



CMS-EMS
Center for Mathematical Studies in Economics
And Management Science

Discussion Paper #1547

Bounded Reasoning and
Higher-Order Uncertainty

Willemien Kets
Northwestern University
March 19, 2012

JEL Classification: C700, C720, D800, D830

Keywords: Bounded rationality, higher-order beliefs, finite depth of reasoning.



NORTHWESTERN
UNIVERSITY



Bounded Reasoning and Higher-Order Uncertainty

Willemien Kets*

This version: March 19, 2012

First version: November 2009

Abstract

The standard framework for analyzing games with incomplete information models players as if they form beliefs about their opponents' beliefs about their opponents' beliefs and so on, that is, as if players have an infinite depth of reasoning. This strong assumption has nontrivial implications, as is well-known. This paper therefore generalizes the type spaces of Harsanyi (1967–1968) to model that players can have a finite depth of reasoning. The innovation is that players can have a coarse perception of the higher-order beliefs of other players, thus formalizing the small-world idea of Savage (1954) in a type-space context. Unlike the case in other models of finite-order reasoning, players with a finite depth of reasoning can have nontrivial higher-order beliefs about certain events. Intuitively, some higher-order events are generated by events of lower orders, making it possible for players to reason about them, even if they have a finite depth of reasoning.

JEL classification: C700, C720, D800, D830

Keywords: Bounded rationality, higher-order beliefs, finite depth of reasoning.

*Kellogg School of Management, Northwestern University. E-mail: w-kets@kellogg.northwestern.edu. Phone: +1-505-204 8012. I am grateful to Adam Brandenburger, Yossi Feinberg, and Matthew Jackson for their guidance and support, and to Adam Brandenburger, Alfredo Di Tillio, Amanda Friedenberg, and Aviad Heifetz for important input. I also benefited from discussions with Songzi Du, Matt Elliott, Jeff Ely, Yossi Feinberg, Ben Golub, Joe Halpern, Sander Heinsalu, Martin Meier, Rosemarie Nagel, Andrés Perea, Tomasz Sadzik, Dov Samet, Burkhard Schipper, Elias Tsakas, and Jonathan Weinstein, as well as from comments from numerous seminar audiences. Much of this research was carried out during visits to Stanford University and the NYU Stern School of Business, and I thank these institutions for their hospitality. Financial support from the Air Force Office for Scientific Research under Grant FA9550-08-1-0389 is gratefully acknowledged.

1. Introduction

Analyzing games of incomplete information requires taking into account not just the beliefs of players, but also their *higher-order beliefs*. If a manager of a firm does not observe its competitors' costs, for example, he needs to form beliefs about the cost structure of all the firms in the industry—the state of nature—to predict his competitors' pricing decisions, so as to optimally set his own price. But the pricing decisions of the firm's competitors in turn depend on *their* beliefs about the state of nature. To decide on his optimal action, the manager therefore needs to form a belief not only about nature (a so-called first-order belief), but also about his competitors' beliefs about nature (a second-order belief, i.e., a belief about a first-order belief). And because his competitors likewise consider his beliefs about nature, the manager has to form a belief about their beliefs about his beliefs about nature (a third-order belief), and so on, ad infinitum (cf. Harsanyi, 1967–1968).

Are “real” players capable of such higher-order reasoning? The answer to this question is not so clear-cut as it may seem. A statement such as “John Dean did not know that Nixon knew that Dean knew that Nixon knew that McCord had burgled O'Brien's office in the Watergate Apartments” is inherently difficult to reason about (Clark and Marshall, 1981). At the same time, other types of higher-order reasoning seem unproblematic. If two players, say, Ann and Bob, sit across the table from each other and have eye contact with each other, then clearly each of them believes that they have eye contact, believes that the other believes that, believes that the other believes that they believe that, and so on. That is, it is common belief between Ann and Bob that they have eye contact.¹ Similarly, a public announcement immediately becomes common belief (Friedell, 1969; Lewis, 1969; Chwe, 2001). These examples suggest that some higher-order events are easier to reason about than others.

Existing models do not take this into account. On the one hand, standard game-theoretic models model players as if they have higher-order beliefs about every possible event, at all orders, i.e., as if players have an infinite depth of reasoning. On the other hand, in models developed in the experimental literature, such as cognitive-hierarchy models or models of level- k reasoning,² players can have a finite depth of reasoning. However, these models assume that a player with a finite depth of reasoning cannot reason about *any* event at sufficiently high order, ruling out, for example, that it can be common belief between Ann and Bob that they have eye contact.

¹We follow the recent literature in game theory in using the terms “belief” and “common belief” rather than “knowledge” and “common knowledge.” The formal distinction is between “knowledge,” which is considered to be always true, and “beliefs,” which may be true or false.

²See, e.g., Nagel (1995), Stahl and Wilson (1995) Ho et al. (1998), Costa-Gomes et al. (2001), Strzalecki (2009), and Heifetz and Kets (2011).

Because beliefs at arbitrarily high order can have a significant impact on economic outcomes,³ it is important to carefully model what higher-order events players with bounded reasoning abilities can hold beliefs about. This paper provides a framework that does just that.

We propose a class of type spaces, called *extended type spaces*, starting from the idea that a player can have a coarse perception of the state of the world, where the state of the world specifies not only the physical reality—the state of nature—, but also players’ beliefs and higher-order beliefs. A player who has a coarse perception does not distinguish states of the world that differ only in certain details, such as the beliefs of other players at very high orders. We show that a player with a coarse perception of the higher-order beliefs of other players has a finite depth of reasoning, in the sense that there are events beyond a certain finite order that a player with that perception cannot form beliefs about. However, a player with a finite depth of reasoning is able to reason about events at higher orders if these events can be reduced to a lower-order event, as in the case where Ann and Bob have eye contact or where there is a public announcement.

As in the type spaces introduced by Harsanyi (1967–1968), each type in an extended type space is associated with a belief (probability measure) over the states of nature and the types of other players. Unlike in a Harsanyi type space, the beliefs of different types of a player can be defined on different σ -algebras.⁴ As a type’s belief assigns a probability only to those subsets of her opponents’ types that are in the type’s σ -algebra, a type with a coarse σ -algebra has a coarse perception of the other players’ types. And because types generate higher-order beliefs, the coarseness of a type’s σ -algebra thus determines what features of the other players’ higher-order beliefs the type can reason about. In particular, if the type cannot form a belief about others’ beliefs at all orders, it has a finite depth of reasoning.

Coarse perceptions thus model small worlds, introduced by Savage (1954) in the context of one-person decision situations. A state in a small world describes the possible uncertainties a decision maker faces in less detail than a state in a large world, by neglecting certain distinctions between states. This means that “a state of the smaller world corresponds not to one state of the larger, but to a *set* of states” (Savage, 1954, p. 9, emphasis added). In the present framework, a player may neglect differences between types for the other player that

³ An action that is optimal for a player given her k th-order belief, for example, may no longer be optimal given her $(k + 1)$ th-order belief, for any finite k (Rubinstein, 1989; Carlsson and van Damme, 1993). Also, beliefs at arbitrarily high order may also determine whether players with a common prior can have different posteriors (Aumann, 1976; Geanakoplos and Polemarchakis, 1982).

⁴A σ -algebra \mathcal{F} on a set X is a collection of subsets of X that contains X and is closed under complements and countable unions. Importantly, if a probability measure on X is defined on a σ -algebra \mathcal{F} , it can assign a probability only to those subsets of X that belong to \mathcal{F} .

differ only in the beliefs they generate at high order, by lumping these types together in his σ -algebra.

Extended type spaces thus generalize the Harsanyi framework: A Harsanyi type space is simply an extended type space in which each type has a σ -algebra that is fine enough for the type to have an infinite depth of reasoning. Extended type spaces can also be seen as a generalization of cognitive-hierarchy and level- k models: Kets (2012) constructs an extended type space such that there is no higher-order event for which players' belief about that event are completely determined by their beliefs about some lower-order event. That means that if a player has a finite depth, then there are no higher-order events she can reason about, as in the cognitive-hierarchy and level- k models.

A natural question is whether types induce well-defined belief hierarchies, as in the Harsanyi case. This is indeed the case, as we demonstrate in Theorem 4.1. We go on to characterize the depth of reasoning of types. A k th-order event is defined to be an event involving the state of nature and the $(k - 1)$ th-order beliefs for the other player; a type is said to have depth (of reasoning) $k < \infty$ if the type induces a belief about every k th-order event, but does not have beliefs about at least some events at higher orders; a type has an infinite depth if it induces a belief about all k th-order events for every k .

If we do not put any restrictions on the σ -algebras that a type can have, then some types may not have a well-defined depth of reasoning, as we demonstrate in Appendix C. However, we impose conditions on the σ -algebras that types can have that ensure that the coarseness of a type's σ -algebra reflects precisely the extent to which a type can reason about the higher-order beliefs of his opponents, and not some other form of coarseness of perception. It then follows immediately that each type induces a belief hierarchy of a well-defined depth (Theorem 4.2).⁵

In principle, the depth of reasoning of a type can be determined by writing out the belief hierarchy it induces, and checking whether there is some finite k such that the belief hierarchy induced by the type does not specify k th-order beliefs about certain events, but this can of course be tedious. Theorem 4.7 therefore characterizes the depth of reasoning of a type in terms of the properties of the type space alone, without making reference to belief hierarchies, under a mild condition on the type space.

Together, these results demonstrate that extended type spaces provide an implicit description of players' finite and infinite hierarchies of beliefs, including higher-order uncertainty

⁵Appendix C shows that the conditions we impose are tight in the sense that if one of the conditions is not satisfied, then there exists an extended type space in which some types do not have a well-defined depth of reasoning, or where types with different σ -algebras have the same depth of reasoning. The latter is akin to endowing different types (for a given player) in a Harsanyi type space with different σ -algebras (while ensuring that all belief maps are measurable). See Appendix C for details.

about others' depth of reasoning, by specifying types, beliefs about types, and a collection of σ -algebras on each type set, just like the Harsanyi type spaces model players' infinite belief hierarchies implicitly, by specifying types and beliefs about types.

Having characterized the bounds on reasoning for types with a coarse perception, we turn to the question what higher-order events a type with a finite depth can reason about. We show that a type with a finite depth d can form beliefs about a k th-order event for $k > d$ if and only if the event is equivalent to an event of sufficiently low order (Theorem 5.1).

The latter result sheds light on the question why the statement about Dean's and Nixon's higher-order beliefs is so hard to reason about, while it can be common belief among players that they have eye contact, even if they have a finite depth of reasoning. If Ann and Bob have eye contact, then their higher-order beliefs are completely determined by the low-order event that they have eye contact: If Ann and Bob have eye contact, then in each state of the world they consider possible, they have eye contact, each of them believes that they have eye contact, and believes that the other believes that, and so on. The low-order event that Ann and Bob have eye contact is thus equivalent to the high-order event that there is common belief that they have eye contact. This means that even if they only distinguish states of the world in which they have eye contact from states in which they do not, the event that they have eye contact can nevertheless be common belief.

The situation is different when it comes to Dean's beliefs about Nixon's beliefs. In that case, there are many states of the world, corresponding to different higher-order beliefs for Dean and Nixon. Typically, these higher-order beliefs will not be fixed by some lower-order event. For example, it is not the case that Dean knows that Nixon knew that Dean knew that Nixon knew of the burglary whenever Nixon knows of the break-in, and does not know it otherwise. That means that a player has to keep track of many different states of the world in order to be able to reason about Dean's higher-order beliefs, and this may be too taxing. Section 2.4 presents an example where nontrivial common belief is possible, yet players cannot reason about certain higher-order events.

The idea that "simple" events can induce (almost) common knowledge is not new; it is central to the conceptualization of common knowledge by the philosopher David Lewis (1969) and it underlies the formalization of common knowledge and approximate common belief in Aumann (1976) and Monderer and Samet (1989), respectively. Indeed, speaking of a belief hierarchy such as the one described above, where Ann believes she and Bob have eye contact, believes that Bob believes that, and so on, Lewis writes: "this is a chain of implications, [it does not represent] steps in anyone's actual reasoning. Therefore, there is nothing improper about its infinite length" (p. 53). Our contribution here is to point out that this idea applies beyond the context of common knowledge, and, more fundamentally, to formalize it in the

context of players with bounded reasoning abilities, building on Savage’s small world idea, and to use this insight to characterize the higher-order events that bounded reasoners can have beliefs about.

The remainder of this paper is organized as follows. The next section illustrates our main results using simple examples. Section 3 formally introduces the notion of an extended type space, and Section 4 characterizes the depth of reasoning of types. Section 5 investigates the higher-order events players with a finite depth can reason about. Section 6 discusses the related literature. All proofs are relegated to the appendices.

2. Examples

2.1. Extended type spaces

We present some examples to introduce our framework, and to illustrate the main results. Throughout this section, we consider a setting in which two players, Ann (a) and Bob (b), are uncertain about the state of nature θ , which can be either high (H) or low (L).

We represent the uncertainty faced by the players, including their uncertainty about the beliefs of the other player, by an *extended type space*. As in the type spaces of Harsanyi (1967–1968), each player $i = a, b$ is endowed with a type space T_i , and each type $t_i \in T_i$ is associated with a belief (probability measure) $\beta_i(t_i)$ about the state of nature and the other player’s type. Unlike in a Harsanyi type space, the beliefs of types in an extended type space can be defined on different σ -algebras. That is, Ann’s type set T_a is endowed with a collection \mathcal{S}_a of σ -algebras, and a type $t_b \in T_b$ for Bob over Ann’s type is defined on some σ -algebra $\Sigma_b(t_b)$ in \mathcal{S}_a ; likewise for Bob’s type set and the beliefs of Ann’s types. The idea is that the σ -algebra on which a type’s belief is defined reflects the extent to which the type “thinks through” the beliefs of the other player. An extended type space is thus a tuple $(T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i=a,b}$. (We will also require that extended type spaces satisfy some additional conditions, but will ignore that in this informal treatment.)

2.1.1. Infinite depth

We illustrate how types generate higher-order beliefs using the type space in Figure 1. The collection \mathcal{S}_a of σ -algebras on Ann’s type set simply consists of the σ -algebra that contains the singletons; likewise for \mathcal{S}_b . Since type sets are finite, it suffices to specify the belief $\beta_a(t_a)$ for a type t_a for Ann on the partition of Bob’s type set that its σ -algebra $\Sigma_a(t_a)$ induces, and similarly for the types for Bob. For example, specifying the belief for type t_a^1 on the pairs (θ, t_b) for every state of nature θ and type t_b for Bob specifies its belief on the full σ -algebra.

$\beta_a(t_a^1)$	H	L	$\beta_a(t_a^2)$	H	L	$\beta_b(t_b^1)$	H	L	$\beta_b(t_b^2)$	H	L
t_b^1	$\frac{1}{2}$	0	t_b^1	0	0	t_a^1	$\frac{1}{2}$	0	t_a^1	0	0
t_b^2	$\frac{1}{2}$	0	t_b^2	0	0	t_a^2	$\frac{1}{2}$	0	t_b^2	0	0
t_b^3	0	0	t_b^3	$\frac{1}{2}$	0	t_a^3	0	0	t_a^3	$\frac{1}{2}$	0
t_b^4	0	0	t_b^4	$\frac{1}{2}$	0	t_a^4	0	0	t_a^4	$\frac{1}{2}$	0

$\beta_a(t_a^3)$	H	L	$\beta_a(t_a^4)$	H	L	$\beta_b(t_b^3)$	H	L	$\beta_b(t_b^4)$	H	L
t_b^1	0	0	t_b^1	0	$\frac{1}{2}$	t_a^1	0	0	t_a^1	0	$\frac{1}{2}$
t_b^2	0	0	t_b^2	0	$\frac{1}{2}$	t_a^2	0	0	t_a^2	0	$\frac{1}{2}$
t_b^3	0	$\frac{1}{2}$	t_b^3	0	0	t_a^3	0	$\frac{1}{2}$	t_a^3	0	0
t_b^4	0	$\frac{1}{2}$	t_b^4	0	0	t_a^4	0	$\frac{1}{2}$	t_a^4	0	0

Figure 1: An extended type space where every type has an infinite depth of reasoning, with the beliefs for types for Ann on the left, and those for Bob on the right; we write x for the singleton $\{x\}$.

The types and their beliefs determine players' higher-order beliefs. For example, type t_a^1 for Ann believes (with probability 1) that the state of nature is H and that Bob believes that the state of nature is H , as it assigns probability 1 to types t_b^1 and t_b^2 for Bob, which both believe that the state of nature is H . Using that Bob's types have a belief about Ann's type, we see that t_a^1 assigns probability $\frac{1}{2}$ to the event that Bob believes that Ann believes that the state is L (as t_a^1 assigns probability $\frac{1}{2}$ to type t_b^2 for Bob, which in turn assigns probability 1 to the event that Ann has type t_a^3 or t_a^4 , which both put probability 1 on $\theta = L$). Going further, we can derive Ann's beliefs about Bob's beliefs about Ann's beliefs about Bob's beliefs about θ , and so on.

Thus, each type induces an infinite belief hierarchy: a belief about nature, a belief about the other player's belief about nature, and so on; we say that types have an infinite depth (of reasoning) in this case. Indeed, it can be checked that the type space in Figure 1 is a regular Harsanyi type space. In general, a Harsanyi type space is an extended type space in which the beliefs of the types for a player are all defined on a σ -algebra that is sufficiently fine for types to induce beliefs at all orders. Extended type spaces thus generalize the Harsanyi framework.

2.2. Finite depth

What if types have a coarser perception? In the extended type space in Figure 2, the σ -algebra $\Sigma_a(t_a)$ associated with a type t_a for Ann is generated by the partition $\{\{t_b^1, t_b^2\}, \{t_b^3, t_b^4\}\}$ of Bob's type set $\{t_b^1, t_b^2, t_b^3, t_b^4\}$, and likewise for the σ -algebra associated with the types for

Bob.

$\beta_a(t_a^1)$	H	L	$\beta_a(t_a^2)$	H	L	$\beta_b(t_b^1)$	H	L	$\beta_b(t_b^2)$	H	L
$\{t_b^1, t_b^2\}$	1	0	$\{t_b^1, t_b^2\}$	0	0	$\{t_a^1, t_a^2\}$	1	0	$\{t_a^1, t_a^2\}$	0	0
$\{t_b^3, t_b^4\}$	0	0	$\{t_b^3, t_b^4\}$	1	0	$\{t_a^3, t_a^4\}$	0	0	$\{t_a^3, t_a^4\}$	1	0
$\beta_a(t_a^3)$	H	L	$\beta_a(t_a^4)$	H	L	$\beta_b(t_b^3)$	H	L	$\beta_b(t_b^4)$	H	L
$\{t_b^1, t_b^2\}$	0	0	$\{t_b^1, t_b^2\}$	0	1	$\{t_a^1, t_a^2\}$	0	0	$\{t_a^1, t_a^2\}$	0	1
$\{t_b^3, t_b^4\}$	0	1	$\{t_b^3, t_b^4\}$	0	0	$\{t_a^3, t_a^4\}$	0	1	$\{t_a^3, t_a^4\}$	0	0

Figure 2: An extended type space with coarse σ -algebras.

Each type t_a for Ann has a first-order belief, that is, it assigns a probability to each event concerning the state of nature. Type t_a^1 , for example, believes that the state of nature is H . Each type also induces a second-order belief, that is, a belief about any belief Bob may have about θ . Type t_a^1 , for example, assigns probability 1 to the event that Bob has type t_b^1 or t_b^2 (i.e., to $\{t_b^1, t_b^2\}$), and thus to the event that Bob believes that the state of nature is high (as both t_b^1 and t_b^2 assign probability 1 to H). Likewise for the other types. Every type in this type space therefore has a well-defined second-order belief, even though every type has a coarse σ -algebra, in this case, a σ -algebra that is coarser than the σ -algebra generated by the singletons.⁶

2.3. Depth of reasoning

What are the conditions under which types have well-articulated beliefs at a given order? We saw that in the type space in Figure 2, each type could reason about the beliefs of the other player about nature. If we take a closer look at the σ -algebra $\mathcal{F}_b := \Sigma_a(t_a)$ associated with a type t_a for Ann in that type space, we notice that it contains the subsets of Bob's types that hold a particular belief about θ . That is, the σ -algebra \mathcal{F}_b lumps together the types for Bob whenever they coincide in their beliefs about θ , and separates the types otherwise. For example, the σ -algebra \mathcal{F}_b separates the types for Bob that put probability 1 on H (types t_b^1 and t_b^2) from those that put probability 1 on L (types t_b^3 and t_b^4).

More formally, a type for Ann has a second-order belief if its σ -algebra on Bob's type set contains the subsets

$$\{t_b \in T_b : \text{the marginal of } \beta_b(t_b) \text{ on } \{H, L\} \text{ assigns probability at least } p \text{ to } E\}$$

⁶For the structure described here to be an extended type space, the collection \mathcal{S}_i of σ -algebras has to include the trivial σ -algebra $\{T_i, \emptyset\}$ for each player i (by Condition 3), but we can ignore that here.

for every event $E \subseteq \{H, L\}$ and $p \in [0, 1]$. In that case, a type for Ann can assign a probability to the event that Bob assigns probability at least p to any first-order event E , for every p . The σ -algebra \mathcal{F}_b clearly satisfies this condition. If a type induces a second-order belief, as the types in the type space in Figure 2 do, then we say that it has a *depth (of reasoning)* of at least 2.

What is needed, then, for a type to induce a third-order belief, so that it has a depth of reasoning of at least 3? Extending the argument above, a type for Bob induces a third-order belief if its σ -algebra separates the types for Ann that differ in their second-order belief, i.e., in their belief about Bob's belief about θ . Building on the argument above, we can easily characterize the types for Ann that have the same second-order belief: the types for Ann that share the same belief about Bob's beliefs about θ are precisely the types whose beliefs coincide on the σ -algebra \mathcal{F}_b . Because this σ -algebra lumps together exactly those types for Bob that have the same beliefs about θ , two types for Ann whose beliefs coincide on \mathcal{F}_b will have the same beliefs about Bob's beliefs about θ .

To make this more formal, say that a subset $E \subseteq \{H, L\} \times T_b$ is *expressible* in the σ -algebra \mathcal{F}_b if it belongs to the product of the usual σ -algebra on $\{H, L\}$ and \mathcal{F}_b . A type for Bob then induces a third-order belief if its σ -algebra \mathcal{F}_a on Ann's type set contains the subset

$$\{t_a \in T_a : \beta_a(t_a) \text{ assigns probability at least } p \text{ to } E\}$$

for every event E that is expressible in \mathcal{F}_b and for every $p \in [0, 1]$. We say that the σ -algebra \mathcal{F}_a *dominates* the σ -algebra \mathcal{F}_b in this case.

