

Epistemic equivalence of extended belief hierarchies*

ELIAS TSAKAS[†]

Department of Economics, Maastricht University

February 26, 2014

Abstract

In this paper, we introduce a notion of epistemic equivalence between hierarchies of conditional beliefs and hierarchies of lexicographic beliefs, thus extending the standard equivalence results of Halpern (2010) and Brandenburger et al. (2007) to an interactive setting, and we show that there is a Borel surjective function, mapping each conditional belief hierarchy to its epistemically equivalent lexicographic belief hierarchy. Then, using our equivalence result we construct a terminal type space model for lexicographic belief hierarchies. Finally, we show that whenever we restrict attention to full-support beliefs, epistemic equivalence between a lexicographic belief hierarchy and a conditional belief hierarchy implies that an arbitrary Borel event is commonly assumed under the lexicographic belief hierarchy if and only if it is commonly strongly believed under the conditional belief hierarchy. This is the first result in the literature directly linking common assumption in rationality (Brandenburger et al., 2008) with common strong belief in rationality (Battigalli and Siniscalchi, 2002).

KEYWORDS: Epistemic game theory, conditional belief hierarchies, lexicographic belief hierarchies, type spaces, epistemic equivalence, common strong belief, common assumption.

JEL CLASSIFICATION: C70, D80, D81, D82.

*I am extremely indebted to Amanda Friedenberg for numerous conversations on this paper. I would also like to thank Adam Brandenburger, Willemien Kets, Andrés Perea, two anonymous referees and an associate editor for very helpful comments, as well as Christian Bach, Erik Balder, Matthew Embrey, Dimitris Georgiou, Friederike Mengel, Jeff Steif, Mark Voorneveld and the seminar participants at Maastricht University for useful discussions. I would also like to thank the Olin School of Business, at Washington University in St. Louis for its hospitality while working on the first draft of this paper. Financial support from the Marie Curie Fellowship (PIEF-GA-2009-237614) is gratefully acknowledged.

[†]P.O. Box 616, 6200 MD, Maastricht, The Netherlands; Tel: +31-43-38 83649; Fax: +31-43-38 84878; E-mail: e.tsakas@maastrichtuniversity.nl

1. Introduction

A belief hierarchy describes an agents's beliefs, beliefs about every other agent's beliefs, and so on. Belief hierarchies are an integral part of modern economic theory, often used for analyzing games with incomplete information (Harsanyi, 1967-68), as well as for providing epistemic characterizations for several solution concepts, such as rationalizability (Brandenburger and Dekel, 1987; Tan and Werlang, 1988), Nash equilibrium (Aumann and Brandenburger, 1995) and correlated equilibrium (Aumann, 1987), just to mention a few.¹

A well-known problem of standard belief hierarchies is that they fail to capture conditional beliefs given zero probability events, and therefore they are not sufficiently rich to characterize solution concepts, such as iterated admissibility in normal form games or rationalizability in extensive form games, where unlikely yet possible events play a significant role. This difficulty has been circumvented in the literature by extending the notion of beliefs in two different ways that were developed independently.

- According to the first approach, beliefs are captured by a *lexicographic probability system (LPS)*, which consists of a sequence of Borel probability measures, else called theories (Blume et al., 1991a). The primary theory coincides with the standard beliefs, the secondary theory captures the beliefs once the agent has for some reason discarded the primary theory, and so on. Extending this construction to an interactive setting gives rise to a hierarchy of lexicographic beliefs (\mathcal{L} -hierarchy): The first order lexicographic beliefs consist of an LPS over the underlying space of uncertainty, the second order lexicographic beliefs consist of an LPS over the underlying space of uncertainty and the opponents' first order lexicographic beliefs, and so on. Hierarchies of lexicographic beliefs have been used to epistemically characterize several solution concepts in normal form games, such as iterated admissibility (Brandenburger et al., 2008), self-admissible sets (Brandenburger and Friedenberg, 2010), perfect equilibrium and proper equilibrium in two-players normal form games (Blume et al., 1991b) and proper rationalizability (Asheim, 2001; Perea, 2011).
- According to the second approach, beliefs are captured by a *conditional probability system (CPS)*, which consists of a collection of conditioning events and a conditional Borel probability measure given each conditioning event, in a way such that Bayes rule is satisfied whenever possible. In dynamic games, a conditioning event typically corresponds to an information set. Extending this idea to an interactive setting induces a hierarchy of conditional beliefs

¹For an overview of the epistemic game theory literature we refer to the textbook by Perea (2012) or the review article by Brandenburger (2007).

(\mathcal{C} -hierarchy): The first order conditional beliefs are described by a CPS over the underlying space of uncertainty, the second order conditional beliefs consist of a CPS over the underlying space of uncertainty and the opponents' first order conditional beliefs, and so on. Conditional belief hierarchies have been widely used to characterize solution concepts in dynamic games, such as extensive form rationalizability (Battigalli and Siniscalchi, 2002) and extensive form best response sets (Battigalli and Friedenberg, 2012).

Several authors have studied the relationship between the two models (Brandenburger et al., 2007; Halpern, 2010). As it turns out, the two approaches are epistemically equivalent, in the sense that there exists a surjective mapping from the space of CPS's onto the space of LPS's (Brandenburger et al., 2007).²

In this paper, we extend this idea to an interactive setting, by introducing a notion of epistemic equivalence between \mathcal{L} -hierarchies and \mathcal{C} -hierarchies, thus providing a stepping stone for understanding the relationship between solution concepts whose epistemic characterizations use different types of belief hierarchies. The importance of establishing a notion of epistemic equivalence has been already pointed out in a different context (Brandenburger and Friedenberg, 2010).

Our extension is far from trivial, as the previously defined notion of epistemic equivalence relates CPS's and LPS's that are defined on the same space. However, second order conditional beliefs are described by a CPS over the underlying space of uncertainty and the opponents' first order conditional beliefs, whereas second order lexicographic beliefs are described by an LPS over the underlying space of uncertainty and the opponents' first order lexicographic beliefs. Thus, in order to introduce a notion of epistemic equivalence between second order beliefs, we first need to translate each Borel event in the space of first order lexicographic beliefs to a Borel event in the space of first order conditional beliefs. In fact, we do this by showing that the surjective function that maps CPS's onto LPS's is Borel measurable (Lemma 2). Then, second order conditional beliefs are mapped surjectively onto second order lexicographic beliefs via a Borel function, which allows us in turn to define epistemic equivalence between third order beliefs. Continuing inductively, we show that there is a Borel surjective function, mapping each conditional belief hierarchy to a lexicographic belief hierarchy (Theorem 1).

Using our main equivalence result, together with existence of a universal type space for \mathcal{C} -hierarchies (Battigalli and Siniscalchi, 1999, Prop. 2), we indirectly construct a terminal type space model for lexicographic belief hierarchies, i.e., an LPS-based type space model that induces all \mathcal{L} -hierarchies (Theorem 3). To the best of our knowledge, this is the first such result in the literature, and it provides a Bayesian foundation for hierarchies of lexicographic beliefs.

²For a precise definition of epistemic equivalence, see Definition 5.

The natural analogue of probability-1 belief in a CPS is strong belief (Battigalli and Siniscalchi, 1999), while in an LPS the corresponding notion is assumption (Brandenburger et al., 2008). One natural question arising then is whether our concept of epistemic equivalence also implies equivalence between common strong belief and common assumption, i.e., if a Borel event is commonly strongly believed under a conditional belief hierarchy, is it also commonly assumed under the epistemically equivalent lexicographic belief hierarchy, and vice versa? Brandenburger et al. (2007) have already shown that if we restrict attention to full-support beliefs in a single-agent framework, a Borel event is strongly believed under a CPS if and only if it is assumed under the epistemically equivalent LPS. Generalizing this result to our interactive setting, we prove that if beliefs are full-support, a Borel event is indeed commonly strongly believed under a \mathcal{C} -hierarchy if and only if it is commonly assumed under the epistemically equivalent \mathcal{L} -hierarchy (Proposition 4). Notice that this result applies not only to events in the underlying space of uncertainty, but also to events that involve beliefs, such as rationality, implying that it has interesting implications for solution concepts in games, e.g., for clarifying the relationship between common assumption in rationality (Brandenburger et al., 2008) and common strong belief in rationality (Battigalli and Siniscalchi, 2002).

The paper is structured as follows: Section 2 contains some necessary mathematical preliminaries; Section 3 introduces the notions of conditional probability systems and lexicographic probability systems, and presents the existing notion of epistemic equivalence between the two; Section 4 defines conditional belief hierarchies and lexicographic belief hierarchies; Section 5 presents our main results; Section 6 concludes with a discussion. All proofs are relegated to the Appendix.

2. Preliminaries

We begin with some definitions and the basic notation.³ A topological space is called Polish whenever it is separable and completely metrizable. Examples of Polish spaces include countable sets with the discrete topology and \mathbb{R}^n with the usual topology. A closed subset of a Polish space, the countable product of Polish spaces and the topological sum of countably many Polish spaces are also Polish.

Let X be a Polish space, with the Borel σ -algebra \mathcal{F} , and denote the space of Borel probability measures on X by $\Delta(X)$. As usual, endow $\Delta(X)$ with the topology of weak convergence, which is the coarsest topology that makes $\mu \mapsto \int f d\mu$ continuous for every bounded and continuous real-valued f . If X is Polish, $\Delta(X)$ is also Polish. For each $\mu \in \Delta(X)$, let $\text{Supp}(\mu)$ denote the support of μ , i.e., the smallest closed subset of X that receives probability 1 by μ . Recall that the support is unique if

³For a more detailed presentation of the following concepts we refer to standard textbooks, such as Aliprantis and Border (1994) or Kechris (1995).

X is separable and metrizable.

3. Extended probability systems

3.1. Lexicographic probability systems

Let $\mathcal{N}_n(X) := \prod_{k=1}^n \Delta(X)$ denote the space of sequences of Borel probability measures of length n together with the product topology of weak convergence.⁴ As a finite product of Polish spaces, $\mathcal{N}_n(X)$ is Polish. Let $\mathcal{N}(X) := \bigoplus_{n=1}^{\infty} \mathcal{N}_n(X)$ be the topological sum of all $\mathcal{N}_n(X)$, i.e., it is the space of finite sequences of Borel probability measures. Since $\mathcal{N}(X)$ is the topological sum of countably many Polish spaces, it is Polish too. For an arbitrary $\tilde{\mu} = (\mu^1, \dots, \mu^n) \in \mathcal{N}_n(X)$, let $\text{Supp}(\tilde{\mu}) := \bigcup_{m=1}^n \text{Supp}(\mu^m)$ denote the support of $\tilde{\mu}$. Let $\mathcal{N}^+(X)$ denote the space of full-support sequences, i.e., we write $\tilde{\mu} \in \mathcal{N}^+(X)$ whenever it is the case that $\text{Supp}(\tilde{\mu}) = X$.

Definition 1. A *lexicographic probability system (LPS)* over a measurable space X is a finite sequence of Borel probability measures $\tilde{\mu} := (\mu^1, \dots, \mu^n) \in \mathcal{N}_n(X)$ for some positive integer n , such that

(L_1) there exist Borel events $A_1, \dots, A_n \subseteq X$ with $\mu^m(A_m) = 1$, and $\mu^\ell(A_m) = 0$ for all $\ell \neq m$.

We retain the usual interpretation of an LPS. Accordingly, μ^1 is the individual's primary theory over X , μ^2 the secondary theory, μ^3 the tertiary theory, and so on. Intuitively, the primary theory, which is the most important one, describes the agent's usual beliefs about X . The secondary theory describes the agent's beliefs once she has decided for some reason to discard her primary theory. Likewise, the tertiary theory describes the agent's beliefs once she has discarded both the primary and the secondary theory, and so on. Obviously, standard beliefs are a special case of an LPS with length 1. The notion of the LPS was first introduced for finite spaces by Blume et al. (1991a), and further studied by Brandenburger et al. (2008) and Halpern (2010).⁵

The property (L_1) is called *mutual singularity*. For the time being let us focus on a countable X with the discrete topology, in which case mutual singularity has a more natural interpretation. First observe that in this case, mutual singularity is equivalent to requiring that the supports of the different theories are disjoint. Now, in order to understand the intuition behind such a restriction of non-overlapping supports, let us first interpret an LPS as a way to capture the notion of “a state

⁴For each pair $((\mu^1, \dots, \mu^n), (\nu^1, \dots, \nu^n)) \in \mathcal{N}_n(X) \times \mathcal{N}_n(X)$, take the pair (μ^m, ν^m) with the maximum Prohorov distance, thus inducing a metric in $\mathcal{N}_n(X)$. Since the Prohorov metric induces the topology of weak convergence in $\Delta(X)$, the metric defined above is simply the product topology of weak convergence.

⁵Blume et al. (1991a) use the term “lexicographic conditional probability systems”, whereas Halpern (2010) uses the term “mutually singular lexicographic probability systems” to describe the object introduced in Definition 1.

being infinitely more likely than another state". In particular, $x \in X$ is said to be deemed infinitely more likely than $y \in X$ whenever it is the case that $x \in \text{Supp}(\mu^m)$ and $y \in \text{Supp}(\mu^\ell)$ with $m < \ell$.⁶ Then, mutual singularity ensures that it cannot be the case that x is deemed infinitely more likely than x itself. Thus, mutual singularity allows the agent to partition X into classes of most likely elements, second most likely elements, and so on.⁷

Note that this type of reasoning seems mostly plausible in cases where the agent can directly observe verifiable events (e.g., see Heifetz et al., 2010). In particular, the primary theory corresponds to the unconditional beliefs over X , with $\text{Supp}(\mu^1)$ being the states deemed most likely by the agent. The secondary theory then corresponds to the conditional beliefs given that the agent has learned that no state in $\text{Supp}(\mu^1)$ occurs, and so on.

Lexicographic probability systems (that satisfy mutual singularity) are axiomatized by Blume et al. (1991a, Sect. 5), who present a preference-based foundation that relaxes the standard Archimedean property.⁸ Accordingly, the agent's preferences over two acts, are represented by first comparing the unconditional expected utilities, i.e., the expected utilities under the primary theory. If these are equal, then we compare conditional expected utilities given the set of states that receive zero probability by the primary theory, i.e., we compare the expected utilities under the secondary theory, and so on. In the same paper, Blume et al. (1991a) also point out that in the existence of mutual singularity, a link between LPS's and CPS's – which will be formally defined in the next section – can be established.

Of course, while we find this way of modeling beliefs in general appealing, we still need to recognize that mutual singularity can be rather restrictive in some cases. For instance, if we restrict attention to mutually singular lexicographic beliefs, there exist games with the property that not every self admissible set could arise under rationality and common assumption in rationality (e.g., see Brandenburger et al., 2008, Fig. 2.6).⁹ Moreover, mutual singularity implicitly postulates a rather radical way of forming alternative theories. Namely, in the presence of mutual singularity, once the agent discards her primary theory, she rules out all states that were deemed possible by this theory.

Despite these limitations, in this paper we maintain mutual singularity for two reasons. First, as we have already mentioned above, mutual singularity allows us to draw a relationship between (hierarchies of) lexicographic beliefs and (hierarchies of) conditional beliefs, which is the main aim of this project. In addition, we would like to stay in line with the natural predecessor of this paper

⁶Later in this section, we generalize this definition to capture the idea of a Borel event being infinitely more likely than another Borel event in an arbitrary topological space.

⁷A similar type of reasoning is modeled with a plausibility ordering (e.g., see Perea, 2013).

⁸In the same paper, they also provide a weaker axiomatization for LPS's that does not necessarily satisfy mutual singularity.

⁹I would like to thank one of the referees for pointing this out.

(Brandenburger et al., 2008). We further discuss mutual singularity in Section 6.

