u}_i(\omega, (\sigma_j(\tau_j(\omega)))_{j\in I}).$$

An equilibrium of a given game of asymmetric information is a Nash equilibrium of the corresponding game in strategic form $\langle I, (\Sigma_i, U_i)_{i\in I}\rangle$,

[33] Or a class of situations of the same kind.

[34] As in Section 8.5.4, it is also interesting to consider the possibility that payoffs do not depend on the random variable.

[35] In line with our definition of "transparency" in Section 8.3, this means that it is true and commonly believed that, for each $i \in I$, the subjective prior belief of i is p_i.

where $\Sigma_i = A_i^{T_i}$ for each $i \in I$.

At this point, one may ask: How is this different from a game of incomplete information, modeled as a Bayesian game? If we just look at the mathematical formulation without considering how it is supposed to map to the real world, the two situations may seem indistinguishable. At the interim stage of the asymmetric-information game in which all players have received their private signals about the chance move, the interactive situation is indeed very similar to one with incomplete information: the way payoffs depend on actions is not commonly known, and each player has private information about such dependence, along with probabilistic beliefs regarding the private information and (implicitly) the beliefs of other players.

We can establish an explicit formal relationship between the two models, deriving a Bayesian game as per Definition 39 from the mathematical structure representing a game with asymmetric information about the realization of an initial chance move. For the sake of simplicity, we restrict our attention to the case in which the joint signal function $\tau = (\tau_i)_{i\in I} : \Omega \to \times_{i\in I} T_i$ is one-to-one (injective), that is, $\omega' \neq \omega''$ implies that there is at least one player i for whom $\tau_i(\omega') \neq \tau_i(\omega'')$. Such simplification rules out any residual uncertainty about payoffs. With this, we formally obtain a Bayesian game by defining $\Theta_i = T_i$, letting ϑ_i be the identity function on T_i, so that $\tau(\Omega) \subseteq \Theta = \times_{i\in I} \Theta_i$, and defining state-dependent payoffs as follows:

$$\forall \theta \in \tau(\Omega),\; u_i(\theta, a) = \hat{u}_i(\tau^{-1}(\theta), a). \tag{8.6.1}$$

Note that this is enough, because it is common belief that realizations θ outside $\tau(\Omega)$ are impossible, therefore, the definition of u_i on $\Theta\backslash\tau(\Omega)$ is irrelevant for expected payoff computations. Keeping in mind Remark 23, it is easy to verify that *a profile of functions* $(\sigma_i)_{i\in I}$ *is a Bayesian equilibrium of this Bayesian game if and only if it is an equilibrium of the asymmetric-information game.*

Recall that in order to introduce the constituent elements of the Bayesian game structure we used the metaphor of the *ex ante* stage in describing interactive situations with incomplete information. The use of such metaphor and the formal similarity between Bayesian games and the mathematical description of games with asymmetric information (about an initial chance move) has induced many scholars to neglect the differences between games with incomplete information and games with asymmetric information. We should stress, however, that they are indeed different situations. In games with incomplete information the *ex ante* stage *does not exist*, it is just a useful theoretical fiction: Ω only represents a set of possible states of the world, where by "state of the world" we mean a configuration of payoff state, information, and subjective exogenous beliefs.[36] The profile $(p_i)_{i\in I} \in [\Delta(\Omega)]^I$ of so called prior beliefs is simply a useful *mathematical tool* to determine (along with the functions ϑ_i and τ_i) players' interactive beliefs in a given state of the world. Instead, in the interactive decision problems that we just called "games with asymmetric information" the *ex ante* stage is real and the players' priors represent the expectations they hold at that stage.

These differences in the interpretation of the formal structure are not innocuous. Indeed, the interpretation (asymmetric information *vs.* incomplete information) determines to what extent some given assumptions are meaningful and plausible, and what solution concepts are appropriate. For instance, for the case of asymmetric information about the realization of an initial chance move, it is meaningful and often plausible to assume that there exists a common prior given by an objective distribution over Ω. On the other hand, in the case of incomplete information the meaning of the common prior assumption is not obvious and its plausibility is even less obvious. Furthermore, we will see that different notions of self-confirming equilibrium are appropriate for situations with incomplete information and for situations with asymmetric information about an initial chance move, even though the two situations can be represented by the same mathematical structure. In the appendix of this chapter, we show that similar considerations apply to the appropriate notion of rationalizability for Bayesian games, as long as one is willing to assume some independence restrictions on players' conjectures.

[36] Recall, we called "state of nature" the configuration $\theta = \left(\theta_0, (\theta_i)_{i\in I}\right)$ of payoff state and private information interpreted as a fixed imperfectly known parameter. A "state of the world" also include a specification of subjective beliefs.

8.7 Self-Confirming Equilibrium and Incomplete Information

As previously argued, the self-confirming (or conjectural) equilibrium concept characterizes patterns of behavior that are stable with respect to learning processes in situations of recurrent interaction, taking into account players' information feedback. The appropriate definition of self-confirming equilibrium for games with incomplete information depends on the relevant scenario, and in particular on the interpretation of the mathematical structure we use to represent incomplete information.

The first question to address is whether the set of interacting players and their characteristics are (a) fixed once and for all (long-run interaction), or (b) randomly determined in each period via draws from large populations (anonymous recurrent interaction). In both cases it is necessary to specify what information players are able to obtain ex post about the actions and characteristics of co-players. As in Section 6.3, we describe this with feedback functions. As we did earlier, we assume for the sake of simplicity that the players' goal is to maximize the expected payoff of the current period, without worrying about future payoffs.

8.7.1 Long-Run Interaction

This is the easiest case to analyze. It is sufficient to consider the game with payoff uncertainty

$$\hat{G} = \langle I, \Theta_0, (\Theta_i, A_i, u_i : \Theta \times A \to \mathbb{R})_{i \in I} \rangle .$$

The state of nature (profile of parameters) θ is *fixed* once and for all; obviously, we continue to assume that every player i knows only θ_i, that the sets of possible values for θ is $\Theta = \Theta_0 \times \left(\bigtimes_{i \in I} \Theta_i \right)$, and that *ex post* i gets further information about θ_0, θ_{-i} and a_{-i} according to a feedback function f_i. In general, the signal received may depend also on θ, that is $f_i : \Theta \times A \to M_i$; for instance, the signal could be the payoff obtained by the player.[37] Every player i holds a belief $\mu^i \in \Delta(\Theta_0 \times \Theta_{-i} \times A_{-i})$ about the residual uncertainty θ_0 and the opponents' information-types and actions.

[37] However, we need to be aware that the payoff u_i does not necessarily represent a material gain which can be cashed. Therefore in the theoretical analysis we are not forced to assume that the realization of u_i is observed by i (observed payoffs).

The question is whether, for any *given* θ, a profile of actions and beliefs $(a_i, \mu^i)_{i\in I} \in \bigtimes_{i\in I} (A_i \times \Delta(\Theta_0 \times \Theta_{-i} \times A_{-i}))$ satisfies the conditions of rationality and confirmed conjectures. These conditions can be obtained as a simple generalization of Definition 26. As in Section 6.3, we give the definition of self-confirming equilibrium for games with payoff uncertainty and feedback $(\hat{G}, f)$, where $f = (f_i)_{i\in I}$ is the profile of feedback functions.

Definition 41. *Fix a game with payoff uncertainty and feedback* $(\hat{G}, f)$ *and a state of nature* $\theta^* \in \Theta$. *A profile of actions and beliefs* $(a_i^*, \mu^i)_{i\in I} \in \bigtimes_{i\in I} (A_i \times \Delta(\Theta_0 \times \Theta_{-i} \times A_{-i}))$ *is a* ***self-confirming equilibrium (SCE) at*** θ^* *if for every* $i \in I$:[38]
*(1) (*rationality*)* $a_i^* \in r_i(\mu^i, \theta_i^*)$,
*(2) (*confirmed conjectures*):*

$$\mu^i\left(\{(\theta_0, \theta_{-i}, a_{-i}) : f_i(\theta_0, \theta_i^*, \theta_{-i}, a_i^*, a_{-i}) = f_i(\theta^*, a^*)\}\right) = 1.$$

$(a_i^*)_{i\in I}$ *is an SCE action profile at* θ^* *if there is some profile of conjectures* $(\mu^i)_{i\in I}$ *such that* $(a_i^*, \mu^i)_{i\in I}$ *is an SCE at* θ^*.

In what follows, for each game with payoff uncertainty

$$\hat{G} = \langle I, \Theta_0, (\Theta_i, A_i, u_i : \Theta \times A \to \mathbb{R})_{i\in I}\rangle$$

and state of nature $\theta \in \Theta$, we let $\hat{G}_\theta$ denote the game

$$\hat{G}_\theta = \langle I, (A_i, u_{i,\theta} : A \to \mathbb{R})_{i\in I}\rangle .$$

Thus, for a game with payoff uncertainty and feedback $(\hat{G}, f)$ (where $f_i : \Theta \times A \to M_i$ for each $i \in I$) and fixed θ, $(\hat{G}_\theta, f_\theta)$ denotes the "true" game with feedback at state of nature θ.

Remark 26. *Fix a game with payoff uncertainty* $\hat{G}$. *For each state of nature* $\theta^* \in \Theta$ *and each profile of feedback functions* f, *every Nash equilibrium of* $\hat{G}_{\theta^*}$ *is an SCE action profile at* θ^* *of the game with payoff uncertainty and feedback* $(\hat{G}, f)$.

[38] Recall that $r_i(\mu^i, \theta_i)$ is the set of actions of i that maximize her expected payoff given θ_i and μ^i.

The following remark highlights the fact that in the private-values case (when u_i does not depend on θ_0 and θ_{-i}) we get the notion of self-confirming equilibrium already introduced in Section 6.3, coherently with the claim made there according to which the self-confirming equilibrium concept only presumes that a player knows his own payoff function, not those of the co-players.

Remark 27. *Fix a game with payoff uncertainty and feedback* $(\hat{G}, f)$*, a state of nature* θ^* *and an action profile* a^**, and suppose that the game with payoff uncertainty* $\hat{G}$ *has* private values*. Then* a^* *is an SCE action profile at* θ^* *if and only if* a^* *is an SCE action profile of the game with feedback* $(\hat{G}_{\theta^*}, f_{\theta^*})$*.*

8.7.2 Anonymous Recurrent Interaction

Assume for simplicity that the given game with payoff uncertainty is *finite* and consider the following scenario: There are n heterogeneous populations, $i \in I = \{1, ..., n\}$; the fraction of agents with characteristic θ_i is $q_i(\theta_i) > 0$; n agents are drawn at random, one from each population i, and play a static game with payoffs determined by the parameterized functions $u_i : \Theta \times A \to \mathbb{R}$ $(i \in I)$, where Θ_0 represents a set of possible values of a random shock that affect agents' utility: each time agents play the game, $\theta_0 \in \Theta_0$ is drawn with probability $q_0(\theta_0) > 0$ independently of previous plays. This is precisely the scenario that motivated the definition of equilibrium for a "simple" Bayesian game (with type-independent beliefs)

$$BG = \langle I, \Theta_0, q_0, (\Theta_i, q_i, A_i, u_i)_{i \in I} \rangle \tag{8.7.1}$$

where $q_0 \in \Delta(\Theta_0)$, $q_i \in \Delta(\Theta_i)$, $u_i : \Theta \times A \to \mathbb{R}$ $(i \in I)$.

