The Context of the Game *

Amanda Friedenberg†   Martin Meier‡

First Draft: October 2007
This Draft: October 2014

Abstract

Here, we study games of incomplete information and argue that it is important to correctly specify the "context" within which hierarchies of beliefs lie. We consider a situation where the players understand more than the analyst: It is transparent to the players—but not to the analyst—that certain hierarchies of beliefs are precluded. In particular, the players' type structure can be viewed as a strict subset of the analyst's type structure. How does this affect a Bayesian equilibrium analysis? One natural conjecture is that this doesn't change the analysis—i.e., every equilibrium of the players' type structure can be associated with an equilibrium of the analyst's type structure. We show that this conjecture is wrong. Bayesian equilibrium may fail an Extension Property. This can occur even in the case where the game is finite and the analyst uses the so-called universal structure (to analyze the game)—and, even, if the associated Bayesian game has an equilibrium. We go on to explore specific situations in which the Extension Property is satisfied.

†Department of Economics, W.P. Carey School of Business, Arizona State University, amanda.friedenberg@asu.edu, http://www.public.asu.edu/ afrieden.

‡Institut für Höhere Studien, Vienna, meier@ihs.ac.at

1 Introduction

This paper introduces a novel robustness question for the analysis of incomplete information games. We focus on a situation where the analyst correctly specifies the exogenous parameters of the game and the players' hierarchies of beliefs, but misspecifies the context within which these hierarchies lie. We ask: Are the analyst's predictions robust to misspecifying the context of the game?

What is the Context of the Game? Suppose that Nature tosses a coin, whose realization is either $\bar{\theta}$ or $\underline{\theta}$. The realization of this toss results in distinct payoff functions. Each of two players, resp. Izzy (i) and Joe (j), face uncertainty about the realization of this coin toss. What choices should Izzy and Joe make here? Presumably, Izzy's choice will depend on her belief about the realization of the coin toss. But, presumably, Izzy's choice will also depend on what she thinks about Joe's belief about the realization of the coin toss. After all, Joe's belief (about the realization of the coin toss) should influence his action, too. And, Izzy is concerned not only with what matrix is being played, but also with what choice Joe is making within the matrix.

To analyze the situation, we add to the description of the game, so that it also reflects these hierarchies of beliefs. In particular, we append a type structure to the game. One such type structure is given in Figure 1.1. Here, there are two possible types of Izzy, viz. t_i and u_i, and one possible type of Joe, viz. t_j. Type t_i (resp. u_i) of Izzy assigns probability one to Nature choosing $\bar{\theta}$ (resp. $\underline{\theta}$) and Joe's type being t_j. Type t_j of Joe assigns probability $\frac{1}{2}$ to "Nature choosing $\bar{\theta}$ and Izzy being type t_i" and probability $\frac{1}{2}$ to "Nature choosing $\underline{\theta}$ and Izzy being type u_i." So, type t_j of Joe assigns probability $\frac{1}{2}$ to "Nature choosing $\bar{\theta}$ and Izzy assigning probability one to $\bar{\theta}$" and probability $\frac{1}{2}$ to "Nature choosing $\underline{\theta}$ and Izzy assigning probability one to $\underline{\theta}$." And so on.

$\beta_i(\cdot)$	$(\bar{\theta}, t_j)$	$(\underline{\theta}, t_j)$
t_i	1	0
u_i	0	1

$\beta_j(\cdot)$	$(\bar{\theta}, t_i)$	$(\underline{\theta}, t_i)$	$(\bar{\theta}, u_i)$	$(\underline{\theta}, u_i)$
t_j	$\frac{1}{2}$	0	0	$\frac{1}{2}$

Figure 1.1: Type Structure

This type structure describes a situation where there are only two possible hierarchies of beliefs that Izzy can hold and only one possible hierarchy of beliefs that Joe can hold. In particular, it does not induce all hierarchies of beliefs. What is the rationale for limiting the type structure in this way? We view the specified game as only one part of the picture—a small piece of a larger story. The game sits within a broader strategic situation. That is, there is a history to the game, and this history influences the players. As Brandenburger, Friedenberg and Keisler (2008, p. 319) put it:

> We think of a particular ... structure as giving the "context" in which the game is played. In line with Savage's Small-Worlds idea in decision theory, who the players are in the given game can be seen as a shorthand for their experiences before the game. The players' possible characteristics—including their possible types—then reflect the prior history or context.

Under this view, the type structure, taken as a whole, reflects the context of the game. (Section 6a expands on this point and discusses the relationship to other views of game theory.)

Misspecifying the Context of the Game Consider the following scenario: The analyst looks at the strategic situation and the history. Perhaps the analyst even deduces that certain hierarchies are inconsistent with the history. But, to the players, it is transparent that other— that is, even more—hierarchies are inconsistent with the history. Put differently, players rule out hierarchies the analyst hasn't ruled out.

$\beta_i(\cdot)$	$(\bar{\theta}, t_j)$	$(\underline{\theta}, t_j)$	$(\bar{\theta}, u_j)$	$(\underline{\theta}, u_j)$		$\beta_j(\cdot)$	$(\bar{\theta}, t_i)$	$(\underline{\theta}, t_i)$	$(\bar{\theta}, u_i)$	$(\underline{\theta}, u_i)$
t_i	1	0				t_j	$\frac{1}{2}$	0	0	$\frac{1}{2}$
u_i	0	1				u_j				

Figure 1.2: The Analyst's Type Structure

Return to the earlier example and suppose the players' type structure is as given in Figure 1.1. Suppose the analyst misspecifies the type structure and instead studies the structure in Figure 1.2. But, it contains one extra type of Joe, viz. u_j. Type u_j is associated with some belief, distinct from type t_j's belief. The particular belief is immaterial. What is important is that each of Izzy's types assigns zero probability to this type of Joe. More to the point, each of Izzy's types is associated with the exact same beliefs as in the players' type structure. So, the players' type structure can be viewed as a subset (or substructure) of the analyst's type structure.

How does this affect an analysis? Take the solution concept of Bayesian Equilibrium applied to a Bayesian game associated with the type structure in Figure 1.2. For a given Bayesian Equilibrium, the analyst will have a prediction associated with the type u_j—i.e., a type that the players have ruled out. But the analyst will also have a prediction for the types t_i, u_i, and t_j. These are types in the players' structure, namely Figure 1.1.

The question is: How does the analyst's predictions for these types relate to the predictions he would have, if he had analyzed the game using the players' type structure? Presumably, the analyst's predictions shouldn't change. After all, the beliefs associated with t_i, u_i, and t_j have not changed at all. So, we can associate any equilibrium of the players' actual type structure with an equilibrium of the analyst's type structure, and vice versa.

3

Implicit in the above is that Bayesian Equilibrium satisfies Extension and Pull-Back Properties: Fix a type structure, viz. $\mathcal{T}$, associated with type sets T_i and T_j. We will think of $\mathcal{T}$ as the players' type structure. Now, consider another type structure $\mathcal{T}^*$, associated with type sets T_i^* and T_j^*. Suppose there is a map $h_i : T_i \to T_i^*$ (resp. $h_j : T_j \to T_j^*$) so that each t_i and $h_i(t_i)$ (resp. t_j and $h_j(t_j)$) induce the same hierarchies of beliefs. We will think of $\mathcal{T}^*$ as the analyst's structure. Now, we can state the Extension and Pull-back Properties.

> **The Equilibrium Extension Problem (Preliminary Version).** Fix an equilibrium of $\mathcal{T}$. Does there exist an equilibrium of $\mathcal{T}^*$ so that each $h_i(t_i) \in T_i^*$ and each $h_j(t_j) \in T_j^*$ plays the same strategy as do t_i and t_j (under the original equilibrium of $\mathcal{T}$)?

> **The Equilibrium Pull-Back Problem (Preliminary Version).** Fix an equilibrium of $\mathcal{T}^*$. Does there exist an equilibrium of $\mathcal{T}$ so that each $t_i \in T_i$ and each $t_j \in T_j$ plays the same strategy as do $h_i(t_i)$ and $h_j(t_j)$ (under the original equilibrium of $\mathcal{T}^*$)?

Return to the question of whether the analyst can study the Bayesian game in Figure 1.2. The answer is yes, provided that the analyst won't lose any predictions and won't introduce any new predictions. The question of losing predictions is the Extension Problem. The question of introducing new predictions is the Pull-Back Problem.

What is Already Known? While the robustness question is new to this paper, examples and results in the literature appear to speak to the Extension and Pull-Back Problems—at least as we have formalized these ideas, thus far. We begin with two examples.

Example 1.1. [1] Suppose Nature chooses the single parameter from $\Theta = \{\theta\}$. Type structures $\mathcal{T}$ and $\mathcal{T}^*$, in Figures 1.3-1.4, describe Izzy's and Joe's hierarchies of beliefs about Θ. Observe that these two type structures induce exactly the same set of hierarchies of beliefs about Nature's choice from Θ: In each type structure, each type of each player assigns probability 1 to θ. In each each type structure, each type of each player assigns probability one to "the other player assigns probability 1 to θ." And so on.

$\beta_i(\cdot)$	(θ, t_j)	(θ, u_j)
t_i	1	0
u_i	0	1

$\beta_j(\cdot)$	(θ, t_i)	(θ, u_i)
t_j	1	0
u_j	0	1

Figure 1.3: Type Structure $\mathcal{T}$

---

[1] We thank Pierpaolo Battigalli for suggesting this example.

4

$\beta_i^*(\cdot)$	(θ, t_j^*)
t_i^*	1

$\beta_j^*(\cdot)$	(θ, t_i^*)
t_j^*	1

Figure 1.4: Type Structure $\mathcal{T}^*$

Let Nature's choice of the parameter θ result in the payoff matrix in Figure 1.5. While the type structures $\mathcal{T}$ and $\mathcal{T}^*$ induce the same set of hierarchies of beliefs, we will see that the Bayesian game associated with the type structure $\mathcal{T}$ has equilibrium predictions that cannot be induced by the Bayesian game associated with the type structure $\mathcal{T}^*$.

First focus on the Bayesian game associated with type structure $\mathcal{T}^*$. The Bayesian equilibria of this game correspond exactly to the Nash equilibria of the game in Figure 1.5, i.e., either types t_i^* and t_j^* play (Up, Left), play (Down, Right), or assign $\frac{1}{2} : \frac{1}{2}$ to Up : Down and Left : Right.

Joe

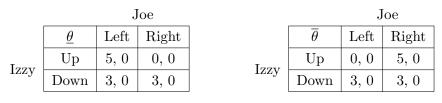|		Left	Right
Izzy	Up	1,1	0,0
	Down	0,0	1,1

Figure 1.5

Next, focus on the Bayesian game associated with type structure $\mathcal{T}$. There are Bayesian equilibria where the types of both players coordinate on a Nash equilibrium of the game in Figure 1.5, e.g., where both t_i, u_i play Up and t_j, u_j play Left. But there is also a Bayesian equilibrium where, say, (t_i, t_j) play (Up, Left) and (u_i, u_j) play (Down, Right).

Observe that this already points to a problem with the Equilibrium Extension Property, at least as we have defined it. For instance, take $\mathcal{T}$ to be the players' type structure and $\mathcal{T}^*$ to be the analyst's type structure. Then there are mappings $h_i(t_i) = h_i(u_i) = t_i^*$ and $h_j(t_j) = h_j(u_j) = t_j^*$ that preserve hierarchies of beliefs. Yet, there is an equilibrium of the players' type structure that cannot be extended to an equilibrium of the analyst's type structure: The types t_i and u_i in the players type structure are mapped to the same type t_i^* in the analyst's type structure. But, there is some equilibrium where these types, i.e., t_i and u_i, choose different actions. □

Example 1.2 (Battigalli and Siniscalchi, 2003; Ely and Peski, 2006; Dekel, Fudenberg and Morris, 2007; Liu, 2009)**.** Suppose Nature chooses a parameter from $\Theta = \{\underline{\theta}, \overline{\theta}\}$. This choice determines the players' payoff functions, as specified in Figure 1.6.

Type structures $\mathcal{T}$ and $\mathcal{T}^*$, in Figures 1.7-1.8, describe Izzy's and Joe's hierarchies of beliefs about Θ. Observe that these two type structures induce exactly the same set of hierarchies of beliefs about the parameter: In each type structure, each type of each player assigns $\frac{1}{2} : \frac{1}{2}$ to $\underline{\theta} : \overline{\theta}$. So, in each type structure, each type of each player assigns probability one to "the other player

5

	Joe				Joe	

<table>
<tr><td rowspan="3">Izzy</td><td align="center">Joe</td></tr>
</table>

	Joe	
$\underline{\theta}$	Left	Right
Up	5, 0	0, 0
Down	3, 0	3, 0

(Izzy)

	Joe	
$\overline{\theta}$	Left	Right
Up	0, 0	5, 0
Down	3, 0	3, 0

(Izzy)

Figure 1.6

assigns $\frac{1}{2} : \frac{1}{2}$ to $\underline{\theta} : \overline{\theta}$." And so on.

$\beta_i(\cdot)$	$(\underline{\theta}, t_j)$	$(\overline{\theta}, t_j)$	$(\underline{\theta}, u_j)$	$(\overline{\theta}, u_j)$
t_i	$\frac{1}{2}$	0	0	$\frac{1}{2}$

$\beta_j(\cdot)$	$(\underline{\theta}, t_i)$	$(\overline{\theta}, t_i)$
t_j	$\frac{1}{2}$	$\frac{1}{2}$
u_j	$\frac{1}{2}$	$\frac{1}{2}$

Figure 1.7: Type Structure $\mathcal{T}$

$\beta_i^*(\cdot)$	$(\underline{\theta}, t_j^*)$	$(\overline{\theta}, t_j^*)$
t_i^*	$\frac{1}{2}$	$\frac{1}{2}$

$\beta_j^*(\cdot)$	$(\underline{\theta}, t_i^*)$	$(\overline{\theta}, t_i^*)$
t_j^*	$\frac{1}{2}$	$\frac{1}{2}$

Figure 1.8: Type Structure $\mathcal{T}^*$

While the type structures $\mathcal{T}$ and $\mathcal{T}^*$ induce the same set of hierarchies of beliefs, the Bayesian game associated with the type structure $\mathcal{T}$ has equilibrium predictions that cannot be induced by the Bayesian game associated with the type structure $\mathcal{T}^*$. In any equilibrium of Bayesian game associated with $\mathcal{T}^*$, t_i^* plays Down. But, there exists a Bayesian equilibrium associated with the type structure $\mathcal{T}$ where t_i plays Up, t_j plays Left, and u_j plays Right.

This example also points to a problem with the Equilibrium Extension and Pull-Back Properties, as we have defined them. For instance, take $\mathcal{T}$ to be the players' type structure and $\mathcal{T}^*$ to be the analyst's type structure. Then there are mappings $h_i(t_i) = t_i^*$ and $h_j(t_j) = h_j(u_j) = t_j^*$ that preserve hierarchies of beliefs. Yet, there is an equilibrium of the players' type structure that cannot be extended to an equilibrium of the analyst's type structure. Likewise, if we take $\mathcal{T}$ to be the analyst' type structure and $\mathcal{T}^*$ to be the player's type structure, then there is an equilibrium of the analyst's type structure that cannot be pulled-back to an equilibrium of the players' type structure. □

Examples 1.1-1.2 illustrate failures of preliminary versions of the Extension and Pull-Back Properties. But, they do *not* address the robustness question we are interested in. To see this, begin with Example 1.1, where the true parameter θ is common belief amongst the players. We

described the players' type structure as $\mathcal{T}$, where two types of Izzy (resp. Joe) induce the same single hierarchy of beliefs about θ. These types specify information that is not available to the players when the only set of *ex ante* uncertainty is about the parameter $\Theta = \{\theta\}$. They provide additional information—not only about Θ—but also about the realization of some external signal. In effect, the name of these types, i.e., t_i vs. u_i (resp. t_j vs. u_j), specifies the information about the realization of these external signals. Indeed, the new equilibrium of the Bayesian game (associated with $\mathcal{T}$) can be obtained as an objective correlated equilibrium of the game matrix in Figure 1.5. (See Aumann, 1987.)

A similar idea is at play in Example 1.2. There, the type structure $\mathcal{T}$ provides information that is not available to Joe when the only set of *ex ante* uncertainty is hierarchies of beliefs about the parameter Θ. Two types of Joe, viz. t_j and u_j induce the same hierarchies of beliefs about Θ. The fact that Joe's action can vary with these types reflects the idea that Joe has obtained different information about the realization of some external signal.

In both examples, we have a type structure $\mathcal{T}$, where two types of a player induce the same hierarchies of beliefs about the parameter Θ. This type structure is **redundant**. Redundant structures can provide information above and beyond hierarchies of beliefs about the parameter Θ; they can also provide information about an external signal. (See Liu, 2009 for a formal statement.)

We are interested in the case where the analyst correctly specifies the parameters of the game (including the set of actual signals), correctly specifies players' hierarchies of beliefs about these parameters, but simply also considers possible that players may have 'additional hierarchies,' i.e., hierarchies ruled out by the players themselves. Examples 1.1-1.2 illustrate that, as stated, a failure of the Extension or Pull-Back Property need not reflect this robustness criterion—it may instead reflect a failure to correctly specify the parameter set (or signal space) in the game. Thus, we will need to amend the statement of the Extension and Pull-Back Properties to reflect our robustness question. One simple way to do so is by restricting Extension and Pull-Back to non-redundant type structures. Once we introduce the main formalism, we will see that we can, in fact, state the Extension and Pull-Back Problems somewhat more generally.

Are the Extension and Pull-Back Problems Satisfied? We will see that the Pull-Back Problem is indeed satisfied. This fact is 'in the air' so to speak. Thus the focus of this paper will be on the Extension Problem. It is easy to construct simple pathological examples where the Extension Property fails:

Example 1.3. Suppose Nature chooses a parameter from $\Theta = \{\underline{\theta}, \overline{\theta}\}$. In either case, players' actions are $0, 1, 2, 3, \ldots$. If the true parameter is $\overline{\theta}$, then Izzy (resp. Joe) obtains a payoff of 1 if she (resp. he) chooses 0 and obtains a payoff of 0 otherwise. If the true parameter is $\underline{\theta}$, then Izzy (resp. Joe) obtains a payoff of 100 if she (resp. he) chooses an action strictly higher than Joe's (resp. Izzy's) action; otherwise, she (resp. he) obtains a payoff of 0.

$\beta_i(\cdot)$	$(\underline{\theta}, t_j)$	$(\bar{\theta}, t_j)$		$\beta_j(\cdot)$	$(\underline{\theta}, t_i)$	$(\bar{\theta}, t_i)$
t_i	0	1		t_j	0	1

Figure 1.9: Players' Type Structure $\mathcal{T}$

Take the player's type structure to be $\mathcal{T}$ described in Figure 1.9. Here, there is a unique Bayesian equilibrium where the single type Izzy and Joe both play the action 0. Now, take the analyst's type structure to be $\mathcal{T}^*$ as in Figure 1.10. This adds a type for each player. This new Bayesian game does not have a equilibrium. Thus, we cannot extend a Bayesian equilibrium from the players' structure to a Bayesian equilibrium from the analyst's type structure. □

$\beta_i^*(\cdot)$	$(\underline{\theta}, t_j^*)$	$(\bar{\theta}, t_j^*)$	$(\underline{\theta}, u_j^*)$	$(\bar{\theta}, u_j^*)$		$\beta_j^*(\cdot)$	$(\underline{\theta}, t_i^*)$	$(\bar{\theta}, t_i^*)$	$(\underline{\theta}, u_i^*)$	$(\bar{\theta}, u_i^*)$
t_i^*	0	1	0	0		t_j^*	0	1	0	0
u_i^*	0	.1	.9	0		u_j^*	0	.1	.9	0

Figure 1.10: Type Structure $\mathcal{T}^*$

Example 1.3 illustrates that we may have a failure of the Equilibrium Extension Property. The reason for this failure is that the game of incomplete information is itself pathological. As a consequence, we have a situation where there is no Bayesian equilibrium of the analyst's game.

But the failure of Equilibrium Extension need *not* be an artifact of such pathologies. We will build an example of an Extension Failure from, arguably, "standard" ingredients—that is, ingredients which are well-understood and for which we would very much expect no problem to arise. Let us point to some features of the construction:

- The parameter set Θ is finite.

- The game Γ has a finite number of players and each player has a finite number of choices.

- For any associated type structure, there is an equilibrium of the associated Bayesian game. So, in particular, there will be an equilibrium of the analyst's Bayesian game.

- The players' type structure $\mathcal{T}$ has (at most) a countable number of types. There are no further restrictions on the structure—so, for instance, we can take it to arise from a common prior.

- The analyst's type structure $\mathcal{T}^*$ is the canonical construction of the universal type structure based on the (finite) parameter set Θ.

8

So we have a finite parameter set, a finite game, a finite or countable players' type structure, a universal analyst's structure, and existence in the analyst's Bayesian game—standard ingredients.