Let $\mathcal{L}(X)$ denote the space of lexicographic probability systems, while $\mathcal{L}_n(X) := \mathcal{L}(X) \cap \mathcal{N}_n(X)$ denotes the space of lexicographic probability systems of length n . Throughout the paper, whenever an LPS $\tilde{\mu}$ is of length n , we write $\Lambda(\tilde{\mu}) = n$. Furthermore, let $\mathcal{L}^+(X) := \mathcal{L}(X) \cap \mathcal{N}^+(X)$ denote the space of full-support LPS's. Several papers in the literature restrict focus only to full-support LPS's in order to capture the idea of cautious agents who deem every state possible (e.g., Brandenburger et al., 2008; Heifetz et al., 2010). However, in general an LPS may still have null states that are deemed unlikely by each theory. In the present paper, we allow for such LPS's. Finally, recall from Brandenburger et al. (2008, Cor. C.2) that $\mathcal{L}(X)$ and $\mathcal{L}^+(X)$ are Borel in $\mathcal{N}(X)$.

Now, we recall the notion of *assumption*, first introduced by Brandenburger et al. (2008). Fix an LPS, $\tilde{\mu} = (\mu^1, \dots, \mu^n) \in \mathcal{L}(X)$ and two disjoint Borel events $A, B \in \mathcal{F}$. We say that A is *infinitely more likely than* B under $\tilde{\mu}$ whenever

- (a) for each open $T \subseteq X$ with $A \cap T \neq \emptyset$, there is some $m \in \{1, \dots, n\}$ such that $\mu^m(A \cap T) > 0$,
- (b) if $\mu^m(A \cap T_1) > 0$ and $\mu^\ell(B \cap T_2) > 0$ for any open $T_1, T_2 \subseteq X$, then $m < \ell$.

Intuitively, A is infinitely more likely than B if it is the case that B is deemed likely only after A has already been ruled out, e.g., if X is finite, every element of A appears before every element of B in the LPS that describes the agent's beliefs. Notice that the idea of an event being infinitely more likely than another event is conceptually related to the notion of mutual singularity, in the sense that A cannot be infinitely more likely than A itself. Similar notions of “infinitely more likely” have appeared in several papers in the literature (e.g., Blume et al., 1991a,b; Battigalli, 1996; Asheim and Dufwenberg, 2003; Brandenburger et al., 2008; Heifetz et al., 2010). Though, most of these papers consider full support LPS's, it does not need to be the case.

Definition 2. A Borel event $A \in \mathcal{F}$ is *assumed* under $\tilde{\mu}$ whenever A is infinitely more likely than its complement $X \setminus A$ under $\tilde{\mu}$.

This is the natural extension of probability-1 belief to cases where the agent reasons according to an LPS. Throughout the paper, let $\mathcal{A}(A) \subseteq \mathcal{L}(X)$ denote the LPS's that assume A , and then we naturally define $\mathcal{A}^+(A) := \mathcal{A}(A) \cap \mathcal{L}^+(X)$. Brandenburger et al. (2008, Cor. 3) showed that $\mathcal{A}^+(A)$ is a Borel subset of $\mathcal{L}(X)$ whenever A is Borel.

3.2. Conditional probability systems

Let $\mathcal{B} \subseteq \mathcal{F} \setminus \{\emptyset\}$ be a collection of non-empty, Borel conditioning events.

Definition 3. A *conditional probability system (CPS)* on $(X, \mathcal{F}, \mathcal{B})$ is a function $\pi : \mathcal{F} \times \mathcal{B} \rightarrow [0, 1]$ that satisfies

$$(C_1) \quad \pi(B|B) = 1, \text{ if } B \in \mathcal{B},$$

$$(C_2) \quad \pi(\cdot|B) \text{ is a probability measure over } (X, \mathcal{F}) \text{ for every } B \in \mathcal{B},$$

$$(C_3) \quad \pi(A|C) = \pi(A|B) \cdot \pi(B|C), \text{ if } A \subseteq B \subseteq C, \text{ with } A \in \mathcal{F} \text{ and } B, C \in \mathcal{B}.$$

According to the usual interpretation, a CPS describes the agent's beliefs upon having observed each of the conditioning hypotheses in \mathcal{B} . Standard (unconditional) beliefs are a special case of a CPS with $\mathcal{B} = \{X\}$. Conditional probability systems were first introduced by Rényi (1955), and later further studied by Myerson (1986) and Battigalli and Siniscalchi (1999).

In general, $\mathcal{B} \cup \{\emptyset\}$ does not need to form an algebra of events. Indeed, typically \mathcal{B} corresponds to the collection of information sets in a finite dynamic game with perfect recall, which do not necessarily form an algebra in the space of the opponents' strategies.

Now, in the presence of perfect recall, \mathcal{B} is restricted to satisfy certain structural properties, viz., two arbitrary conditioning events either have an empty intersection or one is a subset of the other (see Halpern, 2010, Sect. 3.3). Every CPS satisfying these restrictions can be extended (in a belief-preserving way) to another CPS with the collection of conditioning events forming an algebra, i.e., the extended collection of conditioning events will be the closure of the original one with respect to finite unions, finite intersections and complements. Of course, we should point out that such an extension is not in general unique, e.g., if we start with $X = \{x_1, x_2, x_3\}$ and $\mathcal{B} = \{\{x_1\}\}$, then the extended collection of conditioning events should also contain $\{x_2, x_3\}$ and X itself. In this case, the conditional beliefs given $\{x_2, x_3\}$ are not uniquely defined, and the same applies for the unconditional beliefs given X . Nevertheless, every extension would preserve the original conditional beliefs at every information set where the player is active, and therefore from a game-theoretic point of view the conditional beliefs given the “new” information sets are irrelevant. Furthermore, notice that if we assume that players reason at all histories in the game (including those where nature moves) – similarly to what is assumed in Battigalli and Siniscalchi (2002) – rather than just at those where they are active, the extension would be unique as the players would have beliefs at every history about the relative likelihood of any two of their own information sets. In this paper, we also adopt this last point of view, and therefore assuming that $\mathcal{B} \cup \{\emptyset\}$ is an algebra is done without loss of generality.¹⁰ We further discuss this issue in the next section after Proposition 1.

Whenever $\mathcal{B} \cup \{\emptyset\}$ is a finite algebra on X , we say that \mathcal{B} is *finitely generated*. For a fixed \mathcal{B} , let $\Delta^{\mathcal{B}}(X)$ denote the space of conditional probability systems on $(X, \mathcal{F}, \mathcal{B})$. If \mathcal{B} is finitely generated,

¹⁰This assumption is also present in Brandenburger et al. (2007).

we say that $\pi \in \Delta^{\mathcal{B}}(X)$ is *finitary*. For some countable class $\mathfrak{E} \subseteq \mathfrak{F}$ of finitely generated collections of conditioning events, let $\mathcal{C}(X, \mathfrak{E}) := \bigoplus_{\mathcal{B} \in \mathfrak{E}} \Delta^{\mathcal{B}}(X)$. If X is a countable set, \mathfrak{F} is also countable. In this case, let $\mathcal{C}(X) := \mathcal{C}(X, \mathfrak{F})$ denote the space of all finitary CPS's.

Recall that if \mathcal{B} is a collection of clopen events, $\Delta^{\mathcal{B}}(X)$ is Polish (Battigalli and Siniscalchi, 1999, Lem. 1). Therefore, if X is a countable set with the discrete topology, $\Delta^{\mathcal{B}}(X)$ is Polish for every $\mathcal{B} \in \mathfrak{F}$, and therefore $\mathcal{C}(X)$ is also Polish. Throughout the paper, unless explicitly stated otherwise, whenever we consider a countable space of uncertainty we assume that it is endowed with the discrete topology. This assumption is discussed in Section 6.

Following Brandenburger et al. (2007, Def. 4), we say that a CPS $\pi \in \Delta^{\mathcal{B}}(X)$ is full-support, and we write $\pi \in \mathcal{C}^+(X)$, whenever $B \subseteq \text{Supp}(\pi(\cdot|B))$ for each $B \in \mathcal{B}$. Notice that if X is countable with the discrete topology, $B \subseteq \text{Supp}(\pi(\cdot|B))$ is equivalent to $B = \text{Supp}(\pi(\cdot|B))$.

For some finitary $\pi \in \Delta^{\mathcal{B}}(X)$, define the sub-collection $\mathcal{B}_\pi := \{B_\pi^1, \dots, B_\pi^\nu\} \subseteq \mathcal{B}$ of conditioning events as follows: First, let $\mathcal{P}_\mathcal{B}$ denote the finite partition of X from which \mathcal{B} is generated, i.e., $\mathcal{P}_\mathcal{B}$ is the finest \mathcal{B} -measurable partition, and let

$$\mathcal{P}_\pi^1 := \mathcal{P}_\mathcal{B} \qquad B_\pi^1 := \bigcup_{P \in \mathcal{P}_\pi^1} P$$

Obviously, since \mathcal{P}_π^1 is a partition of X it follows directly that B_π^1 coincides with the entire space X , and therefore $\pi(\cdot|B_\pi^1)$ describes the agent's unconditional beliefs. Then, for each integer $m > 1$ inductively define

$$\mathcal{P}_\pi^m := \{P \in \mathcal{P}_\pi^{m-1} : \pi(P|B_\pi^{m-1}) = 0\} \qquad B_\pi^m := \bigcup_{P \in \mathcal{P}_\pi^m} P$$

Intuitively, B_π^2 is the largest conditioning event for which the conditional beliefs are not given by applying Bayes rule, i.e., it is the largest $B \in \mathcal{B}$ with $\pi(B|B_\pi^1) = 0$, implying that (C_3) cannot be used to derive $\pi(\cdot|B_\pi^2)$. In other words, $\pi(\cdot|B_\pi^2)$ describes the conditional beliefs once the agent has ruled out all conditioning hypotheses deemed likely by $\pi(\cdot|B_\pi^1)$. Notice already the conceptual similarity between $\pi(\cdot|B_\pi^2)$ and the secondary theory of an LPS, in that both describe the agent's beliefs once the usual unconditional beliefs have been discarded. We come back to the relationship between the two notion in the next section. Likewise, B_π^3 is the largest $B \in \mathcal{B}$ with $\pi(B|B_\pi^1) = 0$ and $\pi(B|B_\pi^2) = 0$, i.e., $\pi(\cdot|B_\pi^3)$ describes the conditional beliefs once the agent has ruled out all conditioning hypotheses deemed likely by $\pi(\cdot|B_\pi^1)$ or by $\pi(\cdot|B_\pi^2)$, and so on.

Observe that by construction there exists some integer $\nu > 0$ such that

$$B_\pi^1 \supsetneq B_\pi^2 \supsetneq \dots \supsetneq B_\pi^\nu \neq \emptyset = B_\pi^{\nu+1} = B_\pi^{\nu+2} = \dots$$

in which case we say that π is of length ν , and we write $\Lambda(\pi) = \nu$. The latter follows directly from the fact that \mathcal{B} is finitely generated.

Lemma 1. Consider $\pi \in \Delta^{\mathcal{B}}(X)$ and $\rho \in \Delta^{\mathcal{B}}(X)$ with $\Lambda(\pi) = \Lambda(\rho) = n$, such that $\pi(A|B_{\pi}^m) = \rho(A|B_{\rho}^m)$ for every $A \in \mathcal{F}$ and for all $m = 1, \dots, n$. Then, $\pi(A|B) = \rho(A|B)$ for all $A \in \mathcal{F}$ and all $B \in \mathcal{B}$.

The previous result implies that every $\pi \in \Delta^{\mathcal{B}}(X)$ is determined by the conditional beliefs given the events in \mathcal{B}_{π} . Henceforth, we call \mathcal{B}_{π} the *collection of π -relevant conditioning events*.

Now, we recall the notion of strong belief, first introduced in the literature by Battigalli and Siniscalchi (1999).

Definition 4. A Borel event $A \in \mathcal{F}$ is *strongly believed* under π whenever for each $B \in \mathcal{B}$ with $A \cap B \neq \emptyset$ it is the case that $\pi(A|B) = 1$.

Intuitively, the agent is certain of A given B , unless B contradicts A . In this respect, strong belief is a generalization of probability-1 belief to cases where the agent reasons according to a CPS. Throughout the paper, let $\mathcal{SB}(A) \subseteq \mathcal{C}(X)$ denote the set of finitary CPS's that strongly believe A . We also define $\mathcal{SB}^+(A) := \mathcal{SB}(A) \cap \mathcal{C}^+(X)$.

3.3. Epistemic equivalence of extended probability systems

In this section, we introduce our notion of epistemic equivalence between a CPS and LPS, according to which the two are equivalent whenever they share the same length, and also for every integer m smaller or equal than the length, the m -th theory of the LPS coincides with the conditional beliefs given the m -th relevant conditioning event.

Definition 5. Consider some $\tilde{\mu} \in \mathcal{L}(X)$ and some $\pi \in \Delta^{\mathcal{B}}(X)$. We say that $\tilde{\mu}$ is *epistemically equivalent* to π whenever it is the case that

$$(E_1) \quad \Lambda(\tilde{\mu}) = \Lambda(\pi) = n,$$

$$(E_2) \quad \mu^m(A) = \pi(A|B_{\pi}^m), \text{ for all Borel } A \subseteq X \text{ and for each } m = 1, \dots, n.$$

This definition of epistemic equivalence is already present, though not formally stated, in Brandenburger et al. (2007). A similar definition of equivalence is also used by Halpern (2010), who relates LPS's with a special class of CPS's, the so-called Popper spaces.

Before moving forward, let us first elaborate on the two conditions that constitute our notion of epistemic equivalence. Firstly, observe that we require the primary hypothesis of the LPS to coincide with the unconditional beliefs in the CPS. We find this requirement natural, as both models are generalizations of the standard probabilistic beliefs. Moreover, as we have already discussed in the previous section, $\pi(\cdot|B_{\pi}^m)$ describes the conditional beliefs once the agent has discarded the beliefs

described by $\pi(\cdot|B_\pi^1)$ and $\pi(\cdot|B_\pi^2) \dots$ and $\pi(\cdot|B_\pi^{m-1})$. On the other hand, the m -th theory μ^m of an LPS describes the agent's beliefs once she has discarded the theories μ^1 and $\mu^2 \dots$ and μ^{m-1} . In this respect, $\pi(\cdot|B_\pi^m)$ and μ^m capture exactly the same idea, i.e., they both describe the beliefs once the agent has already discarded her beliefs $m - 1$ times. Finally, the respective lengths describe the number of times the agent may change her mind due to the fact that her previously-formed beliefs have been discarded, implying that a notion of epistemic equivalence should require them to be equal.

The following result is due to Brandenburger et al. (2007) and associates every finitary CPS with an epistemically equivalent LPS.¹¹

Proposition 1 (Brandenburger et al., 2007, Prop. 1). *There exists a surjective function $\beta : \mathcal{C}(X) \rightarrow \mathcal{L}(X)$ such that π is epistemically equivalent to $\beta(\pi)$ for all $\pi \in \mathcal{C}(X)$.*

First, note that β is the only function mapping each CPS to an epistemically equivalent LPS, as shown in the proof of the previous result in the Appendix. Moreover, though β is surjective – implying that every LPS is the image of some finitary CPS – it is not necessarily injective, i.e., there may exist more than one CPS's which are epistemically equivalent to the same LPS (see Brandenburger et al., 2007, Ex. 1). However, this can only be the case if the two CPS's have a different collection of conditioning events, i.e., if $\pi_1 \in \Delta^{\mathcal{B}_1}(X)$ and $\pi_2 \in \Delta^{\mathcal{B}_2}(X)$ are such that $\pi_1 \neq \pi_2$ and $\beta(\pi_1) = \beta(\pi_2)$, then it is the case that $\mathcal{B}_1 \neq \mathcal{B}_2$. The latter follows directly from Lemma 1.