Thus, a type for Bob has depth at least 3 if its σ -algebra dominates a σ -algebra, viz., \mathcal{F}_b , that corresponds to a depth of reasoning of at least 2. The results in Section 4 imply that this condition is also necessary, and that it holds for any depth of reasoning: A type for Bob has depth $k < \infty$ if and only if its σ -algebra (on Ann's type set) dominates a σ -algebra that corresponds to depth $k - 1$ (on Bob's type set), but does not dominate a σ -algebra of depth k . Likewise for the types for Ann.

Going back to the extended type space in Figure 2, we see that a σ -algebra \mathcal{F}_a on Ann's type set dominates \mathcal{F}_b only if it contains the singletons. For example, if we take the event $E' := (H, \{t_b^1, t_b^2\})$ that the state of nature is H and that Bob believes that, then the only type that assigns probability 1 to E' is t_a^1 , so that a σ -algebra that dominates \mathcal{F}_b needs to contain the singleton $\{t_a^1\}$. Since there is no type for Bob whose σ -algebra contains the singletons, all types have depth 2: they induce a second-order belief, but not a third-order belief; the same, in fact, holds for Ann's types.

Extended type spaces can thus model situations in which players have a finite depth of reasoning. By endowing different types for a player with different σ -algebras, we can also

allow for uncertainty about depth: players may be uncertain about the depth of reasoning of their opponent. The type space in Example 1 in Section 4.3 provides an instance of this.

2.4. Higher-order beliefs

So far, we have focused on what players with a coarse perception *cannot* reason about. We now turn to the question what higher-order events types *can* think about, despite their finite depth. We start with an example. The set of states of nature is now $\Theta = \{eb, en, xb, xn\}$, and each player $i = a, b$ has eight types: $t_i^1, t_i^2, t_i^3, t_i^4, t_i^5, t_i^6, t_i^7, t_i^8$. The beliefs for the types for Ann are given in Figure 3; the beliefs for the types for Bob are obtained by exchanging the indices a and b .

Suppose eb means that Ann and Bob have eye contact and that a burglary has taken place, while en means that they have eye contact, but no break-in has occurred; xb stands for a situation in which there is no eye contact, but there has been a burglary, and xn is a situation without eye contact and without a burglary.

In the state in which Ann and Bob have types t_a^1 and t_b^1 , respectively, there is common belief that they have eye contact: t_a^1 believes that there is eye contact (i.e., $\theta \in \{eb, xb\}$), believes that Bob believes that, and so on, and similarly for type t_b^1 for Bob. Thus, even if players have a finite depth of reasoning, they can nevertheless hold nontrivial higher-order beliefs about certain events. The key is that any event concerning the higher-order beliefs of players about their having eye contact is equivalent to the lower-order event that they have eye contact. A public announcement can similarly induce common belief in its content.

On the other hand, players only have limited beliefs about the other's beliefs about the burglary: each believes that a burglary has taken place, believes that the other believes that, but no type can reason about the other's beliefs about the burglary at higher-order. Type t_a^1 , for example, believes that there has been a burglary and believes that Bob believes that (as both t_b^1 and t_b^2 believe that there has been a break-in), but cannot reason about what Bob believes that she believes about the break-in: she cannot distinguish types for Bob that differ in their beliefs about her beliefs about the burglary. The problem here is that events concerning other players' higher-order beliefs about the burglary cannot be expressed in terms of an event of sufficiently low order. Intuitively, even if a player believes that a burglary has taken place, and that the other believes that, the other may or may not believe that she believes that the burglary has occurred.

Thus, players with a finite depth of reasoning can have beliefs about a nontrivial higher-order event provided that it is equivalent to an event of sufficiently low order, but not otherwise. Theorem 5.1 shows that this holds generally: A player of depth $d < \infty$ can form beliefs about

an event F of order k if and only if F is equivalent to an event of order at most d .

$\beta_a(t_a^1)$	<i>eb</i>	<i>en</i>	<i>xb</i>	<i>xn</i>	$\beta_a(t_a^2)$	<i>eb</i>	<i>en</i>	<i>xb</i>	<i>xn</i>
$\{t_b^1, t_b^2\}$	1	0	0	0	$\{t_b^1, t_b^2\}$	0	0	0	0
$\{t_b^3, t_b^4\}$	0	0	0	0	$\{t_b^3, t_b^4\}$	1	0	0	0
$\{t_b^5, t_b^6\}$	0	0	0	0	$\{t_b^5, t_b^6\}$	0	0	0	0
$\{t_b^7, t_b^8\}$	0	0	0	0	$\{t_b^7, t_b^8\}$	0	0	0	0
$\beta_a(t_a^3)$	<i>eb</i>	<i>en</i>	<i>xb</i>	<i>xn</i>	$\beta_a(t_a^4)$	<i>eb</i>	<i>en</i>	<i>xb</i>	<i>xn</i>
$\{t_b^1, t_b^2\}$	0	1	0	0	$\{t_b^1, t_b^2\}$	0	0	0	0
$\{t_b^3, t_b^4\}$	0	0	0	0	$\{t_b^3, t_b^4\}$	0	1	0	0
$\{t_b^5, t_b^6\}$	0	0	0	0	$\{t_b^5, t_b^6\}$	0	0	0	0
$\{t_b^7, t_b^8\}$	0	0	0	0	$\{t_b^7, t_b^8\}$	0	0	0	0
$\beta_a(t_a^5)$	<i>eb</i>	<i>en</i>	<i>xb</i>	<i>xn</i>	$\beta_a(t_a^6)$	<i>eb</i>	<i>en</i>	<i>xb</i>	<i>xn</i>
$\{t_b^1, t_b^2\}$	0	0	0	0	$\{t_b^1, t_b^2\}$	0	0	0	0
$\{t_b^3, t_b^4\}$	0	0	0	0	$\{t_b^3, t_b^4\}$	0	0	0	0
$\{t_b^5, t_b^6\}$	0	0	1	0	$\{t_b^5, t_b^6\}$	0	0	0	0
$\{t_b^7, t_b^8\}$	0	0	0	0	$\{t_b^7, t_b^8\}$	0	0	1	0
$\beta_a(t_a^7)$	<i>eb</i>	<i>en</i>	<i>xb</i>	<i>xn</i>	$\beta_a(t_a^8)$	<i>eb</i>	<i>en</i>	<i>xb</i>	<i>xn</i>
$\{t_b^1, t_b^2\}$	0	0	0	0	$\{t_b^1, t_b^2\}$	0	0	0	0
$\{t_b^3, t_b^4\}$	0	0	0	0	$\{t_b^3, t_b^4\}$	0	0	0	0
$\{t_b^5, t_b^6\}$	0	0	0	1	$\{t_b^5, t_b^6\}$	0	0	0	0
$\{t_b^7, t_b^8\}$	0	0	0	0	$\{t_b^7, t_b^8\}$	0	0	0	1

Figure 3: The beliefs for Ann's types in an extended type space with types of a finite depth and nontrivial higher-order beliefs

3. Extended type spaces

We now begin the formal treatment. Section 3.1 defines the class of extended type spaces, and Section 3.2 demonstrates that each Harsanyi type space can be seen as an extended type space.

3.1. Definition

There is a set of two players, denoted by N ; the results can be extended to any finite number of players. Players are uncertain about the the *state of nature*, which is drawn from a nonempty set Θ . A state of nature $\theta \in \Theta$ could for instance specify the payoff functions of players, or their actions. The set Θ of states of nature is endowed with some σ -algebra \mathcal{F}_Θ . Throughout this paper, if we fix a player i , then the player other than i is denoted by j and vice versa, i.e., $j \neq i$.

A (Θ -based) *extended type space* is a structure

$$(T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N},$$

that satisfies Conditions 1–3 below, where for each player i , T_i is a nonempty set of *types* for player i ; \mathcal{S}_i is a collection of σ -algebras on T_i , assumed to be nonempty and countable. The function Σ_i maps the types in T_i to a σ -algebra $\Sigma_i(t_i) \in \mathcal{S}_j$ on T_j , and β_i maps each type t_i to a probability measure on the product σ -algebra $\mathcal{F}_\Theta \otimes \Sigma_i(t_i)$.

The function β_i is the *belief map* for player i , and the probability measure $\beta_i(t_i)$ is the *belief* of $t_i \in T_i$ over the set of states of nature Θ and the other's type set T_j . The σ -algebra $\Sigma_i(t_i)$ for a type $t_i \in T_i$ for player i on the type set T_j of the other player represents the perception of type t_i of the other's type set. A *state (of the world)* is a tuple (θ, t) , where $t := (t_i)_{i \in N}$ is a type profile, with $t_i \in T_i$ for all i , and $\theta \in \Theta$ is a state of nature.

As noted above, extended type spaces satisfy Conditions 1–3. Condition 1 states that the σ -algebras associated with the types of a player can be completely ordered in terms of set inclusion. That is, the collection of σ -algebras for a player form a filtration:

Condition 1. [Filtration] For each player $i \in N$ and each pair $\mathcal{F}_i, \mathcal{F}'_i \in \mathcal{S}_i$ of σ -algebras, either $\mathcal{F}_i \subseteq \mathcal{F}'_i$ or $\mathcal{F}'_i \subseteq \mathcal{F}_i$.

To state the other two conditions, we need some more notation. Given a player $i \in N$, say that a σ -algebra \mathcal{F}_i on T_i *dominates* a σ -algebra \mathcal{F}_j on T_j if for each $E \in \mathcal{F}_\Theta \otimes \mathcal{F}_j$ and $p \in [0, 1]$,

$$\{t_i \in T_i : E \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i), \beta_i(t_i)(E) \geq p\} \in \mathcal{F}_i.$$

If \mathcal{F}_i dominates \mathcal{F}_j , then we write $\mathcal{F}_i \succ \mathcal{F}_j$; we write $\mathcal{F}_i \not\succeq \mathcal{F}_j$ if \mathcal{F}_i does not dominate \mathcal{F}_j . If \mathcal{F}_i is the coarsest σ -algebra that dominates \mathcal{F}_j , we write $\mathcal{F}_i \succ^* \mathcal{F}_j$. Two σ -algebras \mathcal{F}_i and \mathcal{F}_j on T_i and T_j , respectively, that dominate each other will be called a *mutual-dominance pair*.

Condition 2 states that there is at most one one such pair:

Condition 2. [Unique pair] Fix $i \in N$, and let $\mathcal{F}_i \in \mathcal{S}_i$, $\mathcal{F}_j \in \mathcal{S}_j$. If $(\mathcal{F}_i, \mathcal{F}_j)$ is a mutual-dominance pair, then there is no σ -algebra $\mathcal{F}'_i \neq \mathcal{F}_i$ in \mathcal{S}_i such that $(\mathcal{F}'_i, \mathcal{F}'_j)$ is a mutual-dominance pair for some σ -algebra \mathcal{F}'_j on T_j .

The last condition says that any nontrivial σ -algebra dominates some σ -algebra on the type set of the other player.

Condition 3. [Dominance] For any player $i \in N$ and σ -algebra $\mathcal{F}_i \in \mathcal{S}_i$ such that $\mathcal{F}_i \neq \{T_i, \emptyset\}$, there is $\mathcal{F}_j \in \mathcal{S}_j$ such that

- (a) $(\mathcal{F}_i, \mathcal{F}_j)$ is a mutual-dominance pair; or
- (b) \mathcal{F}_i is the coarsest σ -algebra that dominates \mathcal{F}_j , i.e., $\mathcal{F}_i \succ^* \mathcal{F}_j$.

Each of the Conditions 1–3 is necessary to ensure that types generate a belief hierarchy of a well-defined depth, and to rule out redundant σ -algebras, as we demonstrate in Appendix C.⁷

3.2. Harsanyi type spaces

As the name suggests, extended type spaces generalize Harsanyi type spaces, in the sense that any Harsanyi type space can be viewed as an extended type space. A (Θ -based) *Harsanyi type space* is a structure $(T_i^{\mathcal{H}}, \mathcal{F}_i^{\mathcal{H}}, \beta_i^{\mathcal{H}})_{i \in N}$, where $T_i^{\mathcal{H}}$ is a nonempty set of types for player i , and $\mathcal{F}_i^{\mathcal{H}}$ is a σ -algebra on $T_i^{\mathcal{H}}$. The function $\beta_i^{\mathcal{H}}$ maps each type $t_i \in T_i^{\mathcal{H}}$ into a probability measure $\beta_i^{\mathcal{H}}(t_i)$ on the product σ -algebra $\mathcal{F}_\Theta \otimes \mathcal{F}_j^{\mathcal{H}}$. The set of probability measures on $\mathcal{F}_\Theta \otimes \mathcal{F}_j^{\mathcal{H}}$ is endowed with the σ -algebra that is generated by sets of the form

$$\{\mu : \mu(E) \geq p\} : \quad E \in \mathcal{F}_\Theta \otimes \mathcal{F}_j^{\mathcal{H}}, p \in [0, 1].$$

That is, the set of probability measures on $\mathcal{F}_\Theta \otimes \mathcal{F}_j^{\mathcal{H}}$ is endowed with the coarsest σ -algebra that contains all sets $\{\mu : \mu(E) \geq p\}$ of probability measures that assign probability at least p to the event E , for any $E \in \mathcal{F}_\Theta \otimes \mathcal{F}_j^{\mathcal{H}}$ and $p \in [0, 1]$. The belief maps $\beta_i^{\mathcal{H}}$ are assumed to be measurable. (This specifications covers most of the alternative definitions in the literature, such as where the type sets are required to be separable metrizable or Polish, and the belief maps are assumed to be Borel measurable or continuous.)

Given a Harsanyi type space $(T_i^{\mathcal{H}}, \mathcal{F}_i^{\mathcal{H}}, \beta_i^{\mathcal{H}})_{i \in N}$, we can define a structure $(T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$, where for each player i , $T_i := T_i^{\mathcal{H}}$, and $\beta_i := \beta_i^{\mathcal{H}}$. Also, let $\mathcal{S}_i := \{\mathcal{F}_i^{\mathcal{H}}\}$, so that $\Sigma_i(t_i) = \mathcal{F}_i^{\mathcal{H}}$ for all $t_i \in T_i$. Then:

Proposition 3.1. *The structure $(T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$ derived from a Harsanyi type space satisfies Conditions 1–3. Therefore, any Harsanyi type space can be seen as an extended type space.*

As there are clearly extended type spaces that are not Harsanyi type spaces, such as the one in Figure 2, extended type spaces generalize the Harsanyi framework.

⁷The examples in Appendix C also show that Conditions 1–3 are logically independent: For each of the conditions, there is a structure that satisfies all but one of them. So, none of the conditions is implied by the others.

4. The depth of reasoning of types

This section demonstrates how the σ -algebra associated with a type determines its depth of reasoning. The depth of reasoning of a type is inherently a property of the belief hierarchy that the type induces, where a belief hierarchy summarizes the higher-order beliefs that a type generates. Section 4.1 therefore constructs the space of belief hierarchies generated by types from extended type spaces. Section 4.2 shows that each type has a well-defined depth of reasoning, and provides a first characterization of a type's depth of reasoning. Section 4.3 characterizes a type's depth of reasoning directly in terms of the properties of the type space.

4.1. Belief hierarchies

The first step in characterizing a type's depth of reasoning is to map a type into its belief hierarchy. To that aim, we construct the space of belief hierarchies, where a belief hierarchy specifies a k th-order belief for each k . Unlike in constructions of the space of belief hierarchies generated by types from Harsanyi type spaces (e.g. Mertens and Zamir, 1985; Brandenburger and Dekel, 1993), we need to take into account that types can have a finite depth of reasoning. To accommodate that, the set of k th-order beliefs for player i is endowed with a *collection* of σ -algebras, as opposed to a single one, with different σ -algebras capturing different perceptions for player j of i 's k th-order beliefs.

To construct such a space of belief hierarchies, we need some more notation. Given a set X and a (nonempty) collection \mathcal{S} of σ -algebras on X , let $\Delta(X, \mathcal{S})$ be the set of probability measures on some σ -algebra \mathcal{F} in \mathcal{S} . If μ is a probability measure in $\Delta(X, \mathcal{S})$, then $\Sigma(\mu) \in \mathcal{S}$ is the σ -algebra on which μ is defined. The set $\Delta(X, \mathcal{S})$ is endowed with the σ -algebra $\mathcal{F}_{\Delta(X, \mathcal{S})}$ generated by the sets

$$\{\mu \in \Delta(X, \mathcal{S}) : E \in \Sigma(\mu), \mu(E) \geq p\} : \quad E \in \mathcal{F}, \mathcal{F} \in \mathcal{S}, p \in [0, 1].$$

This σ -algebra naturally separates beliefs (probability measures) according to the probability they assign to events; this choice of σ -algebra makes it possible to talk about “beliefs about beliefs,” and so on (cf. Heifetz and Samet, 1998b).

We are now ready to construct the space of belief hierarchies. The first step is to construct a sequence B_i^0, B_i^1, \dots of spaces for each player i that describe the higher-order beliefs for that player. Formally, for each player $i \in N$, fix an arbitrary “seed” μ_i^0 . Then define

$$B_i^0 := \{\mu_i^0\},$$

and let $\mathcal{F}_{B_i^0} := \{B_i^0, \emptyset\}$ be the trivial σ -algebra on B_i^0 . For every extended type space $\mathcal{T} =$

$(T_n, \mathcal{S}_n, \Sigma_n, \beta_n)_{n \in N}$, define the function $h_i^{\mathcal{T},0}$ from T_i into B_i^0 in the obvious way: $h_i^{\mathcal{T},0}(t_i) := \mu_i^0$ for $t_i \in T_i$.

For $k > 0$, suppose that for each $i \in N$ and $\ell \leq k - 1$ the set B_i^ℓ has been defined, and that $\mathcal{F}_{B_i^\ell}$ is a σ -algebra on B_i^ℓ . Also assume that $h_i^{\mathcal{T},\ell}$ is a function from T_i into B_i^ℓ for each extended type space \mathcal{T} . Then, for each player i , define

$$\mathcal{S}_i^k := \{ \mathcal{F}_\Theta \otimes \mathcal{A}_j^{k-1} : \mathcal{A}_j^{k-1} \text{ is a sub-}\sigma \text{ algebra of } \mathcal{F}_{B_j^{k-1}} \},$$

and

$$B_i^k = B_i^{k-1} \times \Delta(\Theta \times B_j^{k-1}, \mathcal{S}_i^k).$$

Define the σ -algebra $\mathcal{F}_{B_i^k}$ on B_i^k to be the product σ -algebra $\mathcal{F}_{B_i^{k-1}} \otimes \mathcal{F}_{\Delta(\Theta \times B_j^{k-1}, \mathcal{S}_i^k)}$. Finally, for each extended type space $\mathcal{T} = (T_n, \mathcal{S}_n, \Sigma_n, \beta_n)_{n \in N}$, define the function $h_i^{\mathcal{T},k}$ from T_i into B_i^k by:

$$h_i^{\mathcal{T},k}(t_i) = (h_i^{\mathcal{T},k-1}(t_i), \beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T},k-1})^{-1}).$$

Theorem 4.1 below shows that $h_i^{\mathcal{T},k}$ is well-defined.

The interpretation is that for any $k \geq 1$, $h_i^{\mathcal{T},k}(t_i)$ is the belief hierarchy of order k induced by t_i : it specifies the higher-order beliefs for t_i up to order k . The first term, $h_i^{\mathcal{T},k-1}(t_i)$ is the belief hierarchy of order $k - 1$ generated by t_i (whenever $k > 1$). The second term, the probability measure $\beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T},k-1})^{-1}$, is the k th-order belief for type t_i : it gives the belief of type t_i over Θ and the set B_j^{k-1} of belief hierarchies of order $k - 1$ for j . The lower-order beliefs of a type can be obtained from its k th-order belief by appropriate marginalization, owing to the recursive construction. A player thus has a belief about another player's k th-order belief hierarchy if and only if he has a belief about the other's k th-order belief.

A (full) belief hierarchy for player i is then simply a sequence of probability measures that specify the higher-order beliefs for player i at each order. Formally, define the function $h_i^{\mathcal{T}}$ from T_i to $\times_{k=1}^\infty \Delta(\Theta \times B_j^{k-1}, \mathcal{S}_i^k)$ by

$$h_i^{\mathcal{T}}(t_i) := (\beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T},0})^{-1}, \beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T},1})^{-1}, \dots).$$

Then, $h_i^{\mathcal{T}}(t_i)$ is the *belief hierarchy* generated by type t_i in type space \mathcal{T} , with $\beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T},k-1})^{-1}$ the k th-order belief of type t_i .

The next result shows that every type induces a well-defined belief hierarchy

Theorem 4.1. *Let $\mathcal{T} = (T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$ be an extended type space. Then, for each player $i \in N$ and type $t_i \in T_i$,*

- (a) *for each k , $h_i^{\mathcal{T},k}(t_i) \in B_i^k$, that is, t_i induces a k th-order belief hierarchy;*

(b) $h_i^{\mathcal{T}}(t_i) \in \times_{k=1}^{\infty} \Delta(\Theta \times B_j^{k-1}, \mathcal{S}_i^k)$, that is, t_i induces a full belief hierarchy.

The present construction is closely related to the construction of Heifetz and Samet (1998b) of the space of belief hierarchies induced by types from Harsanyi type spaces. Indeed, the set of belief hierarchies for player i induced by types from Harsanyi type spaces is equivalent to the set of belief hierarchies constructed by Heifetz and Samet for the Harsanyi case, and forms a proper subset of the set of belief hierarchies induced by extended type spaces. The critical difference between the present construction and that of Heifetz and Samet is that the k th-order beliefs for player can be defined on different σ -algebras here, to represent different depths of reasoning. The next section defines the depth of reasoning of types, and shows that each type has a well-defined depth.

4.2. Depth of reasoning

We first define the depth of reasoning of a belief hierarchy, before defining the depth of types. Informally, a belief hierarchy has an infinite depth of reasoning if it specifies a well-defined belief about each k th-order event, for every k , where a k th-order event concerns the state of nature and the $(k-1)$ th-order belief of the other player. A belief hierarchy has a finite depth of reasoning d if it has well-defined beliefs about every k th-order event for $k \leq d$, but there exist higher-order events it cannot assign a probability to.