Along the way, we will construct an example of a second extension failure—one that satisfies the above requirements with one notable exception: it need not be the case that, for any type structure, there is an equilibrium of the associated Bayesian game. In particular, in this second construction, there will not be an equilibrium of the analyst's universal Bayesian game.[2] This alternate construction also implies a failure of Equilibrium Extension. But, it is *not* built from "standard" ingredients. In our main example, we have a failure of Extension, despite the fact that the players' Bayesian game and the analyst's universal Bayesian game *both have an equilibrium.* Indeed, precisely because the analyst's universal Bayesian game does have an equilibrium, the analyst may be misled into thinking that he has captured all possible predictions, when he has not. By contrast, if there is no equilibrium of the analyst's structure, he will presumably not be misled in this way.

The case of a universal type structure is of particular interest. It is often presumed that the analyst should necessarily take the universal structure to applications, even if the current state of applied work does not do so. See, e.g., Morris and Shin (2003) who say "optimal strategic behavior should be analyzed in the space of all possible infinite hierarchies of beliefs." The Extension failure tells us that—while perhaps appealing—such an general principle may, in fact, be problematic.

Positive Results   The negative results raise the question: Are there situations in which the analyst can be guaranteed that his analysis will not fail the Extension property? We provide two sets of conditions under which the answer is yes. First, we can extend any universally measurable equilibrium in compact continuous games, provided there are (at most) a countable number of types that are in the analyst's structure but not the players' structure. (See Definition 2.9 for the concept of universal measurability.) Second, we have an Extension property if the analyst's structure satisfies a common prior plus a positivity requirement. See Sections 5.1-5.2.

Going Forward   These positive results get at—but do not answer—an important question. To what extent do the Bayesian games studied in applications satisfy or fail Extension? The positive results tell us that, for certain applications, we do indeed satisfy Extension. But, they do not cover all applications. At the theoretical level, addressing this question requires answering a more fundamental question: Can we characterize the set of Bayesian games that satisfy or fail Extension? We don't know the answer and leave this as an open question.

Absent such a characterization, how can the analyst proceed (when the sufficient conditions do not obtain)? One idea is to modify the Bayesian equilibrium concept and use, instead, what Sadzik (2011) calls, Local Bayesian Equilibrium (LBE). Under an LBE analysis, the analyst does not stop at characterizing the set of Bayesian equilibria for a given Bayesian game. Instead, the

---

[2]This construction uses an important result due to Hellman (2014). However, it is not a Corollary of Hellman. It also makes use of Lemma 2.2 and the Pull-Back Property below. To the best of our knowledge, this is the first example of a finite game so that the associated universal Bayesian game does not have *any* Bayesian equilibrium.

analyst looks across all sub-Bayesian games and characterizes all equilibria across this class. In effect, LBE requires that the analyst analyze every possible players' type structure.

Sadzik (2011) introduced the LBE concept to 'get around' a different robustness question: robustness to misspecifying external signals. (Refer back to the discussion on pages 6-7.) One might have thought that, when the analyst can correctly specify the parameters of the game (including the set of external signals) but is uncertain of the context, studying the LBE concept—as opposed to the Bayesian equilibrium concept—is 'overkill.' However, our negative result shows that this is not the case: Even if the game of incomplete information is finite and there exists an equilibrium of every associated Bayesian game, using Bayesian equilibrium to analyze the universal type structure may be insufficient to capture predictions associated with every players' type structure.

The paper proceeds as follows. Section 2 sets up preliminaries. The Extension and Pull-Back Properties are formally defined in Section 3. There, we also show the Pull-Back result. Section 4 shows the negative results. Sections 5.1-5.2 provide positive results—conditions on the game and on the type structure which guarantee the Extension property. Finally, Section 6 concludes by discussing some conceptual and formal aspects of the paper.

2 Bayesian Games

Throughout the paper, we adopt the following conventions. We will endow the product of topological spaces with the product topology, and a subset of a topological space with the induced topology. Given a metrizable space Ω, endow Ω with the Borel sigma-algebra $\mathcal{B}(\Omega)$ unless otherwise stated. In this case, write $\Delta(\Omega)$ for the set of probability measures on Ω and endow $\Delta(\Omega)$ with the topology of weak convergence. If Ω is Polish, so is $\Delta(\Omega)$. Given some $\mu \in \Delta(\Omega)$, write $\mathcal{B}(\Omega; \mu)$ for the completion of the Borel sigma-algebra with respect to μ.

Given a finite index set I, write $\Omega = \prod_{i \in I} \Omega_i$, $\Omega_{-i} = \prod_{j \in I \setminus \{i\}} \Omega_j$, and $\Omega_{-i-j} = \prod_{k \in I \setminus \{i,j\}} \Omega_k$. Write ω (resp. ω_{-i}) for a typical element of Ω (resp. Ω_{-i}). Given maps $f_i : \Omega_i \to \Phi_i$, for each $i \in I$, write f for the product map from Ω to Φ, given by $f(\omega_1, \ldots, \omega_{|I|}) = (f_1(\omega_1), \ldots, f_{|I|}(\omega_{|I|}))$. Define f_{-i} analogously.

Fix measure spaces $(\Omega_1, \mathcal{S}(\Omega_1))$ and $(\Omega_2, \mathcal{S}(\Omega_2))$, where $\mathcal{S}(\Omega_1)$ and $\mathcal{S}(\Omega_2)$ are arbitrary sigma-algebras on Ω_1 and Ω_2. A function $f : \Omega_1 \to \Omega_2$ is $(\mathcal{S}(\Omega_1), \mathcal{S}(\Omega_2))$-**measurable** if, for each $E_2 \in \mathcal{S}(\Omega_2)$, $f^{-1}(E_2) \in \mathcal{S}(\Omega_1)$. A function $f : \Omega_1 \to \Omega_2$ is (**Borel**) **measurable** if it is $(\mathcal{B}(\Omega_1), \mathcal{B}(\Omega_2))$-measurable and μ-**measurable** if it is $(\mathcal{B}(\Omega_1; \mu), \mathcal{B}(\Omega_2))$-measurable.

Say f is μ-**integrable** if it is Lebesgue integrable with respect to the measure μ. A standard fact that we will make use of is that a bounded function $f : \Omega \to \mathbb{R}$ is μ-integrable if and only if it is μ-measurable. (See, e.g., Bogachev, 2006, pages 118, 121-122.)

Call $f : \Omega_1 \to \Omega_2$ **universally measurable** if it is μ-measurable for all $\mu \in \Delta(\Omega_1)$ or, equivalently, if it is $(\mathcal{B}_{\mathrm{UM}}(\Omega_1), \mathcal{B}(\Omega_2))$-measurable where $\mathcal{B}_{\mathrm{UM}}(\Omega_1) = \bigcap_{\mu \in \Delta(\Omega_1)} \mathcal{B}(\Omega_1; \mu)$. Sets in $\bigcap_{\mu \in \Delta(\Omega)} \mathcal{B}(\Omega; \mu)$ are called **universally measurable sets**. The set of universally measurable

sets contains the set of Borel sets.

We will make use of some facts about universally measurable functions: Any measurable function is universally measurable. A function $f : \Omega \to \Phi$ is universally measurable if and only if it is $(\mathcal{B}_{\mathrm{UM}}(\Omega), \mathcal{B}_{\mathrm{UM}}(\Phi))$-measurable. (See Fremlin, 2000, page 188.) A consequence is that the composition of two universally measurable functions is universally measurable. If Ω_1, Ω_2 are separable metrizable, $f_1 : \Omega_1 \to \Phi_1$ and $f_1 : \Omega_2 \to \Phi_2$ are universally measurable if and only if the product map $f : \Omega \to \Phi$ is universally measurable. (See Lemma A.1.)

Given a measurable mapping $f : \Omega \to \Phi$, write $\underline{f} : \Delta(\Omega) \to \Delta(\Phi)$ for the map that takes each measure $\mu \in \Delta(\Omega)$ to its image measure under f, i.e., $\underline{f}(\mu)(E) = \mu(f^{-1}(E))$ for each $E \in \mathcal{B}(\Phi)$.

Bayesian Games

Let Θ be a Polish set, to be interpreted as a **parameter set** or a set of **states of Nature**. Throughout, we fix a finite player set I and label players as $1, \ldots, |I|$. Write i for a particular player from I. A **Θ-based game** is then some $\Gamma = ((C_i, \pi_i)_{i \in I})$. Here, C_i is a **choice** or an **action** set for player i, which is taken to be Polish. A **payoff function** for player i is a bounded measurable map $\pi_i : \Theta \times C \to \mathbb{R}$. Extend π_i to $\Theta \times \prod_{j \in I} \Delta(C_j)$ in the usual way; the extended functions are again bounded and measurable. A special case will be of particular interest—namely, a **finite game**, i.e., a game where the parameter set Θ and each of the choice sets C_i are each finite.

To analyze the Θ-based game, we will need to append to the game a Θ-based type structure.

Definition 2.1. *A **Θ-based type structure** is some $\mathcal{T} = (\Theta, (T_i, \beta_i)_{i \in I})$, where each T_i is a (non-empty) Polish **type set** for player i and each β_i is a measurable **belief map** $\beta_i : T_i \to \Delta(\Theta \times T_{-i})$ for player i.*

Say $\mathcal{T} = (\Theta, (T_i, \beta_i)_{i \in I})$ is **countable** if each T_i is at most countable.

Definition 2.2. *A **Θ-based Bayesian game** consists of a pair $(\Gamma, \mathcal{T})$, where Γ is a Θ-based game and $\mathcal{T}$ is a Θ-based type structure.*

The Bayesian game induces strategies. A strategy for i, viz. s_i, is a map from T_i to $\Delta(C_i)$. Let S_i be the set of strategies for player i.

Bayesian Equilibrium

It will be convenient to introduce the following notation: Fix some $c_i \in C_i$ and write $\hat{\pi}_i[c_i] : \Theta \times \prod_{j \in I \setminus \{i\}} \Delta(C_j) \to \mathbb{R}$ for the mapping with $\hat{\pi}_i[c_i](\theta, \sigma_{-i}) = \pi_i(\theta, c_i, \sigma_{-i})$ where $\sigma_{-i} \in \prod_{j \in I \setminus \{i\}} \Delta(C_j)$. So $\hat{\pi}_i[c_i]$ specifies the payoff of playing action c_i, as a function of the payoff parameter and mixed actions of the other players. Each $\hat{\pi}_i[c_i]$ is a bounded measurable function. Given some $c_i \in C_i$ and a strategy profile $s_{-i} : T_{-i} \to \prod_{j \in I \setminus \{i\}} \Delta(C_j)$, define $\Pi_i[c_i, s_{-i}] : \Theta \times T_{-i} \to \mathbb{R}$ so that

$$\Pi_i[c_i, s_{-i}](\theta, t_{-i}) = \pi_i(\theta, c_i, s_1(t_1), \ldots, s_{i-1}(t_{i-1}), s_{i+1}(t_{i+1}), \ldots, s_{|I|}(t_{|I|})).$$

Note, $\Pi_i[c_i, s_{-i}] = \hat{\pi}_i[c_i] \circ (\text{id} \times s_{-i})$, where $\text{id} : \Theta \to \Theta$ denotes the identity map. So, $\Pi_i[c_i, s_{-i}]$ specifies the payoff of playing action c_i when the others play the strategy profile s_{-i}, as a function of the payoff parameter and types of the other players.

Definition 2.3. *Say* $(s_1, \ldots, s_{|I|})$ *is a **Bayesian equilibrium** if, for each* $i \in I$*, each* $t_i \in T_i$ *and each* $c_i \in C_i$*, the following hold:*

(i) $\Pi_i[c_i, s_{-i}]$ *is* $\beta_i(t_i)$*-integrable; and*

(ii) $\int_{\Theta \times T_{-i}} \Pi_i[s_i(t_i), s_{-i}] d\beta_i(t_i) \geq \int_{\Theta \times T_{-i}} \Pi_i[c_i, s_{-i}] d\beta_i(t_i).$

Condition (i) says that each type can compute her expected payoffs for each possible action c_i she may choose, given that all other players choose the equilibrium strategy. It is automatically satisfied for a Bayesian game $(\Gamma, \mathcal{T})$ with $\mathcal{T}$ countable. But, more generally, it must be stated explicitly. Condition (ii) requires that each type maximize its expected payoffs, given its associated belief.

Note, Condition (i) of Definition 2.3 is satisfied if and only if each $\Pi_i[c_i, s_{-i}]$ is $\beta_i(t_i)$-measurable. Since $\hat{\pi}_i[c_i] : \Theta \times \prod_{j \in I \setminus \{i\}} \Delta(C_j) \to \mathbb{R}$ is measurable, this will in turn be satisfied if the mapping

$$\text{id} \times s_{-i} : \Theta \times T_{-i} \to \Theta \times \prod_{j \in I \setminus \{i\}} \Delta(C_j)$$

given by $(\text{id} \times s_{-i})(\theta, t_{-i}) = (\theta, s_{-i}(t_{-i}))$ is $\beta_i(t_i)$-measurable.[3] In fact, in a particular class of games, $\beta_i(t_i)$-measurability of $\text{id} \times s_{-i}$ is also necessary for $\beta_i(t_i)$-integrability of $\Pi_i[c_i, s_{-i}]$.

Definition 2.4. *Call a* Θ*-based game* Γ ***injective*** *if, for each* i *and each* $c_i \in C_i$*,* $\hat{\pi}_i[c_i]$ *is injective.*

Note carefully that a Θ-based game may be injective even if some player's payoff function π_i is not injective. Many games of interest fail the injectivity condition. But, when the injectivity condition is met, we have the following characterization of a Bayesian equilibrium.[4]

Lemma 2.1. *Fix a Bayesian Game* $(\Gamma, \mathcal{T})$*, where* Γ *is injective. For each* $(c_i, s_{-i}) \in C_i \times S_{-i}$ *and each* $\mu_i \in \Delta(\Theta \times T_{-i})$*,* $\Pi_i[c_i, s_{-i}]$ *is* μ_i*-measurable if and only if* $(\text{id} \times s_{-i})$ *is* μ_i*-measurable.*

Corollary 2.1. *Fix a Bayesian Game* $(\Gamma, \mathcal{T})$*, where* Γ *is injective. Then,* $(s_1, \ldots, s_{|I|})$ *is a Bayesian equilibrium if and only if, for each* i*, each* $t_i \in T_i$ *and each* $c_i \in C_i$*, the following hold:*

(i) $(\text{id} \times s_{-i})$ *is* $\beta_i(t_i)$*-measurable; and*

(ii) $\int_{\Theta \times T_{-i}} \Pi_i[s_i(t_i), s_{-i}] d\beta_i(t_i) \geq \int_{\Theta \times T_{-i}} \Pi_i[c_i, s_{-i}] d\beta_i(t_i).$

---

[3]A consequence is that Condition (i) is automatically satisfied if each s_{-i} is measurable. That is, the requirement that the equilibrium strategy be measurable is sufficient for Condition (i). Thus, some papers replace Condition (i) with a measurability requirement. Restricting attention to measurable equilibrium suffices for a positive result, i.e., to establish existence or to characterize certain behavior as consistent with equilibrium. But, to establish a negative result, as we have here, it is important to rule out more than simply 'a sufficient condition cannot be satisfied.' We thank Jeff Ely pointing us to this distinction and thereby to push us toward establishing a significantly stronger result than in a previous version of the paper.

[4]Proofs not found in the main text can be found in the Appendices.

Universal Bayesian Games and Equilibria

We will want to consider a special case, where the analyst studies a Θ-based Bayesian game $(\Gamma, \mathcal{T}^*)$, where $\mathcal{T}^*$ induces all hierarchies of beliefs. Mertens and Zamir (1985), Brandenburger and Dekel (1993), and Heifetz and Samet (1998) each provide (different) canonical constructions of (Θ-based) type structures that contain all hierarchies of beliefs. Here, we will not need to make use of the details of a particular construction. Instead, we can focus on certain properties that each of these constructions satisfy. To state a key property, we will need to introduce some terminology.

Definition 2.5 (Mertens and Zamir, 1985). *Fix two Θ-based structures $\mathcal{T} = (\Theta, (T_i, \beta_i)_{i \in I})$ and $\mathcal{T}^* = (\Theta, (T_i^*, \beta_i^*)_{i \in I})$ and measurable maps $h_1, \ldots, h_{|I|}$, where each $h_i : T_i \to T_i^*$. Call $(h_1, \ldots, h_{|I|})$ a **type morphism (from $\mathcal{T}$ to $\mathcal{T}^*$)** if, for each i, $\underline{\mathrm{id} \times h_{-i}} \circ \beta_i = \beta_i^* \circ h_i$.*
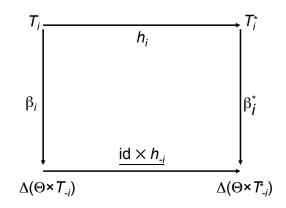

Figure 2.1

Definition 2.5 says that $(h_1, \ldots, h_{|I|})$ is a type morphism if it preserves the belief maps $\beta_1, \ldots, \beta_{|I|}$. Specifically, it requires that the diagram in Figure 2.1 commutes. Proposition 5.1 in Heifetz and Samet (1998) shows that each type morphism is a mapping that preserves hierarchies of beliefs, i.e., a hierarchy morphism.

Definition 2.6. *Fix a player set I and a parameter set Θ. Call a Θ-based type structure, viz. $\mathcal{T}^*$, **terminal** if, for each Θ-based structure $\mathcal{T}$, there is a type morphism from $\mathcal{T}$ to $\mathcal{T}^*$.*

Definition 2.7. *Fix a player set I and a parameter set Θ. Call a Θ-based type structure, viz. $\mathcal{T}^*$, **universal** if:*

(i) $\mathcal{T}^*$ it is terminal; and

*(ii) $\mathcal{T}^*$ is **non-redundant**, i.e., no two types induce the same hierarchies of beliefs.[5]*

---

[5] We will not need a formal definition of non-redundancy, since we will only make use of a consequence of the property.

Each of the canonical constructions of a "universal type structure" in Mertens and Zamir (1985), Brandenburger and Dekel (1993), Heifetz and Samet (1998) satisfy Definition 2.7. At times we will write $\mathcal{U}(\Theta)$ to indicate that the particular Θ-based type structure is universal.

Definition 2.8. *Call a Θ-based Bayesian game, viz. $(\Gamma, \mathcal{T})$, a **universal Bayesian game** if $\mathcal{T}$ is universal.*

Fix some injective game Γ. Bayesian equilibria of a universal Bayesian game $(\Gamma, \mathcal{T})$ will necessarily have a nice measurability property.

Definition 2.9. *Call a Bayesian equilibrium, viz. $(s_1, \ldots, s_{|I|})$, a **universally measurable equilibrium** if, for each i and $\mu \in \Delta(T_i)$, s_i is μ-measurable.*

Lemma 2.2. *Fix a Θ-based universal Bayesian game, viz. $(\Gamma, \mathcal{U}(\Theta))$.*

    (i) Fix a Bayesian equilibrium $(s_1, \ldots, s_{|I|})$ of $(\Gamma, \mathcal{U}(\Theta))$. For each i and each $c_i \in C_i$, $\Pi_i[c_i, s_{-i}]$ is universally measurable.

    (ii) If Γ is injective, any Bayesian equilibrium of $(\Gamma, \mathcal{U}(\Theta))$ is a universally measurable equilibrium.

Proof. Begin with part (i): Fix some universal Bayesian game $(\Gamma, \mathcal{U}(\Theta))$ and write $\mathcal{U}(\Theta) = (\Theta, (U_i, \gamma_i)_{i \in I})$. Note, $\mathcal{U}(\Theta)$ is complete in the sense of Brandenburger (2003), i.e., for each $\mu_i \in \Delta(\Theta \times U_{-i})$, there is a type $u_i \in U_i$ with $\gamma_i(u_i) = \mu_i$. (See Proposition 4.1 in Friedenberg, 2010 or, alternatively, use Theorem 4 in Meier, 2012.) So, by Condition (i) of Definition 2.3, if $(s_1, \ldots, s_{|I|})$ is a Bayesian equilibrium, then for each $c_i \in C_i$ and each $\mu_i \in \Delta(\Theta \times U_{-i})$, $\Pi_i[c_i, s_{-i}]$ is μ_i-measurable.

Turn to part (ii): By part (i) and Lemma 2.1, id $\times s_{-i}$ is universally measurable for each i. Then, by Lemma A.1, s_i is universally measurable for each i. ∎

3 The Extension and Pull-Back Properties

We now turn to formalize the extension and pull back properties. To do so, fix two Θ-based structures $\mathcal{T} = (\Theta, (T_i, \beta_i)_{i \in I})$ and $\mathcal{T}^* = (\Theta, (T_i^*, \beta_i^*)_{i \in I})$. We want to map $\mathcal{T}$ to $\mathcal{T}^*$ in a way that preserves hierarchies of beliefs, i.e., for each player i and each type t_i in T_i, there is a type t_i^* in T_i^* that induces the same hierarchy of beliefs. As we have seen, the type morphism concept allows us to capture this idea without explicitly describing hierarchies of beliefs. We will state the Extension and Pull-Back properties relative to the type morphism concept. Below we explain why.