In this case, however, it is only assumed that each agent in each population i knows θ_i and $u_{i,\theta_i} : \Theta_0 \times \Theta_{-i} \times A \to \mathbb{R}$, while the interactive situation represented by BG is *not* assumed to be common knowledge. Agents play the game recurrently, each time with a different θ_0 drawn from Θ_0 and with different co-players randomly drawn from the respective populations. If the play stabilizes, at least in a statistical sense, for every j, θ_j, a_j, the fraction of agents in sub-population j with characteristic θ_j that choose action a_j remains constant. Let $\alpha_j(a_j|\theta_j)$ denote this fraction and write

$$\alpha_{-i} = (\alpha_j(\cdot|\theta_j))_{j\neq i,\theta_j\in\Theta_j} \in \bigtimes_{j\neq i} \Delta\left(A_j\right)^{\Theta_j}$$

to denote the profile of fractions for the populations different from i. By random matching (and the law of large numbers), the long-run frequency of each profile $(\theta_0, \theta_{-i}, a_{-i})$ is $q_0\left(\theta_0\right)\prod_{j\neq i}\alpha_j(a_j|\theta_j)q_j(\theta_j)$.

The ex post information feedback of an agent playing in role i is described by a feedback function $f_i : \Theta \times A \to M_i$. Hence, in a steady state, an agent of population i with characteristic θ_i that (always) chooses a_i, observes that the long-run frequency of each message m_i is

$$\mathbb{P}^{f_i}_{a_i,\alpha_{-i},q_{-i}}\left(m_i|\theta_i\right) := \sum_{(\theta_0,\theta_{-i},a_{-i}):f_i(\theta_0,\theta_i,\theta_{-i},a_i,a_{-i})=m_i} q_0\left(\theta_0\right)\prod_{j\neq i}\alpha_j(a_j|\theta_j)q_j(\theta_j).$$

Given θ_i and a_i, conjecture $\mu^i \in \Delta(\Theta_0\times\Theta_{-i}\times A_{-i})$ assigns to each message m_i the probability

$$\mathbb{P}^{f_i}_{a_i,\mu^i}\left(m_i|\theta_i\right) := \sum_{(\theta_0,\theta_{-i},a_{-i}):f_i(\theta_0,\theta_i,\theta_{-i},a_i,a_{-i})=m_i} \mu^i(\theta_0, \theta_{-i}, a_{-i}).$$

A conjecture is confirmed in the long run for a player of type θ_i who keeps playing a_i if the subjective probability of each message is equal to the observed frequency. If $a_i \in r_i(\mu^i, \theta_i)$ and μ^i is confirmed, then the agents of type θ_i choosing a_i have no reason to switch to another action. In the following definition we let $\mu^i_{\theta_i,a_i}$ denote a conjecture that justifies action a_i for an agent with characteristic (information-type) θ_i.

Definition 42. *Fix a simple Bayesian game with feedback* (BG, f)*. A profile of mixed actions and conjectures:*

$$\left(\alpha_i(\cdot|\theta_i), (\mu^i_{(a_i,\theta_i)})_{a_i\in\text{supp}\alpha_i(\cdot|\theta_i)}\right)_{i\in I,\theta_i\in\Theta_i}$$

(a mixed action $\alpha_i(\cdot|\theta_i)$ *for each* i *and* θ_i*, and a conjecture for each* i*,* θ_i *and* $a_i \in \text{supp}\alpha_i(\cdot|\theta_i)$*) is an* ***anonymous self-confirming equilibrium*** *of* (BG, f) *if for every* $i \in I$*,* $\theta_i \in \Theta_i$*,* $a_i \in A_i$*, the following conditions hold:*
*(1) (*rationality*) if* $\alpha_i(a_i|\theta_i) > 0$*, then* $a_i \in r_i(\mu^i_{(a_i,\theta_i)}, \theta_i)$*,*
*(2) (*confirmed conjectures*)* $\mathbb{P}^{f_i}_{a_i,\mu^i}\left(\cdot|\theta_i\right) = \mathbb{P}^{f_i}_{a_i,\alpha_{-i},q_{-i}}\left(\cdot|\theta_i\right).$

Observation 8. *The definition needs to be modified and made more stringent if the distributions* q_j *are known. In this case equilibrium conjectures must be consistent with* $q = (q_j)_{j \in I_0}$*. For instance, in the two-person case,* $\text{marg}_{\Theta_j} \mu^i = q_j$*.*[39]

Remark 28. *The definition of self-confirming equilibrium* at θ^* *(Definition 41) is obtained as a special case when* $q(\theta^*) = 1$*.*

Remark 29. *Every mixed equilibrium of the interim strategic form of BG is a self-confirming equilibrium profile of mixed actions.*

Adapting in the obvious way the definitions of observed payoffs and own-action independence of feedback of Section 6.3, we obtain the analog of Theorem 21: if (BG, f) satisfies observed payoffs and own-action independence of feedback, then every (anonymous) self-confirming equilibrium profile of mixed actions $(\alpha_i(\cdot|\theta_i))_{i \in I, \theta_i \in \Theta_i}$ is also a mixed Bayes-Nash equilibrium of BG, therefore the two equilibrium concepts coincide under these assumptions about feedback.

Analogous considerations apply to defining SCE in games with asymmetric information about the realization of an initial chance move. Also in this case, the confirmed-conjectures condition must require that, for each player, the observed distribution of messages coincides with the subjectively expected distribution of messages.

8.8 Appendix

The cleanest and easiest extension of the rationalizability idea to incomplete-information environments is the notion of (directed) rationalizability for games with payoff uncertainty studied in Section 8.2 (and 8.3). However, most theorists have analyzed incomplete information restricting their attention to Bayesian games. Therefore, the

[39] A special case of equilibrium of this kind obtains when actions are observed ($f_i(\theta, a) = a$) and conjectures are "naive" as they do not take into account that opponents' actions depend on their type. For example in the two-agents case without residual uncertainty, $\mu^i(\theta_j, a_j) = q_j(\theta_j)\nu^i(a_j)$ for some $\nu^i \in \Delta(A_j)$. This is called *cursed equilibrium*. Such behavior may explain, for example, the so called "winners' curse" in common value auctions. This refinement of self-confirming equilibrium was proposed by Eyster and Rabin [43], although the link to the self-confirming equilibrium concept was not made explicit.

rationalizability idea has been first extended to such games. Thus, rather than defining what is rationalizable for an information-type θ_i, theorists first tried to define what is rationalizable for a Harsanyi type t_i; but the meaning of "action a_i is rationalizable for type t_i" is not entirely transparent, because Harsanyi's notion of "type" is self-referential, hence elusive. Let us make a preliminary observation. Often the Bayesian games considered in applications *conflate Harsanyi types t_i with information-types θ_i*. Formally, this means that, for each player i, $\Theta_i = T_i$ and the information-type map ϑ_i is the identity on Θ_i (equivalently, Θ_i and T_i are isomorphic, so that ϑ_i is a bijection). For such simple Bayesian games, we just have to consider a special case of Directed Rationalizability: for each i and each θ_i, let

$$\Delta_{i,\theta_i} = \left\{ \mu^i \in \Delta\left(\Theta_0 \times \Theta_{-i} \times A_{-i}\right) : \mathrm{marg}_{\Theta_0 \times \Theta_{-i}} \mu^i = p_i^1\left(\cdot|\theta_i\right) \right\},$$

where $p_i^1\left(\cdot|\theta_i\right)$ is the exogenous first-order belief of type θ_i. Let $\Delta = \left(\Delta_{i,\theta_i}\right)_{i \in I, \theta_i \in \Theta_i}$ denote the corresponding profile of belief restrictions. With this, we compute what actions are Δ-rationalizable for each type θ_i of each player i.

What we explain below applies instead to *general* Bayesian games. We first delve into the conceptual subtleties of defining rationalizability for Bayesian games and then we illustrate the concepts with the so called "electronic mail game." To simplify the analysis, we restrict our attention to games with *finite* action sets and finite or *countable* type sets.

8.8.1 Rationalizability in Bayesian Games

A game with payoff uncertainty, i.e., the mathematical structure

$$\hat{G} = \langle I, \Theta_0, (\Theta_i, A_i, u_i)_{i \in I} \rangle,$$

does not specify the exogenous interactive beliefs of players, that is, their beliefs about private information and beliefs of the opponents (recall that we call such beliefs "exogenous" because they are not explained, nor partially determined, by strategic reasoning). A specification of exogenous beliefs is necessary to extend the traditional definition of equilibrium due to Nash.[40] The richer structure

$$BG = \langle I, \Omega, \Theta_0, \vartheta_0, (\Theta_i, T_i, A_i, \tau_i, \vartheta_i, p_i, u_i)_{i \in I} \rangle ,$$

known as "Bayesian game," was introduced with this purpose, allowing the definition of Bayesian equilibrium.

It was also pointed out that, for traditional equilibrium analysis, it makes no difference whether the given mathematical structure BG represents a game situation where there is no common knowledge of the payoff functions (incomplete information), or a situation where there is asymmetric information about an initial chance move. The interpretation of BG may make some specific assumptions about interactive beliefs more or less plausible (think about the common prior assumption), but the computation of equilibria is not affected.[41]

However, the interpretation seems to matter when it comes to define the appropriate concept of rationalizability, assuming that the interactive situation represented by BG is transparent to the players. The crucial point is the following: in an incomplete information game, the rationality and strategic reasoning of a player need to be evaluated according to his type. Specifically, we consider the so called *interim* stage and, for each action a_i and type t_i, we ask: Is a_i a justifiable choice for t_i?[42] On the other hand, in a game with asymmetric information about an initial chance move, rationality and strategic reasoning need to be assessed at the *ex ante* stage. For each decision function $\sigma_i : T_i \to A_i$, we ask: Is σ_i a justifiable decision function for player i? By answering these different questions, game theorists obtained two different concept of rationalizability for Bayesian games.

