Robust Predictions in Dynamic Games

A Thesis Submitted to the Faculty of The Graduate School of The University of Minnesota

by

DAVID URBANO RUIZ GOMEZ

in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Advisor: Aldo Rustichini

August, 2018

Copyright 2018 ©David Ruiz G.

All Rights Reserved

Acknowledgement

I owe infinite thanks to my advisors: Aldo Rustichini and David Rahman, for everything: their time, patience, countless discussions, but especially for all the encouragement and dedication during the preparation of my dissertation. I also want to thank Jan Werner for his continuous support, dedication, and delightful conversations about the topic of my thesis; and Peter Polacik, who provide me with guidance during my time as a graduate student in mathematics.

I am very thankful to the economics department staff at the University of Minnesota, especially to Wendy Williamson and Caty Bach for all their help and support throughout my graduate study. I extend my gratitude to Bonny Fleming.

I want to thank my family. My admiration is to them. Mom, for your endless love; Dad, for teaching me that we never stop learning; Dani and Joaco, for showing me their passions for life, their laughter, and for being here with us; and to Caro, because my heart and soul belongs to you.

I dedicate my work to my family:

My beloved parents Verónica Gómez F. and Fernando Ruiz M., my siblings Daniela Ruiz G., and Joaquín Ruiz G, and my wife Carolina Ortiz C.

My infinite love is to them.

They are the reason for who I am; they are the reason to know where I go.

Abstract of the Dissertation

Robust Predictions in Dynamic Games

This dissertation is about understanding the robustness property of predictions to misspecification of higher-order beliefs in dynamic games with payoff uncertainty. In particular, it asks: Which simplifying assumptions about beliefs provide robust predictions in dynamic games? The most important result of this dissertation, presented in the second chapter, is to show that lack of robustness is a generic property of predictions consistent with (interim) sequential rationalizability (*ISR*) unless the prediction is unique. I consider this to be an essential and novel contribution to the literature of robustness in game theory since it challenges the validity of the standard approach to modeling uncertainty in dynamic games because it gives rise, for almost every model of uncertainty, to spurious predictions.

Typically, when analyzing a model, different parameters represent different assumptions of the model, and therefore, predictions from the model are sensitive to the specification of those parameters. For example, it is well known that in the standard Bayesian approach to games with incomplete information, a crucial parameter that requires to be specified, and at the same time is neither observable nor verifiable without any error from the point of view of researchers, is players beliefs and hierarchies of beliefs; hence, because of the previous observation, it happens that in many applications hierarchies of beliefs encode strong (informational) assumptions, and as I already mentioned, behavioral predictions (e.g., in the form of Perfect Bayesian Equilibrium, Interim Sequential Rationalizability, among others) depend on those assumptions; moreover, in some cases, this dependence can be very sensitive at the tails of the hierarchies of beliefs specified in the model.

The robustness property refers, in this case, to the possibility of guarantee that slight changes in the specification of the parameters do not lead to significant changes in predictions, since at least from a methodological point of view, if the property holds generically it would provide a justification for the validity of the standard approach to model uncertainty.

One approach to understanding the robustness property of set-valued solution concepts in games is to ask: Which predictions remain valid after all common certain-belief assumptions are relaxed?

Penta (2012) have shown that in (finite) dynamic games with incomplete information the only predictions that remain valid after relaxing all the assumptions about beliefs and hierarchies of beliefs are those consistent with (interim) sequential rationalizability. In other words, *ISR* is the strongest solution concept such that, for every model of beliefs it is possible to guarantee that an outcome that was ruled out by *ISR* is ruled out for every approximation of the model. This result implies lack of robustness of any refinement of *ISR* as, for example, any of the familiar

equilibrium concepts.

In this dissertation, a stronger notion of robustness is considered, that is if in addition, it is possible to guarantee that there are no spurious predictions, in the sense that for every predicted outcome of a solution concept there is no approximation ruling that outcome out. This last notion is formalized through a notion of full continuity of predictions with respect to beliefs and hierarchies of beliefs.

An approach to this question for static games was given in Ely and Peski (2011). They introduced the concept of critical types as precisely those assumptions on beliefs that are vulnerable to misspecification, that is, as those types to which there are spurious predictions consistent with rationalizability. They showed that critical types are non-generic (rare). The key argument in Ely and Peski's result exploits the fact that, in static games, rationalizability does not depend on the timing of the arrival of players' information.

However, in dynamic games, *ISR* does depend on the timing of information. In particular, players beliefs are restricted only at the beginning of the game, and via conditioning whenever possible. However, at zero probability events, conditional beliefs are unrestricted. I exploit this observation to show that Ely and Peski's result does not hold in dynamic settings: lack of robustness is a generic property of *ISR* whenever it delivers multiple predictions.

As *ISR* often delivers multiple predictions in applications, this result casts doubts on the interpretation and validity of solution concepts such as Perfect Bayesian Equilibrium, Sequential Equilibrium, and *ISR* itself. By acknowledging model misspecification of higher-order beliefs, there is no type in Harsanyi's framework at which a researcher can guarantee that no slight perturbation on the modeling assumptions exists which rules some prediction out unless the prediction is unique.

Finally, I propose an ongoing research agenda in the problem of robust predictions in dynamic games. In particular, we consider dynamic games with payoff uncertainty and, as in Siniscalchi (2016a,b), assume that players in the game choose strategies according to structural rationality. Players with structural preferences induce, at the beginning of the game, a collection of alternative hypothesis about how the game is going to unfold; and rank any two strategies depending on the expected payoff under those alternative priors in a lexicographic way. A strategy is structurally rational if it is maximal. We propose to study general properties of (weak) interim structural rationalizability (*IStR*), a solution concept that characterizes the behavioral implications of common certainty in structural rationality.

In the case of Bayesian dynamic games with incomplete information, we are particularly interested in the robustness properties of *IStR* to perturbations of higher-order beliefs. As illustrated by an example, three results are conjecture: a structure theorem of structural rationalizability, a characterization of critical types, and a non-generic result of the set of critical types.

Contents

List of Figures

List of Tables

Chapter 1

Robust Predictions in Game Theory

Game-theoretic models are typically used to represent an idealized strategic interaction and to provide reasonable behavioral predictions for the situation of interest. It involves formalizing different simplifying assumptions that typically are neither observable nor verifiable without any error. For example, researchers use "*type spaces*" to represent assumptions on players beliefs about the state, as well as to represent assumptions on players higher-order beliefs about the state; i.e., to specify how players interactively reason about payoff-relevant states. At the same time, the set of conditions that describes reasonable behavioral predictions impose additional restrictions on how players reason, in addition, about strategic uncertainty, that is: about the behavior of a player, about what a player believes about the behavior of his opponents, and so on. Rationalizability (see Bernheim (1984), Pearce (1984), and Dekel et al. (2007)), for example, characterizes the set of actions that are consistent with the event of "*rationality and common belief in rationality*". That is, those actions taken by a rational player, who believes that his opponents are rational, believes (with probability one) that his opponents believe that their opponents are rational, and so on.

Most applications in game theory consider simple type spaces, as described in Harsanyi's (1967-68) model of beliefs based on "*types*", and formalized later by Mertens and Zamir (1985). For each player, a type is an object that summarizes all payoff-relevant parameters and beliefs about the physical aspect of a game, and implicitly encodes assumptions about players' infinite hierarchy of beliefs. For example, in a two-period reputation game (Kreps and Wilson (1982b), and Milgrom and Roberts (1982b)) firm 1 has two potential types: "*sane*" and "*crazy*". Both types differ on some payoff-relevant parameter even though each type of firm 1 believes with probability one that firm 2 is of a unique type. Firm 2, on the other hand, is of a unique type that believes with probability p that firm 1 is "*sane*", and probability $1-p$ that firm 1 is "*crazy*". Alternatively, in a repeated public-good game two players simultaneously decide whether to contribute to the period-t public good. While benefits of the public good are common knowledge, a type of a player is its cost of contributing, which is private information, and every player's type believes that the costs are drawn independently from the same distribution.

Therefore, many models rely on reasonable approximations to the situation of interest by specifying only finitely many levels of players hierarchy of beliefs. However, due to unavoidable approximation errors of the underlying informational assumptions, researchers have a particular interest in models of beliefs that provide robust predictions in the following sense: slight relaxations of those implicit common-knowledge assumptions encoded in a type space do not lead to drastic changes in behavioral predictions. The robustness property of predictions to misspecification of higher-order beliefs is a fundamental criterion to justify the validity of a model of uncertainty within a theory of strategic interactions. As long as players' actual beliefs and higher-order beliefs about payoffs cannot be observed without any error, the robustness property guarantees that similarity in beliefs corresponds to a similarity in predicted behavior, while ruling out the conceptual problem that some spurious prediction may arise. Therefore, this paper asks: for which type spaces Interim Sequential Rationalizability (*ISR*)[1] provides robust behavioral predictions in the class of dynamic games with incomplete information? How can an analyst know in advance if predictions from *ISR* depend too much on the exact details of the type space chosen? Moreover, for which type spaces *ISR* delivers predictions that are too sensitive to these misspecifications?

One approach to understand the robustness property in Bayesian dynamic games of incomplete information was given in Penta (2012). In particular, he asked which predictions are still valid when all common knowledge assumptions about payoffs are relaxed. His main result, the structure theorem for *ISR*, shows that if all common knowledge assumptions on payoffs are relaxed, then *ISR* is the strongest solution concept among the ones consistent with the event of *RCIBR* whose predictions survive perturbations of hierarchies of beliefs. This result, which relies on a richness condition based on conditional dominance and the fact that *ISR* is upper hemicontinuous (*uhc*) in the universal type space, explain why in this paper we focus on *ISR*.

Even though the structure theorem implies a lack of robustness of any stronger solution concept than *ISR* (e.g., equilibrium concepts such as Perfect Bayesian Equilibrium, or Sequential Equilibrium) I will argue later that it is still a weak criterion of robustness, especially for dynamic

[1] See Penta (2012). *ISR* is a solution concept that characterizes the set of *all* the strategies that are consistent with the event of *sequential rationality and common initial belief of sequential rationality (RCIBR)* (see Ben-Porath (1997), Battigalli and Siniscalchi (1999, 2003)).

games. Observe that from the structure theorem you can guarantee that every prediction that is ruled out for the original type space, it is ruled out for every approximation to the original type space. In other words, that for a given type in a model there are not too few predictions compared to predictions of approximated models: for every approximated type, every prediction for the approximation is a prediction for the original type. Is in this sense that any refinement of *ISR* is not robust to perturbations of higher order beliefs: an analyst cannot guarantee that a prediction that is ruled out by the refinement of *ISR* in a model, is ruled out in every approximation.

A stronger criterion is for predictions to satisfy a notion of strategic robustness at a type space, in the sense that it is possible to guarantee that similarity in beliefs corresponds to a similarity in predicted behavior. In this paper we are interested in this last property, since is the one we consider to be a fundamental criterion for a theory of strategic interactions. As long as the property holds, the use of type spaces to formalize non-verifiable and non-observable assumptions on players beliefs is justified.

Following the argument from above, strategic robustness would hold if, in addition, it is possible to guarantee that for every prediction for a type there is no approximation ruling that prediction out when all common knowledge assumptions about payoffs are relaxed. This last intuition corresponds to a notion of lower hemicontinuity. Note that failure of strategic robustness raises the conceptual problem that some spurious prediction may arise, since there is an approximated model that rules that prediction out, which clearly cast doubts on the interpretation and validity of a given prediction for a type.

If *ISR* were to satisfy strategic robustness at *almost every type space*, then a researcher would be able to guarantee that behavioral predictions consistent with *ISR* do not depend on specific details of the type space chosen. In other words, *ISR* would be a satisfactory solution concept to analyze dynamic games with incomplete information, in the sense of providing predictions that are robust to misspecification of higher-order of beliefs. On the other hand, we say that a players' type in a type space is critical if *ISR* does not satisfy strategic robustness at that type. That is, if an analyst cannot guarantee that for every prediction of that type there is no arbitrary small perturbation of hierarchies of beliefs ruling that prediction out. The first result in the next chapter provides a characterization of critical types, describing sufficient and necessary conditions for which simplifying assumptions about beliefs imply a spurious prediction.

The second result of my dissertation shows that the set of critical types for a given dynamic game is not negligible[2] and, under certain conditions, it is dense relative to the set of types with multiple *ISR*-strategies. In other words, lack of strategic robustness is a generic property of *ISR*:

[2]Technically, it is shown not to be meager. Mathematically, a set is nowhere dense if its closure has empty interior. A set is meager (or of first-category), if it is a countable union of nowhere dense sets. This concept formalizes the idea of a set being negligible relative to an ambient space.

for almost every model of beliefs for which *ISR* delivers multiple predictions, a researcher cannot guarantee that those predictions do not depend on the exact details of the model. The above means that, by acknowledging model misspecification of higher-order beliefs, the only robust predictions are those derived from assumptions with only one *ISR* prediction. The results from this paper imply that *ISR* is not a satisfactory theory of strategic interaction in dynamic settings since it is extremely irregular and highly sensitive to misspecification of higher-order beliefs in most interesting cases. To appreciate the importance of these results, we need to see where it stands in the context of the existing literature. Then, we introduce an example to motivate the original question in this paper and illustrate the fundamental aspects of the results.

1.1 Related Literature

The problem of informational robustness about beliefs goes back to the early 90's where several papers provide a partial answer. The main result was that predictions from equilibrium solution concepts are highly sensitive to informational assumptions, in particular, to relaxations of common knowledge (see, for example, Rubinstein (1989), and the seminal paper on global games of Carlsson and Van Damme (1993)). It was not until Weinstein and Yildiz (2007) that the field reached a better understanding of this issue. WY showed that in static games the only predictions that are still valid when informational assumptions are relaxed, are the ones provided by rationalizability, as defined in Dekel et al. (2007). Both Chen (2012), and Penta (2012), extend that result to dynamic (extensive-form) games. These last few papers justify why to focus on rationalizability concepts when it comes to studying questions about perturbations of beliefs.

Ely and Peski (2011) introduced to this literature the concept of critical types. This particular paper is of fundamental importance for many of the ideas discussed here. They study *(Interim) Correlated Rationalizability (ICR)* and focus on static games, which is a significant restriction on the physical description of the game. In static games, rationalizability does not depend on the timing of the arrival of players' information for decision making. Their main contribution exploits this fact to show that, not only for a given game but also after quantifying across all games, *ICR* satisfies a notion of strategic robustness for almost every type, in the sense that non-critical types are generic in the universal type space, and we can consider the set of critical types to be negligible.

However, many social and economic interactions in real-life situations occur sequentially, there is a transmission of information, and players learn about the state as the game unfolds. Hence, in dynamic games, *ISR* does depend on the timing of arrival of new information for decision making. I exploit this observation to show that Ely and Peski's result does not hold in dynamic settings: lack of robustness is a generic property of *ISR* whenever it delivers multiple predictions.

I have emphasized that type spaces with critical types are problematic for a researcher because predictions would be highly sensitive to the implied informational assumptions. We learn from Ely and Peski (2011) that when it comes to describing static strategic interactions, the set of critical types (those carrying the problematic assumptions for a researcher) is negligible. However, this paper shows that when it comes to describing dynamic strategic interactions critical types represent a non-negligible set, and in particular, it is dense relative to those types with multiplicity, so that Ely and Peski's result is overturned. We offer an interpretation of this result in Section (4) of the paper, contrasting the implications discussed in the last paragraph concerning Ely and Peski's result.

1.2 (Interim) Sequential Rationalizability: An Example with common certainty.

The following example will help us introduce the important objects under study: dynamic games in extensive-form, conditional probability systems, type spaces, *ISR*, and the distinction between critical types in dynamic versus static games, as well as to give some intuition for the results of this dissertation.

Example 1: Consider the following game in extensive form, taken from Penta (2012). There are two players, named 1 and 2. Player 1 chooses, at the beginning of the game, either a_1 after which the game is over, or *In* after which he faces player 2 in a simultaneous move game. There are two possible states of nature $\Theta_0 = \{\theta_0, \theta_0^*\}$. Figure 1.1 describes the rules of this game and payoffs for each player at θ_0.

To have a complete formal description of a dynamic game with incomplete information, an analyst needs to include a type space. Formally, a type space is a structure $\mathbf{Z} = \{\Theta_0, (Z_i, \beta_i)_{i \in N}\}$, where Z_i is a set of types for player i, and $\beta_i : Z_i \to \Delta(\Theta_0 \times Z_{-i})$ assigns to each type of player i a probability measure over states of nature and opponents' type. It describes, at the beginning of the game, exogenous assumptions on players subjective beliefs about payoff-relevant states. As it is explained later in the paper, each type induces an explicit hierarchy of beliefs: a belief about the state of nature (*first-order belief*), beliefs about the state of nature and opponents' beliefs about the state of nature (*second-order beliefs*), and so on. For example, a particular type space is the one representing a situation of complete information, that is, when there is common certain belief, at the beginning of the game, that the state of nature is θ_0: each player believes with probability one that the state is θ_0, believes with probability one that the state is θ_0 and his opponent believes with probability one that the state is θ_0, and so on. Table 1 describes a type space that induces a situation of common knowledge over the state θ_0. For each player there is only one type, say z_i, and the probability measure associated with each type assigns probability one to the state of nature being θ_0 and opponents' type being z_{-i}.

Figure 1.1: Top figure, extensive-form representation with payoff at θ_0. Bottom figures, normal-form representations conditional on each history.

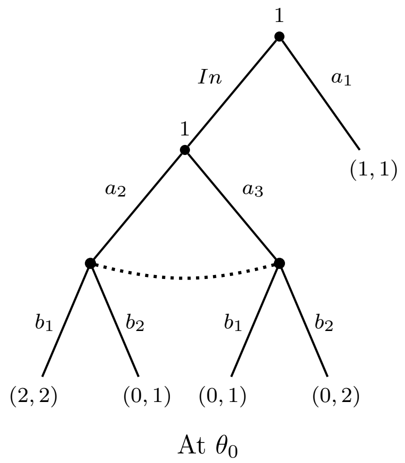

$[\emptyset]$ - $\{1\}$	b_1	b_2
a_1	(1,1)	(1,1)
(In, a_2)	(2,2)	(0,1)
(In, a_3)	(0,1)	(0,2)

$[In]$ - $\{1,2\}$	b_1	b_2
(In, a_2)	(2,2)	(0,1)
(In, a_3)	(0,1)	(0,2)

Table 1.1: A type space representing complete information over θ_0

$Z_1 = \{z_1\}$	(θ_0, z_2)	$Z_2 = \{z_2\}$	(θ_0, z_1)
$\beta_1[z_1]$	1	$\beta_2[z_2]$	1

The game is analyzed by studying behavioral predictions derived from *ISR*, a solution concept defined player-by-player, and for each type, as a procedure of iterated elimination: at each round of the procedure, a strategy survives if it is a sequential best reply to a conditional conjecture that satisfies: Bayes updating whenever possible; it is supported, at the beginning of the game, by opponents' strategies that haven't been eliminated in previous rounds; and, it coincides, at the beginning of the game, with type's beliefs.

Back to the example, at the first round, strategy (In, a_3) is strictly dominated by a_1 for player 1 of type z_1. Given this, player 2's initial conjectures needs to put probability zero on (In, a_3). However, as it was described above, no further restrictions are imposed on conditional conjectures, besides Bayes updating whenever possible. For example, let an initial conjecture be $\hat{\mu}_2[\emptyset](\theta_0, z_1, a_1) = 1$, so there is zero probability on reaching node $[In]$. Since Bayes updating is not possible, conditional conjectures of player 2 are unrestricted and $\hat{\mu}_2[In](\theta_0, z_1, (In, a_3)) = 1$ makes b_2 a sequential best reply to $\hat{\mu}_2$. Alternatively, let $\hat{\mu}_2[\emptyset](\theta_0, z_1, (In, a_2)) > 0$. Bayes updating implies that $\hat{\mu}_2[In](\theta_0, z_1, (In, a_2)) = 1$, which makes b_1 sequentially best reply for player 2.

Given this, note that when player 1 reasons about players 2 behavior in the next round, player 1 is actually reasoning about players 2 conditional conjectures. If player 1 is certain that player 2 initial conjecture is such that $\hat{\mu}_2[\emptyset](\theta_0, z_1, (In, a_2)) > 0$, it follows from Bayes updating that player 1 is certain that player 2 chooses b_1 conditional on reaching history $[In]$, making (In, a_2) a sequential best reply. However, how could player 1 reason about player 2 conditional conjectures, when player 1 is certain that player 2 initial conjecture is $\hat{\mu}_2[\emptyset](\theta_0, z_1, a_1) = 1$? Since a deviation of player 1 to $[In]$ would contradict player 2 initial understanding of the game, the logic of *ISR* is completely silent about how player 2 would make inferences on those situations. Therefore, player 1 conditional beliefs about player 2 conditional conjectures need to consider every possible way of reasoning of player 2, even conditional conjectures for player 2 that are certain that player 1 would do something irrational, like choosing a_3. If player 1 is certain about this last scenario, then player 1 is certain that player 2 would choose b_2 conditional on reaching $[In]$, making a_1 a sequential best reply. The procedure stops at this stage with $ISR_1^\infty(z_1) = \{a_1, (In, a_2)\}$ and $ISR_2^\infty(z_2) = \{b_1, b_2\}$.

Once predictions for this game have been derived, it is possible to determine whether *ISR* satisfies strategic robustness at complete information types of Table 1. It will be defined that failure of strategic robustness requires that there is $\epsilon > 0$, a strategy profile and a perturbed type space, with the interpretation of representing a situation approximated to complete information, for which the strategy profile belongs to the *ISR*-set in the limit, but is ruled out as a prediction of ϵ-*ISR* for every perturbed type in the approximation.

Figure 1.2: Top figure, extensive-form representation with payoff at θ_0^*. Bottom figures, normal-form representations conditional on each history.

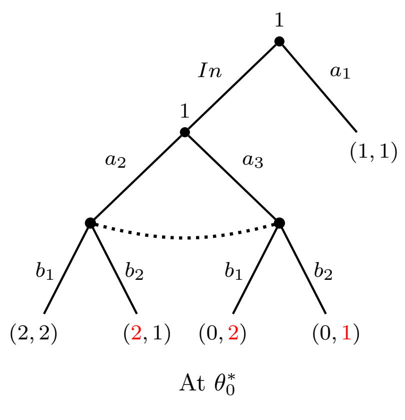

$[\emptyset]$ - $\{1\}$	b_1	b_2
a_1	(1,1)	(1,1)
(In, a_2)	(2,2)	(2,1)
(In, a_3)	(0,2)	(0,1)

$[In]$ - $\{1,2\}$	b_1	b_2
(In, a_2)	(2,2)	(2,1)
(In, a_3)	(0,2)	(0,1)

Table 1.2: A type space representing complete information over θ_0^*

$\bar{Z}_1 = \{\bar{z}_1\}$	$(\theta_0^*, \bar{z}_2)$	$\bar{Z}_2 = \{\bar{z}_2\}$	$(\theta_0^*, \bar{z}_1)$
$\beta_1[\bar{z}_1]$	1	$\beta_2[\bar{z}_2]$	1

Table 1.3: A type space of perturbed types, for $\rho \in (0,1]$

$Z_1^\rho = \{\bar{z}_1, (z_1^\rho)\}$	(θ_0, z_2^ρ)	$(\theta_0^*, \bar{z}_2)$	$Z_2^\rho = \{\bar{z}_2, (z_2^\rho)\}$	(θ_0, z_1^ρ)	$(\theta_0^*, \bar{z}_1)$
$\beta_1[z_1^\rho]$	$1-\rho$	ρ	$\beta_2[z_2^\rho]$	$1-\rho$	ρ

Figure 1.2 describes payoffs for each player at θ_0^*, while Table 1.2 describes a type space that induces a situation of common knowledge over the state θ_0^*. At the first stage of elimination, strategies a_1 and (In, a_3) are strictly dominated by (In, a_2) for player 1. Given this, and $\epsilon \in (0,1)$, strategy b_1 is the unique ϵ-sequentially best reply to conjectures that necessarily put probability one to (In, a_2). Hence, b_2 is eliminated at the second stage for player 2, and the procedure stops with $ISR_1^\infty(\bar{z}_1|\epsilon) = \{(In, a_2)\}$, and $ISR_2^\infty(\bar{z}_2|\epsilon) = \{b_1\}$.

Consider now the type space specified in Table 1.3. For each $\rho \in (0,1]$, each players' type puts probability ρ to the state θ_0^*, and $1-\rho$ to the state θ_0. Moreover, each players' type is, in addition, uncertain about opponents' type: puts probability ρ to θ_0^* and that his opponent has common knowledge that the state is θ_0^*; and with probability $1-\rho$ believes that the state is θ_0 and his opponent is uncertain about the state; and so on. When ρ is small, the type space is meant to represent a situation approximated to the complete information case as described by the type space of Table 1.1.

For an $\epsilon \in (0,1)$, at the first stage of elimination, strategy (In, a_3) cannot be made an ϵ-best reply to any conjecture of every type z_1^ρ of player 1. Given this, any initial conjecture for player 2 of type z_2^ρ needs to put at least probability $\rho > 0$ to player 1 choosing (In, a_2), and zero probability to (In, a_3). Hence, by Bayes updating every conditional conjecture at node $[In]$ for every type z_2^ρ of player 2 puts probability zero to (In, a_3), so that b_1 is the unique ϵ-sequential best reply of every type z_2^ρ. Given this, in the next round, every initial conjecture of every type z_1^ρ of player 1 needs to put probability 1 to player 2 choosing b_1, which makes (In, a_2) the unique ϵ-sequential best reply of every type z_1^ρ of player 1, and the process stops with $ISR_1^\infty(z_1^\rho|\epsilon) = \{(In, a_2)\}$ and $ISR_2^\infty(z_2^\rho|\epsilon) = \{b_1\}$.