Definition 3.1. *Fix Θ-based type structures $\mathcal{T}$ and $\mathcal{T}^*$. Say $\mathcal{T}$ can be **mapped to $\mathcal{T}^*$** (**via $\mathbf{h_1}, \ldots, \mathbf{h_{|I|}}$**) if $(h_1, \ldots, h_{|I|})$ is a type morphism from $\mathcal{T}$ to $\mathcal{T}^*$.*

Given a Θ-based game Γ, write s_i for a strategy of player i in the Bayesian game $(\Gamma, \mathcal{T})$, and write s_i^* for a strategy of player i in the Bayesian game $(\Gamma, \mathcal{T}^*)$. We begin by stating the Pull-Back Property.

Definition 3.2. *Let $\mathcal{T}$ and $\mathcal{T}^*$ be Θ-based type structures, so that $\mathcal{T}$ can be mapped to $\mathcal{T}^*$. Say a Bayesian Equilibrium, viz. $(s_1^*, \ldots, s_{|I|}^*)$, of $(\Gamma, \mathcal{T}^*)$ can be **pulled-back** to a Bayesian Equilibrium of $(\Gamma, \mathcal{T})$ if, for every type morphism, viz. $(h_1, \ldots, h_{|I|})$, from $\mathcal{T}$ to $\mathcal{T}^*$, $(s_1^* \circ h_1, \ldots, s_{|I|}^* \circ h_{|I|})$ is a Bayesian Equilibrium of $(\Gamma, \mathcal{T})$.*

The pair $\langle \mathcal{T}, \mathcal{T}^* \rangle$ satisfies the **Equilibrium Pull-Back Property for the Θ-based game Γ** if each Bayesian Equilibrium of $(\Gamma, \mathcal{T}^*)$ can be pulled-back to a Bayesian Equilibrium of $(\Gamma, \mathcal{T})$.

To understand Definition 3.2, refer back to Example 1.2. We saw that there exists a hierarchy preserving map from the structure $\mathcal{T}^*$ to the structure $\mathcal{T}$. We argued above that the structure $\mathcal{T}$ specified external signals that were not included in the structure $\mathcal{T}^*$. This was not the idea we sought to capture. Indeed, this is why we require the mapping, in Definition 3.2, to be a type morphism and not simply any hierarchy preserving map: There is no type morphism from the structure $\mathcal{T}^*$ to the structure $\mathcal{T}$.

To state the Extension Property we will need more. Refer back to Example 1.1. We saw there exists a hierarchy preserving map from the structure $\mathcal{T}$ to the structure $\mathcal{T}^*$. Indeed, that hierarchy preserving map is a type morphism, despite the fact that the structure $\mathcal{T}$ specifies external signals that are not included in the type structure $\mathcal{T}^*$. While the type morphism concept is sufficient to 'rule out' external signals associated with the analyst's type structure, it is insufficient to 'rule out' external signals associated with the players' type structure. To understand how these signals will be ruled out, observe that the type structure $\mathcal{T}$ was redundant, in the sense that there were two types that induce the same hierarchies of beliefs. To 'rule out' external signals associated with the players' type structure, it suffices to assume that the players' type structure is non-redundant.

We will not need a formal definition of non-redundancy. Instead, we will make use of a single property that follows from it—namely, if $\mathcal{T}$ can be mapped to $\mathcal{T}^*$ via $(h_1, \ldots, h_{|I|})$ and $\mathcal{T}$ is non-redundant, then $(h_1, \ldots, h_{|I|})$ is injective.[6]

Definition 3.3. *Let $\mathcal{T}$ and $\mathcal{T}^*$ be two Θ-based type structures, so that $\mathcal{T}$ can be mapped to $\mathcal{T}^*$ via injective maps $(h_1, \ldots, h_{|I|})$. Say a Bayesian Equilibrium, viz. $(s_1, \ldots, s_{|I|})$, of $(\Gamma, \mathcal{T})$ can be **extended** to a Bayesian Equilibrium of $(\Gamma, \mathcal{T}^*)$ if there exists a Bayesian Equilibrium, viz. $(s_1^*, \ldots, s_{|I|}^*)$, of $(\Gamma, \mathcal{T}^*)$ so that $(s_1^* \circ h_1, \ldots, s_{|I|}^* \circ h_{|I|}) = (s_1, \ldots, s_{|I|})$.*

The pair $\langle \mathcal{T}, \mathcal{T}^* \rangle$ **satisfies the Equilibrium Extension Property for the Θ-based game Γ** if each Bayesian Equilibrium of $(\Gamma, \mathcal{T})$ can be extended to a Bayesian Equilibrium of $(\Gamma, \mathcal{T}^*)$. The pair $\langle \mathcal{T}, \mathcal{T}^* \rangle$ **fails the Equilibrium Extension Property for the Θ-based game Γ** if $\mathcal{T}$ can be mapped to $\mathcal{T}^*$ via injective maps $(h_1, \ldots, h_{|I|})$, but there is some Bayesian Equilibrium of $(\Gamma, \mathcal{T})$ that cannot be extended to a Bayesian Equilibrium of $(\Gamma, \mathcal{T}^*)$.

---

[6]Redundancy implies injectivity but the converse does not hold.

Observe that, if $\mathcal{T}$ cannot be mapped to $\mathcal{T}^*$ via injective maps, then the pair $\langle \mathcal{T}, \mathcal{T}^* \rangle$ neither satisfies nor fails the Equilibrium Extension Property. The Equilibrium Extension Property is **not defined** when one type structure cannot be mapped to the second via injective maps. So, in particular, Examples 1.1-1.2 do not show a failure of the Equilibrium Extension Property. In both examples, the unique map from $\mathcal{T}$ to $\mathcal{T}^*$ was not injective. Indeed, there we explained why those examples do not capture the question of invariance to misspecifying the context of the game. To address our question, we need to be able to view $\mathcal{T}$ as a substructure of $\mathcal{T}^*$. The injectivity requirement allows us to do just that.

Remark 3.1. *Let $\mathcal{T}$ and $\mathcal{T}^*$ be two Θ-based type structures, so that $\mathcal{T}$ can be mapped to $\mathcal{T}^*$ via injective maps $(h_1, \ldots, h_{|I|})$. Then, by Purves's (1966) Theorem, each h_i is bimeasurable, i.e., the image of each measurable set is itself measurable. So, by definition, each h_i is an embedding. In light of this, we will say that $\mathcal{T}$ can be **embedded** into $\mathcal{T}^*$ via $(h_1, \ldots, h_{|I|})$ if $\mathcal{T}$ can be mapped to $\mathcal{T}^*$ via $(h_1, \ldots, h_{|I|})$, and each h_i is injective.*

We conclude this section by pointing out some basic observations about the Equilibrium Pull-Back and Extension Properties, which will be useful in the subsequent analysis. First, the Equilibrium Pull-Back Property is satisfied. Indeed, we can also obtain a Measurable Pull-Back Property. (This property will be useful in arguments to come.) The pair $\langle \mathcal{T}, \mathcal{T}^* \rangle$ satisfies the **Measurable Equilibrium Pull-Back Property for the Θ-based game Γ** if any universally measurable equilibrium of $(\Gamma, \mathcal{T}^*)$ can be pulled-back to a universally measurable equilibrium of $(\Gamma, \mathcal{T})$.[7]

Proposition 3.1. *Fix Θ-based type structures, $\mathcal{T}$ and $\mathcal{T}^*$, so that $\mathcal{T}$ can be mapped to $\mathcal{T}^*$.*

   (i) For any Θ-based game Γ, the pair $\langle \mathcal{T}, \mathcal{T}^* \rangle$ satisfies the Equilibrium Pull-Back Property for Γ.

   (ii) For any Θ-based game Γ, the pair $\langle \mathcal{T}, \mathcal{T}^* \rangle$ satisfies the Measurable Equilibrium Pull-Back Property for Γ.

Second, return back to the idea of a universal Bayesian game. We have the following property.

Lemma 3.1. *Fix a Θ-based Bayesian game $(\Gamma, \mathcal{T})$ so that $\mathcal{T}$ can be embedded into $\mathcal{U}(\Theta)$. Then the following are equivalent:*

   (i) The pair $\langle \mathcal{T}, \mathcal{U}(\Theta) \rangle$ satisfies the Equilibrium Extension Property for Γ.

   (ii) For every Θ-based structure $\mathcal{T}^*$ so that $\mathcal{T}$ can be embedded into $\mathcal{T}^*$, the pair $\langle \mathcal{T}, \mathcal{T}^* \rangle$ satisfies the Equilibrium Extension Property for Γ.

---

[7]Observe, the Measurable Equilibrium Pull-Back Property is not a logical implication of the Equilibrium Pull-Back Property.

4 An Equilibrium Extension Failure

This section shows an example of a finite Θ-based game Γ and Θ-based type structures $\mathcal{T}^*, \mathcal{T}^{**}$, so that $\langle \mathcal{T}^*, \mathcal{T}^{**} \rangle$ fails the Equilibrium Extension Property for Γ. The game Γ is built off an important example in Hellman (2014). Section 6b elaborates on the connection, as well as the connection to an important paper by Simon (2003).

Example of Extension Failure

$\underline{\theta}$	L_2	M_2	R_2
L_1	3, 3	2, 2	1, -1
M_1	2, 2	3, 3	1, -1
R_1	-1, 1	-1, 1	1, 1

$\overline{\theta}$	L_2	M_2	R_2
L_1	4, 4	6, 7	1, -1
M_1	7, 6	4, 4	1, -1
R_1	-1, 1	-1, 1	1, 1

θ^*	L_2	M_2	R_2
L_1	1, 1	1, 1	0, 0
M_1	1, 1	1, 1	0, 0
R_1	0, 0	0, 0	1, 1

Figure 4.1: A Game that Fails Extension

Figure 4.1 describes a finite Θ-based game Γ. We point to several features of the game: The parameter set is $\Theta = \{\underline{\theta}, \overline{\theta}, \theta^*\}$. When the parameter is $\underline{\theta}$, L_i (resp. M_i) uniquely maximizes i's payoffs if and only if i's co-player, viz. $-i$, plays L_{-i} (resp. M_{-i}). When the parameter is $\overline{\theta}$, L_i (resp. M_i) uniquely maximizes i's payoffs if and only if i's co-player, viz. $-i$, plays M_{-i} (resp. L_{-i}). For any given parameter, R_i maximizes i's payoffs only if i's co-player, viz. $-i$, plays R_{-i}. When $\theta \in \{\underline{\theta}, \overline{\theta}\}$, R_i maximizes i's payoffs if and only if both L_i and M_i also maximize i's payoffs.

Remark 4.1. *For any Θ-based structure $\mathcal{T}$, $(\Gamma, \mathcal{T})$ has a Bayesian equilibrium. In particular, for each $(\Gamma, \mathcal{T})$, there is a Bayesian equilibrium where each type of each player chooses R_i with probability one.*

Nonetheless, we construct Θ-based structures $\mathcal{T}^*$ and $\mathcal{T}^{**}$, where some equilibrium of $(\Gamma, \mathcal{T}^*)$ cannot be extended to an equilibrium of $(\Gamma, \mathcal{T}^{**})$. Of course, there will be another equilibrium of $(\Gamma, \mathcal{T}^*)$ that can be extended to an equilibrium of $(\Gamma, \mathcal{T}^{**})$, i.e., the one just mentioned above.

We will take the players' type structure $\mathcal{T}^* = (\Theta, T_1^*, T_2^*, \beta_1^*, \beta_2^*)$ to be an arbitrary countable type structure, i.e., where T_1^* and T_2^* are countable. Then we have:

Remark 4.2. *For any Θ-based countable type structure $\mathcal{T}^*$, $(\Gamma, \mathcal{T}^*)$ has a Bayesian equilibrium, viz. (s_1^*, s_2^*), where for each i and each $t_i^* \in T_i^*$, $s_i^*(t_i^*)$ assigns probability one to $\{L_i, M_i\}$.*

To see this, take $\hat{\Gamma}$ to be the Θ-based game that differs from Γ only in restricting the action sets to $\{L_i, M_i\}$. Then the remark follows from the existence of a Bayesian equilibrium for $(\hat{\Gamma}, \mathcal{T}^*)$ and the fact that R_i is not optimal under any belief that assigns probability one to $\{L_{-i}, M_{-i}\}$.[8]

We will take the analyst's type structure $\mathcal{T}^{**}$ to be some arbitrary universal Θ-based type structure, viz. $\mathcal{U}(\Theta) = (\Theta, U_1, U_2, \gamma_1, \gamma_2)$. This has the property that $\mathcal{T}^*$ can be mapped to $\mathcal{U}(\Theta)$. In keeping with Remark 4.1, there is a Bayesian equilibrium of $(\Gamma, \mathcal{T}^*)$ that can be extended to a Bayesian equilibrium of $(\Gamma, \mathcal{U}(\Theta))$, specifically the equilibrium where each type of each player i assigns probability one to R_i. But, we will show:

Theorem 4.1.

  (i) *There is a Bayesian equilibrium, viz. (s_1^*, s_2^*), of $(\Gamma, \mathcal{T}^*)$ that cannot be extended to a Bayesian equilibrium of $(\Gamma, \mathcal{U}(\Theta))$.*

  (ii) *There is a Bayesian equilibrium of $(\Gamma, \mathcal{U}(\Theta))$.*

Part (i) of Theorem 4.1 says that the pair $\langle \mathcal{T}^*, \mathcal{U}(\Theta) \rangle$ fails the Equilibrium Extension Property. Part (ii) says that this failure of Equilibrium Extension is non-trivial. In the statement of Theorem 4.1, we will take (s_1^*, s_2^*) to be any equilibrium where, for each player i, each type t_i^* assigns probability one to $\{L_i, M_i\}$.

Before coming to the proof, it is worth repeating that Theorem 4.1 delivers a failure of Equilibrium Extension with standard ingredients. In particular, the Θ-based game Γ is finite. Moreover, for any Θ-based type structure $\mathcal{T}$, there exists an equilibrium of the Bayesian game $(\Gamma, \mathcal{T})$, a fortiori for the universal Bayesian game $(\Gamma, \mathcal{U}(\Theta))$.

Proof of Main Theorem

We now turn to show Theorem 4.1. To do so, we make use of Lemma 3.1. By that Lemma, it suffices to show:

Lemma 4.1. *There exists a Θ-based type structure $\mathcal{T} = (\Theta, T_1, T_2, \beta_1, \beta_2)$, so that:*

  (i) *$\mathcal{T}^*$ can be embedded into $\mathcal{T}$, and*

  (ii) *the pair $\langle \mathcal{T}^*, \mathcal{T} \rangle$ fails the Equilibrium Extension Property for Γ.*

We turn to construct a Θ-based type structure $\mathcal{T} = (\Theta, T_1, T_2, \beta_1, \beta_2)$ satisfying the requirements of Lemma 4.1. To do so, we begin by constructing a subset of the parameters, viz. $\overline{\Theta} = \{\underline{\theta}, \overline{\theta}\}$. Let $\mathcal{U}(\overline{\Theta}) = (\overline{\Theta}, \overline{U}_1, \overline{U}_2, \overline{\gamma}_1, \overline{\gamma}_2)$ be a $\overline{\Theta} = \{\underline{\theta}, \overline{\theta}\}$-based universal type structure. We use the $\overline{\Theta}$-based structure $\mathcal{U}(\overline{\Theta})$ to build the Θ-based structure $\mathcal{T}$.

Each T_i is the disjoint union of T_i^* and $\overline{U}_i$, where we endow T_i with the disjoint union topology. For the map β_i, refer to Figure 4.2: For each $t_i^* \in T_i^* \subseteq T_i$, take $\beta_i(t_i^*)(E_{-i}) = \beta_i^*(t_i^*)(E_{-i} \cap (\Theta \times$

---

[8]It is often taken for granted that there exists a Bayesian equilibrium for some $(\Gamma^*, \mathcal{T}^*)$ where Γ is finite and $\mathcal{T}^*$ is countable. Takahashi (2009) has written a proof of this claim making use of Glicksberg's (1952) Theorem.

T^*_{-i})). For each type $\overline{u}_i \in \overline{U}_i \subseteq T_i$, define $\beta_i(\overline{u}_i)$ as follows. Fix some $p \in (0,1)$ and $t^*_{-i} \in T^*_{-i}$. (Note, p and t^*_{-i} will be chosen to be the same for each type $\overline{u}_i \in \overline{U}_i$.) Take $\beta_i(\overline{u}_i)(E_{-i}) = p\overline{\gamma}_i(\overline{u}_i)(E_{-i} \cap (\overline{\Theta} \times \overline{U}_{-i})) + (1-p)$ if $(\theta^*, t^*_{-i}) \in E_{-i}$. Take $\beta_i(\overline{u}_i)(E_{-i}) = p\overline{\gamma}_i(\overline{u}_i)(E_{-i} \cap (\overline{\Theta} \times \overline{U}_{-i}))$ if $(\theta^*, t^*_{-i}) \notin E_{-i}$.
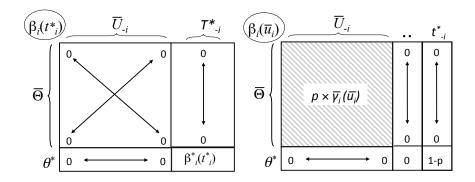

Figure 4.2: Constructed Type Structure

Property 4.1. *The structure $\mathcal{T}^*$ can be embedded into $\mathcal{T}$ via* $(\mathrm{id}_1, \mathrm{id}_2)$*, where* $\mathrm{id}_i : T^*_i \to T_i$ *denotes the identity map.*

Of course, there is a Bayesian equilibrium of $(\Gamma, \mathcal{T})$, i.e., where each type of i plays R_i with probability one. But, we will show that there is no Bayesian equilibrium of $(\Gamma, \mathcal{T})$, viz. (s_1, s_2), where each $t_i \in T^*_i \subseteq T_i$ has $s_i(t_i)(\{L_i, M_i\}) = 1$.

Why is this the case? Derive an $\overline{\Theta} = \{\underline{\theta}, \overline{\theta}\}$-based game, viz. $\overline{\Gamma} = (\overline{\Theta}, \overline{C}_1, \overline{C}_2, \overline{\pi}_1, \overline{\pi}_2)$, from Γ: Take each $\overline{C}_i = \{L_i, M_i\}$ and take each $\overline{\pi}_i$ to be the restriction of π_i to $\overline{\Theta} \times \overline{C}_1 \times \overline{C}_2$. The game $\overline{\Gamma}$ is given in Figure 4.3.

Now we will see two seemingly contradictory facts: First, there is no Bayesian equilibrium of the game $(\overline{\Gamma}, \mathcal{U}(\overline{\Theta}))$. Second, if there is a Bayesian equilibrium of $(\Gamma, \mathcal{T})$ so that, for each i, $t^*_i \in T^*_i \subseteq T_i$ plays $\{L_i, M_i\}$ with probability one, then there is a Bayesian equilibrium of $(\overline{\Gamma}, \mathcal{U}(\overline{\Theta}))$. Putting these two together, we get that there is no equilibrium of $(\Gamma, \mathcal{T})$ so that, for each i, $t^*_i \in T^*_i \subseteq T_i$ plays $\{L_i, M_i\}$ with probability one. As such, we cannot extend an equilibrium of $(\Gamma, \mathcal{T}^*)$ to an equilibrium of $(\Gamma, \mathcal{T})$. This establishes Lemma 4.1 and completes the proof of Theorem 4.1.

$\underline{\theta}$	L_2	M_2
L_1	3, 3	2, 2
M_1	2, 2	3, 3

$\overline{\theta}$	L_2	M_2
L_1	4, 4	6, 7
M_1	7, 6	4, 4

Figure 4.3: The Restricted Game

We now turn to the two stated steps. First:

Proposition 4.1. *There is no equilibrium of the (finite) $\overline{\Theta}$-based Bayesian game $(\overline{\Gamma}, \mathcal{U}(\overline{\Theta}))$.*[9]

To see why this is the case, let us recall an important result from Hellman (2014): He studies the $\overline{\Theta}$-based game $\overline{\Gamma}$ and shows that there is some associated Bayesian game $(\overline{\Gamma}, \overline{\mathcal{T}})$ with no universally measurable Bayesian equilibrium.

Proposition 4.2. *There exists a $\overline{\Theta}$-based structure $\overline{\mathcal{T}}$ so that:*

   (i) [Simon, 2003] there is an equilibrium of the Bayesian game $(\overline{\Gamma}, \overline{\mathcal{T}})$, but

   (ii) [Hellman, 2014] there is no universally measurable equilibrium of the Bayesian game $(\overline{\Gamma}, \overline{\mathcal{T}})$.

Part (i) follows from Proposition 1 in Simon (2003). We don't make use of this part; we only include it for completeness. We will make explicit use of part (ii). The Online Appendix reviews this result with an eye toward explaining two facts: First, the payoffs in Figure 4.3 are formally different from those found in Hellman; we verify that Hellman's (2014) arguments apply to this case. (His proof is robust to changing payoffs, so long as three conditions are satisfied. We change his example so that it satisfies injectivity.) Second, we show how Hellman's result can be translated to our framework. (His notion of a type structure is different from that here.)