It is important to point out that the previous result relies on the collection of conditioning events satisfying certain properties (Halpern, 2010, Prop. 3.9), which by the way trivially hold when the \mathcal{B} corresponds to the information sets in a game with perfect recall, and therefore also hold when \mathcal{B} is finitely generated. Otherwise, it follows from Halpern (2010, Ex. 3.8) that there is a CPS with no epistemically equivalent LPS.¹²

Proposition 2. *If X is countable, then β is Borel.*

4. Hierarchies of extended beliefs

Let Θ be a countable fundamental space. In a game, Θ denotes the different values that a set of relevant parameters can take, i.e., each $\theta \in \Theta$ corresponds to a payoff vector (Harsanyi, 1967-68), or a strategy profile (Brandenburger et al., 2008; Battigalli and Siniscalchi, 2002), or a combination of the two. Let $I = \{a, b\}$ be the set of agents, with typical elements i and j .¹³

¹¹For the sake of completeness we also provide a proof in the Appendix.

¹²I am indebted to one of the referees for suggesting this particular comment.

¹³Our analysis can be directly generalized to any finite set of agents, in which case we obviously allow for correlated beliefs, as usual.

4.1. Hierarchies of lexicographic beliefs

Each agent forms lexicographic beliefs about Θ (first order lexicographic beliefs), lexicographic beliefs about Θ and the opponent's first order lexicographic beliefs (second order lexicographic beliefs), and so on. Formally, consider the sequence

$$\begin{aligned}\Theta_0 &:= \Theta \\ \Theta_1 &:= \Theta_0 \times \mathcal{L}(\Theta_0) \\ &\vdots \\ \Theta_{k+1} &:= \Theta_k \times \mathcal{L}(\Theta_k) \\ &\vdots\end{aligned}$$

Definition 6. A *lexicographic belief hierarchy* (\mathcal{L} -*hierarchy*) is sequence of LPS's $(\tilde{\mu}_1, \tilde{\mu}_2, \dots) \in \mathcal{L}_n(\Theta_0) \times \mathcal{L}_n(\Theta_1) \times \dots$ for some integer $n > 0$. Let $T_0^L := \bigoplus_{n=1}^{\infty} (\mathcal{L}_n(\Theta_0) \times \mathcal{L}_n(\Theta_1) \times \dots)$ denote the space of all \mathcal{L} -hierarchies.

The LPS $\tilde{\mu}_k = (\mu_k^1, \dots, \mu_k^n) \in \mathcal{L}_n(\Theta_{k-1})$ denotes the k -th order lexicographic beliefs, with $\mu_k^m \in \Delta(\Theta_{k-1})$ being the m -th theory of the k -th order beliefs. Observe that by definition the \mathcal{L} -hierarchy consists of a sequence of LPS's with same length, henceforth called the length of the lexicographic belief hierarchy, and denoted by $\Lambda(\tilde{\mu}_1, \tilde{\mu}_2, \dots)$.¹⁴

As usual, we restrict \mathcal{L} -hierarchies to satisfy the standard coherency property, which roughly speaking says that higher order beliefs do not contradict lower order beliefs. Formally, a lexicographic belief hierarchy is coherent whenever it belongs to

$$T_1^L := \{ (\tilde{\mu}_1, \tilde{\mu}_2, \dots) \in T_0^L : \text{marg}_{\Theta_{k-2}} \tilde{\mu}_k = \tilde{\mu}_{k-1}, \text{ for all } k > 1 \},$$

where $\text{marg}_{\Theta_{k-2}} \tilde{\mu}_k := (\text{marg}_{\Theta_{k-2}} \mu_k^1, \dots, \text{marg}_{\Theta_{k-2}} \mu_k^n)$. Observe that by definition, if $(\tilde{\mu}_1, \tilde{\mu}_2, \dots)$ is coherent, it is the case that $\text{marg}_{\Theta} \tilde{\mu}_k$ is an LPS over Θ . Throughout, the paper for an arbitrary measurable space Y , let $\mathcal{L}_{\Theta}(\Theta \times Y) := \{ \tilde{\mu} \in \mathcal{L}(\Theta \times Y) : \text{marg}_{\Theta} \tilde{\mu} \in \mathcal{L}(\Theta) \}$. We also adopt the notational convention that $\mathcal{L}_{\Theta}(\Theta) := \mathcal{L}(\Theta)$.

¹⁴This is a standard assumption, even though it is rarely explicitly stated in the literature. The reason is that most papers start with a lexicographic type space model – which is formally defined later in this paper – from which they derive the lexicographic belief hierarchies (e.g., Brandenburger et al., 2008). However, in this case each hierarchy consists by construction of a sequence of LPS's with the same length, which is the length of the corresponding type. To our knowledge, the only paper that begins by explicitly constructing lexicographic belief hierarchies is Catonini (2012), but even there the property of all orders of beliefs being of the same length is implicitly imposed via the standard coherency restriction, that we also introduce below.

Notice our definition of an \mathcal{L} -hierarchy requires every order of beliefs to be an LPS, and therefore to satisfy mutual singularity. As we have already discussed in Section 3.1, imposing mutual singularity on the underlying space of uncertainty can be justified on the basis of agent i receiving direct verifiable evidence about Θ . On the other hand, i cannot receive direct verifiable evidence regarding j 's beliefs, and therefore mutual singularity on the opponent's beliefs may not be as natural (Heifetz et al., 2010). However, once we have assumed that $(\tilde{\mu}_1, \tilde{\mu}_2, \dots)$ is coherent, it follows directly that $\tilde{\mu}_2$ satisfies mutual singularity on $\Theta \times \mathcal{L}(\Theta)$, even if two different theories of $\tilde{\mu}_2$ have exactly the same beliefs over $\mathcal{L}(\Theta)$. Therefore, mutual singularity on higher order beliefs is a direct consequence of mutual singularity over the underlying space of uncertainty together with coherency, thus implying that it is not as restrictive as it might seem at first. We further elaborate on this issue in Section 6.

Now, we further restrict \mathcal{L} -hierarchies to satisfy common certainty in coherency.¹⁵ Formally, for every $\ell > 1$, we inductively define

$$T_\ell^L := \{ (\tilde{\mu}_1, \tilde{\mu}_2, \dots) \in T_1^L \ : \ \mu_{k+2}^m(\Theta \times \text{Proj}_{\mathcal{L}(\Theta_0) \times \dots \times \mathcal{L}(\Theta_k)} T_{\ell-1}^L) = 1, \\ \text{for all } k \geq 0 \text{ and for every } m = 1, \dots, \Lambda(\tilde{\mu}_{k+2}) \} .$$

For instance, T_2^L contains those \mathcal{L} -hierarchies which have the property that for every $k > 0$ all theories of the k -th order lexicographic beliefs assign probability 1 to the opponent's lower order lexicographic beliefs not contradicting each other. Then, the \mathcal{L} -hierarchies satisfying coherency and common certainty in coherency are those in

$$T^L := \bigcap_{\ell=1}^{\infty} T_\ell^L .$$

Henceforth, unless explicitly stated otherwise, the term *lexicographic belief hierarchies* refers to \mathcal{L} -hierarchies that satisfy coherency and common certainty in coherency.

Hierarchies of lexicographic beliefs are typically represented by type space models (Brandenburger et al., 2008). This is a natural extension of the usual representation of belief hierarchies as introduced by Harsanyi (1967-68).

Definition 7. We define a *lexicographic type space model* as a tuple $(\Theta, T_a, T_b, \tilde{\lambda}_a, \tilde{\lambda}_b)$, where T_i is a Polish space of \mathcal{L} -types, and $\tilde{\lambda}_i : T_i \rightarrow \mathcal{L}_\Theta(\Theta \times T_j)$ is a Borel function.

An \mathcal{L} -type, $t_i \in T_i$, is a complete description of agent i 's state of mind, thus inducing an \mathcal{L} -hierarchy, as illustrated below. Let $\tilde{\lambda}_i(t_i) := (\lambda_i^1(t_i), \dots, \lambda_i^n(t_i)) \in \mathcal{L}_\Theta(\Theta \times T_j)$.

The first order lexicographic beliefs are given by $\tilde{\mu}_1(t_i) = (\mu_1^1(t_i), \dots, \mu_1^n(t_i)) \in \mathcal{N}_n(\Theta_0)$, where for each Borel event $E_0 \subseteq \Theta_0$,

$$\mu_1^m(t_i)(E_0) = \int_{\{(\theta, t_j) : \theta \in E_0\}} d\lambda_i^m(t_i)$$

¹⁵We say that an event is certain under an LPS if it receives probability 1 by every theory of the LPS.

is the probability that E_0 receives by the m -th theory of $\tilde{\mu}_1(t_i)$. In order to show that $\tilde{\mu}_1(t_i)$ is an LPS, it suffices to prove that (L_1) holds, i.e., we would like to show that there exist Borel events $A_1, \dots, A_n \subseteq \Theta_0$ such that $\mu_1^m(t_i)(A_m) = 1$ and $\mu_1^\ell(t_i)(A_m) = 0$. The latter follows directly from Definition 7, and more specifically from the fact that $\text{marg}_\Theta \tilde{\lambda}_i(t_i) \in \mathcal{L}(\Theta)$.

Likewise, the k -th order lexicographic beliefs are given by $\tilde{\mu}_k(t_i) = (\mu_k^1(t_i), \dots, \mu_k^n(t_i)) \in \mathcal{N}_n(\Theta_{k-1})$, where for each Borel event $E_{k-1} \subseteq \Theta_{k-1}$,

$$\mu_k^m(t_i)(E_{k-1}) = \int_{\{(\theta, t_j): (\theta, \tilde{\mu}_1(t_j), \dots, \tilde{\mu}_{k-1}(t_j)) \in E_{k-1}\}} d\lambda_i^m(t_i)$$

is the probability that E_{k-1} receives by the m -th theory of $\tilde{\mu}_k(t_i)$. Once again in order to prove that $\tilde{\mu}_k(t_i)$ is an LPS, we use the fact that $\tilde{\lambda}_i(t_i) \in \mathcal{L}_\Theta(\Theta \times T_j)$. More specifically, observe that $\mu_k^m(t_i)(A_m \times \mathcal{L}(\Theta_0) \times \dots \times \mathcal{L}(\Theta_{k-2})) = 1$ while at the same time $\mu_k^\ell(t_i)(A_m \times \mathcal{L}(\Theta_0) \times \dots \times \mathcal{L}(\Theta_{k-2})) = 0$, implying that (L_1) holds and therefore $\tilde{\mu}_k(t_i)$ is an LPS over Θ_{k-1} .

Notice that the previous definition is slightly different compared to the usual one, in that we require $\tilde{\lambda}_i(t_i)$ to be an element of $\mathcal{L}_\Theta(\Theta \times T_j)$, instead of simply being an LPS in $\mathcal{L}(\Theta \times T_j)$. The reason we impose this additional restriction is that otherwise we would have types whose first order beliefs would violate (L_1) , implying that they would not correspond to a lexicographic belief hierarchy as this is defined above. Formally, if a type t_i is mapped via $\tilde{\lambda}_i$ to an LPS outside $\mathcal{L}_\Theta(\Theta \times T_j)$, then it is the case that $\tilde{\mu}_1(t_i) = \text{marg}_\Theta \tilde{\lambda}_i(t_i) \notin \mathcal{L}(\Theta)$, implying that $(\tilde{\mu}_1(t_i), \tilde{\mu}_2(t_i), \dots)$ is not an \mathcal{L} -hierarchy.¹⁶ The following example illustrates such a case.

Example 1. Consider the tuple $(\Theta, T_a, T_b, \tilde{\lambda}_a, \tilde{\lambda}_b)$, such that $\Theta = \{\theta_1, \theta_2\}$, $T_a = \{t_a^0\}$, $T_b = \{t_b^1, t_b^2\}$, and suppose that $\tilde{\lambda}_a(t_a^0) = (\lambda_a^1(t_a^0), \lambda_a^2(t_a^0))$ is such that $\lambda_a^1(t_a^0)(\theta_1, t_b^1) = 1$ and $\lambda_a^2(t_a^0)(\theta_1, t_b^2) = 1$. Then, observe that the first order beliefs $\tilde{\mu}_1(t_a^0) = (\mu_1^1(t_a^0), \mu_1^2(t_a^0))$ violate (L_1) , since it is the case that $\mu_1^1(t_a^0)(\theta_1) = \mu_1^2(t_a^0)(\theta_1) = 1$. Hence, $\tilde{\mu}_1(t_a^0)$ is not an LPS over Θ , implying that $(\Theta, T_a, T_b, \tilde{\lambda}_a, \tilde{\lambda}_b)$ is not a lexicographic type space. \triangleleft

Observe that by construction every \mathcal{L} -type is associated with a lexicographic belief hierarchy that satisfies coherency and common certainty in coherency, and therefore it is an element of T^L . However, this does not imply common certainty in the event that everybody's \mathcal{L} -hierarchy has the same length, e.g., it may be the case that $\tilde{\lambda}(t_i) \in \mathcal{L}_n(\Theta \times T_j)$, and still for some $(\theta, t_j) \in \text{Supp}(\tilde{\lambda}_i(t_i))$ it is the case that $\tilde{\lambda}_j(t_j) \in \mathcal{L}_m(\Theta \times T_i)$ with $m \neq n$.

¹⁶Several papers in the literature, including Brandenburger et al. (2008), do not require mutual singularity on the underlying space of uncertainty, but rather only on $\Theta \times T_j$, in which case the image of each type does not need to be in $\mathcal{L}_\Theta(\Theta \times T_j)$, but merely in $\mathcal{L}(\Theta \times T_j)$. However, if one wants to impose mutual singularity on Θ , like for instance in Heifetz et al. (2010) or in this paper, $\tilde{\lambda}_i(t_i)$ should necessarily belong to $\mathcal{L}_\Theta(\Theta \times T_j)$.

Finally, note that T^L should not be viewed as a set of types, but merely as a subset of belief hierarchies. The reason is that types are in principle abstract objects which have a meaning only within a type space model, whereas belief hierarchies are self-contained. In either case, later in the paper, we construct a type space model that induces all belief hierarchies in T^L .

4.2. Hierarchies of conditional beliefs

Each agent is endowed with a finitely generated collection of conditioning hypotheses, and forms conditional beliefs about Θ (first order conditional beliefs), conditional beliefs about Θ and the opponent's first order conditional beliefs (second order conditional beliefs), and so on. Before moving forward, let us already point out that our analysis generalizes the standard one by Battigalli and Siniscalchi (1999), in that we allow players to be uncertain about the opponent's collection of conditioning events. We further elaborate on this issue later in the paper.