Formally, define H_i to be the set of belief hierarchies generated by some type for i in some extended type space \mathcal{T} , so that H_i is a subset of $\times_{k=1}^{\infty} \Delta(\Theta \times B_j^{k-1}, \mathcal{S}_i^k)$. It follows from Theorem 4.1(b) that H_i is nonempty. Then, fix an extended type space $\mathcal{T} = (T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$, and for each player i , let $B_i^{\mathcal{T}, k}$ be the image of the k th-order hierarchy map $h_i^{\mathcal{T}, k}$ in B_i^k for $k = 0, 1, \dots$, that is, $B_i^{\mathcal{T}, k} := h_i^{\mathcal{T}, k}(T_i)$. By Theorem 4.1(a), $B_i^{\mathcal{T}, k}$ is a (nonempty) subset of B_i^k , and we endow $B_i^{\mathcal{T}, k}$ with the relative σ -algebra $\mathcal{F}_{B_i^{\mathcal{T}, k}}$ induced by $\mathcal{F}_{B_i^k}$. A k th-order event for a player i is then an element of $\mathcal{F}_{\Theta} \otimes \mathcal{F}_{B_j^{\mathcal{T}, k-1}}$, that is, an event that involves the state of nature and the $(k-1)$ th-order belief for j .

Given a player $i \in N$, we define the σ -algebra on $B_i^{\mathcal{T}, k-1}$ that is generated by d th-order events:

$$\mathcal{A}_i^{\mathcal{T}, k-1, d-1} := \left\{ \left(\pi_{B_i^{\mathcal{T}, d-1}}^{B_i^{\mathcal{T}, k-1}} \right)^{-1} (B) : B \in \mathcal{F}_{B_i^{\mathcal{T}, d-1}} \right\}$$

where $d \leq k$ and where $\pi_{B_i^{\mathcal{T}, d-1}}^{B_i^{\mathcal{T}, k-1}}$ is the projection from $B_i^{\mathcal{T}, k-1}$ into $B_i^{\mathcal{T}, d-1}$. That is, $\mathcal{A}_i^{\mathcal{T}, k-1, d-1}$ is the σ -algebra on the space $B_i^{\mathcal{T}, k-1}$ of $(k-1)$ th-order belief hierarchies induced by types for i in \mathcal{T} , formed by lumping together the belief hierarchies that coincide up to order $d-1$. As a player's m th-order belief specifies his beliefs at all lower orders (see Section 4.1), $\mathcal{A}_i^{\mathcal{T}, k-1, m}$

is a refinement of $\mathcal{A}_i^{\mathcal{T},k-1,m-1}$ for all m . In particular, $\mathcal{A}_i^{\mathcal{T},k-1,k-1} = \mathcal{F}_{B_i^{\mathcal{T},k-1}}$ is a refinement of any $\mathcal{A}_i^{\mathcal{T},k-1,m}$.

The depth of reasoning d of a belief hierarchy for player i is the order up to which it distinguishes between different belief hierarchies for player j . Intuitively, if the belief for a type t_i for player i over j 's $(d-1)$ th-order beliefs is defined on the σ -algebra $\mathcal{A}_j^{\mathcal{T},d-1,d-1} = \mathcal{F}_{B_j^{\mathcal{T},d-1}}$, while for any $k > d$, its belief over j 's $(k-1)$ th-order beliefs are defined on $\mathcal{A}_j^{\mathcal{T},k-1,d-1} \subsetneq \mathcal{A}_j^{\mathcal{T},k-1,d}$, then t_i has depth d : the type can reason about all d th-order events in \mathcal{T} , but not about some higher-order events. The formal definition takes into account that a belief hierarchy can be generated by different types, potentially from different type spaces:

Definition. Let $(\mu_i^1, \mu_i^2, \dots) \in H_i$ be a belief hierarchy for player i . Then:

- (a) $(\mu_i^1, \mu_i^2, \dots)$ has *depth* $d = \infty$ if for each extended type space $\mathcal{T} = (T_n, \mathcal{S}_n, \Sigma_n, \beta_n)_{n \in N}$ and type $t_i \in T_i$ such that $h_i^{\mathcal{T}}(t_i) = (\mu_i^1, \mu_i^2, \dots)$,

$$(\text{Id}_{\Theta}, h_j^{\mathcal{T},k-1})^{-1}(E) \in \mathcal{F}_{\Theta} \otimes \Sigma_i(t_i)$$

for every $k = 1, 2, \dots$ and $E \in \mathcal{F}_{\Theta} \otimes \mathcal{F}_{B_j^{\mathcal{T},k-1}}$;

- (b) $(\mu_i^1, \mu_i^2, \dots)$ has *depth* $d = 1, 2, \dots$ if for each extended type space $\mathcal{T} = (T_n, \mathcal{S}_n, \Sigma_n, \beta_n)_{n \in N}$ and type $t_i \in T_i$ such that $h_i^{\mathcal{T}}(t_i) = (\mu_i^1, \mu_i^2, \dots)$, the following hold:

(b1) $(\text{Id}_{\Theta}, h_j^{\mathcal{T},k-1})^{-1}(E) \in \mathcal{F}_{\Theta} \otimes \Sigma_i(t_i)$ for $k = 1, \dots, d-1$ and event $E \in \mathcal{F}_{\Theta} \otimes \mathcal{F}_{B_j^{\mathcal{T},k-1}}$;

(b2) $(\text{Id}_{\Theta}, h_j^{\mathcal{T},k-1})^{-1}(E) \in \mathcal{F}_{\Theta} \otimes \Sigma_i(t_i)$ for $k > d-1$ and $E \in \mathcal{F}_{\Theta} \otimes \mathcal{A}_j^{\mathcal{T},k-1,d-1}$; and for every $m > d-1$,

$$\mathcal{A}_j^{\mathcal{T},k-1,d-1} \subsetneq \mathcal{A}_j^{\mathcal{T},k-1,m}.$$

The last part of (b2) is important: it implies that there are k th-order events that μ_i^k cannot assign a probability to, for any $k > d$. Every belief hierarchy has a well-defined depth:

Theorem 4.2. *For every belief hierarchy $(\mu_i^1, \mu_i^2, \dots) \in H_i$, there is a unique $d \in \mathbb{N} \cup \{\infty\}$ such that $(\mu_i^1, \mu_i^2, \dots)$ has depth d .*

We will say that a type t_i for player i has depth (of reasoning) $d_i(t_i) = d$ if it generates a belief hierarchy of depth d , where $d = 1, 2, \dots$ or $d = \infty$. Theorem 4.2 implies that the depth of reasoning of each type is well-defined.

Given a type, how do we determine its depth of reasoning? A first result is that types with a finer σ -algebra have a greater depth of reasoning:

Proposition 4.3. *For any pair of types t_i, t'_i for player $i \in N$ in an extended type space $\mathcal{T} = (T_n, \mathcal{S}_n, \Sigma_n, \beta_n)_{n \in N}$,*

$$\Sigma_i(t_i) \subsetneq \Sigma_i(t'_i)$$

if and only if $d_i(t_i) < d_i(t'_i)$.

Intuitively, a type of depth k induces a k th-order belief μ^k ; and any ℓ th-order belief μ^ℓ for $\ell \leq k$ can be obtained from μ^k by taking the appropriate marginal, as discussed in Section 4.1. Thus, a k th-order belief contains more information than a belief of lower order, and this translates into a finer σ -algebra for types that have a k th-order belief as opposed to types that only have beliefs at lower orders.

A second result is that types from Harsanyi type spaces have an infinite depth of reasoning, as we would expect:

Proposition 4.4. *Let $(\mu_i^1, \mu_i^2, \dots) \in H_i^{\mathcal{H}}$ be a belief hierarchy for player i induced by a type from a Harsanyi type space. Then $(\mu_i^1, \mu_i^2, \dots)$ has an infinite depth of reasoning.*

Proposition 4.3 gives a partial characterization of a type's depth in terms of the properties of its σ -algebra, and Proposition 4.4 gives a full characterization of a type's depth for a special class of type spaces. But how do we determine a type's depth in arbitrary extended type spaces? In principle, the type's depth be determined by writing out the belief hierarchy it induces: if a type has a well-defined belief over each k th-order event for each k , then the type has an infinite depth of reasoning; otherwise, there is some finite k such that the type has well-defined beliefs over each ℓ th-order event for $\ell \leq k$, but there exist $(k + 1)$ th-order events (and thus m th-order events for $m > k + 1$) it cannot reason about. This method of identifying a type's depth can of course be cumbersome. In the next section, we therefore develop a method to determine the depth of reasoning of a type directly from the properties of the type space.

4.3. Rank and depth

In this section we show that the depth of reasoning of a type can be determined in terms of the properties of the type space. This characterization is complete provided that the set of σ -algebras is sufficiently rich, a requirement that is typically implied by Conditions 1–3. To characterize a type's depth in this way, we first need to introduce the rank of a σ -algebra. The rank of a σ -algebra on the type set for player i characterizes the set of σ -algebras on the type set for player j that the σ -algebra dominates.

Formally, fix an extended type space $\mathcal{T} = (T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$. For each player $i \in N$, define

$$\mathcal{R}_i^1 := \{ \mathcal{F}_i \in \mathcal{S}_i : \text{there is no } \mathcal{F}_j \in \mathcal{S}_j \text{ such that } \mathcal{F}_i \succ \mathcal{F}_j \}.$$

For $\ell = 2, 3, \dots$, let

$$\mathcal{R}_i^\ell := \{ \mathcal{F}_i \in \mathcal{S}_i : \text{there is } \mathcal{F}_j \in \mathcal{R}_j^{\ell-1} \text{ such that } \mathcal{F}_i \succ \mathcal{F}_j, \\ \text{and there is no } \mathcal{F}'_j \in \mathcal{S}_j \setminus \bigcup_{m < \ell} \mathcal{R}_j^m \text{ such that } \mathcal{F}_i \succ \mathcal{F}'_j \}.$$

That is, \mathcal{R}_i^1 is the set of σ -algebras on T_i that do not dominate any σ -algebras on T_j ; \mathcal{R}_i^ℓ for $\ell > 1$ is the set of σ -algebras that dominate some σ -algebra in $\mathcal{R}_j^{\ell-1}$, but no σ -algebra in \mathcal{S}_j that does not belong to any \mathcal{R}_j^m for $m < \ell$. Also, let

$$\mathcal{R}_i^\infty := \{ \mathcal{F}_i \in \mathcal{S}_i : \text{for each } \ell \in \mathbb{N}, \text{ there is } \mathcal{F}_j \in \bigcup_{m \geq \ell} \mathcal{R}_j^m \text{ such that } \mathcal{F}_i \succ \mathcal{F}_j \}.$$

The subsets $\mathcal{R}_i^\infty, \mathcal{R}_i^1, \mathcal{R}_i^2, \dots$ thus classify the σ -algebras in \mathcal{S}_i according to the σ -algebras they dominate. Each σ -algebra in \mathcal{S}_i belongs to precisely one of these subsets, as the next result shows:

Lemma 4.5. *For each player $i \in N$, the sets $\mathcal{R}_i^\infty, \mathcal{R}_i^1, \mathcal{R}_i^2, \dots$ partition \mathcal{S}_i .*

Say that a σ -algebra $\mathcal{F}_i \in \mathcal{S}_i$ for player $i \in N$ has *rank* k if $\mathcal{F}_i \in \mathcal{R}_i^k$, where $k = 1, 2, \dots$ or $k = \infty$. Lemma 4.5 guarantees that the rank of each σ -algebra is well-defined. With some abuse of terminology, we say that a type t_i for i has rank $r_i(t_i) = k$ if its σ -algebra $\Sigma_i(t_i) \in \mathcal{S}_j$ has rank k .

The rank of a type's σ -algebra can be used to characterize a type's depth of reasoning. A first result is that a type's rank is a lower bound on the depth of reasoning:

Proposition 4.6. *For each $i \in N$ and $t_i \in T_i$, the depth of t_i is at least its rank: $d_i(t_i) \geq r_i(t_i)$. In particular, if a type has rank $r_i(t_i) = \infty$, then it has an infinite depth of reasoning.*

To understand the result, consider the following simple case. Suppose a type $t_i \in T_i$ for player i is endowed with the trivial σ -algebra on T_j , i.e., $\Sigma_i(t_i) = \{T_j, \emptyset\}$. If $\Sigma_i(t_i)$ does not dominate the trivial σ -algebra $\{T_i, \emptyset\}$, then t_i has rank 1. Moreover, its σ -algebra $\Sigma_i(t_i)$ does not separate the types for j that differ in their beliefs on $\mathcal{F}_\theta \otimes \{T_i, \emptyset\}$, i.e., the types for j that differ in their belief about θ , or, equivalently, their first-order belief. Thus, t_i has depth 1. In turn, if there is a type t_j whose σ -algebra $\Sigma_j(t_j)$ dominates $\{T_j, \emptyset\}$, but no σ -algebra that is finer, then \mathcal{F}_j has rank $k = 2$. Also, $\Sigma_j(t_j)$ separates types for i that differ in their first-order beliefs, so t_j has depth 2. Proposition 4.6 shows that this intuition generalizes.

Proposition 4.6 already gives us a useful characterization of the depth of reasoning of a type in terms of the rank of its σ -algebra. However, the bound in Proposition 4.6 need not be tight, as the following two examples demonstrate:

$\beta_a(t_a^1)$	H	L
$\{t_b^1, t_b^2\}$	1	0

$\beta_b(t_b^1)$	H	L
t_a^1	1	0
t_a^2	0	0

$\beta_a(t_a^2)$	H	L
t_b^1	0	$\frac{1}{2}$
t_b^2	0	$\frac{1}{2}$

$\beta_b(t_b^2)$	H	L
t_a^1	0	0
t_a^2	1	0

Figure 4: The extended type space in Example 1.

$\beta_a(q_a^1)$	H	L
$\{q_b^1, q_b^2\}$	1	0

$\beta_b(q_b^1)$	H	L
q_a^1	1	0
q_a^2	0	0

$\beta_a(q_a^2)$	H	L
$\{q_b^1, q_b^2\}$	0	1

$\beta_b(q_b^2)$	H	L
q_a^1	0	0
q_a^2	0	1

Figure 5: The extended type space in Example 2.

Example 1. We follow the notation in Section 2, with two players, Ann (a) and Bob (b), and two states of nature, H and L . Ann has two types, t_a^1, t_a^2 and likewise for Bob. The collection \mathcal{S}_a of σ -algebras on Ann's type set is the singleton $\{\mathcal{F}_a\}$, where \mathcal{F}_a is generated by the singletons. The σ -algebras in \mathcal{S}_b are \mathcal{F}_b^0 and \mathcal{F}_b , where \mathcal{F}_b^0 is the trivial σ -algebra (i.e., the σ -algebra generated by the full type set $\{t_b^1, t_b^2\}$), and \mathcal{F}_b is generated by the singletons. Type t_a^1 has σ -algebra $\Sigma_a(t_a^1) = \mathcal{F}_b^0$, and type t_a^2 has $\Sigma_a(t_a^2) = \mathcal{F}_b$. Each type t_b for Bob is endowed with the σ -algebra $\Sigma_b(t_b) = \mathcal{F}_a$. The type space is depicted in Figure 4.

It is easy to check that $\Sigma_a(t_a^1) = \mathcal{F}_b^0$ does not dominate any σ -algebra in $\mathcal{S}_a = \{\mathcal{F}_a\}$, so that t_a^1 has rank 1. However, $\Sigma_a(t_a^1)$ does separate the types for Bob according to their beliefs about nature, so t_a^1 has depth 2. \triangleleft

Example 2. Ann and Bob both have two types, denoted by q_i^1, q_i^2 , where $i = a, b$. The collections of σ -algebras for this type space are $\mathcal{S}'_a = \{\mathcal{F}'_a\}$ and $\mathcal{S}'_b = \{\mathcal{F}'_b\}$, with \mathcal{F}'_a the σ -algebras generated by the singletons, and \mathcal{F}'_b the trivial σ -algebras. The type space is given in Figure 5.

Here, \mathcal{F}'_b has rank 1, and \mathcal{F}'_a has rank 2. But \mathcal{F}'_a distinguishes the types for Ann perfectly, so the types for Bob with σ -algebra \mathcal{F}'_a have infinite depth. \triangleleft

Examples 1 and 2 illustrate the two reasons why a type's rank can be strictly lower than

its depth. First, in Example 1, the σ -algebra that separates the types for Bob based on their first-order beliefs (i.e., dominates the trivial σ -algebra on Ann's type set) coincides with the σ -algebra \mathcal{F}_b^0 that does not separate the types for Bob at all. Because the collection \mathcal{S}_a of σ -algebras does not include the trivial σ -algebra on Ann's type set, the rank of \mathcal{F}_b^0 does not take into account that \mathcal{F}_b^0 separates the types for Bob according to their beliefs about nature, or, equivalently, that \mathcal{F}_b separates the types for Bob that have different beliefs on the product σ -algebra $\mathcal{F}_\Theta \otimes \{\{t_a^1, t_a^2\}, \emptyset\}$. The problem is thus that \mathcal{F}_b^0 dominates some σ -algebra, viz., the trivial σ -algebra on Ann's type set, that corresponds to some finite depth of reasoning, but that is not included in the collection \mathcal{S}_a of σ -algebras on Ann's type set.

Second, in Example 2, the σ -algebra \mathcal{F}'_b not only dominates the trivial σ -algebra on Ann's type set (and thus separates the types for Bob based on their first-order beliefs), but it also dominates the σ -algebra on Ann's type set that is generated by the singletons, which we denote by \mathcal{F}''_a . That σ -algebra dominates \mathcal{F}'_b in turn, and it is easy to see that the σ -algebras in such a mutual-dominance pair correspond to an infinite depth of reasoning.⁸ Thus, the σ -algebra that separates the types for Bob at all orders coincides with the σ -algebra \mathcal{F}'_b that separates the types for Bob based on their first-order beliefs. Because the collection \mathcal{S}'_a of σ -algebras on Ann's type set does not include the σ -algebra \mathcal{F}''_a that corresponds to an infinite depth of reasoning, \mathcal{F}'_b is assigned a finite rank.

This suggests that if for any σ -algebra, the σ -algebras are included that correspond to a finite depth and that are dominated by that σ -algebra, or with which it forms a mutual-dominance pair, then the depth of each type is given by its rank. This motivates the following condition:

Condition 4. Fix $i \in N$ and suppose $\mathcal{F}_i \in \mathcal{S}_i$ is such that $(\mathcal{F}_i, \mathcal{F}_j)$ is not a mutual-dominance pair for any $\mathcal{F}_j \in \mathcal{S}_j$.⁹ If there exists a σ -algebra \mathcal{G}_j on T_j such that

(a) $(\mathcal{F}_i, \mathcal{G}_j)$ is a mutual-dominance pair; or

(b) it holds that

$$\mathcal{F}_i \succ^* \mathcal{G}_j \succ^* \mathcal{G}_i^1 \succ^* \mathcal{G}_j^1 \succ^* \dots \succ^* \mathcal{G}_n^m = \{T_n, \emptyset\}$$

for some $n \in N$ and σ -algebras $\mathcal{G}_i^\ell, \mathcal{G}_j^\ell$, $\ell \leq m$, on T_i, T_j , respectively;

then $\mathcal{G}_j \in \mathcal{S}_j$.

⁸To see this, note that because \mathcal{F}'_b has depth at least 1, and \mathcal{F}''_a separates the types for Ann according to their beliefs on \mathcal{F}'_b , \mathcal{F}''_a has depth at least 2. By a similar argument, \mathcal{F}'_b has depth at least 2. Reiterating this argument gives that \mathcal{F}'_b distinguishes the types for i according to their k th-order beliefs for any k , and likewise for \mathcal{F}''_a . See Lemma A.8 for the formal statement.

⁹This qualification ensures that extended type spaces derived from Harsanyi type spaces satisfy Condition 4.

The next result says that, with this additional condition, a type’s rank fully characterizes its depth:

Theorem 4.7. *Suppose $(T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$ is an extended type space that satisfies Condition 4. Then for each $i \in N$ and $t_i \in T_i$, the depth $d_i(t_i)$ for t_i equals the rank $r_i(t_i)$ of its σ -algebra.*

The importance of this result is that it makes it possible to characterize the depth of reasoning of a type in terms of the properties of a type space alone. That is, we can determine the depth of reasoning of a type without having to write out its belief hierarchy explicitly.

How strong is Condition 4? For the type spaces in Section 2, Condition 4 is implied by Conditions 1–3. This is generally the case. The problem with Examples 1 and 2 is that there, a σ -algebra on Bob’s type set that corresponds to depth k coincides with a σ -algebra that corresponds to some lower depth k' . In such a case, the beliefs at order k' completely determine the beliefs at order k , so that, in a sense, there is not enough richness in Bob’s beliefs at different orders. If a player’s beliefs are not sufficiently rich, then the conditions in Section 3 do not guarantee that the collection of σ -algebras on the other player’s type set includes a σ -algebra that corresponds to depth $k - 1$ (if k is finite) or k (if $k = \infty$). As can be seen from comparing the proofs of Proposition 4.6 and Theorem 4.7, if there is sufficient variation in players’ beliefs at different order, the types’ σ -algebra completely characterizes a type’s depth. On the other hand, if a type’s rank is lower than its depth, the extended type space can be transformed into a type space that satisfies Condition 4, by adding the relevant σ -algebras. In that sense, a completeness condition such as Condition 4 is not a strong one.

5. Higher-order beliefs

The previous section showed that players with a coarse perception have a bounded depth of reasoning. Here we explore the idea that players with a finite depth of reasoning may nevertheless be able to reason about events of arbitrarily high orders under certain conditions. Intuitively, while there are events of order greater than d that a type with a depth of reasoning $d < \infty$ cannot reason about, the type can have beliefs about a higher-order event if the event can be reduced to an event of sufficiently low order. The next section characterizes the higher-order events that a player of finite depth can reason about. Section 5.2 presents the conditions under which common belief and high-order mutual belief can be attained when players have a finite depth of reasoning.

5.1. Characterization

We are interested in the question what higher-order events a player with a finite depth can reason about. To investigate this question, we fix an extended type space $\mathcal{T} = (T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$. The next result states that a type for a player i of finite depth d can reason about a k th-order event if and only if the event is equivalent to a d th-order event.

Theorem 5.1. *Fix a player i and a type $t_i \in T_i$. Suppose $d_i(t_i) = d < \infty$. Then for each event $B \subseteq H_j$ concerning the full belief hierarchies for j , it holds that*

$$\{t_j \in T_j : h_j^{\mathcal{T}}(t_j) \in B\} \in \Sigma_i(t_i)$$

if and only if there is $B_{d-1} \in \mathcal{F}_{B_j^{\mathcal{T}, d-1}}$ such that

$$\{t_j \in T_j : h_j^{\mathcal{T}}(t_j) \in B\} = \{t_j \in T_j : h_j^{\mathcal{T}, d-1}(t_j) \in B_{d-1}\}$$

This result implies that a type for player i that has a finite depth of reasoning d can reason about a k th-order event $E \subseteq \Theta \times B_j^{\mathcal{T}, k-1}$ for $k > d$ if and only if the event is equivalent to an event of sufficiently low order, i.e., an event of order at most d . Thus, we see why certain higher-order events can be hard to reason about, while players with a finite depth can effortlessly reason about others, as discussed in Section 1 and illustrated in Section 2.4. In the type space in Section 2.4, players could reason at all orders about the event that they had eye contact, but Ann was unable to think about the question whether Bob believed that she believed that a burglary had taken place. The reason was that players' higher-order beliefs about the event that they had eye contact were completely determined by the event that there was eye contact, while their higher-order beliefs about the burglary were not pinned down by any lower-order event.