Part (ii) is used to show:

Proposition 4.3. *The finite $\overline{\Theta}$-based universal Bayesian game $(\overline{\Gamma}, \mathcal{U}(\overline{\Theta}))$ does not have a Bayesian equilibrium.*

Proof. Suppose, contra hypothesis, that there is an equilibrium of $(\overline{\Gamma}, \mathcal{U}(\overline{\Theta}))$. Note, $\overline{\Gamma}$ is injective. So, by Lemma 2.2, the equilibrium is universally measurable. But then, using the Measurable Pull-Back property (Proposition 3.1), there is a universally measurable equilibrium of $(\overline{\Gamma}, \overline{\mathcal{T}})$, where $\overline{\mathcal{T}}$ is as in Proposition 4.2(ii). But this contradicts Hellman's (2014) result (i.e., Proposition 4.2). ∎

This shows the first step. Now we turn to show the second step:

Lemma 4.2. *Suppose there is an equilibrium $(\Gamma, \mathcal{T})$, viz. (s_1, s_2), so that, for each i and each $t_i^* \in T_i^* \subseteq T_i$, $s_i(t_i^*)$ assigns probability one to $\{L_i, M_i\}$. Then there is a Bayesian equilibrium of $(\overline{\Gamma}, \mathcal{U}(\overline{\Theta}))$.*

Proof. Fix a Bayesian Equilibrium, viz. (s_1, s_2), of $(\Gamma, \mathcal{T})$ so that, for each i and each $t_i^* \in T_i^* \subseteq T_i$, $s_i(t_i^*)(\{L_i, M_i\}) = 1$. We will first show that, for each i and each $\overline{u}_i \in \overline{U}_i \subseteq T_i$, $s_i(\overline{u}_i)(\{L_i, M_i\}) = 1$. Then, we will use this fact to construct a Bayesian equilibrium $(\overline{s}_1, \overline{s}_2)$ of $(\overline{\Gamma}, \mathcal{U}(\overline{\Theta}))$.

Fix some $\overline{u}_i \in \overline{U}_i \subseteq T_i$. For this type, the expected payoffs from choosing some $c_i \in \{L_i, M_i\}$ are

$$\int_{\overline{\Theta} \times \overline{U}_{-i}} \Pi_i[c_i, s_{-i}] d\beta_i(\overline{u}_i) = p \int_{\overline{\Theta} \times \overline{U}_{-i}} \Pi_i[c_i, s_{-i}] d\overline{\gamma}_i(\overline{u}_i) + (1 - p).$$

---

[9]While, as a step, we will make use of a Bayesian game without an equilibrium, the final product will not simply be a corollary of this fact. See Section 6c for a discussion.

(Recall, $\Pi_i[c_i, s_{-i}] = \hat{\pi}_i[c_i] \cdot (\text{id} \times s_{-i})$.) This type's expected payoffs from choosing R_i are

$$\int_{\overline{\Theta} \times \overline{U}_{-i}} \Pi_i[R_i, s_{-i}] d\beta_i(\overline{u}_i) = p \int_{\overline{\Theta} \times \overline{U}_{-i}} \Pi_i[R_i, s_{-i}] d\overline{\gamma}_i(\overline{u}_i).$$

Note, for any given $(\theta, \overline{u}_{-i}) \in \overline{\Theta} \times \overline{U}_{-i}$ and any given $c_i \in \{L_i, M_i\}$,

$$\Pi_i[c_i, s_{-i}](\theta, \overline{u}_{-i}) \geq \Pi_i[R_i, s_{-i}](\theta, \overline{u}_{-i}).$$

Since $1 - p > 0$, it follows that

$$\max\left\{\int_{\overline{\Theta} \times \overline{U}_{-i}} \Pi_i[L_i, s_{-i}] d\beta_i(\overline{u}_i), \int_{\overline{\Theta} \times \overline{U}_{-i}} \Pi_i[M_i, s_{-i}] d\beta_i(\overline{u}_i)\right\} > \int_{\overline{\Theta} \times \overline{U}_{-i}} \Pi_i[R_i, s_{-i}] d\beta_i(\overline{u}_i).$$

Thus, it follows from condition (ii) of a Bayesian equilibrium that, for each $\overline{u}_i \in \overline{U}_i$, $s_i(\overline{u}_i)(\{L_i, M_i\}) = 1$.

Now turn to the $\overline{\Theta} = \{\underline{\theta}, \overline{\theta}\}$-based game Bayesian $(\overline{\Gamma}, \mathcal{U}(\overline{\Theta}))$. Construct strategies $\overline{s}_1$ and $\overline{s}_2$ as follows: For each $\overline{u}_i \in \overline{U}_i$ and each $E_i \subseteq \{L_i, R_i\}$, $\overline{s}_i(\overline{u}_i)(E_i) = s_i(\overline{u}_i)(E_i)$. (This is well defined since $s_i(\overline{u}_i)(\{L_i, M_i\}) = 1$, as we have shown above.) We will show that $(\overline{s}_1, \overline{s}_2)$ is an equilibrium of the Bayesian game $(\overline{\Gamma}, \mathcal{U}(\overline{\Theta}))$.

To show condition (i): Fix some $c_i \in \{L_i, M_i\}$ and some $\overline{u}_i \in \overline{U}_i$. Let $f_{-i} : \overline{\Theta} \times \overline{U}_{-i} \to \Theta \times T_{-i}$ be the identity mapping. Note that $\overline{\Pi}_i[c_i, \overline{s}_{-i}] = \Pi_i[c_i, s_{-i}] \circ f_{-i}$. Then, by Corollary B.1, $\overline{\Pi}_i[c_i, \overline{s}_{-i}]$ is $\overline{\gamma}_i(\overline{u}_i)$-integrable provided that $\Pi_i[c_i, s_{-i}]$ is $\underline{f}_{-i}(\overline{\gamma}_i(\overline{u}_i))$-integrable. (Note, $\underline{f}_{-i}(\overline{\gamma}_i(\overline{u}_i))$ is the image measure of $\overline{\gamma}_i(\overline{u}_i)$ under f_{-i}.) The fact that $\Pi_i[c_i, s_{-i}]$ is $\underline{f}_{-i}(\overline{\gamma}_i(\overline{u}_i))$-integrable follows from the fact that $\Pi_i[c_i, s_{-i}]$ is $\beta_i(\overline{u}_i)$-integrable and Lemma B.2.

To show condition (ii): Fix some $\overline{u}_i \in \overline{U}_i$ and some $c_i \in \{L_i, M_i\}$. Then,

$$\int_{\Theta \times T_{-i}} \Pi_i[c_i, s_{-i}] d\beta_i(\overline{u}_i) = p \int_{\overline{\Theta} \times \overline{U}_{-i}} \overline{\Pi}_i[c_i, \overline{s}_{-i}] d\overline{\gamma}_i(\overline{u}_i) + (1 - p).$$

Since (s_1, s_2) is a Bayesian equilibrium with $s_i(\overline{u}_i)(\{L_i, M_i\}) = 1$ (and $\overline{s}_i(\overline{u}_i) = s_i(\overline{u}_i)$), it follows that

$$p \int_{\overline{\Theta} \times \overline{U}_{-i}} \overline{\Pi}_i[\overline{s}_i(\overline{u}_i), \overline{s}_{-i}] d\overline{\gamma}_i(\overline{u}_i) + (1 - p) \geq p \int_{\overline{\Theta} \times \overline{U}_{-i}} \overline{\Pi}_i[c_i, \overline{s}_{-i}] d\overline{\gamma}_i(\overline{u}_i) + (1 - p),$$

for each $c_i \in \{L_i, M_i\}$. From this,

$$\int_{\overline{\Theta} \times \overline{U}_{-i}} \overline{\Pi}_i[\overline{s}_i(t_i), \overline{s}_{-i}] d\overline{\gamma}_i(\overline{u}_i) \geq \int_{\overline{\Theta} \times \overline{U}_{-i}} \overline{\Pi}_i[c_i, \overline{s}_{-i}] d\overline{\gamma}_i(\overline{u}_i),$$

for each $c_i \in \{L_i, M_i\}$. ∎

5 Positive Results

In Section 4, we saw that the Equilibrium Extension Property may fail. Now we ask: Are there (interesting) situations where the Extension Property does obtain? To do so, it will suffice to focus on situations where the players' type structure can be seen as a strict substructure of the analyst's type structure. We formalize this idea below.

Definition 5.1. *Say $\mathcal{T}$ can be **properly embedded** into $\mathcal{T}^*$ if $\mathcal{T}$ can be embedded into $\mathcal{T}^*$ but $\mathcal{T}^*$ cannot be embedded into $\mathcal{T}$.*

Fix Θ-based structures $\mathcal{T}$. Say $\prod_{i \in I} E_i$ is a **belief-closed subset** of $T = \prod_{i \in I} T_i$ if each E_i is measurable in T_i and, for each $t_i \in E_i$, $\beta_i(t_i)(\Theta \times E_{-i}) = 1$.

Lemma 5.1. *Fix Θ-based structures $\mathcal{T}$ and $\mathcal{T}^*$. If $\mathcal{T}$ can be embedded into $\mathcal{T}^*$ via $(h_1, \ldots, h_{|I|})$, then $\prod_{i \in I} h_i(T_i)$ is a belief-closed subset of $T^* = \prod_{i \in I} T_i^*$.*

Now, by the Pull-Back Property, we have:

Corollary 5.1. *Let $\mathcal{T}$ and $\mathcal{T}^*$ be two Θ-based type structures, so that $\mathcal{T}$ can be embedded into $\mathcal{T}^*$ and $\mathcal{T}^*$ can be embedded into $\mathcal{T}$. Then, the pair $\langle \mathcal{T}, \mathcal{T}^* \rangle$ satisfies the Equilibrium Extension Property.*

In light of Corollary 5.1, we will focus on the case in which $\mathcal{T}$ can be properly embedded into $\mathcal{T}^*$. In this case, we have the following:

Lemma 5.2. *Fix Θ-based structures $\mathcal{T}$ and $\mathcal{T}^*$, so that $\mathcal{T}$ can be properly embedded into $\mathcal{T}^*$ via $(h_1, \ldots, h_{|I|})$. Then, for some $i = 1, \ldots, |I|$, $h_i(T_i) \subsetneq T_i^*$.*

5.1 Compact and Continuous Games

In Section 4, we saw that, even in the case of a finite game, the Extension property may fail. Note two features of the example: First, because the players' type structure was (at most) countable, each equilibrium of the players' Bayesian game was a universally measurable equilibrium. Second, there were an uncountable number of types that are in the analyst's structure but not in the players' structure—so, the analyst's structure is "large" relative to the players' type structure.

This section picks up on these two features. In particular, fix a finite game, viz. Γ, and an associated "players' Bayesian game," viz. $(\Gamma, \mathcal{T})$. We will see that, if there are (at most) a countable number of types that are in the analyst's structure but not the players' structure, then we will be able to extend any universally measurable equilibrium of the players' structure to the analyst's structure. Thus, for a finite game, we can only have an extension failure if either (i) there are an uncountable number of types that are in the analyst's structure but not the players' structure, or (ii) the original equilibrium (i.e., which we are trying to extend) is not universally measurable.

Definition 5.2. *Say a Θ-based game, viz. $\Gamma = (\Theta, (C_i, \pi_i)_{i \in I})$, is **compact and continuous** if each C_i is compact and each π_i is continuous.*

(Note, there is no requirement that Θ be compact.)

Proposition 5.1. *Fix Θ-based structures $\mathcal{T}$ and $\mathcal{T}^*$, so that $\mathcal{T}$ can be properly embedded into $\mathcal{T}^*$ via $(h_1, \ldots, h_{|I|})$ and each $T_1^* \backslash h_1(T_1), \ldots, T_{|I|}^* \backslash h_{|I|}(T_{|I|})$ is (at most) countable. If Γ is compact and continuous, then any universally measurable equilibrium of $(\Gamma, \mathcal{T})$ can be extended to a universally measurable equilibrium of $(\Gamma, \mathcal{T}^*)$.*

The proof can be found in Appendix D. Here, we give the idea. In so doing, we will see the role of the requirements that each $T_i^* \backslash h_i(T_i)$ is countable and that each s_i is universally measurable.

Suppose $\mathcal{T}$ can be embedded into $\mathcal{T}^*$ via $(h_1, \ldots, h_{|I|})$. Fix a universally measurable equilibrium $(s_1, \ldots, s_{|I|})$ of the Bayesian game $(\Gamma, \mathcal{T})$. We want to show that there is an equilibrium of the Bayesian Game $(\Gamma, \mathcal{T}^*)$, viz. $(s_1^*, \ldots, s_{|I|}^*)$, that extends the equilibrium $(s_1, \ldots, s_{|I|})$, i.e., that satisfies $(s_1, \ldots, s_{|I|}) = (s_1^* \circ h_1, \ldots, s_{|I|}^* \circ h_{|I|})$.

We will begin by constructing a certain game of complete information, viz. G, that depends on the game Γ and the equilibrium $(s_1, \ldots, s_{|I|})$. There will be many players in this game, each corresponding to a type in $T_i^* \backslash h_i(T_i)$ for some player i. As such, there are (at most) a countable number of players in this game. Each such player $t_i \in T_i^* \backslash h_i(T_i)$ gets to make a choice from C_i, as in Γ. The payoff functions will be constructed in a specific way. In particular, they will depend on Γ and the equilibrium $(s_1, \ldots, s_{|I|})$.

The complete information game G is compact and continuous. Compactness follows from the fact that the underlying game is compact. Continuity uses the fact that the underlying game is continuous, but it does not follow immediately from this fact. There are two issues: First, the payoff functions depend on the equilibrium and the equilibrium may be discontinuous. Second, there may be an infinite (but countable) number of players in the game and, when there are a countable number of players, payoff functions may be discontinuous even if the choice set is finite. See Peleg (1969).[10]

Now we have a compact and continuous complete information game G, with a countable number of players. As such, we can apply Glicksberg's (1952) fixed-point theorem to show that there exists a mixed-strategy equilibrium of G.

Finally, we return to the Bayesian game $(\Gamma, \mathcal{T}^*)$. We consider strategies that extend the equilibrium $(s_1, \ldots, s_{|I|})$ of $(\Gamma, \mathcal{T})$. We show that, in a certain sense, these strategies correspond to the mixed strategies of the complete information game G. As such, we can use the fact that there is a mixed strategy equilibrium of G to show that there is an equilibrium of $(\Gamma, \mathcal{T}^*)$ that extends the equilibrium $(s_1, \ldots, s_{|I|})$ of $(\Gamma, \mathcal{T})$.

---

[10]Satoru Takahashi pointed us to the fact that, if a game with a countable number of players is (in a sense) "generated" by a compact and continuous game of incomplete information, then the payoff functions are nonetheless continuous. In so doing, Takahashi generalized a result in a previous version of this paper. We are very much indebted to Satoru for this contribution.

Notice, it is important, for this argument, that we begin with a universally measurable equilibrium of the Bayesian game $(\Gamma, \mathcal{T})$. To see why, suppose that we begin instead with an equilibrium, viz. $(s_1, \ldots, s_{|I|})$, that is not universally measurable. Then, by Lemma A.1, some $\mathrm{id} \times s_{-i}$ is not μ-measurable for some $\mu \in \Delta(\Theta \times T_{-i})$. Notice, there may be a type $t_i^* \in T_i^* \backslash h_i(T_i)$ so that $\beta_i^*(t_i^*)(\Theta \times h_{-i}(T_{-i})) = 1$ and $\beta_i^*(t_i^*)(\cdot | \Theta \times h_{-i}(T_{-i})) = \mu$. Then for any extension of s_{-i}, viz. s_{-i}^*, $\mathrm{id} \times s_{-i}^*$ is not $\beta_i^*(t_i^*)$-measurable. In this case, we cannot associate an equilibrium of the complete information game with an equilibrium of the Bayesian game $(\Gamma, \mathcal{T}^*)$.[11] Of course, for a type structure with (at most) a countable number of types, all strategies are measurable. As such, an equilibrium is a universally measurable equilibrium. With this, we have the following corollary:

Corollary 5.2. *Fix Θ-based structures $\mathcal{T}$ and $\mathcal{T}^*$, so that $\mathcal{T}$ can be properly embedded into $\mathcal{T}^*$. If Γ is compact and continuous and $\mathcal{T}^*$ is countable, then any equilibrium of $(\Gamma, \mathcal{T})$ can be extended to an equilibrium of $(\Gamma, \mathcal{T}^*)$.*

5.2 The Common Prior Assumption

In this section, we will see that, if the analyst's structure satisfies the common prior assumption, then we have the Extension Property. Note, this holds independent of the underlying game Γ, i.e., even if Γ is not compact or continuous.

To see why, let us begin by reviewing the analysis in Section 4. There, we had Θ-based structures $\mathcal{T}^* = (\Theta, T_1^*, T_2^*, \beta_1^*, \beta_2^*)$ and $\mathcal{U}(\Theta) = (\Theta, U_1, U_2, \gamma_1, \gamma_2)$. The structure $\mathcal{T}^*$ can be viewed as a belief-closed subset of $\mathcal{U}(\Theta)$. Write $h_1(T_1^*) \times h_2(T_2^*) \subseteq U_1 \times U_2$ for this belief closed subset. Note, types in this belief closed subset impose an equilibrium restriction on (some) types outside of this subset. This is because there are types in $U_i \backslash h_i(T_i^*)$ that assign strictly positive probability to types in $h_{-i}(T_{-i}^*)$. This problem would not arise if the only types in the analyst's structure that assigned positive probability to types in $h_{-i}(T_{-i}^*)$ are types that are in $h_i(T_i^*)$. (Of course, this is not the case in a universal type structure.)

With the above in mind, consider the following scenario: Suppose we instead have a type structure, viz. $\mathcal{T}^*$, that can be viewed as the union of two type structures. For a given game, can we extend an equilibrium associated with one of these structures to an equilibrium associated with $\mathcal{T}^*$? The answer will be yes if and only if there exists an equilibrium associated with the other structure.

Let us first formalize the idea that a type structure $\mathcal{T}^*$ can be viewed as the union of some structure $\mathcal{T}$ and some 'remaining structure,' which we'll call the difference structure.

Definition 5.3. *Fix Θ-based structures $\mathcal{T}$ and $\mathcal{T}^*$. Say $\mathcal{T}$ **induces a decomposition of $\mathcal{T}^*$ (via $(\mathbf{h_1}, \ldots, \mathbf{h_{|I|}})$)** if $\mathcal{T}$ can be embedded into $\mathcal{T}^*$ via $(h_1, \ldots, h_{|I|})$ and $\prod_{i \in I} T_i^* \backslash h_i(T_i)$ is a belief-closed subset of T^*.*

---

[11]There is the question of whether we can instead begin with the weaker requirement that each $\Pi_i[c_i, s_{-i}]$ is universally measurable. We do not know. In Appendix D, we point out that our proof of continuity breaks down with this weaker assumption.

Note, if $\mathcal{T}$ induces a decomposition of $\mathcal{T}^*$ (via $(h_1, \ldots, h_{|I|})$), then both $\prod_{i \in I} h_i(T_i)$ and $\prod_{i \in I}(T_i^* \backslash h_i(T_i))$ are belief-closed subsets of T^*. Any belief-closed subset of $\mathcal{T}^*$ induces, what we will call, a separable metrizable type structure, i.e., a structure that differs from Definition 2.1 only in the fact that the type sets may not be complete. (See Lemma E.1.) So, we can view the Θ-based structure $\mathcal{T}^*$ as the union of two metrizable Θ-based structures: **the structure induced by** $\mathcal{T}$ (which corresponds to the belief-closed set $\prod_{i \in I} h_i(T_i)$) and **the difference structure** (which corresponds to the belief-closed set $\prod_{i \in I}(T_i^* \backslash h_i(T_i))$). Write

$$(\mathcal{T}^* \backslash \mathcal{T}) = (\Theta, (T_i^* \backslash h_i(T_i), \gamma_i^{\triangledown})_{i \in I}),$$

for this difference structure. (Here, $\gamma_i^{\triangledown}(t_i^*)(E^{\triangledown}) = \beta_i^*(t_i^*)(E^{\triangledown})$ for each event $E^{\triangledown}$ in $\Theta \times \prod_{j \neq i}(T_j^* \backslash h_j(T_j))$. Again, refer to Lemma E.1 for details.)

Now, we can state the result.

Lemma 5.3. *Fix Θ-based structures $\mathcal{T}$ and $\mathcal{T}^*$, so that $\mathcal{T}$ induces a decomposition of $\mathcal{T}^*$. Fix, also, a Θ-based game Γ so that $(\Gamma, \mathcal{T})$ has an equilibrium. Then, $\langle \mathcal{T}, \mathcal{T}^* \rangle$ satisfies the Equilibrium Extension Property for Γ if and only if there is an equilibrium of the difference game $(\Gamma, (\mathcal{T}^* \backslash \mathcal{T}))$.*

As a consequence of Lemma 5.3 and the Pull-Back Property, we have the following:

Proposition 5.2. *Fix Θ-based structures $\mathcal{T}$ and $\mathcal{T}^*$, so that $\mathcal{T}$ induces a decomposition of $\mathcal{T}^*$. Fix, also, a Θ-based game Γ, so that $(\Gamma, \mathcal{T})$ has an equilibrium. Then, $\langle \mathcal{T}, \mathcal{T}^* \rangle$ satisfies the Equilibrium Extension Property for Γ if and only if there is an equilibrium for the Bayesian game $(\Gamma, \mathcal{T}^*)$.*

In contrast to Proposition 5.1, Proposition 5.2 does not restrict attention to universally measurable equilibria nor does it require Γ to be compact and continuous. That said, it imposes restrictions on the players' and analyst's type structures.