Formally, for an arbitrary measurable space Y and a finitary collection \mathcal{B} of conditioning events in Θ , let $\mathcal{B} \times Y := \{B \times Y | B \in \mathcal{B}\}$ denote the cylinders generated by \mathcal{B} . Each cylinder in $\mathcal{B} \times Y$ corresponds to a conditioning event in $\Theta \times Y$. Moreover, define

$$\begin{aligned}\mathfrak{F}_Y &:= \{ \mathcal{B} \times Y \mid \mathcal{B} \in \mathfrak{F} \} \\ &= \{ \{B \times Y | B \in \mathcal{B}\} \mid \mathcal{B} \in \mathfrak{F} \}.\end{aligned}$$

Each element of \mathfrak{F}_Y is a different collection of conditioning events in $\Theta \times Y$. Observe that \mathfrak{F}_Y is countable, whenever Θ is countable. Now, let

$$\begin{array}{ll}\Psi_0 := \Theta & \mathfrak{F}_0 := \mathfrak{F} \\ \Psi_1 := \Psi_0 \times \mathcal{C}(\Psi_0, \mathfrak{F}_0) & \mathfrak{F}_1 := \mathfrak{F}_{\mathcal{C}(\Psi_0, \mathfrak{F}_0)} \\ \vdots & \vdots \\ \Psi_{k+1} := \Psi_k \times \mathcal{C}(\Psi_k, \mathfrak{F}_k) & \mathfrak{F}_{k+1} := \mathfrak{F}_{\mathcal{C}(\Psi_0, \mathfrak{F}_0) \times \dots \times \mathcal{C}(\Psi_k, \mathfrak{F}_k)} \\ \vdots & \vdots\end{array}$$

Definition 8. A *hierarchy of conditional beliefs* (\mathcal{C} -*hierarchy*) is a sequence of CPS's $(\pi_1, \pi_2, \dots) \in \Delta^{\mathcal{B}}(\Psi_0) \times \Delta^{\mathcal{B} \times \mathcal{C}(\Psi_0, \mathfrak{F}_0)}(\Psi_1) \times \dots$ for some $\mathcal{B} \in \mathfrak{F}$. Let $T_0^{\mathcal{C}} := \bigoplus_{\mathcal{B} \in \mathfrak{F}} (\Delta^{\mathcal{B}}(\Psi_0) \times \Delta^{\mathcal{B} \times \mathcal{C}(\Psi_0, \mathfrak{F}_0)}(\Psi_1) \times \dots)$ denote the space of all \mathcal{C} -hierarchies.

The CPS $\pi_k \in \Delta^{\mathcal{B} \times \mathcal{C}(\Psi_0, \mathfrak{F}_0) \times \dots \times \mathcal{C}(\Psi_{k-2}, \mathfrak{F}_{k-2})}(\Psi_{k-1})$ denotes the k -th order conditional beliefs, with $\pi_k(\cdot | B \times \mathcal{C}(\Psi_0, \mathfrak{F}_0) \times \dots \times \mathcal{C}(\Psi_{k-2}, \mathfrak{F}_{k-2})) \in \Delta(\Psi_k)$ denoting the k -th order conditional beliefs given the conditioning event $B \times \mathcal{C}(\Psi_0, \mathfrak{F}_0) \times \dots \times \mathcal{C}(\Psi_{k-2}, \mathfrak{F}_{k-2}) \in \mathcal{B} \times \mathcal{C}(\Psi_0, \mathfrak{F}_0) \times \dots \times \mathcal{C}(\Psi_{k-2}, \mathfrak{F}_{k-2})$. Observe that by definition the \mathcal{C} -hierarchy consists of a sequence of CPS's with the property that

the collections of conditioning events are cylinders generated by some fixed $\mathcal{B} \in \mathfrak{F}$. That is, if an agent's conditioning events on the fundamental space of uncertainty are those in $\mathcal{B} \in \mathfrak{F}$, then her collection of conditioning events on Ψ_{k+1} contains those in $\mathcal{B} \times \mathcal{C}(\Psi_0, \mathfrak{F}_0) \times \cdots \times \mathcal{C}(\Psi_k, \mathfrak{F}_k)$, which is obviously an element of \mathfrak{F}_{k+1} . This restriction is already present in Battigalli and Siniscalchi (1999), with the difference that they only consider cases where \mathfrak{F}_0 contains a unique collection of conditioning hypotheses for each player, and this collection is common knowledge. Henceforth, for notation simplicity, we write $\Delta^{\mathcal{B}}(\Psi_{k-1}) := \Delta^{\mathcal{B} \times \mathcal{C}(\Psi_0, \mathfrak{F}_0) \times \cdots \times \mathcal{C}(\Psi_{k-2}, \mathfrak{F}_{k-2})}(\Psi_{k-1})$ and for each $B \in \mathcal{B}$ we write $\pi_k(\cdot|B) := \pi_k(\cdot | B \times \mathcal{C}(\Psi_0, \mathfrak{F}_0) \times \cdots \times \mathcal{C}(\Psi_{k-2}, \mathfrak{F}_{k-2}))$.

As usual, we restrict attention to conditional belief hierarchies that satisfy coherency, i.e., we require higher order beliefs not to contradict lower order beliefs. Formally, the coherent hierarchies are those in

$$T_1^C := \bigcup_{B \in \mathfrak{F}} \left\{ (\pi_1, \pi_2, \dots) \in \prod_{k=0}^{\infty} \Delta^{\mathcal{B}}(\Psi_k) : \text{marg}_{\Psi_{k-2}} \pi_k = \pi_{k-1}, \text{ for all } k > 1 \right\},$$

where $\text{marg}_{\Psi_{k-2}} \pi_k := (\text{marg}_{\Psi_{k-2}} \pi(\cdot|B) ; B \in \mathcal{B})$.

Lemma 2. *There exists a homeomorphism $f : T_1^C \rightarrow \mathcal{C}(\Theta \times T_0^C, \mathfrak{F}_{T_0^C})$.*

The previous result is a generalization of the standard result by Battigalli and Siniscalchi (1999, Prop. 1). Accordingly, each coherent \mathcal{C} -hierarchy $(\pi_1, \pi_2, \dots) \in T_1^C$ is identified by a CPS on $(\Theta \times T_0^C, \mathcal{B} \times T_0^C)$, with $f^B(\pi_1, \pi_2, \dots) \in \Delta(\Theta \times T_0^C)$ denoting the conditional beliefs given the conditioning hypothesis $B \times T_0^C \in \mathcal{B} \times T_0^C$. The previous homeomorphism is a natural one, in that $\text{marg}_{\Psi_{k-1}} f^B(\pi_1, \pi_2, \dots) = \pi_k(\cdot|B)$ for all $k > 0$ and every $B \in \mathcal{B}$.

We further restrict conditional belief hierarchies to satisfy not only coherency, but also common certainty in coherency. Formally, for each $\ell > 1$, define

$$T_\ell^C := \{t \in T_1^C : f^B(t)(\Theta \times T_{\ell-1}^C) = 1, \text{ for all } B \in \mathcal{B}\}.$$

For instance, T_2^C contains those \mathcal{C} -hierarchies which have the property that for every $k > 0$ the k -th order conditional beliefs given each conditioning event $B \times \mathcal{C}(\Psi_0, \mathfrak{F}_0) \times \cdots \times \mathcal{C}(\Psi_{k-2}, \mathfrak{F}_{k-2})$ assign probability 1 to the event that the opponent's lower order conditional beliefs will not contradict each other. In other words, $(\pi_1, \pi_2, \dots) \in \prod_{k=0}^{\infty} \Delta^{\mathcal{B}}(\Psi_k)$ belongs to T_2^C whenever for each $B \in \mathcal{B}$, $f^B(\pi_1, \pi_2, \dots)$ attaches probability 1 to the opponent being coherent. Then, the conditional belief hierarchies satisfying coherency and common certainty in coherency are those in

$$T^C := \bigcap_{\ell=1}^{\infty} T_\ell^C.$$

Henceforth, unless explicitly stated otherwise the term *conditional belief hierarchies* refers to \mathcal{C} -hierarchies that satisfy coherency and common certainty in coherency.

Hierarchies of conditional beliefs are typically represented by type space models (Battigalli and Siniscalchi, 1999). This is the natural extension of Harsanyi's construction to the case of conditional belief hierarchies.

Definition 9. We define a *conditional type space model* as a tuple $(\Theta, T_a, T_b, g_a, g_b)$, where T_i is a Polish space of \mathcal{C} -types, and $g_i : T_i \rightarrow \mathcal{C}(\Theta \times T_j, \mathfrak{F}_{T_j})$ is a continuous function.

A \mathcal{C} -type, $t_i \in T_i$, is a complete description of the agent's epistemic state, in that it is associated with a \mathcal{C} -hierarchy, as shown below. Let $g_i(t_i) := (g_i^B(t_i) ; B \in \mathcal{B}) \in \Delta^{\mathcal{B}}(\Theta \times T_j)$, with \mathcal{B} being the collection of conditioning hypotheses associated with t_i .

The first order conditional beliefs are given by $\pi_1(t_i) \in \Delta^{\mathcal{B}}(\Psi_0)$, where for each Borel event $E_0 \subseteq \Psi_0$,

$$\pi_1(t_i)(E_0|B) = \int_{\{(\theta, t_j) : \theta \in E_0\}} dg_i^B(t_i)$$

is the probability that E_0 receives given the conditioning hypothesis B . Verifying that $\pi_1(t_i)$ is a CPS is straightforward.

Likewise, the k -th order conditional beliefs are given by $\pi_k(t_i) \in \Delta^{\mathcal{B}}(\Psi_{k-1})$, where for each Borel event $E_{k-1} \subseteq \Psi_{k-1}$

$$\pi_k(t_i)(E_{k-1}|B) = \int_{\{(\theta, t_j) : (\theta, \pi_1(t_j), \dots, \pi_{k-1}(t_j)) \in E_{k-1}\}} dg_i^B(t_i)$$

is the probability that E_{k-1} receives given the conditioning event $B \times \mathcal{C}(\Psi_0, \mathfrak{F}_0) \times \dots \times \mathcal{C}(\Psi_{k-2}, \mathfrak{F}_{k-2})$. Once again, it is straightforward verifying that $\pi_k(t_i)$ is a CPS. Moreover, observe that by construction the \mathcal{C} -hierarchy $(\pi_1(t_i), \pi_2(t_i), \dots)$ satisfies coherency and common certainty in coherency. However, this does not imply common certainty in the event that every player has the same collection of conditioning events, as illustrated in the next example.

Example 2. Consider the conditional type space $(\Theta, T_a, T_b, g_a, g_b)$, where $\Theta = \{\theta_1, \theta_2\}$. Notice that $\mathfrak{F} = \{\mathcal{B}_1, \mathcal{B}_2\}$ where $\mathcal{B}_1 = \{ \{\theta_1, \theta_2\} \}$ and $\mathcal{B}_2 = \{ \{\theta_1\}, \{\theta_2\}, \{\theta_1, \theta_2\} \}$. Suppose that $T_a = \{t_a^1\}$ and $T_b = \{t_b^1, t_b^2\}$ are such that $g_i(t_i^k) \in \Delta^{\mathcal{B}_k}(\Theta \times T_j)$, and assume that $g_a^{\{\theta_1, \theta_2\}}(t_a^1)$ assigns probability 1/4 to each $(\theta, t_b) \in \Theta \times T_b$. Then, observe that t_a^1 is not certain about b 's collection of conditioning events, as both \mathcal{B}_1 and \mathcal{B}_2 are deemed equally likely (given the unique conditioning event $\{\theta_1, \theta_2\}$). \triangleleft

Battigalli and Siniscalchi (1999) showed that if it is the case that $\mathfrak{F}_0 = \{\mathcal{B}\}$, then there exists a conditional type space model such that for every conditional belief hierarchy there exists a unique type associated with it, while at the same time the function $g_i : T_i \rightarrow \Delta^{\mathcal{B}}(\Theta \times T_j)$ is a homeomorphism, implying that every conditional belief is associated with some type. Their result naturally extends the standard result by Mertens and Zamir (1985) and Brandenburger and Dekel (1993) to

the space of conditional beliefs. Below, we further generalize Battigalli and Siniscalchi's construction by introducing uncertainty about the other agent's collection of conditioning events.

Proposition 3. *There exists a homeomorphism $g : T^C \rightarrow \mathcal{C}(\Theta \times T^C, \mathfrak{F}_{T^C})$.*

For every $B \in \mathcal{B}$, let $g^B(\pi_1, \pi_2, \dots) \in \Delta(\Theta \times T^C)$ denote the conditional beliefs given $B \times T^C$. Once again, g is a natural homeomorphism, in that $\text{marg}_{\Psi_{k-1}} g^B(\pi_1, \pi_2, \dots) = \pi_k(\cdot|B)$ for all $k > 0$ and every $B \in \mathcal{B}$.

5. Epistemic equivalence of hierarchies of extended beliefs

5.1. Definition and existence

Recall the notion of epistemic equivalence between a CPS and an LPS, formally introduced in Definition 5. Accordingly, the two are equivalent whenever they share the same length, and the m -th theory of the LPS coincides with the conditional probability measure given the m -th relevant conditioning hypothesis of the CPS.

This notion of epistemic equivalence naturally applies to first order beliefs, thus inducing the concept of first order epistemic equivalence between first order conditional beliefs and first order lexicographic beliefs. Formally, $\tilde{\mu}_1 \in \mathcal{L}(\Theta_0)$ is *first order epistemically equivalent* to $\pi_1 \in \mathcal{C}(\Psi_0, \mathfrak{F}_0)$ whenever it is the case that $\beta(\pi_1) = \tilde{\mu}_1$.

However, the previous concept of epistemic equivalence cannot be directly extended to higher order beliefs, as the k -th order lexicographic beliefs are defined on a different space than the k -th order conditional beliefs. For instance, the second order lexicographic beliefs are described by an LPS over Θ_1 whereas the second order conditional beliefs are described by a CPS over Ψ_1 . Thus, in order to define second order epistemic equivalence between second order lexicographic beliefs and second order conditional beliefs, we first need to translate each event in Θ_1 to an event in Ψ_1 . In order to do this, we define the Borel surjective function $\beta_1 : \Psi_1 \rightarrow \Theta_1$ by $\beta_1(\theta, \pi_1) = (\theta, \beta(\pi_1))$. Then, we say that $\tilde{\mu}_2 \in \mathcal{L}(\Theta_1)$ is *second order epistemically equivalent* to $\pi_2 \in \mathcal{C}(\Psi_1, \mathfrak{F}_1)$ given β_1 whenever $\Lambda(\tilde{\mu}_2) = \Lambda(\pi_2) = n$, and $\mu_2^m(E_1) = \pi_2(\beta_1^{-1}(E_1) \mid B_{\pi_2}^m)$, for all Borel events $E_1 \subseteq \Theta_1$ and for every $m = 1, \dots, n$.

Notice that in order to define second order epistemic equivalence, it is necessary to first introduce the Borel function β_1 . Otherwise, the Borel event $E_1 \subseteq \Theta_1$ could not be expressed as a Borel event in Ψ_1 , and therefore second order conditional beliefs about E_1 could not be expressed in the first place.

Higher order epistemic equivalence is inductively defined in an analogous way. Firstly, some more notation is needed: Let $T_{\Psi_k}^C := \text{Proj}_{\Psi_k}(\Theta \times T^C)$ and $T_{\Theta_k}^L := \text{Proj}_{\Theta_k}(\Theta \times T^L)$, and notice that specifically for $k = 1$ it is the case that $T_{\Psi_1}^C = \Psi_1$ and $T_{\Theta_1}^L = \Theta_1$.

Definition 10. Fix an arbitrary $k > 0$, and consider $(\theta, \tilde{\mu}_1, \dots, \tilde{\mu}_{k+1}) \in T_{\Theta_{k+1}}^L$ and $(\theta, \pi_1, \dots, \pi_{k+1}) \in T_{\Psi_{k+1}}^C$. Suppose that there exists a Borel surjective function $\beta_k : T_{\Psi_k}^C \rightarrow T_{\Theta_k}^L$. Then, we say that $\tilde{\mu}_{k+1}$ is *(k+1)-th order epistemically equivalent* to π_{k+1} given β_k whenever

$$(E_1^k) \quad \Lambda(\tilde{\mu}_{k+1}) = \Lambda(\pi_{k+1}) = n,$$

$$(E_2^k) \quad \mu_{k+1}^m(E_k) = \pi_{k+1}(\beta_k^{-1}(E_k) \mid B_{\pi_{k+1}}^m), \text{ for all Borel events } E_k \subseteq T_{\Theta_k}^L \text{ and every } m = 1, \dots, n.$$

Furthermore, if it is also the case that $\beta_k(\theta, \pi_1, \dots, \pi_k) = (\theta, \tilde{\mu}_1, \dots, \tilde{\mu}_k)$, we say that $(\theta, \tilde{\mu}_1, \dots, \tilde{\mu}_{k+1})$ is *up to (k+1)-th order epistemically equivalent* to $(\theta, \pi_1, \dots, \pi_{k+1})$ given β_k .