There are also type spaces in which types with a finite depth d cannot reason about any event at higher order; see Kets (2012) for an example. Intuitively, this is the case if there are many different beliefs that types think possible at each order. In that case, no higher-order belief is pinned down by a belief at lower order: at each order, there is a myriad of possible higher-order beliefs.

5.2. Common belief

In the remainder of this section, we will focus on the question to what extent players with a finite depth of reasoning can think about statements of the form “ i believes with high probability that j believes with high probability that...the other player believes with high probability that the state of nature is x ,” i.e., to what extent there can be (approximate)

mutual belief at high order or even common belief. Thus, we investigate the conditions under which there is a lower-order event B_{d-1} that is equivalent to a higher-order event B , as in the statement of Theorem 5.1, for the special case where the event B refers to high-order mutual belief or common belief. Before we can present the results, we need to introduce a belief operator for our setting. This is done in the next section.

5.2.1. Belief operators

We extend the standard notions of mutual p -belief and common p -belief to the context of extended type spaces. These concepts were introduced by Monderer and Samet (1989) in the context of belief spaces.¹⁰ We adapt the definition slightly, to take into account that some types have a finite depth of reasoning.

Throughout this section, we fix an extended type space $\mathcal{T} = (T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$. It will be convenient to write T and t for the set $\times_{i \in N} T_i$ of type profiles and a particular type profile $(t_i)_{i \in N} \in T$, respectively. Let $E \subseteq \Theta \times T$. For a player $i \in N$ and a type $t_i \in T_i$, define

$$E_{t_i} := \{(\theta, t_j) : (\theta, t_i, t_j) \in E\}$$

and for any $p \in [0, 1]$, define

$$B_{i,p}(E) := \{(\theta, t) \in \Theta \times T : E \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i), \beta_i(t_i)(E) \geq p\},$$

to be the event that player i believes E with probability at least p . Then,

$$B_p(E) := \bigcap_{i \in N} B_{i,p}(E)$$

is the event that all players assign probability at least p to E . That is, $B_p(E)$ is the event that E is *mutual p -belief*. The set of states at which E is *k th-order mutual p -belief* is $B_p^k(E)$, where $B_p^1(E) := B_p(E)$, and $B_p^k = B_p \circ B_p^{k-1}$.

Common p -belief in an event obtains if all players believe an event with probability at least p , believe with probability at least p that all players believe the event with probability at least p , and so on. That is, E is *common p -belief* at the state (of the world) (θ, t) if

$$(\theta, t) \in \bigcap_{k=1}^{\infty} B_p^k(E).$$

Note that these definitions deviate from the standard ones in that we can have that there is an event E that some types can reason about (i.e., $E_{t_i} \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i)$ for some t_i). This implies that a type t_i need not put probability greater than $1 - p$ on an event E whenever it does

¹⁰Also see Brandenburger and Dekel (1987) and Fagin and Halpern (1988b).

not believe its complement with probability at least p : it could be that the event E_{t_i} is not in the σ -algebra of t_i , even if other types can reason about E . The belief operator nevertheless satisfies many of the standard properties of belief operators, as shown in Appendix B.

5.2.2. Result

In this section we characterize the conditions under which events can be common p -belief, even if players have a bounded depth. The first step is to extend a result of Monderer and Samet (1989), who characterized the conditions under which an event can be approximate common belief in a Harsanyi type space.¹¹ To state the result, we need one more definition. An event $E \subseteq \Theta \times T$ is an *evident p -belief event* if

$$E \subseteq B_{i,p}(E)$$

for all $i \in N$. That is, an event is evident p -belief if every player believes it with probability at least p whenever it occurs.

Proposition 5.2. *Let $F \subseteq \Theta \times T$ and $p \in (0, 1]$. Then, F is common p -belief at (θ, t) if and only if*

- (1) *there is an evident p -belief event E such that $(\theta, t) \in E$,*

and for all $i \in N$,

- (2) *$E \subseteq B_{i,p}(F)$;*

- (3) *for every type t_i such that $E_{t_i} \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i)$ and $\beta_i(t_i)(E_{t_i}) \geq p$,*

$$[B_p^k(F)]_{t_i} \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i)$$

for every $k = 1, 2, \dots$

Conditions (1) and (2) are essentially the conditions in Monderer and Samet (1989). The conditions require that whenever the evident event E occurs, all players believe F with probability at least p .

Condition (3) says that whenever a player believes the evident event E , he can reason about all higher-order beliefs involving F at arbitrarily high orders. This does not always hold: the online appendix presents an example in which there is an evident event that implies that each player believes another event F (i.e., conditions (1) and (2) above are satisfied), yet F is not commonly believed at any state, because condition (3) is not satisfied.

¹¹Monderer and Samet (1989) stated their result in the context of belief spaces, but their result can be recast in the Harsanyi framework.

On the other hand, condition (3) is naturally satisfied in the case of public announcements, or in situations such as the one in Section 2.4, where there is eye contact (cf. Monderer and Samet, 1989). In that case, there is a single event that generates beliefs at all orders. More generally, if there is a finite k such that k th-order belief in an event implies $(k + 1)$ th-order belief in the event, i.e., if $B_p^k(F) \subseteq B_p^{k+1}(F)$, then common belief can be attained whenever there is a type for each player that has k th-order belief in the event: Because $B_p^{\ell+1}(F) \subseteq B_p^\ell(F)$ for all ℓ , as can easily be verified, we then have that $C_p(F) = \bigcap_\ell B_p^\ell = B_p^k(F)$. That is, it is sufficient that players can reason about the events $B_p(F), \dots, B_p^{k-1}(F)$ in this case.

Thus, there are cases in which common belief is attained with a finite set of events that indicate players' beliefs at higher-orders. The next result shows that this is not the case if there is merely mutual belief at arbitrarily high orders, but not common belief:

Proposition 5.3. *Suppose $B_p^k(F) \neq \emptyset$ for all k , and that for each player $i \in N$, there is a finite set \mathcal{B}_i of events in $\mathcal{F}_{B_i^{\mathcal{T}, d_i-1}}$ for some $d_i < \infty$ such that for each k ,*

$$\{t_i \in T_i : [B_p^k(F)]_{t_i} \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i), \beta_i(t_i)([B_p^k(F)]_{t_i}) \geq p\} = \{t_i \in T_i : h_i^{\mathcal{T}, d_i-1}(t_i) \in B_i\}$$

for some $B_i \in \mathcal{B}_i$. Then $C_p(F) \neq \emptyset$.

Thus, if player i has depth at most $d_j < \infty$, i.e., $d_i(t_i) \leq d_j$ for all $t_i \in T_i$, then mutual belief at arbitrarily high order implies common belief, unless there is an infinite set \mathcal{B}_i of lower-order events for each player i that indicate mutual belief in F at all higher orders. In that sense, mutual belief at arbitrarily high orders is harder to attain than common belief. Clark and Marshall (1978) give a number of examples that illustrate how contrived the lower-order events are that indicate k th-order belief but not $(k + 1)$ th-order belief in an event. The online appendix gives an example of a type space with types of depth at most 2 in which there is mutual belief at arbitrarily high orders.

6. Related literature

6.1. Limited sophistication and unawareness

A possible interpretation of a player's inability to assign probabilities to a given event is that this event is ambiguous to her. Ahn (2007) introduces a class of type spaces with ambiguous beliefs, where ambiguous beliefs are represented by sets of probability measures. If the beliefs of a type are given by a set of probability measures, then the type's beliefs are not pinned down for events on which these probability measures differ; see Halpern (2003, Thm. 2.3.3) for a result along these lines. However, there seems to be a difference at the intuitive

level between a player with a finite depth of reasoning and a player with ambiguous beliefs about higher-order beliefs. In particular, it is not clear that the assumption of ambiguity aversion would be natural in the present context.

Alternatively, one could view a player with a bounded depth of reasoning as a player whose preferences are incomplete when it comes to states that differ only in the beliefs of the other player at very high orders. Di Tillio (2008) considers a class of type structures that allow for incomplete preferences, but does not investigate this issue further.

Finally, the inability of players to distinguish certain states on the basis of others' higher-order beliefs can be interpreted in terms of unawareness: a player may simply not realize that he could distinguish the states more finely. This is a different form of unawareness than the one commonly studied in the literature, which assumes that players may be unaware of certain aspects of the state of nature; see Geanakoplos (1989), Feinberg (2004), and Heifetz et al. (2006), among others.¹² This may not be the most natural interpretation, however: if a player reasons about the third-order beliefs of another player, but not about his fourth-order beliefs, then this seems to be more a matter of limited computational powers than of unawareness. Indeed, the present framework differs from standard models of unawareness in that the belief operator we define satisfies negative introspection: if a player does not believe an event, then she believes that she does not believe it (Appendix B.2).

6.2. Type spaces of finite order

At the technical level, there is a connection between this paper and the literature on what we will call “finite-order Harsanyi type spaces” (Morris et al., 1995; Heifetz and Samet, 1998a; Qin and Yang, 2010).¹³ While this literature does not model players with a finite depth of reasoning, the higher-order beliefs induced by these spaces take a particularly simple form: there is some finite order k such that for each player, conditional on his k th-order beliefs, his higher-order beliefs (about every possible event) are common knowledge. The connection with our work is that a type of a finite depth d can reason about an event of order greater than d

¹²Fagin and Halpern (1988a) propose a number of logic models that can capture various forms of unawareness. Heinsalu (2011) proves the existence of a universal type space for a class of type spaces in which players may be unaware of certain states of nature or players, and where players may also be bounded in their reasoning about others in the sense of level- k or cognitive-hierarchy models (so that types with a finite depth of reasoning cannot have beliefs about certain higher-order events, unlike here).

¹³The literature refers to these spaces as belief spaces of finite depth, or, alternatively, rank, where “rank” and “depth” are a property of the belief space as a whole, not of (the σ -algebras associated with) individual types, as is the case here. To avoid any confusion, we use the terms rank and depth only in the sense used in the rest of the paper.

if and only if the event can be expressed in terms of a d th-order event (Theorem 5.1). This implies that in an extended type space of a finite order k , a type either has depth at most k , or it has an infinite depth of reasoning. Intuitively, if a type can reason about the k th-order beliefs of his opponents (i.e., has depth at least $k + 1$), then it can reason about his opponents' beliefs at all orders, as these beliefs are completely transparent to the type.

The focus of that literature is thus very different. While the existing literature concentrates on characterizing the “order” of a type space as a whole (i.e., the set of states of the world), assuming that all players have an infinite depth of reasoning, we are interested in the “order” of individual events (i.e., proper subsets of the set of states of the world), to characterize the set of events a type of a finite depth can reason about. Also, methodologically, we identify conditions on the type space (Conditions 1–3) such that the σ -algebras lump together precisely the belief hierarchies that coincide up to some order k , whereas the literature is concerned with belief hierarchies only.

6.3. Measurable structures on type sets

One insight of the present paper is that, by choosing the measurable sets on which a type's belief is defined, we can get types that can reason about only finitely many orders of beliefs. This idea fits in with a broader literature that the measurable structure associated with types in Harsanyi type spaces implicitly imposes restrictions on reasoning, i.e., on belief hierarchies Brandenburger and Keisler (2006); Friedenberg and Meier (2012); Friedenberg (2010); Friedenberg and Keisler (2011); see Friedenberg and Keisler (2011) for a detailed discussion.

6.4. A universal extended type space?

Type spaces are of course just a convenient device to model the higher-order beliefs of players. We therefore want to make sure that any belief hierarchy can be modeled by some type in some (extended) type space. Here, as in Heifetz and Samet (1998b), this is tautologically true: the set of belief hierarchies is the set of belief hierarchies generated by some type in some type space. However, unlike in the Harsanyi case, the space of belief hierarchies constructed here cannot easily be transformed into a universal extended type space \mathcal{T}^* such that for every (coherent) belief hierarchy, there is a type in \mathcal{T}^* with that hierarchy. Following an approach analogous to that of Heifetz and Samet (1998b) would give a structure that is not an extended type space, as the structure would violate Condition 1. The problem is that types from different type spaces that have a finite depth of reasoning may lump together different subsets

of belief hierarchies. Whether this problem can be overcome is an open question.¹⁴

Appendix A Proofs for Sections 3 and 4

A.1 Preliminary results

A.1.1 Basic results

We start with some basic results. The first result says that taking inverse images preserves σ -algebras:

Lemma A.1. *Let $f : X \rightarrow Y$ be a function from X into Y , and let \mathcal{E} be a nonempty collection of subsets of Y . Then,*

$$\sigma(\{f^{-1}(E) : E \in \mathcal{E}\}) = \{f^{-1}(E) : E \in \sigma(\mathcal{E})\},$$

where $\sigma(\mathcal{E})$ is the σ -algebra generated by \mathcal{E} .

The proof is standard (e.g., Aliprantis and Border, 2005, Lemma 4.23), and thus omitted.

A useful result is that measurability of the belief maps is equivalent to mutual dominance of the relevant σ -algebras. To state the result, we need some more notation. Fix an extended type space $\mathcal{T} = (T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$, and suppose that for each $i \in N$, there is $\mathcal{F}_i \in \mathcal{S}_i$ such that $\mathcal{F}'_i \subseteq \mathcal{F}_i$ for all $\mathcal{F}'_i \in \mathcal{S}_i$. That is, \mathcal{F}_i is the finest σ -algebra in \mathcal{S}_i . For $i \in N$, let

$$\mathcal{S}_i := \{\mathcal{F}_\Theta \otimes \mathcal{F}'_j : \mathcal{F}'_j \in \mathcal{S}_j\}$$

be the collection of product σ -algebras on $\Theta \times T_j$ on which the belief of a type for i can be defined. Denote the set of probability measures on a σ -algebra in \mathcal{S}_i by $\Delta(\Theta \times T_j, \mathcal{S}_i)$, and denote the σ -algebra that is generated by sets of the form

$$\{\mu \in \Delta(\Theta \times T_j, \mathcal{S}_i) : \mu(E) \geq p\} : \quad E \in \mathcal{F}_\Theta \otimes \mathcal{F}_j, p \in [0, 1]$$

by $\mathcal{F}_{\Delta(\Theta \times T_j, \mathcal{S}_i)}$.

Lemma A.2. *Let $(T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$ be an extended type space, and suppose that for $i \in N$, there is $\mathcal{F}_i \in \mathcal{S}_i$ such that $\mathcal{F}'_i \subseteq \mathcal{F}_i$ for all $\mathcal{F}'_i \in \mathcal{S}_i$. The following are equivalent:*

¹⁴Pintér and Udvari (2011) show that a universal type space can be constructed for a class of type spaces that (strictly) includes the extended type spaces (in particular, Conditions 1–3 need not hold). Since these universal spaces are not extended type spaces, it appears that the depth of reasoning of a type need not be defined; cf. Appendix C.

(a) for each player i , the belief map β_i is $\mathcal{F}_i/\mathcal{F}_{\Delta(\Theta \times T_j, \mathcal{S}_i)}$ -measurable;

(b) $(\mathcal{F}_i, \mathcal{F}_j)$ is a mutual-dominance pair.

Proof. Fix $i \in N$ and, as always, let $j \neq i$. It follows from Lemma A.1 that the belief map β_i is $\mathcal{F}_i/\mathcal{F}_{\Delta(\Theta \times T_j, \mathcal{S}_i)}$ -measurable if and only if for every $E \in \mathcal{F}_\Theta \otimes \mathcal{F}_j$ and $p \in [0, 1]$,

$$\{t_i \in T_i : E \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i), \beta_i(t_i)(E) \geq p\} \in \mathcal{F}_i,$$

which is equivalent to $\mathcal{F}_i \succ \mathcal{F}_j$. Likewise, β_j is $\mathcal{F}_j/\mathcal{F}_{\Delta(\Theta \times T_i, \mathcal{S}_j)}$ -measurable if and only if $\mathcal{F}_j \succ \mathcal{F}_i$. \square

The next result is technical in nature, but will be used in various proofs. It puts restrictions on the type of dominance relations a σ -algebra can be involved in, if the σ -algebra is finer than a σ -algebra that is part of a mutual-dominance pair.

Lemma A.3. Fix an extended type space $\mathcal{T} = (T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$. Let $i \in N$ and $\mathcal{F}_i \in \mathcal{S}_i$. Suppose there exist σ -algebras \mathcal{G}_i and \mathcal{G}_j on T_i and T_j , respectively, such that $(\mathcal{G}_i, \mathcal{G}_j)$ is a mutual-dominance pair. If $\mathcal{F}_i \supsetneq \mathcal{G}_i$, then there is no $\mathcal{F}_j \in \mathcal{S}_j$ such that $\mathcal{F}_i \succ^* \mathcal{F}_j$ and $\mathcal{F}_j \not\prec \mathcal{F}_i$.

Proof. Suppose by contradiction that there exist $\mathcal{F}_i \in \mathcal{S}_i$ and $\mathcal{F}_j \in \mathcal{S}_j$ such that $\mathcal{F}_i \supsetneq \mathcal{G}_i$, $\mathcal{F}_i \succ^* \mathcal{F}_j$, and $\mathcal{F}_j \not\prec \mathcal{F}_i$.

First consider the case that $\mathcal{F}_j = \{T_j, \emptyset\}$. In that case, $\mathcal{G}_j \supseteq \mathcal{F}_j$. As $\mathcal{G}_i \succ \mathcal{G}_j$, it follows that $\mathcal{G}_i \succ \mathcal{F}_j$. But this contradicts that \mathcal{F}_i is the coarsest σ -algebra that dominates \mathcal{F}_j .

Next consider the case that $\mathcal{F}_j \neq \{T_j, \emptyset\}$. By Condition 3, there is $\mathcal{F}'_i \in \mathcal{S}_i$ such that $\mathcal{F}_j \succ \mathcal{F}'_i$. If $\mathcal{F}'_i \succ \mathcal{F}_j$, then $\mathcal{F}'_i \supseteq \mathcal{F}_i$, as \mathcal{F}_i is the coarsest σ -algebra that dominates \mathcal{F}_j . But then $\mathcal{F}_j \succ \mathcal{F}'_i$ implies that $\mathcal{F}_j \succ \mathcal{F}_i$, a contradiction.

So, by Condition 3 and the above, $\mathcal{F}_j \succ^* \mathcal{F}'_i$, and $\mathcal{F}'_i \not\prec \mathcal{F}_j$. Repeating this argument gives that there is a finite chain

$$\mathcal{F}'_i \supsetneq \mathcal{F}_j \supsetneq \mathcal{F}'_i \supsetneq \mathcal{F}_j \supsetneq \dots \supsetneq \mathcal{F}'_n = \{T_n, \emptyset\}$$

for some $m = 1, 2, \dots$ and $n \in N$, where for all $r \in N, s \neq r$ and $\ell \leq m$, $\mathcal{F}'_r \in \mathcal{S}_n$, and $\mathcal{F}'_r \not\prec \mathcal{F}'_s \supsetneq \mathcal{F}'_r$.

Consequently, there is $\ell \leq m - 1, r \in N$, and $s \neq r$ such that $\mathcal{F}'_r \supsetneq \mathcal{G}_r$ and $\mathcal{F}'_s \supsetneq \mathcal{G}_s$. As $\mathcal{F}'_s \supsetneq \mathcal{G}_s$, $\mathcal{G}_r \succ \mathcal{G}_s$ implies that $\mathcal{G}_r \succ \mathcal{F}'_s$. But this contradicts that \mathcal{F}'_r is the coarsest σ -algebra that dominates \mathcal{F}'_s . \square

A.1.2 A key lemma

Lemma A.4 below plays a central role in many of the results in Section 4. We prove the lemma for a more general class of structures than the class of extended type spaces, as the more general result will be useful when we examine the role of Conditions 1–3 in Appendix C. That is, we consider structures of the form $(T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$ which need not satisfy Conditions 1–3.

To state the result, we need some additional notation. Given a structure $\mathcal{T} = (T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$, a player $i \in N$ and $k = 0, 1, \dots$, let $\sigma(h_i^{\mathcal{T}, k})$ be the σ -algebra generated by $h_i^{\mathcal{T}, k}$, i.e.,

$$\sigma(h_i^{\mathcal{T}, k}) := \{(h_i^{\mathcal{T}, k})^{-1}(B) : B \in \mathcal{F}_{B_i^{\mathcal{T}, k}}\},$$

whenever $h_i^{\mathcal{T}, k}$ is well-defined. It is easy to see that if $h_i^{\mathcal{T}, k}$ and $h_i^{\mathcal{T}, k+1}$ are well-defined, then

$$\sigma(h_i^{\mathcal{T}, k}) \subseteq \sigma(h_i^{\mathcal{T}, k+1}),$$

a result we use without explicit mention. We also use the notation $B_i^{\mathcal{T}, k}$ for the image $h_i^{\mathcal{T}, k}(T_i)$ of T_i in B_i^k , whenever $h_i^{\mathcal{T}, k}$ is well-defined. In that case, we endow $B_i^{\mathcal{T}, k}$ with the relative σ -algebra induced by $\mathcal{F}_{B_i^k}$.

Lemma A.4. *Let $\mathcal{T} = (T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$ be a structure that satisfies Condition 3. Then,*

- (a) *for each $i \in N$ and $k = 0, 1, \dots$, $h_i^{\mathcal{T}, k}(t_i) \in B_i^k$ for each $t_i \in T_i$;*
- (b) *for each $i \in N$ and $k = 0, 1, \dots$, $\sigma(h_i^{\mathcal{T}, k+1}) \succ^* \sigma(h_i^{\mathcal{T}, k})$;*
- (c) *for each $\mathcal{F}_i \in \mathcal{S}_i$, one of the following is the case:*
 - *there is $m = 0, 1, \dots$ such that $\mathcal{F}_i = \sigma(h_i^{\mathcal{T}, m})$; or*
 - *$\mathcal{F}_i \supseteq \sigma(h_i^{\mathcal{T}, m})$ for all m .*

Proof. Clearly, $h_i^{\mathcal{T}, 0}(t_i) \in B_i^0$ for $i \in N$ and $t_i \in T_i$. Let $i \in N$ and $t_i \in T_i$. We want to show that $h_i^{\mathcal{T}, 1}(t_i) \in B_i^1$. This holds if and only if $\beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T}, 0})^{-1}$ is a probability measure in $\Delta(\Theta \times B_j^0, \mathcal{S}_i^1)$. As $\Sigma_i(t_i) \supseteq \sigma(h_j^{\mathcal{T}, 0}) = \{T_j, \emptyset\}$, this holds, so that $h_i^{\mathcal{T}, 1}(t_i) \in B_i^1$.