Taken together, Propositions 3.1 and 5.2 say: If $\mathcal{T}$ induces a decomposition of $\mathcal{T}^*$, then $\langle \mathcal{T}, \mathcal{T}^* \rangle$ satisfies the Equilibrium Extension Property for Γ if and only if either both $(\Gamma, \mathcal{T})$ and $(\Gamma, \mathcal{T}^*)$ have an equilibrium or both $(\Gamma, \mathcal{T})$ and $(\Gamma, \mathcal{T}^*)$ do not have an equilibrium. So, here, we cannot have the Extension failure when there is an equilibrium for $(\Gamma, \mathcal{T}^*)$.

Let's now ask: Is it of interest to consider the case where $\mathcal{T}$ induces a decomposition of $\mathcal{T}^*$? We will show that there is a notable case in which $\mathcal{T}$ does induce a decomposition of $\mathcal{T}^*$. Specifically, if the analyst's structure, viz. $\mathcal{T}^*$, admits a common prior, then the players' structure must induce a decomposition of the analyst's structure.

Why is this the case? Recall, the common prior assumption (CPA) reflects the idea that differences in beliefs reflect only differences in information. That is, if an outside observer looks at the situation, he can understand the different beliefs (i.e., associated with different types) as reflecting some underlying belief, common to both players. Each type of each player reflects the conditional of this belief on certain information.

Under a common prior, what does Izzy think Joe thinks about Izzy? Can a type of Izzy consider it possible that Joe considers that type of Izzy impossible? The answer would seem to be no. In particular, this appears to require that Izzy considers it possible that Joe has learned certain information that is inconsistent with the information she herself learned. This suggests that, if a type structure satisfies the CPA, then it also satisfies a mutual absolute continuity condition—i.e., if a type t_i^* of Izzy considers a type t_j^* of Joe possible (i.e., if $\beta_i^*(t_i^*)(\Theta \times \{t_j^*\} \times T_{-i-j}^*) > 0$), then that type t_j^* of Joe also considers the given type t_i^* of Izzy possible (i.e., then $\beta_j^*(t_j^*)(\Theta \times \{t_i^*\} \times T_{-i-j}^*) > 0$). Note, here, we write T_{-i-j}^* for $\prod_{k \neq i,j} T_k^*$.

Going back to the structures $\mathcal{T}$ and $\mathcal{T}^*$, suppose the analyst's structure satisfies the common prior assumption. We have just argued that it also satisfies a mutual absolute continuity condition. Consider a type t_i^* that is not contained in the structure induced by $\mathcal{T}$. Can the type t_i^* assign strictly positive probability to a type of Joe in the structure induced by $\mathcal{T}$? No. The structure induced by $\mathcal{T}$ is a belief-closed subset. So, types in this structure cannot assign positive probability to the type t_i^*, which is what mutual absolute continuity would require. As such, the type t_i^* must assign probability one to types in (what will be) the difference structure. That is, $\mathcal{T}$ induces a decomposition of $\mathcal{T}^*$.

Let us state these facts formally: Fix a Θ-based type structure $\mathcal{T} = (\Theta, (T_i, \beta_i)_{i \in I})$. Write $[t_i]$ for the event $\Theta \times \{t_i\} \times T_{-i}$. Given a measure $\mu \in \Delta(\Theta \times T)$ with $\mu([t_i]) > 0$, write $\mu(\cdot || [t_i])$ for conditional of μ on $[t_i]$ and write $\text{marg}_{\Theta \times T_{-i}} \mu$ for the marginal of μ on $\Theta \times T_{-i}$.

Definition 5.4. *Fix a Θ-based type structure $\mathcal{T} = (\Theta, (T_i, \beta_i)_{i \in I})$. Call $\mu \in \Delta(\Theta \times T)$ a **common prior (for $\mathcal{T}$)** if $\mathcal{T}$ is countable and, for each player i and each $t_i \in T_i$,*

(i) $\mu([t_i]) > 0$,

(ii) $\beta_i(t_i) = \text{marg}_{\Theta \times T_{-i}} \mu(\cdot || [t_i])$.

*Say the structure $\mathcal{T}$ **admits a common prior** if there is a common prior for $\mathcal{T}$.*

Definition 5.5 (Stuart, 1997). *Say a Θ-based type structure $\mathcal{T} = (\Theta, (T_i, \beta_i)_{i \in I})$ is **mutually absolutely continuous** if $\mathcal{T}$ is countable and, for every pair of (distinct) players $i, j \in I$, $\beta_i(t_i)(\Theta \times \{t_j\} \times T_{-i-j}) > 0$ implies $\beta_j(t_j)(\Theta \times \{t_i\} \times T_{-i-j}) > 0$.*

Now, the connections.

Lemma 5.4. *Fix a Θ-based type structure $\mathcal{T} = (\Theta, (T_i, \beta_i)_{i \in I})$, where $\mathcal{T}$ admits a common prior. Then, $\mathcal{T}$ is mutually absolutely continuous.*

Lemma 5.5. *Fix non-redundant Θ-based structures $\mathcal{T}$ and $\mathcal{T}^*$, so that $\mathcal{T}$ can be properly embedded into $\mathcal{T}^*$ via $(h_1, \ldots, h_{|I|})$. If $\mathcal{T}^*$ is mutually absolutely continuous, then $\mathcal{T}$ induces a decomposition of $\mathcal{T}^*$ via $(h_1, \ldots, h_{|I|})$.*

Now, as a Corollary of Lemma 5.5 and Proposition 5.2, we have:

Corollary 5.3. *Fix Θ-based structures $\mathcal{T}$ and $\mathcal{T}^*$, so that $\mathcal{T}$ can be properly embedded into $\mathcal{T}^*$ and so that $\mathcal{T}^*$ satisfies mutual absolute continuity. Fix, also, a Θ-based game Γ, so that $(\Gamma, \mathcal{T})$ has an equilibrium. Then, $\langle \mathcal{T}, \mathcal{T}^* \rangle$ satisfies the Equilibrium Extension Property for Γ if and only if there is an equilibrium for the Bayesian game $(\Gamma, \mathcal{T}^*)$.*

And, as a corollary of Lemma 5.4 and Corollary 5.3, we have[12]:

Proposition 5.3. *Fix Θ-based structures $\mathcal{T}$ and $\mathcal{T}^*$, so that $\mathcal{T}$ can be properly embedded into $\mathcal{T}^*$ and so that $\mathcal{T}^*$ admits a common prior. Fix, also, a Θ-based game Γ, so that $(\Gamma, \mathcal{T})$ has an equilibrium. Then, $\langle \mathcal{T}, \mathcal{T}^* \rangle$ satisfies the Equilibrium Extension Property for Γ if and only if there is an equilibrium for the Bayesian game $(\Gamma, \mathcal{T}^*)$.*

This says that, if the analyst's structure satisfies the common prior assumption, then the only way we can have an Extension failure is if there is an equilibrium of the players' Bayesian game but not the analyst's Bayesian game.

Notice, if the analyst's structure satisfies the common prior assumption, then both the players and analyst's structure have (at most) a countable number of types. So, in this case, if Γ is compact and continuous, then both the players' and analyst's Bayesian game do have an equilibrium. (See Footnote 8.) This is a special case of Corollary 5.2.

6 Discussion

This section discusses the relationship to the literature, and further discusses the results.

a. The Context of the Game: There are two distinct views of a game. Under the first view, the game itself is a complete description of all interactions past, present, and future. See, for instance, the discussion in Kohlberg and Mertens (1986). Under the second view, it is impractical to write down "the big game." Instead, the game studied represents a snapshot of the strategic situation. This is a game-theoretic analog to Savage's (1972) Small Worlds view in decision theory.

Our position is that each of these views is of interest—both deserve to be studied. Here, we focus on the second view, where there is a history prior to the given game. As such, it seems natural to consider the case where the history influences which hierarchies of beliefs players can hold. That is, it seems natural to consider the case in which the history determines the context of the game.

In this case, two robustness questions arise. First, what if the players know more than the analyst? This is the question we focused on here, i.e., the Extension Problem. But we can also address a second question. What if the analyst rules out more hierarchies than the players? This

---

[12]Note that in the special case where Γ is compact and continuous, Corollary 5.3 and Proposition 5.3 follow directly from Proposition 5.1: All involved type sets are at most countable, and therefore all equilibria are measurable.

corresponds to the Pull-Back Problem. In this case, the analyst will not lose any prediction, but may instead introduce extraneous predictions.

b. Proof of Extension Failure: Section 4 constructs a non-pathological example of two Θ-based Bayesian games so that we cannot extend an equilibrium of the players' Bayesian game to the analyst's Bayesian game, despite the fact that there is an equilibrium of the analyst's Bayesian game. To do so, it first constructs a pathological example of a $\overline{\Theta}$-based extension failiure. In particular, it constructs a $\overline{\Theta}$-based Bayesian game $(\overline{\Gamma}, \mathcal{U}(\overline{\Theta}))$—where $\overline{\Gamma}$ is a finite $\overline{\Theta}$-based game and $\mathcal{U}(\overline{\Theta})$ is a $\overline{\Theta}$-based universal structure—so that there is no Bayesian equilibrium of $(\overline{\Gamma}, \mathcal{U}(\overline{\Theta}))$. To the best of our knowledge, this is the first such example in the literature.

The example of $(\overline{\Gamma}, \mathcal{U}(\overline{\Theta}))$ builds on a result of Hellman (2014): There exists a finite $\overline{\Theta}$-based Bayesian game $(\overline{\Gamma}, \overline{\mathcal{T}})$, so that any Bayesian equilibrium $(\overline{s}_1, \overline{s}_2)$ of $(\overline{\Gamma}, \overline{\mathcal{T}})$ has some player i with $\overline{s}_i$ that is not μ_i measurable for some $\mu_i \in \Delta(\overline{T}_i)$. There is an analogous earlier result of Simon (2003). We explicitly use Hellman's example and not Simon's. Below we review these results and why we make this choice.

First, let's start with Hellman's (2014) result—that any Bayesian equilibrium of $(\overline{\Gamma}, \overline{\mathcal{T}})$, viz. $(\overline{s}_1, \overline{s}_2)$, has some $\overline{s}_i$ that is not μ_i-measurable for a $\mu_i \in \Delta(\overline{T}_i)$. Note, the measure μ_i does not correspond to the belief of any type of $-i$ in $\overline{\mathcal{T}}$. So, we cannot conclude "if a strategy profile is a proposed Bayesian equilibrium, some type in the model will not be able to compute its expected payoffs." And, indeed, there is a Bayesian equilibrium of Hellman's example; all types can, of course, compute their expected payoffs under the Bayesian equilibrium. (Again, this is a consequence of Proposition 1 in Simon, 2003.)

Certainly, formally, non-existence in the universal Bayesian game $(\overline{\Gamma}, \mathcal{U}(\overline{\Theta}))$ is not a far leap from Hellman's result. Nonetheless, we note that Hellman does not make this leap.[13] Moreover, making the leap appears to require either a proof distinct from Hellman (2014) or a modification of Hellman's example: Note, Hellman's is about non-measurability of a strategy mapping—namely $\overline{s}_i$. To find a violation of a Bayesian equilibrium, we need non-integrability of a payoff mapping—namely $\overline{\Pi}_i[c_i, \overline{s}_{-i}]$. We modify Hellman's to have an example that satisfies the injectivity. Then we use Lemma 2.1 (plus standard properties of $\mathcal{U}(\overline{\Theta})$) to draw the connection and reach the desired conclusion.

Simon's example is also about non-measurability of a strategy mapping. We don't know if, in the context of his example, it implies non-measurability (and so non-integrability) of a payoff mapping and, likewise, we don't know if we can modify his example to obtain non-measurability of a payoff mapping. As such, we make use of Hellman's example instead.

c. Ingredients of the Extension Failure: The negative result in Section 4 makes use of a Bayesian game without an equilibrium. But, it is important to note that the result is not simply

---

[13]After learning of our result, Hellman added an example, modifying his original example. There is no Bayesian equilibrium of this modified Bayesian game. This modified example is *not* an example of a universal Bayesian game.

a corollary of the fact that there is some Bayesian game that does not have an equilibrium. In particular, we have seen an example of a (finite) Θ-based Bayesian game $(\Gamma, \mathcal{T})$ where we cannot extend an equilibrium of this game to an equilibrium of $(\Gamma, \mathcal{U}(\Theta))$, despite the fact that there is an equilibrium of this Bayesian game. Moreover, each belief closed subset of $\mathcal{U}(\Theta)$ induces a Bayesian game that does have an equilibrium.

To better understand the connection between Extension failures and non-existence, it may be useful to compare this analysis to another solution concept, namely correlated rationalizability. There are games—albeit, perhaps, pathological games—for which the set of rationalizable strategies is empty. (See Example 2 in Dufwenberg and Stegeman, 2002.) Yet, we have the following result.

> *Result:* Fix Θ-based structures, $\mathcal{T}$ and $\mathcal{T}^*$, so that $\mathcal{T}$ can be properly embedded into $\mathcal{T}^*$. Fix also a Θ-based game Γ. If the rationalizable strategies are non-empty in both the Bayesian games $(\Gamma, \mathcal{T})$ and $(\Gamma, \mathcal{T}^*)$, then $\langle \mathcal{T}, \mathcal{T}^* \rangle$ satisfies rationalizable extension and pull-back properties for Γ. [14]

As such, a non-existence example (as in Dufwenberg and Stegeman, 2002) cannot be used to get an Extension failure where the analyst's Bayesian game satisfies existence.

To sum up: Certainly, we can have an Extension failure that stems from the fact that there is a prediction associated with the players' Bayesian game but not the analyst's Bayesian game. Such an Extension failure necessarily stems from an Existence problem. But, the case of interest is the case where there is a prediction associated with the analyst's game. In this case, whether we do vs. do not have such an Extension failure depends on the particular solution concept studied. In particular, for Bayesian equilibrium there is such an Extension failure while for correlated rationalizability there is no such Extension failure—this is the case despite the fact that, for both solution concepts, there are Bayesian games that fail existence.

d. The Common Prior Assumption: Definition 5.4 states that the CPA reflects two requirements, a common prior requirement and a positivity requirement.

Consider the sets $[t_i] = \Theta \times \{t_i\} \times T_{-i}$ and note that these sets form a partition of $\Theta \times T$. Write τ_i for the subalgebra generated by this partition. Given a measure $\mu \in \Delta(\Theta \times T)$ and an event E in $\Theta \times T$, write $\mu(E, \cdot \| \tau_i)$ for a version of μ-conditional probability of E given τ_i. (Note, since the conditioning events for Izzy and Joe are distinct, the versions of conditional probability will also be distinct.) The **common prior requirement** is: There exists a measure $\mu \in \Delta(\Theta \times T)$ and a version of μ-conditional probability of E given τ_i so that, for any type t_i and any event E in $[t_i]$, $\beta_i(t_i)(\text{proj}_{\Theta \times T_{-i}} E) = \mu(E, [t_i] \| \tau_i)$. (Note, $\mu(E, \cdot \| \tau_i)$ is constant on $[t_i]$.) **Positivity** requires that, in addition, $\mu([t_i]) > 0$, for each type $t_i \in T_i$.

---

[14]A proof is available upon request. Dekel, Fudenberg and Morris (2007) show an analogous result, when the parameter and action sets are finite.

The positivity requirement is important for Proposition 5.3. To see this, return to the type structure $\mathcal{T}^*$ in Section 4 and note that this structure satisfies the common prior requirement. In particular, the measure $\mu \in \Delta(\Theta \times T^*)$ with $\mu(\{(\theta_3, t_1, t_2, t_3)\}) = 1$ is a common prior for $\mathcal{T}^*$. Of course, it is not positive. Thus, we can see that the common prior requirement alone does not suffice for Proposition 5.3. We also need the positivity requirement.

The need for the positivity requirement is important from the perspective of generalizing Proposition 5.3. In particular, if T_i is uncountably infinite, there is no probability measure that assigns strictly positive probability to each event $[t_i]$. This suggests a limitation to Proposition 5.3. Alternatively, this might suggest that other tools are needed to study the case of uncountably infinite spaces—i.e., lexicographic probability systems (Blume, Brandenburger and Dekel, 1991), conditional probability systems (Rényi, 1955), or non-standard probabilities.

There is an interesting connection to be made at the conceptual level. Does a non-positive common prior fit with the CPA? Arguably not. Recall, the idea of the CPA is that differences in probabilities only reflect differences in information. As a consequence, the only personalistic features of probability should come from informational differences. But, there may be many (regular and proper) versions of conditional probability. Given this, the common prior requirement (as specified above) need not pin down the beliefs (i.e., each $\beta_i(t_i)$). Indeed, in the example above, there are many Θ-based structures $\mathcal{T}$ corresponding to the common prior μ. In fact, choosing distinct probabilities p gives just such structures.

Appendix A   Proofs for Section 2

Lemma A.1. *Fix separable metrizable spaces $\Omega_1, \Omega_2, \Phi_1$, and Φ_2. Let $f_1 : \Omega_1 \to \Phi_1$, $f_2 : \Omega_2 \to \Phi_2$ and write $f : \Omega \to \Phi$ for the associated product map.*

(i) If f_1 and f_2 are universally measurable, then f is universally measurable.

(ii) If f is universally measurable, then f_1 and f_2 are universally measurable.

Proof. Begin with part (i). Assume f_1 and f_2 are universally measurable. Since Ω_1, Ω_2 are separable and metrizable, $\mathcal{B}(\Phi_1 \times \Phi_2) = \mathcal{B}(\Phi_1) \times \mathcal{B}(\Phi_2)$. Thus, to show f is $(\mathcal{B}_{\mathrm{UM}}(\Omega_1 \times \Omega_2), \mathcal{B}(\Phi_1 \times \Phi_2))$-measurable, it suffices to show that, for each $E_1 \times E_2 \in \mathcal{B}(\Phi_1) \times \mathcal{B}(\Phi_2)$, $f^{-1}(E_1 \times E_2) \in \mathcal{B}_{\mathrm{UM}}(\Omega_1 \times \Omega_2)$. By universal measurability of f_1, f_2, $f^{-1}(E_1 \times E_2) = f_1^{-1}(E_1) \times f_2^{-1}(E_2) \in \mathcal{B}_{\mathrm{UM}}(\Omega_1) \times \mathcal{B}_{\mathrm{UM}}(\Omega_2)$. Since $\mathcal{B}_{\mathrm{UM}}(\Omega_1) \times \mathcal{B}_{\mathrm{UM}}(\Omega_2) \subseteq \mathcal{B}_{\mathrm{UM}}(\Omega)$ (see, e.g., Fremlin, 2000, page 202), the conclusion follows.

Now assume that f is universally measurable. Fix some $\nu_1 \in \Delta(\Omega_1)$ and some $E_1 \in \mathcal{B}(\Phi_1)$. We will show that there are Borel sets $F_1, G_1 \in \mathcal{B}(\Omega_1)$ so that $F_1 \subseteq f_1^{-1}(E_1) \subseteq G_1$ and $\nu_1(F_1) = \nu_1(G_1)$. Thus, f_1 is universally measurable. (And, analogously, for f_2.)

Fix $\omega_2^* \in \Omega_2$ and define $k : \Omega_1 \to \Omega_1 \times \Omega_2$ so that $k(\omega_1) = (\omega_1, \omega_2^*)$. Certainly, k is measurable. Define μ as the image measure of ν_1 under k. Since f is universally measurable, there are

Borel sets $F, G \subseteq \mathcal{B}(\Omega_1 \times \Omega_2)$ so that $F \subseteq f^{-1}(E_1 \times \Phi_2) \subseteq G$ and $\mu(F) = \mu(G)$. Note that $f^{-1}(E_1 \times \Phi_2) = f_1^{-1}(E_1) \times \Omega_2$.