Therefore, to conclude the argument, for an $\epsilon \in (0,1)$, strategy profile $s = (a_1, b_2)$, and the sequence of perturbed types z_i^ρ which converges to[3] z_i as $\rho \downarrow 0$ we have that $(a_1, b_2) \in ISR^\infty(z)$ and $(a_1, b_2) \notin ISR^\infty(z^\rho|\epsilon)$ for every $\rho > 0$, showing the lack of strategic robustness of *ISR*. As we mention in the introduction, lack of strategic robustness cast doubts on the interpretation and validity of the strategy profile (a_1, b_2) as a prediction for the situation with complete information of Table 1. In particular, note that the logic of *ISR* is such that, once (In, a_3) was eliminated, player 2 cannot infer that move a_1, at the beginning of the game, is not a rational thing to do. As long as player 1 believes that player 2 is certain that the game is going to stop after a_1, then whenever the observed evidence contradicts player 2 initial conjecture, the theory does not discipline what player 1 believes about player 2 conditional conjectures. When this is the case, even the initial hypothesis of player 2 that player 1 is sequentially rational has been contradicted, and therefore, any inference from player 1 about how player 2 is going to play is possible. Hence, if the set of (conditional) preference is rich enough, any behavior of player 2 can be justified by some inference of player 1. This is the reason for why player 2 cannot infer that move a_1, at the beginning of the game, is not rational; the event just described unfolds a sequence of events in which the theory does not discipline higher-orders of belief that in the limit allow for the profile (a_1, b_2) to be predicted by *ISR*.

Sequential rationality is a stronger form of optimal behavior than rationality, and is meant to reject some unreasonable outcomes whose optimality does not depend on the timing of the arrival of players' information. This is an important consideration for the results of the paper, that contrast with Ely and Peski's result for static games. In particular, it is not hard to see that *ICR* (which is defined similarly to *ISR*, but it is required only to be a best reply to an initial conjecture) has the same set of predictions as *ISR*, that is, $ICR_1^\infty(z_1) = \{a_1, (In, a_2)\}$ and $ICR_2^\infty(z_2) = \{b_1, b_2\}$. However, for this game, *ICR* does satisfy strategic robustness at complete information types of Table 1. In particular, for every $\epsilon > 0$, and every sequence of types for player 2, $(z_2^m)_m$, that converges to z_2 of Table 1, is such that $b_2 \in ICR_2^\infty(z_2^m|\epsilon)$ for m large enough. Hence, the conclusion follows. Therefore, when the timing of the arrival of information is an essential characteristic of the game, the use of type spaces to model players subjective beliefs might not be justified based on robustness considerations.

[3]Convergence of types is defined in terms of convergence of hierarchies of beliefs with respect to the product topology, as is formalized in section 2.1.3. However, note that for ρ small enough, first-order beliefs of type z_i^ρ puts probability closed to one to θ_0; puts probability closed to one that the state is θ_0 and his opponents type puts probability closed to one to the state θ_0, and so on; so that when $\rho \downarrow 0$ the hierachy of beliefs induced by types z_i^ρ converges to the hierarchy of belief induced by z_i.

1.3 Future research on the Robustness of Structural Rationalizability: The role of information and higher order beliefs.

As the previous example illustrated, the fundamental behavioral assumption in dynamic games, namely sequential rationality, requires an objective interpretation of a strategy (*what a player would actually do at every information set*) which is described independently of players beliefs. Moreover, it puts no restriction on the action and conditional beliefs that a player would take at information sets that (s)he does not expect to reach and that indeed are not reached in the observed play of the game. Therefore, interactive reasoning about sequential rationality carries over a collection of implicit assumptions that impose no testable restrictions over observable outcomes, neither for players in the game nor for an outside researcher. I will argue that it is this feature, implicit not only in *ISR* but most familiar solution concepts for dynamic games, the primary reasons for the result of critical types being generic. The following example illustrates these ideas.

Figure 1.3: The role of information and the weakness of sequential rationality

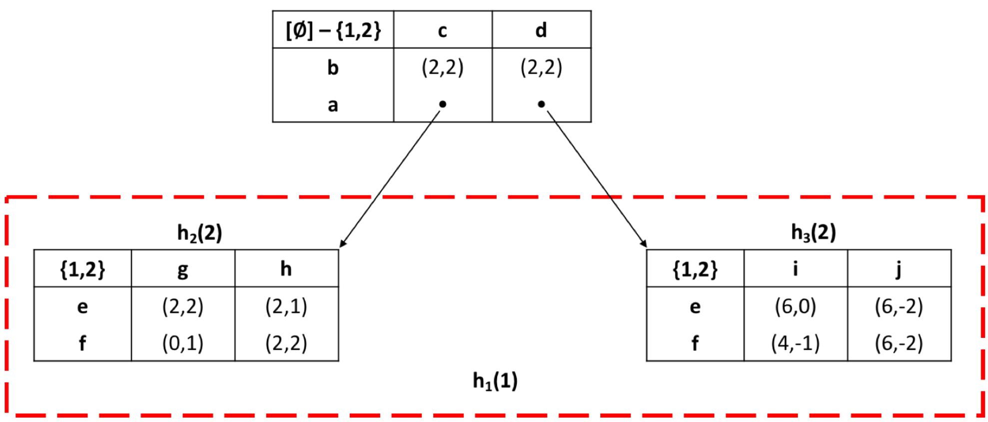

Strategy sets of each player:
$S_1 = \{(a,e),(a,f),(b,\cdot)\}$,
$S_2 = \{(c,g,i),(c,g,j),(c,h,i),(c,h,j),(d,g,i),(d,g,j),(d,h,i),(d,h,j)\}$,

Information sets of each player:
$\mathcal{H}_1 = \{[\emptyset],[h_1]\}$, with $[\emptyset] = [h_1] = S_2$
$\mathcal{H}_2 = \{[\emptyset],[h_2],[h_3]\}$, with $[\emptyset] = S_1$, $[h_2] = [h_3] = \{(a,e),(a,f)\}$

Example 2: Consider the following two-player game with imperfect information. At the beginning of the game, players choose simultaneously. If at the beginning of the game, player 1 chooses b, the game is over, and payoffs are realized. However, if player 1 chooses a at the beginning of the game, they move to the second stage of the game. Depending on the choice of player 2 between c and d at the beginning of the game, they face one of two possible simultaneously-move games in the second stage. At the second stage, player 2 perfectly knows which game is it, while player 1 does not know. A graphical representation of the situation is represented in figure 1.3. To facilitate the exposition, we have written at the bottom of the figure the strategy and informations sets of each player[4].

As we noted in the previous section, sequential rationality is identified by the first step in the iterative procedure of *ISR*. It requires strategies that survive the first step to be sequentially best responses to players conjectures. In particular, requires that for each information set that the strategy leads to, the player can get a conditional expected payoff at least as high as any other alternative strategy that could also lead to that information set. Consider, for example, any of the outcome-equivalent strategies of player 2, $(d, \cdot, i)$, and a conjecture given by $\mu = (\mu[\emptyset], \mu[h_2], \mu[h_3]) \in \Delta^{\mathcal{H}}(S_1)$, satisfying the standard properties of conditional probability systems[5]. Given that the only information sets that the strategy leads to are $\mathcal{H}((d, \cdot, i)) = \{[\emptyset], [h_3]\}$, then for $(d, \cdot, i)$ to be a sequential best response, such a conjecture must satisfy:

i_2^0) $\mu[\emptyset](b, \cdot) = 1$;

ii_2^0) No additional restrictions on $\mu[h_3]$, because $(d, \cdot, i)$ is conditionally strictly dominant; and,

iii_2^0) No additional restrictions on $\mu[h_2]$, because $[h_2] \notin \mathcal{H}_2((d, \cdot, i))$.

I would like to emphasize that there are important differences between condition ii_2^0) and iii_2^0) when it comes to understand how player 1 reasons about the event that player 2 is sequentially rational. In particular, given this and from the perspective of player 1, in the second step of *ISR* consider strategy (a, e) and a conjecture $\mu = (\mu[\emptyset], \mu[h_1]) \in \Delta^{\mathcal{H}}(S_2)$ satisfying $\mu[\emptyset]((d, \cdot, i)) = \mu[h_1]((d, \cdot, i)) = 1$, which makes it a sequential best response. Moreover, for such a conjecture, player 1 is certain at the beginning of the game, of the event:

$[t_2^1]$: *there is* $\mu \in \Delta^{\mathcal{H}}(S_1)$ *for p2 for which* $(d, \cdot, i) \in SR_2$, *satisfying* $i^0.2), ii^0.2) and, iii^0.2)$.

In other words, player 1 is certain at the beginning of the game of the following facts: there is a conjecture for player 2 that makes $(d, \cdot, i)$ a sequentially best response, and for which player 2 is certain at the beginning of the game that player 1 will end the game by choosing b; however, would player 2 be surprised by the choice of a, which happens to be player 1 actual choice, then player 1 is certain at the beginning of the game, that any conditional conjecture at $[h_2]$ for player

[4] In the case of player 1, the strategy $(b, \cdot)$ refers to the equivalence class of outcome-equivalent strategies (b, e) and (b, f)

[5] For more detail, see Chapter 2, Section 2.1.2

2 about what player 1 might do between e and f, would make $(d, \cdot, i)$ a conditional best reply. This is the content of ii_2^0). However, since the observed play does not lead to the information set $[h_2]$, then player 1 cannot deduce any reasonable fact about neither conditional conjectures at that information set for player 2 nor about how player 2 would act at such counterfactual event. This is the content of iii_2^0).

The weakness of sequential rationality as a behavioral assumption for dynamic games comes from situations like the one illustrated here. We emphasize that sequential rationality puts no restriction on the action that a player would take at information sets that he does not expect to reach, and that indeed are not reached in the observed play of the game; and it also puts no restriction on how players reason about opponents conditional beliefs at information sets that his opponents do not expect to reach, and that indeed are not reached in the observed play of the game. I consider this feature to be essential in understanding the robustness properties of predictions in dynamic games.

In what follows, I propose to study interactive reasoning of a stronger behavioral assumption for dynamic games, named *Structural Rationality*, introduced by Siniscalchi (2016a,b), for which strategies are evaluated at the beginning of the game, with the intended interpretation of a subjective plan describing what the player thinks (s)he would choose, at every information set, should that information set occur, and such evaluation of strategies takes into account players' conditional beliefs, in particular, the timing of arrival of new information.

Let me go back to the example to illustrate the properties of structural rationality. To introduce the concept, we need to define three important notions: a plausibility relation under conjectures μ on information sets, a basis for a conjecture μ; and, the structurally preferred relation on strategies.

We start with the plausibility relation under conjectures μ on information sets, denoted $\trianglerighteq_\mu$ and to be read as the "at least as plausible as" relation on information sets. As an example consider the following case: Fix $\mu \in \Delta^{\mathcal{H}}(S_1)$ for player 2 such that $\text{supp}(\mu[\emptyset]) = S_1$. Then, clearly, $\mu[h_2]$, and $\mu[h_3]$ are given by conditioning, and,

1. since $\mu[h_2]([\emptyset]) > 0$, we say $p2$ considers $[\emptyset] \trianglerighteq_\mu [h_2]$; and,

2. since $\mu[\emptyset]([h_2]) > 0$, we say that $p2$ considers $[h_2] \trianglerighteq_\mu [\emptyset]$;

Hence, we consider both information sets to be equally plausible, that is $[\emptyset] \backsim_\mu [h_2]$. Similarly we have $[\emptyset] \backsim_\mu [h_3]$. Now in general consider the following definition:

Definition 1. *Fix a conjecture μ on S_{-i}, indexed by $\mathcal{H}_i$. For every $[h], [h'] \in \mathcal{H}_i$, we say that: "$[h]$ is at least as plausible as $[h']$ under μ", denoted $[h] \trianglerighteq_\mu [h']$, if there are $[h_1], ..., [h_M]$ in $\mathcal{H}_i$ with $[h_1] = [h']$ and $[h_M] = [h]$ and for all $m = 1, ..., M$,*

$\mu[h_m]([h_{m+1}]) > 0$.

The relation $\triangleright_\mu$ on $\mathcal{H}_i$ is a non-complete, reflexive and transitive preorder. It follows from the definition that we should also be able to order the information sets $[h_2]$ and $[h_3]$ as follows: since there is $[\emptyset] \in \mathcal{H}_2$ such that $\mu[h_2]([\emptyset]) > 0$, and $\mu[\emptyset]([h_3]) > 0$, we have $[h_3] \triangleright_\mu [h_2]$; and, since there is $[\emptyset] \in \mathcal{H}_2$ such that $\mu[h_3]([\emptyset]) > 0$, and $\mu[\emptyset]([h_2]) > 0$, we have $[h_2] \triangleright_\mu [h_3]$. Therefore, $[\emptyset] \backsim_\mu [h_2] \backsim_\mu [h_3]$. Note that in our example, as long as $\mu[\emptyset](a, \cdot) > 0$, we have $[\emptyset] \backsim_\mu [h_2] \backsim_\mu [h_3]$.

Consider now a second case as follows: Fix $\mu \in \Delta^{\mathcal{H}}(S_1)$ for player 2 such that $\mu[\emptyset](b, \cdot) = 1$. Then, $\mu[h_2]$, and $\mu[h_3]$ are not given by conditioning. Then it can be easily check that we have,

1. $[\emptyset] \triangleright_\mu [h_2]$, $[\emptyset] \triangleright_\mu [h_3]$, and $[h_2] \backsim_\mu [h_3]$;

2. $([h_2], [\emptyset]) \notin \triangleright_\mu$, and $([h_3], [\emptyset]) \notin \triangleright_\mu$.

As it is illustrated, for a given conjecture μ, the plausibility relation on information sets refines in a nontrivial way the natural order of the extensive-form given by the precedence relation on information sets.

The second notion that we need to introduce is that of a basis for a given conjecture. For each conjecture μ we consider a collection $p^\mu = (p_h)_{h \in \mathcal{H}}$ where each p_h is a probability measure over opponent strategy set, with the intended interpretation of being alternative beliefs, at the beginning of the game, about how the game is going to unfold, and that generate μ by conditioning. More formally,

Definition 2. *Fix a conjecture μ on $(S_{-i}, \mathcal{H}_i)$. A basis for μ is a collection $\boldsymbol{p}^\mu = (p_h)_{h \in \mathcal{H}} \subset \Delta(S_{-i})$ that satisfies the following conditions:*

1. *for every $[h], [h']$ in $\mathcal{H}_i$, $p_h = p_{h'}$ iff both $[h] \triangleright_\mu [h']$ and $[h'] \triangleright_\mu [h]$.*

2. *for every $[h]$ in $\mathcal{H}_i$, $p_h(\cup\{[h'] \in \mathcal{A}_i : [h] \triangleright_\mu [h'], [h'] \triangleright_\mu [h]\}) = 1$.*

3. *for every $[h]$ in $\mathcal{H}_i$, $p_h([h]) > 0$, and for every (measurable) subset $B \subset S_{-i}, \mu[h](B \cap [h]) = \frac{p_h(B \cap [h])}{p_h([h])}$.*

As an example, for the conjecture of player 2 given in the first case before, note that a basis for such a conjecture exists and is given by $(\mu[\emptyset], p_{h_2}, p_{h_3})$ such that $\mu[\emptyset] = p_{h_2} = p_{h_3}$. However, for the second case, we have that a basis for such a conjecture exists as long as both $\mu[h_2]$ and $\mu[h_3]$ have full support, and it is given by $(\mu[\emptyset], p_{h_2}, p_{h_3})$ such that $p_{h_2} = p_{h_3}$, but $\mu[\emptyset] \neq p_{h_2}$, and $\mu[\emptyset] \neq p_{h_3}$. It was shown in Siniscalchi (2016a,b) that not every conjecture admits a basis. For example, let a conjecture for player 2 be given by $\mu[\emptyset](b, \cdot) = \mu[h_2](a, e) = \mu[h_3](a, f) = 1$. As in the second case before, we note that it must be the case that $p_{h_2} = p_{h_3}$, but $\mu[\emptyset] \neq p_{h_2}$, and $\mu[\emptyset] \neq p_{h_3}$ which comes from the plausibility order. However, an inconsistency is generated since it is impossible that a unique probability measure would generate such a conditioning. We emphasize that we only allow players to consider conjectures that admits a basis, and this is common knowledge.

Finally, we arrived at the most important concept in this section, the notion of structural preferences. Formally,

Definition 3. *Fix a conjecture μ that admits a basis $\boldsymbol{p}^\mu$. For any $s_i, s_i' \in S_i$, we say that s_i is (weakly) structurally preferred to s_i' given u_i and $\boldsymbol{p}^\mu$, denoted $s_i \succeq^{u_i,\mu} s_i'$, iff for every $[h] \in \mathcal{H}_i$ such that*

$$\mathbb{E}_{p_h}[U_i(s_i, \cdot)] < \mathbb{E}_{p_h}[U_i(s_i', \cdot)]$$

there is $[h'] \in \mathcal{H}_i$ such that $[h'] \rhd_\mu [h]$ and

$$\mathbb{E}_{p_{h'}}[U_i(s_i, \cdot)] > \mathbb{E}_{p_{h'}}[U_i(s_i', \cdot)]$$

We say that s_i is strictly structurally preferred to s_i' given u_i and $\boldsymbol{p}^\mu$, denoted $s_i \succ^{u_i,\mu} s_i'$, iff $s_i \succeq^{u_i,\mu} s_i'$, and not $s_i' \succeq^{u_i,\mu} s_i$.

It follows that a strategy $s_i \in S_i$ is structurally rational for (u_i, μ), if there is no strategy $s_i' \in S_i$ such that $s_i' \succ^{u_i,\mu} s_i$. It is not hard to see that a structural preference is a reflexive and transitive partial order on strategies that, in general, is not complete; that it clearly reduces to expected utility when the game is static, or when every information set receives positive probability, but it only differs from expected utility if at least one observable event is not expected to be reached. Finally, Theorem 1 in Siniscalchi (2016a) shows that if a strategy is structurally rational, then it is sequentially rational. Is in this sense that we consider structural rationality to be a stronger behavioral assumption than sequential rationality.

What we emphasize in this section is that structural rationality precludes conjectures with arbitrary conditional beliefs at information sets that are not expected to be reached, and in fact are not reached in the observed play. This is essential to understand the problem of robustness. We close this section by illustrating how the concept is applied to our example in figure 1.3. We showed that the outcome-equivalent strategy $(d, \cdot, i)$ is sequentially rational for player 2. However, in what follows we show that neither (d, g, i) nor (d, h, i) are structurally rational. In other words, player 2 has a strategy that is sequentially rational, but is not structurally rational. We only need to consider two relevant cases. First, consider a conjecture as in the first case given before, that is, $\mu[\emptyset](a, \cdot) > 0$, and note the following:

1. Since $[\emptyset] \backsim_\mu [h_2] \backsim_\mu [h_3]$, a basis for μ is sth. $\mu[\emptyset] = p_{h_2} = p_{h_3}$.

2. $\mathbb{E}_{\mu[\emptyset]}\left[U_2(s_2', \cdot)\right] \in [1, 2]$, for $s_2' \in \{(c, g, i), (c, g, j), (c, h, i), (c, h, j)\}$.

3. Let $s_2 \in \{(d, g, i), (d, h, i)\}$, and $s_2' \in \{(c, g, i), (c, g, j), (c, h, i), (c, h, j)\}$.

   $\mathbb{E}_{\mu[\emptyset]}\left[U_2(s_2, \cdot)\right] \geq \mathbb{E}_{\mu[\emptyset]}\left[U_2(s_2', \cdot)\right]$ iff $\mu[\emptyset](b, \cdot) = 1$, which contradicts the requirement $\mu[\emptyset](a, \cdot) > 0$.

Second, consider a conjecture as in the second case given before that admits a basis. We have:

1. Since $[\emptyset] \rhd_\mu [h_2]$, $[\emptyset] \rhd_\mu [h_3]$, and $[h_2] \backsim_\mu [h_3]$; but, $([h_2], [\emptyset]) \notin \rhd_\mu$, and $([h_3], [\emptyset]) \notin \rhd_\mu$, a basis for μ is sth. $\mu[\emptyset] \neq p_{h_2} = p_{h_3}$

2. $\mathbb{E}_{\mu[\emptyset]}\left[U_2(s_2, \cdot)\right] = 2$, for all $s_2 \in S_2$.

3. $\mathbb{E}_{p_h}\left[U_2(s_2', \cdot)\right] \in [1, 2]$, for $s_2' \in \{(c, g, i), (c, g, j), (c, h, i), (c, h, j)\}$.

4. $\mathbb{E}_{p_h}\left[U_2(s_2, \cdot)\right] < 0$, for $s_2 \in \{(d, g, i), (d, h, i)\}$.

which completes the proof of the claim, that is, we have that neither (d, g, i) nor (d, h, i) are in the set of structurally rational strategies for player 2, which is given by $StR_2^0 = \{(c, g, i), (c, g, j), (c, h, i), (c, h, j)\}$.

Finally, we are interested in a formal treatment of interactive reasoning of structural rationality, that is, we ask: What are the behavioral implications of the following line of reasoning:

StR_1^0: $p1$ is StR.	StR_2^0: $p2$ is StR.
StR_1^1: StR_1^0 and $p1$ is certain (at $[\emptyset]$) of StR_2^0.	StR_2^1: StR_2^0 and $p2$ is certain (at $[\emptyset]$) of StR_1^0.
...	...
StR_1^k: StR_1^{k-1} and $p1$ is certain (at $[\emptyset]$) of StR_2^{k-1}.	StR_2^k: StR_2^{k-1} and $p2$ is certain (at $[\emptyset]$) of StR_1^{k-1}.
and so on..	and so on..

The formal treatment considers dynamic games with payoff uncertainty and, as in Siniscalchi (2016a,b), assume that players in the game choose strategies according to structural rationality. We study the general properties (e.g. non-emptiness, measurability, etc.) of (Weak) interim structural rationalizability (*IStR*), a solution concept that characterizes the behavioral implications of common certainty in structural rationality, as presented above.

In the case of Bayesian dynamic games with incomplete information, we are particularly interested in the robustness properties of *IStR* to perturbations of higher-order beliefs. We conjecture three results: a structure theorem of structural rationalizability, a characterization of critical types, and a non-generic result of the set of critical types.

Chapter 2

Critical Types in Dynamic Games

2.1 Framework

I consider finite dynamic games with payoff uncertainty and with observable actions. The basic structure of the game allow us to describe, for each player, a space of external states together with a collection of observable events. An external state is a description of a payoff-relevant state and a description of how players' would act at each observable event. Conditional probability systems would be used in order to describe not only players beliefs about the external states, but also how players update their belief as new information is revealed.

2.1.1 Dynamic games with Payoff Uncertainty

The basic structure is described as follows. Let $M < \infty$ and $\Gamma = \{N_0, \Theta, (A_i)_{i \in N}, (\Gamma_M, \geq)\}$, where $N_0 = \{0\} \cup N$ is the set of players (including chance), and for each player $i \in N_0$, A_i is a (finite) set of actions[1]. A history $h \in A^m$, for some $1 \leq m < \infty$, is a finite concatenation of action profiles of length m. Let $H = \cup_{m \in \mathbb{N}} A^m$ be the set of all histories. A (finite) game tree of length M is an ordered set $(\Gamma_M, \geq)$, where Γ_M is a (finite) subset of the set of all histories, such that Γ_M is closed under initial segments[2], and the length of each history is at most M. The order of the nodes in Γ_M is given by $(a^\tau)_{\tau=1}^t \geq (a'^\tau)_{\tau=1}^k$ if $t \leq k$ and $a^\tau = a'^\tau$ for all $\tau = 1, .., t$, and by $\emptyset \geq h$ for all $h \in \Gamma_M$, that is, $\emptyset$ is the root of the game tree $(\Gamma_M, \geq)$. The game tree Γ_M can be partitioned into the set of terminal histories $\mathcal{Z}$, that is, histories that are not the initial segment of any other history in Γ_M, and the set of partial histories $\mathcal{H}$ (by definition, $\emptyset \in \mathcal{H}$)[3].

[1] The standard notation about product spaces applies. That is, let $A = \bigtimes_{i \in N_0} A_i$, let $A_{-i} = \bigtimes_{j \neq i} A_j$, and let $A^m = \bigtimes_1^m A$.

[2] That is, if $h \in \Gamma_M$ and h' is an initial segment of h (i.e for $h = (a^1, ..., a^s, ..., a^m)$, and $h|_s = (a^1, ..., a^s)$, we have $h' = h|_s$), then $h' \in \Gamma_M$.

[3] Note that, since the game is finite each terminal history can be identified uniquely with a path in Γ_M. A path is a maximal chain in the ordered structure $(\Gamma_M, \geq)$, a subset of histories that is totally ordered, and is maximal for this property.