Notice that the previous definition restricts attention to hierarchies that satisfy coherency and common certainty in coherency, i.e., the function β_k maps elements of $T_{\Psi_k}^C$ to elements of $T_{\Theta_k}^L$, rather than elements of Ψ_k to elements of Θ_k . This is a natural restriction as – at least to the best of our knowledge – the entire literature focuses exclusively on agents whose belief hierarchies have this type of internal consistency.

Before moving forward, let us elaborate on the previous definition. First, observe that higher order epistemic equivalence is defined inductively. As we have already mentioned, this is necessary in order to make sure that k -th order lexicographic beliefs can be translated to k -th order conditional beliefs, before associating $(k + 1)$ -th order lexicographic beliefs with $(k + 1)$ -th order conditional beliefs. Then, once we have introduced the function β_k , our notion of $(k + 1)$ -th order epistemic equivalence follows the same logic as the notion of epistemic equivalence that was introduced and discussed in Section 3.

Since $(k + 1)$ -th order epistemic equivalence relies on the existence of a Borel surjective function β_k , it is important to make sure that such a Borel surjective function exists for every $k > 0$.

Lemma 3. *There exists a sequence of Borel surjective functions $\{\beta_k : T_{\Psi_k}^C \rightarrow T_{\Theta_k}^L\}_{k=1}^\infty$ such that*

$$(i) \quad \beta_1(\theta, \pi_1) = (\theta, \beta(\pi_1)) \text{ for all } (\theta, \pi_1) \in T_{\Psi_1}^C,$$

$$(ii) \quad (\theta, \pi_1, \dots, \pi_k) \text{ is up to } k\text{-th order epistemically equivalent to } \beta_k(\theta, \pi_1, \dots, \pi_k) \text{ given } \beta_{k-1} \text{ for all } (\theta, \pi_1, \dots, \pi_k) \in T_{\Psi_k}^C, \text{ and for each } k > 1.$$

Henceforth, whenever we refer to a sequence $\{\beta_k\}_{k=1}^\infty$ of Borel surjective functions, we implicitly assume that it satisfies conditions (i) – (ii) of the previous lemma. Then, following the previous result, we can now define a notion of epistemic equivalence between a hierarchy of lexicographic beliefs and a hierarchy of conditional beliefs.

Definition 11. We say that an \mathcal{L} -hierarchy $(\tilde{\mu}_1, \tilde{\mu}_2, \dots) \in T^L$ is *epistemically equivalent* to a \mathcal{C} -hierarchy $(\pi_1, \pi_2, \dots) \in T^C$ if for some sequence $\{\beta_k\}_{k=1}^\infty$ of Borel surjective functions, it is the case that $\beta_k(\theta, \pi_1, \dots, \pi_k) = (\theta, \tilde{\mu}_1, \dots, \tilde{\mu}_k)$ for each $k > 0$.

Then, the following result follows directly from Lemma 3. The proof relies on the fact that for every $k > 0$ if $(\pi_1, \pi_2, \dots) \in T^C$ then $(\theta, \pi_1, \dots, \pi_k) \in T_{\Psi_k}^C$, and likewise if $(\tilde{\mu}_1, \tilde{\mu}_2, \dots) \in T^L$ then $(\theta, \tilde{\mu}_1, \dots, \tilde{\mu}_k) \in T_{\Theta_k}^L$.

Theorem 1. *There exists a Borel surjective function $h : T^C \rightarrow T^L$ such that (π_1, π_2, \dots) is epistemically equivalent to $h(\pi_1, \pi_2, \dots)$ for all $(\pi_1, \pi_2, \dots) \in T^C$.*

The fact that the function h is Borel implies that we can map Borel events from the space of lexicographic belief hierarchies to Borel events in the space of conditional belief hierarchies, implying that the language used by players whose reasoning is modeled with lexicographic probability systems can be translated to the language used by players whose reasoning is modeled with conditional probability systems.

It follows from h being a surjective function that every \mathcal{L} -hierarchy is the image of some epistemically equivalent \mathcal{C} -hierarchy, even if the \mathcal{L} -hierarchy is such that agent is uncertain about the length of the opponent's first order lexicographic beliefs. Then, it may be the case that the only epistemically equivalent \mathcal{C} -hierarchy has the property that the agent is uncertain about the opponent's collection of conditioning events. The latter illustrates why we have extended the standard construction of Battigalli and Siniscalchi (1999) in a way that permits the agent to be uncertain about the opponent's $\mathcal{B} \in \mathfrak{F}$. We further elaborate on this issue in Section 6.

5.2. Equivalence of types spaces

Recall the standard way of representing \mathcal{C} -hierarchies and \mathcal{L} -hierarchies via the corresponding type space models. A natural question that arises at this point is whether there exist conditions imposed directly on the type space models which lead to epistemic equivalence between the conditional belief hierarchy associated with a \mathcal{C} -type and the lexicographic belief hierarchy associated with an \mathcal{L} -type. The following result shows that this is in fact possible.

Theorem 2. *Consider two Polish spaces T_a and T_b , and let $(\Theta, T_a, T_b, g_a, g_b)$ and $(\Theta, T_a, T_b, \tilde{\lambda}_a, \tilde{\lambda}_b)$ be a conditional type space model and a lexicographic type space model respectively, such that the LPS $\tilde{\lambda}_i(t_i) \in \mathcal{L}_\Theta(\Theta \times T_j)$ is epistemically equivalent to the CPS $g_i(t_i) \in \mathcal{C}(\Theta \times T_j, \mathfrak{F}_{T_j})$ for all $t_i \in T_i$ and all $i \in \{a, b\}$. Then, $(\tilde{\mu}_1(t_i), \tilde{\mu}_2(t_i), \dots) \in T^L$ is epistemically equivalent to $(\pi_1(t_i), \pi_2(t_i), \dots) \in T^C$ for every $t_i \in T_i$.*

In the previous theorem, notice that for each player i the space of \mathcal{C} -types coincides with the space of \mathcal{L} -types. This is not necessary, as long as the two are homeomorphic. In either case, a type is an abstract object which has some specific meaning only in the context of a type space model. Therefore, using the same type space of types in the two models or simply using different (homeomorphic) type spaces does not affect our analysis.

An interesting consequence of the previous result is that we do not need to use the functions $\{\beta_k\}_{k=1}^\infty$ in order to verify that an \mathcal{L} -hierarchy is epistemically equivalent to a \mathcal{C} -hierarchy.

5.3. Terminal lexicographic type space

So far in the literature there is no general result about the existence of a terminal lexicographic type space, i.e., one inducing all \mathcal{L} -hierarchies from T^L .¹⁷ One of the difficulties with this construction is the fact that there is no general result about the topological structure of $\mathcal{L}(X)$, even when X is a Polish space.¹⁸ In this section, we use the equivalence results established above, together with the homeomorphism of Proposition 3, to prove the existence of such a large type space.

Theorem 3. *There exists a Polish type space T and a Borel surjective function $\tilde{\lambda} : T \rightarrow \mathcal{L}_\Theta(\Theta \times T)$ such that for each $(\tilde{\mu}_1, \tilde{\mu}_2, \dots) \in T^L$ there exists some $t \in T$ with $(\tilde{\mu}_1(t), \tilde{\mu}_2(t), \dots) = (\tilde{\mu}_1, \tilde{\mu}_2, \dots)$.*

Notice that the function $\tilde{\lambda}$ is not injective, implying that there may exist redundant types in T . This is due to the fact that the previous result uses the equivalence result established in Theorem 1, according to which the function $h : T^C \rightarrow T^L$ is surjective but not injective. An open question for future research is to construct a complete lexicographic type space model without redundant types.¹⁹

5.4. Common assumption versus common strong belief

In the standard model of belief hierarchies, one concept of particular interest is common belief. Accordingly, an event is commonly believed if everybody believes it, everybody believes that everybody

¹⁷There is no consensus in the literature regarding a term that describes a type space with the property that for every belief hierarchy there exists a type inducing this hierarchy. In this paper, we follow Friedenbergr (2010) and Perea (2012, p. 131) who call such a type space terminal.

¹⁸Recall that $\mathcal{L}(X)$ is Borel (Brandenburger et al., 2008, Cor. C.1). However, we do not know whether it is closed in $\mathcal{N}(X)$ or not, implying that $\mathcal{L}(X)$ may not be Polish, and therefore we cannot directly apply the Kolmogorov extension theorem.

¹⁹In a slightly different framework that allows first order beliefs to violate mutual singularity, Catonini (2012) constructs a non-redundant terminal type space model for lexicographic belief hierarchies. The difference is that in his setting, lexicographic beliefs are not restricted to satisfy (L_1) , implying that they form a Polish space. Therefore, the Kolmogorov extension Theorem can be used, similarly to Brandenburger and Dekel (1993) and unlike this paper.

believes it, and so on. The analogue of common belief in an environment where agents reason according to lexicographic belief hierarchies is common assumption, whereas if agents reason according to conditional belief hierarchies the corresponding notion is common strong belief.

The question that naturally arises at this point is whether our notion of epistemic equivalence between \mathcal{C} -hierarchies and \mathcal{L} -hierarchies implies equivalence between common strong belief and common assumption. In other words, does common strong belief of a Borel event $E \subseteq \Theta \times T_a \times T_b$ at some state imply common assumption of the same event, and vice versa? In general, the answer is negative, as according to Brandenburger et al. (2007, Ex. 3), an event may be strongly believed by a (non-full-support) CPS and still not assumed by the equivalent LPS. Thus, we restrict attention to full-support beliefs.

Proposition 4. *Consider a countable Θ and two Polish spaces T_a and T_b , and let $(\Theta, T_a, T_b, g_a, g_b)$ and $(\Theta, T_a, T_b, \tilde{\lambda}_a, \tilde{\lambda}_b)$ be a conditional type space model and a lexicographic type space model such that $\tilde{\lambda}_i(t_i) \in \mathcal{L}_\Theta^+(\Theta \times T_j)$ is epistemically equivalent to $g_i(t_i) \in \mathcal{C}^+(\Theta \times T_j, \mathfrak{F}_{T_j})$ for all $t_i \in T_i$ and every $i \in \{a, b\}$. Consider some Borel event $E \subseteq \Theta \times T_a \times T_b$. Then, at some epistemic state (t_a, t_b) , the event E is commonly strongly belief if and only if it is commonly assumed.*

The previous result is the first one in the literature – at least to our knowledge – directly relating common assumption and common strong belief. Given that the two notions are extensively used for the characterization of various solution concepts, we believe that Proposition 4 can have important implications for the analysis of games, e.g., if we fix a conditional type space model and a lexicographic type space model satisfying the equivalence conditions stated in the previous result, “rationality and common strong belief of rationality” (RCSBR) coincides with “rationality and common assumption in rationality” (RCAR). Of course, our result does not resolve the tension between the impossibility result in Brandenburger et al. (2008, Thm. 10.1) and the positive result in Battigalli and Siniscalchi (2002, Prop. 6), as the type space models in Proposition 4 differ from the ones considered in these two papers.²⁰ Moreover, we do not provide conditions under which RCAR and RCSBR are empty events. What we do instead, is to provide conditions under which the two concepts are essentially equivalent. We further discuss the relationship between RCAR and RCSBR in the next section.

²⁰Brandenburger et al. (2008) showed that in complete continuous lexicographic type space models, RCAR is an empty event. On the other hand, Battigalli and Siniscalchi (2002) showed that in complete continuous conditional type space models, RCSBR is non-empty. For the precise definitions of complete continuous type space models, we refer to the corresponding papers, while for a discussion on the relationship between the two results, we refer to the supplementary material of Brandenburger et al. (2008).

6. Discussion

6.1. Mutual singularity

Recall that mutual singularity captures the idea that the supports of the different theories of an LPS do not overlap “much”. This assumption has been already motivated for countable spaces with the discrete topology, in which case it is equivalent to assuming that the supports of the different theories are disjoint. This equivalence would not hold in general Polish spaces, in which case the condition of disjoint supports is only sufficient for mutual singularity, but not necessary. In fact, generalizing the decision-theoretic foundation of Blume et al. (1991a) for LPS’s that satisfy (L_1) in an arbitrary space remains an open problem.

Nonetheless, throughout this paper, we only consider \mathcal{L} -hierarchies $(\tilde{\mu}_1, \tilde{\mu}_2, \dots) \in T^L$ such that for every $k > 0$ it is the case that the different theories $(\mu_k^1, \dots, \mu_k^n)$ of $\tilde{\mu}_k$ have disjoint supports. This follows from the fact that (i) Θ has been assumed to be a countable space with the discrete topology,²¹ and (ii) coherency is imposed. In particular, as we have already mentioned in Section 3, the fact that Θ is countable implies that $\tilde{\mu}_1 \in \mathcal{L}(\Theta)$ satisfies the condition of non-overlapping supports. Furthermore, it follows from the standard coherency assumption that for every $k > 1$ and for every $m = 1, \dots, n$, it is the case that $\text{Supp}(\mu_k^m) \subseteq \text{Supp}(\mu_1^m) \times \mathcal{L}(\Theta_0) \times \dots \times \mathcal{L}(\Theta_{k-2})$. Intuitively, this is because, as we have already pointed out in Section 4, in practice we only impose mutual singularity on the beliefs about the underlying space of uncertainty, and not about the opponents beliefs.

Finally, let us stress that so far, there is no consensus in the literature when it comes to assuming mutual singularity or not, i.e., some game-theoretic applications impose it (e.g., Brandenburger et al., 2007, 2008; Heifetz et al., 2010; Keisler and Lee, 2011), while others do not (e.g., Blume et al., 1991b).

6.2. Cardinality and topology of the underlying space of uncertainty

Throughout the paper, we have mostly focused on a countable underlying space of uncertainty Θ endowed with the discrete topology. The reason for doing so is a technical one, and stems from the fact that if Θ was assumed to be uncountable, the collection of finite algebras, and therefore \mathfrak{F} would not be a countable set. Hence, the space of finitary CPS’s would become the topological sum of uncountably many Polish spaces, and therefore it would not be necessarily Polish. As a consequence our analysis would not be valid any more.

At the same time, assuming that Θ is endowed with the discrete topology is also necessary, as otherwise there would exist some collections $\mathcal{B} \in \mathfrak{F}$ that would not consist of clopen conditioning

²¹See the next section for further discussion on this assumption.

events, and the latter is crucial for our results, as well as the ones by Battigalli and Siniscalchi (1999). For instance, in such a case we would not be able to establish the existence of the homeomorphism g from Proposition 3. In any case, though we recognize that there are certain applications in economics with an uncountable space of parameters, we still find that our results sufficiently general to accommodate most interesting cases.

6.3. Uncertainty about the opponent's collection of conditioning events

As it has been already stated, T^C contains hierarchies of conditional beliefs that express uncertainty about the opponent's collection of conditioning events. On the one hand, such a generalization of Battigalli and Siniscalchi (1999) is interesting from a game-theoretic point of view, as it provides a framework for studying dynamic games with players who are uncertain of the opponents' information partition (e.g., Zuazo-Garin, 2013). On the other hand, as it has been already mentioned, extending the model in a way that allows players to have this additional type of uncertainty is necessary for a technical reason. Namely, our main aim in this paper is to relate lexicographic belief hierarchies and conditional belief hierarchies, which we do in Theorem 2, and this would not have been possible without allowing this additional layer of uncertainty. The reason is that there are $t_\ell \in T^L$ such that every epistemically equivalent $t_c \in h^{-1}(t_\ell)$ is associated with a CPS over $\Theta \times T^C$ that deems possible \mathcal{C} -types of the opponent with different collections of conditioning events. This is often the case, when t_ℓ is associated with an LPS that deems possible \mathcal{L} -types of the opponent with different lengths. Let us illustrate this with an example.