Since $\mu(\Omega_1 \times \{\omega_2^*\}) = 1$, we have that $\mu(F \cap (\Omega_1 \times \{\omega_2^*\})) = \mu(F) = \mu(G) = \mu(G \cap (\Omega_1 \times \{\omega_2^*\}))$, and $F \cap (\Omega_1 \times \{\omega_2^*\}) \subseteq f_1^{-1}(E_1) \times \{\omega_2^*\} \subseteq G \cap (\Omega_1 \times \{\omega_2^*\})$. Define $F_1 = \mathrm{proj}_{\Omega_1}(F \cap (\Omega_1 \times \{\omega_2^*\})$ and $G_1 = \mathrm{proj}_{\Omega_1}(G \cap (\Omega_1 \times \{\omega_2^*\})$. Then, $F_1, G_1 \in \mathcal{B}(\Omega_1)$ (Aliprantis and Border, 2007, Theorem 4.44 and Lemma 4.46) with $F_1 \subseteq f_1^{-1}(E_1) \subseteq G_1$. Moreover, since $F \cap (\Omega_1 \times \{\omega_2^*\}) = F_1 \times \{\omega_2^*\}$, $G \cap (\Omega_1 \times \{\omega_2^*\}) = G_1 \times \{\omega_2^*\}$, and $\mu(\Omega_1 \times \{\omega_2^*\}) = 1$, we have $\nu_1(F_1) = \mu(F_1 \times \{\omega_2^*\}) = \mu(G_1 \times \{\omega_2^*\}) = \nu_1(G_1)$. ■

Proof of Lemma 2.1 . Fix a Bayesian Game $(\Gamma, \mathcal{T})$, where Γ is injective. Fix also a strategy profile $s_{-i} : T_{-i} \to \prod_{j \in I \setminus \{i\}} \Delta(C_j)$. We will show that, for each $c_i \in C_i$ and $\mu_i \in \Delta(\Theta \times T_{-i})$, $(\mathrm{id} \times s_{-i})$ is μ_i-measurable if and only if $\Pi_i[c_i, s_{-i}]$ is μ_i-measurable.

First suppose that $(\mathrm{id} \times s_{-i})$ is μ_i-measurable. Then, using the fact that $\Pi_i[c_i, s_{-i}] = \hat{\pi}_i[c_i] \circ (\mathrm{id} \times s_{-i})$ and $\hat{\pi}_i[c_i]$ is measurable, it follows that $\Pi_i[c_i, s_{-i}]$ is μ_i-measurable.

Next suppose that $\Pi_i[c_i, s_{-i}]$ is μ_i-measurable. Let $F \in \mathcal{B}(\Theta \times \prod_{j \in I \setminus \{i\}} \Delta(C_j))$. We want to show that $(\mathrm{id} \times s_{-i})^{-1}(F) \in \mathcal{B}(\Theta \times T_{-i}; \mu_i)$. Note that $(\Pi_i[c_i, s_{-i}])^{-1}(\cdot) = (\mathrm{id} \times s_{-i})^{-1}((\hat{\pi}_i[c_i])^{-1}(\cdot))$, and therefore

$$(\mathrm{id} \times s_{-i})^{-1}(F) = (\mathrm{id} \times s_{-i})^{-1}((\hat{\pi}_i[c_i])^{-1}(\hat{\pi}_i[c_i](F))) = (\Pi_i[c_i, s_{-i}])^{-1}(\hat{\pi}_i[c_i](F)).$$

By Purves' Theorem (Purves, 1966) and the fact that $\hat{\pi}_i[c_i]$ is injective and measurable, $\hat{\pi}_i[c_i](F) \in \mathcal{B}(\mathbb{R})$. Thus, since $\Pi_i[c_i, s_{-i}]$ is μ_i-measurable, $(\mathrm{id} \times s_{-i})^{-1}(F) = (\Pi_i[c_i, s_{-i}])^{-1}(\hat{\pi}_i[c_i](F)) \in \mathcal{B}(\Theta \times T_{-i}; \mu_i)$, as required. ■

Appendix B Proofs for Sections 3 and 4

Lemma B.1. *Fix Polish spaces Ω, Ω^* and a Borel measurable mapping $f : \Omega \to \Omega^*$. Let $g : \Omega \to \mathbb{R}$ and $g^* : \Omega^* \to \mathbb{R}$ be such that $g = g^* \circ f$. For each $\mu \in \Delta(\Omega)$, if g^* is $\underline{f}(\mu)$-measurable, then g is μ-measurable.*

Proof. Assume g^* is $\underline{f}(\mu)$-measurable. Then, for $E \subseteq \mathbb{R}$ Borel, $(g^*)^{-1}(E) \in \mathcal{B}(\Omega^*; \underline{f}(\mu))$. This says that there are Borel sets $X^*, Y^* \subseteq \Omega^*$ with $X^* \subseteq (g^*)^{-1}(E) \subseteq Y^*$ and $\underline{f}(\mu)(X^*) = \underline{f}(\mu)(Y^*)$. Then, $f^{-1}(X^*), f^{-1}(Y^*)$ are Borel subsets of Ω with $f^{-1}(X^*) \subseteq f^{-1}((g^*)^{-1}(E)) \subseteq f^{-1}(Y^*)$ and $\mu(f^{-1}(X^*)) = \mu(f^{-1}(Y^*))$. Note, $f^{-1}((g^*)^{-1}(E)) = g^{-1}(E)$ so that $f^{-1}(X^*) \subseteq g^{-1}(E) \subseteq f^{-1}(Y^*)$. From this $g^{-1}(E) \in \mathcal{B}(\Omega; \mu)$, as required. ■

Corollary B.1. *Fix Polish spaces Ω, Ω^* and a Borel measurable mapping $f : \Omega \to \Omega^*$. Let $g : \Omega \to \mathbb{R}$ and $g^* : \Omega^* \to \mathbb{R}$ be bounded functions such that $g = g^* \circ f$. For each $\mu \in \Delta(\Omega)$, if g^* is $\underline{f}(\mu)$-integrable, then g is μ-integrable.*

Proof of Proposition 3.1. Begin with part (i): Fix a Bayesian equilibrium, viz. $(s_1^*, \ldots, s_{|I|}^*)$,

of $(\Gamma, \mathcal{T}^*)$. We will show that $(s_1, \ldots, s_{|I|}) = (s_1^* \circ h_1, \ldots, s_{|I|}^* \circ h_{|I|})$ is a Bayesian Equilibrium of $(\Gamma, \mathcal{T})$.

Begin with condition (i) of Definition 2.3. Fix some $t_i \in T_i$ and c_i and note that, by the fact that $(s_1^*, \ldots, s_{|I|}^*)$ is an equilibrium $\Pi_i^*[c_i, s_{-i}^*]$ is $\beta_i^*(h_i(t_i))$-integrable. Note $\beta_i^*(h_i(t_i))$ is the image measure of $\beta_i(t_i)$ under $\mathrm{id} \times h_{-i}$. Moreover, $\Pi_i^*[c_i, s_{-i}^*] \circ (\mathrm{id} \times h_{-i}) = \Pi_i[c_i, s_{-i}]$. Thus, it follows from Corollary B.1 that $\Pi_i[c_i, s_{-i}]$ is $\beta_i(t_i)$-integrable.

Now turn to Condition (ii) of Definition 2.3. Fix some type $t_i \in T_i$ and some choice c_i of the Bayesian game $(\Gamma, \mathcal{T})$. We have that

$$
\begin{aligned}
\int_{\Theta \times T_{-i}} \Pi_i[s_i^*(h_i(t_i)), s_{-i}^* \circ h_{-i}] d\beta_i(t_i) &= \int_{\Theta \times T_{-i}} \pi_i(\theta, s_i^*(h_i(t_i)), s_{-i}^*(h_{-i}(t_{-i}))) d\beta_i(t_i) \\
&= \int_{\Theta \times T_{-i}^*} \pi_i(\theta, s_i^*(h_i(t_i)), s_{-i}^*(t_{-i}^*)) d\beta_i^*(h_i(t_i)) \\
&\geq \int_{\Theta \times T_{-i}^*} \pi_i(\theta, c_i, s_{-i}^*(t_{-i}^*)) d\beta_i^*(h_i(t_i)) \\
&= \int_{\Theta \times T_{-i}} \pi_i(\theta, c_i, s_{-i}^*(h_{-i}(t_{-i}))) d\beta_i(t_i) \\
&= \int_{\Theta \times T_{-i}} \Pi_i[c_i, s_{-i}^* \circ h_{-i}] d\beta_i(t_i),
\end{aligned}
$$

where the second and fourth lines use the Change of Variables Theorem (e.g., Billingsley, 2008, Theorem 16.13) plus the fact that $(h_1, \ldots, h_{|I|})$ is a type morphism and the third line uses the fact that $(s_1^*, \ldots, s_{|I|}^*)$ is a Bayesian equilibrium of $(\Gamma, \mathcal{T}^*)$. This establishes condition (ii) of Definition 2.3.

Now turn to part (ii): Fix a universally measurable Bayesian equilibrium, viz. $(s_1^*, \ldots, s_{|I|}^*)$, of $(\Gamma, \mathcal{T}^*)$. By part (i), $(s_1, \ldots, s_{|I|}) = (s_1^* \circ h_1, \ldots, s_{|I|}^* \circ h_{|I|})$ is a Bayesian equilibrium. We will show that each s_i is μ-measurable, for each $\mu \in \Delta(T_i)$. For this, fix μ and let $\mu^* = \underline{h}_i(\mu)$. Since the Bayesian equilibrium $(s_1^*, \ldots, s_{|I|}^*)$ is universally measurable, s_i^* is μ^*-measurable. Then, by Lemma B.1, s_i is μ-measurable. ∎

Proof of Lemma 3.1. Fix a Θ-based Bayesian game $(\Gamma, \mathcal{T})$ where $\mathcal{T}$ can be embedded into $\mathcal{U}(\Theta)$. First suppose that, for each Θ-based structure $\mathcal{T}^*$, so that $\mathcal{T}$ can be embedded into $\mathcal{T}^*$, the pair $\langle \mathcal{T}, \mathcal{T}^* \rangle$ satisfies the Equilibrium Extension Property for Γ. Then certainly this is the case when $\mathcal{T}^* = \mathcal{U}(\Theta)$. We show that, if the pair $\langle \mathcal{T}, \mathcal{U}(\Theta) \rangle$ satisfies the Equilibrium Extension Property for Γ, then the pair $\langle \mathcal{T}, \mathcal{T}^* \rangle$ also satisfies the Equilibrium Extension Property for Γ, where $\mathcal{T}^*$ is some Θ-based structure so that $\mathcal{T}$ can be embedded into $\mathcal{T}^*$.

To show this, it will be useful to begin with properties of the mappings between these structures. By assumption, there exists an injective type morphism, viz. $(h_1, \ldots, h_{|I|})$, from $\mathcal{T}$ to $\mathcal{T}^*$. Since $\mathcal{U}(\Theta)$ is terminal, there is also a (not necessarily injective) type morphism $(l_1, \ldots, l_{|I|})$ from $\mathcal{T}^*$ to $\mathcal{U}(\Theta)$. Note, the map $(l_1 \circ h_1, \ldots, l_{|I|} \circ h_{|I|})$ is a type morphism from

$\mathcal{T}$ to $\mathcal{U}(\Theta) = (\Theta, (U_i, \gamma_i)_{i \in I})$. To see this, fix an event E in $\Theta \times U_{-i}$ and note that

$$\begin{aligned} \gamma_i(l_i(h_i(t_i)))(E) &= \beta_i^*(h_i(t_i))((\mathrm{id} \times l_{-i})^{-1}(E)) \\ &= \beta_i(t_i)((\mathrm{id} \times h_{-i})^{-1}((\mathrm{id} \times l_{-i})^{-1}(E))), \end{aligned}$$

where the first line uses the fact that $(l_1, \ldots, l_{|I|})$ is a type morphism from $\mathcal{T}^*$ to $\mathcal{U}(\Theta)$ and the second line uses the fact that $(h_1, \ldots, h_{|I|})$ is a type morphism from $\mathcal{T}$ to $\mathcal{T}^*$. So, $\gamma_i(l_i(h_i(t_i)))$ is the image measure of $\beta_i(t_i)$ under $(\mathrm{id} \times l_{-i}) \circ (\mathrm{id} \times h_{-i}) = \mathrm{id} \times (l_{-i} \circ h_{-i})$, as required. An implication is that $\mathcal{T}$ can be mapped to $\mathcal{U}(\Theta)$ via $(l_1 \circ h_1, \ldots, l_{|I|} \circ h_{|I|})$.

Now observe that, by assumption, we have an injective type morphism $(k_1, \ldots, k_{|I|})$ from $\mathcal{T}$ to $\mathcal{U}(\Theta)$. We also have that $(l_1 \circ h_1, \ldots, l_{|I|} \circ h_{|I|})$ is a type morphism from $\mathcal{T}$ to $\mathcal{U}(\Theta)$. Since $\mathcal{U}(\Theta)$ is non-redundant and type morphisms preserve hierarchies of beliefs, it follows that $(l_1 \circ h_1, \ldots, l_{|I|} \circ h_{|I|}) = (k_1, \ldots, k_{|I|})$, i.e., $(l_1 \circ h_1, \ldots, l_{|I|} \circ h_{|I|})$ is injective.

Let $(s_1, \ldots, s_{|I|})$ be a Bayesian equilibrium of $(\Gamma, \mathcal{T})$. Since $\langle \mathcal{T}, \mathcal{U}(\Theta) \rangle$ satisfies the Extension Property for Γ, there exists an equilibrium $(r_1, \ldots, r_{|I|})$ of $(\Gamma, \mathcal{U}(\Theta))$ so that $(s_1, \ldots, s_{|I|}) = (r_1 \circ l_1 \circ h_1, \ldots, r_{|I|} \circ l_{|I|} \circ h_{|I|})$. The Pull-Back Property (Proposition 3.1) gives that $(s_1^*, \ldots, s_{|I|}^*) = (r_1 \circ l_1, \ldots, r_{|I|} \circ l_{|I|})$ is an equilibrium of $(\Gamma, \mathcal{T}^*)$. Thus, $(s_1, \ldots, s_{|I|}) = (s_1^* \circ h_1, \ldots, s_{|I|}^* \circ h_{|I|})$, as required. ∎

Lemma B.2. *Fix some $\mu \in \Delta(\Omega)$ and some μ-measurable $f : \Omega \to \mathbb{R}$. Suppose there is $\omega^* \in \Omega$ so that $\mu(\{\omega^*\}) \in (0, 1)$. Then, f is $(\mu(\cdot | \Omega \backslash \{\omega^*\}))$-measurable.*

Proof. Fix some E Borel in $\mathbb{R}$. Since f is μ-measurable, there are Borel sets $F, G \subseteq \Omega$ with $F \subseteq f^{-1}(E) \subseteq G$ and $\mu(F) = \mu(G)$. We will show that $F \subseteq f^{-1}(E) \subseteq G$ and $\mu(F | \Omega \backslash \{\omega^*\}) = \mu(G | \Omega \backslash \{\omega^*\})$: Since $\mu(F) = \mu(G)$ and $\mu(\{\omega^*\}) > 0$, it must be that either $\omega^* \notin F \cup G$ or $\omega^* \in F \cap G$. From this the claim follows. ∎

Appendix C   Proofs for Section 5

Proof of Lemma 5.1. Fix Θ-based structures $\mathcal{T}$ and $\mathcal{T}^*$, so that $\mathcal{T}$ can be embedded into $\mathcal{T}^*$ via $(h_1, \ldots, h_{|I|})$. Then, each $h_i(T_i) \in \mathcal{B}(T_i)$ and so each $\Theta \times \prod_{j \in I \backslash \{i\}} h_j(T_j) \in \mathcal{B}(\Theta \times T_{-i})$. So, by definition of a type morphism, for any $h_i(t_i) \in h_i(T_i)$,

$$\beta_i^*(h_i(t_i))(\Theta \times \prod_{j \in I \backslash \{i\}} h_j(T_j)) = \beta_i(t_i)(\Theta \times T_{-i}) = 1,$$

as desired. ∎

Proof of Lemma 5.2. Fix Θ-based structures $\mathcal{T}$ and $\mathcal{T}^*$, so that $\mathcal{T}$ can be embedded into $\mathcal{T}^*$ via $(h_1, \ldots, h_{|I|})$. Suppose, for each i, $h_i(T_i) = T_i^*$. Write $g_i = (h_i)^{-1} : T_i^* \to T_i$. (This is well

defined since, by assumption, each h_i is bijective.) We will show that $\mathcal{T}^*$ can be embedded into $\mathcal{T}$ via $(g_1, \ldots, g_{|I|})$.

Certainly, each g_i is an injective bimeasurable map. It suffices to show that $(g_1, \ldots, g_{|I|})$ is a type morphism from $\mathcal{T}^*$ to $\mathcal{T}$: For this, fix some $g_i(t_i^*) \in T_i$ and some Borel $E_{-i} \subseteq \Theta \times T_{-i}$. Then,

$$
\begin{aligned}
\beta_i(g_i(t_i^*))(E_{-i}) &= \beta_i(g_i(t_i^*))((\text{id} \times h_{-i})^{-1}((\text{id} \times g_{-i})^{-1}(E_{-i}))) \\
&= \beta_i^*(h_i(g_i(t_i^*)))((\text{id} \times g_{-i})^{-1}(E_{-i})) \\
&= \beta_i^*(t_i^*)((\text{id} \times g_{-i})^{-1}(E_{-i}))
\end{aligned}
$$

where the first and last lines are by definition and the second line uses the fact that $(h_1, \ldots, h_{|I|})$ is a type morphism from $\mathcal{T}$ to $\mathcal{T}^*$. This establishes the desired conclusion. ∎

Appendix D   Proofs for Section 5.1

This appendix is devoted to proving Proposition 5.1. Throughout, we make use of the following notational conventions: Given sets $\Omega_1, \ldots, \Omega_{|I|}$ and some subset $K \subseteq \{1, \ldots, |I|\}$, write $\Omega_K = \prod_{k \in K} \Omega_k$ and write ω_K for a profile in Ω_K. Likewise, given maps $f_1, \ldots, f_{|I|}$, where each $f_i : \Omega_i \to \Phi_i$, write $f_K : \Omega_K \to \Phi_K$ for the associated product map.

Fix two (non-redundant) Θ-based type structures $\mathcal{T} = (\Theta, (T_i, \beta_i)_{i \in I})$ and $\mathcal{T}^* = (\Theta, (T_i^*, \beta_i^*)_{i \in I})$. Suppose, further, that $\mathcal{T}$ can be properly embedded into $\mathcal{T}^*$ via $(h_1, \ldots, h_{|I|})$, so that each $T_1^* \backslash h_1(T_1), \ldots, T_{|I|}^* \backslash h_{|I|}(T_{|I|})$ is (at most) countable (and possibly empty). By Lemma 5.2, there is some $i = 1, \ldots, |I|$, so that $T_i^* \backslash h_i(T_i)$ is non-empty. Order players so that (a) for each $i = 1, \ldots, J$, $T_i^* \backslash h_i(T_i) \neq \emptyset$ and (b) for each $i = J+1, \ldots, |I|$, $T_i^* \backslash h_i(T_i) = \emptyset$ (if $J < |I|$). For each $i = 1, \ldots, J$, write $M(i)$ for the cardinality of $T_i^* \backslash h_i(T_i)$ and $m(i)$ for some element of $T_i^* \backslash h_i(T_i)$. By assumption, $M(i)$ is (at most) countable.

Consider a Θ-based compact and continuous game $\Gamma = (\Theta, (C_i, \pi_i)_{i \in I})$. Throughout this appendix, we fix a universally measurable equilibrium of the Bayesian game $(\Gamma, \mathcal{T})$, viz. $(s_1, \ldots, s_{|I|})$. We want to show that there is a universally measurable equilibrium of the Bayesian game $(\Gamma, \mathcal{T}^*)$, viz. $(s_1^*, \ldots, s_{|I|}^*)$, with $(s_1, \ldots, s_{|I|}) = (s_1^* \circ h_1, \ldots, s_{|I|}^* \circ h_{|I|})$.

Section 5.1 gives the idea of the proof. In particular, we begin by constructing the game of complete information, namely G. The game has a finite or countable number of players, corresponding to $\bigcup_{i=1}^J T_i^* \backslash h_i(T_i)$. The choice set for a player $m(i) \in T_i^* \backslash h_i(T_i)$ is C_i. Write $\mathcal{C}_i$ for the set $[C_i]^{M(i)}$, so that $\mathcal{C} = \prod_{i=1}^J \mathcal{C}_i$ is the set of choice profiles in this game. Note, we can think of $\overrightarrow{c}_i = (c_i^1, c_i^2, \ldots) \in \mathcal{C}_i$ as a mapping $\overrightarrow{c}_i : T_i^* \backslash h_i(T_i) \to C_i$. So, when we write $\overrightarrow{c}_i(t_i^*)$ we mean the t_i^*-th component of $\overrightarrow{c}_i = (c_i^1, c_i^2, \ldots)$. Likewise, given a subset of players $K \subseteq \{1, \ldots, J\}$, we can think of the mapping $\overrightarrow{c}_K : \prod_{i \in K}(T_i^* \backslash h_i(T_i)) \to C_K$. Write $\overrightarrow{c}_K(t_K^*)$ for the profile in C_K with $\overrightarrow{c}_K(t_K^*) = (\overrightarrow{c}_i(t_i^*) : i \in K)$. Note, we endow $T_i^* \backslash h_i(T_i)$ with the discrete topology and so

the mapping $\overrightarrow{c}_i$ is continuous.

We now want to define a payoff function $u_{m(i)} : \mathcal{C} \to \mathbb{R}$ for player $m(i)$ (in the game G). To do so, it will be useful to first define auxiliary (payoff) functions for $m(i)$ that depend on subsets of players. The function $u_{m(i)}$ will be, effectively, the sum of these auxiliary functions.