We make explicit one additional property on $(\Gamma_M, \geq)$:

Assumption 1. *For all $h, h' \in \Gamma_M$, if $h > h'$ then there exists $h'' \in \Gamma_M$ such that $h > h''$ and neither $h' \geq h''$ nor $h'' \geq h'$.*

(A1) eliminates irrelevant nodes in the game tree, that is, nodes in which there is no decision to make. At every stage of the game, the history h that has just occurred becomes public information and is perfectly recalled by all the players. For each $h \in \mathcal{H}$ and $i \in N$, let $A_i(h)$ denote the nonempty and finite set of actions available to player i at history h[4].

Let Θ denote the set of payoff-relevant states. Let $\{\Theta_0, (\Theta_i, u_i)_{i \in N}\}$ be the payoff-information structure associated to Θ; it is assumed to be common knowledge. Players payoffs are represented by functions $u_i : \Theta \times \mathcal{Z} \to \mathbb{R}$ for each $i \in N$. We assume that the payoff-information structure is represented by partitions of Θ with a product structure, that is, $\Theta = \bigtimes_{i \in N_0} \Theta_i$, where for each $i \in N$, Θ_i is the set of player's i payoff types. Let Θ_0 denote the set of states of nature. It represents the residual uncertainty after pooling everybody else information.

At the payoff state $(\theta_0, \theta_1, ..., \theta_n) \in \Theta$, player i's payoff type is θ_i which is informed only to and observed only by player i at the beginning of the game. Therefore, payoff types represents players' private information about the payoff-relevant states: Player i knows θ_i only if the realized state is in $\Theta_0 \times \{\theta_i\} \times \Theta_{-i}$, that is, if θ_i is true.

Given a (finite) game tree $(\Gamma_M, \geq)$ and the set Θ of payoff-relevant states, we consider the collection of game trees $(\Theta \times \Gamma_M, \geq)$, given by $\Theta \times \Gamma_M = \{(\theta, h);\ \theta \in \Theta, h \in \Gamma_M\}$, and an order such that for any two $(\theta, h), (\theta', h') \in \Theta \times \Gamma_M$, we define $(\theta, h) \geq (\theta', h')$ if and only if $\theta = \theta'$ and $h \geq h'$[5].

A strategy (or plan of action) for player i is a mapping s_i that assigns to each partial history $h \in \mathcal{H}$ an available choice $s_i(h) \in A_i(h)$, unless h cannot be reached due to some choice $s_i(h')$ at an earlier history $h' \in \mathcal{H}$. Let S_i denote the set of strategies of player i. Since the game is finite, each profile of strategies $s \in S$ induces a unique terminal history $\rho(s) \in \mathcal{Z}$.

We introduce now three mappings relating nodes and strategies. First, for each $h \in \mathcal{H}$, let $S(h)$ be the set of profiles $s \in S$ that allow h to be reached[6]; its projection on S_i and S_{-i} are denoted by $S_i(h)$ and $S_{-i}(h)$. Next, let $\mathcal{H}(s_i) = \{h \in \mathcal{H};\ s_i \in S_i(h)\}$ denote the set of partial histories not discarded by s_i.

Finally, fix a history $h \in \mathcal{H}$, and a strategy profile $s \in S$. Define the set of discarded nodes at h by strategy profile s, $D_h(s)$, as the set of nodes $h' \in \Gamma_M$ for which h is a predecessor, that is $h > h'$, and for which there are $\hat{h} \in \mathcal{H}$ with $h \geq \hat{h} > h'$, $a'_i \in A_i(\hat{h})$ for every $i \in N$, such that

[4]Formally, for every $h \in \Gamma_M$, $A_i(h) = \text{proj}_{A_i}\{a \in A; (h, a) \in \Gamma_M\}$. Whenever $A_i(h)$ contains just one element we will say that agent i is inactive.

[5]We abuse notation on the definition of the partial order on $\Theta \times \Gamma_M$, but no confusion should arise.

[6]Formally, for $h \in \mathcal{H}$ let $S(h) = \{s \in S;\ h = \rho(s)|_m \text{ for some } 0 \leq m \leq M\}$.

$s_i(\hat{h}) \neq a'_i$ for some $i \in N$, where $(\hat{h}, (a'_i)_{i\in N})$ is an initial segment of h'. Let $U_h(s)$ denote the set of undiscarded nodes at h, as the set of nodes $h' \in \Gamma_M$ with $h \geq h'$ such that $h' \notin D_h(s)$. Of special interest are the sets of discarded nodes, $D_\emptyset(s) \subset \Gamma_M \setminus \{\emptyset\}$, and the set of undiscarded nodes $U_\emptyset(s) = \Gamma_M \setminus D_\emptyset(s)$.

We denote by $\Sigma_i := \Theta_i \times S_i$ the space of strategy-payoff type pairs for player i. Define the space of external states to be $\Sigma = \bigtimes_{i\in N_0} \Sigma_i$, with $\Sigma_0 = \Theta_0$. However, since at the beginning of the game each player $i \in N$ is informed about his payoff-type $\theta_i \in \Theta_i$, from the point of view of player i, who has to decide which strategy $s_i \in S_i$ to carry out, he faces $X_i = \Sigma_{-i}$ as a basic space of uncertainty. In addition, for each (partial) history $h \in \mathcal{H}$, let $\Sigma_i(h) := \Theta_i \times S_i(h)$ be the set of strategy-payoff type pairs consistent with history h. The interpretation is as follows: as the game unfolds, player i realizes that at history h his opponents can only be playing a strategy combination s_{-i} that leads to h, for some payoff type profile, i.e. that the "true"external state must be an element in $\Sigma_{-i}(h)$. Therefore, from the point of view of player i, his collection of observable events is given by $\mathcal{A}_i = \{\Sigma_{-i}(h);\, h \in \mathcal{H}\}$.

We can define a strategic-form representation of the game Γ_M: For each player i, let $U_i : \Sigma_i \times \Sigma_{-i} \to \mathbb{R}$ denote the strategic-form payoff function of player i defined, for each $\sigma_i = (\theta_i, s_i) \in \Sigma_i$, and $\sigma_{-i} = (\theta_{-i}, s_{-i}) \in \Sigma_{-i}$, by $U_i(\sigma_i, \sigma_{-i}) = u_i(\theta, \rho(s))$.

2.1.2 Conditional Probability Systems

In order to analyze players' reasoning as the game unfolds, we consider *conditional belief vectors*, a notion of endogenous probabilistic beliefs represented by a system of conditional probabilities. Given a player $i \in N$, and a basic space of uncertainty X_i, let $\mathcal{A}_i$ denote a collection of observable events.

Definition 4. *A conditional probability system for player i on $(X_i, \mathcal{A}_i)$ is a collection $\mu_i : \mathcal{A}_i \to \Delta(X_i)$ satisfying the following conditions*[7]*:*

1. *For all $A \in \mathcal{A}_i$, $\mu_i[A] \in \Delta(X_i)$ such that $\mu_i[A](A) = 1$.*
2. *For every measurable set $B \subset X_i$, and $A, A' \in \mathcal{A}_i$ such that $B \subset A \subset A'$*

$$\mu_i[A](B) \cdot \mu_i[A'](A) = \mu_i[A'](B)$$

The set of conditional belief vectors on $(X_i, \mathcal{A}_i)$ is denoted by $\Delta^{\mathcal{A}_i}(X_i)$, a closed subspace

[7]For an arbitrary set X, denote by $\Delta(X)$ the set of probability measures on X; whenever X is a topological space, probability measures are defined over the Borel subsets of X, that is, on $\mathcal{B}(X)$, and we endow $\Delta(X)$ with the weak*-topology. In addition, whenever the space X is Polish, the Borel measurable subsets of $\Delta(X)$ are generated by the maps $\nu \to \nu(A)$, $A \in \mathcal{B}(X)$.

of the topological set $[\Delta(X_i)]^{\mathcal{A}_i}$ endowed with the product topology, which makes $\Delta^{\mathcal{A}_i}(X_i)$ a compact metrizable space.

In the context of a game as defined above we can associate each observable event of player i with a history $h \in \mathcal{H}$. We shall write $\Delta^{\mathcal{H}}(\Sigma_{-i})$ (or as in the next section $\Delta^{\mathcal{H}}(\Omega_{-i})$) for $\Delta^{\mathcal{A}_i}(X_i)$[8]. Recall that at every stage the history h becomes public information. Condition (1) states that players are always certain about what they know as the game unfolds. Condition (2) says that players conditional beliefs are consistent with Bayes updating whenever possible. Therefore, $\Delta^{\mathcal{H}}(\Sigma_{-i})$ describe the set of player i's *first-order conditional beliefs* about strategy-payoff type pairs of his opponents in the game tree $(\Gamma_M, \geq)$ with payoff-information structure $\{\Theta_0, (\Theta_i, u_i)_{i\in N}\}$.

2.1.3 Bayesian Games in Extensive Form

In order to complete the formal description of a game an analyst append to the game tree $(\Gamma_M, \geq)$ and the information structure $\{\Theta_0, (\Theta_i, u_i)_{i\in N}\}$, a model $\mathbf{Z}$ describing, at the beginning of the game, players' subjective beliefs about payoff-relevant states (so-called *first-order beliefs*), beliefs about payoff-relevant states and opponents' beliefs about payoff-relevant states (that is, *second-order beliefs*), and so on. This triplet defines what is usually called a *Bayesian game in extensive form*, a basic structure behind all classical solution concepts for games of incomplete information in extensive-form as Sequential Equilibrium, Perfect Bayesian Equilibrium, among other refinements, and it is known as the standard approach to analyze game-theoretic models[9]. Formally,

Definition 5. *Given $\{\Theta_0, (\Theta_i, u_i)_{i\in N}\}$, a Θ-based belief space (or a model) is a tuple $\mathbf{Z} = \{\Theta_0, (Z_i, \hat{\theta}_i, \beta_i)_{i\in N}\}$, such that for each player $i \in N$, Z_i is a compact (metrizable) set of types*[10]*, $\hat{\theta}_i : Z_i \to \Theta_i$ is a surjective continuous function, and $\beta_i : Z_i \to \Delta(\Theta_0 \times Z_{-i})$ is a continuous function.*

For each player i, each type in Z_i induces a hierarchy of belief about payoff-types on the universal type space[11] Z_i^* as follows: first-order beliefs of type $z_i \in Z_i$ over $\Theta_0 \times \Theta_{-i}$ are generated by the map $\hat{\beta}_i^1 : Z_i \to \Delta(\Theta_0 \times \Theta_{-i})$ defined, for every measurable set $Y \subset \Theta_0 \times \Theta_{-i}$,

[8] We identify, for each player i, history h with the event $\Sigma_{-i}(h)$.

[9] In many applications, the choice of a model varies from simple and tractable ones, e.g. models from a common-prior, or arbitrary finite models; to the study of very complex models, as for example, the universal belief space. In any case, models are considered as data in the analysis, and as such, they are approximations to some idealized strategic situation.

[10] Note the difference, we call $\theta_i \in \Theta_i$ a payoff-type for player i, an element $\psi_i \in \Psi_i$ is called an epistemic type, and an element $z_i \in Z_i$ is called a type.

[11] For a formal presentation see the Appendix. See also, Mertenz and Zamir (1985), or Brandenburger and Dekel (1993).

by

$$\hat{\beta}_i^1[z_i](Y) = \beta_i[z_i]\left(\left\{(\theta_0, z_{-i}) \in \Theta_0 \times Z_{-i};\ \left(\theta_0, \hat{\theta}_{-i}(z_{-i})\right) \in Y\right\}\right)$$

For $k > 1$, let Z_i^{k-1} be the set of coherent belief hierarchies of order $k-1$. Hence, the k-order belief can be induced by a map $\hat{\beta}_i^k : Z_i \to \Delta(\Theta_0 \times \Theta_{-i} \times Z_{-i}^{k-1})$ as follows: for every measurable set $Y \subset \Theta_0 \times \Theta_{-i} \times Z_{-i}^{k-1}$

$$\hat{\beta}_i^k[z_i](Y) = \beta_i[z_i]\left(\left\{(\theta_0, z_{-i}) \in \Theta_0 \times Z_{-i};\ \left(\theta_0, \hat{\theta}_{-i}(z_{-i}), \hat{\beta}_{-i}^{k-1}(z_{-i})\right) \in Y\right\}\right)$$

Hence, the mapping $\hat{\beta}_i : Z_i \to Z_i^*$ defined by $\hat{\beta}_i[z_i] = \left(\hat{\beta}_i^1[z_i], \hat{\beta}_i^2[z_i], ...\right)$ associates to each type z_i on the Θ-based belief space $\mathbf{Z}$ a corresponding hierarchy of beliefs about payoff relevant states. Let $\hat{\beta}_i^* : Z_i \to \Theta_i \times Z_i^*$ be defined by $\hat{\beta}_i^*[z_i] = \left(\hat{\theta}_i(z_i), \hat{\beta}_i[z_i]\right)$, and consider the profile $\hat{\beta}^* = \left(\hat{\beta}_i^*\right)_{i\in N}$ such that for each profile of types $z \in Z$ from $\mathbf{Z}$, it assigns a profile of payoff-types and of hierarchy of beliefs about payoff-states, that is, $\hat{\beta}^*[z] = \left(\hat{\beta}_i^*[z_i]\right)_{i\in N} \in \Theta \times Z^*$.

Mertens and Zamir (1986) showed that in fact there exists a natural homeomorphism $\eta_i^* : Z_i^* \to \Delta(\Theta_0 \times \Theta_{-i} \times Z_{-i}^*)$ that preserves beliefs of all orders: for each $z_i = (\nu_1, \nu_2, ...) \in Z_i^*$,

$$\underset{\Theta_0 \times \Theta_{-i} \times Z_i^{k-1}}{\mathrm{marg}} \eta_i^*(z_i) = \nu_k \tag{2.1}$$

That η_i^* is one-to-one means that different types have different belief over $\Theta_0 \times \Theta_{-i}$ and others players types. That η_i^* is onto means that any belief about $\Theta_0 \times \Theta_{-i}$ and others players types is held by some type of player i.

The tuple $\mathcal{T}^* = \{\Theta, (T_i^*, \theta_i^*, \beta_i^*)_{i\in N}\}$, where $T_i^* := \Theta_i \times Z_i^*$, and the functions $\theta_i^* : T_i^* \to \Theta_i$ and $\beta_i^* : T_i^* \to \Delta(\Theta_0 \times T_{-i}^*)$ are given by $\theta_i^*(\theta_i, z_i) = \theta_i$ and $\beta_i^*(\theta_i, z_i) = \eta_i^*(z_i)$, is a canonical representation of the universal belief space. A belief space $\mathbf{Z}$ is said not to have redundant types if $\forall\ z_i, z_i' \in Z_i,\ z_i \neq z_i'$ implies that $\hat{\beta}_i^*[z_i] \neq \hat{\beta}_i^*[z_i']$. For any non-redundant belief space, the set $\hat{\beta}^*(Z)$ is said to be a *belief-closed* subset of Z^*, if for every $\hat{\beta}_i(z_i) \in \hat{\beta}_i(Z_i)$ we have that $\eta_i^*(\hat{\beta}_i(z_i))[\Theta_0 \times \hat{\beta}_{-i}^*(Z_{-i})] = 1$. A type $z_i \in Z_i^*$ is finite if it belongs to a finite belief-closed subspace of Z_i^*. Let $\hat{Z}$ denote the set of finite types.

A given model $\mathbf{Z}$ is considered as exogenous assumptions describing a player's subjective view, at the beginning of the game, about the strategic situation. As remark by Penta (2012), *It is given such (exogenous) beliefs that we can apply game theoretic reasoning to make predictions about players behavior (the endogenous variables).*

2.1.4 Solution Concepts

As we emphasized earlier in this paper, we are concern about the robustness properties of predictions of Bayesian games in extensive form, in the sense that reasonable behavioral predictions do not change too much to small perturbations of beliefs and higher-order beliefs about payoff-relevant states. Reasonable behavioral predictions refer to strategies consistent with rationality and common initial belief in rationality (*RCIBR*), characterized via *Interim Sequential Rationalizability (*ISR*)* (Penta, 2012), a solution concept for Bayesian games in extensive form, and robustness refers to a notion of strategic continuity of *ISR*.

In this section, we formally define two solution concepts: Extensive-Form Best Reply Sets (*EFBRS*, Battigalli and Friedenberg 2012), which identifies *some* strategy-payoff type pairs consistent with *RCIBR*; and *ISR*, which identifies *all* strategy-payoff type pairs consistent with *RCIBR* in a belief-complete epistemic type space. In section 2.2 we consider the question about robustness, while in section 2.3 we justify the consideration of *EFBRS* along with *ISR*.

Interim Sequential Rationalizability (ISR)

Given a game tree $(\Gamma_M, \geq)$, a payoff-information structure $\{\Theta_0, (\Theta_i, u_i)_{i\in N}\}$, and a model $\mathbf{Z}$, let $\hat{u}_i : \Theta_0 \times Z \times \mathcal{Z} \to \mathbb{R}$ be such that, for each $(\theta_0, z, \bar{h}) \in \Theta_0 \times Z \times \mathcal{Z}$, $\hat{u}_i(\theta_0, z, \bar{h}) = u_i(\theta_0, \hat{\theta}(z), \bar{h})$. In the same way, we can extend the domain of the strategic-form payoff function U_i to include the payoff irrelevant higher-order belief. We will continue using U_i for the later case.

Given $\epsilon \geq 0$, we assume that each player $i \in N$, of type $z_i \in Z_i$ with $\hat{\theta}_i(z_i) = \theta_i^{z_i} \in \Theta_i$, chooses a strategy $s_i \in S_i$ that is ϵ-optimal at every history $h \in \mathcal{H}$ that s_i leads to, for a given conjecture $\hat{\mu}_i \in \Delta^{\mathcal{H}}(\Theta_0 \times Z_{-i} \times S_{-i})$. Formally,

Definition 6. *Let $\hat{\mu}_i \in \Delta^{\mathcal{H}}(\Theta_0 \times Z_{-i} \times S_{-i})$, and $(z_i, s_i) \in Z_i \times S_i$. For any $\epsilon \geq 0$, a strategy s_i is a sequential ϵ-best reply to $\hat{\mu}_i$ for z_i if and only if, for every $h \in \mathcal{H}(s_i)$, and every $s_i' \in S_i(h)$,*

$$\int_{\Theta_0 \times Z_{-i} \times S_{-i}} U_i(z_i, s_i, \theta_0, z_{-i}, s_{-i}) d\hat{\mu}_i[h] \geq \mathbb{E}_{\hat{\mu}_i[h]} \left[U_i(z_i, s_i', \theta_0, z_{-i} s_{-i}) \right] - \epsilon \tag{2.2}$$

For any $\hat{\mu}_i \in \Delta^{\mathcal{H}}(\Theta_0 \times Z_{-i} \times S_{-i})$, let $br_{i,z_i}(\hat{\mu}_i|\epsilon)$ denote the set of sequential ϵ-best replies to $\hat{\mu}_i$ for z_i, and let $br_i(\hat{\mu}_i|\epsilon) = \{(z_i, s_i) \in Z_i \times S_i;\ s_i \in br_{i,z_i}(\hat{\mu}_i|\epsilon)\}$.

A strategy $s_i \in S_i$ is ϵ-sequentially rational for type z_i, written $s_i \in r_i(z_i|\epsilon)$, if there exists a $\hat{\mu}_i \in \Delta^{\mathcal{H}}(\Theta_0 \times Z_{-i} \times S_{-i})$ such that $s_i \in br_{i,z_i}(\hat{\mu}_i|\epsilon)$.

It is well known that a strategy is sequentially rational, with $\epsilon = 0$, for a type z_i if and only if there is no history that the strategy leads to at which it is strictly dominated in that history. For each player $i \in N$ and each type $z_i \in Z_i$, define the set of *consistent conditional belief vectors* for

z_i as those that agree, at the beginning of the game, with type z_i (subjective) beliefs. Formally,

$$\varPhi_i(z_i) = \{\hat{\mu} \in \Delta^{\mathcal{H}}(\Theta_0 \times Z_{-i} \times S_{-i});\ \operatorname*{marg}_{\Theta_0 \times Z_{-i}} \hat{\mu}[\emptyset] = \beta_i[z_i]\} \tag{2.3}$$

For each pair $(z_i, s_i) \in Z_i \times S_i$, and (measurable) $R_{-i} \subseteq Z_{-i} \times S_{-i}$, let

$$\mathcal{B}_i^{\mathcal{H}}\left[(z_i, s_i); R_{-i}, \epsilon\right] = \{\hat{\mu}_i \in \varPhi(z_i);\ s_i \in br_{i,z_i}(\hat{\mu}_i|\epsilon) \text{ and } \operatorname{supp}(\hat{\mu}[\emptyset]) \subseteq \Theta_0 \times R_{-i}\} \tag{2.4}$$

Of particular interest are those strategies that are ϵ-sequentially rational for type $z_i \in Z_i$ for a consistent conditional belief vector that, at the beginning of the game, are concentrated on some measurable subset $R_{-i} \subseteq Z_{-i} \times S_{-i}$. That is, we define $s_i \in r_i(z_i|R_{-i}, \epsilon)$ if and only if $\mathcal{B}_i^{\mathcal{H}}\left[(z_i, s_i); R_{-i}, \epsilon\right] \neq \emptyset$. We also define for each $h \in \mathcal{H}$,

$$\begin{aligned}
\mathcal{B}_i^{h}\left[(z_i, s_i); R_{-i}, \epsilon\right] &= \operatorname*{proj}_h \mathcal{B}_i^{\mathcal{H}}\left[(z_i, s_i); R_{-i}, \epsilon\right] \\
&= \left\{\mu \in \Delta(\Theta_0 \times Z_{-i} \times S_{-i}(h));\ \exists \hat{\mu} \in \mathcal{B}_i^{\mathcal{H}}\left[(z_i, s_i); R_{-i}, \epsilon\right] \text{ sth. } \operatorname*{proj}_h \hat{\mu} = \hat{\mu}[h] = \mu\right\}
\end{aligned}$$

Note that the set $\mathcal{B}_i^{[\emptyset]}\left[(z_i, s_i); R_{-i}, \epsilon\right]$ describes those initial conjectures, with support in R_{-i}, of consistent conditional belief vectors that justifies s_i as an ϵ-sequential best reply for z_i. Finally, let $\mathcal{B}_i^{[\emptyset]}\left[z_i; R_{-i}, \epsilon\right] = \bigcup_{s_i} \mathcal{B}_i^{[\emptyset]}\left[(z_i, s_i); R_{-i}, \epsilon\right]$[12].

ϵ-Interim Sequential Rationalizability (ϵ-ISR) is a solution concept for Bayesian games in extensive form, defined as an iterated deletion procedure for each type of each player, as follows:

Definition 7. *Let $\epsilon \geq 0$. For each player $i \in N$, let $ISR_i^0(\epsilon) = Z_i \times S_i$. Recursively, for $l = 1, 2, ...,$ and $z_i \in Z_i$, let*

$$\begin{aligned}
ISR_{-i}^{l-1}(\epsilon) &= \bigtimes_{j \neq i} ISR_j^{l-1}(\epsilon), \\
ISR_i^{l}(z_i;\epsilon) &= \left\{\hat{s}_i \in ISR_i^{l-1}(z_i;\epsilon);\ \exists \hat{\mu}_i \in \varPhi_i(z_i)\ sth.\ (i)\ \hat{s}_i \in br_{i,z_i}(\hat{\mu}_i|\epsilon);\ (ii)\ \operatorname{supp}(\hat{\mu}_i[\emptyset]) \subseteq \Theta_0 \times ISR_{-i}^{l-1}(\epsilon)\right\} \\
ISR_i^{l}(\epsilon) &= \left\{(z_i, s_i) \in Z_i \times S_i;\ s_i \in ISR_i^{l}(z_i;\epsilon)\right\} \\
ISR^{l}(\epsilon) &= \bigtimes_{i \in N} ISR_i^{l}(\epsilon)
\end{aligned}$$

Let $ISR^{\infty}(\epsilon) := \bigcap_{l \geq 0} ISR^{l}(\epsilon)$

[12] A geometric interpretation comes from the fact that for each type z_i, and at each history h (including the beginning of the game $[\emptyset]$), each strategy defines a vector of payoff for player i, and the collection of all these vectors, one for each strategy, are included in a closed convex set. Hence, the set $\mathcal{B}_i^{[\emptyset]}\left[(z_i, s_i); R_{-i}\right]$, with $\epsilon = 0$, is a particular subset of the normal cone of that closed convex set at the vector of payoff induced by strategy s_i. A more detailed discussion can be found in section 2.3.

That is, $(z_i, s_i) \in Z_i \times S_i$ survives the lth round of elimination if and only if

$$ISR_i^l(\epsilon) = \left\{(z_i, s_i) \in ISR_i^{l-1}(\epsilon);\ \mathcal{B}_i^{\mathcal{H}}\left[(z_i, s_i); ISR_{-i}^{l-1}(\epsilon), \epsilon\right] \neq \emptyset\right\} \tag{2.5}$$

that is, if it is justified by a consistent conditional belief vector that, at the beginning of the game, is concentrated on pairs (z_{-i}, s_{-i}) consistent with previous rounds of elimination.

Penta (2012) argued that *ISR* is an extensive-form analogue of *Interim Correlated Rationalizability* (see Dekel, et al. (2007)), with the following properties: *i)* It coincides with *ICR* whenever the game is static, but due to sequential rationality, *ISR* refines *ICR* in dynamic environments; *ii)* *ISR* is belief-space invariant[13]; *iii)* *ISR* is upper hemicontinuous in the universal type space[14]; *iv)* *ISR* is an EFBRS.[15] It follows from *ii)* that there is no loss of generality in considering the universal type space Z^* as the domain of the correspondence ISR^∞. In what follows, we will identify a type by its belief hierarchy.