Definition 43. *Let BG be a Bayesian game. An action $a_i \in A_i$ is* interim rationalizable *for type $t_i \in T_i$ of player $i \in I$, if a_i is rationalizable for t_i in the* interim *strategic form $\mathcal{IS}(BG)$. A decision function $\sigma_i : T_i \to A_i$ is* ex ante rationalizable *if σ_i is a rationalizable decision function of the* ex ante *strategic form $\mathcal{AS}(BG)$.*

[40] As we argue in Section 8.7, this does not hold for other definitions of equilibrium that characterize the rest points of adaptive processes.

[41] As we pointed out, this is not true for self-confirming equilibrium.

[42] Recall that we say that a choice is "*justifiable*" if it is a best reply to some belief. A choice is "*rationalizable*" if it survives the iterated elimination of non justifiable choices.

Lemma 24. *If a decision function σ_i^* is justifiable in the* ex ante *strategic form, then for every type t_i the action $a^*_{t_i} = \sigma_i^*(t_i)$ is justifiable in the* interim *strategic form.*

Proof. First observe that, for each type t_i, according to eq. (8.5.2), the actions of the other types t_i' of player i are completely irrelevant for the determination of the *interim* payoff. Hence, one can re-define the conjecture of type t_i in the *interim* strategic form as a probability distribution over the action profiles of the types of players $j \neq i$. Notice that such action profiles coincide with the profiles of decision functions of the opponents of player i in the *ex ante* strategic form:

$$(a_{t_j})_{j\neq i, t_j \in T_j} \in \bigtimes_{j\neq i} \left(A^{T_j}\right) = \bigtimes_{j\neq i} \Sigma_j = \Sigma_{-i}.$$

Hence, the set of conjectures of any type/player t_i in the *interim* strategic form coincides with the sets of conjectures of player i in the *ex ante* strategic form. To be consistent with the notation used for games of complete information, we denote by $U_i(\sigma_i, \mu^i)$ and $\overline{u}_{t_i}(a_i, \mu^i)$ the expected payoffs of player i and of type t_i in the corresponding strategic forms, given conjecture $\mu^i \in \Delta(\Sigma_{-i})$.

Let us fix arbitrarily a decision function σ_i^* and a conjecture μ^i. We now prove that σ_i^* is a best reply to μ^i if and only if, for every type $t_i \in T_i$, action $\sigma_i^*(t_i)$ is a best reply to μ^i. This implies the thesis. Fix $\mu^i \in \Delta(\Sigma_{-i})$. For every decision function σ_i, the expected payoff of σ_i given μ^i is

$$U_i(\sigma_i, \mu^i) = \sum_{\sigma_{-i}} \mu^i(\sigma_{-i}) U_i(\sigma_i, \sigma_{-i})$$

$$\begin{aligned}
&= \sum_{\sigma_{-i}} \mu^i(\sigma_{-i}) \sum_{t_i} p_i\left(t_i\right) \sum_{\theta_0, t_{-i}} p_i\left(\theta_0, t_{-i}|t_i\right) u_i(\theta_0, \vartheta_i(t_i), \vartheta_{-i}(t_{-i}), \sigma_i(t_i), \sigma_{-i}(t_{-i})) \\
&= \sum_{t_i} p_i\left(t_i\right) \sum_{\sigma_{-i}} \mu^i(\sigma_{-i}) \sum_{\theta_0, t_{-i}} p_i\left(\theta_0, t_{-i}|t_i\right) u_i(\theta_0, \vartheta_i(t_i), \vartheta_{-i}(t_{-i}), \sigma_i(t_i), \sigma_{-i}(t_{-i})) \\
&= \sum_{t_i} p_i(t_i) \overline{u}_{t_i}(\sigma_i(t_i), \mu^i).
\end{aligned}$$

Recall: $p_i\left(t_i\right) > 0$ for each $t_i \in T_i$. It follows that

$$\sigma_i^* \in \arg\max_{\sigma_i} U_i(\sigma_i, \mu^i) = \arg\max_{\sigma_i} \sum_{t_i} p_i\left(t_i\right) \overline{u}_{t_i}(\sigma_i(t_i), \mu^i)$$

if and only if

$$\forall t_i \in T_i,\ \sigma_i^*(t_i) \in \arg\max_{a_i} \overline{u}_{t_i}(a_i, \mu^i).$$

The average of the expected payoffs for the different types of i is maximized if and only if[43] the expected payoff of every type of i is maximized. ■

Corollary 5. *If a decision function $\sigma_i \in \Sigma_i$ is rationalizable in the* ex ante *strategic form, then for every type $t_i \in T_i$ the action $\sigma_i(t_i)$ is rationalizable for t_i in the* interim *strategic form.*

From the previous considerations it looks like interim rationalizability is the appropriate solution concept if we interpret the mathematical structure BG as an interactive situation with incomplete information, where the *ex ante* stage is only a useful notational device and therefore the states ω represent only possible configurations of private information and beliefs. *Ex ante* rationalizability is instead the appropriate solution concept if we interpret BG as an interactive situation with asymmetric information over some initial chance move (which directly affects the payoffs only through the payoff-relevant components of θ).

Corollary 5 states that *ex ante* rationalizability is at least as strong a solution concept as *interim* rationalizability. It can be shown by example that it is actually stronger. The formal reason may be understood already from the proof of Lemma 24. Assume that, for every t_i, action $\sigma_i^*(t_i)$ is justifiable. Then there exists a profile of conjectures $(\mu^{t_i})_{t_i \in T_i}$ (with $\mu^{t_i} \in \Delta(\Sigma_{-i})$) such that $\sigma_i^*(t_i)$ is a best reply to μ^{t_i} for each t_i. But this does not imply that there exists a *unique* conjecture $\mu^i \in \Delta(\Sigma_{-i})$ such that $\sigma_i^*(t_i)$ is a best reply to μ^i for every t_i. The difference between the two solutions concept is clearly illustrated by the following numerical example.

Example 39. Let $\Omega = \{\omega', \omega''\}$; player 1 (Rowena) knows the true state, $\tau_1(\omega') \neq \tau_1(\omega'')$, whereas player 2 (Colin) does not, $\tau_2(\omega') = \tau_2(\omega'')$, and considers the two states equally likely: $p_2(\omega') = p_2(\omega'') = \frac{1}{2}$. The payoff functions are as in the matrixes below (note that the payoff of Rowena is independent of the state, this simplifies the example):

[43] Of course, the "only if" part holds because we assume that $p_i(t_i) > 0$ for each $t_i \in T_i$.

ω'	c	d
a	$3,3$	$0,2$
m	$2,0$	$2,2$
b	$0,0$	$3,2$

ω''	c	d
a	$3,0$	$0,2$
m	$2,0$	$2,2$
b	$0,3$	$3,2$

First, observe that every action by Colin is justifiable. In particular, c is justifiable by the conjecture that Rowena chooses a if ω' and b if ω''. It is easy to verify that for each type of Rowena all the actions are justifiable by some conjecture about Colin. Hence, if choices are evaluated at the *interim* stage it is not possible to exclude any action, everything is *interim* rationalizable. Instead if choices are evaluated at the *ex ante* stage, then the decision function $\sigma_1(\omega') = a$, $\sigma_1(\omega'') = b$, denoted by ab, can be excluded. Indeed, since Rowena's payoff does not depend on the state, ab could be a best reply to some conjecture μ^1 only if both a and b were best replies to μ^1. But if Rowena believes $\mu^1(c) \geq \frac{1}{2}$, then b is **not** a best reply; if Rowena believes $\mu^1(c) \leq \frac{1}{2}$, then a is **not** a best reply. Therefore ab is not justifiable by any μ^1. The same argument shows that ba is not justifiable either. The *ex ante* justifiable decision functions of Rowena are: *aa,am,ma,mm,bb,bm,mb*.[44] Given any belief μ^2 such that $\mu^2(ab) = 0$, action c yields an expected payoff less or equal than $\frac{3}{2}$ (check this); action d instead yields $2 > \frac{3}{2}$. Then, the only rationalizable action for Colin in the *ex ante* strategic form is d. It follows that the only rationalizable decision function of Rowena in the *ex ante* strategic form is bb. ▲

Although the gap between *ex-ante* and *interim* rationalizability was at first just accepted as a fact, it should be disturbing. Conceptually, rationalizability is meant to capture the assumptions of rationality (i.e., expected utility maximization) and common belief in rationality given the background transparency of the Bayesian Game BG. Since both strategic forms are based on the same Bayesian game and since *ex-ante* maximization is equivalent to *interim* maximization,[45] why does not *ex-*

[44]Given that Row's payoff does not depend on ω, those decision functions that select different actions for the two states are justifiable only by beliefs that make Row indifferent between these two actions; the decision functions am and ma are among the best replies to the conjecture $\mu^1(c) = \frac{2}{3}$, the decision functions bm and mb are among the best replies to the conjecture $\mu^1(c) = \frac{1}{3}$.

[45]This equivalence holds under the assumptions that for every player i, all types t_i

ante rationalizability coincide (in terms of behavioral predictions) with *interim* rationalizability? Which features of the strategic form create the gap highlighted in Example 39? Before providing an answer to these questions, we introduce another puzzling feature of rationalizability in Bayesian environments: the dependence of *interim* rationalizability on apparently irrelevant details of the state space. To understand this problem, consider the following example.

Example 40. [Dekel et al. [40]] Rowena and Colin are involved in a betting game. Each player can decide to bet (action B) or not to bet (action N). There is an unknown parameter $\theta_0 \in \Theta_0 = \{\theta_0', \theta_0''\}$ and players have no private information ($\Theta_i = \left\{\bar{\theta}_i\right\}$ for every agent i). Rowena (Colin) wins if both players bet and $\theta_0 = \theta_0'$ ($\theta_0 = \theta_0''$). The decision to bet entails a deadweight loss of \$4 independently of the opponent's action; if both agents bet, the loser gives \$12 to the winner. The corresponding game with payoff uncertainty $\hat{G}$ can be represented as follows:

θ_0'	B	N
B	$8, -16$	$-4, 0$
N	$0, -4$	$0, 0$

θ_0''	B	N
B	$-16, 8$	$-4, 0$
N	$0, -4$	$0, 0$

Let us assume that there is common belief that each agent assigns probability $\frac{1}{2}$ to each payoff state. Thus, each player i has only one hierarchy of beliefs about θ_0.

The simplest state space representing this situation is the following: $\Omega = \{\omega', \omega''\}$,,$\vartheta_0(\omega') = \theta_0'$, $\vartheta_0(\omega'') = \theta_0''$ and for every agent i, $T_i = \{\bar{t}_i\}$, functions ϑ_i and τ_i are trivially defined and $p_i(\omega') = \frac{1}{2}$. In this case, the *ex-ante* and the *interim* strategic forms coincide and are represented by the following game:

	B	N
B	$-4, -4$	$-4, 0$
N	$0, -4$	$0, 0$

---

have positive probability.