Fix some player i and consider a subset K of players not containing i, i.e., some $K \subseteq \{1, \ldots, J\} \backslash \{i\}$. Write $K^c = \{1, \ldots, I\} \backslash (K \cup \{i\})$, i.e., all players that are not in $K \cup \{i\}$. Let us give the loose idea: We will construct a function $v_{m(i)}[K]$ that takes choice profiles for members of K and a choice for $m(i)$, and maps it into a payoff for player $m(i)$. When we do so, we will assume that players in K^c (if there are any) play according to the equilibrium profile. For instance, if $I = J = 3$ and $i = 1$, then we can have K be either $\emptyset$, $\{2\}$, $\{3\}$, or $\{2,3\}$. Consider the case of $K = \{2\}$. We will have $v_{m(1)}[\{2\}] : C_1 \times C_2 \to \mathbb{R}$, so that we are computing expected payoffs for $m(1)$ when types for player 2 are in $T_2^* \backslash h_2(T_2)$ and types for player 3 are in $h_3(T_3)$. Because (for this subset K) types for player 2 are in $T_2^* \backslash h_2(T_2)$, $v_{m(1)}[\{2\}]$ maps a choice for player $m(1)$ plus choices players in $T_2^* \backslash h_2(T_2)$, i.e., $C_1 \times C_2$, into a payoff. Because (for this subset K) types for player 3 are in $h_3(T_3)$, we assume they play according to the given equilibrium.

Once we have the functions $v_{m(i)}[K]$ for all subsets $K \subseteq \{1, \ldots, J\} \backslash \{i\}$, we can extend these functions to a function $u_{m(i)} : \mathcal{C} \to \mathbb{R}$. Specifically, set $u_{m(i)} = \sum_{K \subseteq J} [v_{m(i)}[K] \circ \text{proj}_{C_i \times \mathcal{C}_K}]$, where we write $\text{proj}_{C_i \times \mathcal{C}_K} : \mathcal{C} \to C_i \times \mathcal{C}_K$ for the projection map. The functions $u_{m(i)}$ are the payoff functions for the game G.

Now, let's specify the functions $v_{m(i)}[K]$. To do so, it will be useful to recall that, for each $j = 1, \ldots, |I|$, $h_j : T_j \to T_j^*$ is injective and bimeasurable. As such, we can define a bimeasurable map $g_j : h_j(T_j) \to T_j$ so that $g_j(h_j(t_j)) = t_j$. Now, fix a $K \subseteq \{1, \ldots, J\} \backslash \{i\}$. Let $v_{m(i)}[K] : C_i \times \mathcal{C}_K \to \mathbb{R}$ be such that

$$v_{m(i)}[K](c_i, \overrightarrow{c}_K) = \int_{\Theta \times \prod_{j \in K}(T_j^* \backslash h_j(T_j)) \times h_{K^c}(T_{K^c})} \pi_i(\theta, c_i, \overrightarrow{c}_K(t_K^*), s_{K^c}(g_{K^c}(t_{K^c}^*))) d\beta_i^*(m(i)).$$

(Note, if $K = \emptyset$, then we take the convention that $\Theta \times \prod_{j \in K}(T_j^* \backslash h_j(T_j)) \times h_{K^c}(T_{K^c}) = \Theta \times h_{K^c}(T_{K^c})$ so that $v_{m(i)}[K]$ reduces to a mapping from C_i to $\mathbb{R}$. If $K^c = \emptyset$, then we take the convention that $\Theta \times \prod_{j \in K}(T_j^* \backslash h_j(T_j)) \times h_{K^c}(T_{K^c}) = \Theta \times \prod_{j \in K}(T_j^* \backslash h_j(T_j))$, so that $v_{m(i)}[K]$ reduces with $s_{K^c}(g_{K^c}(t_{K^c}^*))$ no longer being a factor.)

We begin by showing that each $v_{m(i)}[K]$ is continuous. For this, we will need a mathematical result.

Lemma D.1. *Fix metrizable spaces Ω_1, Ω_2. Let $\mu \in \Delta(\Omega_2)$ and $f : \Omega_1 \times \Omega_2 \to \mathbb{R}$ be a bounded function so that each $f(\omega_1, \cdot) : \Omega_2 \to \mathbb{R}$ is μ-measurable and each $f(\cdot, \omega_2) : \Omega_1 \to \mathbb{R}$ is continuous. Define $F : \Omega_1 \to \mathbb{R}$ so that*

$$F(\omega_1) = \int_{E_2} f(\omega_1, \omega_2) d\mu,$$

where $E_2 \in \mathcal{B}(\Omega_2)$. Then, F is a bounded continuous function.

Proof. The fact that F is bounded follows directly from the fact that f is bounded and $\mu(E_2) \leq 1$. We focus on showing that F is continuous. For this, fix a sequence $(\omega_1^n : n = 1, 2, \ldots)$ contained in Ω_1 and suppose $\omega_1^n \to \omega_1^*$. To show that F is continuous, it suffices to show that $F(\omega_1^n) \to F(\omega_1^*)$.

Write $f^*(\cdot) : \Omega_2 \to \mathbb{R}$ for the ω_1^*-section of the map f. Also, for each n, write $f^n(\cdot) : \Omega_2 \to \mathbb{R}$ for the ω_1^n-section of the map f. By assumption, each of $f^*, f^1, f^2, \ldots$ is μ-measurable. Moreover, since f is bounded, f^* is bounded and the sequence $(f^n : n = 1, 2, \ldots)$ is uniformly bounded. Given this, it suffices to show that pointwise $f^n \to f^*$ (that is, $f^n(\omega_2) \to f^*(\omega_2)$, for all $\omega_2 \in \Omega_2$) If so, then, by the Dominated Convergence Theorem, $F(\omega_1^n) \to F(\omega_1^*)$. (See Aliprantis and Border, 2007, page 407.)

To show that $f^n \to f^*$: Note that $\omega_1^n \to \omega_1^*$. It follows from the fact that each $f(\cdot, \omega_2)$ is continuous that $f^n \to f^*$. ∎

Lemma D.2. *For each $m(i) \in T_i^* \backslash h_i(T_i)$ and each $K \subseteq \{1, \ldots, J\} \backslash \{i\}$, $v_{m(i)}[K] : C_i \times \mathcal{C}_K \to \mathbb{R}$ is continuous.*

Proof. Define a mapping $f_i[K] : C_i \times \mathcal{C}_K \times \Theta \times \prod_{j \in K}(T_j^* \backslash h_j(T_j)) \times h_{K^c}(T_{K^c}) \to \mathbb{R}$ so that

$$f_i[K](c_i, \overrightarrow{c}_K, \theta, t_K^*, t_{K^C}^*) = \pi_i(\theta, c_i, \overrightarrow{c}_K(t_K^*), s_{K^c}(g_{K^c}(t_{K^c}^*))).$$

Certainly, then, $f_i[K]$ is bounded. We will show that each $f_i[K](c_i, \overrightarrow{c}_K, \cdot)$ is universally measurable and each $f_i[K](\cdot, \theta, t_K^*, t_{K^C}^*)$ is continuous. Then the result follows from Lemma D.1 and the fact that

$$v_{m(i)}[K](c_i, \overrightarrow{c}_K) = \int_{\Theta \times (\prod_{j \in K}(T_j^* \backslash h_j(T_j))) \times h_{K^c}(T_{K^c})} f_i[K](c_i, \overrightarrow{c}_K, \theta, t_K^*, t_{K^C}^*) d\beta_i^*(m(i)).$$

First we show that, for each $(c_i, \overrightarrow{c}_K)$, $f_i[K](c_i, \overrightarrow{c}_K, \cdot)$ is universally measurable: Write

$$F_i[c_i, \overrightarrow{c}_K] : \Theta \times \prod_{j \in K}(T_j^* \backslash h_j(T_j)) \times h_{K^c}(T_{K^c}) \to \Theta \times \{c_i\} \times C_K \times \prod_{j \in K^C} \Delta(C_j)$$

for the mapping $(\theta, t_K^*, t_{K^c}^*) \mapsto (\theta, c_i, \overrightarrow{c}_K(t_K^*), s_{K^c}(g_{K^c}(t_{K^c}^*)))$. To show that $f_i[K](c_i, \overrightarrow{c}_K, \cdot)$ is universally measurable, it suffices to show that $F_i[c_i, \overrightarrow{c}_K]$ is universally measurable: Then $f_i[K](c_i, \overrightarrow{c}_K, \cdot) = \pi_i \circ F_i[c_i, \overrightarrow{c}_K]$ is the composite of universally measurable maps and so universally measurable.

To see that $F_i[c_i, \overrightarrow{c}_K]$ is universally measurable: Applying Lemma A.1 and the fact that each s_j is universally measurable, s_{K^c} is universally measurable. So, the restriction of s_{K^c} to the domain $h_{K^c}(T_{K^c})$ is universally measurable. Now, note that $F_i[c_i, \overrightarrow{c}_K]$ is the product of universally measurable maps, each of which has a separable metrizable domain. So, again applying Lemma A.1, $F_i[c_i, \overrightarrow{c}_K]$ is universally measurable.

Next we show that, for each $(\theta, t_K^*, t_{K^C}^*)$, $f_i[K](\cdot, \theta, t_K^*, t_{K^C}^*)$ is continuous: For this, suppose that $(c_i^n, \overrightarrow{c}_K^n) \to (c_i, \overrightarrow{c}_K)$. Then, note that $(\theta, c_i^n, \overrightarrow{c}_K^n(t_K^*), s_{K^c}(g_{K^c}(t_{K^c}^*))) \to (\theta, c_i, \overrightarrow{c}_K(t_K^*), s_{K^c}(g_{K^c}(t_{K^c}^*)))$.

So, using the continuity of π_i, $f_i[K](c_i^n, \overrightarrow{c}_K^n, \theta, t_K^*, t_{K^C}^*) \to f_i[K](c_i, \overrightarrow{c}_K, \theta, t_K^*, t_{K^C}^*)$, as required. ∎

Note, the proof of Lemma D.2 explicitly uses the fact that each s_i is universally measurable. We do not know if it would attain, if we instead assumed that each $\Pi_i[c_i, s_{-i}]$ is universally measurable.

Lemma D.3. *The map $u_{m(i)}$ is continuous.*

Proof. Note, each $\text{proj}_{C_i \times \mathcal{C}_K}$ is a continuous function. With this and Lemma D.2, each $v_{m(i)}[K] \circ \text{proj}_{C_i \times \mathcal{C}_K}$ is a continuous function. It follows that $u_{m(i)}$ is a finite sum of continuous functions and so continuous. ∎

Write $\mathcal{D}_i$ for the set $[\Delta(C_i)]^{M(i)}$ write $\overrightarrow{\sigma}_i$ for an arbitrary element of $\mathcal{D}_i$. Take $\mathcal{D} = \prod_{i=1}^{J} \mathcal{D}_i$ and write $\overrightarrow{\sigma} = (\overrightarrow{\sigma}_1, \ldots, \overrightarrow{\sigma}_j)$ for an arbitrary element of $\mathcal{D}$. For a given player $m(i)$, take $\mathcal{D}_{-m(i)}$ to be $[\Delta(C_i)]^{(M(i)-1)} \times \prod_{j \neq i} \mathcal{D}_j$ if $M(i)$ is finite and $\mathcal{D}$ otherwise. Note, if $M(i)$ is (countably) infinite $\mathcal{D}_{-m(i)} = \mathcal{D}$. An arbitrary element of $\mathcal{D}_{-m(i)}$ will be denoted as $\overrightarrow{\sigma}_{-m(i)}$.

Extend payoff functions to $u_{m(i)} : \mathcal{D} \to \mathbb{R}$ in the usual way. Note, the extended functions remain continuous. (Use, e.g., Fristedt and Gray, 1996, Theorem 20, Chapter 18 and the definition of weak convergence.)

Lemma D.4. *There exists some mixed choice equilibrium for the game G.*

Proof. For each player $m(i)$, define a best response correspondence $\text{BR}_{m(i)} : \mathcal{D}_{-m(i)} \twoheadrightarrow \Delta(C_i)$ so that

$$\text{BR}_{m(i)}(\overrightarrow{\sigma}_{-m(i)}) = \{\sigma_{m(i)} \in \arg\max u_{m(i)}(\cdot, \overrightarrow{\sigma}_{-m(i)})\}.$$

Extend this correspondence to a best response correspondence $\mathbb{BR}_{m(i)} : \mathcal{D} \twoheadrightarrow \mathcal{D}$ so that

$$\mathbb{BR}_{m(i)}(\sigma_{m(i)}, \overrightarrow{\sigma}_{-m(i)}) = \text{BR}_{m(i)}(\overrightarrow{\sigma}_{-m(i)}) \times \mathcal{D}_{-m(i)}.$$

Define $\mathbb{BR} : \mathcal{D} \twoheadrightarrow \mathcal{D}$ so that $\mathbb{BR}(\overrightarrow{\sigma}) = \bigcap_{i=1}^{J} \bigcap_{m(i)=1}^{M(i)} \mathbb{BR}_{m(i)}(\overrightarrow{\sigma})$. To show that there is a mixed strategy equilibrium of the game G, it suffices to show that there is a fixed point of $\mathbb{BR}$.

To show that there is a fixed point of $\mathbb{BR}$, we will apply the Glicksberg's (1952) Theorem. For this, it suffices to show that $\mathcal{D}$ is a non-empty, compact, convex subset of a convex Hausdorff linear topological space and that $\mathbb{BR}$ has a closed graph and is non-empty convex valued.

Note that each $\Delta(C_i)$ is a non-empty, compact, convex subset of a convex Hausdorff linear topological space. It follows that $\mathcal{D}$ satisfies the desired conditions. As such, we focus on the properties of $\mathbb{BR}$.

First, we show that $\mathbb{BR}$ has a closed graph: By Berge's Maximum Theorem (see 17.31 in Aliprantis and Border, 2007), for each $m(i)$, $\text{BR}_{m(i)}$ is compact valued and upper-hemicontinuous. It follows that $\mathbb{BR}_{m(i)}$ is a compact valued and upper-hemicontinuous correspondence to a Hausdorff space. So, applying Theorem 17.10 in Aliprantis and Border (2007), it follows that $\mathbb{BR}_{m(i)}$

has a closed graph. It now follows from Theorem 17.25 in Aliprantis and Border (2007) that $\mathbb{BR}$ has a closed graph.

Next we show that $\mathbb{BR}$ is non-empty convex valued: By Berge's Maximum Theorem (see 17.31 in Aliprantis and Border, 2007), for each $m(i)$, $\mathrm{BR}_{m(i)}$ is non-empty valued. It is standard that $\mathrm{BR}_{m(i)}$ is convex valued. (This follows from the fact that payoffs are linear in mixtures of probabilities of choices.) It follows from construction then that $\mathbb{BR}_{m(i)}$ and $\mathbb{BR}$ are non-empty valued. $\blacksquare$

In what follows, we fix strategies r_i^* of $(\Gamma, \mathcal{T}^*)$ satisfying $r_i^* \circ h_i = s_i$. Note, such strategies are well defined since h_i is injective. If $T_i^* \backslash h_i(T_i) \neq \emptyset$, then given some r_i^* we write $\overrightarrow{r}_i^*$ for $(r_i^*(1), r_i^*(2), \ldots)$, i.e., the associated element of $\mathcal{D}_i$ played by types in $T_i^* \backslash h_i(T_i)$ under r_i^*. A standard argument establishes the next remark.

Remark D.1. *Fix some $m(i) \in T_i^* \backslash h_i(T_i)$. For any $(r_1^*, \ldots, r_{|I|}^*)$ with $(r_1^* \circ h_1, \ldots, r_{|I|}^* \circ h_{|I|}) = (s_1, \ldots, s_{|I|})$,*

$$\int_{\Theta \times T_{-i}^*} \Pi_i^*[r_i^*(m(i)), r_{-i}^*] d\beta_i^*(m(i)) = u_{m(i)}(\overrightarrow{r}_1^*, \ldots, \overrightarrow{r}_j^*).$$

Conversely, given some $(\overrightarrow{\sigma}_1, \ldots, \overrightarrow{\sigma}_j) \in \mathcal{D}$, there is a unique strategy profile $(r_1^*, \ldots, r_{|I|}^*)$ with $(\overrightarrow{r}_1^*, \ldots, \overrightarrow{r}_j^*) = (\overrightarrow{\sigma}_1, \ldots, \overrightarrow{\sigma}_j)$ and $(r_1^* \circ h_1, \ldots, r_{|I|}^* \circ h_{|I|}) = (s_1, \ldots, s_{|I|})$. In this case,

$$\int_{\Theta \times T_{-i}^*} \Pi_i^*[r_i^*(m(i)), r_{-i}^*] d\beta_i^*(m(i)) = u_{m(i)}(\overrightarrow{\sigma}_1, \ldots, \overrightarrow{\sigma}_j).$$

Lemma D.5. *Let Ω, Ω^* be Polish. If $f : \Omega \to \Omega^*$ is an embedding, then f maps sets in $\mathcal{B}_{UM}(\Omega)$ to sets in $\mathcal{B}_{UM}(\Omega^*)$.*

Proof. Fix some $E \in \mathcal{B}_{\mathrm{UM}}(\Omega)$ and some $\mu^* \in \Delta(\Omega^*)$. We will show that $f(E)$ is μ^*-measurable. Note that $f(\Omega) \in \mathcal{B}(\Omega^*)$, since f is an embedding. Thus, if $\mu^*(f(\Omega)) = 0$, then $\emptyset \subseteq f(E) \subseteq f(\Omega)$ with $\mu^*(\emptyset) = \mu^*(f(\Omega)) = 0$, i.e., $f(E) \in \mathcal{B}(\Omega^*; \mu^*)$. As such, we focus on the case where $\mu^*(f(\Omega)) > 0$.

Define $\mu(G) = \frac{\mu^*(f(G))}{\mu^*(f(\Omega))}$, for each $G \in \mathcal{B}(\Omega)$. (Since f is bimeasurable, this is well-defined.) Using the fact that f is injective, this defines a probability measure $\mu \in \Delta(\Omega)$. Given that $E \in \mathcal{B}_{\mathrm{UM}}(\Omega)$, there exist $X, Y \in \mathcal{B}(\Omega)$ so that $X \subseteq E \subseteq Y$ and $\mu(X) = \mu(Y)$. Since f is bimeasurable, $f(X), f(Y) \in \mathcal{B}(\Omega^*)$. By construction, $f(X) \subseteq f(E) \subseteq f(Y)$ with $\mu^*(f(X)) = \mu^*(f(Y))$, as required. $\blacksquare$

Proof of Proposition 5.1. Fix a universally measurable equilibrium $(s_1, \ldots, s_{|I|})$ of the Bayesian game $(\Gamma, \mathcal{T})$. As above, construct the game G (based on $(s_1, \ldots, s_{|I|})$). By Lemma D.4, there exists a mixed choice profile, viz. $(\overrightarrow{\sigma}_1, \ldots, \overrightarrow{\sigma}_j)$, that is an equilibrium for the game G. Now, by Remark D.1, we can find a strategy profile $(s_1^*, \ldots, s_{|I|}^*)$ so that $(\overrightarrow{s}_1^*, \ldots, \overrightarrow{s}_j^*) = (\overrightarrow{\sigma}_1, \ldots, \overrightarrow{\sigma}_j)$ and $(s_1^* \circ h_1, \ldots, s_{|I|}^* \circ h_{|I|}) = (s_1, \ldots, s_{|I|})$. We will show that $(s_1^*, \ldots, s_{|I|}^*)$ is a universally equilibrium for the Bayesian game $(\Gamma, \mathcal{T}^*)$.

First we show that each s_i^* is universally measurable. Fix a Borel E_i in $\Delta(C_i)$ and note that

$$(s_i^*)^{-1}(E_i) = h_i((s_i)^{-1}(E_i)) \cup \{t_i^* \in T_i^* \backslash h_i(T_i) : s_i^*(t_i^*) \in E_i\}.$$

Note, since s_i is universally measurable, $(s_i)^{-1}(E_i)$ is a universally measurable set and, so, using the fact that h_i is an embedding and Lemma D.5, $h_i((s_i)^{-1}(E_i))$ is a universally measurable set. Next, notice that $T_i^* \backslash h_i(T_i)$ is countable (and possibly empty); so $\{t_i^* \in T_i^* \backslash h_i(T_i) : s_i^*(t_i^*) \in E_i\}$ is Borel. It follows that $(s_i^*)^{-1}(E_i)$ is the union of two universally measurable sets and so universally measurable.