The starting point for the results in the next section is *The Structure Theorem for ISR* which shows that if all common knowledge assumptions on payoff are relaxed (that is, if $\{\Theta_0, (\Theta_i, u_i)_{i \in N}\}$ satisfies the Richness Condition, an important assumption defined later), then for any finite type profile $\hat{z} \in Z^*$, and any $\hat{s} \in ISR^\infty(\hat{z})$, there is a sequence $\{z^m\}_m$ converging to $\hat{z}$ and such that $\{\hat{s}\} = ISR^\infty(z^m)$ for every m. This result has two main consequences.

First, under the RC on Θ, it implies that any refinement of *ISR* that includes the epistemic condition of common initial belief in sequential rationality, is not uhc. In other words, given a model **Z**, any refinement of *ISR* may rule out strategies that can be carried out by types that approximate arbitrary close to the types in the model. Hence, *ISR* is the strongest solution concept whose predictions survive perturbations of beliefs and hierarchies of beliefs about payoffs, among those that satisfies the epistemic assumption of RCIBR.

Second, under the RC, *ISR* is a generically unique and a locally constant solution concept, that delivers multiple strategies at, and only at, the intersection of boundary points of disjoint open sets where *ISR* changes the selected strategy. Therefore, we can partition the universal type space as follows: For each $s \in S$, let $U^s = \{z \in Z^*; ISR^\infty(z) = \{s\}\}$. We have,

$$Z^*_{(+)} = \{z \in Z^*; |ISR(z)| > 1\}, \quad Z^*_u = \{z \in Z^*; |ISR(z)| = 1\} \tag{2.6}$$

[13] That is, if from two models we have types $z_i \in Z_i$ and $z_i' \in Z_i'$, such that $\hat{\beta}_i^*(z_i) = \hat{\beta}_i^*(z_i') \in \Theta_i \times Z_i^*$, then $ISR_i^\infty(z_i) = ISR_i^\infty(z_i')$.

[14] That is, for each $z \in Z^*$ and sequence $\{z^m\}$, such that $z^m \to z$, and for $\{s^m\} \subset S$ such that $s^m \to \hat{s}$ and $s^m \in ISR^\infty(z^m)$ for every m, $\hat{s} \in ISR^\infty(z)$.

[15] Formally, an epistemic characterization of *ISR* goes as follows: in a Γ_M-based belief-complete epistemic type structure, $(z, s) \in ISR^\infty$ iff $\exists \omega \in RCIBR^\infty$ such that $\nu^*(\omega) = \hat{\beta}^*(z)$, where $\nu^* : \Omega \to T^*$ describes beliefs and hierarchies of beliefs about payoffs, at the beginning of the game, of the epistemic types in ω. Hence, the collection $D \subset \Sigma$, where for each $i \in N$, $D_i = \cup_{z_i \in Z_i} \left\{(\hat{\theta}_i(z_i), ISR_i^\infty(z_i))\right\}$ is an EFBRS by Proposition (1).

Clearly, $Z_u^* = \cup_{s\in S} U^s$, and $z \in Z^*_{(+)}$ if and only if there exists $s, s' \in ISR^\infty(z)$ such that $s \neq s'$, $z \in bd(U^s) \cap bd(U^{s'})$, where $bd(U^s) = cl(U^s) \setminus U^s$. It should be clear that such partition depends on the Bayesian game in extensive form. It follows from Theorem 2 in Penta (2012), that $Z^*_{(+)}$ is a closed subset of a compact metrizable space. By the Baire Category Theorem, it is a Baire space[16]. Hence, $Z^*_{(+)}$ with its relative topology, will be considered as an ambient space for the study of critical types.

EFBRS

Given a game tree $(\Gamma_M, \geq)$, a payoff-information structure $\{\Theta_0, (\Theta_i, u_i)_{i\in N}\}$, and a model $\mathbf{Z}$, consider for every player $i \in N$, and every type $z_i \in Z_i$ with $\hat{\theta}_i(z_i) = \theta_i^{z_i} \in \Theta_i$, a subset of strategies $D_{i,z_i} \subset S_i$ for type z_i. Let $D = \Theta_0 \times \bigtimes_{i\in N} D_i$, and $D_{-i} = \Theta_0 \times \bigtimes_{j\neq i} D_j$ where for each $j \in N$, $D_j = \cup_{z_j \in Z_j} \{z_j\} \times D_{j,z_j}$. Let $\epsilon \geq 0$,

Definition 8. *The collection $(D_{i,z_i})_{i\in N, z_i \in Z_i}$ is called an **extensive-form ϵ-best reply set (ϵ-EFBRS)** if for every $i \in N$, and each $(z_i, s_i) \in D_i$ there is a consistent conditional belief vector $\hat{\mu}_{(z_i,s_i)} \in \Phi_i(z_i)$ satisfying,*

1. $(z_i, s_i) \in br_i\left(\hat{\mu}_{(z_i,s_i)}|\epsilon\right)$
2. $\hat{\mu}_{(z_i,s_i)}[\emptyset](D_{-i}) = 1$
3. $D_{i,z_i} \supseteq br_{i,z_i}(\hat{\mu}_{(z_i,s_i)}|\epsilon)$

The intuition is as follows: for any player $i \in N$ of type z_i, each strategy $s_i \in D_{i,z_i}$ is ϵ-undominated with respect to D_{-i}, the sets D_{i,z_i} are maximal with respect to this property, and there is common certainty, at the beginning of the game, of that event. It is not hard to see that, for given $\epsilon \geq 0$ and a Bayesian game in extensive form, there might be several ϵ-EFBRS, that the strategies in D are also in $ISR^\infty(\epsilon)$ which, at the same time, is the largest ϵ-EFBRS. These ideas are illustrated with the following example

Example 1 Continued. Consider again Example 1, where the rules of the game and payoff at θ_0 are described in Figure 1, while the type space, describing a situation with complete information about θ_0 was described in Table 1. We will show that, for $\epsilon \in (0,1)$, the collection $D_1 = \{a_1\}$ and $D_2 = \{b_2\}$ is an ϵ-EFBRS.

Recall that, at the beginning of the game, (In, a_3) is strictly dominated by a_1, so it can not be an ϵ-sequential best reply for player 1 to any consistent conditional belief of player 1. Since for $\hat{\mu}_{(z_1,a_1)}[\emptyset](\theta_0, z_2, b_2) = 1$, (In, a_2) cannot be a ϵ-sequential best reply for player 1 to such belief, we have that $\{a_1\} = br_{1,z_i}(\hat{\mu}_{(z_1,a_1)}|\epsilon)$.

[16] A topological space is a Baire space, if every intersection of countably many dense open sets in it is dense.

For player 2, $\hat{\mu}_{(z_2,b_2)}[\emptyset](D_1) = 1$ implies that $\hat{\mu}_{(z_2,b_2)}[\emptyset](\theta_0, z_1, a_1) = 1$. Hence, for $\hat{\mu}_{(z_2,b_2)}[In](0, z_1, (In, a_3)) = 1$, we have that $\{b_2\} = br_{2,z_2}(\hat{\mu}_{(z_2,b_2)}|\epsilon)$. In a similar way, it can be shown that the collection $D_1 = \{(In, a_2)\}$ and $D_2 = \{b_1\}$ is also an ϵ-*EFBRS*. □

In section 2.3 it will be discussed the relation between *EFBR* and *ISR* from an epistemic perspective.

2.2 Robustness and Critical Types

To capture the notion that similarity in beliefs corresponds to similarity in behavior, we consider the following two properties:

1. For each type z_i, strategy $s_i \in S_i$, and sequence $\{z_i^m\}_m$, sth. $z_i^m \to z_i$, and $s_i \in ISR_i^\infty(z_i^m)$, implies $s_i \in ISR_i^\infty(z_i)$

2. For each type z_i and strategy $s_i \in ISR_i^\infty(z_i)$ we have that, for every $\epsilon > 0$ and every $\{z_i^m\}_m$ sth. $z_i^m \to z_i$, there is an N_ϵ such that for all $m \geq N_\epsilon$, $s_i \in ISR_i^\infty(z_i^m|\epsilon)$.

Observe that (1) captures a notion of robustness at which an analyst can guarantee that every prediction that is rule out for a type, is rule out for every approximation. However, (2) corresponds, in this setup, to a stronger notion of robustness, in the sense that an analyst can guarantee that, for every prediction of a type, there is no perturbation ruling that prediction out, even when we relaxed the assumption of sequential rationality to ϵ-sequentially rational, for every $\epsilon > 0$. Since *ISR* satisfies (1), we say that predictions satisfy a notion of strategic robustness at a type space, if (2) holds too.

Since models of beliefs are considered as data in the analysis of games, if *ISR* were to satisfies strategic robustness for *almost every* model (in a topological sense), we can guarantee that models of beliefs are "good approximations to any idealized strategic situation", in the sense that predictions would not depend on specific details of the model. In particular, that is the case for *ICR* in static games.

However, I will show that for a non-negligible set of types, *ISR* does not satisfies (2) on $Z^*_{(+)}$, as defined at the end of section 2.1.4. I consider this an important result, specially when compared to the existing literature. Ely and Peski (2011), for example, showed that in static games, where all decisions are made simultaneously, Γ_0-critical types are non-generic.

In what follows, we work with the following definitions and assumption on Θ,

Definition 9. *A type $z_i \in Z_i^*$ is said to be*[17]*,*

[17] See Ely and Peski (2011).

1. Γ_M-*critical if there exist an* $\epsilon > 0$*, a strategy* $s_i \in S_i$*, and a sequence of types* $\{z_i^m\}$ *with* $z_i^m \to z_i$ *such that* $s_i \in ISR_i^\infty(z_i)$ *and* $s_i \notin ISR_i^\infty(z_i^m; \epsilon)$ *for every* m*.*

2. *Critical, if there is a game* Γ_M *for which it is* Γ_M*-critical.*

3. Γ_M*-regular (regular), if it is not* Γ_M*-critical (critical).*

Definition 10. *Strategy* s_i *is conditionally dominant at* θ *if* $\forall h \in \mathcal{H}(s_i)$*,* $\forall s_i' \in S_i(h)$*,* $\forall s_{-i} \in S_{-i}(h)$*, we have*

$$s_i(h) \neq s_i'(h) \Rightarrow u_i\left(\theta, \rho(s_i, s_{-i})\right) > u_i\left(\theta, \rho(s_i', s_{-i})\right) \tag{2.7}$$

Assumption 2. ***Richness Condition.*** *Let* $\{\Theta_0, (\Theta_i, u_i)_{i \in N}\}$ *be such that:*

i) $\forall\ \sigma = (\theta, s) \in \Sigma$*,* $\exists\ \theta^s \in \Theta$*:* $\forall i$ s_i *is conditionally dominant at* θ^s*, and* $u_i(\theta_0, \theta_i, \theta_{-i}, \rho(s)) = u_i(\theta_0, \theta_i, \theta_{-i}^s, \rho(s))$*,*

ii) $\forall i \in N$*,* Θ_i *is convex.*

Our first result, Proposition 1, provides an answer to the question, how can an analyst know in advance if predictions from *ISR* are to sensitive to the exact details of the type space chosen? It gives an epistemic foundation of Γ_M-critical types for *ISR*, as it will be discussed in section 2.3, since it is expressed only in terms of *EFBRS* sets. Example 1 suggests that whenever there are Γ_M-critical types in a model $\mathbf{Z}$, there exists an $\epsilon > 0$, a collection of ϵ-*EFBRS* sets, and a type $z_i \in Z_i$ such that $s_i \in ISR_i^\infty(z_i) \setminus D_{i,z_i}$. Proposition 1 shows that in general this is a necessary condition for Γ_M-critical types. Moreover, the condition is also sufficient.

Proposition 1. *Suppose that* $\{\Theta_0, (\Theta_i, u_i)_{i \in N}\}$ *satisfies the RC, and let* $\mathbf{Z}$ *be a type space.*

i) If there exists a collection $(D_{i,z_i})_{i \in N, z_i \in Z_i}$ *of* ϵ-EFBRS *with* $\epsilon > 0$ *and a type* $z_i \in Z_i$ *with* $\hat{\theta}_i(z_i) = \theta_i^*$ *such that* $s_i \in ISR_i^\infty(z_i) \setminus D_{i,z_i}$*, then type* z_i *is* Γ_M*-critical.*

ii) If the type space $\mathbf{Z}$ *is finite and a type* z_i *of player* i *with* $\hat{\theta}_i(z_i) = \theta_i^*$ *is* Γ_M*-critical, then there exists a collection* $(D_{i,z_i})_{i \in N, z_i \in Z_i}$ *of* ϵ-EFBRS *with* $\epsilon > 0$ *and a strategy* $s_i \in S_i$ *for player* i*, such that* $s_i \in ISR_i^\infty(z_i) \setminus D_{i,z_i}$*.*

In the Appendix, we provide the proof of this result, and in section 2.3 we give a detailed discussion of the interpretation of Proposition 1. However, the next example will help to illustrate how useful this result is to determine in advanced if types are critical or not, and in particular, it will provide some insights on the proof of the sufficient condition.

Example 3: A Signaling Game. Figure 2.1 describes a signaling game with two players. The payoff information structure $\{\Theta_0, (\Theta_i, u_i)_{i \in N}\}$ is given by $\Theta_0 = \{-3, 1, 3\}$, $\Theta_1 = \{\theta_s, \theta_w\}$,

$\Theta_2 = [0,2]$. Once each player learns about its own payoff type, at the beginning of the game, player 1 decides either to have Beer (B) or Quiche (Q) for breakfast. Player 2 observes player 1 move and chooses either to duel (d) or not (n), after which the game is over, and payoff realized. As discussed before, we make use of a type space $\mathbf{Z}$ to approximate an idealized strategic situation and study behavioral predictions from *ISR*.

Figure 2.1: Extensive-form representation for a Signaling Game with payoff for the model $\mathbf{Z}$.

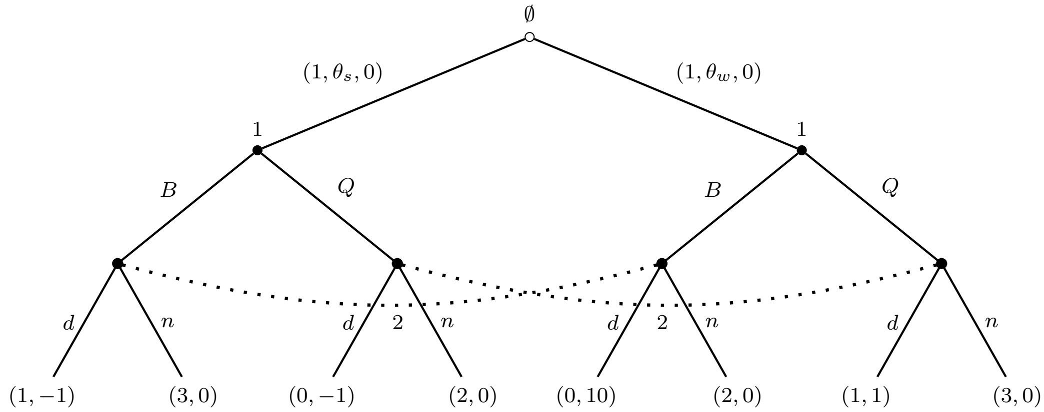

Table 2.1: A Normal form representation for a Signaling Game at $(\theta_0,\theta_s,\theta_2) = (1,\theta_s,0)$

$[\emptyset]-\{1\}$	(d_B,d_Q)	(d_B,n_Q)	(n_B,d_Q)	(n_B,n_Q)
$[B]-\{2\}$	$(1,-1)$	$(1,-1)$	$(3,0)$	$(3,0)$
$[Q]-\{2\}$	$(0,-1)$	$(2,0)$	$(0,-1)$	$(2,0)$

Table 2.2: A Normal form representation for a Signaling Game at $(\theta_0,\theta_w,\theta_2) = (1,\theta_w,0)$

$[\emptyset]-\{1\}$	(d_B,d_Q)	(d_B,n_Q)	(n_B,d_Q)	(n_B,n_Q)
$[B]-\{2\}$	$(0,1)$	$(0,1)$	$(2,0)$	$(2,0)$
$[Q]-\{2\}$	$(1,1)$	$(3,0)$	$(1,1)$	$(3,0)$

Table 2.3: For Player 1

$Z_1=\{z_s,z_w\}$	$(1,z_2)$		Θ_1
$\beta_1[z_s]$	1	$\hat{\theta}_1(z_s)$	θ_s
$\beta_1[z_w]$	1	$\hat{\theta}_1(z_w)$	θ_w

Table 2.4: For Player 2, with $p > 0.5$.

$Z_2=\{z_2\}$	$(1,z_s)$	$(1,z_w)$		Θ_2
$\beta_2[z_2]$	p	$1-p$	$\hat{\theta}_2(z_2)$	0

Tables 2.3 and 2.4 describe a model $\mathbf{Z} = \{\tilde{\Theta}_0, (Z_i,\hat{\theta}_i,\beta_i)_{i\in\{1,2\}}\}$ for this game with $\tilde{\Theta}_0 = \{1\}$, and for each player a set of types, $Z_1 = \{z_s, z_w\}$, and $Z_2 = \{z_2\}$[18].

We start by determining the set of *ISR* strategies for this game. It is not difficult to see that after the first stage of elimination for player 2 of type z_2, strategy (d_B,d_Q) does not survive[19]: at the beginning of the game, expected payoff from playing (d_B,d_Q) is equal to $1-2p<0$ for any $\hat{\mu}_2 \in \Phi_2(z_2)$, while the expected payoff from playing (n_B,n_Q) is always zero.

[18] Note that the model $\mathbf{Z}$ implicitly assumes that: (*i*) Its is common knowledge (CK) that for any given action of player 2, the strong type prefers beer whereas the weak type prefers quiche. (Let $\tilde{\Theta}_0 = \{1\}$ denote a taste-parameter of player 1); (*ii*) It is CK that player 1 does not like to duel, no matter what type he is.; (*iii*) It is CK that player 2 profits from challenging if and only if he faces the weak type. Hence, a concern about the robustness properties of *ISR* with respect to perturbation of belief and hierarchies of belief about payoff states is in order.

[19] We denote a strategy $s_i \in S_i$ by $s_i = (s_h)_{h\in\mathcal{H}(s_i)}$.

Therefore, after two iterations of *ISR*, we have that $\{B, Q\} = ISR_1^\infty(z_s) = ISR_1^\infty(z_w)$, and $\{(d_B, n_Q), (n_B, d_Q), (n_B, n_Q)\} = ISR_2^\infty(z_2)$.

For this example we will like to show that all types in the model $\mathbf{Z}$ from Tables 2.3 and 2.4 are Γ_M-critical types. That is, that there is an $\epsilon > 0$, strategies $\tilde{s}_i \in S_i$ for each $i \in \{1, 2\}$, and a sequence of types $\{z_i^m\}_m$ with $z_i^m \to z_i$ such that $\tilde{s}_i \in ISR_i^\infty(z_i)$ and $\tilde{s}_i \notin ISR_i^\infty(z_i^m; \epsilon)$ for every m. Let $\tilde{s}_1 = \{B\}$, and $\tilde{s}_2 = (n_B, d_Q)$, be a strategy profile such that $\tilde{s}_2 \in ISR_2^\infty(z_2)$, and $\tilde{s}_1 \in ISR_1^\infty(z_j)$ for $j \in \{s, w\}$. From proposition 1.i), the existence of a collection $(D_{i,z_i})_{i \in N, z_i \in Z_i}$ of ϵ-*EFBRS* with $\epsilon > 0$ such that $\tilde{s}_i \in ISR_i^\infty(z_i) \setminus D_{i,z_i}$ is enough to show that a type is Γ_M-critical. Let $\epsilon \in (0, 2p - 1)$, with $p > 0.5$, and consider the collection $(D_{i,z_i})_{i \in \{1,2\}, z_i \in Z_i}$ of strategies given by $D_{1,z_s} = D_{1,z_w} = \{Q\}$, and $D_{2,z_2} = \{(d_B, n_Q)\}$. We need to show that it is a collection of ϵ-*EFBRS* sets.

Consider player 2, with $s_2 = (d_B, n_Q) \in D_{2,z_2}$ and a conditional belief vector $\hat{\mu}_{(z_2,s_2)} \in \Phi_2(z_2)$ such that, at the beginning of the game, $\hat{\mu}_{(z_2,s_2)}[\emptyset](1, z_s, B) = \hat{\mu}_{(z_2,s_2)}[\emptyset](1, z_w, B) = 0$, $\hat{\mu}_{(z_2,s_2)}[\emptyset](1, z_s, Q) = p$, and $\hat{\mu}_{(z_2,s_2)}[\emptyset](1, z_w, Q) = 1 - p$, satisfying 2 of definition 8. Then, the expected payoff from playing s_2, at the beginning of the game, is zero which is ϵ-optimal with respect to the other three strategies. At history $[Q]$, player 2 observation of player 1 move does not contradict player 2 initial understanding of the game, hence from Bayes updating, $\hat{\mu}_{(z_2,s_2)}[Q](1, z_s, Q) = p$, and $\hat{\mu}_{(z_2,s_2)}[Q](1, z_w, Q) = 1 - p$. The expected gain at $[Q]$ from choosing n is ϵ-optimal with respect to d. However, at history $[B]$, the observed evidence contradicts player 2 initial understanding of the game, and therefore conditional beliefs are unrestricted for player 2. Hence, let $\hat{\mu}_{(z_2,s_2)}[B](1, z_w, B) = 1$, such that, at $[B]$, the expected gain from choosing d is ϵ-optimal with respect to n. This argument shows that condition 1 of definition 8 is satisfied: $s_2 = (d_B, n_Q)$ is an ϵ-sequential best reply to a consistent conditional conjecture $\hat{\mu}_{(z_2,s_2)} \in \Phi_2(z_2)$.

In order to satisfy 3 of definition 8, we show that actually s_2 is the unique ϵ-sequential best reply to $\hat{\mu}_{(z_2,s_2)}$. Recall that at the beginning of the game player 2 conjectures are such that player 1 moves Q with probability one. Since $\epsilon < 2p - 1$, it follows that (d_B, d_Q) and (n_B, d_Q), that is, choosing d after Q can not be an ϵ-sequentially best reply: by following any of those strategies player 2 gets $1 - 2p$ which is ϵ-suboptimal to playing s_2. However, note that conditional on history $[B]$, player 2 is certain that player 1 is of type z_w, an therefore moving n after B is ϵ-suboptimal to playing s_2. Finally, we have $\{s_2\} = br_{2,z_2}(\hat{\mu}_{(z_2,s_2)}|\epsilon) = D_{2,z_2}$.

Consider now player 1 of type z_j for $j \in \{s, w\}$, with a consistent conditional belief vector $\hat{\mu}_{(z_j,Q)} \in \Phi_1(z_j)$ such that, at the beginning of the game, $\hat{\mu}_{(z_j,Q)}[\emptyset](1, z_2, s_2) = 1$. Since $\epsilon < 2p - 1 \leq 1$, we have that $\{Q\} = br_{1,z_j}(\hat{\mu}_{(z_j,Q)}|\epsilon)$, for every $j \in \{s, w\}$.

Hence, it follows from proposition 1 that types in the type space $\mathbf{Z}$ are all critical. In what follows, we discuss some insights about the general proof of this result and provide some intuition on why the statement is true.

Figure 2.2: Extensive-form representation for a Signaling Game with general payoff structure.