We next show that $\sigma(h_i^{\mathcal{T}, 1}) \succ^* \sigma(h_j^{\mathcal{T}, 0})$. To see this, note that

$$\sigma(h_i^{\mathcal{T}, k}) = \left\{ \{t_i \in T_i : \beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T}, 0})^{-1} \in B\} : B \in \mathcal{F}_{\Delta(\Theta \times B_j^0, \mathcal{S}_i^1)} \right\},$$

so that by Lemma A.1, $\sigma(h_i^{\mathcal{T}, 1})$ is generated by sets of the form

$$\{t_i \in T_i : E \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i), \beta_i(t_i)(E) \geq p\} : \quad E \in \mathcal{F}_\Theta \otimes \sigma(h_j^{\mathcal{T}, 0}), p \in [0, 1].$$

That is, $\sigma(h_i^{\mathcal{T}, 1}) \succ^* \sigma(h_j^{\mathcal{T}, 0})$.

Finally, suppose $\mathcal{F}_i \subsetneq \sigma(h_i^{\mathcal{T},1})$. We claim that $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},0})$. Suppose by contradiction that $\mathcal{F}_i \neq \sigma(h_i^{\mathcal{T},0}) = \{T_i, \emptyset\}$. Then, by Condition 3, there is $\mathcal{F}_j \in \mathcal{S}_j$ such that $\mathcal{F}_i \succ \mathcal{F}_j$, and it follows that $\mathcal{F}_i \succ \sigma(h_j^{\mathcal{T},0})$, as $\sigma(h_j^{\mathcal{T},0}) = \{T_j, \emptyset\}$ is the coarsest σ -algebra on T_j . But then $\mathcal{F}_i \supseteq \sigma(h_i^{\mathcal{T},1})$ (because $\sigma(h_i^{\mathcal{T},1}) \succ^* \sigma(h_j^{\mathcal{T},0})$), which gives a contradiction.

For $k > 1$, suppose that for each $i \in N$, the following hold:

$$(a') \quad h_i^{\mathcal{T},k-1}(t_i) \in B_i^k \text{ for each } t_i \in T_i;$$

$$(b') \quad \sigma(h_i^{\mathcal{T},k-1}) \succ^* \sigma(h_j^{\mathcal{T},k-2});$$

$$(c') \quad \text{for each } \mathcal{F}_i \in \mathcal{S}_i \text{ such that } \mathcal{F}_i \subsetneq \sigma(h_i^{\mathcal{T},k-1}), \text{ there is } m < k-1 \text{ such that } \mathcal{F}_i = \sigma(h_i^{\mathcal{T},m}).$$

The proof follows from the following four claims. The first claim says that even though the σ -algebras need not form a filtration (i.e., \mathcal{T} need not satisfy Condition 1), there is some ordering of the σ -algebras.

Claim 1. For $i \in N$ and $\mathcal{F}_i \in \mathcal{S}_i$, either $\mathcal{F}_i \subseteq \sigma(h_i^{\mathcal{T},k-1})$ or $\mathcal{F}_i \supseteq \sigma(h_i^{\mathcal{T},k-1})$.

Proof of Claim 1. Let $i \in N$ and $\mathcal{F}_i \in \mathcal{S}_i$. If $\mathcal{F}_i = \{T_i, \emptyset\}$, then $\mathcal{F}_i \subseteq \sigma(h_i^{\mathcal{T},k-1})$, so suppose $\mathcal{F}_i \not\subseteq \{T_i, \emptyset\}$. By Condition 3, there is $\mathcal{F}_j^1 \in \mathcal{S}_j$ such that $\mathcal{F}_i \succ \mathcal{F}_j^1$. It will be convenient to write $\mathcal{F}_i^0 := \mathcal{F}_i$.

We first note that there is no finite chain of the form

$$\mathcal{F}_i^0 \succ^* \mathcal{F}_j^1 \succ^* \mathcal{F}_i^2 \succ^* \dots \succ^* \mathcal{F}_n^m$$

for some $m = 1, 2, \dots$ and $n \in N$, where for all $r \in N, s \neq r$ and $\ell \leq m$, $\mathcal{F}_r^\ell \in \mathcal{S}_n$, and $\mathcal{F}_r^\ell \not\prec \mathcal{F}_s^{\ell-1}$, and

$$\mathcal{F}_n^m \succ \mathcal{F}_k, \mathcal{F}_k \succ \mathcal{F}_n^m,$$

where $k \neq n$, and $\mathcal{F}_k \in \mathcal{S}_k$. For suppose not. As $\mathcal{F}_k^{m-1} \succ^* \mathcal{F}_n^m$, it holds that $\mathcal{F}_k \supseteq \mathcal{F}_k^{m-1}$. But then $\mathcal{F}_n^m \succ \mathcal{F}_k$ implies that $\mathcal{F}_n^m \succ \mathcal{F}_k^{m-1}$, a contradiction.

Applying Condition 3 repeatedly then gives that one of the following is the case:

(i) there is a chain of the form

$$\mathcal{F}_i^0 \succ^* \mathcal{F}_j^1 \succ^* \mathcal{F}_i^2 \succ^* \dots \succ^* \mathcal{F}_n^m = \{T_n, \emptyset\} \tag{A.1}$$

for some $m = 1, 2, \dots$ and $n \in N$, where for all $r \in N, s \neq r$ and $\ell \leq m$, $\mathcal{F}_r^\ell \in \mathcal{S}_n$, and $\mathcal{F}_r^\ell \not\prec \mathcal{F}_s^{\ell-1}$;

(ii) there is $\mathcal{F}_j \in \mathcal{S}_j$ such that $(\mathcal{F}_i^0, \mathcal{F}_j)$ is a mutual-dominance pair;

(iii) there is a (potentially infinite) cycle of the form

$$\dots \succ^* \mathcal{F}_i^{-2} \succ^* \mathcal{F}_j^{-1} \succ^* \mathcal{F}_i^0 \succ^* \mathcal{F}_j^1 \succ^* \mathcal{F}_i^2 \succ^* \dots, \quad (\text{A.2})$$

where $\mathcal{F}_n^\ell \in \mathcal{S}_n$ for all $n \in N$ and ℓ .

First suppose there is a chain of the form (A.1). If $m \leq k - 1$, then the induction hypothesis (b') gives that $\mathcal{F}_i^0 \subseteq \sigma(h_i^{\mathcal{T},k-1})$. So suppose $m > k - 1$. In that case, it follows from the induction hypothesis (b') that

$$\mathcal{F}_i^0 \succ^* \mathcal{F}_j^1 \succ^* \mathcal{F}_i^2 \succ^* \dots \succ^* \mathcal{F}_i^{m-(k-1)} \succ^* \sigma(h_j^{\mathcal{T},k-2})$$

or

$$\mathcal{F}_i^0 \succ^* \mathcal{F}_j^1 \succ^* \mathcal{F}_i^2 \succ^* \dots \succ^* \mathcal{F}_j^{m-(k-1)} \succ^* \sigma(h_i^{\mathcal{T},k-2}).$$

We claim that $\mathcal{F}_i^0 \supseteq \sigma(h_i^{\mathcal{T},k-1})$. We prove the result for the first case; the proof for the second case is similar. By the induction hypothesis (b'), $\mathcal{F}_i^{m-(k-1)} \succ^* \sigma(h_j^{\mathcal{T},k-2})$ implies that $\mathcal{F}_i^{m-(k-1)} = \sigma(h_i^{\mathcal{T},k-1})$. Therefore, $\mathcal{F}_j^{m-(k-1)-1} \succ^* \mathcal{F}_i^{m-(k-1)}$ implies that $\mathcal{F}_j^{m-(k-1)-1} \succ \sigma(h_i^{\mathcal{T},k-2})$. It follows that $\mathcal{F}_j^{m-(k-1)-1} \supseteq \sigma(h_j^{\mathcal{T},k-1})$. Repeating this argument gives that $\mathcal{F}_i^0 \supseteq \sigma(h_j^{\mathcal{T},k-1})$.

Next suppose there is $\mathcal{F}_j \in \mathcal{S}_j$ such that $\mathcal{F}_i^0 \succ \mathcal{F}_j$ and $\mathcal{F}_j \succ \mathcal{F}_i^0$. As $\mathcal{F}_j \supseteq \sigma(h_j^{\mathcal{T},0}) = \{T_j, \emptyset\}$ and $\sigma(h_i^{\mathcal{T},1}) \succ^* \sigma(h_j^{\mathcal{T},0})$, it follows that $\mathcal{F}_i^0 \supseteq \sigma(h_i^{\mathcal{T},1})$. Likewise, $\mathcal{F}_j \supseteq \sigma(h_j^{\mathcal{T},1})$. Repeating this argument (using the induction hypothesis (b')), we get that $\mathcal{F}_i^0 \supseteq \sigma(h_i^{\mathcal{T},k-1})$.

Finally, suppose there is a cycle of the form (A.2). Again, $\mathcal{F}_j^1 \supseteq \sigma(h_j^{\mathcal{T},0})$, so that $\mathcal{F}_i^0 \supseteq \sigma(h_i^{\mathcal{T},1})$. Similarly, using that $\mathcal{F}_i^2 \supseteq \sigma(h_i^{\mathcal{T},0})$, we obtain that $\mathcal{F}_j^1 \supseteq \sigma(h_j^{\mathcal{T},1})$. Repeatedly applying this argument gives that $\mathcal{F}_i^0 \supseteq \sigma(h_i^{\mathcal{T},k-1})$. \square

Before proceeding to the next step, note that if $\mathcal{F}_i \subseteq \sigma(h_i^{\mathcal{T},k-1})$, then $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},m})$ for some $m \leq k - 1$ (by the induction hypothesis (c')).

Claim 2. For $i \in N$ and $t_i \in T_i$, $h_i^{\mathcal{T},k}(t_i) \in B_i^k$.

Proof of Claim 2. Fix $i \in N$ and $t_i \in T_i$. By the induction hypothesis (a'), $h_i^{\mathcal{T},k}(t_i) \in B_i^k$ if and only if $\beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T},k-1})^{-1}$ is a probability measure in $\Delta(\Theta \times B_j^{k-1}, \mathcal{S}_i^k)$. That is, there is a sub- σ algebra $\mathcal{G}_j^{k-1} \subseteq \mathcal{F}_{B_j^{k-1}}$ such that for each $E \in \mathcal{G}_j^{k-1}$, $(h_j^{\mathcal{T},k-1})^{-1}(E) \in \Sigma_i(t_i)$, and there is no sub- σ algebra $\tilde{\mathcal{G}}_j^{k-1} \subseteq \mathcal{F}_{B_j^{k-1}}$ such that $\tilde{\mathcal{G}}_j^{k-1} \supsetneq \mathcal{G}_j^{k-1}$ with $(h_j^{\mathcal{T},k-1})^{-1}(\tilde{E}) \in \Sigma_i(t_i)$ for every $\tilde{E} \in \tilde{\mathcal{G}}_j^{k-1}$.

By Claim 1, it holds that $\Sigma_i(t_i) \subsetneq \sigma(h_j^{\mathcal{T},k-1})$ or $\Sigma_i(t_i) \supseteq \sigma(h_j^{\mathcal{T},k-1})$. First suppose that $\Sigma_i(t_i) \subsetneq \sigma(h_j^{\mathcal{T},k-1})$. By the induction hypothesis (c'), $\Sigma_i(t_i) = \sigma(h_j^{\mathcal{T},m})$ for some $m < k - 1$; if there are multiple $m < k - 1$ for which this holds (i.e., $\Sigma_i(t_i) = \sigma(h_j^{\mathcal{T},m}) = \sigma(h_j^{\mathcal{T},m'}$) for $m \neq m'$), we can take the largest one.

Using that $\mathcal{F}_{B_j^{\mathcal{T},m}}$ is the relative σ -algebra on $B_j^{\mathcal{T},m} = h_j^{\mathcal{T},m}(T_j)$ induced by $\mathcal{F}_{B_j^m}$, we obtain

$$\begin{aligned}\sigma(h_j^{\mathcal{T},m}) &= \left\{ (h_j^{\mathcal{T},m})^{-1}(B) : B \in \mathcal{F}_{B_j^{\mathcal{T},m}} \right\} \\ &= \left\{ (h_j^{\mathcal{T},m})^{-1}(B \cap B_j^{\mathcal{T},m}) : B \in \mathcal{F}_{B_j^m} \right\} \\ &= \left\{ (h_j^{\mathcal{T},m})^{-1}(B) : B \in \mathcal{F}_{B_j^m} \right\}.\end{aligned}\tag{A.3}$$

Consequently, $\beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T},k-1})^{-1}$ is a probability measure on the product σ -algebra $\mathcal{F}_\Theta \otimes \left\{ (\pi_{B_j^m}^{B_j^{k-1}})^{-1}(B) : B \in \mathcal{F}_{B_j^{m-1}} \right\}$, where $\pi_{B_j^m}^{B_j^{k-1}}$ is the projection of B_j^{k-1} into B_j^m . As the projection function is measurable, this is a sub- σ algebra of $\mathcal{F}_\Theta \otimes \mathcal{F}_{B_j^{k-1}}$.

Finally, if $\Sigma_i(t_i) \supseteq \sigma(h_j^{\mathcal{T},k-1})$, $\beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T},k-1})^{-1}$ is a probability measure on $\mathcal{F}_\Theta \otimes \mathcal{F}_{B_j^{k-1}}$. Hence, $h_i^{\mathcal{T},k}(t_i) \in B_i^k$. \square

Claim 3. For $i \in N$ and $j \neq i$, $\sigma(h_i^{\mathcal{T},k}) \succ^* \sigma(h_j^{\mathcal{T},k-1})$.

Proof of Claim 3. To show that $\sigma(h_i^{\mathcal{T},k})$ is the coarsest σ -algebra that dominates $\sigma(h_j^{\mathcal{T},k-1})$, we first show that $\sigma(h_i^{\mathcal{T},k})$ is generated by sets of the form

$$\begin{aligned}\left\{ t_i \in T_i : (\text{Id}_\Theta, h_j^{\mathcal{T},k-1})^{-1}(E) \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i), \beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T},k-1})^{-1}(E) \geq p \right\} : \\ E \in \mathcal{F}_\Theta \otimes \mathcal{F}_{B_j^{\mathcal{T},k-1}}, p \in [0, 1].\end{aligned}\tag{A.4}$$

Notice that

$$\sigma(h_i^{\mathcal{T},k}) = \left\{ (h_i^{\mathcal{T},k})^{-1}(B) : B \in \mathcal{F}_{B_i^{\mathcal{T},k-1}} \otimes \mathcal{F}_{\Delta(\Theta \times B_j^{k-1}, \mathcal{S}_i^k)} \right\}.$$

Because every σ -algebra in \mathcal{S}_i^k is a sub- σ algebra of $\mathcal{F}_\Theta \otimes \mathcal{F}_{B_j^{k-1}}$, the σ -algebra $\mathcal{F}_{\Delta(\Theta \times B_j^{k-1}, \mathcal{S}_i^k)}$ is generated by sets of the form

$$\left\{ \mu_i^k \in \Delta(\Theta \times B_j^{k-1}, \mathcal{S}_i^k) : E \in \Sigma(\mu_i^k), \mu_i^k(E) \geq p \right\} : \quad E \in \mathcal{F}_\Theta \otimes \mathcal{F}_{B_j^{k-1}}, p \in [0, 1].$$

Using Lemma A.1, it follows that $\sigma(h_i^{\mathcal{T},k})$ is generated by the sets in $\sigma(h_j^{\mathcal{T},k-1})$ and sets of the form

$$\begin{aligned}\left\{ t_i \in T_i : (\text{Id}_\Theta, h_j^{\mathcal{T},k-1})^{-1}(E) \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i), \beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T},k-1})^{-1}(E) \geq p \right\} : \\ E \in \mathcal{F}_\Theta \otimes \mathcal{F}_{B_j^{k-1}}, p \in [0, 1].\end{aligned}\tag{A.5}$$

The next step is to show that each of the generating sets in $\sigma(h_i^{\mathcal{T},k-1})$ can be written as a set of the form (A.5), so that $\sigma(h_i^{\mathcal{T},k})$ is generated by sets of the form (A.5).

By the induction hypothesis (b'), $\sigma(h_i^{\mathcal{T},k-1})$ is generated by sets of the form

$$\begin{aligned}\left\{ t_i \in T_i : (\text{Id}_\Theta, h_j^{\mathcal{T},k-2})^{-1}(E) \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i), \beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T},k-2})^{-1}(E) \geq p \right\} : \\ E \in \mathcal{F}_\Theta \otimes \mathcal{F}_{B_j^{\mathcal{T},k-2}}, p \in [0, 1].\end{aligned}$$

Fix $E \in \mathcal{F}_\Theta \otimes \mathcal{F}_{B_j^{\mathcal{T},k-2}}$ and $p \in [0, 1]$. Then $E' := (\pi_{\Theta \times B_j^{\mathcal{T},k-1}}^{\Theta \times B_j^{\mathcal{T},k-1}})^{-1}(E) \in \mathcal{F}_\Theta \times \mathcal{F}_{B_j^{\mathcal{T},k-1}}$, and

$$\begin{aligned} \{t_i \in T_i : (\text{Id}_\Theta, h_j^{\mathcal{T},k-2})^{-1}(E) \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i), \beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T},k-2})^{-1}(E) \geq p\} = \\ \{t_i \in T_i : (\text{Id}_\Theta, h_j^{\mathcal{T},k-1})^{-1}(E') \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i), \beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T},k-1})^{-1}(E') \geq p\}. \end{aligned}$$

Consequently, $\sigma(h_i^{\mathcal{T},k})$ is generated by sets of the form (A.5). Rewriting (using a similar transformation as in (A.3)) gives that $\sigma(h_i^{\mathcal{T},k})$ is generated by sets of the form (A.4).

The proof that $\sigma(h_i^{\mathcal{T},k})$ is the coarsest σ -algebra that dominates $\sigma(h_j^{\mathcal{T},k-1})$ is now immediate. Because $\sigma(h_i^{\mathcal{T},k})$ is the coarsest σ -algebra on T_i that contains the sets (A.4) for $E \in \mathcal{F}_\Theta \otimes \mathcal{F}_{B_j^{\mathcal{T},k-1}}$ and $p \in [0, 1]$, it is the coarsest σ -algebra that includes the sets

$$\{t_i \in T_i : E \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i), \beta_i(t_i)(E) \geq p\}$$

for $E \in \mathcal{F}_\Theta \otimes \sigma(h_j^{\mathcal{T},k-1})$ and $p \in [0, 1]$. That is, $\sigma(h_i^{\mathcal{T},k}) \succ^* \sigma(h_j^{\mathcal{T},k-1})$. \square

Claim 4. For $i \in N$ and $\mathcal{F}_i \in \mathcal{S}_i$ such that $\mathcal{F}_i \subsetneq \sigma(h_i^{\mathcal{T},k})$, there is $m < k$ such that $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},m})$.

Proof of Claim 4. By the induction hypothesis (b'), the claim holds if $\mathcal{F}_i \subsetneq \sigma(h_i^{\mathcal{T},k-1})$, and it clearly holds if $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},k-1})$. So suppose $\sigma(h_i^{\mathcal{T},k-1}) \subsetneq \mathcal{F}_i \subsetneq \sigma(h_i^{\mathcal{T},k})$.

As $\mathcal{F}_i \supseteq \sigma(h_i^{\mathcal{T},0}) = \{T_i, \emptyset\}$, it follows from Condition 3 that there is $\mathcal{F}_j^1 \in \mathcal{S}_j$ such that $\mathcal{F}_i \succ \mathcal{F}_j^1$. As in the proof of Claim 1, applying Condition 3 repeatedly gives that $(\mathcal{F}_i, \mathcal{F}_j^1)$ is a mutual dominance pair, or \mathcal{F}_i and \mathcal{F}_j^1 are part of a finite chain of the form (A.1), or they are part of a cycle of the form (A.2).

If $(\mathcal{F}_i, \mathcal{F}_j^1)$ is a mutual-dominance pair, then an argument similar to the one in Claim 1 can be employed to show that $\mathcal{F}_i \supseteq \sigma(h_i^{\mathcal{T},k})$ (using that $\sigma(h_i^{\mathcal{T},k}) \succ^* \sigma(h_j^{\mathcal{T},k-1})$, by Claim 2), a contradiction. Similarly, if \mathcal{F}_i and \mathcal{F}_j^1 belong to a cycle of the form (A.2), then $\mathcal{F}_i \supseteq \sigma(h_i^{\mathcal{T},k})$, again giving a contradiction. Finally, suppose \mathcal{F}_i is part of a chain of the form (A.1). If $m \geq k$, we again obtain a contradiction. If $m < k$, then $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},m})$, a contradiction. So we cannot have that $\sigma(h_i^{\mathcal{T},k-1}) \subsetneq \mathcal{F}_i \subsetneq \sigma(h_i^{\mathcal{T},k})$. \square

This completes the induction. We thus have that for each $i \in N$ and $k = 0, 1, \dots$, $h_i^{\mathcal{T},k}$ is well-defined. Moreover, $\sigma(h_i^{\mathcal{T},k+1}) \succ^* \sigma(h_i^{\mathcal{T},k})$.