Now we show Condition (ii) of Definition 2.3: First, fix some type $h_i(t_i) \in h_i(T_i)$. Notice that, for each $c_i \in C_i$,

$$
\begin{aligned}
\int_{\Theta \times T_{-i}^*} \pi_i(\theta, s_i^*(h_i(t_i)), s_{-i}^*(t_{-i}^*)) d\beta_i^*(h_i(t_i)) &= \int_{\Theta \times T_{-i}} \pi_i(\theta, s_i^*(h_i(t_i)), s_{-i}^*(h_{-i}(t_{-i}))) d\beta_i(t_i) \\
&\geq \int_{\Theta \times T_{-i}} \pi_i(\theta, c_i, s_{-i}^*(h_{-i}(t_{-i}))) d\beta_i(t_i) \\
&= \int_{\Theta \times T_{-i}^*} \pi_i(\theta, c_i, s_{-i}^*(t_{-i}^*)) d\beta_i^*(h_i(t_i)),
\end{aligned}
$$

where the first and last lines use the Change of Variables Theorem (e.g., Billingsley (2008, Theorem 16.13) plus the fact that h_{-i} is injective, and the second line uses the fact that the fact that $(s_1, \ldots, s_{|I|}) = (s_1^* \circ h_1, \ldots, s_{|I|}^* \circ h_{|I|})$ is an equilibrium for the Bayesian game $(\Gamma, \mathcal{T})$.

Next, fix some type $t_i^* \in T_i^* \backslash h_i(T_i)$, if one exists. Here, Condition (ii) follows from Remark D.1 and the fact that $(\overrightarrow{\sigma}_1, \ldots, \overrightarrow{\sigma}_i)$ is an equilibrium of the constructed strategic form game G.

Thus, we have that $(s_1^*, \ldots, s_{|I|}^*)$ is a universally measurable Bayesian equilibrium of $(\Gamma, \mathcal{T}^*)$. Moreover, $(s_1^* \circ h_1, \ldots, s_{|I|}^* \circ h_{|I|}) = (s_1, \ldots, s_{|I|})$, as required. ∎

Appendix E Proofs for Section 5.2

For the purpose of this Appendix, we will need to extend the notion of a type structure. Call $\mathcal{T} = (\Theta, (C_i, \pi_i)_{i \in I})$ a **Θ-based separable metrizable type structure** if is satisfies the conditions of Definition 2.1, with the exception that each T_i may only be a separable metrizable set. The definitions in Sections 2-3 apply by replacing a Θ-based type structure with a Θ-based separable metrizable type structure.

Lemma E.1. *Fix a Θ-based type structure $\mathcal{T} = (\Theta, (T_i, \beta_i)_{i \in I})$. Let $\prod_{i=1}^{|I|} \overline{T}_i$ be a belief-closed subset of T. Then, there is a Θ-based separable metrizable type structure*

$$\overline{\mathcal{T}} = (\Theta, (\overline{T}_i, \overline{\beta}_i)_{i \in I}),$$

where, for each $t_i \in \overline{T}_i$ and each event E_{-i} in $\Theta \times \overline{T}_{-i}$, $\overline{\beta}_i(t_i)(E_{-i}) = \beta_i(t_i)(E_{-i})$.

Proof. Since we endow each $\overline{T}_i$ with the relative topology, we have that each $\overline{T}_i$ is separable metrizable. Also note that $\overline{\beta}_i(t_i)$ is indeed a probability measure on $\Theta \times \overline{T}_{-i}$. To see this, recall that $\prod_{i=1}^{|I|} \overline{T}_i$ is a belief-closed subset of T, so that each $\overline{T}_{-i}$ is Borel in T_{-i} with $\beta_i(t_i)(\Theta \times \overline{T}_{-i}) = 1$. So, if E_{-i} is an event in $\Theta \times \overline{T}_{-i}$, it is also an event in $\Theta \times T_{-i}$ and $\overline{\beta}_i(t_i)$ forms a probability measure.

Finally, we show that each $\overline{\beta}_i$ is measurable. Fix some F Borel in $\Delta(\Theta \times \overline{T}_{-i})$. Define $H \subseteq \Delta(\Theta \times T_{-i})$ so that $\nu \in H$ if and only if there exists some $\mu \in F$ so that $\mu(E_{-i}) = \nu(E_{-i})$ for each event E_{-i} in $\Theta \times \overline{T}_{-i}$. It follows from Lemma 14.4 in Aliprantis and Border (2007) that H is Borel in $\Delta(\Theta \times T_{-i})$. It is immediate from the construction that $(\overline{\beta}_i)^{-1}(F) = (\beta_i)^{-1}(H) \cap \overline{T}_i$. So, using the fact that β_i is measurable, $(\overline{\beta}_i)^{-1}(F)$ is measurable, as required. ∎

Take the Θ-based type structure $\mathcal{T} = (\Theta, (T_i, \beta_i)_{i \in I})$ and the constructed separable metrizable structure, viz. $\overline{\mathcal{T}} = (\Theta, (\overline{T}_i, \overline{\beta}_i)_{i \in I})$, from Lemma E.1. Write $(\overline{\mathrm{id}}_i : \overline{T}_i \to T_i)$ for the identity map. Then, $\overline{\mathcal{T}}$ can be embedded into $\mathcal{T}$ via $(\overline{\mathrm{id}}_1, \ldots, \overline{\mathrm{id}}_{|I|})$.

Lemma E.2. *Fix a Θ-based Bayesian games $(\Gamma, \mathcal{T})$ and $(\Gamma, \mathcal{T}^*)$ where*

(i) *$\mathcal{T} = (\Theta, (T_i, \beta_i)_{i \in I})$ is a separable metrizable type structure and*

(ii) *$\mathcal{T}$ can be embedded in $\mathcal{T}^* = (\Theta, (T_i^*, \beta_i^*)_{i \in I})$ via $(h_1, \ldots, h_{|I|})$.*

Let $(s_1, \ldots, s_{|I|})$ (resp. $(s_1^*, \ldots, s_{|I|}^*)$) be a strategy profile of $(\Gamma, \mathcal{T})$ (resp. $(\Gamma, \mathcal{T}^*)$) so that $s_i = s_i^* \circ h_i$ for each i. If $\Pi_i[c_i, s_{-i}]$ is $\beta_i(t_i)$-measurable, then $\Pi_i^*[c_i, s_{-i}^*]$ is $\beta_i^*(h_i(t_i))$ measurable.

Proof. Fix some $h_i(t_i) \in h_i(T_i) \subseteq T_i^*$. Fix also some $c_i \in C_i$ and some Borel $E \subseteq \mathbb{R}$. We will show that $(\Pi_i^*[c_i, s_{-i}^*])^{-1}(E) \in \mathcal{B}(\Theta \times T_{-i}^*; \beta_i^*(h_i(t_i)))$.

By assumption, $(\Pi_i[c_i, s_{-i}])^{-1}(E)$ is in $\mathcal{B}(\Theta \times T_{-i}; \beta_i(t_i))$. That is, there exists $F_{-i}, G_{-i} \in \mathcal{B}(\Theta \times T_{-i})$ so that

$$F_{-i} \subseteq (\Pi_i[c_i, s_{-i}])^{-1}(E) \subseteq G_{-i}$$

and $\beta_i(t_i)(F_{-i}) = \beta_i(t_i)(G_{-i})$. Since $(\mathrm{id} \times h_{-i})$ is bimeasurable, $(\mathrm{id} \times h_{-i})(F_{-i}), (\mathrm{id} \times h_{-i})(G_{-i}) \in \mathcal{B}(\Theta \times T_{-i}^*)$. Moreover,

$$(\mathrm{id} \times h_{-i})(F_{-i}) \subseteq (\mathrm{id} \times h_{-i})((\Pi_i[c_i, s_{-i}])^{-1}(E)) \subseteq (\mathrm{id} \times h_{-i})(G_{-i})$$

and, using the fact that $(\mathrm{id} \times h_{-i})$ is injective,

$$\beta_i^*(h_i(t_i))((\mathrm{id} \times h_{-i})(F_{-i})) = \beta_i(t_i)(F_{-i}) = \beta_i(t_i)(G_{-i}) = \beta_i^*(h_i(t_i))((\mathrm{id} \times h_{-i})(G_{-i})).$$

Now notice that

$$(\mathrm{id} \times h_{-i})((\Pi_i[c_i, s_{-i}])^{-1}(E)) = (\Pi_i^*[c_i, s_{-i}^*])^{-1}(E) \cap (\Theta \times h_{-i}(T_{-i})).$$

This allows us to conclude that $(\Pi_i^*[c_i, s_{-i}^*])^{-1}(E) \cap (\Theta \times h_{-i}(T_{-i}))$ is $\beta_i^*(h_i(t_i))$ measurable.

Since $(\Pi_i^*[c_i, s_{-i}^*])^{-1}(E) \cap (\Theta \times h_{-i}(T_{-i}))$ is $\beta_i^*(h_i(t_i))$ measurable, there exists F_{-i}^*, G_{-i}^* in $\mathcal{B}(\Theta \times T_{-i}^*)$ with

$$F_{-i}^* \subseteq (\Pi_i^*[c_i, s_{-i}^*])^{-1}(E) \cap (\Theta \times h_{-i}(T_{-i})) \subseteq G_{-i}^*$$

and $\beta_i^*(h_i(t_i))(F_{-i}^*) = \beta_i^*(h_i(t_i))(G_{-i}^*)$. Take $H_{-i}^* = \Theta \times (T_{-i}^* \backslash h_{-i}(T_{-i})))$. Since (id $\times h_{-i}$) is bimeasurable, $G_{-i}^* \cup H_{-i}^*$ is Borel. Thus,

$$F_{-i}^* \subseteq (\Pi_i^*[c_i, s_{-i}^*])^{-1}(E) \subseteq G_{-i}^* \cup H_{-i}^*.$$

Moreover, since H_{-i}^* is $\beta_i^*(h_i(t_i))$-null, $\beta_i^*(h_i(t_i))(F_{-i}^*) = \beta_i^*(h_i(t_i))(G_{-i}^* \cup H_{-i}^*)$. ∎

Proof of Lemma 5.3. Suppose $\mathcal{T}$ induces a decomposition of $\mathcal{T}^*$ via $(h_1, \ldots, h_{|I|})$. Then, both $\prod_{i \in I} h_i(T_i)$ and $\prod_{i \in I}(T_i^* \backslash h_i(T_i))$ are belief-closed subsets of T^*. So, using Lemma E.1, each of these induce a Θ-based separable metrizable type structure. Write

$$\mathcal{T}(h_1, \ldots, h_{|I|}) = (\Theta, (h_i(T_i), \overline{\beta}_i)_{i \in I})$$

for the structure induced by $\mathcal{T}$, and write

$$(\mathcal{T}^* \backslash \mathcal{T}) = (\Theta, (T_i^* \backslash h_i(T_i), \overline{\beta}_i^{\triangledown})_{i \in I})$$

for the difference structure.

Fix a Θ-based game Γ and an equilibrium $(s_1, \ldots, s_{|I|})$ for the Bayesian Game $(\Gamma, \mathcal{T})$. Suppose there exists an equilibrium for the difference game $(\Gamma, (\mathcal{T}^* \backslash \mathcal{T}))$, viz. $(s_1^{\triangledown}, \ldots, s_{|I|}^{\triangledown})$. Construct a strategy, viz. s_i^*, for $(\Gamma, \mathcal{T}^*)$, as follows. For each $t_i \in T_i$, let $s_i^*(h_i(t_i)) = s_i(t_i)$. (This is well-defined since each h_i is injective.) For each $t_i^* \in T_i^* \backslash h_i(T_i)$, let $s_i^*(t_i^*) = s_i^{\triangledown}(t_i^*)$. We now show that the constructed $(s_1^*, \ldots, s_{|I|}^*)$ is a Bayesian equilibrium for $(\Gamma, \mathcal{T}^*)$.

Condition (i) follows from Lemma E.2. Thus we focus on Condition (ii).

First, fix a type $h_i(t_i) \in h_i(T_i)$. For an action $c_i \in C_i$, the Change of Variables Theorem (e.g., Billingsley, 2008, Theorem 16.13) gives that

$$\int_{\Theta \times T_{-i}^*} \pi_i(\theta, c_i, s_{-i}^*(t_{-i}^*)) d\beta_i^*(h_i(t_i)) = \int_{\Theta \times T_{-i}} \pi_i(\theta, c_i, s_{-i}(t_{-i})) d\beta_i(t_i).$$

So, using the fact that $(s_1, \ldots, s_{|I|})$ is a Bayesian Equilibrium of $(\Gamma, \mathcal{T})$,

$$\int_{\Theta \times T_{-i}^*} \pi_i(\theta, s_i^*(h_i(t_i)), s_{-i}^*(t_{-i}^*)) d\beta_i^*(h_i(t_i)) \geq \int_{\Theta \times T_{-i}^*} \pi_i(\theta, c_i, s_{-i}^*(t_{-i}^*)) d\beta_i^*(h_i(t_i)), \quad (1)$$

for all c_i.

Likewise, given a type $t_i^* \in T_i^* \backslash h_i(T_i)$ and a choice $c_i \in C_i$,

$$\int_{\Theta \times T_{-i}^*} \pi_i(\theta, c_i, s_{-i}^*(t_{-i}^*)) d\beta_i^*(t_i^*) = \int_{\Theta \times \prod_{j \neq i} T_j^* \backslash h_j(T_j)} \pi_i(\theta, c_i, s_{-i}^\triangledown(t_{-i}^*)) d\beta_i^\triangledown(t_i^*).$$

So, using the fact that $(s_1^\triangledown, \ldots, s_{|I|}^\triangledown)$ is a Bayesian Equilibrium of $(\Gamma, \mathcal{T}^* \backslash \mathcal{T})$,

$$\int_{\Theta \times T_{-i}^*} \pi_i(\theta, s_i^*(t_i^*), s_{-i}^*(t_{-i}^*)) d\beta_i^*(t_i^*) \geq \int_{\Theta \times T_{-i}^*} \pi_i(\theta, c_i, s_{-i}^*(t_{-i}^*)) d\beta_i^*(t_i^*), \tag{2}$$

for all strategies c_i.

Taking Equations 1-2, $(s_1^*, \ldots, s_{|I|}^*)$ is an equilibrium for the Bayesian game $(\Gamma, \mathcal{T}^*)$. The converse follows immediately from the Pull-Back Property (Proposition 3.1). ■

Proof of Proposition 5.2. If $\langle \mathcal{T}, \mathcal{T}^* \rangle$ satisfies the Extension Property for Γ, then it is immediate that there is an equilibrium for the Bayesian game $(\Gamma, \mathcal{T}^*)$. Conversely, suppose there is an equilibrium for the Bayesian game $(\Gamma, \mathcal{T}^*)$. By the Pull-Back Property (Proposition 3.1), there is an equilibrium for the difference game $(\Gamma, (\mathcal{T}^* \backslash \mathcal{T}))$. Now, using Lemma 5.3, $\langle \mathcal{T}, \mathcal{T}^* \rangle$ satisfies the Equilibrium Extension Property for Γ. ■

Proof of Lemma 5.4. Now, let μ be a common prior for $\mathcal{T}$. Fix distinct players i and j and note that
$$\beta_i(t_i)(\Theta \times \{t_j\} \times T_{-i-j}) = \frac{\mu(\Theta \times \{t_i\} \times \{t_j\} \times T_{-i-j})}{\mu(\Theta \times \{t_i\} \times T_{-i-j})}.$$

So, $\beta_i(t_i)(\Theta \times \{t_j\} \times T_{-i-j}) > 0$ if and only if $\mu(\Theta \times \{t_i\} \times \{t_j\} \times T_{-i-j}) > 0$. But, an analogous argument for j gives that $\beta_j(t_j)(\Theta \times \{t_i\} \times T_{-i-j}) > 0$ if and only if $\mu(\Theta \times \{t_i\} \times \{t_j\} \times T_{-i-j}) > 0$. This establishes the result. ■

Proof of Lemma 5.5. By Lemma 5.2, there exists some i with $T_i^* \backslash h_i(T_i)$ non-empty. In particular, fix $t_i^* \in T_i^* \backslash h_i(T_i)$. Recall, since $\mathcal{T}^*$ is mutually absolutely continuous, it is countable. As such, for each player $j \neq i$, we can find some $t_j^* \in T_j^*$ with $\beta_i^*(t_i^*)(\Theta \times \{t_j^*\} \times T_{-i-j}^*) > 0$. Again using the fact that $\mathcal{T}^*$ is mutually absolutely continuous, we also have that, for each such t_j^*, $\beta_j^*(t_j^*)(\Theta \times \{t_i^*\} \times T_{-i-j}^*) > 0$. This implies that $t_j^* \in T_j^* \backslash h_j(T_j)$. (If $t_j^* \in h_j(T_j)$, then there is some $t_j \in T_j$ with $\beta_j(t_j)(\Theta \times (h_i)^{-1}(\{t_i^*\}) \times T_{-i-j}) > 0$, contradicting the fact that $(h_i)^{-1}(\{t_i^*\}) = \emptyset$.)

Now, note that, since each h_j is bimeasurable, each $T_j^* \backslash h_j(T_j)$ is Borel. So, for each j,

$\beta_i^*(t_i^*)(\Theta \times (T_j^* \backslash h_j(T_j)) \times T_{-i-j}^*) = 1$. Since this holds for each $j \neq i$, we have

$$
\begin{aligned}
1 &= \beta_i^*(t_i^*)(\Theta \times \cap_{j \neq i}((T_j^* \backslash h_j(T_j)) \times T_{-i-j}^*)) \\
&= \beta_i^*(t_i^*)(\Theta \times \prod_{j \neq i}(T_j^* \backslash h_j(T_j))),
\end{aligned}
$$

as required.

Finally, note that we showed that, for each $j \neq i$, $T_j^* \backslash h_j(T_j)$ is non-empty. So, applying the same argument to each $t_j^* \in T_j^* \backslash h_j(T_j)$, we get the desired result. ∎

References

Aliprantis, C.D. and K.C. Border. 2007. *Infinite dimensional analysis: a hitchhiker's guide.* Springer Verlag.

Aumann, R.J. 1987. "Correlated equilibrium as an expression of Bayesian rationality." *Econometrica: Journal of the Econometric Society* pp. 1–18.

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

Billingsley, P. 2008. *Probability and measure.* Wiley-India.

Blume, L., A. Brandenburger and E. Dekel. 1991. "Lexicographic probabilities and choice under uncertainty." *Econometrica* 59(1):61–79.

Bogachev, V.I. 2006. *Measure Theory.* Vol. 1 Springer Verlag.

Brandenburger, A. 2003. "On the Existence of a "Complete" Possibility Structure." *Cognitive Processes and Economic Behaviour* p. 30.

Brandenburger, A., A. Friedenberg and H.J. Keisler. 2008. "Admissibility in Games." *Econometrica* 76(2):307.

Brandenburger, A. and E. Dekel. 1993. "Hierarchies of Beliefs and Common Knowledge." *Journal of Economic Theory* 59:189–189.

Dekel, E., D. Fudenberg and S. Morris. 2007. "Interim correlated rationalizability." *Theoretical Economics* 2(1):15–40.

Dufwenberg, M. and M. Stegeman. 2002. "Existence and uniqueness of maximal reductions under iterated strict dominance." *Econometrica* 70(5):2007–2023.

Ely, J.C. and M. Peski. 2006. "Hierarchies of belief and interim rationalizability." *Theoretical Economics* 1(1):19–65.

Fremlin, D.H. 2000. *Measure theory.* Vol. 4.

Friedenberg, A. 2010. "When do Type Structures Contain all Hierarchies of Beliefs?" *Games and Economic Behavior* 68(1):108–129.

Fristedt, B.E. and L.F. Gray. 1996. *A modern approach to probability theory.* Birkhäuser Boston.

Glicksberg, IL. 1952. "A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points." *Proceedings of the American Mathematical Society* pp. 170–174.

Heifetz, A. and D. Samet. 1998. "Topology-free typology of beliefs." *Journal of Economic Theory* 82(2):324–341.

Hellman, Z. 2014. "A Game with No Approximate Bayesian Equilibrium." *Journal of Economic Theory* 153:138–151.

Kohlberg, E. and J.F. Mertens. 1986. "On the Strategic Stability of Equilibria." *Econometrica* 54(5):1003–1037.

Liu, Q. 2009. "On redundant types and Bayesian formulation of incomplete information." *Journal of Economic Theory* 144(5):2115–2145.

Meier, M. 2012. "An infinitary probability logic for type spaces." *Israel Journal of Mathematics* .

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.

Morris, S. and H.S. Shin. 2003. Global games: theory and applications. In *advances in Economics and Econometrics: theory and applications, Eighth world Congress.* Vol. 1 pp. 56–114.

Peleg, B. 1969. "Equilibrium points for games with infinitely many players." *Journal of the London Mathematical Society* 1(1):292.

Purves, R. 1966. "Bimeasurable functions." *Fundam. Math.* 58:149–157.

Rényi, A. 1955. "On a new axiomatic theory of probability." *Acta Mathematica Hungarica* 6(3):285–335.

Sadzik, T. 2011. "Beliefs Revealed in Bayesian-Nash Equilibrium." *New York University* .

Savage, L.J. 1972. *The foundations of statistics.* Dover Pubns.

Simon, R.S. 2003. "Games of incomplete information, ergodic theory, and the measurability of equilibria." *Israel journal of mathematics* 138(1):73–92.

Stuart, H.W. 1997. "Common Belief of Rationality in the Finitely Repeated Prisoners' Dilemma." *Games and Economic Behavior* 19(1):133–143.

Takahashi, S. 2009. private communication.