Example 3. Let $\Theta = \{\theta_1, \theta_2\}$ and suppose that a deems equally likely that (i) b 's first order lexicographic beliefs have only one theory assigning probability 1 to $\{\theta_1\}$, and (ii) b 's first order lexicographic beliefs have two theories with the primary one assigning probability 1 to $\{\theta_1\}$ and the secondary one assigning probability 1 to $\{\theta_2\}$. Now, the only way to construct b 's first order conditional beliefs which are epistemically equivalent to b 's first order lexicographic beliefs in (i), is to let b 's collection of conditioning hypotheses be $\{\{\theta_1, \theta_2\}\}$, as otherwise the length of the LPS capturing these beliefs would be equal to 2. At the same time, in order to construct b 's first order conditional beliefs which are epistemically equivalent to b 's first order lexicographic beliefs in (ii), we need to let b 's collection of conditioning hypotheses be $\{\{\theta_1, \theta_2\}, \{\theta_1\}, \{\theta_2\}\}$. Thus, in order to construct a \mathcal{C} -hierarchy of a that is epistemically equivalent to a 's \mathcal{L} -hierarchy that contains these second order lexicographic beliefs, it is necessary to allow a to be uncertain about b 's collection of conditioning events. ◁

6.4. Common strong belief in rationality versus common assumption in rationality

One of the recent questions within epistemic game theory is to further clarify the relationship between the impossibility result from Brandenburger et al. (2008, Thm. 10.1), according to which the event RCAR is empty in a complete continuous lexicographic type space model, with the positive result in Battigalli and Siniscalchi (2002, Prop. 6), according to which the event RCSBR is non-empty in a complete continuous conditional type space model. As we have already mentioned, our Proposition 4 does not resolve the tension between these two results, but merely provides conditions under which the two concepts are equivalent, and therefore yield the same predictions. However, it is important to point out that many of our conditions are not present in either Brandenburger et al. (2008), or in Battigalli and Siniscalchi (2002).

Firstly, notice that we allow for uncertainty over the opponent's collection of conditioning events, contrary to Battigalli and Siniscalchi (2002). Secondly, unlike Brandenburger et al. (2008), we restrict attention to lexicographic type spaces with the property that the marginal distributions over the underlying space of uncertainty form an LPS. Thirdly, observe that Battigalli and Siniscalchi (2002) do not restrict attention to full-support \mathcal{C} -types, implying that common strong belief of rationality may not imply common assumption of rationality in the corresponding \mathcal{L} -type space. Finally, the negative result of Brandenburger et al. (2008) relies on the lexicographic type space being complete, implying that there exists at least one type for each player associated with a non-full support LPS, which is not the case in Proposition 4.

Concluding, although further research is needed in order to clearly understand the relationship between the two notions, we can conclude that RCAR and RCSBR are not equivalent in general, but only under additional conditions. Still, our result provides a new tool for studying solution concepts, e.g., sufficient epistemic conditions for a solution concept in terms of RCSBR in a \mathcal{C} -type space with full-support types, would immediately yield the corresponding sufficient conditions for this solution concept in terms of RCAR.

A. Proofs of Section 3

Proof of Lemma 1. Consider arbitrary $A \in \mathcal{F}$ and $B \in \mathcal{B}$. If $\pi(A \cap B|B_\pi^m) = \rho(A \cap B|B_\rho^m) = 0$ for all $m = 1, \dots, n$, then it follows directly that $\pi(A|B) = \rho(A|B) = 0$. Suppose instead that m is the smallest number in $\{1, \dots, n\}$ such that $\pi(A \cap B|B_\pi^m) = \rho(A \cap B|B_\rho^m) > 0$. Observe that $A \cap B \in \mathcal{F}$ and $B_\pi^m \in \mathcal{B}$, while at the same time $A \cap B \subseteq B \subseteq B_\pi^m$. Then, it follows by (C_3) that $\pi(A \cap B|B_\pi^m) = \pi(A \cap B|B) \cdot \pi(B|B_\pi^m)$, implying that $\pi(A|B) = \pi(A \cap B|B_\pi^m) / \pi(B|B_\pi^m)$. Likewise, we show that $\rho(A|B) = \rho(A \cap B|B_\rho^m) / \rho(B|B_\rho^m)$.

Observe that the right-hand sides of the previous two equations are equal, implying that $\pi(A|B) = \rho(A|B)$, which completes the proof. \square

Proof of Proposition 1. Firstly, we show that for every $\pi \in \mathcal{C}(X)$ there is a unique $\tilde{\mu} \in \mathcal{L}(X)$ such that $\tilde{\mu}$ is equivalent to π : Let $\pi \in \Delta^{\mathcal{B}}(X)$ for some $\mathcal{B} \in \mathfrak{F}$. Define, $\tilde{\mu} = (\mu^1, \dots, \mu^n) \in \mathcal{N}_n(X)$ by $\mu^m := \pi(\cdot|B_\pi^m)$, and observe that the collection $\{(B_\pi^1 \setminus B_\pi^2), \dots, (B_\pi^{n-1} \setminus B_\pi^n), B_\pi^n\}$ of Borel events satisfies (L_1) , implying that $\tilde{\mu} \in \mathcal{L}(X)$. Moreover, notice that by construction $\Lambda(\tilde{\mu}) = \Lambda(\pi)$. Suppose that there exists some other $\tilde{\nu} \in \mathcal{L}_n(X)$ which is equivalent to π . Then, it follows by definition that $\nu^m = \pi(\cdot|B_\pi^m)$ for all $m = 1, \dots, n$, implying that $\mu^m = \nu^m$. Secondly, we show that for every $\tilde{\mu} \in \mathcal{L}_n(X)$ there is some equivalent $\pi \in \mathcal{C}(X)$: Consider some $\tilde{\mu} = (\mu^1, \dots, \mu^n) \in \mathcal{L}_n(X)$. Then, by (L_1) , there is a collection of Borel events $\{A_1, \dots, A_n\}$ such that $\mu^m(A_m) = 1$ and $\mu^m(A_\ell) = 0$. Define the partition $\mathcal{P} := \{B_1, \dots, B_n\}$ by $B_m := A_m \setminus (\bigcup_{\ell \neq m} A_\ell)$ for each $m = 1, \dots, n-1$ and $B_n := A_n \cup (X \setminus (B_1 \cup \dots \cup B_{n-1}))$. Let the finitely generated collection \mathcal{B} be the closure of \mathcal{P} with respect to union, and define $\mathcal{B}_\pi = \{B_\pi^1, \dots, B_\pi^n\}$ by $B_\pi^m := X \setminus (\bigcup_{\ell < m} B_\ell)$. Now, define $\pi : \mathcal{F} \times \mathcal{B} \rightarrow [0, 1]$ such that $\pi(\cdot|B_\pi^m) := \mu^m$ for all $m = 1, \dots, n$. For every $B \in \mathcal{B}$, consider the smallest $m \in \{1, \dots, n\}$ such that $\pi(B|B_\pi^m) > 0$, which by construction always exists. Then, for every $A \in \mathcal{F}$, let $\pi(A|B) := \pi(A \cap B|B_\pi^m)/\pi(B|B_\pi^m)$. Finally, verify that $\pi \in \Delta^{\mathcal{B}}(X)$, with \mathcal{B}_π being the collection of π -relevant events, which completes the proof. \square

Proof of Proposition 2. First, we show that $\Delta_\pi^{\mathcal{B}}(X) := \{\rho \in \Delta^{\mathcal{B}}(X) : \mathcal{B}_\rho = \mathcal{B}_\pi\}$ is closed in $[\Delta(X)]^{\mathcal{B}}$. Observe that $\Delta_\pi^{\mathcal{B}}(X)$ can be rewritten as

$$\Delta^{\mathcal{B}}(X) \cap \underbrace{\left(\{p \in \Delta(X) : p(B_\pi^1 \setminus B_\pi^2) \geq 1\} \times \dots \times \{p \in \Delta(X) : p(B_\pi^n) \geq 1\} \right)}_{\text{given } \pi\text{-relevant events}} \times \underbrace{\Delta(X) \times \dots \times \Delta(X)}_{\text{given other events in } \mathcal{B}}.$$

Since X is countable, it follows that \mathcal{B} is a collection of clopen events, and therefore $\Delta^{\mathcal{B}}(X)$ is closed in $[\Delta(X)]^{\mathcal{B}}$ (Battigalli and Siniscalchi, 1999, Lem. 1). Moreover, it follows from Aliprantis and Border (1994, Cor. 15.6) that $\{p \in \Delta(X) : p(B_\pi^m \setminus B_\pi^{m+1}) \geq 1\}$ is closed in $\Delta(X)$ for every $m = 1, \dots, n-1$, and therefore $\{p \in \Delta(X) : p(B_\pi^1 \setminus B_\pi^2) \geq 1\} \times \dots \times \{p \in \Delta(X) : p(B_\pi^n) \geq 1\} \times \Delta(X) \times \dots \times \Delta(X)$ is also closed in $[\Delta(X)]^{\mathcal{B}}$, which proves our claim. Now, let $E \subseteq \mathcal{L}(X)$ be Borel, and notice that

$$\begin{aligned} \beta^{-1}(E) &= \bigcup_{n=1}^{\infty} \beta^{-1}(E \cap \mathcal{N}_n(X)) \\ &= \bigcup_{n=1}^{\infty} \bigcup_{\mathcal{B} \in \mathfrak{F}} \bigcup_{\Delta_\pi^{\mathcal{B}}(X) \subseteq \Delta^{\mathcal{B}}(X)} \left(\beta^{-1}((E \cap \mathcal{N}_n(X)) \cap \Delta_\pi^{\mathcal{B}}(X)) \right) \\ &= \bigcup_{n=1}^{\infty} \bigcup_{\mathcal{B} \in \mathfrak{F}} \bigcup_{\Delta_\pi^{\mathcal{B}}(X) \subseteq \Delta^{\mathcal{B}}(X)} \{ \rho \in \Delta_\pi^{\mathcal{B}}(X) : \beta(\rho) \in E \cap \mathcal{N}_n(X) \} \end{aligned} \quad (1)$$

Fix some $n > 0$ and observe that $E \cap \mathcal{N}_n(X)$ is a Borel subset of $\mathcal{N}_n(X)$. Since X is separable and metrizable, $E \cap \mathcal{N}_n(X)$ is generated by Borel rectangles of the form of

$$\{p \in \Delta(X) : p(A_1) \geq \alpha_1\} \times \dots \times \{p \in \Delta(X) : p(A_n) \geq \alpha_n\} \subseteq \mathcal{N}_n(X).$$

where $A_m \subseteq X$ is Borel and $\alpha_m \in [0, 1]$ for each $m = 1, \dots, n$. Therefore, for any $\mathcal{B} \in \mathfrak{F}$ and an arbitrary $\pi \in \Delta^{\mathcal{B}}(X)$, the event $\{ \rho \in \Delta_{\pi}^{\mathcal{B}}(X) : \beta(\rho) \in E \cap \mathcal{N}_n(X) \}$ is generated by events of the form

$$\begin{aligned} & \{ \rho \in \Delta_{\pi}^{\mathcal{B}}(X) : \beta(\rho) \in \{p \in \Delta(X) : p(A_1) \geq \alpha_1\} \times \dots \times \{p \in \Delta(X) : p(A_n) \geq \alpha_n\} \} \\ &= \{ \rho \in \Delta_{\pi}^{\mathcal{B}}(X) : \rho(A_1|B_{\pi}^1) \geq \alpha_1 \} \cap \dots \cap \{ \rho \in \Delta_{\pi}^{\mathcal{B}}(X) : \rho(A_n|B_{\pi}^n) \geq \alpha_n \} \\ &= \Delta_{\pi}^{\mathcal{B}}(X) \cap \underbrace{\left(\underbrace{\{p \in \Delta(X) : p(A_1) \geq \alpha_1\} \times \dots \times \{p \in \Delta(X) : p(A_n) \geq \alpha_n\}}_{\text{beliefs given } \pi\text{-relevant events}} \times \underbrace{\Delta(X) \times \dots \times \Delta(X)}_{\text{given other events in } \mathcal{B}} \right)}_{\text{subset of } [\Delta(X)]^{\mathcal{B}}}. \end{aligned}$$

Since X is separable and metrizable it follows that $\{p \in \Delta(X) : p(A_m) \geq \alpha_m\}$ is Borel in $\Delta(X)$, and therefore $\{p \in \Delta(X) : p(A_1) \geq \alpha_1\} \times \dots \times \{p \in \Delta(X) : p(A_n) \geq \alpha_n\} \times \Delta(X) \times \dots \times \Delta(X)$ is Borel in $[\Delta(X)]^{\mathcal{B}}$, which together with the fact that $\Delta_{\pi}^{\mathcal{B}}(X)$ is closed, implies that $\{ \rho \in \Delta_{\pi}^{\mathcal{B}}(X) : \beta(\rho) \in E \cap \mathcal{N}_n(X) \}$ is generated by Borel events, and therefore it is Borel itself. Finally, it follows from X being countable that \mathfrak{F} is also countable. Hence, by Eq. (1), $\beta^{-1}(E)$ is Borel which completes the proof. \square