It is immediate to see that betting is not rationalizable. Thus, N is the only *interim* rationalizable action.
Now, consider an alternative state space: $\Omega = \left\{\omega^1, \omega^2, \ldots, \omega^8\right\}$,

$$\vartheta_0(\omega) = \begin{cases} \theta_0', & \text{if } \omega \in \left\{\omega^1, \omega^2, \omega^3, \omega^4\right\}, \\ \theta_0'', & \text{if } \omega \in \left\{\omega^5, \omega^6, \omega^7, \omega^8\right\}, \end{cases}$$

for every agent i, $T_i = \{t_i', t_i''\}$, ϑ_i is trivially defined,

$$\tau_1(\omega) = \begin{cases} t_1', & \text{if } \omega \in \left\{\omega^1, \omega^2, \omega^5, \omega^6\right\}, \\ t_1'', & \text{if } \omega \in \left\{\omega^3, \omega^4, \omega^7, \omega^8\right\}, \end{cases}$$

$$\tau_2(\omega) = \begin{cases} t_2', & \text{if } \omega \in \left\{\omega^1, \omega^3, \omega^5, \omega^7\right\}, \\ t_2'', & \text{if } \omega \in \left\{\omega^2, \omega^4, \omega^6, \omega^8\right\}, \end{cases}$$

and $p_1 = p_2 = p$:

State of the world, ω	ω_1	ω_2	ω_3	ω_4	ω_5	ω_6	ω_7	ω_8
$p[\omega]$	$\frac{1}{4}$	0	0	$\frac{1}{4}$	0	$\frac{1}{4}$	$\frac{1}{4}$	0

One can easily check that this common prior induces a distribution over residual uncertainty and types which can be represented as follows:

θ_0'	t_2'	t_2''
t_1'	$\frac{1}{4}$	0
t_1''	0	$\frac{1}{4}$

θ_0''	t_2'	t_2''
t_1'	0	$\frac{1}{4}$
t_1''	$\frac{1}{4}$	0

and that for both types of both players the following holds: (i) the player has no private information, (ii) he assigns probability $\frac{1}{2}$ to each state θ_0, and (iii) there is common belief of this. Thus, this second state space represents exactly the same belief hierarchies than the former one: each i assigns probability $\frac{1}{2}$ to each θ_0, each i is certain that $-i$ assigns probability $\frac{1}{2}$ to each θ_0, and so on. If we write down the *ex-ante* strategic form derived from this state space, we get the following game:

	BB	BN	NB	NN
BB	$-4, -4$	$-4, -2$	$-4, -2$	$-4, 0$
BN	$-2, -4$	$1, -5$	$-5, 1$	$-2, 0$
NB	$-2, -4$	$-5, 1$	$1, -5$	$-2, 0$
NN	$0, -4$	$0, -2$	$0, -2$	$0, 0$

Notice that, in this game, strategy *NB* (*BN*) is a best response for Rowena to the conjecture that Colin is playing *NB* (*BN*), while *NB* (*BN*) is Colin's best response to the conjecture that Rowena is playing *BN* (*NB*). Thus, the set $\{BN,NB\}\times\{BN,NB\}$ has the best response property and, consequently, betting is *ex-ante* and, thanks to Corollary 5, *interim* rationalizable. ▲

In Example 40, the two state spaces generate two different Bayesian Games representing the same information and belief hierarchies over Θ; thus, they seem to represent the same background common knowledge and beliefs concerning the game with payoff uncertainty $\hat{G}$. Then, should not we expect them to lead to the same *interim* rationalizable prediction? Why is this intuition false? Why does the choice of one type space over the other matter?

It turns out that these two *puzzles* are related: they both *depend on independence assumptions that are implicitly made in the construction of the strategic forms.* Once these assumptions are removed and the proper (*ex ante* and *interim*) notions of *correlated rationalizability* are defined, the puzzles disappear: (i) the set of *interim correlated rationalizable* actions is not affected by the choice among different state spaces representing the same information and belief hierarchies over Θ, and (ii) there is *no gap* between *interim correlated* and *ex-ante correlated rationalizability.*

First, let us focus on the puzzle highlighted by Example 40. Consider the second state space and recall that in such state space each player has two types representing the same belief hierarchy over Θ.[46] In this case, if type t_1' conjectures that t_2' plays B and t_2'' plays N, this conjecture, together with the common belief prior p, will lead t_1' to believe that Colin will bet only when the residual uncertainty is θ_0'; this belief

[46] In Example 40, players have no private information.

will justify the decision of type t_1' to bet.[47] To put it differently, in this state space, Rowena's types hold beliefs in which Colin's type is correlated with residual uncertainty and this can, in turn, introduce correlation between Colin's behavior, $\sigma_2(\cdot)$, and residual uncertainty, θ_0. On the contrary, this correlation is not possible with the first state space in which each belief hierarchy is represented only by one type; this happens because the construction of the *interim* strategic form imposes an implicit independence constraint on players' beliefs: each player has to regard his opponents' behavior, $\sigma_2(\cdot)$, as independent of the residual uncertainty, θ_0, conditional on the opponents' types. If there are relatively few types representing the same information and belief hierarchy, this independence constraint may bind, preventing players from holding correlated beliefs and restricting the set of rationalizable actions.[48] To address this problem, Dekel et al. [40] define a solution concept, **Interim Correlated Rationalizability**, in which beliefs are explicitly allowed to exhibit correlation between opponents' actions and residual uncertainty. Since the independence restriction imposed by the construction of the *interim* strategic form is not related to the assumption of rationality, we can immediately conclude that *interim correlated rationalizability* represents a step forward in the search for a solution concept *characterized by Rationality and Common Belief in Rationality (RCBR) given the background transparency of the given Bayesian game BG*. Actually, it can be shown that *interim* correlated rationalizability fully characterizes the behavioral implications of these epistemic assumptions. To provide a formal definition of this solution concept, fix a Bayesian Game $BG = \langle I, \Omega, \Theta_0, \vartheta_0, (\Theta_i, T_i, A_i, \tau_i, \vartheta_i, p_i, u_i)_{i\in I}\rangle$. Then, for every i, t_i and $\mu^i \in \Delta(\Omega \times A_{-i})$ let

$$r_i\left(t_i, \mu^i\right) = \arg\max_{a_i\in A_i} \sum_{\omega, a_{-i}} \mu^i(\omega, a_{-i})\, u_i\left(\vartheta_0(\omega), \vartheta_i(t_i), \vartheta_{-i}(\tau_{-i}(\omega)), a_i, a_{-i}\right).$$

The set of *interim* correlated rationalizable actions is iteratively defined

[47] It is easy to see that a similar reasoning holds for any other type.

[48] This problem would not arise if for every player i and every information and belief hierarchy he may hold, there were at least $|A_i|$ types representing this information and belief hierarchy. Stricter requirement can be provided, but this would go beyond the scope of this textbook.

as follows: for every i and t_i, let $ICR_i^{0,BG}(t_i) = A_i$ and for each $k \geq 1$, let

$$ICR_i^{k,BG}(t_i) = \left\{ \begin{array}{c} a_i \in A_i : \exists \mu^i \in \Delta(\Omega \times A_{-i}), \ \exists \varphi_{-i} : \Theta_0 \times T_{-i} \to \Delta(A_{-i}), \\ 1) \ a_i \in r_i\left(t_i, \mu^i\right) \\ 2) \ \forall (\theta_0, t_{-i}), \ \varphi_{-i}(\theta_0, t_{-i})(a_{-i}) > 0 \Rightarrow a_{-i} \in ICR_{-i}^{k-1,BG}(t_{-i}) \\ 3) \ \forall (\omega, a_{-i}), \ \mu^i(\omega, a_{-i}) = p(\omega|t_i) \cdot \varphi_{-i}(\vartheta_0(\omega), \tau_{-i}(\omega))(a_{-i}) \end{array} \right\},$$

where $ICR_{-i}^{k-1,BG}(t_{-i}) = \times_{j \neq i} ICR_j^{k-1,BG}(t_j)$.

In the previous definition, function $\varphi_{-i}(\theta_0, t_{-i})$ represents the conjecture of player i concerning the behavior of his opponents given their types are t_{-i} and residual uncertainty is θ_0. Since such functions depend both on θ_0 and on t_{-i}, representing conjectures with such functions allows to introduce correlation between a_{-i} and θ_0 even after conditioning on t_{-i}; this correlation is not allowed by *interim* rationalizability.[49]

Definition 44. *Fix a Bayesian game BG. For player $i \in I$, action $a_i \in A_i$ and type $t_i \in T_i$, a_i is* interim *correlated rationalizable for t_i if $a_i \in ICR_i^{BG}(t_i) = \bigcap_{k \geq 0} ICR_i^{k,BG}(t_i)$.*

Since *interim* correlated rationalizability allows for correlation between opponents' behavior and residual uncertainty, it does not impose any implicit independence restriction and it captures the assumptions of players' rationality and common belief in rationality given the transparency of the Bayesian game BG (as usual, the k-th step in the iterative definition of interim correlated rationalizability corresponds to the $k-1$-th level of mutual beliefs in rationality). As a consequence, given a game with payoff uncertainty $\hat{G}$, the set of interim correlated rationalizable actions of type t_i depends only on the information and *belief hierarchy* entailed by t_i. From now on, the latter will be denoted by $\hat{p}_i(t_i)$. The following Theorem formalizes this result.

Theorem 35. *Fix $\hat{G} = \left\langle I, \Theta_0, (\Theta_i, A_i, u_i)_{i \in I} \right\rangle$ and consider two Bayesian Games based on it: $BG' = \langle I, \Omega', \Theta_0, \vartheta_0', (\Theta_i, T_i', A_i, \tau_i', \vartheta_i', p_i', u_i)_{i \in I} \rangle$ and $BG'' = \langle I, \Omega'', \Theta_0, \vartheta_0'', (\Theta_i, T_i'', A_i, \tau_i'', \vartheta_i'', p_i'', u_i)_{i \in I} \rangle$. Fix $i \in I$, $t_i' \in T_i'$ and $t_i'' \in T_i''$ arbitrarily. Then, if $\hat{p}_i(t_i') = \hat{p}_i(t_i'')$, $ICR_i^{BG'}(t_i') = ICR_i^{BG''}(t_i'')$.*

[49]Indeed, it is possible to show that the set of *interim* rationalizable actions for type t_i is equivalent to the set of actions obtained through an iterative construction similar to the one we just described, but in which $\varphi_{-i} : T_{-i} \to \Delta(A_{-i})$.

Proof. Take $t_i' \in T_i'$ and $t_i'' \in T_i''$ and suppose $\hat{p}_i(t_i') = \hat{p}_i(t_i'')$. We will prove by induction that, for every i and for every $k \geq 0$, $ICR_i^{k,BG'}(t_i') = ICR_i^{k,BG''}(t_i'')$. To simplify notation, we will not specify the dependence of the set of *interim* correlated rationalizable actions on Bayesian Game (notice that this is also justified by the statement of the Theorem). The result is trivially true for $k = 0$. Now, suppose that $ICR_i^s(t_i') = ICR_i^s(t_i'')$ for every $s \leq k-1$. We need to show that $ICR_i^k(t_i') = ICR_i^k(t_i'')$.