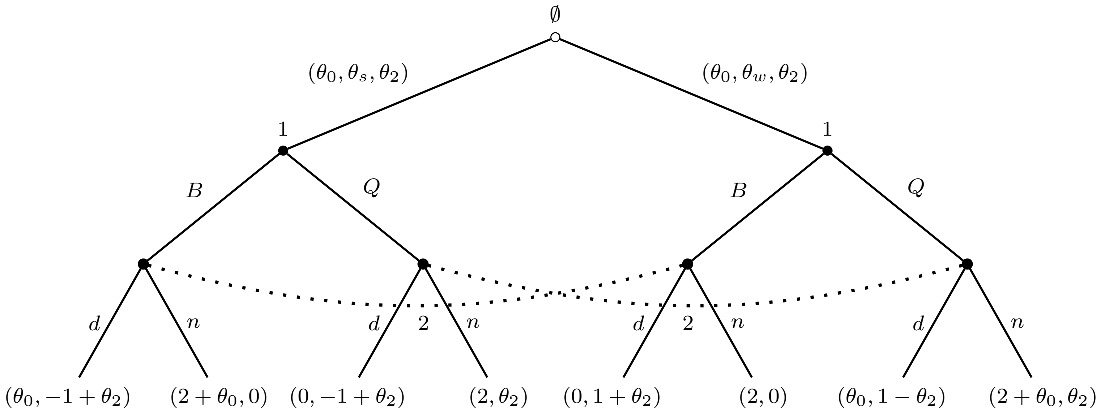

Note:

Table 2.5: A Normal form representation for a Signaling Game at $(\theta_0, \theta_s, \theta_2)$

$[\emptyset]-\{1\}$	(d_B, d_Q)	(d_B, n_Q)	(n_B, d_Q)	(n_B, n_Q)
$[B]-\{2\}$	$(\theta_0, -1+\theta_2)$	$(\theta_0, -1+\theta_2)$	$(2+\theta_0, 0)$	$(2+\theta_0, 0)$
$[Q]-\{2\}$	$(0, -1+\theta_2)$	$(2, \theta_2)$	$(0, -1+\theta_2)$	$(2, \theta_2)$

Table 2.6: A Normal form representation for a Signaling Game at $(\theta_0, \theta_w, \theta_2)$

$[\emptyset]-\{1\}$	(d_B, d_Q)	(d_B, n_Q)	(n_B, d_Q)	(n_B, n_Q)
$[B]-\{2\}$	$(0, 1+\theta_2)$	$(0, 1+\theta_2)$	$(2, 0)$	$(2, 0)$
$[Q]-\{2\}$	$(\theta_0, 1-\theta_2)$	$(2+\theta_0, \theta_2)$	$(\theta_0, 1-\theta_2)$	$(2+\theta_0, \theta_2)$

Recall that the payoff information structure $\{\Theta_0, (\Theta_i, u_i)_{i\in N}\}$ is given by $\Theta_0 = \{-3, 1, 3\}$, $\Theta_1 = \{\theta_s, \theta_w\}$, $\Theta_2 = [0, 2]$. We start the argument by defining a type space of dominant types, described in Tables 2.7 and 2.8. There is only one type for player 2, namely $\bar{z}_2$, with a payoff-type of 2, and mapping to a probability distribution over states and opponents types that puts probability p to $(\theta_0, \bar{z}_s)$ and probability $1-p$ to $(\theta_0, \bar{z}_w)$, as it is shown in Table 2.8. There are four dominant types for player 1, with payoff-type described in the last two columns of Table 2.7. Two of these dominant types for player 1, namely $\bar{z}_s$ and $\bar{z}_s'$, are of the θ_s-type, but type $\bar{z}_s$ is such that it puts probability one on the state $\theta_0 = -3$, while type $\bar{z}_s'$ is such that it puts probability one to the state $\theta_0 = 3$. A similar definition applies to the remaining two dominant types of player 1 in $\bar{Z}_1$. Given the original type space from Tables 2.3 and 2.4, and the collection $(D_{i,z_i})_{i,z_i\in Z_i}$ of ϵ-*EFBRS* given above, the existence of a type space of dominant types is justified by the **RC** on the payoff information structure. Its importance comes from the following: For player 1 of type $\bar{z}_j$, $\{Q\} = br_{1,\bar{z}_j}(\hat{\mu}|\epsilon)$ for every $\hat{\mu} \in \Phi(\bar{z}_j)$, for $j \in \{s, w\}$, that is $\{Q\} = ISR_1^1(\bar{z}_j; \epsilon)$ for $j \in \{s, w\}$; given this, at the next round of elimination, for player 2 of type $\bar{z}_2$, consistent conjectures are such that at the beginning of the game puts probability one on Q, so that conditional on the history $[B]$, any conditional belief player 2 is such that $\{(d_B, n_Q)\} = ISR_2^2(\bar{z}_2;\epsilon)$. Therefore $\{Q\} = ISR_1^\infty(\bar{z}_j;\epsilon)$ for $j \in \{s,w\}$, and $\{(d_B, n_Q)\} = ISR_2^\infty(\bar{z}_2;\epsilon)$. Dominant types are such that they have as unique ϵ-*ISR* the profile of strategies in the collection of ϵ-*EFBRS* from above. At the same time, is not hard to see that $\{B\} = ISR_1^\infty(\bar{z}'_j;\epsilon)$ for $j \in \{s,w\}$.

Table 2.7: A Type Space of dominant types for Player 1

$\bar{Z}_1 = \{\bar{z}_s, \bar{z}_w, \bar{z}'_s, \bar{z}'_w\}$	$(3,\bar{z}_2)$	$(-3,\bar{z}_2)$		Θ_1
$\beta_1[\bar{z}_s]$	0	1	$\hat{\theta}_1(\bar{z}_s)$	θ_s
$\beta_1[\bar{z}_w]$	1	0	$\hat{\theta}_1(\bar{z}_w)$	θ_w
$\beta_1[\bar{z}'_s]$	1	0	$\hat{\theta}_1(\bar{z}'_s)$	θ_s
$\beta_1[\bar{z}'_w]$	0	1	$\hat{\theta}_1(\bar{z}'_w)$	θ_w

Table 2.8: A Type Space of dominant types for Player 2, with $p > 0.5$, and $\theta_0 \neq 1$.

$\bar{Z}_2 = \{\bar{z}_2\}$	$(\theta_0,\bar{z}_s)$	$(\theta_0,\bar{z}_w)$		Θ_2
$\beta_2[\bar{z}_2]$	p	$1-p$	$\hat{\theta}_2(\bar{z}_2)$	2

Consider now the type space of perturbed types described in Table 2.9 and 2.10. The set of perturbed types for player 2, namely Z_2^ρ, consists of the dominant type from Table 2.8, and a collection of perturbed types, indexed by $\rho \in (0,1]$. Payoff-type of each z_2^ρ is a convex combination between the payoff-type of z_2 from Table 2.4, and the payoff-type of $\bar{z}_2$ from Table 2.8: $\hat{\theta}_2(z_2^\rho) = \rho\hat{\theta}_2(\bar{z}_2) + (1-\rho)\hat{\theta}_2(z_2) = 2\rho$; and it maps to a probability measure over states in Θ_0 and opponents types in Z_1^ρ that is itself a mixture: with probability ρ its uniform over the set of dominant types of player 1, and with probability $(1-\rho)$ it is a probability measure that is certain that the state is $\theta_0 = 1$, puts probability p to player 1 type z_s^ρ, and probability $1-\rho$ to z_w^ρ. A similar intuition applies for the set of perturbed types of player 1. The general interpretation is that for ρ small enough, the perturbed type space is approximated closed to the original type space described in Tables 2.3 and 2.4.

Table 2.9: A Type Space of perturbed types for Player 1

$Z_1^\rho = \bar{Z}_1 \cup \{(z_s^\rho),(z_w^\rho)\}$	$(-3,\bar{z}_2)$	$(3,\bar{z}_2)$	$(1,z_2^\rho)$		Θ_1
$\beta_1[z_s^\rho]$	ρ	0	$1-\rho$	$\hat{\theta}_1(z_s^\rho)$	θ_s
$\beta_1[z_w^\rho]$	0	ρ	$1-\rho$	$\hat{\theta}_1(z_w^\rho)$	θ_w

Table 2.10: A Type Space of perturbed types for Player 2, with $p > 0.5$.

$Z_2^\rho = \bar{Z}_2 \cup \{(z_2^\rho)\}$	$(1,z_s^\rho)$	$(1,z_w^\rho)$	$(\theta_0,z_1) \in \Theta_0 \times \bar{Z}_1$		Θ_2
$\beta_2[z_2^\rho]$	$p(1-\rho)$	$(1-p)(1-\rho)$	$\frac{\rho}{4}$	$\hat{\theta}_2(z_2^\rho)$	2ρ

Given $\epsilon \in (0, 2p-1)$, and for each $l \geq 1$, and $\rho \in (0,1]$, we are interested in describing the sets $ISR_1^l(z_s^\rho;\epsilon)$, $ISR_1^l(z_w^\rho;\epsilon)$, and $ISR_s^l(z_2^\rho;\epsilon)$, and finally showing that $z_i^\rho \to z_i$ as $\rho \to \infty$.

For the first round of elimination, consider player 1 of type z_s^ρ. Every consistent conjecture $\hat{\mu}_1 \in \Phi_1(z_s^\rho)$ is such that $\hat{\mu}_1[\emptyset](-3,\bar{z}_2,(d_B,n_Q)) = \rho$, and $\sum_{s_2}\hat{\mu}_1[\emptyset](1,z_2^\rho,s_2) = 1-\rho$. By choosing Q, it gets an expected payoff of at least 2ρ and at most 2, while by choosing B it gets a strictly negative expected payoff, which is ϵ-suboptimal for every $\rho \in (0,1]$. Therefore $ISR_1^l(z_s^\rho;\epsilon) = \{Q\}$, for every $\rho \in (0,1]$.

For the next round of elimination, player 2 initial consistent conjectures $\hat{\mu}_2 \in \Phi_2(z_2^\rho)$ are such that $\hat{\mu}_2[\emptyset](\theta_0, \bar{z}_s, Q) = \hat{\mu}_2[\emptyset](\theta_0, \bar{z}_w, Q) = \hat{\mu}_2[\emptyset](\theta_0, \bar{z}'_s, B) = \hat{\mu}_2[\emptyset](\theta_0, \bar{z}'_w, B) = \frac{\rho}{4}$, $\hat{\mu}_2[\emptyset](1, z_s^\rho, Q) = p(1-\rho)$, and $\hat{\mu}_2[\emptyset](1, z_w^\rho, Q) + \hat{\mu}_2[\emptyset](1, z_w^\rho, B) = (1-p)(1-\rho)$, for every $\rho \in (0,1]$. Note that the construction of the perturbed type space for player 2 is such that every consistent conjecture in $\Phi_i(z_2^\rho)$ is restricted to put, at the beginning of the game, positive probability to every history. Given this it follows that strategy (n_B, d_Q) does not survive this round for every $\rho \in (0,1]$: Expected payoff at the beginning of the game from following this strategy are at most $(1-\rho)(1-2\rho)(1-2p) < 0$, and at history $[Q]$ that amount is scaled by the initial total probability of reaching that node, which is strictly positive. However, by following (d_B, n_Q) player 2 gets at least 2ρ, and since $\epsilon \in (0, 2p-1)$, (n_B, d_Q) is ϵ-suboptimal to (d_B, n_Q) for every $\rho \in (0,1]$, that is $(1-\rho)(1-2\rho)(1-2p) < 2\rho - \epsilon$. A similar argument holds for the ϵ-suboptimality of (d_B, d_Q) to (d_B, n_Q). In other words, given that initially there is a strictly positive probability of reaching history $[Q]$, at that node, moving d is ϵ-suboptimal to moving n for player 2, for every $\rho \in (0,1]$. Therefore, $ISR_2^2(z_2^\rho;\epsilon) = \{(d_B, n_Q), (n_B, n_Q)\}$.

For the next round of elimination, note that for player 1 of type z_w^ρ, every consistent conjecture $\hat{\mu}_1 \in \Phi_1(z_w^\rho)$ is such that $\hat{\mu}_1[\emptyset](3, \bar{z}_2, (d_B, n_Q)) = \rho$, and $\hat{\mu}_1[\emptyset](1, z_2^\rho, (d_B, n_Q)) + \hat{\mu}_1[\emptyset](1, z_2^\rho, (n_B, n_Q)) = 1-\rho$. Therefore, by choosing Q, it gets an expected payoff of $3+2\rho$, while by choosing B it gets at most $2(1-\rho)$, which makes B ϵ-suboptimal to Q for every $\rho \in (0,1]$. Therefore $ISR_1^3(z_1^\rho;\epsilon) = \{Q\}$, for every $\rho \in (0,1]$.

Nothing else gets eliminated on the rounds that follows, and the procedure stops with $ISR_1^\infty(z_j^\rho;\epsilon) = \{Q\}$, for $j \in \{s,w\}$ and $ISR_2^\infty(z_2^\rho;\epsilon) = \{(d_B, n_Q), (n_B, n_Q)\}$, for every $\rho \in (0,1]$. Finally, is not hard to see that $z_i^\rho \to z_i$ for $i \in \{s,w,2\}$ as $\rho \downarrow 0$. It follows by induction over the k-order belief $\hat{\beta}_i^k(z_i^\rho)$, and that $\hat{\beta}_i^k(z_i^\rho) \to \hat{\beta}_i^k(z_i)$ as $\rho \downarrow 0$, for every $i \in \{s,w,2\}$, and $k \geq 0$.

Hence, we have shown that the types in the original type space from tables 2.3 and 2.4 are critical; that is, for $\epsilon \in (0, 2p-1)$, strategies $\tilde{s}_1 = \{B\}$, and $\tilde{s}_2 = \{(n_B, d_Q)\}$, the sequence of perturbed types (z_i^ρ) with $z_i^\rho \to z_i$ as $\rho \downarrow 0$, such that $\tilde{s}_1 \in ISR_1^\infty(z_j)$ for $j \in \{s,w\}$, $\tilde{s}_2 \in ISR_2^\infty(z_2)$, and $\tilde{s}_1 \notin ISR_1^\infty(z_j^\rho;\epsilon)$ for $j \in \{s,w\}$, $\tilde{s}_2 \in ISR_2^\infty(z_2^\rho;\epsilon)$ for every $\rho \in (0,1]$.

□

The example suggest that if there is any hope to identify Γ_M-critical types, we need a further exploration on the subsets of $Z^*_{(+)}$, since every type in Z^*_u is clearly a Γ_M-regular type. In addition, from Theorem 1 there is another partition of the universal type space. Let $C_i^{\Gamma_M}$ denote the set of Γ_M-critical types of player i, and let $R_i^{\Gamma_M}$ denote the set of types which are not Γ_M-critical. It is clear that $C^{\Gamma_M} \subseteq Z^*_{(+)}$. Also, it is clear that types that express uniqueness can not be critical, hence, $Z_u^* \subseteq R^{\Gamma_M}$.

Our next result describes a special property of these partitions. In particular, we use our previous characterization result to show that, under certain conditions, the set of Γ_M-critical types is dense relative to the space of types with multiplicity. Before we present the formal proof, I attempt to provide some intuition on why in dynamic environments the topological properties of such a set are significantly different than in a static set up by looking at some "special cases"of types with multiplicity. We start by defining some objects that would make the presentation of our result more clear.

Let W denote the set of all path w in Γ_M[20]. For each path $w \in W$, and each $h \in w$, $\Theta_0 \times Z_{-i} \times S_{-i}(h)$ is a closed subset of $\Theta_0 \times Z_{-i} \times S_{-i}$. Consider the family of upper semicontinuous maps[21] $\hat{O}_h : \Delta(\Theta_0 \times Z_{-i} \times S_{-i}) \to [0,1]$ defined by $\hat{O}_h(\mu) = \mu(\Theta_0 \times Z_{-i} \times S_{-i}(h)) \in [0,1]$. Given $\rho > 0$, define open sets in $\Delta(\Theta_0 \times Z_{-i} \times S_{-i})$,

$$\hat{O}(w;\rho) = \bigcup_{h \in w} \left\{\mu \in \Delta(\Theta_0 \times Z_{-i} \times S_{-i});\ \hat{O}_h(\mu) < \rho\right\} \tag{2.8}$$

$$\hat{O}(\rho) = \bigcup_{w \in W} \hat{O}(w;\rho) \tag{2.9}$$

and let $\hat{O} = \bigcap_{\rho>0} \hat{O}(\rho)$. Since for given $\rho > 0$, $\mu \in \hat{O}(\rho)$ if and only if there is a path $w \in W$ and $h \in w$ such that $\mu(\Theta_0 \times Z_{-i} \times S_{-i}(h)) < \rho$, it follows that for each $\mu \in \hat{O}$ there exists a path $w \in W$ and $h \in w$ such that for each $(\theta_0, z_{-i}, s_{-i}) \in \operatorname{supp}(\mu)$ we have $(\theta_0, z_{-i}, s_{-i}) \notin \Theta_0 \times Z_{-i} \times S_{-i}(h)$.

Finally, (2.6) suggests the definition of the following map: Given that $Z_{-i,u} \subset Z_{-i}^*$ is an open subset, define a lower semicontinuous map $\hat{u} : \Delta(\Theta_0 \times Z_{-i}^*) \to [0,1]$ by $\hat{u}(\nu) = \nu(\Theta_0 \times Z_{-i,u}) \in [0,1]$. For each $n \in \mathbb{N}$, consider the collection of open subsets of types for player i,

$$Z_{i,(+),u}^n = \left\{z_i \in Z_{i,(+)};\ \hat{u} \circ \beta_i(z_i) > 1 - \frac{1}{n}\right\} \tag{2.10}$$

In addition, let $Z_{i,(+),u} = \{z_i \in Z_{i,(+)};\ \hat{u} \circ \beta_i(z_i) = 1\}$. Hence, a type $z_i \in Z_{i,(+),u}$, with payoff-type $\hat{\theta}_i(z_i) = \theta_i^{z_i}$, is such that beliefs assign probability 1 to opponents types with a unique *ISR* strategy. Since for such a type there are multiple *ISR* strategies, for each $s_i^l \in ISR_i^\infty(z_i)$, there are conjectures $\hat{\mu}_i^l \in \Delta^{\mathcal{H}}(\Theta_0 \times Z_{-i} \times S_{-i})$ such that:

1. $s_i^l \in br_{i,z_i}(\hat{\mu}_i^l)$;

2. $\beta_i(z_i) = \operatorname{marg}_{\Theta_0 \times Z_{-i}} \hat{\mu}_i^l[\emptyset]$; with $\operatorname{supp}\left(\hat{\mu}_i^l[\emptyset]\right) \subseteq \Theta_0 \times ISR_{-i}^\infty$,

Since for each $(\theta_0, z_{-i}) \in \operatorname{supp}(\beta_i(z_i))$ with $\hat{\theta}_{-i}(z_{-i}) = \theta_{-i}^{z_{-i}}$, there exists a unique profile $s_{-i}^{z_{-i}} \in S_{-i}$ such that $ISR_{-i}^\infty(z_{-i}) = \{s_{-i}^{z_{-i}}\}$, it should be clear that there is a unique $\mu^{z_i} \in$

[20]cf. footnote 6. Since the game is finite, we can identify the set W with the set of terminal nodes $\mathcal{Z}$.
[21]The reader should review footnote 9 and Kechris, 1995, chapter 17.

$\Delta(\Theta_0 \times ISR^{\infty}_{-i})$, such that $\hat{\mu}^l_i[\emptyset] = \mu^{z_i}$ for every conjecture justifying any strategy in $ISR^{\infty}_i(z_i)$: that is, for $z_i \in Z_{i,(+),u}$, $\mathcal{B}^{[\emptyset]}_i\left[z_i; ISR^{\infty}_{-i}\right] = \{\mu^{z_i}\}$.

Consider, now, a path $w \in W$ with the following property: for every $h \in w$, there exists $(\theta_0, z_{-i}, s_{-i}) \in \text{supp}(\mu^{z_i})$ such that $(\theta_0, z_{-i}, s_{-i}) \in \Theta_0 \times Z_{-i} \times S_{-i}(h)$. Then, along the path w, each conjecture $\hat{\mu}^l_i[h]$ is constructed from μ^{z_i} via Bayes rule. Therefore, it must be the case that along w there is a payoff-tie at h between any $s^l_i \in ISR^{\infty}_i(z_i)$ for which $h \in \mathcal{H}(s^l_i)$, that is, for every $\hat{s}_i \in S_i(h)$

$$\mathbb{E}_{\hat{\mu}^l_i[h]}[U_i(z_i, s^l_i, \theta_0, z_{-i}, s_{-i})] = \mathbb{E}_{\hat{\mu}^l_i[h]}[U_i(z_i, s^{l'}_i, \theta_0, z_{-i}, s_{-i})] \geq \mathbb{E}_{\hat{\mu}^l_i[h]}[U_i(z_i, \hat{s}_i, \theta_0, z_{-i}, s_{-i})]$$

If the game were static, then for every player every strategy has a unique decision node that cannot be discarded under any consistent conjecture, and is a common decision node among all strategies of every player. Hence, it must be the case that for types with multiple rationalizable actions that assign probability 1 to opponents with uniqueness, there is a payoff-tie across all rationalizable actions. By an application of Proposition 1, they cannot be critical. For every $\epsilon > 0$, and every sequence of types converging to such a limit, it must be the case that for sufficiently closed types, rationalizable actions in the limit must be ϵ-rationalizable for those closed enough types.

However, in dynamic games different strategies may have different decision nodes, or even if they have the same decision nodes some of those may be discarded under some conditional conjecture. Therefore, we have,

Proposition 2. *Suppose that $\{\Theta_0, (\Theta_i, u_i)_{i \in N}\}$ satisfies the RC and, for a finite model $\boldsymbol{Z}$, let $z_i \in Z_{i,(+),u}$. If $\mathcal{B}^{[\emptyset]}_i\left[z_i; ISR^{\infty}_{-i}\right] \in \hat{O}$, then $z_i \in C^{\Gamma_M}_i$.*

That is, $\hat{Z}_{i,(+),u} := \left\{z_i \in Z_{i,(+),u};\ \mathcal{B}^{[\emptyset]}_i\left[z_i; ISR^{\infty}_{-i}\right] \in \hat{O}\right\} \subset C^{\Gamma_M}_i$. The next proposition presents our main result on the size of $C^{\Gamma_M}_i$. It says that the set $\hat{Z}_{i,(+),u}$ is comeager in $Z_{i,(+)}$. As we mentioned before, since $Z_{i,(+)}$ is a Baire space, every comeager set is dense, so the set of critical types, relative to $Z_{i,(+)}$, contains a comeager set, and it is therefore dense in $Z_{i,(+)}$.

Theorem 1. *Suppose that $\{\Theta_0, (\Theta_i, u_i)_{i \in N}\}$ satisfies the RC. Then $\hat{Z}_{i,(+),u}$ is dense in $Z_{i,(+)}$, where $Z_{i,(+)}$ is endow with the relative product topology.*

The proof of both results can be found in the appendix. We leave to the next section a discussion on the interpretation and importance of this result. However, let me give some intuition for the proof of Proposition (2). Recall our geometric interpretation of the set $\mathcal{B}^{[\emptyset]}_i\left[(z_i, s_i); ISR^{\infty}_{-i}\right]$, as the normal cone of a closed convex set. When the type space is finite, such a set is finite-dimensional. For a type in $\hat{Z}_{i,(+),u}$, the condition $\mathcal{B}^{[\emptyset]}_i\left[z_i; ISR^{\infty}_{-i}\right] \in \hat{O}$ says that it is possible to break ties at histories that at the beginning of the game have probability zero under the unique conjecture in $\mathcal{B}_i^{[\emptyset]}\left[z_i; ISR_{-i}^{\infty}\right]$. The reason that makes it possible to break ties comes from the fact that at such history the dimension of $\mathcal{B}_i^h\left[(z_i, s_i); Z_{-i}\times S_{-i}\right]$ changes with respect to $\mathcal{B}_i^{[\emptyset]}\left[(z_i, s_i); ISR_{-i}^{\infty}\right]$, since the support of the conditional conjecture is not restricted at all at histories with zero probability at the beginning of the game.

2.3 Discussion

2.3.1 Epistemic Type Structures

First-order conditional beliefs, as defined in section 2.1.2, do not exhaust all the uncertainty that a player might face in a dynamic strategic situation. To analyze how a player reasons about other players' reasoning, it must be specified a description of player's higher order conditional beliefs. We follow Battigalli and Siniscalchi (1999, 2002) to define a Γ_M-based epistemic type structure as a model of players hierarchies of conditional beliefs[22].

Definition 11. *A Γ_M-based epistemic type structure is a tuple $(\Omega_0, (\Omega_i, \Psi_i, \varkappa_i)_{i\in N})$ such that $\Omega_0 = \Theta_0 \times \{\psi_0\}$ and, for every $i \in N$, Ψ_i is a compact metrizable space of epistemic types, $\Omega_i \subseteq \Sigma_i \times \Psi_i$ is a closed subset with $\operatorname{proj}_{\Sigma_i}\Omega_i = \Sigma_i$, endowed with its Borel σ-algebra $\mathcal{B}(\Omega_i)$. Finally, let $\varkappa_i : \Psi_i \to \Delta^{\mathcal{H}}(\Omega_{-i})$ be a continuous mapping.*
A Γ_M-based epistemic type structure is belief-complete if, in addition, for each i, $\Omega_i = \Sigma_i \times \Psi_i$ and $\varkappa_i : \Psi_i \to \Delta^{\mathcal{H}}(\Omega_{-i})$ is onto.

Elements $\omega_i = (\theta_i, s_i, \psi_i) \in \Omega_i$ are called states of player i, describing player i's payoff-type, θ_i, a strategy s_i, and an epistemic type ψ_i. An epistemic type implicitly encodes a hierarchy of conditional beliefs about: the opponents' strategy-payoff types pairs; about the conditional belief vectors that every opponent has about his opponents' strategy-payoff types pairs; about the conditional belief vectors that every opponent has about the conditional belief vectors that each of his opponents has about his opponents' strategy-payoff types pairs; and so on; all of that summarized as a conditional belief vector over other players' states.

In particular, for an epistemic type $\psi_i \in \Psi_i$ of player i, let $f_i : \Psi_i \to [\Delta(\Sigma_{-i})]^{\mathcal{H}}$ denote the *first-order belief mapping*, defined by

$$f_{i,h}(\psi_i) = \operatorname*{marg}_{\Sigma_{-i}} \varkappa_{i,h}(\psi_i) \quad \forall h \in \mathcal{H} \tag{2.11}$$

[22]Conditions for existence of a canonical belief-complete epistemic type structure, which satisfies to be universal in the sense that it contains all reasonable hierarchies of conditional beliefs, (where reasonably refers to the property of a hierarchy of belief being coherent, and this being common belief by the players), and additional technical properties were provided by Battigalli and Siniscalchi (1999).

It is not hard to see that the mapping is continuous and, in general, $f_i(\Psi_i) \subseteq \Delta^{\mathcal{H}}(\Sigma_{-i})$. In a similar fashion it is possible to derive a hierarchy of conditional belief, that is, an infinite collection of conditional belief vectors, with each element in the collection representing a k-order conditional belief.