B. Proofs of Section 4

Proof of Lemma 2. First, we inductively show that every Ψ_k is Polish. It holds by assumption that Ψ_0 is Polish, and suppose that Ψ_k is also Polish. Observe that Ψ_{k+1} can be rewritten as $\Psi_k \times (\bigoplus_{\mathcal{B} \in \mathfrak{F}} \Delta^{\mathcal{B}}(\Psi_k))$. Recall that, since \mathcal{B} contains only clopen events, $\Delta^{\mathcal{B}}(\Psi_k)$ is Polish (Battigalli and Siniscalchi, 1999, Lem. 1). Furthermore, notice that \mathfrak{F} is countable, implying that $\bigoplus_{\mathcal{B} \in \mathfrak{F}} \Delta^{\mathcal{B}}(\Psi_k)$ is Polish, and therefore so is Ψ_{k+1} . Now, consider an arbitrary $(\pi_1, \pi_2, \dots) \in T_1^C$ such that $\pi_k \in \Delta^{\mathcal{B}}(\Psi_{k-1})$, and fix some $B \in \mathcal{B}$. Then, it follows from Kolmogorov extension theorem (Aliprantis and Border, 1994, Cor. 15.27) that there exists a unique probability measure $\pi(\cdot|B) \in \Delta(\Theta \times T_0^C)$ such that $\text{marg}_{\Psi_{k-1}} \pi(\cdot|B) = \pi_k(\cdot|B)$ for all $k > 0$. Now consider the collection of probability measures $\pi := (\pi(\cdot|B) ; B \in \mathcal{B}) \in [\Delta(\Theta \times T_0^C)]^{\mathcal{B}}$, and we show that π is a CPS. First, notice that (C_1) and (C_2) are trivially satisfied. Then, consider a Borel event $A \subseteq \Theta \times T_0^C$, and two conditioning events $B, C \in \mathcal{B}$ such that $B \subseteq C$. Observe that $A = \bigcap_{k=0}^{\infty} (\text{Proj}_{\Psi_k} A \times \prod_{\ell=k}^{\infty} \mathcal{C}(\Psi_{\ell}, \mathfrak{F}_{\ell}))$, thus yielding

$$\begin{aligned} \pi(A|C \times T_0^C) &= \pi \left(\bigcap_{k=0}^{\infty} \left(\text{Proj}_{\Psi_k} A \times \prod_{\ell=k}^{\infty} \mathcal{C}(\Psi_{\ell}, \mathfrak{F}_{\ell}) \right) \mid C \times T_0^C \right) \\ &= \lim_{k \rightarrow \infty} \pi \left(\text{Proj}_{\Psi_k} A \times \prod_{\ell=k}^{\infty} \mathcal{C}(\Psi_{\ell}, \mathfrak{F}_{\ell}) \mid C \times T_0^C \right) \\ &= \lim_{k \rightarrow \infty} \pi_{k+1} \left(\text{Proj}_{\Psi_k} A \mid C \times \prod_{\ell=0}^k \mathcal{C}(\Psi_{\ell}, \mathfrak{F}_{\ell}) \right) \\ &= \lim_{k \rightarrow \infty} \pi_{k+1} \left(\text{Proj}_{\Psi_k} A \mid B \times \prod_{\ell=0}^k \mathcal{C}(\Psi_{\ell}, \mathfrak{F}_{\ell}) \right) \cdot \pi_{k+1} \left(B \times \prod_{\ell=0}^k \mathcal{C}(\Psi_{\ell}, \mathfrak{F}_{\ell}) \mid C \times \prod_{\ell=0}^k \mathcal{C}(\Psi_{\ell}, \mathfrak{F}_{\ell}) \right) \\ &= \pi(A|B \times T_0^C) \cdot \pi(B \times T_0^C|C \times T_0^C), \end{aligned}$$

implies that π satisfies (C_3) . The latter induces a function $f : T_1^C \rightarrow \mathcal{C}(\Theta \times T_0^C, \mathfrak{F}_{T_0^C})$. Then, we show that f is an injection: Suppose that $(\pi_1, \pi_2, \dots) \in T_1^C$ and $(\pi'_1, \pi'_2, \dots) \in T_1^C$ are such that $f(\pi_1, \pi_2, \dots) = f(\pi'_1, \pi'_2, \dots)$. Recall that

$$\begin{aligned} \pi_{k+1} &= \text{marg}_{\Psi_k} f(\pi_1, \pi_2, \dots) \\ &= \text{marg}_{\Psi_k} f(\pi'_1, \pi'_2, \dots) \\ &= \pi'_{k+1} \end{aligned}$$

for all $k \geq 0$, implying that $(\pi_1, \pi_2, \dots) = (\pi'_1, \pi'_2, \dots)$, thus proving that f is injective. Then, we show that f is a surjection: Let $\pi \in \Delta^{\mathcal{B}}(\Theta \times T_0^C)$ for some $\mathcal{B} \in \mathfrak{F}$, and define the sequence $(\text{marg}_{\Psi_0} \pi, \text{marg}_{\Psi_1} \pi, \dots)$. Observe that by construction $\text{marg}_{\Psi_k} \pi \in \Delta^{\mathcal{B}}(\Psi_k)$, and furthermore $(\text{marg}_{\Psi_0} \pi, \text{marg}_{\Psi_1} \pi, \dots)$ satisfies coherency, implying that $f(\text{marg}_{\Psi_0} \pi, \text{marg}_{\Psi_1} \pi, \dots) = \pi$, which proves that f is surjective. Finally, proving that f and f^{-1} are continuous is done identically to Battigalli and Siniscalchi (1999, Prop. 1). \square

Proof of Proposition 3. Let $T_{\mathcal{B}}^C := \{t \in T_1^C : f(t) \in \Delta^{\mathcal{B}}(\Theta \times T_0^C)\}$, and observe that $T^C = \bigcup_{\mathcal{B} \in \mathfrak{F}} \{t \in T_{\mathcal{B}}^C : f^{\mathcal{B}}(t)(\Theta \times T^C) = 1, \forall \mathcal{B} \in \mathcal{B}\}$. The remainder of the proof replicates the one of Battigalli and Siniscalchi (1999, Prop. 2). \square

C. Proof of Section 5

Proof of Lemma 3. We proceed inductively. Firstly observe that since $T_{\Psi_1}^C = \Psi_1$ and $T_{\Theta_1}^L = \Theta_1$, it follows directly from Propositions 1 and 2 that there exists a Borel surjective function $\beta_1 : T_{\Psi_1}^C \rightarrow T_{\Theta_1}^L$ defined as in (i). Secondly, consider an arbitrary $(\theta, \pi_1, \pi_2) \in T_{\Psi_2}^C$ with $\Lambda(\pi_1) = \Lambda(\pi_2) = n$, and define $(\theta, \tilde{\mu}_1, \tilde{\mu}_2) \in \Theta \times \mathcal{N}_n(\Theta_0) \times \mathcal{N}_n(\Theta_1)$ such that $(\theta, \tilde{\mu}_1) = \beta_1(\theta, \pi_1)$ and

$$\mu_2^m(E_1) = \pi_2(\beta_1^{-1}(E_1) \mid B_{\pi_2}^m)$$

for all Borel $E_1 \subseteq T_{\Theta_1}^L$ and every $m = 1, \dots, n$. Observe that since we already have the function β_1 , such a $(\theta, \tilde{\mu}_1)$ exists, and moreover it belongs to $T_{\Theta_1}^L$. Then, we move on to show that $\tilde{\mu}_2 \in \mathcal{L}(\Theta_1)$: It follows from $(\theta, \pi_1, \pi_2) \in T_{\Psi_2}^C$ that the π_2 -relevant events $\mathcal{B}_{\pi_2} = \{B_{\pi_2}^1, \dots, B_{\pi_2}^n\} = \{B^1 \times \mathcal{C}(\Psi_0, \mathfrak{F}_0), \dots, B^n \times \mathcal{C}(\Psi_0, \mathfrak{F}_0)\}$ are such that $\{B^1, \dots, B^n\} = \mathcal{B}_{\pi_1}$. Then, define the collection $\mathcal{F}_1 := \{A_1^1, \dots, A_1^n\}$ of subsets of Θ_1 by

$$\begin{aligned} A_1^\ell &:= (B^\ell \setminus B^{\ell+1}) \times \mathcal{L}(\Theta_0) \\ &= \beta_1((B^\ell \setminus B^{\ell+1}) \times \mathcal{C}(\Psi_0, \mathfrak{F}_0)) \\ &= \beta_1(B_{\pi_2}^\ell \setminus B_{\pi_2}^{\ell+1}) \end{aligned}$$

for each $\ell = 1, \dots, n$, and with the convention that $B^{n+1} = B_{\pi_2}^{n+1} = \emptyset$. Now, notice that $\mu_2^m(A_2^m) = \pi_2(B_{\pi_2}^m \setminus B_{\pi_2}^{m+1} \mid B_{\pi_2}^m) = 1$ for each $m = 1, \dots, n$, and $\mu_2^m(A_2^\ell) = \pi_2(B_{\pi_2}^\ell \setminus B_{\pi_2}^{\ell+1} \mid B_{\pi_2}^m) = 0$ for each $\ell \neq m$, implying that $\tilde{\mu}_2 = (\mu_2^1, \dots, \mu_2^n) \in \mathcal{L}(\Theta_1)$. Furthermore, showing that $\tilde{\mu}_2$ is the unique LPS in Θ_1 that is

second order epistemically equivalent to π_2 is straightforward by replicating the corresponding part of the proof of Proposition 1. Thirdly, we show that $\text{marg}_\Theta \tilde{\mu}_2 = \tilde{\mu}_1$, which would then imply that $(\theta, \tilde{\mu}_1, \tilde{\mu}_2) \in T_{\Theta_1}^L$. Consider an arbitrary $E \subseteq \Theta$, and observe that

$$\begin{aligned} \text{marg}_\Theta \mu_2^m(E) &= \text{marg}_\Theta \pi_2(E|B_{\pi_2}^m) \\ &= \pi_1(E|B_{\pi_1}^m) \\ &= \mu_1^m(E), \end{aligned}$$

which proves our claim. Hence, there exists a function $\beta_2 : T_{\Psi_2}^C \rightarrow T_{\Theta_2}^L$ such that every $(\theta, \pi_1, \pi_2) \in T_{\Psi_2}^C$ is up to second order epistemically equivalent to $\beta_2(\theta, \pi_1, \pi_2) \in T_{\Theta_2}^L$. Now, we show that β_2 is surjective. In order to do that, we first need to show that there exists a Borel section of β_1 , i.e., a Borel function $\delta_1 : T_{\Theta_1}^L \rightarrow T_{\Psi_1}^C$ such that the composition $\beta_1 \circ \delta_1$ is the identity function. Fix an arbitrary $(\theta, \tilde{\mu}_1) \in T_{\Theta_1}^L$ and observe that since β_1 is surjective, there exists at least one $(\theta, \pi_1) \in T_{\Psi_1}^C$ such that $\beta_1(\theta, \pi_1) = (\theta, \tilde{\mu}_1)$. Moreover, if there exists another $(\theta, \pi_1') \in T_{\Psi_1}^C$ with $\beta_1(\theta, \pi_1') = (\theta, \tilde{\mu}_1)$, then π_1 and π_1' have a different collection of events (by Lemma 1). Hence, it follows from \mathfrak{F} being countable that there are at most countably many elements of $T_{\Psi_1}^C$ that are equivalent to $(\theta, \tilde{\mu}_1)$. Therefore, we can select one of them, thus constructing a section δ_1 . Proving that δ_1 is a Borel function follows similar steps as those in the proof of Proposition 2: Take an arbitrary Borel event $E_1 \subseteq T_{\Psi_1}^C$, and observe that

$$\begin{aligned} \delta_1^{-1}(E_1) &= \bigcup_{\mathcal{B} \in \mathfrak{F}} \delta_1^{-1}(E_1 \cap (\Theta \times \Delta^{\mathcal{B}}(\Psi_0))) \\ &= \bigcup_{\mathcal{B} \in \mathfrak{F}} \{ (\theta, \tilde{\mu}_1) \in T_{\Theta_1}^L : \delta_1(\theta, \tilde{\mu}_1) \in E_1 \cap (\Theta \times \Delta^{\mathcal{B}}(\Psi_0)) \}. \end{aligned}$$

Now, observe that $E_1 \cap (\Theta \times \Delta^{\mathcal{B}}(\Psi_0)) \subseteq \Theta \times [\Delta(\Psi_0)]^{\mathcal{B}}$ is generated by Borel rectangles of the form

$$E_0 \times \underbrace{\{p \in \Delta(\Psi_0) : p(A_1) \geq \alpha_1\} \times \cdots \times \{p \in \Delta(\Psi_0) : p(A_N) \geq \alpha_N\}}_{\text{corresponding to the events in } \mathcal{B}}$$

where $A_m \subseteq \Psi_0$ is Borel and $\alpha_m \in [0, 1]$. Therefore, the event $\{ (\theta, \tilde{\mu}_1) \in T_{\Theta_1}^L : \delta_1(\theta, \tilde{\mu}_1) \in E_1 \cap (\Theta \times \Delta^{\mathcal{B}}(\Psi_0)) \}$ is generated by events of the form

$$\{ (\theta, \tilde{\mu}_1) \in T_{\Theta_1}^L : \delta_1(\theta, \tilde{\mu}_1) \in E_0 \times \{p \in \Delta(\Psi_0) : p(A_1) \geq \alpha_1\} \times \cdots \times \{p \in \Delta(\Psi_0) : p(A_N) \geq \alpha_N\} \}$$

which can be rewritten as the union over all $\Delta_\pi^{\mathcal{B}}(\Psi_0) \subseteq \Delta^{\mathcal{B}}(\Psi_0)$ of the events

$$\{ (\theta, \tilde{\mu}_1) \in T_{\Theta_1}^L : \delta_1(\theta, \tilde{\mu}_1) \in E_0 \times (\Delta_\pi^{\mathcal{B}}(\Psi_0) \cap (\{p \in \Delta(\Psi_0) : p(A_1) \geq \alpha_1\} \times \cdots \times \{p \in \Delta(\Psi_0) : p(A_N) \geq \alpha_N\})) \}.$$

The previous event contains only $(\theta, \tilde{\mu}_1) \in T_{\Theta_1}^L$ that are mapped via δ_1 to $(\theta, \pi_1) \in T_{\Psi_1}^C$ such that $\pi_1 \in \Delta_\pi^{\mathcal{B}}(\Psi_0)$, implying that they are determined by the conditional probabilities given the π -relevant events. Therefore, the previous event can be rewritten as

$$\{ (\theta, \tilde{\mu}_1) \in T_{\Theta_1}^L : \delta_1(\theta, \tilde{\mu}_1) \in E_0 \times \underbrace{(\Delta_\pi^{\mathcal{B}}(\Psi_0) \cap (\{p \in \Delta(\Psi_0) : p(A_1) \geq \alpha_1\} \times \cdots \times \{p \in \Delta(\Psi_0) : p(A_n) \geq \alpha_n\}))}_{\text{corresponding to the events in } \mathcal{B}_\pi} \},$$

which in turn is equal to

$$T_{\Theta_1}^L \cap (E_0 \times \underbrace{\{p \in \Delta(\Theta_0) : p(A_1) \geq \alpha_1\} \times \cdots \times \{p \in \Delta(\Theta_0) : p(A_1) \geq \alpha_1\}}_{\text{subset of } \mathcal{N}_n(\Theta_0)}).$$

Observe that the latter is Borel, and therefore so is the union of these events over all $\Delta_{\pi}^{\mathcal{B}}(\Psi_0) \subseteq \Delta^{\mathcal{B}}(\Psi_0)$. Hence, $\{(\theta, \tilde{\mu}_1) \in T_{\Theta_1}^L : \delta_1(\theta, \tilde{\mu}_1) \in E_1 \cap (\Theta \times \Delta^{\mathcal{B}}(\Psi_0))\}$ is also Borel, since it is generated by countably many Borel events. Finally, take the union over $\mathcal{B} \in \mathfrak{F}$, thus obtaining the Borel event $\delta_1^{-1}(E_1)$, which proves that δ_1 is a Borel section. Then, it follows that every Borel measure on $T_{\Theta_1}^L$ can be pushed forward via δ_1^{-1} to $T_{\Psi_1}^C$, i.e., for every Borel probability measure $q \in \Delta(T_{\Theta_1}^L)$ take $q \circ \delta_1^{-1} \in \Delta(T_{\Psi_1}^C)$ (Aldaz and Render, 2000). Therefore, for every $(\theta, \tilde{\mu}_1, \tilde{\mu}_2) \in T_{\Theta_2}^L$ there exists an up to second order epistemically equivalent $(\theta, \pi_1, \pi_2) \in T_{\Psi_2}^C$, which proves that β_2 is surjective. Finally, the proof that β_2 is Borel replicates the one of Proposition 2. Replicating the previous steps for each $k > 1$ proves by induction that a Borel surjective function β_k exists for every $k > 1$. \square