Take any $a_i \in ICR_i^k(t_i')$. By definition, we can find $\mu_{t_i'}^i \in \Delta(\Omega' \times A_{-i})$ and $\varphi_{t_i'} : \Theta_0 \times T_{-i}' \to \Delta(A_{-i})$ such that (i) $a_i \in r_i\left(t_i', \mu_{t_i'}^i\right)$, (ii) $\varphi_{t_i'}(\theta_0, t_{-i})(a_{-i}) > 0$ implies $a_{-i} \in ICR_{-i}^{k-1}(t_{-i})$ and (iii) for every pair (ω, a_{-i}),

$$\mu_{t_i'}^i(\omega, a_{-i}) = p\left(\omega | t_i'\right) \cdot \varphi_{t_i'}\left(\vartheta_0'(\omega), \tau_{-i}'(\omega)\right)(a_{-i}).$$

Let

$$P_{-i}^{k-1} = \left\{\left(\vartheta_j'(t_j), p_j^{k-1}(t_j)\right)_{j \neq i} : t_j \in T_j'\right\}$$

be the set of possible profiles of private information and $(k-1)$-th order beliefs that agents other than i can have. For every $\theta_0 \in \Theta_0$, $\pi_{-i}^{k-1} \in P_{-i}^{k-1}$ and $a_{-i} \in A_{-i}$, define[50]

$$[\theta_0]_{BG'} = \left\{(\omega, a_{-i}) \in \Omega' \times A_{-i} : \vartheta_0(\omega) = \theta_0\right\},$$

$$[a_{-i}]_{BG'} = \left\{\left(\omega, a_{-i}'\right) \in \Omega' \times A_{-i} : a_{-i}' = a_{-i}\right\},$$

and

$$\left[\pi_{-i}^{k-1}\right]_{BG'} = \left\{(\omega, a_{-i}) \in \Omega' \times A_{-i} : \left(\vartheta_j'(\tau_j(\omega)), p_j^{k-1}(\tau_j(\omega))\right)_{j \neq i} = \pi_{-i}^{k-1}\right\}.$$

To ease notation, we also write

$$\left[\theta_0, \pi_{-i}^{k-1}, a_{-i}\right]_{BG'} = [\theta_0]_{BG'} \cap \left[\pi_{-i}^{k-1}\right]_{BG'} \cap \left[a_{-i}'\right]_{BG'},$$

$$\left[\theta_0, \pi_{-i}^{k-1}\right]_{BG'} = [\theta_0]_{BG'} \cap \left[\pi_{-i}^{k-1}\right]_{BG'},$$

[50] As in the previous sections, we let symbols in square brackets denote corresponding events in the relevant state space.

and

$$\left[\theta_0, a'_{-i}\right]_{BG'} = [\theta_0]_{BG'} \cap \left[a'_{-i}\right]_{BG'}.$$

Finally, for every $\omega \in \Omega'$, let

$$\hat{\pi}_{-i}^{k-1,BG'}(\omega) = \left(\vartheta'_j\left(\tau'_j(\omega)\right), p_j^{k-1}\left(\tau'_j(\omega)\right)\right),$$

and, similarly, for every $\omega \in \Omega''$, let

$$\hat{\pi}_{-i}^{k-1,BG''}(\omega) = \left(\vartheta''_j\left(\tau''_j(\omega)\right), p_j^{k-1}\left(\tau''_j(\omega)\right)\right).$$

Then for every $\theta_0 \in \Theta_0$ and $\pi_{-i}^{k-1} \in P_{-i}^{k-1}$, define

$$\mu_{t'_i}\left(\left[\theta_0, \pi_{-i}^{k-1}\right]\right) = \sum_{(\omega,a_{-i})\in\left[\theta_0,\pi_{-i}^{k-1}\right]_{BG'}} \mu_{t'_i}(\omega, a_{-i}).$$

Note that $\mu_{t'_i}\left(\left[\theta_0, \pi_{-i}^{k-1}\right]\right)$ is the probability that type t'_i assigns to the event "θ_0 is the case and other agents have private information and $(k-1)$-th order beliefs given by π_{-i}^{k-1}." If $\mu_{t'_i}\left(\left[\theta_0, \pi_{-i}^{k-1}\right]\right) > 0$, let

$$\varphi_{t_i}\left(\theta_0, \pi_{-i}^{k-1}\right)\left(a'_{-i}\right) = \frac{\sum_{(\omega,a_{-i})\in\left[\theta_0,\pi_{-i}^{k-1},a'_{-i}\right]_{BG'}} \mu_{t'_i}(\omega, a_{-i})}{\mu_{t'_i}\left(\left[\theta_0, \pi_{-i}^{k-1}\right]\right)};$$

if $\mu_{t'_i}\left(\left[\theta_0, \pi_{-i}^{k-1}\right]\right) = 0$, take any $t'_{-i} = \left(t'_j\right)_{j\neq i} \in T'_{-i}$ such that $\left(\vartheta'_j\left(t'_j\right), p_j^{k-1}\left(t'_j\right)\right)_{j\neq i} = \pi_{-i}^{k-1}$ and let

$$\varphi_{t'_i}\left(\theta_0, \pi_{-i}^{k-1}\right)(a_{-i}) = \begin{cases} \frac{1}{\left|ICR_{-i}^{k-1}(t'_{-i})\right|}, & \text{if } a_{-i} \in ICR_{-i}^{k-1}\left(t'_{-i}\right), \\ 0, & \text{otherwise} \end{cases}$$

(by the inductive hypothesis, the definition of $\varphi_{t'_i}$ does not depend on the actual choice of t'_{-i}).

For every $(\omega, a_{-i}) \in \Omega'' \times A_{-i}$, define

$$\mu_{t''_i}(\omega, a_{-i}) = p\left(\omega|t''_i\right)\varphi_{t'_i}\left(\vartheta''_0(\omega), \hat{\pi}_{-i}^{k-1,BG''}(\omega)\right)(a_{-i}),$$

where the previous expression is well defined since BG' and BG'' are based on the same $\hat{G}$ and $\hat{p}_i(t_i') = \hat{p}_i(t_i'')$. Moreover, since $\hat{p}_i(t_i') = \hat{p}_i(t_i'')$,

$$\begin{aligned}
\mu_{t_i'}\left(\left[\theta_0, \pi_{-i}^{k-1}\right]\right) &= \sum_{(\omega,a_{-i})\in\left[\theta_0,\pi_{-i}^{k-1}\right]_{BG'}} \mu_{t_i'}(\omega, a_{-i}) \\
&= p\left(\left\{\omega\in\Omega' : \vartheta'(\omega)=\theta_0,\ \hat{\pi}_{-i}^{k-1,BG'}(\omega)=\pi_{-i}^{k-1}\right\}|t_i'\right) \\
&= p\left(\left\{\omega\in\Omega'' : \vartheta''(\omega)=\theta_0,\ \hat{\pi}_{-i}^{k-1,BG''}(\omega)=\pi_{-i}^{k-1}\right\}|t_i''\right).
\end{aligned}$$

Then, for every $\left(\theta_0, a_{-i}'\right)$, we have

$$\sum_{(\omega,a_{-i})\in\left[\theta_0,a_{-i}'\right]_{BG''}} \mu_{t_i''}(\omega, a_{-i}) = \sum_{\omega:\vartheta_0''(\omega)=\theta_0} p\left(\omega|t_i''\right)\varphi_{t_i'}\left(\vartheta_0''(\omega), \hat{\pi}_{-i}^{k-1}\left(\tau_{-i}''(\omega)\right)\right)\left(a_{-i}'\right)$$

$$= \sum_{\pi_{-i}^{k-1}\in P_{-i}^{k-1}} p\left(\left\{\omega : \vartheta_0''(\omega)=\theta_0,\ \hat{\pi}_{-i}^{k-1,BG''}(\omega)=\pi_{-i}^{k-1}\right\}\mid t_i''\right)\varphi_{t_i'}\left(\theta_0, \pi_{-i}^{k-1}\right)\left(a_{-i}'\right)$$

$$= \sum_{\pi_{-i}^{k-1}\in P_{-i}^{k-1}} \mu_{t_i'}\left(\left[\theta_0, \pi_{-i}^{k-1}\right]\right)\frac{\sum_{(\omega,a_{-i})\in\left[\theta_0,\pi_{-i}^{k-1},a_{-i}'\right]_{BG'}}\mu_{t_i}(\omega,a_{-i})}{\mu_{t_i'}\left(\left[\theta_0,\pi_{-i}^{k-1}\right]\right)}$$

$$= \sum_{(\omega,a_{-i})\in\left[\theta_0,a_{-i}'\right]_{BG'}} \mu_{t_i'}(\omega, a_{-i}),$$

where the definition of $[\theta_0, a_{-i}]_{BG''}$ is analogous to the one of $[\theta_0, a_{-i}]_{BG'}$. Thus, $\mu_{t'}$ and $\mu_{t''}$ have the same marginal distribution over $\Theta_0 \times A_{-i}$ implying that $a_i \in r_i\left(t_i'', \mu_{t_i''}^i\right)$. Furthermore, we already constructed a function $\varphi_{-i} : \Theta_0 \times T_{-i}'' \to \Delta(A_{-i})$ (namely, $\varphi_{t_i'}$) such that for every (ω, a_{-i}), $\mu_{t_i''}^i(\omega, a_{-i}) = p(\omega|t_i'')\varphi_{-i}\left(\vartheta_0''(\omega), \tau_{-i}''(\omega)\right)(a_{-i})$. Finally, by construction and by the inductive hypothesis, $\mu_{t_i''}(\omega, a_{-i}) > 0$ implies $a_{-i} \in ICR_{-i}^{k-1}\left(\tau_{-i}''(\omega)\right)$. We conclude that $a_i \in ICR_i^k(t_i'')$. The statement of the Theorem follows by induction. ■

Notice that, by definition, the conditional independence restriction implicit in the construction of the *interim* strategic form has no bite when there is distributed knowledge of the state (Θ_0 is a singleton); in this case it is easy to verify that *interim* correlated rationalizability is equivalent to *interim* rationalizability. On the contrary, suppose that there is some residual uncertainty ($|\Theta_0| > 1$) and that we insist on using *interim* rationalizability. A natural question arises: can we at least provide an *expressible* characterization of this solution concept and, in particular, of the independence restriction implied by it? A characterization is deemed **expressible** *if it can be stated in a language based on primitives (that is, elements contained in the description of the game with payoff uncertainty,* $\hat{G}$*) and terms derived from them (such as, hierarchies of beliefs over these primitives).*

The answer to the previous question is affirmative only in some particular cases and the reason goes to the very heart of Harsanyi's approach. To understand why, recall that *interim* rationalizability requires players to regard opponents' actions as independent of the residual uncertainty *conditional on their types.* But what is a type? A *type* is a self-referential object: it *is a private information and a belief over residual uncertainty and other players' types.* Thus, unless we can establish a one-to-one mapping between types and agents' information and belief hierarchies over primitives, the conditional independence assumption is not expressible. Obviously, in order to assess the existence of such a mapping, we can use both payoff-relevant and payoff-irrelevant primitives. Thus, even though two types represent the same private information and belief hierarchy over payoff-relevant parameters, they can still be distinguished by the information and beliefs hierarchies over payoff-irrelevant elements that they capture. Whenever this is the case, the conditional independence assumption is expressible and we can show that, given the transparency of BG, *interim* rationalizability characterizes the behavioral implications of the following epistemic assumptions: (R) players are rational, (CI) their beliefs satisfy *independence* between of the opponents' behavior and the residual uncertainty *conditional* on the information and belief hierarchy over primitives of the opponents, and ($\text{CB}(R \cap CI)$) there is common belief of R and CI.