Given a Γ_M-based epistemic type structure, we define belief operators as follows: For each player $i \in N$, history $h \in \mathcal{H}$, and event $E \in \mathcal{B}(\Omega_{-i})$, let

$$B_{i,h}(E) = \{(\sigma_i, \psi_i) \in \Omega_i;\ \varkappa_{i,h}(\psi_i)[E] = 1\} \in \mathcal{B}(\Omega_i) \tag{2.12}$$

define the event *"player i would be certain of E, were he to observe that history h was reached"*. Finally, we define recursively the n-fold composition on an operator C of an event E by: $C^0(E) = E$, and for $n \geq 1$, $C^n(E) = C(C^{n-1}(E))$.

An arbitrary Γ_M-based epistemic type structure[23], describes the set $\Omega = \bigtimes_{i \in N_0} \Omega_i$ of possible *states of the world*. In it, it is commonly believed which beliefs players consider possible and which ones are not. This is of considerable importance since it directly influences the behavioral implications and the interpretation of predictions in game theoretic models. Battigalli and Friedenberg (2012) considered this as describing a context to the game within which any additional restriction about beliefs may be, or not, transparent to the players (see Battigalli and Prestipino (2012)). Therefore, an epistemic type structure encodes a context to the game, and *"this context influences what beliefs players do versus do not consider possible"* (Battigalli and Friedenberg (2012), page 79).

Under this interpretation, an analyst might consider (potentially) additional assumptions that restrict only first-order conditional beliefs. Then, a context to the game can be constructed as an epistemic type structure in which the assumptions on first-order conditional beliefs are transparent to the players: a context at which not only the restrictions are true, but there is also common belief in the restrictions at every node.

Formally, fix an underlying epistemic type structure $(\Omega_0, (\Omega_i^*, \Psi_i^*, \varkappa_i^*)_{i \in N})$, not necessarily required to be belief-complete. Let a profile $\Delta = (\Delta_{\theta_i})_{i \in N, \theta_i \in \Theta_i}$ be such that, for each player $i \in N$, and each payoff-type $\theta_i \in \Theta_i$, $\Delta_{\theta_i} \subseteq \Delta^{\mathcal{H}}(\Sigma_{-i})$ is a non-empty set of assumptions about first-order conditional belief exogenously determined, eg. by an analyst. Given Δ, for every $n \geq 0$, let $\Psi_0^n[\Delta] = \Theta_0 \times \{\psi_0\}$, and for every $i \in N$, define the event

$$\Psi_i^0[\Delta] = \{(\theta_i, s_i, \psi_i) \in \Omega_i^*;\ f_i(\psi_i) \in \Delta_{\theta_i}\} \tag{2.13}$$

[23]For example, an epistemic type structure which is not belief complete, in which additional assumptions about beliefs are imposed by the analyst, relative to which events are defined.

describing those individual states in which the restrictions hold for player i; and for $n \geq 1$

$$\Psi_i^n[\Delta] = \left\{ (\theta_i, s_i, \psi_i) \in \Omega_i^*;\ \forall h \in \mathcal{H},\ \varkappa_{i,h}^*(\psi_i) \left[\bigtimes_{j \neq i} \Psi_j^{n-1}[\Delta] \right] = 1 \right\} \tag{2.14}$$

is the event describing those individual states in which player i believes at every node that opponents believes at every node,...(n-times), that the restrictions are true. Let $\bar{\Psi}_i[\Delta] = \bigcap_{n \geq 0} \Psi_i^n[\Delta]$ for every $i \in N_0$, and define a context of the game consistent with assumptions Δ by $\bar{\Psi}[\Delta] := \bigtimes_{i \in N_0} \bar{\Psi}_i[\Delta] \subseteq \Omega^*$. We say that assumptions Δ on first-order conditional beliefs are transparent at state $\omega \in \Omega^*$ whenever $\omega \in \bar{\Psi}[\Delta]$.

In particular, for a given type space $\mathbf{Z}$, let Δ be such that for each player i, they coincide with the first-order subjective beliefs about payoff-states induced by the type. Then, in any context in which those assumptions are transparent we obtain, at the beginning of the game, the same subjective beliefs and hierarchies of belief about payoff-states as those described by the model $\mathbf{Z}$.

2.3.2 Sequential Rationality and Common Initial Belief in Sequential Rationality (RCIBR)

In this section we formally define RCIBR as an event in an underlying space of states of the world, and consider two solution concepts: Extensive-Form Best Reply Sets (*EFBRS*, Battigalli and Friedenberg 2012) as defined in section 8, that identifies *some* strategy-payoff type pairs consistent with RCIBR in an arbitrary epistemic type structure; and *ISR*, which identifies *all* strategy-payoff type pairs consistent with RCIBR in a belief-complete epistemic type space. The discussion from this section justify the consideration of *EFBRS* along with *ISR*.

We assume that each player, given his payoff type $\theta_i \in \Theta_i$ and first-order conditional beliefs $\mu_i \in \Delta^{\mathcal{H}}(\Sigma_{-i})$, chooses a strategy $s_i \in S_i$ that is optimal at every history $h \in \mathcal{H}$ that s_i leads to. Formally,

Definition 12. *Let $\mu_i \in \Delta^{\mathcal{H}}(\Sigma_{-i})$, and $(\theta_i, s_i) \in \Sigma_i$. A strategy s_i is a sequential best reply to μ_i for θ_i if and only if, for every $h \in \mathcal{H}(s_i)$, and every $s_i' \in S_i(h)$,*

$$\mathbb{E}_{\mu_i[h]}\left[U_i(\theta_i, s_i, \sigma_{-i})\right] \geq \mathbb{E}_{\mu_i[h]}\left[U_i(\theta_i, s_i', \sigma_{-i})\right] \tag{2.15}$$

For any $\mu_i \in \Delta^{\mathcal{H}}(\Sigma_{-i})$, let $br_{i,\theta_i}(\mu_i)$ denote the set of sequential best replies to μ_i for θ_i, and let $br_i(\mu_i) = \{(\theta_i, s_i) \in \Sigma_i;\ s_i \in br_{i,\theta_i}(\mu_i)\}$.

A strategy $s_i \in S_i$ is sequentially rational for payoff-type θ_i, written $s_i \in r_i(\theta_i)$, if there exists a $\mu_i \in \Delta^{\mathcal{H}}(\Sigma_{-i})$ such that $s_i \in br_{i,\theta_i}(\mu_i)$.

For a given Γ_M-based epistemic type structure define the (measurable) event *"player i is sequentially rational"*[24]

$$R_i := \{(\sigma_i, \psi_i) \in \Omega_i;\ \sigma_i \in br_i(f_i(\psi_i))\} \tag{2.16}$$

Sequential rationality provides a criteria for those reasonable strategies. However, reasonable behavioral predictions requires that reasonable strategies be supported by reasonable first- and higher-order conditional beliefs.

We say that an epistemic type, $\psi_i \in \Psi_i$ expresses *initial belief in opponents sequential rationality* if, at the beginning of the game, first-order conditional beliefs puts positive probability on opponents strategy-payoff type pairs for which the strategy is sequentially rational for the payoff type. Note that the condition above refers to $f_{i,\emptyset}(\psi_i) \in \Delta(\Sigma_{-i})$.

We say that an epistemic type expresses 2-*fold initial belief in opponents sequential rationality*, if in addition, at the beginning of the game, second-order conditional belief puts positive probability on opponents' first-order conditional beliefs that express initial belief in opponents sequential rationality. An epistemic type expresses common initial belief in opponents sequential rationality, if it expresses k-fold initial belief in opponents sequential rationality, for every $k \geq 1$.

Formally, for a given Γ_M-based epistemic type structure $(\Omega_0, (\Omega_i, \Psi_i, \varkappa_i)_{i\in N})$, let for each $i \in N_0$, $R_i^0 = R_i \in \mathcal{B}(\Omega_i)$, and consider the sequence of events, for every $k \geq 1$

$$R_i^k = R_i^{k-1} \cap B_{i,\emptyset}(R_{-i}^{k-1}) \tag{2.17}$$

$$RCIBR_i^{\infty} = \cap_{k\geq 0} R_i^k \tag{2.18}$$

For $k = 1$, and for every $i \in N$, R_i^1 is the event in which R_i holds true and player i believes (with probability 1), at the beginning of the game, that the event R_{-i} is true. A similar interpretation holds for every $k \geq 1$. It follows that $RCIBR^{\infty} = \times_{i\in N_0} RCIBR_i^{\infty}$ represent the event in which $R = \times_{i\in N_0} R_i$ holds, and it is commonly believed (with probability 1) at the beginning of the game. Therefore, $\mathrm{proj}_{\Sigma}\, RCIBR^{\infty}$ describes the behavioral implications of the epistemic condition RCIBR as those strategy-payoff type pairs for which there exist an epistemic type that expresses common initial belief in sequential rationality, and the strategy is sequentially rational for the payoff type, given the first-order conditional belief of the epistemic type.

Is it possible to identify those strategy-payoff type pairs by analyzing only the basic struc-

[24]For each i and $\theta_i \in \Theta_i$, br_{i,θ_i} is a non-empty-valued upper-hemicontinuous correspondence by standard dynamic programming arguments. Since f_i is continuous, we have that R_i is homeomorphic to the graph of an upper-hemicontinous correspondence $br_i \circ f_i$, it follows by the closed graph theorem that R_i is closed, hence it is in $\mathcal{B}(\Omega_i)$.

ture of the game? Lemma 1 below, adapted from Battigalli and Prestipino (2012), answers to the characterization question mentioned above. In particular, it identifies an *EFBRS* as the behavioral implications of RCIBR in some epistemic type structure.

Lemma 1. *For every* $(\Gamma_M, \geq)$ *and* $\{\Theta_0, (\Theta_i, u_i)_{i\in N}\}$*:*

1. *In any* Γ_M*-based epistemic type structure* $\operatorname{proj}_\Sigma RCIBR^\infty$ *is an EFBRS;*
2. *For every EFBRS, D, there is a* Γ_M*-based epistemic type structure such that* $D = \operatorname{proj}_\Sigma RCIBR^\infty$*. In particular, there is a profile* Δ*, such that* $\bar{\Psi}[\Delta]$ *is a* Γ_M*-based epistemic type structure within which* $D = \operatorname{proj}_\Sigma RCIBR^\infty[\Delta]$.

In other words, *EFBRS* captures the behavioral implication of RCIBR across all epistemic type structures. When the Δ restrictions are derived from a type space $\mathbf{Z}$, lemma 1 justify definition 8, and describes the behavioral implications of RCIBR across all epistemic type structures in which the restrictions from the type space $\mathbf{Z}$ are transparent.

Note that condition (1) and (2) of definition 8 are necessary conditions for the epistemic characterization result: given some Γ_M-based epistemic type structure in which $(\theta_i, s_i) \in \operatorname{proj}_{\Sigma_i} RCIBR_i^\infty$, there is an epistemic type with first-order conditional belief vector $\mu_{(\theta_i,s_i)}$ for which (1) hold. And, since the epistemic type initially believes in R_{-i}^k, for every $k \geq 0$, then it initially believes in $\cap_k R_{-i}^k$. Therefore, its first-order conditional belief is such that it assigns, at the beginning of the game, probability one to $D_{-i} = \operatorname{proj}_{\Sigma_{-i}} RCIBR_{-i}^\infty$, so that (2) hold.

However, (1) and (2) are not sufficient for the characterization. In particular, they would characterize sequential rationality and at most initial belief in opponents rationality in some epistemic type structure. Condition (3) of definition 8 requires the sets D to be maximal, in the sense of including all sequential best replies to a given (first-order) conditional belief. Hence, from the perspective of any opponent that believes on player i's rationality, every possible strategy in that event is considered. By including condition (3), we can guarantee a characterization of sequential rationality and k-fold initial belief in opponents sequential rationality, for every $k \geq 1$.

2.3.3 An Epistemic Argument of Proposition 1 and 2

Given a game tree $(\Gamma_M, \geq)$, and a payoff-information structure $\Theta = \{\Theta_0, (\Theta_i, u_i)\}$ define, for each $(\theta_i, s_i) \in \Sigma_i$, the mapping $\bar{U}_i[\theta_i, s_i] : \Sigma_{-i} \to \mathbb{R}$ given by $\bar{U}_i[\theta_i, s_i](\sigma_{-i}) = U_i((\theta_i, s_i), \sigma_{-i})$. Hence, given $\theta_i \in \Theta_i$ and for each history $h \in \mathcal{H}$, let $\bar{U}_{i,h}[\theta_i] = \{\bar{U}_i[\theta_i, s_i] : s_i \in S_i(h)\}$. In particular, at the beginning of the game, $\bar{U}_{i,[\emptyset]}[\theta_i] = \{\bar{U}_i[\theta_i, s_i] : s_i \in S_i\}$.

Given a type space $\mathbf{Z}$, and a type for player i with $z_i \in Z_{i,(+)}$ we would like to understand why is it that there are multiple *ISR* strategies. Note that Corollary 3 in Penta (2012) provides a partial answer to that question. However, it follows from the epistemic characterization of *ISR* that given a type $z_i \in Z_{i,(+)}$ with $\hat{\theta}_i(z_i) = \theta_i^{z_i}$, it holds that $(z_i, s_i^1), (z_i, s_i^2) \in ISR_i^\infty$, with $s_i^1 \neq s_i^2$, whenever there are states of player i, $\omega_i^1 = (\theta_i^{z_i}, s_i^1, \psi_i^1)$, and $\omega_i^2 = (\theta_i^{z_i}, s_i^2, \psi_i^2)$ such that $\omega_i^l \in RCIBR_i$ and $\hat{\beta}_i^*(z_i) = \nu_i^*(\omega_i^l)$ for $l = 1, 2$. Given $\sigma_i^l = (\theta_i^{z_i}, s_i^l)$, define,

$$RCIBR_i(\sigma_i^l) = \{\psi_i \in \Psi_i^*;\ (\sigma_i^l, \psi_i) \in RCIBR_i \text{ and } \hat{\beta}_i^*(z_i) = \nu_i^*(\sigma_i^l, \psi_i)\} \tag{2.19}$$

Since the sets $RCIBR_i(\sigma_i^l)$ can be identify with a particular subset of the normal cone of a convex set containing $\bar{U}_{i,[\emptyset]}[\theta_i^{z_i}]$, at the vector $\bar{U}_i[\theta_i^{z_i}, s_i^l]$, there are two mutually exclusive alternatives: either

(A) $RCIBR_i(\theta_i^{z_i}, s_i^1) \setminus RCIBR_i(\theta_i^{z_i}, s_i^2) \neq \emptyset$ and $RCIBR_i(\theta_i^{z_i}, s_i^2) \setminus RCIBR_i(\theta_i^{z_i}, s_i^1) \neq \emptyset$

or

(B) $RCIBR_i(\theta_i^{z_i}, s_i^1) = RCIBR_i(\theta_i^{z_i}, s_i^2)$

That is, (A) refers to a situation in which there are different epistemic justifications to different strategies satisfying the exogenous restrictions imposed by z_i, at the beginning of the game. (B) refers to a situation in which every epistemic justification of one strategy is an epistemic justification of the other, and vice versa. Given our characterization result of Γ_M-critical types, if (A) were true for type z_i as described above, then z_i is a good candidate to be a Γ_M-critical type. This is actually the content of Proposition 1: A type is critical if and only if additional informational restrictions describing a context to the game exclude some strategies that would have been predicted by *ISR* in a context-free situation. When a researcher is concerned about misspecifications of (subjective) beliefs and higher-order beliefs, it is necessary to give up on the hypothesis that belief restrictions are fully transparent to both the players and the researcher, since subjective beliefs encoded in nearby types may be consistent with players beliefs in different contexts. Therefore, it becomes essential to investigate the behavioral implications of *RCIBR* across **all** reasonable contexts to the game, that is, all reasonable ways of how players interactively reason about uncertainty.

If the game were static, it should be clear that (A) holds if and only if z_i is such that it assigns positive probability to opponents types with multiple rationalizable strategies, while (B) holds if and only if z_i assigns positive probability only to opponents types with a unique rationalizable strategy. It basically refers to a payoff-tie of multiple strategies to a unique first-order belief. The above explains why, for static games, (B) holds if and only if $z_i \in Z_{i,(+),u}$, and it suggest that no type in $Z_{i,(+),u}$ can be Γ_0-critical.

In dynamic environments the above sufficient and necessary conditions for either (A) or (B) to hold are not true, and that is the content of Proposition 2. As long as a type in $Z_{i,(+),u}$ has an initial valid conjecture that assigns probability zero to a node in a path in which a decision needs to be made, there is no way how to discipline a player beliefs at that node, specially when all beliefs are allowed, as it is implicit on the epistemic characterization of ISR. Therefore, under the RC, basically, arbitrary behavior is possible following that node.

Then, if it is common knowledge to the players (while not transparent to the researcher) that some beliefs about the strategic uncertainty are not possible, while still being consistent with the exogenous restrictions on subjective beliefs about payoff-relevant states, there are approximating types with subjective beliefs arbitrarily close, for which among the players it still is common knowledge that some beliefs about the strategic uncertainty are not possible. Therefore, for that approximation, some ISR-strategies of the original type must be excluded.

2.3.4 On the Implications of Theorem 1

Typically, researchers know that they do not know in full detail any informational aspect of a game, for example, a player beliefs, and hierarchies of beliefs about payoff-relevant states, and also about the strategic uncertainty that they face. Therefore, game-theoretic model relies implicitly on common knowledge assumptions, in order to approximate an ideal situation. Hence, it is natural to ask when a researcher can guarantee that predictions from a solution concept do not depend on the details of those implicit assumptions. I consider this a basic criteria for a theory in order to be useful.

A particular form of those common knowledge assumptions occurs when the context of the game is commonly known not only to the players but also to the researcher. We show that, when it comes to the problem of robustness as we did in this paper, misspecification of the context for dynamic games has large consequences on the behavioral predictions derive from *ISR*, and is actually what is behind the fact that *ISR* is highly sensitive to perturbations of higher order beliefs about payoff-relevant states.

The above has strong implications on the modeling technique that considers the possibility of modeling beliefs about payoff relevant states independently of beliefs about strategic uncertainty. The possibility for a researcher to consider only subjective beliefs about payoff relevant states to determine rationalizable behavior in dynamic settings, allows players to consider arbitrary ways of reasoning in terms of strategic uncertainty, in particular on how the game is going to unfold.

Chapter 3

Bibliography

1. Battigalli, P., and A. Friedenberg. (2012), Forward Induction Reasoning Revisited. Theoretical Economics 7:57-98.

2. Battigalli, P., and A. Prestipino. (2012), Transparent restrictions on beliefs and forward induction reasoning in games with asymmetric information. The BE Journal of Theoretical Economics 13(1), 79-130.

3. Battigalli, P., and M. Siniscalchi. (1999), Hierarchies of Conditional Beliefs and Interactive Epistemology in Dynamic Games. Journal of Economic Theory 88:188-230.

4. Battigalli, P., and M. Siniscalchi. (2002), Strong Belief and Forward Induction Reasoning. Journal of Economic Theory 106:356-91.

5. Battigalli, P. and Siniscalchi, M. (2003), Rationalization and Incomplete Information, Advances in Theoretical Economics, 3(1), Article 3.

6. Ben-Porath, E. (1997), Rationality, Nash Equilibrium and Backwards Induction in Perfect Information Games, Review of Economic Studies, 64, 23-46.

7. Bernheim, D. (1984), Rationalizable Strategic Behavior, Econometrica, 52, 1007-1028.

8. Carlsson, H., and E. Van Damme (1993), Global Games and Equilibrium Selection, Econometrica, 61, 989-1018.

9. Dekel, E., D. Fudenberg, and S. Morris (2007). Interim correlated rationalizability. Theoretical Economics 2, 15-40.

10. Ely, J. C. and M. Peski (2011). Critical types. Review of Economic Studies 78, 907-937.

11. Harsanyi, J. (1967-1968). Games with incomplete information played by bayesian players, I-III. Management Science 14(3), 159-182.

12. Mertens, J.-F. and S. Zamir (1985). Formulation of Bayesian analysis for games with incomplete information. International Journal of Game Theory 14(1), 1-29.

13. Penta, A. (2012a), Higher Order Uncertainty and Information: Static and Dynamic Games, Econometrica, 80, 631-660.

14. Pearce, D. (1984). Rationalizable Strategic Behavior and the Problem of Perfection. Econometrica, 52, 1029-1050.

15. Rubinstein, A. (1989). The electronic mail game: A game with almost common knowledge. American Economic Review 79, 389-391.

16. Weinstein, J. and M. Yildiz (2007). A structure theorem for rationalizability with application to robust predictions of refinements. Econometrica 75(2), 365-400.

Chapter 4

Appendix to Chapter 2

4.1 Preliminaries

This section intends to provide proofs of the results presented in the paper. We start with the presentation of important concepts for completeness of the results.

4.1.1 Θ-based belief space and the Universal Belief Space

For each player $i \in N$, let $Z_i^0 = \Theta_0 \times \Theta_{-i}$, and for $k \geq 1$ $Z_i^k = Z_i^{k-1} \times \Delta(Z_{-i}^{k-1})$. The collection $\{\Delta(Z_i^k)_{k\geq 0}\}$ induces the space of coherent hierarchies of beliefs, represented as an element $z_i = (\mu_1, \mu_2, ...) \in Z_i^* \subset \bigtimes_{k\geq 1} \Delta(Z_i^k)$ for which the condition that for each k, the belief hierarchy of order $k+1$ is an extension of the belief hierarchy of order k holds true. That is, for every k,

$$\begin{aligned} \hat{Z}_i^1 &:= \Delta(\Theta_0 \times \Theta_{-i}) \\ \hat{Z}_i^2 &:= \left\{ \mu_2 = (\mu_1, \nu_1) \in \hat{Z}_i^1 \times \Delta(\Theta_0 \times \Theta_{-i} \times \hat{Z}_{-i}^1);\ \underset{\Theta_0\times\Theta_{-i}}{\operatorname{marg}}\ \nu_1 = \mu_1 \right\} \end{aligned}$$

and for $k \geq 2$,

$$\hat{Z}_i^{k+1} := \left\{ \mu_{k+1} = (\mu_k, \nu_k) \in \hat{Z}_i^k \times \Delta(\Theta_0 \times \Theta_{-i} \times \hat{Z}_{-i}^k);\ \underset{\Theta_0\times\Theta_{-i}\times\hat{Z}_{-i}^{k-1}}{\operatorname{marg}}\ \nu_k = \nu_{k-1}, \text{where}\, \mu_k = (\mu_{k-1}, \nu_{k-1}) \right\}$$

An element $\mu_k \in \hat{Z}_i^k$ is called a coherent belief hierarchy of order k. z_i is called a type of player i. Z_i^* is called the *universal type space.*

It can be shown that each Z_i^* is compact, and in fact, there exists a natural homeomorphism $\eta_i^* : Z_i^* \rightarrow \Delta(\Theta_0 \times \Theta_{-i} \times Z_{-i}^*)$ that preserves beliefs of all orders; that is, for each $z_i =$

$(\mu_1, \mu_2, ...) \in Z_i^*$,

$$\underset{\Theta_0\times\Theta_{-i}\times Z_i^{k-1}}{\text{marg}} \eta_i^*(z_i) = \mu_k \tag{4.1}$$

That η_i^* is one-to-one means that different types have different belief over $\Theta_0 \times \Theta_{-i}$ and others players types. That η_i^* is onto means that any belief about $\Theta_0 \times \Theta_{-i}$ and others players types is held by some type of player i. Let's say that the tuple $\mathcal{T}^* = \{\Theta, (T_i^*, \theta_i^*, \beta_i^*)_{i\in N}\}$, where $T_i^* := \Theta_i \times Z_i^*$, and the functions $\theta_i^* : T_i^* \to \Theta_i$ and $\beta_i^* : T_i^* \to \Delta(\Theta_0 \times T_{-i}^*)$ are given by $\theta_i^*(\theta_i, z_i) = \theta_i$ and $\beta_i^*(\theta_i, z_i) = \eta_i^*(z_i)$, is a canonical representation of the universal belief space.

An important step in the interpretation of our results is provided by the fact that each hierarchy of conditional belief in a belief-complete Γ_M- based epistemic type structure $(\Omega_0, (\Omega_i, \Psi_i^*, \varkappa_i)_{i\in N})$ induces naturally a hierarchy of belief, at the beginning of the game, in a Θ-based type structure. For this, we define a map $\nu^* : \Omega_i \to Z_i^*$ by $\nu^* = m_i^* \circ \rho_{i,\emptyset}^*$, where $m_i^* : \Omega_{i,\emptyset} \to Z_i^*$, and $\rho_{i,\emptyset}^* : \Omega_i \to \Omega_{i,\emptyset}$ are defined recursively as follows:

Step 0: Let $X_i^0 = \Sigma_{-i} = \Theta_0 \times \Theta_{-i} \times S_{-i}$, and define, recursively, for every $k \geq 1$, $X_i^k = X_i^{k-1} \times \Delta^{\mathcal{H}}(X_{-i}^{k-1})$. The collection $\{\Delta^{\mathcal{H}}(X_i^k)\}_k$ induces the set of hierarchies of CPS, and as before, we let $\Psi_i^* \subset \bigtimes_k \Delta^{\mathcal{H}}(X_i^k)$ to denote the set of i's collectively coherent CPS-hierarchies.

In addition, for each player i, let $Y_i^0 = \Theta_0 \times \Theta_{-i} \times S_{-i}$, and for each $k \in \mathbb{N}$ define Y_i^k and Y_i^* in the same way as Z_i^* was constructed.