The last step is to show (c). Let $i \in N$. Clearly, $\mathcal{F}_i \supseteq \sigma(h_i^{\mathcal{T},0}) = \{T_i, \emptyset\}$ for every $\mathcal{F}_i \in \mathcal{S}_i$. Combining this with Claim 1, for every $k = 0, 1, \dots$, $\mathcal{F}_i \in \mathcal{S}_i$,

$$\mathcal{F}_i \subseteq \sigma(h_i^{\mathcal{T},k}) \quad \text{or} \quad \mathcal{F}_i \supseteq \sigma(h_i^{\mathcal{T},k}).$$

By Claim 4, if $\mathcal{F}_i \subsetneq \sigma(h_i^{\mathcal{T},k})$ for some $k = 0, 1, \dots$, then there is $m \leq k$ such that $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},m})$. \square

A.2 Proof of Proposition 3.1

We need to show that the structure $(T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$ satisfies Conditions 1–3. Conditions 1 and 2 hold trivially, as $\mathcal{S}_i = \{\mathcal{F}_i^{\mathcal{H}}\}$ is a singleton for each player $i \in N$. That Condition 3 holds follows from Lemma A.2 (with $\mathcal{S}_i := \{\mathcal{F}_\Theta \otimes \mathcal{F}_j^{\mathcal{H}}\}$). \square

A.3 Proof of Theorem 4.1

As an extended type space satisfies Condition 3, the result is immediate from Lemma A.4(a). \square

A.4 Proof of Theorem 4.2

We use the following auxiliary result:

Lemma A.5. *Fix $i \in N$ and let $(\mu_i^1, \mu_i^2, \dots) \in H_i$ be a belief hierarchy for i . Suppose there exist extended type spaces $\mathcal{T} = (T_n, \mathcal{S}_n^{\mathcal{T}}, \Sigma_n^{\mathcal{T}}, \beta_n^{\mathcal{T}})_{n \in N}$ and $\mathcal{Q} = (Q_n, \mathcal{S}_n^{\mathcal{Q}}, \Sigma_n^{\mathcal{Q}}, \beta_n^{\mathcal{Q}})_{n \in N}$ and types $t_i \in T_i$ and $q_i \in Q_i$ such that*

$$h_i^{\mathcal{T}}(t_i) = h_i^{\mathcal{Q}}(q_i) = (\mu_i^1, \mu_i^2, \dots).$$

Then for each $k = 0, 1, \dots$,

$$(a) \Sigma_i^{\mathcal{T}}(t_i) \supseteq \sigma(h_j^{\mathcal{T},k}) \text{ if and only if } \Sigma_i^{\mathcal{Q}}(q_i) \supseteq \sigma(h_j^{\mathcal{Q},k});$$

$$(b) \Sigma_i^{\mathcal{T}}(t_i) \subsetneq \sigma(h_j^{\mathcal{T},k}) \text{ if and only if } \Sigma_i^{\mathcal{Q}}(q_i) \subsetneq \sigma(h_j^{\mathcal{Q},k}).$$

Proof. We prove (a); the proof of (b) is similar and thus omitted. By (A.3),

$$\sigma(h_j^{\mathcal{T},k}) = \{(h_i^{\mathcal{T},k})^{-1}(B) : B \in \mathcal{F}_{B_i^k}\}.$$

Thus, if $\Sigma_i^{\mathcal{T}}(t_i) \supseteq \sigma(h_j^{\mathcal{T},k})$, then μ_i^k is defined on $\mathcal{F}_\Theta \otimes \mathcal{F}_{B_i^k}$. Consequently,

$$\Sigma_i^{\mathcal{Q}}(q_i) \supseteq \sigma(h_j^{\mathcal{Q},k}).$$

Of course, the converse holds by symmetry. \square

We can now prove Theorem 4.2. By Lemma A.4(c), for each extended type space $\mathcal{T} = (T_n, \mathcal{S}_n^{\mathcal{T}}, \Sigma_n^{\mathcal{T}}, \beta_n^{\mathcal{T}})_{n \in N}$, player $i \in N$ and type $t_i \in T_i$, either

$$\Sigma_i^{\mathcal{T}}(t_i) \supseteq \sigma(h_j^{\mathcal{T},k-1})$$

for every $k = 1, 2, \dots$, or there is $d = 1, 2, \dots$ such that

$$\Sigma_i^{\mathcal{T}}(t_i) = \sigma(h_j^{\mathcal{T}, d-1})$$

(so that $\Sigma_i^{\mathcal{T}}(t_i) \supseteq \sigma(h_j^{\mathcal{T}, k-1})$ for $k < d$) and

$$\Sigma_i^{\mathcal{T}}(t_i) \subsetneq \sigma(h_j^{\mathcal{T}, k-1})$$

for every $k > d$.

Fix a player $i \in N$ and a belief hierarchy $(\mu_i^1, \mu_i^2, \dots) \in H_i$. Thus, there is an extended type space $\mathcal{T} = (T_n, \mathcal{S}_n^{\mathcal{T}}, \Sigma_n^{\mathcal{T}}, \beta_n^{\mathcal{T}})_{n \in N}$ and type $t_i \in T_i$ such that $h_i^{\mathcal{T}}(t_i) = (\mu_i^1, \mu_i^2, \dots)$.

First suppose $\Sigma_i^{\mathcal{T}}(t_i) \supseteq \sigma(h_j^{\mathcal{T}, k-1})$ for each k . By Lemma A.5(a), $\Sigma_i^{\mathcal{Q}}(q_i) \supseteq \sigma(h_j^{\mathcal{Q}, k-1})$ for every extended type space $\mathcal{Q} = (Q_n, \mathcal{S}_n^{\mathcal{Q}}, \Sigma_n^{\mathcal{Q}}, \beta_n^{\mathcal{Q}})_{n \in N}$ and type $q_i \in Q_i$ such that $h_i^{\mathcal{Q}}(q_i) = (\mu_i^1, \mu_i^2, \dots)$. It follows that $(\mu_i^1, \mu_i^2, \dots)$ has depth $d = \infty$.

Next suppose that there is $d = 1, 2, \dots$ such that $\Sigma_i^{\mathcal{T}}(t_i) = \sigma(h_j^{\mathcal{T}, d-1})$ and $\Sigma_i^{\mathcal{T}}(t_i) \subsetneq \sigma(h_j^{\mathcal{T}, k-1})$ for all $k > d$. Again, by Lemma A.5, for any extended type space $\mathcal{Q} = (Q_n, \mathcal{S}_n^{\mathcal{Q}}, \Sigma_n^{\mathcal{Q}}, \beta_n^{\mathcal{Q}})_{n \in N}$ and type $q_i \in Q_i$ are such that $h_i^{\mathcal{Q}}(q_i) = (\mu_i^1, \mu_i^2, \dots)$, it holds that $\Sigma_i^{\mathcal{Q}}(q_i) \supseteq \sigma(h_j^{\mathcal{Q}, d-1})$ and $\Sigma_i^{\mathcal{Q}}(q_i) \subsetneq \sigma(h_j^{\mathcal{Q}, k-1})$ for $k > d$, so that by Lemma A.4(c), $\Sigma_i^{\mathcal{Q}}(q_i) = \sigma(h_j^{\mathcal{Q}, d-1})$.

It is easy to verify that for every extended type space $\mathcal{T} = (T_n, \mathcal{S}_n^{\mathcal{T}}, \Sigma_n^{\mathcal{T}}, \beta_n^{\mathcal{T}})_{n \in N}$ and type $t_i \in T_i$ such that $\Sigma_i^{\mathcal{T}}(t_i) = \sigma(h_j^{\mathcal{T}, d-1})$ and $\Sigma_i^{\mathcal{T}}(t_i) \subsetneq \sigma(h_j^{\mathcal{T}, k-1})$, $\beta_i(t_i) \circ (\text{Id}_{\Theta}, h_j^{\mathcal{T}, k-1})^{-1}$ is a probability measure on $\mathcal{F}_{\Theta} \otimes \mathcal{F}_{B_j^{\mathcal{T}, k-1}}$ for $k \leq d-1$, and on $\mathcal{F}_{\Theta} \otimes \mathcal{A}_j^{\mathcal{T}, k-1, d-1}$ for $k \geq d$, with $\mathcal{A}_j^{\mathcal{T}, k-1, d-1} \subsetneq \mathcal{F}_{B_j^{\mathcal{T}, k-1}}$. Hence, the depth of $(\mu_i^1, \mu_i^2, \dots)$ is d . \square

A.5 Proof of Proposition 4.3

We use some preliminary results. The first result states that a σ -algebra on a given type set dominates more σ -algebras (on the other type set) than another σ -algebra if and only if the former is a strict refinement of the latter.

Lemma A.6. *Fix a player $i \in N$ and let $\mathcal{F}_i, \mathcal{F}'_i \in \mathcal{S}_i$. Then, $\mathcal{F}'_i \supseteq \mathcal{F}_i$ if and only if there is $\mathcal{F}_j \in \mathcal{S}_j$ such that $\mathcal{F}'_i \succ^* \mathcal{F}_j$ and $\mathcal{F}_i \not\succeq \mathcal{F}_j$. Moreover, for all $\mathcal{F}_j \in \mathcal{S}_j$, if $\mathcal{F}_i \succ \mathcal{F}_j$, then $\mathcal{F}'_i \succ \mathcal{F}_j$.*

Proof. Suppose $\mathcal{F}_i \subsetneq \mathcal{F}'_i$. To prove the first claim, note that $\mathcal{F}'_i \supseteq \{T_i, \emptyset\}$, so by Condition 3, there is \mathcal{F}_j such that $\mathcal{F}'_i \succ^* \mathcal{F}_j$, or $(\mathcal{F}'_i, \mathcal{F}_j)$ is a mutual-dominance pair. If \mathcal{F}'_i is the coarsest σ -algebra that dominates \mathcal{F}_j , then in particular \mathcal{F}_i does not dominate \mathcal{F}_j . So suppose $(\mathcal{F}'_i, \mathcal{F}_j)$ is a mutual-dominance pair. We want to show that $\mathcal{F}_i \not\succeq \mathcal{F}_j$. As $\mathcal{F}_j \succ \mathcal{F}'_i$ and $\mathcal{F}'_i \supseteq \mathcal{F}_i$, we have that $\mathcal{F}_j \succ \mathcal{F}_i$. Consequently, if $\mathcal{F}_i \succ \mathcal{F}_j$, $(\mathcal{F}_i, \mathcal{F}_j)$ is a mutual-dominance pair, violating Condition 2.

Conversely, suppose that there is $\mathcal{F}_j \in \mathcal{S}_j$ such that $\mathcal{F}'_i \succ \mathcal{F}_j$ and $\mathcal{F}_i \not\prec \mathcal{F}_j$. By Condition 1, $\mathcal{F}_i \subseteq \mathcal{F}'_i$ or $\mathcal{F}'_i \subseteq \mathcal{F}_i$. It is immediate that we cannot have $\mathcal{F}_i = \mathcal{F}'_i$ or $\mathcal{F}_i \supsetneq \mathcal{F}'_i$. Conclude that $\mathcal{F}_i \subsetneq \mathcal{F}'_i$.

The proof of the last claim is immediate. \square

To state the next result, we need some more notation. For each player i , let

$$\tilde{\mathcal{S}}_i := \{\sigma(h_i^{\mathcal{T},m}) : m = 0, 1, \dots\},$$

where the notation $\sigma(h_i^{\mathcal{T},m})$ was introduced in Appendix A.1.2; it follows from Theorem 4.1 that $\sigma(h_i^{\mathcal{T},m})$ is well-defined for all m . We denote the σ -algebra on T_i that is generated by the subsets in $\sigma(h_i^{\mathcal{T},m})$, $m = 0, 1, \dots$, by $\sigma(\tilde{\mathcal{S}}_i)$.

Lemma A.7. *Fix $i \in N$. Then $\sigma(\tilde{\mathcal{S}}_i)$ is the coarsest σ -algebra that dominates $\sigma(\tilde{\mathcal{S}}_j)$, and, conversely, $\sigma(\tilde{\mathcal{S}}_j)$ is the coarsest σ -algebra that dominates $\sigma(\tilde{\mathcal{S}}_i)$.*

Proof. We want to show that $\sigma(\tilde{\mathcal{S}}_i)$ is generated by sets of the form

$$\{t_i \in T_i : E \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i), \beta_i(t_i)(E) \geq p\} : \quad E \in \mathcal{F}_\Theta \otimes \sigma(\tilde{\mathcal{S}}_j), p \in [0, 1].$$

By Lemma 4.5 of Heifetz and Samet (1998b), it suffices to show that $\sigma(\tilde{\mathcal{S}}_i)$ is generated by sets of the form

$$\{t_i \in T_i : E \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i), \beta_i(t_i)(E) \geq p\} : \quad E \in \bigcup_{m=0}^{\infty} (\mathcal{F}_\Theta \otimes \sigma(h_j^{\mathcal{T},m})), p \in [0, 1]. \quad (\text{A.6})$$

To show this, note that $\sigma(\tilde{\mathcal{S}}_i)$ is the coarsest σ -algebra that contains the sets

$$\{t_i \in T_i : h_i^{\mathcal{T},m}(t_i) \in B\} : \quad B \in \mathcal{F}_{B_i^m}, m = 0, 1, \dots$$

It follows from Lemma A.1 and argument similar to the one in the proof of Claim 3 of Lemma A.4 that $\sigma(\tilde{\mathcal{S}}_i)$ is the coarsest σ -algebra that contains the sets

$$\{t_i \in T_i : (\text{Id}_\Theta, h_j^{\mathcal{T},m})^{-1}(E) \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i), \beta_i(t_i) \circ (\text{Id}_\Theta, h_j^{\mathcal{T},m})^{-1}(E) \geq p\} : \\ E \in \mathcal{F}_\Theta \otimes \mathcal{F}_{B_j^{\mathcal{T},m}}, p \in [0, 1].$$

Rewriting gives that $\sigma(\tilde{\mathcal{S}}_i)$ is generated by sets of the form (A.6). \square

We are now ready to prove Proposition 4.3. It is immediate that $d_i(t_i) < d_i(t'_i)$ implies $\Sigma_i(t_i) \subsetneq \Sigma_i(t'_i)$.

To prove the converse, suppose $\Sigma_i(t_i) \subsetneq \Sigma_i(t'_i)$. For ease of notation, define $\mathcal{F}_j := \Sigma_i(t_i)$ and $\mathcal{F}'_j := \Sigma_i(t'_i)$. We want to show that $d_i(t_i) < d_i(t'_i)$.

By Lemma A.6, there is $\mathcal{F}_i \in \mathcal{S}_i$ such that $\mathcal{F}'_j \succ \mathcal{F}_i$ and $\mathcal{F}_j \not\succeq \mathcal{F}_i$. By Lemma A.4(c), either $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},m})$ for some m , or $\mathcal{F}_i \supseteq \sigma(h_i^{\mathcal{T},m})$ for all m .

First suppose $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},m})$ for some m . By Lemma A.4(b), $\mathcal{F}'_j \succ \mathcal{F}_i$ implies $\mathcal{F}'_j \supseteq \sigma(h_j^{\mathcal{T},m+1})$, while $\mathcal{F}_j \not\succeq \mathcal{F}_i$ implies $\mathcal{F}_i \subsetneq \sigma(h_j^{\mathcal{T},m+1})$. Consequently, $d_i(t'_i) \geq m + 1$, and $d_i(t_i) < m + 1$.

Next suppose $\mathcal{F}_i \supseteq \sigma(h_i^{\mathcal{T},m})$ for all m . As $\mathcal{F}'_j \succ \mathcal{F}_i$, it follows from Lemma A.4(b) that $\mathcal{F}'_j \supseteq \sigma(h_j^{\mathcal{T},m})$ for all m , so that $d_i(t'_i) = \infty$. Suppose by contradiction that $d_i(t_i) = \infty$, i.e., $\mathcal{F}_j \supseteq \sigma(h_j^{\mathcal{T},m})$ for all m .

As $\mathcal{F}_j \supseteq \sigma(h_j^{\mathcal{T},m})$ for all m , it follows from Lemma A.4(b) that $\mathcal{F}_j \succ \sigma(h_i^{\mathcal{T},m})$ for all m . By Lemma 4.5 of Heifetz and Samet (1998b), we have that $\mathcal{F}_j \succ \sigma(\tilde{\mathcal{S}}_i)$, so that by Lemma A.7, $\mathcal{F}_j \supseteq \sigma(\tilde{\mathcal{S}}_j)$. As $\mathcal{F}_j \not\succeq \mathcal{F}_i$, $\mathcal{F}_i \not\supseteq \sigma(\tilde{\mathcal{S}}_i) \supseteq \{T_i, \emptyset\}$.

Noting that $(\sigma(\tilde{\mathcal{S}}_i), \sigma(\tilde{\mathcal{S}}_j))$ is a mutual-dominance pair (by Lemma A.7), it follows from Lemma A.3 that there is no $\tilde{\mathcal{F}}_j \in \mathcal{S}_j$ such that $\mathcal{F}_i \succ^* \tilde{\mathcal{F}}_j$ and $\tilde{\mathcal{F}}_j \not\succeq \mathcal{F}_i$. By Condition 3, therefore, there is $\tilde{\mathcal{F}}_j \in \mathcal{S}_j$ such that $(\mathcal{F}_i, \tilde{\mathcal{F}}_j)$ is a mutual-dominance pair.

Similarly, because $\mathcal{F}'_j \supseteq \mathcal{F}_j \supseteq \sigma(\tilde{\mathcal{S}}_j)$, there is $\tilde{\mathcal{F}}_i \in \mathcal{S}_i$ such that $(\tilde{\mathcal{F}}_i, \mathcal{F}'_j)$ is a mutual-dominance pair. By Condition 2, $\tilde{\mathcal{F}}_i = \mathcal{F}_i$ and $\tilde{\mathcal{F}}_j = \mathcal{F}'_j$, i.e., $(\mathcal{F}_i, \mathcal{F}'_j)$ is a mutual-dominance pair.

As $(\mathcal{F}_i, \mathcal{F}'_j)$ is a mutual-dominance pair, Condition 2 implies that $\mathcal{F}_j \neq \sigma(\tilde{\mathcal{S}}_j)$. Thus, $\mathcal{F}_j \not\supseteq \sigma(\tilde{\mathcal{S}}_j)$. Then, by the same argument as above, there is $\tilde{\mathcal{F}}_i \in \mathcal{S}_i$ such that $(\mathcal{F}_j, \tilde{\mathcal{F}}_i)$ is a mutual-dominance pair. By Condition 2, $\tilde{\mathcal{F}}_i = \mathcal{F}_i$. But this contradicts that $\mathcal{F}_j \not\succeq \mathcal{F}_i$. \square

A.6 Proof of Proposition 4.4

As the extended type space \mathcal{T} is derived from a Harsanyi type space, $\mathcal{S}_i = \{\mathcal{F}_i^{\mathcal{H}}\}$. As before, we write $\sigma(h_i^{\mathcal{T},k})$ for the σ -algebra $\{(h_i^{\mathcal{T},k})^{-1}(B) : B \in \mathcal{F}_{B^{\mathcal{T},k}}\}$.

We use the following preliminary result. We again prove it for a more general class of structures than extended type spaces.

Lemma A.8. *Let $\mathcal{T} = (T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$ be a structure that satisfies Condition 3, and fix $i \in N$. Let $\mathcal{F}_i \in \mathcal{S}_i$ and $\mathcal{F}_j \in \mathcal{S}_j$. If $\mathcal{F}_i \succ \mathcal{F}_j$ and $\mathcal{F}_j \succ \mathcal{F}_i$, then*

$$\mathcal{F}_i \supseteq \sigma(h_i^{\mathcal{T},m})$$

for all $m = 0, 1, \dots$, and similarly for \mathcal{F}_j .

Proof. As $\mathcal{F}_j \supseteq \{T_j, \emptyset\} = \sigma(h_j^{\mathcal{T},0})$, and $\mathcal{F}_i \succ \mathcal{F}_j$, it follows from Lemma A.4(b) that $\mathcal{F}_i \supseteq \sigma(h_i^{\mathcal{T},1})$. By a similar argument, $\mathcal{F}_j \supseteq \sigma(h_j^{\mathcal{T},1})$. Applying Lemma A.4(b) again yields

that $\mathcal{F}_i \supseteq \sigma(h_i^{T,2})$. Repeating this argument gives the desired result. \square

The result is now immediate. By Lemma A.2, $\mathcal{F}_i^{\mathcal{H}} \succ \mathcal{F}_j^{\mathcal{H}}$ for each $i \in N$ and $j \neq i$, so that by Lemma A.8, for each type t_i , $\Sigma_i(t_i) = \mathcal{F}_j^{\mathcal{H}} \supseteq \sigma(h_j^{T,k})$ for all k , that is, t_i has an infinite depth of reasoning. \square

A.7 Proof of Lemma 4.5

We first state some preliminary results. The first result says that σ -algebras that have the same finite rank coincide.

Lemma A.9. *Let $i \in N$ and $\mathcal{F}_i, \mathcal{F}'_i \in \mathcal{S}_i$. If $\mathcal{F}_i, \mathcal{F}'_i \in \mathcal{R}_i^k$ for $k < \infty$, then $\mathcal{F}_i = \mathcal{F}'_i$.*

Proof. The proof is by induction. Fix $i \in N$. Suppose $\mathcal{F}_i, \mathcal{F}'_i \in \mathcal{R}_i^1$, and suppose by contradiction that $\mathcal{F}_i \subsetneq \mathcal{F}'_i$. By Lemma A.6, there is $\mathcal{F}_j \in \mathcal{S}_j$ such that $\mathcal{F}'_i \succ \mathcal{F}_j$, contradicting that $\mathcal{F}'_i \in \mathcal{R}_i^1$. Next suppose that $\mathcal{F}_i \in \mathcal{R}_i^1$ and $\mathcal{F}'_i \in \mathcal{R}_i^m$ for $m > 1$. Then, by Lemma A.6, $\mathcal{F}'_i \supsetneq \mathcal{F}_i$.

For $k > 1$, suppose that for each $i \in N$ and $\ell \leq k - 1$, $\mathcal{F}_i = \mathcal{F}'_i$ whenever $\mathcal{F}_i, \mathcal{F}'_i \in \mathcal{R}_i^\ell$. Also suppose that if $\mathcal{F}_i \in \mathcal{R}_i^\ell$ for $\ell \leq k - 1$ and $\mathcal{F}'_i \in \mathcal{R}_i^m$ for $m > \ell$, then $\mathcal{F}'_i \supsetneq \mathcal{F}_i$.

Let $i \in N$ and suppose $\mathcal{F}_i, \mathcal{F}'_i \in \mathcal{R}_i^k$. Suppose by contradiction that $\mathcal{F}_i \subsetneq \mathcal{F}'_i$. Then by Lemma A.6, there is \mathcal{F}_j such that $\mathcal{F}'_i \succ \mathcal{F}_j$ and $\mathcal{F}_i \not\succeq \mathcal{F}_j$. By the induction hypothesis, $\mathcal{F}_j \notin \bigcup_{\ell \leq k-1} \mathcal{R}_j^\ell$, contradicting that $\mathcal{F}'_i \in \mathcal{R}_i^k$.