Notice that, insofar players may believe that (i) their opponents' behavior may depend on payoff-irrelevant information, and (ii) this information is correlated with θ_0, the information and belief hierarchies over payoff-irrelevant parameters may still be *strategically relevant.* For instance, in Example 40, Rowena may be "superstitious" and believe that the payoff-relevant state, θ_0, is correlated with a particular dream

Colin may have had (payoff-irrelevant information), which could also affect Colin's decision to bet. Similarly, in Example 29, a firm may believe that the quantity produced by its competitor depends on the analysis carried out by its marketing department and that this (payoff-irrelevant) information may be correlated with the actual position of the demand function (θ_0).

As a special case, in which *interim* rationalizability admits an expressible characterization, we can consider **simple Bayesian Games**.[51]

Definition 45. *Fix a game with payoff uncertainty*

$$\hat{G} = \langle I, \Theta_0, (\Theta_i, A_i, u_i)_{i \in I} \rangle .$$

A Bayesian Game

$$BG = \langle I, \Omega, \Theta_0, \vartheta_0, (\Theta_i, T_i, A_i, \tau_i, \vartheta_i, p_i, u_i)_{i \in I} \rangle$$

based on $\hat{G}$ is ***simple*** *if $T_i = \Theta_i$ for each $i \in I$.*[52] *In this case, the functions $(\vartheta_i(\cdot) : T_i \to \Theta_i)_{i \in I}$ are trivially defined: $\vartheta_i = \mathrm{Id}_{\Theta_i}$ for each $i \in I$.*

Focusing on simple Bayesian games, one can show that *interim* rationalizability characterizes the behavioral implications of the following epistemic assumptions: (R) rationality, ($CIPI$) independence between players' actions and residual uncertainty conditional on players' private information, and (CB($R \cap CIPI$)) common belief of R and $CIPI$. Besides being interesting *per se*, simple Bayesian games are also widely used in the study of many relevant problems in information economics, such as adverse selection, moral hazard, mechanism design.

Now, let us turn to the other puzzling result: the gap between *ex-ante* and *interim* rationalizability. As already pointed out, the gap highlighted by Example 39 arises because the *interim* strategic form regards different types as different agents and, as a consequence, different types playing in Rowena's role may hold different beliefs concerning Colin's behavior (one of Rowena's type can believe Colin will play c with probability greater than $\frac{1}{2}$, while the other type can believe that this will happen only with probability

[51] They are slight generalization of "simple Bayesian games with type-independent beliefs" of Section 8.4 because, in the latter, beliefs about θ_{-i} are independent of θ_i, whereas in a simple Bayesian game $p_i \in \Delta(\Theta)$ and $p_i(\cdot|\theta_i)$ may depend on θ_i.

[52] More formally, the requirement is that T_i and Θ_i are isomorphic.

lower than $\frac{1}{2}$); since Rowena's conjecture concerning Colin's behavior can vary with her own type, her beliefs can exhibit correlation between the state of the world, ω (which determines residual uncertainty and players' type) and Colin's action. Instead, the *ex-ante* strategic form implicitly requires Rowena to regard Colin's decision function as independent from the state of the world.[53] To understand this, let us analyze the strategic-form expected payoff of player i when he plays "strategy" σ_i and holds conjecture $\mu^i \in \Delta(\Sigma_{-i})$ about the strategies of other players:

$$U_i(\sigma_i, \mu^i) = \sum_{\sigma_{-i} \in \Sigma_{-i}} \mu^i(\sigma_{-i}) U_i(\sigma_i, \sigma_{-i})$$

$$= \sum_{\sigma_{-i} \in \Sigma_{-i}} \mu^i(\sigma_{-i}) \sum_{\omega \in \Omega} p_i(\omega) u_i(\vartheta_0(\omega), \vartheta_i(\tau_i(\omega)), \vartheta_{-i}(\tau_{-i}(\omega)), \sigma_i(\tau_i(\omega)), \sigma_{-i}(\tau_{-i}(\omega)))$$

$$= \sum_{\omega \in \Omega} \sum_{\sigma_{-i} \in \Sigma_{-i}} p_i(\omega) \mu^i(\sigma_{-i}) \bar{U}_i(\omega, \sigma_i, \sigma_{-i}),$$

where we let $\bar{U}_i(\omega, \sigma_i, \sigma_{-i})$ denote the payoff of i when the state of the world is ω and profile (σ_i, σ_{-i}) is played. As the formula shows, the expected payoff of i is computed under the assumption that the probability of each pair (ω, σ_{-i}) is the product of the marginal probability of ω, $p_i(\omega)$ and the marginal probability of σ_{-i}, $\mu^i(\sigma_{-i})$. This means that i regards ω, the choice of nature, as independent of σ_{-i}, the choice of $-i$.

This implicit independence restriction reduces the set of Rowena's admissible beliefs and makes *ex-ante* rationalizability more demanding (hence, stronger) than *interim* rationalizability. Battigalli et al. [24] address this problem by defining **Ex-Ante Correlated Rationalizability**, a solution concept that allows correlation, according to players' subjective beliefs, between the initial chance move and players' decision functions. To introduce *ex-ante* correlated rationalizability, consider a simple Bayesian game; although such restriction is not necessary from the mathematical point of view, the *ex-ante* strategic form interprets types as actual information received by players concerning the initial chance move and, consequently, the assumption of simple Bayesian games

[53] If we interpret ω as the choice made by an external and neutral player called Nature, the independence restriction on beliefs is equivalent to requiring independence between the strategy chosen by Colin and the one chosen by Nature.

is the most, if not the only, reasonable one. For every i and $\mu^i \in \Delta(\Omega \times \Sigma_{-i})$, let[54]

$$r_i\left(\mu^i\right) = \arg\max_{\sigma_i \in \Sigma_i} \sum_{\omega,\sigma_{-i}} \mu^i\left(\omega,\sigma_{-i}\right) \hat{u}_i\left(\omega, \sigma_i\left(\tau_i\left(\omega\right)\right), \sigma_{-i}\left(\tau_{-i}\left(\omega\right)\right)\right).$$

Then we can recursively define the set of *ex-ante* correlated rationalizable decision functions as follows: for every i, $ACR_i^{0,BG} = \Sigma_i$ and for every $k \geq 1$,

$$ACR_i^{k,BG} = \left\{ \begin{array}{c} \sigma_i \in \Sigma_i : \exists \mu^i \in \Delta\left(\Omega \times \Sigma_{-i}\right), \\ 1)\ \sigma_i \in r_i\left(\mu^i\right),\ 2)\ \mathrm{marg}_\Omega \mu^i = p, \\ 3)\ \mu^i\left(\omega,\sigma_{-i}\right) > 0 \Longrightarrow \sigma_{-i} \in ACR_{-i}^{k-1} \end{array} \right\},$$

where $ACR_{-i}^{k-1} = \times_{j \neq i} ACR_j^{k-1}$.

Definition 46. *Fix a simple Bayesian game BG. A decision function for player i is* ex-ante *correlated rationalizable if* $\sigma_i \in ACR_i^{BG} = \bigcap_{k \geq 0} ACR_i^{k,BG}\left(\Sigma_{-i}\right)$.

From the iterative definition of *ex-ante* correlated rationalizability, it is easy to see that the justifying belief μ^i introduces correlation between the state of the world and opponents' behavior. Since different states ω can be associated with different types $t_i = \tau_i(\omega)$, which can hold different beliefs about the behavior of the other players, this correlation can also be introduced by *interim* correlated rationalizability. On the contrary, such correlation is not allowed by *ex-ante* rationalizability that forces each player to hold the same conjecture over co-players' decision functions independently of her type.

Thus, *interim* correlated rationalizability and *ex-ante* correlated rationalizability avoid any independence restriction implicitly entailed by the *interim* and *ex-ante* strategic form; as a consequence, both these solution concepts are characterized only by the epistemic assumptions of rationality and common belief in rationality given the background common knowledge of BG. Then, it should not come as a surprise that, if we use these solution concepts, the gap highlighted in Example 39 disappears and the following theorem holds.

[54] Recall that function $\hat{u}_i : \Omega \times A \to \mathbb{R}$ was introduced in Section 8.6 to analyze games with asymmetric information about an initial move by chance and that $\hat{u}_i$ can be related to u_i according to 8.6.1.

Theorem 36. *Let* $BG = \langle I, \Omega, \Theta_0, \vartheta_0, (\Theta_i, T_i, A_i, \tau_i, \vartheta_i, p_i, u_i)_{i\in I}\rangle$ *be a simple Bayesian game. Then, for every* i*,*

$$ACR_i^{BG} = \left\{\sigma_i \in \Sigma_i : \forall t_i \in T_i,\ \sigma_i(t_i) \in ICR_i^{BG}(t_i)\right\}.$$

Proof. We will prove that a stronger result holds, namely that for every i and every $k \geq 0$,

$$ACR_i^{k,BG} = \left\{\sigma_i \in \Sigma_i : \forall t_i \in \Theta_i,\ \sigma_i(t_i) \in ICR_i^{k,BG}(t_i)\right\}.$$

The result is trivial for $k = 0$. Now suppose that the result holds for every $s \leq k-1$; we will show that it holds also for k.