Step 1: For any Borel spaces Y, Y', and measurable mapping $f : Y \to Y'$ we define an associated pushforward mapping $\hat{f} : \Delta(Y) \to \Delta(Y')$ by $\hat{f}(\mu)[C'] = \mu(f^{-1}(C'))$ for any measurable $C' \subset Y'$.

Let $\rho_i^0 : X_i^0 \to Y_i^0$ be the identity map and $(\hat{\rho}_{i,h}^0)_{h\in\mathcal{H}} : \Delta^{\mathcal{H}}(X_i^0) \to \times_{h\in\mathcal{H}}\Delta(Y_i^0)$ the collection of pushforward maps. In particular, consider the pushforward of the initial (first-order) beliefs, $\hat{\rho}_{i,\emptyset}^0 : \Delta(X_i^0) \to \Delta(Y_i^0)$. Then recursively,

$$\rho_{i,\emptyset}^k : X_i^k \to Y_i^k : (x_i^{k-1}, \psi_{-i}^{k-1}) \to (\rho_{i,\emptyset}^{k-1}(x_i^{k-1}), \hat{\rho}_{i,\emptyset}^{k-1}(\psi_{-i}^{k-1})) \tag{4.2}$$

Finally, let $\rho_{i,\emptyset}^* : \Omega_i \to Y_i^*$ be defined as $(\theta_i, s_i, \psi_i^1, \psi_i^2, ...) \to (\theta_i, s_i, \hat{\rho}_{i,\emptyset}^1(\psi_i^1), \hat{\rho}_{i,\emptyset}^2(\psi_i^2), ...)$. Hence each individual state $\omega_i \in \Omega_i$ induces a $\Theta \times S$-hierarchy of belief at the beginning of the game.

Step 2: Each $\Theta \times S$-hierarchy of belief induces a Θ-hierarchy of belief, naturally, via recursive marginalization. Let $m_i^0 : Y_i^0 \to Z_i^0$ be the natural projection, and define recursively, $m_i^k : Y_i^k \to Z_i^k$ as $(y_i^{k-1}, \pi_{-i}^k) \to (m_i^{k-1}(y_i^{k-1}), \hat{m}_{-i}^{k-1}(\pi_{-i}^k))$, and let $\hat{m}_i^k : \Delta(Y_i^k) \to \Delta(Z_i^k)$ be the associated pushforward mapping. Finally, define $m_i^* : Y_i^* \to Z_i^*$ as $(\theta_i, s_i, \pi_i^1, \pi_i^2, ...) \to (\theta_i, s_i, \hat{m}_i^0(\pi_i^1), \hat{m}_i^0(\pi_i^2), ...)$, and therefore the mapping $\nu^* : \Omega_i \to Z_i^*$ is a well defined object.

4.2 Proofs of results in Chapter 2

[a.] Decomposition of decision nodes

For any player $i \in N$, any pair $(\theta_i, s_i) \in \Sigma_i$ and CPS $\mu_i \in \Delta^{\mathcal{H}}(\Sigma_{-i})$, we define for any $h \in \Gamma_M$,

$$\begin{aligned} D_h(s_i, \mu_i; \theta_i) \quad = \quad & \{(\tilde{\theta}, \tilde{h}) \in \Theta \times \Gamma_M;\ (\tilde{\theta}, h) > (\tilde{\theta}, \tilde{h}) \text{ with } \tilde{\theta}_i = \theta_i, \text{ and} \\ \text{if} \quad & \tilde{\theta}_{-i} \in \operatorname{supp}\left(\operatorname*{marg}_{\Theta_{-i}} \mu_i[h]\right) \text{ then } \forall s_{-i} \in \operatorname{supp}\left(\mu_i[h](\tilde{\theta}_{-i}, \cdot)\right), \tilde{h} \in D_h(s_i, s_{-i})\} \end{aligned}$$

the set of relevant nodes discarded at h under s_i and μ_i for player $i \in N$ of payoff-type θ_i.

Sequential rationality requires that a strategy be a best reply for a payoff-type θ_i, at every node not discarded by s_i, that is, at every $h \in \mathcal{H}(s_i)$. Hence, when player $i \in N$, of payoff-type θ_i, evaluates a strategy s_i under a (first-order) conditional belief μ_i, the collection of sets $D_h(s_i, \mu_i; \theta_i)$, for every $h \in \Gamma_M$, induces a partition of $\mathcal{H}(s_i)$,

$$\mathcal{H}(s_i) = \left[\mathcal{H}(s_i) \cap \bigcup_{l=\emptyset}^{L} A_{[l]}(s_i, \mu_i)\right] \cup \left[\mathcal{H}(s_i) \setminus \left(\bigcup_{l=\emptyset}^{L} A_{[l]}(s_i, \mu_i)\right)\right] \tag{4.3}$$

as follows: For $h \in \mathcal{H}(s_i)$, we have, $h \in A_{[\emptyset]}(s_i, \mu_i)$ if and only if $((\theta_i, \tilde{\theta}_{-i}), h) \in D_{[\emptyset]}(s_i, \mu_i; \theta_i)$ $\forall \tilde{\theta}_{-i} \in \operatorname{supp}\left(\operatorname{marg}_{\Theta_{-i}} \mu_i[\emptyset]\right)$. Then for every $h' \in \mathcal{H}(s_i)$ such that $h > h'$, we have that $h' \in A_{[\emptyset]}(s_i, \mu_i)$. Therefore, $\mathcal{H}(s_i) = \left(\mathcal{H}(s_i) \cap A_{[\emptyset]}(s_i, \mu_i)\right) \cup \left(\mathcal{H}(s_i) \setminus A_{[\emptyset]}(s_i, \mu_i)\right)$.

Let $h^{[1]} \in \mathcal{H}(s_i) \setminus A_{[\emptyset]}(s_i, \mu_i)$, be such that there is $\tilde{\theta}_{-i} \in \operatorname{supp}(\operatorname{marg}_{\Theta_{-i}} \mu_i[\emptyset])$ and $s_{-i} \in \operatorname{supp}(\mu_i[\emptyset](\tilde{\theta}_{-i}, \cdot))$ such that $((\theta_i, \tilde{\theta}_{-i}), h^{[1]}) \notin D_{[\emptyset]}(s_i, \mu_i; \theta_i)$, and $h^{[1]} \notin D_{[\emptyset]}(s_i, s_{-i})$ and, in addition, for which there is no other relevant node $((\theta_i, \tilde{\theta}_{-i}), h)$ with $h \in \mathcal{H}(s_i) \setminus A_{[\emptyset]}(s_i, \mu_i)$, and $((\theta_i, \tilde{\theta}_{-i}), h) > ((\theta_i, \tilde{\theta}_{-i}), h^{[1]})$. Let $\mathcal{H}^{[1]}(s_i)$ denote the set of all such $h^{[1]} \in \mathcal{H}(s_i) \setminus A_{[\emptyset]}(s_i, \mu_i)$.

At any node $h^{[1]} \in \mathcal{H}^{[1]}(s_i)$, $\mu_i\left[h^{[1]}\right]$ is obtained via Bayes rule from $\mu_i[\emptyset] \in \Delta(\Sigma_{-i})$. Then for $h' \in \mathcal{H}(s_i)$, we have $h' \in A_{[h^{[1]}]}(s_i, \mu_i)$ if and only if $((\theta_i, \tilde{\theta}_{-i}), h') \in D_{h^{[1]}}(s_i, \mu_i; \theta_i)$ $\forall \tilde{\theta}_{-i} \in \operatorname{supp}\left(\operatorname{marg}_{\Theta_{-i}} \mu_i[h^{[1]}]\right)$. For every $h'' \in \mathcal{H}(s_i)$ such that $h' > h''$, we have that $h'' \in A_{[h^{[1]}]}(s_i, \mu_i)$. Let $A_{[1]}(s_i, \mu_i) = \bigcup_{h^{[1]} \in \mathcal{H}^{[1]}(s_i)} A_{[h^{[1]}]}(s_i, \mu_i)$. Therefore, we have,

$$\mathcal{H}(s_i) = \left[\mathcal{H}(s_i) \cap \left(A_{[\emptyset]}(s_i, \mu_i) \cup A_{[1]}(s_i, \mu_i)\right)\right] \cup \left[\mathcal{H}(s_i) \setminus \left(A_{[\emptyset]}(s_i, \mu_i) \cup A_{[1]}(s_i, \mu_i)\right)\right] \tag{4.4}$$

By construction, observe that $A_{[\emptyset]}(s_i, \mu_i)$ and $A_{[1]}(s_i, \mu_i)$ are disjoint. Recursively, given that a collection of disjoint subsets of nodes $(A_{[l]})_{l=\emptyset}^{n}$ has been defined, let $h^{[n+1]} \in \mathcal{H}(s_i) \setminus \bigcup_{i=\emptyset}^{n} A_{[l]}$, be a node in $\mathcal{H}^{[n+1]}(s_i)$. By the same argument as in the previous paragraph, let $A_{[n+1]} := \bigcup_{h^{[n+1]} \in \mathcal{H}^{[n+1]}(s_i)} A_{[h^{[n+1]}]}(s_i, \mu_i)$.

Since the game is finite, there is a finite integer L for which the procedure just described stops. Hence, every node $h' \in \mathcal{H}(s_i)$ either, it belongs to $A_{[l]}(s_i, \mu_i)$, for some $\emptyset \leq l \leq L$, and so is discarded under s_i and μ_i at some history; or, is such that for every $1 \leq l \leq L$ there is $\hat{\theta}_{-i} \in \operatorname{supp}(\operatorname{marg}_{\Theta_{-i}} \mu_i[h^{[l]}])$ and $s_{-i} \in \operatorname{supp}(\mu_i[h^{[l]}](\hat{\theta}_{-i}, \cdot))$ such that $((\theta_i, \hat{\theta}_{-i}), h') \notin D_{h^{[l]}}(s_i, \mu_i; \theta_i)$, and $h' \notin D_{h^{[l]}}(s_i, s_{-i})$.

[b.] Proofs.

By the RC, we define a mapping $g : \Sigma \to \Sigma$ by $g(\theta, s) = (\theta^s, s)$, as described in Assumption (2). In addition, for each $\sigma \in \Sigma$, let $\bar{z}(\sigma) = (\bar{z}_i(\sigma))_{i \in N}$ be such that for each $i \in N$:

$$\hat{\theta}_i\left(\bar{z}_i(\sigma)\right) = \operatorname*{proj}_{\Theta_i} g(\sigma) = \theta_i^s \text{ and } \beta_i\left(\bar{z}_i(\sigma)\right)\left[(\theta_0^s(\sigma), \bar{z}_{-i}(\sigma))\right] = 1 \tag{4.5}$$

where $\theta_0^s(\sigma) = \operatorname{proj}_{\Theta_0} g(\sigma)$. Let $\bar{Z} = \{\bar{z}(\sigma) \in Z^*; \sigma \in \Sigma\}$ denote the set of "dominant type profiles"and, for every $i \in N$, let $\bar{Z}_i$ be the projection on player i's types. Let $\bar{Z}_{-i} = \times_{j \neq i} \bar{Z}_j$.

For each player $i \in N$, and for each $(\theta_i^*, s_i^*) \in \Sigma_i$, define, the following sets:

$$A^{(\theta_i^*, s_i^*)} = \{\sigma \in \Sigma;\ \operatorname{proj}_{\Sigma_i} g(\sigma) = (\theta_i^*, s_i^*)\} \qquad A_{(\theta_i^*, s_i^*)} = \{\sigma \in \Sigma;\ \operatorname{proj}_{\Sigma_i} \sigma = (\theta_i^*, s_i^*)\}$$

$$\bar{Z}^{(\theta_i^*, s_i^*)} = \{\bar{z}(\sigma) \in \bar{Z};\ \sigma \in A^{(\theta_i^*, s_i^*)}\} \qquad \bar{Z}_{(\theta_i^*, s_i^*)} = \{\bar{z}(\sigma) \in \bar{Z};\ \sigma \in A_{(\theta_i^*, s_i^*)}\}$$

Hence, in each profile $\bar{z}(\sigma) \in \bar{Z}^{(\theta_i^*, s_i^*)}$, players' i payoff type is θ_i^*, and s_i^* is conditionally dominant at every element in the support of his beliefs, while each profile $\bar{z}(\sigma) \in \bar{Z}_{(\theta_i^*, s_i^*)}$, is associated with a profile θ^s of payoff-types at which s_i^* is conditionally dominant.

Proof of Proposition 1. [i]: By hypothesis, let $(D_{i,z_i})_{i \in N, z_i \in Z_i}$ be a collection of ϵ-EFBRS sets with $\epsilon > 0$, and let $z_i \in Z_i$ with payoff-type $\hat{\theta}_i(z_i) = \theta_i^*$ be such that $s_i \in ISR_i^\infty(z_i) \setminus D_{i,z_i}$. Consider $s_i^* \in D_{i,z_i}$, for which there is $\hat{\mu}_{(z_i, s_i^*)} \in \Phi_i(z_i)$ satisfying $s_i^* \in br_{i,z_i}\left(\hat{\mu}_{(z_i,s_i^*)}|\epsilon\right) \subseteq D_{i,z_i}$, and $\hat{\mu}_{(z_i,s_i^*)}[\emptyset]\left(D_{-i}\right)$. Moreover, for each $j \in N$, and $z_j \in Z_j$ with $\hat{\theta}_j(z_j) = \theta_j^{z_j}$ we have that for each $s_j \in D_{j,z_j}$, there is $\hat{\mu}_{(z_j,s_j)} \in \Phi_j(z_j)$ satisfying $s_j \in br_{j,z_j}\left(\hat{\mu}_{(z_j,s_j)}|\epsilon\right) \subseteq D_{i,z_j}$, and $\hat{\mu}_{(z_j,s_j)}[\emptyset]\left(D_{-j}\right)$.

Given $\rho \in (0, 1]$, we construct a model of perturbed types in the following way: let Z^ρ be a set types such that $\bar{Z} \subset Z^\rho$; and since to each triplet in the original model $(\theta_0, z, s) \in \Theta_0 \times Z \times D$, with $s_j \in D_{j,\theta_j^{z_j}}$ for each $j \in N$, is associated to a $\sigma^{(\theta_0, z, s)} = (\theta_0, \theta^z, s)$, and a profile of dominant types $\bar{z}(\sigma^{(\theta_0,z,s)})$, we define, for each $j \in N$, a perturbed type $z_j^\rho(z_j, s_j) \in Z_j^\rho$ by:

$$\hat{\theta}_j\left(z_j^\rho(z_j, s_j)\right) = \rho\hat{\theta}_j\left(\bar{z}_j\left(\sigma^{(\theta_0,z,s)}\right)\right) + (1-\rho)\hat{\theta}_j(z_j) \tag{4.6}$$

$$\beta_j\left(z_j^\rho(z_j, s_j)\right) = \rho \cdot Uniform\left(\theta_0^s\left(A_{(\theta_j^{z_j}, s_j)}\right) \times \bar{Z}_{-j,(\theta_j^{z_j}, s_j)}\right) + (1-\rho) \cdot \hat{\mu}_j^{s_j}[\emptyset] \circ (\alpha_j^\rho)^{-1} \tag{4.7}$$

where $\alpha_j^\rho : \Theta_0 \times Z_{-j} \times S_{-j} \to \Theta_0 \times Z_{-j}^\rho$ is given by $\alpha_j^\rho(\theta_0, z_{-j}, s_{-j}) = (\theta_0, z_{-j}^\rho(z_{-j}, s_{-j}))$, with $z_{-j}^\rho(z_{-j}, s_{-j}) = (z_l^\rho(z_l, s_l))_{l \neq j}$. Note that by construction, with probability ρ, type $z_j^\rho(z_j, s_j)$ is certain that s_j is conditionally dominant and puts positive probability on all opponents' dominant types in $\bar{Z}_{-j}(\theta_j^{z_j}, s_j))$.

In what follows we show that for type z_i of player i, with $\hat{\theta}_i(z_i) = \theta_i^*$ and $s_i^* \in D_{i,\theta_i^*}$, $z_i^\rho(z_i, s_i^*) \to z_i$ as $\rho \to 0$ and $ISR_i^\infty(z_i^\rho(z_i, s_i^*)) \subset D_{i,\theta_i^*}$ $\forall \rho > 0$, proving that z_i is a Γ_M-critical type.

For any $\rho > 0$, let player i of type $z_i^\rho(z_i, s_i^*)$, and consider a conjecture $\hat{\mu}_i \in \Phi_i(z_i^\rho(z_i, s_i^*))$ given by: $\hat{\mu}_i[\emptyset] = \beta_i(z_i^\rho(z_i, s_i^*)) \circ (\hat{\alpha}_i^{-1}) \in \Delta(\Theta_0 \times Z_{-i}^\rho \times S_{-i})$,

where, $\hat{\alpha}_i : \Theta_0 \times Z_{-i}^\rho \to \Theta_0 \times Z_{-i}^\rho \times S_{-i}$, is defined by: if $z_{-i}^\rho(z_{-i}, s_{-i}) \in Z_{-i}^\rho \setminus \bar{Z}_{-i}$, then $\hat{\alpha}_i(\theta_0, z_{-i}^\rho(z_{-i}, s_{-i})) = (\theta_0, z_{-i}^\rho(z_{-i}, s_{-i}), s_{-i})$; and if $\bar{z}_{-i}(\sigma) \in \bar{Z}_{-i} \subset Z_{-i}^\rho$ with $\sigma = (\sigma_i, (\theta_{-i}, s_{-i}))$, then $\hat{\alpha}_i(\theta_0, \bar{z}_{-i}(\sigma)) = (\theta_0, \bar{z}_{-i}(\sigma), s_{-i})$.

By construction, and the definition of $\hat{\alpha}_i$, we have that $\bar{Z}_{-i,(\theta_i^{z_i}, s_i^*)} \subset \text{supp}(\text{marg}_{Z_{-i}^\rho} \hat{\mu}_i[\emptyset])$, and that $S_{-i} = \text{supp}(\text{marg}_{S_{-i}} \hat{\mu}_i[\emptyset])$, so that every history h is possible under conjecture $\hat{\mu}_i[\emptyset]$ at the beginning of the game, so that, for every h, $\hat{\mu}_i[h]$ can be obtained via Bayes rule.

Given conjecture $\hat{\mu}_i$ for type $z_i^\rho(z_i, s_i^*)$, we check that s_i^* is a ϵ^*-sequential best response. At each $h \in \mathcal{H}(s_i^*)$ the conditional conjecture $\hat{\mu}_i[h]$ is a mixture: with high enough probability, it coincides with the conditional conjecture $\hat{\mu}_i^{s_i^*}[h]$ which make s_i^* an ϵ-best response at that history; and with a small enough probability it is concentrated on payoff states in which type $z_i^\rho(z_i, s_i^*)$ is certain that s_i^* is conditionally dominant which, given its payoff type $\hat{\theta}_i(z_i^\rho(z_i, s_i^*))$ and high enough $\epsilon' > 0$, breaks all ties in favor of s_i^*, even as the unique ϵ'-best reply at that history. Let $\epsilon^* = \min\{\epsilon, \epsilon'\} > 0$, and we obtain that $\{s_i^*\} = br_{i,z_i^\rho(z_i,s_i^*)}(\hat{\mu}_i|\epsilon^*)$. Note that the same is true for every opponent perturbed type $z_j^\rho(z_j, s_j)$, i.e. $\{s_j\} = br_{j,z_j^\rho(z_j,s_j)}(\hat{\mu}_j|\epsilon^*)$.

Hence, any consistent initial valid conjecture $\bar{\mu}_i \in \Phi_i(z_i^\rho(z_i, s_i^*))$ is such that $\bar{\mu}_i[\emptyset](\theta_0, z_{-i}^\rho(z_{-i}, s_{-i}), s'_{-i}) = 0$ for $s_{-i} \neq s'_{-i}$. Therefore, it follows that $s_i^* \in ISR_i^\infty(z_i^\rho(z_i, s_i^*); \epsilon^*) \subset D_{i,\theta_i^*}$.

Convergence: $z_i^\rho(z_i, s_i^*) \to z_i$ as $\rho \downarrow 0$, follows standard arguments as in Penta (2012).

Proof of Proposition 1. **[ii]:** Since, by hypothesis, $z_i \in Z_i$, in a finite model $\mathbf{Z}$, with $\hat{\theta}_i(z_i) = \theta_i^*$ is a Γ_M-critical type, let $\epsilon > 0$, $s_i^* \in S_i$ and $\{z_i^m\}_m$ with $z_i^m \to z_i$ be such that $s_i^* \in ISR_i^\infty(z_i)$ and $s_i^* \notin ISR_i^\infty(z_i^m; \epsilon)$ $\forall m$. We construct a collection $(D_{i,z_i})_{i \in N, z_i \in Z_i}$ with the desire properties.

First, consider player i of type z_i, with payoff-type $\hat{\theta}_i(z_i) = \theta_i^* \in \Theta_i$, and let

$$D_{i,z_i} = \left\{ s_i \in S_i;\ \exists (s_i^k)_k \subset S_i, \{z_i^{m_k}\}_k \subseteq \{z_i^m\}_m \text{ sth. } s_i^k \to s_i \text{ and } s_i^k \in ISR_i^\infty(z_i^{m_k}; \epsilon)\ \forall k \right\} \tag{4.8}$$

Clearly, $s_i^* \notin D_{i,z_i}$. For types $z'_i \neq z_i$ with payoff-types $\hat{\theta}_i(z'_i) = \theta'_i \neq \theta_i^*$, let $D_{i,z'_i} = ISR_i^\infty(z'_i; \epsilon)$. And now for player $j \neq i$, of type z_j with payoff-type $\hat{\theta}_j(z_j) = \theta_j \in \Theta_j$, let

$D_{j,z_j} = ISR_j^{\infty}(z_j;\epsilon)$.

We want to show that the collection $D \subset Z \times S$ is an ϵ-*EFBRS* such that $s_i^* \in ISR_i^{\infty}(z_i) \setminus D_{i,z_i}$.

First, consider $(z_i, s_i) \in D_i$ such that $s_i \in D_{i,z_i}$. Hence, there is $(s_i^k)_k \subset S_i$ with $s_i^k \to s_i$ such that $s_i^k \in ISR_i^{\infty}(z_i^{m_k};\epsilon)\ \forall k$. Let $\hat{\mu}_i^{m_k} \in \Phi_i(z_i^{m_k})$ for which,

i) $s_i^k \in br_{i,z_i^{m_k}}(\hat{\mu}_i^{m_k}|\epsilon)$

ii) $\text{supp}(\hat{\mu}_i^{m_k}[\emptyset]) \subseteq \Theta_0 \times ISR_{-i}^{\infty}(\epsilon)$

Now, let $\hat{\mu}_{(z_i,s_i)} = \lim_{k\to\infty} \hat{\mu}_i^{m_k} \in \Delta^{\mathcal{H}}(\Theta_0 \times Z_{-i} \times S_{-i})$, and we have,

i′) $\hat{\mu}_{(z_i,s_i)} \in \Phi_i(z_i)$, since $z_i^{m_k} \to z_i$, and $\Phi_i(\cdot)$ has a closed graph.

ii′) $s_i \in br_{i,z_i}\left(\hat{\mu}_{(z_i,s_i)}|\epsilon\right)$, by the *uhc* of $br_{i,z_i}(\cdot|\epsilon)$,

iii′) $\hat{\mu}_{(z_i,s_i)}[\emptyset](D_{-i}) = 1$ by *ii)* above.

i′), *ii′)*, and *iii′)* confirm (1) and (2) from Definition (8). It remains to show (3) from Definition (8), that is, $br_{i,z_i}\left(\hat{\mu}_{(z_i,s_i)}|\epsilon\right) \subset D_{i,z_i}$. So, let $(z_i, \hat{s}_i) \in br_i\left(\hat{\mu}_{(z_i,s_i)}|\epsilon\right)$, with $\hat{s}_i \neq s_i$. We want to show that $\hat{s}_i \in D_{i,z_i}$.

For this, consider the following definition: For fix $(\theta_0, (z_j)_{j\neq i}) \in \Theta_0 \times Z_{-i}$, and $\gamma_i \in S_i$, with $\gamma_i \neq \hat{s}_i$, there exists a profile $(\gamma_j)_{j\neq i}$ with $\gamma_j \in D_{j,z_j}$ for every $j \neq i$, such that

$$U_i((z_i,\hat{s}_i),(\theta_0,(z_j,\gamma_j)_{j\neq i})) - U_i((z_i,\gamma_i),(\theta_0,(z_j,\gamma_j)_{j\neq i})) \geq U_i((z_i,\hat{s}_i),(\theta_0,(z_j,\gamma'_j)_{j\neq i})) - U_i((z_i,\gamma_i),(\theta_0,(z_j,\gamma'_j)_{j\neq i})) \tag{4.9}$$

for all $(\gamma'_j)_{j\neq i}$ with $\gamma'_j \in D_{j,z_j}\ \forall j \neq i$. Let $\bar{\gamma}(\gamma_i, (\theta_0, (z_j)_{j\neq i})) = \{(\gamma_j)_{j\neq i};\ (4.9) \text{ holds}\}$. In words, the above defines a mapping $\bar{\gamma}(\cdot)$, such that for each $(\theta_0, (z_j)_{j\neq i}) \in \Theta_0 \times Z_{-i}$, and $\gamma_i \in S_i$, with $\gamma_i \neq \hat{s}_i$, a (possible, a set of) profile(s) $\bar{\gamma}(\gamma_i, (\theta_0, (z_j)_{j\neq i}))$ of strategies from D_{j,z_j} for each $j \neq i$, which favors $\hat{s}_i$ over γ_i the most.