Proof of Theorem 2. The proof proceeds by induction. Consider an arbitrary $t_i \in T_i$ such that $\Lambda(\pi_k(t_i)) = \Lambda(\tilde{\mu}_k(t_i)) = n$ for each $k > 0$, and observe that by assumption $g_i(t_i)(E | B_{g_i(t_i)}^m) = \lambda_i^m(t_i)(E)$ for every Borel $E \subseteq \Theta \times T_j$ and all $m = 1, \dots, n$. Then, for every Borel $E_0 \subseteq \Theta$ and all $m = 1, \dots, n$ we obtain

$$\begin{aligned} \mu_1^m(t_i)(E_0) &= \int_{\{(\theta, t_j) : \theta \in E_0\}} d\lambda_i^m(t_i) \\ &= \int_{\{(\theta, t_j) : \theta \in E_0\}} dg_i^{B_{g_i(t_i)}^m}(t_i) \\ &= \pi_1(t_i)(E_0 | B_{\pi_1(t_i)}^m) \end{aligned}$$

which proves that $\pi_1(t_i)$ is first order epistemically equivalent to $\tilde{\mu}_1(t_i)$. Then, it follows that for every $E_1 \subseteq T_{\Theta_1}^L$ it is the case that $\{(\theta, t_j) \in \Theta \times T_j : (\theta, \tilde{\mu}_1(t_j)) \in E_1\} = \{(\theta, t_j) \in \Theta \times T_j : (\theta, \pi_1(t_j)) \in \beta_1^{-1}(E_1)\}$. Now suppose that for an arbitrary $k > 1$, every $(\theta, \pi_1(t_i), \dots, \pi_{k-1}(t_i)) \in T_{\Psi_{k-1}}^C$ is up to $(k-1)$ -th order epistemically equivalent to $(\theta, \tilde{\mu}_1(t_i), \dots, \tilde{\mu}_{k-1}(t_i)) \in T_{\Theta_{k-1}}^L$, and also for every $E_{k-1} \subseteq T_{\Theta_{k-1}}^L$ it is the case that $\{(\theta, t_j) : (\theta, \tilde{\mu}_1(t_j), \dots, \tilde{\mu}_{k-1}(t_j)) \in E_{k-1}\} = \{(\theta, t_j) : (\theta, \pi_1(t_j), \dots, \pi_{k-1}(t_j)) \in \beta_{k-1}^{-1}(E_{k-1})\}$. Then, observe that for every $m = 1, \dots, n$

$$\begin{aligned} \mu_k^m(t_i)(E_{k-1}) &= \int_{\{(\theta, \tilde{\mu}_1(t_j), \dots, \tilde{\mu}_{k-1}(t_j)) \in E_{k-1}\}} d\lambda_i^m(t_i) \\ &= \int_{\{(\theta, t_j) : (\theta, \tilde{\mu}_1(t_j), \dots, \tilde{\mu}_{k-1}(t_j)) \in E_{k-1}\}} dg_i^{B_{g_i(t_i)}^m}(t_i) \\ &= \int_{\{(\theta, t_j) : (\theta, \pi_1(t_j), \dots, \pi_{k-1}(t_j)) \in \beta_{k-1}^{-1}(E_{k-1})\}} dg_i^{B_{g_i(t_i)}^m}(t_i) \\ &= \pi_k(t_i)(\beta_{k-1}^{-1}(E_{k-1}) | B_{\pi_k(t_i)}^m), \end{aligned}$$

implying that $(\theta, \pi_1(t_i), \dots, \pi_k(t_i))$ is up to k -th order epistemically equivalent to $(\theta, \tilde{\mu}_1(t_i), \dots, \tilde{\mu}_k(t_i))$. Finally notice that $\{(\theta, t_j) : (\theta, \tilde{\mu}_1(t_j), \dots, \tilde{\mu}_k(t_j)) \in E_k\} = \{(\theta, t_j) : (\theta, \pi_1(t_j), \dots, \pi_k(t_j)) \in \beta_k^{-1}(E_k)\}$ for all $E_k \subseteq T_{\Theta_k}^L$ which completes the proof inductively. \square

Proof of Theorem 3. Recall from Proposition 1 that there exists a surjective function $\beta : \mathcal{C}(\Theta \times T^C) \rightarrow \mathcal{L}(\Theta \times T^C)$. Firstly, we show that $\beta(\mathcal{C}(\Theta \times T^C, \mathfrak{F}_{T^C})) \subseteq \mathcal{L}_\Theta(\Theta \times T^C)$. Consider some $\pi \in \mathcal{C}(\Theta \times T^C, \mathfrak{F}_{T^C})$, implying that there exists some $\mathcal{B} \in \mathfrak{F}$ such that $\pi \in \Delta^{\mathcal{B} \times T^C}(\Theta \times T^C)$. Then, by construction, the collection of π -relevant events would be of the form $\{B^1 \times T^C, \dots, B^n \times T^C\}$, where $B^1, \dots, B^n \in \mathcal{B}$. Then, consider the collection $\{(B^1 \setminus B^2) \times T^C, \dots, (B^n \setminus B^{n+1}) \times T^C\}$ of Borel subsets of $\Theta \times T^C$, where $B^{n+1} := \emptyset$, and observe that $\pi((B^m \setminus B^{m+1}) \times T^C \mid B^m \times T^C) = 1$ for each $m = 1, \dots, n$, while at the same time $\pi((B^\ell \setminus B^{\ell+1}) \times T^C \mid B^m \times T^C) = 0$ for each $\ell \neq m$. Now, let $\tilde{\lambda} = \beta(\pi)$, and notice that

$$\begin{aligned} \text{marg}_\Theta \lambda^m(B^m \setminus B^{m+1}) &= \text{marg}_\Theta \pi((B^m \setminus B^{m+1}) \mid B^m) \\ &= \pi((B^m \setminus B^{m+1}) \times T^C \mid B^m \times T^C) \\ &= 1 \end{aligned}$$

while at the same time

$$\begin{aligned} \text{marg}_\Theta \lambda^m(B^\ell \setminus B^{\ell+1}) &= \text{marg}_\Theta \pi((B^\ell \setminus B^{\ell+1}) \mid B^m) \\ &= \pi((B^\ell \setminus B^{\ell+1}) \times T^C \mid B^m \times T^C) \\ &= 0 \end{aligned}$$

which proves that $\beta(\pi) \in \mathcal{L}_\Theta(\Theta \times T^C)$. Secondly, we show that for each $\tilde{\lambda} \in \mathcal{L}_\Theta(\Theta \times T^C)$ there is some $\pi \in \mathcal{C}(\Theta \times T^C, \mathfrak{F}_{T^C})$ such that $\beta(\pi) = \tilde{\lambda}$. Since $\tilde{\lambda} \in \mathcal{L}_\Theta(\Theta \times T^C)$, there exist Borel events $A_1, \dots, A_n \subseteq \Theta$ such that $\text{marg}_\Theta \lambda^m(A_m) = 1$ for all $m = 1, \dots, n$, and $\text{marg}_\Theta \lambda^m(A_\ell) = 1$ for all $\ell \neq m$. Now, similarly to the proof of Proposition 1, define the partition $\mathcal{P} := \{B_1, \dots, B_n\}$ by $B_m := A_m \setminus (\bigcup_{\ell \neq m} A_\ell)$ for each $m = 1, \dots, n-1$ and $B_n := A_n \cup (\Theta \setminus (B_1 \cup \dots \cup B_{n-1}))$. Let the finitely generated collection \mathcal{B} be the closure of \mathcal{P} with respect to union, and define $\mathcal{B}_\pi = \{B_\pi^1, \dots, B_\pi^n\}$ by $B_\pi^m := \Theta \setminus (\bigcup_{\ell < m} B_\ell)$. Now, define $\pi : \mathcal{F} \times (\mathcal{B} \times T^C) \rightarrow [0, 1]$ such that $\pi(E \mid B_\pi^m \times T^C) := \lambda^m(E)$ for all Borel $E \subseteq \Theta \times T^C$ and all $m = 1, \dots, n$. For every $B \in \mathcal{B}$, consider the smallest $m \in \{1, \dots, n\}$ such that $\pi(B \times T^C \mid B_\pi^m \times T^C) > 0$, which by construction always exists. Then, for every Borel $E \subseteq \Theta \times T^C$, let $\pi(E \mid B \times T^C) := \pi(E \cap (B \times T^C) \mid B_\pi^m \times T^C) / \pi(B \times T^C \mid B_\pi^m \times T^C)$. Finally, verify that $\pi \in \Delta^{\mathcal{B} \times T^C}(\Theta \times T^C)$, with $\mathcal{B}_\pi \times T^C$ being the collection of π -relevant events. Therefore, $\beta : \mathcal{C}(\Theta \times T^C, \mathfrak{F}_{T^C}) \rightarrow \mathcal{L}_\Theta(\Theta \times T^C)$ is surjective. Moreover, since $\Delta^{\mathcal{B} \times T^C}(\Theta \times T^C)$ is Borel for each $\mathcal{B} \in \mathfrak{F}$, and also \mathfrak{F}_{T^C} is countable, it follows that $\mathcal{C}(\Theta \times T^C, \mathfrak{F}_{T^C})$ is Borel in $\mathcal{C}(\Theta \times T^C)$, implying that $\beta : \mathcal{C}(\Theta \times T^C, \mathfrak{F}_{T^C}) \rightarrow \mathcal{L}_\Theta(\Theta \times T^C)$ is a Borel function. Recall from Proposition 3 that there exists a homeomorphism $g : T^C \rightarrow \mathcal{C}(\Theta \times T^C, \mathfrak{F}_{T^C})$. Therefore, $\tilde{\lambda} := \beta \circ g$ is a Borel surjective function from T^C onto $\mathcal{L}_\Theta(\Theta \times T^C)$. In addition, every $t \in T^C$ is by construction such that $g(t) \in \mathcal{C}(\Theta \times T^C, \mathfrak{F}_{T^C})$ is epistemically equivalent with $\tilde{\lambda}(t) \in \mathcal{L}_\Theta(\Theta \times T^C)$, implying by Theorem 2 that the \mathcal{C} -hierarchy associated with $g(t)$ is epistemically equivalent with the \mathcal{L} -hierarchy associated with $\tilde{\lambda}(t)$. Furthermore, it follows from Theorem 1 that every \mathcal{L} -hierarchy is the image of some $t \in T^C$, implying that for each $(\tilde{\mu}_1, \tilde{\mu}_2, \dots) \in T^L$ there exists some $t \in T^C$ such that $(\tilde{\mu}_1(t), \tilde{\mu}_2(t), \dots) = (\tilde{\mu}_1, \tilde{\mu}_2, \dots)$. Finally, let T be a renaming of T^C , which completes the proof. \square

Proof of Proposition 4. For an arbitrary Borel $F \subseteq \Theta \times T_a \times T_b$, let

$$\begin{aligned} A_i(F) &:= \{ t_i \in T_i : \tilde{\lambda}_i(t_i) \in \mathcal{A}(\text{Proj}_{\Theta \times T_j} F) \}, \\ SB_i(F) &:= \{ t_i \in T_i : g_i(t_i) \in \mathcal{SB}(\text{Proj}_{\Theta \times T_j} F) \}. \end{aligned}$$

Now, consider some $t_i \in SB_i(F)$, and observe that, since $g_i(t_i)$ is full support, it follows from Brandenburger et al. (2007, Prop. 3) that $\beta(g_i(t_i)) \in \mathcal{A}^+(\text{Proj}_{\Theta \times T_j} F)$, and therefore $\tilde{\lambda}_i(t_i) \in A_i(F)$. Moreover, since Θ is countable, it follows that every collection of conditioning events in \mathfrak{F}_{T_j} contains only clopen events. Therefore, it follows from Brandenburger et al. (2007, Cor. 3) that if $\tilde{\lambda}(t_i) \in \mathcal{A}^+(\text{Proj}_{\Theta \times T_j} F)$ then $\beta^{-1}(\tilde{\lambda}(t_i)) \in \mathcal{SB}^+(\text{Proj}_{\Theta \times T_j} F)$. Hence, $t_i \in A_i(F)$ implies $t_i \in SB_i(F)$. Therefore, since it is the case that $A_i(F) = SB_i(F)$ for all Borel events $F \subseteq \Theta \times T_a \times T_b$, it follows directly that $E \subseteq \Theta \times T_a \times T_b$ is commonly strongly believed at (t_a, t_b) if and only if it is commonly assumed. \square

References

- ALDAZ, J.M. & RENDER, H. (2000). Borel measure extensions of measures defined on sub- σ -algebras. *Advances in Mathematics* 150, 233-263.
- ALIPRANTIS, C. & BORDER, K. (1994). *Infinite dimensional analysis*. Springer Verlag, Berlin.
- ASHEIM, G.B. (2001). Proper rationalizability in lexicographic beliefs. *International Journal of Game Theory* 30, 453-478.
- ASHEIM, G.B. & DUFWENBERG, M. (2003). Admissibility and common belief. *Games and Economic Behavior* 42, 208-234.
- AUMANN, R.J. (1987). Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55, 1-18.
- AUMANN, R.J. & BRANDENBURGER, A. (1995). Epistemic conditions for Nash equilibrium. *Econometrica* 63, 1161-1180.
- BATTIGALLI, P. (1996). Strategic rationality orderings and the best rationalization principle. *Games and Economic Behavior* 13, 178-200.
- BATTIGALLI, P. & FRIEDENBERG, A. (2012). Forward induction reasoning revisited. *Theoretical Economics* 57-98.
- BATTIGALLI, P. & SINISCALCHI, M. (1999). Hierarchies of conditional beliefs and interactive epistemology in dynamic games. *Journal of Economic Theory* 88, 188-230.
- (2002). Strong belief and forward induction reasoning. *Journal of Economic Theory* 106, 356-391.

- BLUME, L., BRANDENBURGER, A. & DEKEL, E. (1991a). Lexicographic probabilities and choice under uncertainty. *Econometrica* 59, 61–79.
- (1991b). Lexicographic probabilities and equilibrium refinements. *Econometrica* 59, 81–98.
- BRANDENBURGER, A. (2007). The power of paradox: some recent developments in interactive epistemology. *International Journal of Game Theory* 35, 465–492.
- BRANDENBURGER, A. & DEKEL, E. (1987). Rationalizability and correlated equilibria. *Econometrica* 55, 1391–1402.
- (1993). Hierarchies of beliefs and common knowledge. *Journal of Economic Theory* 59, 189–198.
- BRANDENBURGER, A. & FRIEDENBERG, A. (2010). Self-admissible sets. *Journal of Economic Theory* 145, 785–811.
- BRANDENBURGER, A., FRIEDENBERG, A. & KEISSLER, H.J. (2007). Notes on the relationship between strong belief and assumption. *Unpublished manuscript*.
- (2008). Admissibility in games. *Econometrica* 76, 307–352.
- CATONINI, E. (2012). Common assumption of cautious rationality and iterated admissibility. *Unpublished manuscript*.
- FRIEDENBERG, A. (2010). When do type structures contain all hierarchies of beliefs? *Games and Economic Behavior* 68, 108–129.
- HALPERN, J.Y. (2010). Lexicographic probability, conditional probability, and nonstandard probability. *Games and Economic Behavior* 68, 155–179.
- HARSANYI, J. (1967-68). Games with incomplete information played by Bayesian players, I-III. *Management Science* 14, 159–182, 320–334, 486–502.
- HEIFETZ, A., MEIER, M. & SCHIPPER, B. (2010). Comprehensive rationalizability. *Unpublished manuscript*.
- KECHRIS, A. (1995). *Classical descriptive set theory*. Springer Verlag, Berlin.
- KEISLER, H.J. & LEE, B.S. (2011). Common assumption of rationality. *Unpublished manuscript*.
- MERTENS, J.F. & ZAMIR, S. (1985). Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory* 14, 1–29.
- MYERSON, R. (1986). Multistage games with communication. *Econometrica* 54, 323–358.

- PEREA, A. (2011). An algorithm for proper rationalizability. *Games and Economic Behavior* 72, 510–525.
- (2012). *Epistemic game theory: Reasoning and choice*. Cambridge University Press.
- (2013). Plausibility orderings in dynamic games. *Economics and Philosophy* (forthcoming).
- RÊNYI, A. (1955). On a new axiomatic theory of probability. *Acta Mathematica Academiae Scientiarum Hungaricae* 6, 285–335.
- TAN, T. & WERLANG, S. (1988). The Bayesian foundations of solution concepts of games. *Journal of Economic Theory* 45, 370–391.
- ZUAZO-GARIN, P. (2013). Incomplete imperfect information and backward induction. *Unpublished manuscript*.