To complete the induction, suppose that $\mathcal{F}_i \in \mathcal{R}_i^k$ and $\mathcal{F}'_i \in \mathcal{R}_i^m$ for $m > k$. Then, by Lemma A.6, $\mathcal{F}'_i \supsetneq \mathcal{F}_i$. \square

Lemma A.10. *Let $i \in N$ and $\mathcal{F}_i \in \mathcal{S}_i$. If there is $\mathcal{F}_j \in \mathcal{S}_j$ such that $(\mathcal{F}_i, \mathcal{F}_j)$ is a mutual-dominance pair, then $\mathcal{F}_i \in \mathcal{R}_i^\infty$.*

Proof. As $\mathcal{F}_i \succ \mathcal{F}_j$, $\mathcal{F}_i \notin \mathcal{R}_i^1$. Similarly, $\mathcal{F}_j \notin \mathcal{R}_j^1$. Because $\mathcal{F}_i \succ \mathcal{F}_j$ and $\mathcal{F}_j \notin \mathcal{R}_j^1$, it follows from Lemma A.9 that $\mathcal{F}_i \notin \mathcal{R}_i^2$. Likewise, $\mathcal{F}_j \notin \mathcal{R}_j^2$. Repeating this argument gives that for all $k \geq 1$, $\mathcal{F}_j \in \mathcal{S}_j \setminus \bigcup_{\ell \geq k} \mathcal{R}_j^\ell$. Therefore, $\mathcal{F}_i \in \mathcal{R}_i^\infty$. \square

We are now ready to prove Lemma 4.5. Let $i \in N$. A straightforward argument by contradiction establishes that $\mathcal{R}_i^k \cap \mathcal{R}_i^\ell = \emptyset$ whenever $\ell \neq k$. It remains to show that $\mathcal{S}_i = \mathcal{R}_i^\infty \cup \bigcup_{\ell=1}^\infty \mathcal{R}_i^\ell$. Clearly, $\mathcal{S}_i \subseteq \mathcal{R}_i^\infty \cup \bigcup_{\ell=1}^\infty \mathcal{R}_i^\ell$. It remains to show the reverse inclusion.

For any $\mathcal{F}_i \in \mathcal{S}_i$, there is m such that $\mathcal{F}_i = \sigma(h_i^{T,m}) \subsetneq \sigma(h_i^{T,m+1})$, or $\mathcal{F}_i \supseteq \sigma(h_i^{T,m})$ for all m , by Lemma A.4(c).

For $i \in N$ and $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},0}) \subsetneq \sigma(h_i^{\mathcal{T},1})$, there is no $\mathcal{F}_j \in \mathcal{S}_j$ such that $\mathcal{F}_i \succ \mathcal{F}_j$ (by Lemma A.4(b)), so that $\mathcal{F}_i \in \mathcal{R}_i^1$.

For $k > 1$, suppose that for each $i \in N$, $\ell \leq k-1$, and $\mathcal{F}_i \in \mathcal{S}_i$, $\mathcal{F}_i \subseteq \sigma(h_i^{\mathcal{T},\ell-1}) \subsetneq \sigma(h_i^{\mathcal{T},\ell})$ implies that $\mathcal{F}_i \in \mathcal{R}_i^z$ for some $z \leq \ell$.

Let $i \in N$ and $\mathcal{F}_i \in \mathcal{S}_i$. Suppose $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},k-1}) \subsetneq \sigma(h_i^{\mathcal{T},k})$, and suppose $\sigma(h_i^{\mathcal{T},k-1}) \supsetneq \sigma(h_i^{\mathcal{T},0})$. Then by Condition 3 and Lemma A.8, there is $\mathcal{F}_j \in \mathcal{S}_j$ such that $\mathcal{F}_i \succ^* \mathcal{F}_j$ and $\mathcal{F}_j \not\succeq \mathcal{F}_i$.

Define $D_j(\mathcal{F}_i) := \{\mathcal{F}_j \in \mathcal{S}_j : \mathcal{F}_i \succ \mathcal{F}_j\}$, and note that $D_j(\mathcal{F}_i)$ is nonempty. Also, by Lemma A.4(b), there is no $\mathcal{F}_j \in D_j(\mathcal{F}_i)$ such that $\mathcal{F}_j \supseteq \sigma(h_j^{\mathcal{T},k-1})$. By Lemma A.4(c), therefore, there is $\mathcal{F}_j \in D_j(\mathcal{F}_i)$ such that $\mathcal{F}'_j \subseteq \mathcal{F}_j$ for all $\mathcal{F}'_j \in D_j(\mathcal{F}_i)$. That is, $D_j(\mathcal{F}_i)$ has a finest σ -algebra. As $\mathcal{F}_j \subsetneq \sigma(h_j^{\mathcal{T},k-1})$, it follows from Lemma A.4 that $\mathcal{F}_j = \sigma(h_j^{\mathcal{T},\ell}) \subsetneq \sigma(h_j^{\mathcal{T},\ell+1})$ for some $\ell \leq k-2$. By the induction hypothesis, therefore, $\mathcal{F}_j \in \mathcal{R}_j^z$ for some $z \leq \ell+1 \leq k-1$. As there is no $\mathcal{F}'_j \in \mathcal{S}_j$ such that $\mathcal{F}'_j \supsetneq \mathcal{F}_j$ and $\mathcal{F}_i \succ \mathcal{F}_j$, it follows that $\mathcal{F}_i \in \mathcal{R}_i^{z'}$ for some $z' \leq \ell+2 \leq k$.

Hence, if $\mathcal{F}_i \in \mathcal{S}_i$ is such that $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},m}) \subsetneq \sigma(h_i^{\mathcal{T},m+1})$ for some m , then there is $z \leq m$ such that $\mathcal{F}_i \in \mathcal{R}_i^z$.

So suppose $\mathcal{F}_i \in \mathcal{S}_i$ is such that $\mathcal{F}_i \supseteq \sigma(h_i^{\mathcal{T},m})$ for all m . Recall the notation $\tilde{\mathcal{S}}_n$ and $\sigma(\tilde{\mathcal{S}}_n)$, where $n \in N$, from Appendix A.5. We need to consider two cases:

Case 1: $\mathcal{F}_i \supsetneq \sigma(h_i^{\mathcal{T},m})$ for all m Then, $\mathcal{F}_i \supseteq \sigma(\tilde{\mathcal{S}}_i)$, and $\sigma(\tilde{\mathcal{S}}_i) \supsetneq \sigma(h_i^{\mathcal{T},k})$ for all k . Hence, by Lemma A.3, there is no $\mathcal{F}_j \in \mathcal{S}_j$ such that $\mathcal{F}_i \succ^* \mathcal{F}_j$ and $\mathcal{F}_j \not\succeq \mathcal{F}_i$. By Condition 3, therefore, there is $\mathcal{F}_j \in \mathcal{S}_j$ such that $(\mathcal{F}_i, \mathcal{F}_j)$ is a mutual-dominance pair. By Lemma A.10, $\mathcal{F}_i \in \mathcal{R}_i^\infty$.

Case 2: There is $k < \infty$ such that $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},k})$ Without loss of generality, assume that $\mathcal{F}_i \supsetneq \sigma(h_i^{\mathcal{T},k-1})$ if $k \geq 1$. The first step is to show that $\sigma(h_j^{\mathcal{T},\ell}) = \sigma(h_j^{\mathcal{T},k+1})$ for all $\ell \geq k+1$. Suppose not. That is, suppose there is $\ell > k+1$ such that $\sigma(h_j^{\mathcal{T},\ell}) \supsetneq \sigma(h_j^{\mathcal{T},k+1})$. By Lemma A.4(b), $\sigma(h_j^{\mathcal{T},\ell})$ is the coarsest σ -algebra that dominates $\sigma(h_i^{\mathcal{T},\ell-1})$, and $\sigma(h_j^{\mathcal{T},k+1})$ is the coarsest σ -algebra that dominates $\sigma(h_i^{\mathcal{T},k})$. But then $\sigma(h_i^{\mathcal{T},\ell-1}) = \sigma(h_i^{\mathcal{T},k})$ implies that $\sigma(h_j^{\mathcal{T},\ell}) = \sigma(h_j^{\mathcal{T},k+1})$, a contradiction.

Next, suppose that for each $\mathcal{F}_j \in \mathcal{S}_j$, $\mathcal{F}_j \subsetneq \sigma(h_j^{\mathcal{T},k+1})$. Because \mathcal{S}_j is nonempty, and because $\mathcal{F}_j \subsetneq \sigma(h_j^{\mathcal{T},k+1})$ for all $\mathcal{F}_j \in \mathcal{S}_j$, it follows from Lemma A.4(c) that there is $\mathcal{F}_j^* \in \mathcal{S}_j$ such that $\mathcal{F}_j \subsetneq \mathcal{F}_j^*$ for all $\mathcal{F}_j \in \mathcal{S}_j$. Lemma A.4(c) additionally implies that $\mathcal{F}_j^* = \sigma(h_j^{\mathcal{T},\ell}) \supsetneq \sigma(h_j^{\mathcal{T},\ell+1})$ for some $\ell < k+1$. Hence, $\mathcal{F}_j^* \in \mathcal{R}_j^z$ for some $z \leq k+1$. Since there is no $\mathcal{F}_j \in \mathcal{S}_j$ such that $\mathcal{F}_j \supseteq \mathcal{F}_j^*$, $\mathcal{F}_i \in \mathcal{R}_i^y$ for $y \leq k+2$.

To complete the proof, suppose that there is $\mathcal{F}_j \in \mathcal{S}_j$ such that $\mathcal{F}_j \supseteq \sigma(h_j^{\mathcal{T},k+1})$. First suppose $\mathcal{F}_j = \sigma(h_j^{\mathcal{T},k+1})$; note that by Lemma A.3 and Condition 2, there is no $\mathcal{F}'_j \in \mathcal{S}_j$ such that $\mathcal{F}'_j \supsetneq \mathcal{F}_j$. It is immediate that $(\mathcal{F}_i, \mathcal{F}_j)$ is a mutual-dominance pair. It then follows from

Lemma A.10 that $\mathcal{F}_i \in \mathcal{R}_i^\infty$.

Finally, suppose $\mathcal{F}_j \supsetneq \sigma(h_j^{\mathcal{T},k+1})$. By Condition 3 and Lemma A.3, there is $\mathcal{F}_j \in \mathcal{S}_j$ such that $(\mathcal{F}_i, \mathcal{F}_j)$ is a mutual-dominance pair. So, by Lemma A.10, $\mathcal{F}_i \in \mathcal{R}_i^\infty$. \square

A.8 Proof of Proposition 4.6

Let $i \in N$ and $t_i \in T_i$. By Lemma A.4(c), there is k such that $\Sigma_i(t_i) = \sigma(h_j^{\mathcal{T},k}) \subsetneq \sigma(h_j^{\mathcal{T},k+1})$, or $\Sigma_i(t_i) \supseteq \sigma(h_j^{\mathcal{T},m})$ for all m . By the proof of Lemma 4.5, if $\Sigma_i(t_i) = \sigma(h_j^{\mathcal{T},k}) \subsetneq \sigma(h_j^{\mathcal{T},k+1})$, then $\Sigma_i(t_i) \in \mathcal{R}_j^\ell$ for $\ell \leq k$, so that $d_i(t_i) = k \geq \ell = r_i(t_i)$.

If $\Sigma_i(t_i) \supseteq \sigma(h_j^{\mathcal{T},m})$ for all m , then, again by the proof of Lemma 4.5, $\Sigma_i(t_i) \in \mathcal{R}_j^k$ for $k < \infty$, or $\Sigma_i(t_i) \in \mathcal{R}_j^\infty$. That is, $d_i(t_i) = \infty \geq r_i(t_i)$. \square

A.9 Proof of Theorem 4.7

By Lemma A.4(c), for any $i \in N$, $\mathcal{F}_i \in \mathcal{S}_i$, $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},m}) \subsetneq \sigma(h_i^{\mathcal{T},m+1})$ for some m , or $\mathcal{F}_i \supseteq \sigma(h_i^{\mathcal{T},m})$ for all m .

For $i \in N$ and $\mathcal{F}_i \in \mathcal{S}_i$, suppose $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},0}) \subsetneq \sigma(h_i^{\mathcal{T},1})$. That is, \mathcal{F}_i corresponds to depth 1. By Lemma A.4, there is no $\mathcal{F}_j \in \mathcal{S}_j$ such that $\mathcal{F}_i \succ \mathcal{F}_j$, so $\mathcal{F}_i \in \mathcal{R}_i^1$.

For $k > 1$, suppose that for all $i \in N$, $\mathcal{F}_i \in \mathcal{S}_i$, and $\ell \leq k - 1$, it holds that $\mathcal{F}_i \in \mathcal{R}_i^\ell$ whenever $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},\ell-1}) \subsetneq \sigma(h_i^{\mathcal{T},\ell})$. Let $i \in N$, and suppose $\mathcal{F}_i = \sigma(h_i^{\mathcal{T},k-1}) \subsetneq \sigma(h_i^{\mathcal{T},k})$. By Lemma A.4, $\mathcal{F}_i \succ^* \sigma(h_j^{\mathcal{T},k-2})$, and there is no $\mathcal{F}_j \in \mathcal{S}_j$ such that $\mathcal{F}_j \supsetneq \sigma(h_j^{\mathcal{T},k-2})$ and $\mathcal{F}_i \succ \mathcal{F}_j$.

The first step is to show that $\sigma(h_j^{\mathcal{T},k-2}) \subsetneq \sigma(h_i^{\mathcal{T},k-1})$. Suppose by contradiction that $\sigma(h_j^{\mathcal{T},k-2}) = \sigma(h_j^{\mathcal{T},k-1})$. As $\sigma(h_i^{\mathcal{T},k-1}) \succ^* \sigma(h_j^{\mathcal{T},k-2})$ and $\sigma(h_i^{\mathcal{T},k}) \succ^* \sigma(h_j^{\mathcal{T},k-1})$, this implies that $\sigma(h_i^{\mathcal{T},k-1}) = \sigma(h_i^{\mathcal{T},k})$, a contradiction.

Hence, $\sigma(h_j^{\mathcal{T},k-2}) \subsetneq \sigma(h_i^{\mathcal{T},k-1})$, and by Condition 4, $\sigma(h_j^{\mathcal{T},k-2}) \in \mathcal{F}_j$. By the induction hypothesis, therefore, $\sigma(h_j^{\mathcal{T},k-2}) \in \mathcal{R}_j^{k-1}$, and $\mathcal{F}_i \in \mathcal{R}_i^k$.

Thus, we have shown that for any $i \in N$ and $t_i \in T_i$, if $\Sigma_i(t_i) = \sigma(h_j^{\mathcal{T},k-1}) \subsetneq \sigma(h_j^{\mathcal{T},k})$, then $d_i(t_i) = k = r_i(t_i)$.

So suppose $\mathcal{F}_i \supseteq \sigma(h_i^{\mathcal{T},m})$ for all m , that is, \mathcal{F}_i corresponds to an infinite depth of reasoning. Then $\mathcal{F}_i \supseteq \sigma(\tilde{\mathcal{S}}_i)$, where $\sigma(\tilde{\mathcal{S}}_i)$ was defined in Appendix A.5. If we show that there is $\mathcal{F}_j \in \mathcal{S}_j$ such that $(\mathcal{F}_i, \mathcal{F}_j)$ is a mutual-dominance pair, then it follows from Lemma A.10 that $\mathcal{F}_i \in \mathcal{R}_i^\infty$.

If $\mathcal{F}_i = \sigma(\tilde{\mathcal{S}}_i)$, then $(\mathcal{F}_i, \sigma(\tilde{\mathcal{S}}_j))$ is a mutual-dominance pair, by Lemma A.7. By Condition 4, $\sigma(\tilde{\mathcal{S}}_j) \in \mathcal{F}_j$. If $\mathcal{F}_i \supsetneq \sigma(\tilde{\mathcal{S}}_i) \supseteq \{T_i, \emptyset\}$, then it follows from Condition 3 and Lemma A.3 that $(\mathcal{F}_i, \mathcal{F}_j)$ is a mutual-dominance pair for some $\mathcal{F}_j \in \mathcal{S}_j$.

Hence, for any $i \in N$ and $t_i \in T_i$, if $\Sigma_i(t_i) \supseteq \sigma(h_j^{\mathcal{T},m})$ for all m , then $d_i(t_i) = \infty = r_i(t_i)$. \square

Appendix B Proofs for Section 5

B.1 Proof of Theorem 5.1

Because the depth $d_i(t_i)$ of type t_i is $d < \infty$, it follows from Lemma A.4(c) that

$$\Sigma_i(t_i) = \sigma(h_j^{\mathcal{T}, d-1}) = \left\{ \{t_j \in T_j : h_j^{\mathcal{T}, d-1}(t_j) \in B_{d-1}\} : B_{d-1} \in \mathcal{F}_{B_j^{\mathcal{T}, d-1}} \right\}. \quad (\text{B.1})$$

Thus, for any $B \subseteq H_j$, if there is $B_{d-1} \in \mathcal{F}_{B_j^{\mathcal{T}, d-1}}$ such that

$$\{t_j \in T_j : h_j^{\mathcal{T}}(t_j) \in B\} = \{t_j \in T_j : h_j^{\mathcal{T}, d-1}(t_j) \in B_{d-1}\},$$

we have that

$$\{t_j \in T_j : h_j^{\mathcal{T}}(t_j) \in B\} \in \Sigma_i(t_i).$$

Conversely, suppose that

$$\{t_j \in T_j : h_j^{\mathcal{T}}(t_j) \in B\} \in \Sigma_i(t_i)$$

for some $B \subseteq H_j$. By (B.1), this can hold only if there is $B_{d-1} \in \mathcal{F}_{B_j^{\mathcal{T}, d-1}}$ such that

$$\{t_j \in T_j : h_j^{\mathcal{T}}(t_j) \in B\} = \{t_j \in T_j : h_j^{\mathcal{T}, d-1}(t_j) \in B_{d-1}\},$$

that is, if the event that player j has a belief hierarchy in B coincides with the event that he has a $(d-1)$ th-order belief hierarchy in B_{d-1} . \square

B.2 Proof of Proposition 5.2

We first present some properties of the belief operator that we will need to prove Proposition 5.2; the proofs closely mirror the proofs of Monderer and Samet (1989) and are thus relegated to the online appendix.

Lemma B.1. (Positive introspection (1)) *For each $G \subseteq \Theta \times T$, $i \in N$, and $p \in [0, 1]$, $B_{i,p}(G) \subseteq B_{i,p}(B_{i,p}(G))$.*

Lemma B.2. (Positive introspection (2)) *For each $G \subseteq \Theta \times T$, $i \in N$, and $p \in (0, 1]$, $B_{i,p}(B_{i,p}(G)) \subseteq B_{i,p}(G)$.*

Lemma B.3. (Monotonicity) *Let $G, G' \subseteq \Theta \times T$, $i \in N$, and $p \in (0, 1]$. If $G \subseteq G'$ and $G'_{t_i} \in \mathcal{F}_{\Theta \otimes \Sigma_i(t_i)}$ for all t_i such that $(\theta, t_i, t_j) \in B_{i,p}(G)$ for some θ, t_j , then $B_{i,p}(G) \subseteq B_{i,p}(G')$.*

Lemma B.4. *For each $G \subseteq \Theta \times T$ and $p \in (0, 1]$, $(B_p^k(G))_{k \in \mathbb{N}}$ is a decreasing sequence of events.*

Lemma B.5. *Suppose $(E^\ell)_{\ell \in \mathbb{N}}$ is a decreasing sequence of events, where $E^\ell \subseteq \Theta \times T$ for $\ell \in \mathbb{N}$. Then, for each $i \in N$ and $p \in [0, 1]$, it holds that $\bigcap_{\ell \in \mathbb{N}} B_{i,p}(E^\ell) \subseteq B_{i,p}(\bigcap_{\ell} E^\ell)$.*

The final preliminary result says that $C_p(G)$ is an evident p -belief event for any event G :

Lemma B.6. *For each $G \subseteq \Theta \times T$ and $p \in [0, 1]$, it holds that $C_p(G) \subseteq B_{i,p}(C_p(G))$ for all $i \in N$.*

We are now ready to prove Proposition 5.2. First suppose that E is an evident p -belief event. We show by induction that $E \subseteq B_p^k(F)$ for all $k = 1, 2, \dots$, so that $(\theta, t) \in \bigcap_k B_p^k(F) = C_p(F)$ for any $(\theta, t) \in E$. To show this, note that $E \subseteq B_p^1(F)$ by assumption. For $k > 1$, suppose that $E \subseteq B_p^{k-1}(F)$. By assumption, $[B_p^{k-1}(F)]_{t'_i} \in \mathcal{F}_\Theta \otimes \Sigma_i(t'_i)$ for any t'_i such that $(\theta', t'_i, t'_j) \in B_{i,p}(E)$ for some θ', t'_j . Using Lemma B.3, it follows from the induction hypothesis that $B_{i,p}(E) \subseteq B_{i,p}(B_p^{k-1}(F))$. As E is an evident p -belief event, we conclude that $E \subseteq B_p^k(F)$.

To show the converse, note that by Lemma B.6, $C_p(F)$ is an evident p -belief event. Also, it is immediate that $C_p(F) \subseteq B_{i,p}(F)$ for all $i \in N$. It thus remains to show that for every t_i such that $(\theta, t_i, t_j) \in B_{i,p}(C_p(F))$ for some θ, t_j , it holds that $[B_p^k(F)]_{t_i} \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i)$ for all k . Since $p > 0$, $[C_p(F)]_{t_i} = \bigcap_k \{(\theta', t'_j) : (\theta', t_i, t'_j) \in B_p^k(F)\}$ is nonempty. Hence, for all k , there exist θ, t'_j such that $(\theta, t_i, t'_j) \in B_p^k(F)$, so that $(\theta, t_i, t'_j) \in B_{i,p}(B_p^{k-1}(F))$, where $B_p^0(F) := F$. By the definition of $B_{i,p}(B_p^{k-1}(F))$, $[B_p^{k-1}(F)]_{t_i} \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i)$. \square

B.3 Proof of Proposition 5.3

Suppose by contradiction that such a finite set \mathcal{B}_i exists for all $i \in N$, and that $C_p(F) = \emptyset$. Then, for each $i \in N$, there is $K_i < \infty$ such that for each $k > K_i$,

$$\{t_i \in T_i : [B_p^k(F)]_{t_i} \in \mathcal{F}_\Theta \otimes \Sigma_i(t_i), \beta_i(t_i)([B_p^k(F)]_{t_i}) \geq p\} = \{t_i \in T_i : h_i^{T, d_i-1}(t_i) \in B_i\}$$

for some (fixed) $B_i \in \mathcal{B}_i$. Then, if $k > K_i$ for all i , $B_p^{k+1}(F) = B_p^k(F)$, so that $C_p(F) = B_p^{k+1}(F) \neq \emptyset$, a contradiction. \square

Appendix C The role of Conditions 1–3

Section 3 introduced various conditions that extended type spaces satisfy. Structures that satisfy Conditions 1–3 have the property that each type generates a belief hierarchy of a well-defined depth (Theorem 4.2), and that types with different σ -algebras cannot have the same depth of reasoning (Proposition 4.3). Here we show that this is no longer true if one of the conditions is relaxed.