We are now ready to prove that $\hat{s}_i \in D_{i,z_i}$. For each $k \in \mathbb{N}$, let $\hat{\mu}_i^k \in \Phi_i(z_i^{m_k})$ be such that,

1) At the beginning of the game, that is at $[\emptyset]$, for every $(\theta_0, z_{-i}) \in \text{supp}(\beta_i(z_i^{m_k}))$,

$$\hat{\mu}_i^k[\emptyset]\left(\{(\theta_0, z_{-i}, s_{-i})\};\ s_{-i} \in \bar{\gamma}(\gamma_i, (\theta_0, z_{-i})) \text{ for some } \gamma_i \in S_i, \gamma_i \neq \hat{s}_i\right) \geq 0 \tag{4.10}$$

2) $\left(\hat{\mu}_i^k\right)_k$ is such that $\hat{\mu}_i^k \to \hat{\mu}_{(z_i,s_i)}$ as $k \to \infty$.

Note that for every k, $\text{supp}(\hat{\mu}_i^k[\emptyset]) \subseteq \Theta_0 \times ISR_{-i}^{\infty}(\epsilon)$

Claim: $\exists\ \bar{K} \in \mathbb{N}$ large enough, such that, $\hat{s}_i \in ISR_i^{\infty}(z_i^{m_k}|\epsilon)\ \forall k \geq \bar{K}$.

For an arbitrary k, let $\hat{\mu}_i^k$ as described above. Given $\hat{\mu}_i^k$ and $\hat{s}_i$, there is a decomposition of $\mathcal{H}(\hat{s}_i)$ as in (4.3).

For any $\emptyset \leq l \leq L$, let $h \in \mathcal{H}(\hat{s}_i)$ be such that $h \in A_{[l]}(\hat{s}_i, \hat{\mu}_i^k)$, and there is no other $h' \in A_{[l]}(\hat{s}_i, \hat{\mu}_i^k)$ with $h' \geq h$. By the **RC**, $\hat{\mu}_i^k[h]$ can be such that, for every $s' \in S_i(h)$

$$\sum_{(\theta_0, z_{-i}, s_{-i})} \left(U_i((\theta_i^*, \hat{s}_i), (\theta_0, z_{-i}, s_{-i})) - U_i((\theta_i^*, s_i'), (\theta_0, z_{-i}, s_{-i}))\right) \, d\hat{\mu}_i^k[h](\{(\theta_0, z_{-i}, s_{-i})\}) \geq -\epsilon \quad (4.11)$$

without violating the fact that $\hat{\mu}_i^k \in \Phi_i(z_i^{m_k})$. For those nodes $h' \in A_{[l]}(\hat{s}_i, \hat{\mu}_i^k)$ with $h \geq h'$, it either follows by Bayes rule, or a similar argument as above, that an inequality as (4.11) holds, without violating the fact that $\hat{\mu}_i^k \in \Phi_i(z_i^{m_k})$. Since $h \in A_{[l]}(\hat{s}_i, \hat{\mu}_i^k)$ was arbitrary, this checks the condition of $\hat{s}_i$ being an ϵ-best reply at those histories from $\mathcal{H}(\hat{s}_i)$ that happen to be in $A_{[l]}(\hat{s}_i, \hat{\mu}_i^k)$, for every $\emptyset \leq l \leq L$.

Consider now a history $h \in \mathcal{H}(\hat{s}_i) \setminus \left(\bigcup_{l=\emptyset}^L A_{[l]}\right)$, and let $\gamma_i \in S_i(h)$[1]. Define the expected gain of $\hat{s}_i$ for type $z_i^{m_k}$ with respect to γ_i, as

$$EG_{\hat{s}_i}(\gamma_i; z_i^{m_k}) := \sum_{(\theta_0, z_{-i}, s_{-i})} \left(U_i\left[(z_i, \hat{s}_i), (\theta_0, z_{-i}, s_{-i})\right] - U_i\left[(z_i, \gamma_i), (\theta_0, z_{-i}, s_{-i})\right]\right) \, \hat{\mu}_i^k[h](\{(\theta_0, z_{-i}, s_{-i})\}) \tag{4.12}$$

and note that $EG_{\hat{s}_i}(\gamma_i; z_i^{m_k})$ equals

$$\sum_{(\theta_0, z_{-i})} \left(U_i\left[(z_i, \hat{s}_i), (\theta_0, z_{-i}, \bar{\gamma}(\gamma_i, (\theta_0, z_{-i})))\right] - U_i\left[(z_i, \gamma_i), (\theta_0, z_{-i}, \bar{\gamma}(\gamma_i, (\theta_0, z_{-i})))\right]\right) \, \hat{\mu}_i^k[h](\{(\theta_0, z_{-i})\}) \tag{4.13}$$

which follows from (4.10) and the fact that for all $h \in \mathcal{H}(\hat{s}_i) \setminus \left(\bigcup_{l=\emptyset}^L A_{[l]}\right)$, $\hat{\mu}_i^k[h]$ is obtained via Bayes rule.

Since $\hat{s}_i \in br_{i,z_i}(\hat{\mu}_{(z_i,s_i)}|\epsilon)$, the fact that $\hat{\mu}_{(z_i,s_i)}[\emptyset](D_{-i}) = 1$, and by 2) of the definition of $(\hat{\mu}_i^k)_k$, we have that for such an $h \in \mathcal{H}(\hat{s}_i) \setminus \left(\bigcup_{l=\emptyset}^L A_{[l]}\right)$, there exists $K(h, \gamma_i) \in \mathbb{N}$ such that for every $k \geq K(h, \gamma_i)$, we have

$$EG_{\hat{s}_i}(\gamma_i, z_i^{m_k}) \geq \sum_{(\theta_0, z_{-i}, s_{-i})} \left(U_i((z_i, \hat{s}_i), (\theta_0, z_{-i}, s_{-i})) - U_i((z_i, \gamma_i), (\theta_0, z_{-i}, s_{-i}))\right) \hat{\mu}_{(z_i,s_i)}[h](\theta_0, z_{-i}, s_{-i}) \geq -\epsilon \tag{4.14}$$

Let $K(h) = \max_{\gamma_i \in S_i(h)} K(h, \gamma_i)$. Then, for every $k \geq K(h)$ we have that $\hat{s}_i$ is an ϵ-best reply at h for type $z_i^{m_k}$.

Therefore, for each $h \in \mathcal{H}(\hat{s}_i) \setminus \left(\bigcup_{l=\emptyset}^L A_{[l]}\right)$, let $K(h) \in \mathbb{N}$ as before. For $h \in A_{[l]}$ for some $1 \leq l \leq L$, let $K(h) = 1$. Since the game is finite, let $\bar{K} = \max_{h \in \mathcal{H}(\hat{s}_i)} K(h)$. It follows that for every $k \geq \bar{K}$, $\hat{s}_i$ is an ϵ-sequentially best reply for type $z_i^{m_k}$. Hence $\hat{s}_i \in ISR_i^\infty(z_i^{m_k}|\epsilon)$ for every $k \geq \bar{K}$, and we conclude that $\hat{s}_i \in D_{i,z_i}$, as we wanted to show.

By construction, it is clear that for players $j \neq i$, the collections D_j satisfies conditions (1)-(3) of definition (8). And the proof is complete. $\square$

[1] Note that such a history could possible be the root of the game tree, if it is the case that $\emptyset \in \mathcal{H}(\hat{s}_i)$

Proof of Proposition 2.

Let $\mathbf{Z}$ be a finite model, and consider $z_i \in Z_{i,(+),u}$ with $\hat{\theta}_i(z_i) = \theta_i^*$ and $\beta_i[z_i] \in \Delta(\Theta_0 \times Z_{-i})$ such that $\beta_i[z_i](\Theta_0 \times Z_{-i,u}) = 1$. Note that for every $(\theta_0, z_{-i}) \in \text{supp}(\beta_i(z_i))$ with $\hat{\theta}_{-i}(z_{-i}) = \theta_{-i}^{z_{-i}} = (\theta_j^{z_j})_{j \neq i}$, the profile z_{-i} is not critical. Hence, by Proposition 1, for every $\epsilon > 0$, and every collection D of ϵ-*EFBRS* for the model $\mathbf{Z}$, we have that $\{s_{-i}\} = ISR_{-i}^{\infty}(z_{-i}) \subseteq D_{-i,z_{-i}}$. Since the game is finite, there is a collection $(D'_{i,z_i})_{i \in N, z_i \in Z_i}$ of ϵ-*EFBRS* with $\epsilon > 0$ for which D'_{i,z_i} is minimal with respect to set inclusion over subsets of S_i, and satisfies $D'_{i,z_i} \subseteq ISR_i^{\infty}(z_i)$.

In order to show that $z_i \in C_i^{\Gamma_M}$ we use Proposition 1 by showing that for the ϵ-*EFBRS* with $\epsilon > 0$ defined above, there are $s_i, s'_i \in S_i$, such that $s'_i \in D'_{i,z_i}$ (so that, $s'_i \in ISR_i^{\infty}(z_i)$), $s_i \in ISR_i^{\infty}(z_i)$, and $s_i \notin D'_{i,z_i}$.

Let $\hat{\mu}_i \in \Phi_i(z_i)$ be such that $\hat{\mu}_i[\emptyset] = \mu_i^{z_i} \in \Delta(\Theta_0 \times ISR_{-i}^{\infty})$. Now, since $\{\mu^{z_i}\} = \mathcal{B}_i^{[\emptyset]}\left[z_i; ISR_{-i}^{\infty}\right] \in \hat{O}$, let w^* be a path and a node $h^* \in w^*$ for which $(\theta_0, z_{-i}, s_{-i}) \notin \Theta_0 \times Z_{-i} \times S_{-i}(h^*)$ for all $(\theta_0, z_{-i}, s_{-i}) \in \text{supp}(\mu^{z_i})$. Consider strategy $s_i \in S_i$ such that $s_i, s'_i \in S_i(h^*)$, and $s_i(h^*) \neq s'_i(h^*)$, but $s_i(h) = s'_i(h)$ at any other h. For example, let $s_i(h^*) \in A_i(h^*)$ be such that the expected payoff of some opponent $j \neq i$ of type z_j when following strategy $s'_j \in D'_{j,z_j}$ is the lowest when player i of type z_i chooses $s_i(h^*)$.

Now, since there is a path w^* with a node $h^* \in w^*$ for which $\hat{\mu}_i[\emptyset]$ assign probability zero, we have that $\hat{\mu}_i[h^*]$ is unrestricted for $\hat{\mu}_i \in \Phi_i(z_i)$. By the **RC**, let $\hat{\mu}_i[h^*] \in \Delta(\Theta_0 \times Z_{-i} \times S_{-i}(h^*))$ be such that $\text{proj}_{\Theta_0 \times Z_{-i}} \text{supp}(\hat{\mu}_i[h^*]) \subseteq \theta_0^s\left(A^{(\theta_i^*, s_i)}\right) \times \bar{Z}_{-i}^{(\theta_i^*, s_i)}$, so that s_i is conditionally dominant at every element in the support of his beliefs, and therefore the unique sequentially rational strategy. Therefore, $s_i \in ISR_i^{\infty}(z_i)$, but $s_i \notin D'_{i,z_i}$. $\square$

Proof of Theorem 1.

Given our discussion before the presentation of Theorem 1, what we actually prove is that the set $\hat{Z}_{i,(+),u}$ is comeager in $Z_{i,(+)}$. Note that the following holds:

$$\hat{Z}_{i,(+),u} = \bigcap_{n \in \mathbb{N}} \left\{ z_i \in Z^n_{i,(+),u};\ \mathcal{B}_i^{[\emptyset]}\left[z_i; ISR_{-i}^{\infty}\right] \subset \hat{O}(n^{-1}) \right\} \tag{4.15}$$

For each $n \in \mathbb{N}$, the sets $\hat{O}(n^{-1})$ are open. By the uhc of the correspondence $\mathcal{B}_i^{[\emptyset]}\left[\cdot; ISR_{-i}^{\infty}\right]$: $Z_i \to \Delta(\Theta_0 \times Z_{-i} \times S_{-i})$, each of the sets of the intersection are open in $Z_{i,(+)}$. It follows that $\hat{Z}_{i,(+),u}$ is comeager in $Z_{i,(+)}$ whenever it contains the intersection of a countable collection of dense open sets. The next step is to show that each of the sets in the intersection is dense in $Z_{i,(+)}$.

Let $z_i \in Z_{i,(+)}$, with $\hat{\theta}_i(z_i) = \theta_i^*$, and $\beta_i[z_i] \in \Delta(\Theta_0 \times Z_{-i})$ be arbitrary.

Step 1. The initial step is to show that for a type $z_i \in Z_{i,(+)}$ there is at least one initial conjecture for a consistent conditional belief vector that justifies more than one strategies.

Claim: $\exists\ C_i \subset ISR_i^\infty(z_i)$ with $|C_i| > 1$ such that $\mu_{C_i}^{z_i} \in \mathcal{B}_i^{[\emptyset]}\left[(z_i, s_i); ISR_{-i}^\infty\right]\ \forall s_i \in C_i$

Proof of the claim: Suppose not. Let $\Delta(z_i)$ denote the set of initial valid conjectures that are consistent for type z_i, that is,

$\Delta(z_i) = \Big\{\mu \in \Delta(\Theta_0 \times Z_{-i} \times S_{-i});\ \text{supp}(\mu) \subseteq \Theta_0 \times ISR_{-i}^\infty \text{ and } \text{marg}_{\Theta_0 \times Z_{-i}}\,\mu = \beta_i[z_i]\Big\}$

$\Delta(z_i)$ is connected. If not, it would not be path-connected, and there would be a separation I, J of disjoint, open sets such that $\Delta(z_i) = I \cup J$ and any path $f : [0,1] \to \Delta(z_i)$ would be completely contain in one of the two open sets. However, by convexity of $\Delta(z_i)$ every pair of points in $\Delta(z_i)$ can be joined by a path, a contradiction.

By assumption, it follows then that for each s_i we have that,

$$\begin{aligned}\mathcal{B}_i^{[\emptyset]}\left[(z_i, s_i); ISR_{-i}^\infty\right] = \{\mu \in \Delta(\Theta_0 \times Z_{-i} \times S_{-i});\ &\text{supp}(\mu) \subseteq \Theta_0 \times ISR_{-i}^\infty, \text{ and}\\ &\exists \hat{\mu} \in \Phi(z_i) \text{ sth. } \hat{\mu}[\emptyset] = \mu \text{ and } \{s_i\} = br_{i,z_i}(\hat{\mu})\}\end{aligned}$$

and the collection $\Big\{\mathcal{B}_i^{[\emptyset]}\left[(z_i, s_i); ISR_{-i}^\infty\right]\Big\}_{s_i}$ is an open[2], disjoint cover of $\Delta(z_i)$, possibly, with some of them empty. A contradiction.

Step 2.1. From the result of Step 1, let $C_i \subset ISR_i^\infty(z_i)$ with $|C_i| > 1$ and $\mu_{C_i}^{z_i} \in \mathcal{B}_i^{[\emptyset]}\left[(z_i, s_i); ISR_{-i}^\infty\right]$, $\forall s_i \in C_i$ such that $\mu_{C_i}^{z_i} \in \hat{O}$.

From the structure theorem, Penta (2012), we have that for each $(\theta_0, z_{-i}) \in \text{supp}(\beta_i[z_i])$ and $s_{-i} \in ISR_{-i}^\infty(z_{-i})$, there is a sequence profile $\{z_{-i}^m(\theta_0, z_{-i}, s_{-i})\}_{m\in\mathbb{N}} = \{(z_j^m(\theta_0, z_j, s_j))_{j\neq i}\}_{m\in\mathbb{N}}$ such that, for each $j \neq i$, $z_j^m(\theta_0, z_j, s_j) \to z_j$ and $\{s_j\} = ISR_j^\infty(z_j^m(\theta_0, z_j, s_j))$. Define a sequence of types for player i as follows: Let z_i^m, be such that $\hat{\theta}_i(z_i^m) = \theta_i^*$, and $\beta_i[z_i^m] \in \Delta(\Theta_0 \times Z_{-i})$, be given by,

$$\beta_i[z_i^m]\left(\theta_0, z_{-i}^m(\theta_0, z_{-i}, s_{-i})\right) = \mu_{C_i}^{z_i}(\theta_0, z_{-i}, s_{-i}) \tag{4.16}$$

Then, for each $m \in \mathbb{N}$, $C_i \subseteq ISR_i^\infty(z_i^m)$, $z_i^m \in \hat{Z}_{i,(+),u} \subset C_i^{\Gamma_M}$, where the last inclusion follows from Proposition (2). Since we have that $\beta_i[z_i^m] \to \beta_i[z_i]$, it follows that for each $k \geq 1$ the k-th order beliefs are converging, that is, $\hat{\beta}_i^k[z_i^m] \to \hat{\beta}_i^k[z_i]$ as $m \to \infty$, an therefore we have that $z_i^m \to z_i$ as $m \to \infty$. In addition, from (4.15), the sequence $\{z_i^m\}_m$ belongs to each of the sets in the intersection of the right hand side, so, given that the original type $z_i \in Z_{i,(+)}$ was arbitrary, each of those sets is dense in $Z_{i,(+)}$. And the result follows for this case.

Step 2.2. Suppose now that for every subset C_i as in Step 1, there is no $\mu_{C_i}^{z_i} \in \mathcal{B}_i^{[\emptyset]}\left[(z_i, s_i); ISR_{-i}^\infty\right]$, $\forall s_i \in C_i$ such that $\mu_{C_i}^{z_i} \in \hat{O}$.

Recall that for the arbitrary type $z_i \in Z_{i,(+)}$, with $\hat{\theta}_i(z_i) = \theta_i^*$, and $\beta_i[z_i] \in \Delta(\Theta_0 \times Z_{-i})$, we

[2]The uniqueness of the sequential best reply implies that the expected payoff of a given strategy s_i for type z_i, a bilinear map, satisfies inequalities strictly. Openess follows as the inverse of a linear map, by keeping fix expected payoff at s_i, of an open set in $\mathbb{R}$.

have $\hat{\beta}_i^*[z_i] = \left(\hat{\theta}_i(z_i), \hat{\beta}_i[z_i]\right) = \left(\hat{\theta}_i(z_i), \hat{\beta}_i^1[z_i], \hat{\beta}_i^2[z_i], ..., \hat{\beta}_i^k[z_i], ...\right) \in \Theta_i \times Z_i^*$.

Take some $\mu_{C_i}^{z_i}$ as before which, by assumption, $\mu_{C_i}^{z_i} \notin \hat{O}$, and is such that $\text{marg}_{\Theta_0 \times Z_{-i}} \mu_{C_i}^{z_i} = \beta_i[z_i]$. Since *ISR* is uhc in the universal type space, the set $\left\{z_i' \in Z_{i,(+)}; ISR_i^\infty(z_i') \cap C_i \neq \emptyset\right\}$ is a closed subset in $Z_{i,(+)}$.

Consider $z_i(1) \in \left\{z_i' \in Z_{i,(+)}; ISR_i^\infty(z_i') \cap C_i \neq \emptyset\right\}$ defined by $\hat{\theta}_i(z_i(1)) = \hat{\theta}_i(z_i) = \theta_i^*$ and $\hat{\beta}_i^1[z_i(1)] = \hat{\beta}_i^1[z_i]$, that is, for each $(\theta_0, \theta_{-i}) \in \Theta_0 \times \Theta_{-i}$,

$$\hat{\beta}_i^1[z_i(1)](\theta_0, \theta_{-i}) = \hat{\beta}_i^1[z_i](\theta_0, \theta_{-i}) = \sum_{s_{-i}} \mu_{C_i}^{z_i} \left\{(\theta_0', z_{-i}, s_{-i});\, \theta_0' = \theta_0, \text{ and } \hat{\theta}_{-i}(z_{-i}) = \theta_{-i}\right\} \tag{4.17}$$

By step 1, $\exists\, C_i^1 \subset ISR_i^\infty(z_i(1))$ with $|C_i^1| > 1$, and satisfying $\mu_{C_i^1}^{z_i(1)} \in \mathcal{B}_i^{[\emptyset]}\left[(z_i(1), s_i); ISR_{-i}^\infty\right]$ $\forall s_i \in C_i^1$ such that $\mu_{C_i^1}^{z_i(1)} \in \hat{O}(1)$. In what follows, we construct a sequence of types of player i, $\left(z_i^m(1)\right)_m$ with the following properties: $z_i^m(1) \in Z_{i,(+),u}$; $\mathcal{B}_i^{[\emptyset]}\left[z_i^m(1); ISR_{-i}^\infty\right] \subset \hat{O}(1)$; and $z_i^m(1) \to z_i(1)$ as $m \to \infty$.

The construction is similar to the one in step 2.1. From the structure theorem we have that for each $(\theta_0, z_{-i}) \in \text{supp}(\beta_i[z_i(1)])$ and $s_{-i} \in ISR_{-i}^\infty(z_{-i})$, there is a sequence profile $\{z_{-i}^m(\theta_0, z_{-i}, s_{-i})\}_{m \in \mathbb{N}} = \{(z_j^m(\theta_0, z_j, s_j))_{j \neq i}\}_{m \in \mathbb{N}}$ such that, for each $j \neq i$, $z_j^m(\theta_0, z_j, s_j) \to z_j$ and $\{s_j\} = ISR_j^\infty(z_j^m(\theta_0, z_j, s_j))$. Define a sequence of types for player i as follows: Let $z_i^m(1)$, be such that $\hat{\theta}_i(z_i^m(1)) = \theta_i^*$, and $\beta_i[z_i^m(1)] \in \Delta(\Theta_0 \times Z_{-i})$, be given by,

$$\beta_i[z_i^m(1)]\left(\theta_0, z_{-i}^m(\theta_0, z_{-i}, s_{-i})\right) = \mu_{C_i^1}^{z_i(1)}(\theta_0, z_{-i}, s_{-i}) \tag{4.18}$$

Then, for each $m \in \mathbb{N}$, $C_i^1 \subseteq ISR_i^\infty(z_i^m(1))$, $z_i^m(1) \in Z_{i,(+),u}$, and $\mathcal{B}_i^{[\emptyset]}\left[z_i^m(1); ISR_{-i}^\infty\right] \subset \hat{O}(1)$. Since we have that $\beta_i[z_i^m(1)] \to \beta_i[z_i(1)]$, it follows that for each $k \geq 1$ the k-th order beliefs are converging, that is, $\hat{\beta}_i^k[z_i^m(1)] \to \hat{\beta}_i^k[z_i(1)]$ as $m \to \infty$, an therefore we have that $z_i^m(1) \to z_i(1)$ as $m \to \infty$.

Now, for each $t \geq 2$, consider $z_i(t) \in \left\{z_i' \in Z_{i,(+)}; ISR_i^\infty(z_i') \cap C_i \neq \emptyset\right\}$ defined by $\hat{\theta}_i(z_i(t)) = \hat{\theta}_i(z_i) = \theta_i^*$ and $\hat{\beta}_i^k[z_i(t)] = \hat{\beta}_i^k[z_i]$, $\forall k \leq t$. By step 1, $\exists\, C_i^t \subset ISR_i^\infty(z_i(t))$ with $|C_i^t| > 1$, and satisfying $\mu_{C_i^t}^{z_i(t)} \in \mathcal{B}_i^{[\emptyset]}\left[(z_i(t), s_i); ISR_{-i}^\infty\right]$ $\forall s_i \in C_i^t$ such that $\mu_{C_i^t}^{z_i(t)} \in \hat{O}(t^{-1})$. By the similar procedure as the previous paragraph, with $t = 1$, we construct a sequence of types of player i, $\left(z_i^m(t)\right)_m$ with the following properties: $z_i^m(t) \in Z_{i,(+),u}$; $\mathcal{B}_i^{[\emptyset]}\left[z_i^m(t); ISR_{-i}^\infty\right] \subset \hat{O}(t^{-1})$; and $z_i^m(t) \to z_i(t)$ as $m \to \infty$.

Therefore, we have a collection of sequences $\left((z_i^m(t))_{m \in \mathbb{N}}\right)_{t \in \mathbb{N}}$, where for each $t \in \mathbb{N}$, the sequence $\left(z_i^m(t)\right)_{m \in \mathbb{N}}$ has the properties described above. Consider the diagonal sequence $\left(z_i^t(t)\right)_{t \in \mathbb{N}}$, for which should be clear that $z_i^t(t) \to z_i$ as $t \to \infty$. Since the collection of sets $\hat{O}(t^{-1})$ is decreasing in t[3], note that for each $t \in \mathbb{N}$, there is a subsequence $\left(z_i^{t_n}(t_n)\right)_{n\in\mathbb{N}} \subseteq \left(z_i^t(t)\right)_{t\in\mathbb{N}}$ such that $\left(z_i^{t_n}(t_n)\right)_{n\in\mathbb{N}} \subseteq \left\{z_i \in Z^n_{i,(+),u};\ \mathcal{B}_i^{[\emptyset]}\left[z_i; ISR^{\infty}_{-i}\right] \subset \hat{O}(t^{-1})\right\}$ and $z_i^{t_n}(t_n) \to z_i$ as $n \to \infty$; so that each of the sets on the intersection of the right hand side of (4.15) is dense. And the proof is complete.

□

[3]That is, for $t < t'$ we have $\hat{O}(t'^{-1}) \subseteq \hat{O}(t^{-1})$