Take an agent i and an arbitrary $\sigma_i \in ACR_i^{k,BG}$. By definition, we can find a rationalizing belief $\mu^i \in \Delta(\Omega \times \Sigma_{-i})$. To ease notation, we keep denoting with symbols in square brackets the corresponding events in the relevant state space. Thus, let

$$[\theta_0] = \{(\omega, \sigma_{-i}) : \vartheta_0(\omega) = \theta_0\},$$

$$[t] = \{(\omega, \sigma_{-i}) : \tau(\omega) = t\},$$

$$[a_{-i}] = \{(\omega, \sigma_{-i}) : \sigma_{-i}(\tau_{-i}(\omega)) = a_{-i}\},$$

$$[\theta_0, t, a_{-i}] = [\theta_0] \cap [t] \cap [a_{-i}],$$

and

$$p([\theta_0, t]) = p(\{\omega : \vartheta_0(\omega) = \theta_0, \tau(\omega) = t\}).$$

For every $t_i \in T_i$, construct a function $\varphi_{t_i} : \Theta_0 \times T_{-i} \to \Delta(A_{-i})$ in the following way: if $p[\theta_0, t] > 0$,

$$\varphi_{t_i}(\theta_0, t_{-i})(a_{-i}) = \frac{\sum_{(\omega,\sigma_{-i})\in[\theta_0,t,a_{-i}]} \mu(\omega, \sigma_{-i})}{p([\theta_0, t])};$$

if $p([\theta_0, t]) = 0$,

$$\varphi_{t_i}(\theta_0, t_{-i})(a_{-i}) = \begin{cases} \frac{1}{\left|ICR_{-i}^{k-1,BG}(t_{-i})\right|}, & \text{if } a_{-i} \in ICR_{-i}^{k-1,BG}(t_{-i}), \\ 0, & \text{if } a_{-i} \notin ICR_{-i}^{k-1,BG}(t_{-i}). \end{cases}$$

For every t_i and pair (ω, a_{-i}), let

$$\hat{\mu}_{t_i}(\omega, a_{-i}) = p(\omega|t_i) \cdot \varphi_{t_i}(\vartheta_0(\omega), \tau_{-i}(\omega))(a_{-i}).$$

Observe that

$$\begin{aligned}\sum_{(\omega,\sigma_{-i})\in[\theta_0,t,a_{-i}]} \mu(\omega,\sigma_{-i}) &= p([\theta_0,t]) \cdot \varphi_{t_i}(\theta_0, t_{-i})(a_{-i}) \\ &= p(t_i)\, p(\theta_0, t_{-i}|t_i) \cdot \varphi_{t_i}(\theta_0, t_{-i})(a_{-i}).\end{aligned}$$

Thus,

$$\begin{aligned}&\sum_{\omega,\sigma_{-i}} \mu^i(\omega,\sigma_{-i})\, \hat{u}_i(\omega, \sigma_i(\tau_i(\omega)), \sigma_{-i}(\tau_{-i}(\omega))) \\ &= \sum_{\theta_0,t,a_{-i}} \sum_{(\omega,\sigma_{-i})\in[\theta_0,t,a_{-i}]} \mu(\omega,\sigma_{-i})\, \hat{u}_i(\omega, \sigma_i(\tau_i(\omega)), \sigma_{-i}(\tau_{-i}(\omega))) \\ &= \sum_{t_i} p(t_i) \sum_{\theta_0,t_{-i}} p(\theta_0, t_{-i}|t_i) \sum_{a_{-i}} \varphi_{t_i}(\theta_0,t_{-i})(a_{-i}) \cdot u_i(\theta_0, \vartheta_i(t_i), \vartheta_{-i}(t_{-i}), \sigma_i(t_i), a_{-i}).\end{aligned}$$

Since $\sigma_i \in r_i(\mu^i)$ and $p(t_i) > 0$, we conclude that $\sigma_i(t_i) \in r_i(t_i, \hat{\mu}_{t_i})$. Finally, if $p([\theta_0,t]) = 0$, by construction, $\varphi_{t_i}(\vartheta_0(\omega), t_{-i})(a_{-i}) > 0$ only if $a_{-i} \in ICR_{-i}^{k-1,BG}(t_{-i})$. If instead, $p([\theta_0,t]) > 0$, the same result follows from the definition of *ex ante* correlated rationalizability and the inductive hypothesis. Thus, for every t_i, we can conclude that $\sigma_i(t_i) \in ICR_i^{k,BG}(t_i)$ and, consequently,

$$ACR_i^{k,BG} \subseteq \left\{\sigma_i \in \Sigma_i : \forall t_i \in \Theta_i,\ \sigma_i(t_i) \in ICR_i^{k,BG}(t_i)\right\}.$$

Now we will prove the other inclusion. For every $t_i \in T_i$, let (a_{t_i}, t_i) be a pair such that $a_{t_i} \in ICR_i^{k,BG}(t_i)$. Thus for every a_{t_i} we can find a rationalizing belief $\mu_{a_{t_i}} \in \Delta(\Omega \times A_{-i})$ and a function $\varphi_{a_{t_i}} : \Theta_0 \times T_{-i} \to \Delta(A_{-i})$ satisfying the properties stated in the iterative definition of *interim* correlated rationalizability. Now construct belief $\hat{\mu}_i \in \Delta(\Omega \times \Sigma_{-i})$ as follows: for every pair (ω, σ_{-i}),

$$\hat{\mu}_i(\omega,\sigma_{-i}) = \frac{p(\omega) \cdot \varphi_{a_{\tau_i(\omega)}}(\vartheta_0(\omega), \tau_{-i}(\omega))(\sigma_{-i}(\tau_{-i}(\omega)))}{\left|\left\{\sigma'_{-i} \in ACR_{-i}^{k-1,BG} : \sigma'_{-i}(\tau_{-i}(\omega)) = \sigma_{-i}(\tau_{-i}(\omega))\right\}\right|}$$

whenever the denominator is positive,[55] and $\hat{\mu}_i(\omega,\sigma_{-i})=0$ otherwise. It is immediate to see that $\hat{\mu}_i(\omega,\sigma_{-i})>0$ only if $\sigma_{-i}\in ACR_{-i}^{k-1,BG}$. Now, let

$$[\omega,a_{-i}]=\left\{\left(\omega',\sigma_{-i}\right):\omega'=\omega,\ \sigma_{-i}\left(\tau_{-i}\left(\omega\right)\right)=a_{-i}\right\}.$$

With this, if $\left\{\sigma'_{-i}\in ACR_{-i}^{k-1,BG}:\sigma'_{-i}\left(\tau_{-i}\left(\omega\right)\right)=a_{-i}\right\}\neq\emptyset$, then

$$\sum_{(\omega',\sigma_{-i})\in[\omega,a_{-i}]}\hat{\mu}_i\left(\omega',\sigma_{-i}\right)=$$

$$\frac{\sum_{\sigma_{-i}\in ACR_{-i}^{k-1,BG}:\sigma_{-i}(\tau_{-i}(\omega))=a_{-i}}p\left(\omega\right)\cdot\varphi_{a_{\tau_i(\omega)}}\left(\vartheta_0\left(\omega\right),\tau_{-i}\left(\omega\right)\right)\left(a_{-i}\right)}{\left|\left\{\sigma'_{-i}\in ACR_{-i}^{k-1,BG}:\sigma'_{-i}\left(\tau_{-i}\left(\omega\right)\right)=a_{-i}\right\}\right|}=$$

$$p\left(\omega\right)\cdot\varphi_{a_{\tau_i(\omega)}}\left(\vartheta_0\left(\omega\right),\tau_{-i}\left(\omega\right)\right)\left(a_{-i}\right).$$

Also notice that for every (ω,a_{-i}), if $\left\{\sigma'_{-i}\in ACR_{-i}^{k-1,BG}:\sigma'_{-i}\left(\tau_{-i}\left(\omega\right)\right)=a_{-i}\right\}=\emptyset$, then the inductive hypothesis implies that $a_{-i}\notin ICR_{-i}^{k,BG}\left(\tau_{-i}\left(\omega\right)\right)$ and, consequently, $\varphi_{a_{\tau_i(\omega)}}\left(\vartheta_0\left(\omega\right),\tau_{-i}\left(\omega\right)\right)\left(a_{-i}\right)=0$. Thus, for every (ω,a_{-i}),

$$\sum_{(\omega',\sigma_{-i})\in[\omega,a_{-i}]}\hat{\mu}_i\left(\omega,\sigma_{-i}\right)=p\left(\omega\right)\cdot\varphi_{a_{\tau_i(\omega)}}\left(\vartheta_0\left(\omega\right),\tau_{-i}\left(\omega\right)\right)\left(a_{-i}\right)$$

and

$$\begin{aligned}\sum_{\sigma_{-i}}\hat{\mu}_i\left(\omega,\sigma_{-i}\right)&=\sum_{a_{-i}}\sum_{(\omega',\sigma_{-i})\in[\omega,a_{-i}]}\hat{\mu}_i\left(\omega,\sigma_{-i}\right)\\&=\sum_{a_{-i}}p\left(\omega\right)\cdot\varphi_{a_{\tau_i(\omega)}}\left(\vartheta_0\left(\omega\right),\tau_{-i}\left(\omega\right)\right)\left(a_{-i}\right)=p\left(\omega\right).\end{aligned}$$

Notice that we have already proved two of the properties in the definition of *ex-ante* correlated rationalizability. Finally, for every $\sigma_i\in\Sigma_i$

$$\sum_{\omega,\sigma_{-i}}\hat{\mu}^i\left(\omega,\sigma_{-i}\right)\hat{u}_i\left(\omega,\sigma_i\left(\tau_i\left(\omega\right)\right),\sigma_{-i}\left(\tau_{-i}\left(\omega\right)\right)\right)$$

$$=\sum_{\omega,a_{-i}}p\left(\omega\right)\cdot\varphi_{a_{\tau_i(\omega)}}\left(\vartheta_0\left(\omega\right),\tau_{-i}\left(\omega\right)\right)\left(a_{-i}\right)\cdot u_i\left(\vartheta_0\left(\omega\right),\vartheta_i\left(\tau_i\left(\omega\right)\right),\vartheta_{-i}\left(\tau_{-i}\left(\omega\right)\right),\sigma_i\left(\tau_i\left(\omega\right)\right),a_{-i}\right).$$

[55] That is, if $\left\{\sigma'_{-i}\in ACR_{-i}^{k-1,BG}:\sigma'_{-i}\left(\tau_{-i}\left(\omega\right)\right)=\sigma_{-i}\left(\tau_{-i}\left(\omega\right)\right)\right\}\neq\emptyset$.