Before we proceed, we introduce some notation that is used in the examples in this appendix. We label the players by a and b , and refer to them as Ann and Bob, respectively. The set of states of nature is $\Theta = \{H, L\}$. To save space, we use a compact representation of the type spaces. We restrict attention to finite type sets, so that the σ -algebras on the type sets can be represented by the partitions that generates them. For example, if the type set is $\{t_a, t'_a\}$, the σ -algebra $\mathcal{F}_a := \{\{t_a\}, \{t'_a\}, \{t_a, t'_a\}, \emptyset\}$ is generated by the partition $\{\{t_a\}, \{t'_a\}\}$, so that \mathcal{F}_a can be depicted by:

$$\mathcal{F}_a: \{t_a\}, \{t'_a\},$$

and the trivial σ -algebra $\mathcal{F}'_a := \{\{t_a, t'_a\}, \emptyset\}$ is given as

$$\mathcal{F}'_a: \{t_a, t'_a\}.$$

C.1 Depth of reasoning not well-defined

Condition 3 turns out to be crucial in ensuring that each type has a well-defined depth of reasoning. We show here that a structure that satisfies only a weaker version of Condition 3 (as well as Conditions 1 and 2) can contain types that do not have a well-defined depth.

To demonstrate this, recall the definition of the depth of a type. We also use the notation $\sigma(h_i^{\mathcal{T},k})$ introduced in Appendix A.1.2: Given a structure $\mathcal{T} = (T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$ and a player $i \in N$, $\sigma(h_i^{\mathcal{T},k})$ is the σ -algebra on T_i consisting of events of the form $\{t_i \in T_i : h_i^{\mathcal{T},k}(t_i) \in B\}$ for $B \in \mathcal{F}_{B_i^{\mathcal{T},k}}$. Then, a type t_i does not have a well-defined depth of reasoning if there exists $k = 0, 1, \dots$ such that¹⁵

- (1) $\Sigma_i(t_i) \supseteq \sigma(h_j^{\mathcal{T},\ell})$ for $\ell \leq k$, and $\Sigma_i(t_i) \not\supseteq \sigma(h_j^{\mathcal{T},\ell})$ for $\ell > k$; and
- (2) there is $\ell > k$ and $E \in \sigma(h_j^{\mathcal{T},\ell}) \setminus \sigma(h_j^{\mathcal{T},k})$ such that $E \in \Sigma_i(t_i)$.

The first requirement states that t_i can reason about all events up to order k , but not about events at higher order. The second requirement says that there are events at higher orders that the type can reason about, even though they are not expressible in terms of events of order k . Thus, if a type does not have a well-defined depth, then it can reason about certain events that are not “reducible” to lower-order events, but not about other events that require an equal level of sophistication.

The weakening of Condition 3 that we consider is the following:

¹⁵We ignore here that the depth of a belief hierarchy is not allowed to depend on the type space that gives rise to it. The present definition is sufficient for our negative result.

Condition 3'. For any player $i \in N$ and σ -algebra $\mathcal{F}_i \in \mathcal{S}_i$ such that $\mathcal{F}_i \neq \{T_i, \emptyset\}$, there is $\mathcal{F}_j \in \mathcal{S}_j$ such that

(a') $(\mathcal{F}_i, \mathcal{F}_j)$ form a mutual-dominance pair; or

(b') $\mathcal{F}_i \succ \mathcal{F}_j$, and there is no $\mathcal{F}'_i \in \mathcal{S}_i$ such that $\mathcal{F}'_i \succ \mathcal{F}_j$ and $\mathcal{F}'_i \subsetneq \mathcal{F}_i$.

Part (b') in Condition 3' relaxes Condition 3(b), which requires that \mathcal{F}_i is the coarsest σ -algebra among *all* σ -algebras on T_i that dominates \mathcal{F}_j ; rather, there cannot be a coarser σ -algebra *in the collection* \mathcal{S}_i that dominates \mathcal{F}_j .

The next example constructs a structure that satisfies this weaker version of Condition 3, yet some types do not have a well-defined depth:

Example 3. We define a structure $\mathcal{T} = (T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i=a,b}$ that satisfies Conditions 1, 2, and 3', but not Condition 3. Each player $i = a, b$ has four types, labeled $t_i^1, t_i^2, t_i^3, t_i^4$, and the σ -algebras are:

$$\begin{aligned} \mathcal{F}_i: & \{t_i^1, t_i^2, t_i^3, t_i^4\}; \\ \mathcal{F}'_i: & \{t_i^1, t_i^2\}, \{t_i^3\}, \{t_i^4\}; \\ \mathcal{F}''_i: & \{t_i^1\}, \{t_i^2\}, \{t_i^3\}\{t_i^4\}. \end{aligned}$$

All types for player i are endowed with the σ -algebra \mathcal{F}'_i on T_j . The types' beliefs are described by:

$$\begin{aligned} \beta_i(t_i^1)(H, \{t_j^1, t_j^2\}) &= \frac{1}{2}, & \beta_i(t_i^1)(H, \{t_j^3\}) &= \frac{1}{2}; \\ \beta_i(t_i^2)(H, \{t_j^4\}) &= 1; \\ \beta_i(t_i^3)(L, \{t_j^1, t_j^2\}) &= \frac{1}{2}, & \beta_i(t_i^3)(L, \{t_j^3\}) &= \frac{1}{2}; \\ \beta_i(t_i^4)(L, \{t_j^4\}) &= 1. \end{aligned}$$

Condition 1 is clearly satisfied; it is also not hard to check that Conditions 2 and 3' hold. Condition 3 fails, however: while \mathcal{F}'_b is the coarsest σ -algebra in \mathcal{S}_b that dominates \mathcal{F}_a , the σ -algebra generated by the partition $\{\{t_b^1, t_b^2\}, \{t_b^3, t_b^4\}\}$ is a coarsening of \mathcal{F}'_b that dominates \mathcal{F}_a .

Not every type in this structure has a well-defined depth of reasoning. Type t_a^1 , for example, cannot assign a probability to the third-order event that Bob assigns equal probability to the state of nature being H and her believing that $\theta = H$ and to the state of nature being H and her believing that $\theta = L$, because t_b^1 and t_b^2 have different beliefs about this second-order event. On the other hand, t_a^1 is able to form a belief about another third-order event, even though that event cannot be expressed as a second-order event: t_a^1 assigns probability $\frac{1}{2}$ to the event that Bob assigns probability equal probability to the event that $\theta = L$ and that she

believes that $\theta = H$, and to the event that $\theta = L$ and that she believes that $\theta = L$. The reason is that t_a^1 can assign a probability to the subset of types (the singleton $\{t_b^3\}$) that have this second-order belief. There is no corresponding second-order event: all second-order events are generated by the subsets $\{t_b^1, t_b^2\}$ and $\{t_b^3, t_b^4\}$. \triangleleft

C.2 Redundant σ -algebras

Example 3 demonstrated that types need not have a well-defined depth of reasoning if a structure does not satisfy Condition 3. What happens if we relax one of the other conditions, but insist on Condition 3? In that case, it follows from Lemma A.4 that each type generates a belief hierarchy of well-defined depth. However, what can happen if we relax Condition 1 or 2 is that there are types (for a given player) that have different σ -algebras, yet have the same depth of reasoning, as we show now. In the Harsanyi context, this would be akin to endowing types with different σ -algebras (while requiring that the belief maps be measurable, so that every type has an infinite depth of reasoning).

C.2.1 Relaxing Condition 1

We first present an example of a structure that satisfies Conditions 2 and 3, but not Condition 1. In this structure, two types with different σ -algebras both have an infinite depth of reasoning. We then show that this result holds more generally.

Example 4. Each player $i = a, b$ has eight types, labeled $t_i^1, t_i^2, t_i^3, t_i^4, t_i^5, t_i^6, t_i^7, t_i^8$, and the σ -algebras are:

$$\begin{array}{ll}
\mathcal{F}_a: & \{t_a^1, t_a^2, t_a^3, t_a^4, t_a^5, t_a^6, t_a^7, t_a^8\}; \\
\mathcal{F}'_a: & \{t_a^1, t_a^2\}, \{t_a^3, t_a^4\}, \{t_a^5, t_a^6\}, \{t_a^7, t_a^8\}; \\
\mathcal{F}''_a: & \{t_a^1, t_a^3\}, \{t_a^2, t_a^4\}, \{t_a^5, t_a^7\}, \{t_a^6, t_a^8\}; \\
\mathcal{F}'''_a: & \{t_a^1\}, \{t_a^2\}, \{t_a^3\}, \{t_a^4\}, \{t_a^5\}, \{t_a^6\}, \\
& \{t_a^7\}, \{t_a^8\}; \\
\mathcal{F}_b: & \{t_b^1, t_b^2, t_b^3, t_b^4, t_b^5, t_b^6, t_b^7, t_b^8\}; \\
\mathcal{F}'_b: & \{t_b^1, t_b^2\}, \{t_b^3, t_b^4\}, \{t_b^5, t_b^6\}, \{t_b^7, t_b^8\}; \\
\mathcal{F}''_b: & \{t_b^1, t_b^3\}, \{t_b^2, t_b^4\}, \{t_b^5, t_b^7\}, \{t_b^6, t_b^8\}; \\
\mathcal{F}'''_b: & \{t_b^1\}, \{t_b^2\}, \{t_b^3\}, \{t_b^4\}, \{t_b^5\}, \{t_b^6\}, \\
& \{t_b^7\}, \{t_b^8\}.
\end{array}$$

Condition 1 is clearly not satisfied, as the σ -algebra \mathcal{F}'_i for $i = a, b$ is not a subset of \mathcal{F}''_i or vice versa. We now assign beliefs to the types in such a way that Conditions 2 and 3 do hold.

The odd-numbered types t_a for Ann have σ -algebra $\Sigma_a(t_a) = \mathcal{F}'_b$; her even-numbered types have σ -algebra $\Sigma_a(t_a) = \mathcal{F}'''_b$. For Bob, the odd-numbered types t_b are endowed with the σ -algebra $\Sigma_b(t_b) = \mathcal{F}''_a$, and the even-numbered types have σ -algebra $\Sigma_b(t_b) = \mathcal{F}'''_a$. The types' beliefs are given by:

$$\begin{array}{ll}
\beta_a(t_a^1)(H, \{t_b^5, t_b^6\}) = 1; & \beta_b(t_b^1)(H, \{t_a^5, t_a^7\}) = 1; \\
\beta_a(t_a^2)(H, \{t_b^5\}) = 1; & \beta_b(t_b^2)(H, \{t_a^7\}) = 1; \\
\beta_a(t_a^3)(H, \{t_b^7, t_b^8\}) = 1; & \beta_b(t_b^3)(H, \{t_a^6, t_a^8\}) = 1; \\
\beta_a(t_a^4)(H, \{t_b^7\}) = 1; & \beta_b(t_b^4)(H, \{t_a^8\}) = 1; \\
\beta_a(t_a^5)(L, \{t_b^1, t_b^2\}) = 1; & \beta_b(t_b^5)(L, \{t_a^1, t_a^3\}) = 1; \\
\beta_a(t_a^6)(L, \{t_b^1\}) = 1; & \beta_b(t_b^6)(L, \{t_a^3\}) = 1; \\
\beta_a(t_a^7)(L, \{t_b^3, t_b^4\}) = 1; & \beta_b(t_b^7)(L, \{t_a^2, t_a^4\}) = 1; \\
\beta_a(t_a^8)(L, \{t_b^3\}) = 1; & \beta_b(t_b^8)(L, \{t_a^4\}) = 1.
\end{array}$$

It is straightforward to check that this structure satisfies Conditions 2 and 3. This structure has types that have different σ -algebras, yet have the same depth of reasoning. Type t_a^1 for Ann has a σ -algebra that is strictly finer than that of type t_a^2 , but each of these types believes (with probability 1) that $\theta = H$, that Bob believes that $\theta = L$, that Bob believes that Ann believes that $\theta = H$, and so on. That is, t_a^1 and t_a^2 both have an infinite depth of reasoning (and in fact induce the same belief hierarchy). \triangleleft

In the structure in Example 4, there are types for a given player that have different σ -algebras, yet each has an infinite depth of reasoning. This is not an accident, as we show now: For *any* structure that satisfies Conditions 2 and 3 and violates Condition 1, there are types (for a given player) with different σ -algebras that have the *same* depth of reasoning, provided that there is a type whose σ -algebra does not belong to the “usual” filtration.

Formally, fix a structure $\mathcal{T} = (T_i, \mathcal{S}_i, \Sigma_i, \beta_i)_{i \in N}$ that satisfies Conditions 2 and 3, and define

$$\begin{aligned}
\mathcal{S}_i^* := \{ \mathcal{F}_i \in \mathcal{S}_i : & \text{there is } k \text{ s.t. } \mathcal{F}_i = \sigma(h_i^{\mathcal{T}, k}) \text{ or} \\
& \text{there is } \mathcal{F}_j \in \mathcal{S}_j \text{ such that } (\mathcal{F}_i, \mathcal{F}_j) \text{ is a mutual-dominance pair} \}.
\end{aligned}$$

If the structure \mathcal{T} satisfies Condition 1 (as well as the other conditions), then it follows from our earlier results that $\mathcal{S}_i = \mathcal{S}_i^*$. However, Example 4 illustrates that this need not be the case if Condition 1 does not hold.

We can now state the result: If there is a type t_i with a σ -algebra that does not belong to the filtration \mathcal{S}_j^* , then there is a type t'_i with another σ -algebra that has the same depth.

Proposition C.1. *Suppose \mathcal{T} is a structure that satisfies Conditions 2 and 3. If there is $i \in N$ and $t_i \in T_i$ such that $\Sigma_i(t_i) \in \mathcal{S}_j \setminus \mathcal{S}_j^*$, then there is $t'_i \in T_i$ with $\Sigma_i(t'_i) \neq \Sigma_i(t_i)$ such that $d_i(t'_i) = d_i(t_i) = \infty$.*

The proof is relegated to the online appendix. If there are different σ -algebras that correspond to the same depth of reasoning, one of the σ -algebras is redundant: one can always

define the beliefs of a type with that σ -algebra on one of the other σ -algebra, without changing the belief hierarchies. We next show that this form of redundancy also occurs when we relax Condition 2.

C.2.2 Relaxing Condition 2

We first present a simple example of a structure that violates Condition 2 and in which types with different σ -algebras induce the same belief hierarchy.

Example 5. For each player $i = a, b$, the σ -algebras in \mathcal{S}_i are given by:

$$\begin{aligned}\mathcal{F}_i: & \{t_i^1, t_i^2\}; \\ \mathcal{F}'_i: & \{t_i^1\}, \{t_i^2\}.\end{aligned}$$

Types t_a^1 and t_a^2 are endowed with the σ -algebra \mathcal{F}'_b ; types t_b^1 has σ -algebra \mathcal{F}_a and t_b^2 has σ -algebra \mathcal{F}'_a . The beliefs of the types over the state of nature and the type of the other player are given by:

$$\begin{aligned}\beta_a(t_a^1)(H, t_b^1) &= 1, & \beta_b(t_b^1)(H, T_a) &= 1 \\ \beta_a(t_a^2)(H, t_b^2) &= 1, & \beta_b(t_b^1)(H, t_a^1) &= 1\end{aligned}$$

It is readily verified that this structure satisfies Conditions 1 and 3. The structure does not satisfy Condition 2, however: both $(\mathcal{F}_a, \mathcal{F}_b)$ and $(\mathcal{F}'_a, \mathcal{F}'_b)$ are mutual-dominance pairs. Indeed, every type has an infinite depth of reasoning. For example, while type t_b^2 for Bob has a strictly finer σ -algebra than t_b^1 , both types believe that $\theta = H$, believe that Ann believes that, believes that Ann believes that Bob believes that, and so on. \triangleleft

Note that the structure in Example 5 can be viewed as a Harsanyi type space where different types for a player can have different σ -algebras.

The example can readily be adapted to see that a weaker version of Condition 2 does not suffice to rule out redundant σ -algebras. Suppose we adopt the following weakening of Condition 2:

Condition 2'. Fix $i \in N$, and let $\mathcal{F}_i \in \mathcal{S}_i$, $\mathcal{F}_j \in \mathcal{S}_j$. If $(\mathcal{F}_i, \mathcal{F}_j)$ is a mutual-dominance pair, then there is no $\mathcal{F}'_i \neq \mathcal{F}_i$ in \mathcal{S}_i such that $(\mathcal{F}'_i, \mathcal{F}_j)$ is a mutual-dominance pair for some σ -algebra $\mathcal{F}'_j \in \mathcal{S}_j$.

This condition is weaker than Condition 2 in that it requires that \mathcal{F}_j belongs to \mathcal{S}_j . It is satisfied by the structure that is identical to the structure in Example 5 except that the set

\mathcal{S}_a of σ -algebras on T_a is taken to be $\{\mathcal{F}'_a\}$; this structure also satisfies Condition 1 and 3. But types t_b^1 and t_b^2 both have an infinite depth of reasoning, as before.

Again, we can show a general result:

Proposition C.2. *Suppose \mathcal{T} is a structure that satisfies Conditions 1 and 3. If there exist types $t_i, t'_i \in T_i$ for player i and σ -algebras $\mathcal{F}_i, \mathcal{F}'_i \in \mathcal{S}_i$ such that $(\Sigma_i(t_i), \mathcal{F}_i)$ and $(\Sigma_i(t'_i), \mathcal{F}'_i)$ are mutual-dominance pairs, then $d_i(t_i) = d_i(t'_i) = \infty$.*

The proof is immediate, and thus omitted.

References

- Ahn, D. S. (2007). Hierarchies of ambiguous beliefs. *Journal of Economic Theory* 136, 286–301.
- Aliprantis, C. D. and K. C. Border (2005). *Infinite Dimensional Analysis: A Hitchhiker's Guide* (3rd ed.). Berlin: Springer.
- Aumann, R. J. (1976). Agreeing to disagree. *Annals of Statistics* 4(6), 1236–1239.
- Brandenburger, A. and E. Dekel (1987). Common knowledge with probability 1. *Journal of Mathematical Economics* 16, 237–245.
- Brandenburger, A. and E. Dekel (1993). Hierarchies of beliefs and common knowledge. *Journal of Economic Theory* 59, 189–198.
- Brandenburger, A. and H. J. Keisler (2006). An impossibility theorem on beliefs in games. *Studia Logica* 84, 211–240.
- Carlsson, H. and E. van Damme (1993). Global games and equilibrium selection. *Econometrica* 61, 989–1018.
- Chwe, M. S.-Y. (2001). *Rational Ritual: Culture, Coordination, and Common Knowledge*. Princeton, NJ: Princeton University Press.
- Clark, H. H. and C. Marshall (1978). Reference diaries. In *Proceedings of the 1978 workshop on Theoretical Issues in Natural Language Processing*, pp. 57–63.
- Clark, H. H. and C. Marshall (1981). Definite reference and mutual knowledge. In A. K. Joshi, B. L. Webber, and I. A. Sag (Eds.), *Elements of Discourse Understanding*, Chapter 1. Cambridge, UK: Cambridge University Press.

- Costa-Gomes, M., V. P. Crawford, and B. Broseta (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica* 69, 1193–1235.
- Di Tillio, A. (2008). Subjective expected utility in games. *Theoretical Economics* 3, 287–323.
- Fagin, R. and J. Y. Halpern (1988a). Belief, awareness and limited reasoning. *Artificial Intelligence* 34, 39–76.
- Fagin, R. and J. Y. Halpern (1988b). Reasoning about knowledge and probability. In M. Y. Vardi (Ed.), *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, San Mateo, CA, pp. 277–293. Morgan-Kaufmann. Full version published in the *Journal of the ACM*, Vol 41. No. 2, 1994, pp 340–367.
- Feinberg, Y. (2004). Subjective reasoning: Games with unawareness. Working paper, Stanford GSB.
- Friedell, M. (1969). On the structure of shared awareness. *Behavioral Science* 14, 28–39.
- Friedenberg, A. (2010). When do type structures contain all hierarchies of beliefs? *Games and Economic Behavior* 68, 108–129.
- Friedenberg, A. and H. J. Keisler (2011). Iterated dominance revisited. Working paper, Arizona State University.
- Friedenberg, A. and M. Meier (2012). On the relationship between hierarchy and type morphisms. *Economic Theory* 46, 377–399.
- Geanakoplos, J. (1989). Game theory without partitions and applications to speculation and consensus. Working paper, Cowles Foundation for Research in Economics.
- Geanakoplos, J. D. and H. M. Polemarchakis (1982). We can't disagree forever. *Journal of Economic Theory* 28, 192–200.
- Halpern, J. Y. (2003). *Reasoning about Uncertainty*. Cambridge, MA: MIT Press.
- Harsanyi, J. C. (1967–1968). Games on incomplete information played by Bayesian players. Parts I–III. *Management Science* 14, 159–182, 320–334, 486–502.
- Heifetz, A. and W. Kets (2011). All types naive and canny. Working paper.
- Heifetz, A., M. Meier, and B. Schipper (2006). Interactive unawareness. *Journal of Economic Theory* 130, 78–94.

- Heifetz, A. and D. Samet (1998a). Knowledge spaces with arbitrarily high rank. *Games and Economic Behavior* 22, 260–273.
- Heifetz, A. and D. Samet (1998b). Topology-free typology of beliefs. *Journal of Economic Theory* 82, 324–341.
- Heinsalu, S. (2011). Universal type spaces with unawareness. Mimeo, Yale University.
- Ho, T.-H., C. Camerer, and K. Weigelt (1998). Iterated dominance and iterated best response in experimental “*p*-beauty contests”. *American Economic Review* 88, 947–969.
- Kets, W. (2012). Bounded reasoning: A canonical type space. Mimeo.
- Lewis, D. (1969). *Convention: A Philosophical Study*. Cambridge, MA: Harvard University Press.
- Mertens, J. F. and S. Zamir (1985). Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory* 14, 1–29.
- Monderer, D. and D. Samet (1989). Approximating common knowledge with common beliefs. *Games and Economic Behavior* 1, 170–190.
- Morris, S., A. Postlewaite, and H. Shin (1995). Depth of knowledge and the effect of higher order uncertainty. *Economic Theory* 6, 453–467.
- Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review* 85, 1313–1326.
- Pintér, M. and Z. Udvari (2011). Generalized type spaces. Working paper 34107, MPRA.
- Qin, C. and C. Yang (2010). An explicit approach to modeling finite-order type spaces and applications. Working paper, U.C. Santa Barbara.
- Rubinstein, A. (1989). The electronic mail game: Strategic behavior under “almost common knowledge”. *American Economic Review* 79, 385–391.
- Savage, L. J. (1954). *The Foundations of Statistics*. John Wiley & Sons.
- Stahl, D. O. and P. W. Wilson (1995). On players’ models of other players: Theory and experimental evidence. *Games and Economic Behavior* 10, 218–254.
- Strzalecki, T. (2009). Depth of reasoning and higher order beliefs. Working paper, Harvard University.