Since $p\,(t_i) > 0$ for every t_i, the decision function σ_i defined by $\sigma_i\,(t_i) = a_{t_i}$ for every t_i is such that $\sigma_i \in r_i\left(\hat{\mu}^i\right)$. We conclude that $\sigma_i \in ACR_i^{k,BG}$. Since pairs (a_{t_t}, t_i) were chosen arbitrarily, we can write

$$ACR_i^{k,BG} \supseteq \left\{\sigma_i \in \Sigma_i : \forall t_i \in \Theta_i,\ \sigma_i\,(t_i) \in ICR_i^{k,BG}\,(t_i)\right\}.$$

The statement of the theorem follows by induction. ■

8.8.2 The Electronic Mail Game

Consider the following incomplete information game where *player 1* (Rowena) *knows the true payoff matrix*, whereas 2 (Colin) does not know it. The probabilities and the payoffs are represented in the table below:

$(1-\lambda)$:

θ^α	a	b
a	M, M	$1, -L$
b	$-L, 1$	$0, 0$

λ:

θ^β	a	b
a	$0, 0$	$1, -L$
b	$-L, 1$	M, M

$$L > M > 1, \qquad \lambda < 1/2$$

If we had complete information, (a, a) would be the dominant strategy equilibrium for matrix θ^α, and (b, b) would be the Pareto efficient equilibrium for the matrix θ^β. The second matrix has also another equilibrium, (a, a).

Notice that the game has distributed knowledge; thus, our previous analysis implies that the set of *interim* rationalizable actions is equivalent to the set of *interim* correlated rationalizable actions. In particular, it can be easily verified that any Bayesian game representing this situation would have only one *interim* rationalizable profile: (aa, a). As a matter of fact, for type θ^α of Rowena a is dominant. If Colin believes that Rowena is rational, so that type θ^α chooses a, then $a_2 = a$ yields a larger expected utility than $a_2 = b$. More precisely let $\mu^2 \in \Delta(aa, ab, ba, bb)$ be Colin's conjecture about Rowena's strategy and let $u_2(\mu^2, a_2)$ be the expected utility of action $a_2 \in \{a, b\}$ given conjecture μ^2. Since $\mu^2(ba) = \mu^2(bb) = 0$

(as θ^α certainly chooses a), we obtain

$$\begin{aligned} u_2(\mu^2, a) &= (1-\lambda)u_2(\theta^\alpha, a, a) + \lambda\mu^2(aa)u_2(\theta^\beta, a, a) + \lambda\mu^2(ab)u_2(\theta^\beta, b, a) \\ &\geq (1-\lambda)M \\ &> -(1-\lambda)L + \lambda M \\ &\geq (1-\lambda)u_2(\theta^\alpha, a, b) + \lambda\mu^2(aa)u_2(\theta^\beta, a, b) + \lambda\mu^2(ab)u_2(\theta^\beta, b, b) \\ &= u_2(\mu^2, b). \end{aligned}$$

Hence, the only rationalizable action for Colin is a. This implies that also for type θ^β the only *interim* rationalizable choice is a.

Consider now the following more complex variation of the game.[56] Obviously the players would find it convenient to communicate and coordinate on (a, a) if $\theta = \theta^\alpha$ and on (b, b) if $\theta = \theta^\beta$. Let us then consider the following form of communication via electronic mail. Rowen and Colin sit in front of their respective computer screens. If the payoff-type of Rowena is θ^β, her computer, C_1, *automatically* sends a message to the computer of Colin, C_2. Furthermore, if computer C_i receives a message, it *automatically* sends a confirmation message to computer C_{-i}. However, any time a message is sent,it gets *lost* with probability ε and thus it does not reach the receiver.

This information structure yields a corresponding Bayesian game. A generic state of the world is given by a pair of numbers $\omega = (q, r)$ where $q = t_1$ is the number of messages *sent* by C_1 and $r = t_2$ is the number of messages received (and therefore also sent) by C_2. Hence, the set of states of the world is given by

$$\begin{aligned} \Omega &= \{(0,0), (1,0), (1,1), (2,1), (2,2), (3,2), (3,3), ...\} \\ &= \{(q, r) \in \mathbb{N} \times \mathbb{N} : q = r \text{ or } q = r + 1\}. \end{aligned}$$

If the state is $(q, q-1)$ it means that the last message (among those sent either by C_1 or by C_2) has been sent by C_1. If the state is (r, r) (with $r > 0$) it means that the last message sent by C_1 has reached C_2, but the confirmation by C_2 has not reached C_1. The signal functions are $\tau_1(q, r) = q$, $\tau_2(q, r) = r$. The function that determines θ is $\vartheta_1(0) = \theta^\alpha$, $\vartheta_1(q) = \theta^\beta$ if $q > 0$.[57] There is a common prior given by $p(0, 0) = (1 - \lambda)$,

[56] See [73] (or the textbook by Osborne and Rubinstein [65], pp. 81-84). The analysis in terms of *interim* rationalizability is not contained in the original work.

[57] Θ_2 is a singleton, hence ϑ_2 is a constant.

$p(r+1,r) = \lambda(1-\varepsilon)^{2r}\varepsilon$, $p(r+1,r+1) = \lambda(1-\varepsilon)^{2r+1}\varepsilon$ for any $r \geq 0$ (that is, $p(q,r) = \lambda(1-\varepsilon)^{q+r-1}\varepsilon$ for any $(q,r) \in \Omega \backslash \{(0,0)\}$). This information is summed up in the following table:

$t_1{=}q \backslash t_2{=}r$	0	1	2	3	4	5	...
0, θ^α	$1-\lambda$	–	–	–	–	–	...
1, θ^β	$\lambda\varepsilon$	$\lambda(1-\varepsilon)\varepsilon$	–	–	–	–	...
2, θ^β	–	$\lambda(1-\varepsilon)^2\varepsilon$	$\lambda(1-\varepsilon)^3\varepsilon$	–	–	–	...
3, θ^β	–	–	$\lambda(1-\varepsilon)^4\varepsilon$	$\lambda(1-\varepsilon)^5\varepsilon$	–	–	...
4, θ^β	–	–	–	$\lambda(1-\varepsilon)^6\varepsilon$	$\lambda(1-\varepsilon)^7\varepsilon$	–	...
5, θ^β	–	–	–	–	$\lambda(1-\varepsilon)^8\varepsilon$	$\lambda(1-\varepsilon)^9\varepsilon$	...
...	...	...	...	...	...	...	...

The resulting beliefs, or conditional probabilities, are then determined as follows:

$$\begin{aligned}
p_1[(0|0) &= 1, \\
p_1[(r|r+1) &= \frac{\lambda(1-\varepsilon)^{2r}\varepsilon}{\lambda(1-\varepsilon)^{2r}\varepsilon + \lambda(1-\varepsilon)^{2r+1}\varepsilon} = \frac{1}{2-\varepsilon}, \\
p_2[(0|0) &= \frac{1-\lambda}{1-\lambda+\lambda\varepsilon}, \\
p_2[(r+1|r+1) &= \frac{\lambda(1-\varepsilon)^{2r+1}\varepsilon}{\lambda(1-\varepsilon)^{2r+1}\varepsilon + \lambda(1-\varepsilon)^{2r+2}\varepsilon} = \frac{1}{2-\varepsilon} \quad (r > 0).
\end{aligned}$$

In other words, any player—having received a certain number of messages—computes the probability that his confirmation message does not reach the other computer, which is $\frac{1}{2-\varepsilon} > \frac{1}{2}$.

If ε is very small, the Bayesian game is in some sense "close" to the game with common knowledge of θ. First, notice that in all states (q,r) with $r > 0$ both players know that the true state is θ^β. Moreover, for any n and any $\delta > 0$, there always exists an ε sufficiently small such that, given $\theta = \theta^\beta$, the probability that there exists mutual knowledge of degree n of the true value of θ is bigger than $1-\delta$.[58] Given that (b,b) is an equilibrium of the complete information game in which $\theta = \theta^\beta$, we could be induced

[58] For instance, consider $n = 2$. For any $q \geq 2$, in state (q,r) ($r = q$, or $r = q-1$) every player knows that $\theta = \theta^\beta$ and every player knows that the other knows it, therefore

to think that if ε is very small, then action b is *interim* rationalizable for some type of Rowena and Colin. But the following result shows that this intuition is incorrect:

Proposition 2. *In the Electronic Mail game there exists a unique* interim *rationalizable profile: every type of every player chooses action a.*

The crucial point in the proof is to realize that when a player i knows that $\theta = \theta^\beta$, but s/he is sure that $-i$ plays a whenever s/he has not received her/his last message, then i prefers to play a. On the other hand, it is easy to show that when i *does not know* that $\theta = \theta^\beta$ then a is the only rationalizable action. It follows by induction that a is the only rationalizable action in all states. Here is the formal argument:

Proof of Proposition 2. Clearly, the only rationalizable action for $t_1 = 0$ is a, as this type of Rowena knows that in the true payoff matrix a is dominant. An argument similar to the one used for the simple case discussed above shows that the only rationalizable action for $t_2 = 0$ is also a.

Consider now the rationalizable choices for the types $t_i = r > 0$, that is for those types for which player i *knows that* $\theta = \theta^\beta$. We first prove two intuitive intermediate steps:

(i) *If the only rationalizable action for type $t_2 = r - 1$ is a, then the only rationalizable action for type $t_1 = r$ is a.* Any rationalizable action for $t_1 = r$ needs to be justified by a rationalizable conjecture about Colin (see Theorem 2). By assumption, such conjecture must assign probability 1 to the set of profiles $\sigma_2 \in \{a, b\}^{T_2}$ such that $\sigma_2(r-1) = a$. Hence, the expected utility for type $t_1 = r$ if he chooses a is at least 0 (this value is realized if type $t_2 = r$ also chooses a), whereas the expected utility from choosing b is at most $-p_1(r-1|r)\,L + p_1(r|r)\,M$ (this value is realized if type $t_2 = r$ chooses b). Since $p_1(r-1|r) = \frac{1}{2-\varepsilon} > \frac{1}{2}$ and $L > M$, it follows that $0 > -p_1(r-1|r)\,L + p_1(r|r)\,M$.

An analogous argument[59] shows that:

there is mutual knowledge of degree two of the true payoff matrix. The probability of being in one of these states conditional on the event $\theta = \theta^\beta$ is

$$1 - \frac{p(\{(1,0),(1,1)\})}{\lambda} = 1 - 2\varepsilon + \varepsilon^2.$$

To guarantee that such probability is larger than $1 - \delta$ we need $\varepsilon < 1 - \sqrt{1-\delta}$.

[59] Just reverse the roles and modify indexes accordingly.

(ii) *If the only rationalizable action for type* $t_1 = r$ *is* a*, the only rationalizable action for type* $t_2 = r$ *is also* a*.*

From steps (i) and (ii) it follows that, for every $r > 0$, *if the only rationalizable profile in state* $(r-1, r-1)$ *is* (a, a) *then the only rationalizable profile in state* $(r, r-1)$ *is* (a, a)*; if the only rationalizable profile in state* $(r, r-1)$ *is* (a, a) *then the only rationalizable profile in state* (r, r) *is* (a, a). Since we have shown that the only rationalizable profile in state $(0, 0)$ is (a, a), it follows by induction that (a, a) is the only rationalizable profile in every state. ■