

The Existence of Universal Qualitative Belief Spaces*

Satoshi Fukuda[†]

December 17, 2023

Abstract

This paper constructs a canonical representation of players' interactive beliefs, irrespective of natures of beliefs: whether beliefs are qualitative, truthful (i.e., knowledge), or probabilistic (e.g., countably-additive, finitely-additive, or non-additive). The canonical model is the "largest" interactive belief model to which any particular model can be mapped in a unique belief-preserving way. The key insight for the construction is the need to specify players' possible depth of reasoning up to which they can interactively reason about their beliefs (e.g., their beliefs, their beliefs about their beliefs, their beliefs about their beliefs about their beliefs, and so on). The possible depth of reasoning may be a transfinite level (beyond any finite level) when beliefs are qualitative. The specification of possible depth of reasoning also has game-theoretic implications for characterizations of some solution concepts using the canonical space. For instance, for any strategic game with ordinal payoffs, there exists a canonical interactive belief model which characterizes iterated elimination of strictly dominated actions as an implication of common belief in rationality.

JEL Classification: C70; D83

Keywords: Qualitative Belief; Probabilistic Belief; Knowledge; Universal Space; Terminal Space; Depth of Reasoning

*This paper is based on part of the first chapter of my Ph.D. thesis, "The Existence of Universal Knowledge Spaces," submitted to the University of California at Berkeley. I would like to thank David Ahn, William Fuchs, and Chris Shannon for their encouragement, support, and guidance; Robert Anderson, Eric Auerbach, Yu Awaya, Ivan Balbuzanov, Paulo Barelli, Pierpaolo Battigalli, Benjamin Brooks, Christopher Chambers, Aluma Dembo, Eduardo Faingold, Amanda Friedenberg, Drew Fudenberg, Itzhak Gilboa, Benjamin Golub, Sander Heinsalu, Ryuichiro Ishikawa, Yuichiro Kamada, Michihiro Kandori, Shachar Kariv, Ali Khan, Narayana Kocherlakota, Massimo Marinacci, Stephan Morris, Motty Perry, Herakles Polemarchakis, Joseph Root, Burkhard Schipper, Mikkel Sølvsten, Kenji Tsukada, and various conference and seminar participants for comments and discussions. I would also like to thank the editor and two anonymous referees for excellent suggestions that have greatly improved the manuscript. All remaining errors are mine.

[†]Department of Decision Sciences and IGIER, Bocconi University, Milan 20136, Italy.

1 Introduction

Consider a group of players who reason interactively about unknown external values, *states of nature* S , such as strategies in a game. Players reason about states of nature—their strategies. Players also reason about their beliefs about states of nature—their beliefs about each other’s strategies and consequently their rationality. And so on. This paper constructs the first formal framework general enough to represent any conceivable form of interactive beliefs irrespective of specific natures of beliefs: whether beliefs are probabilistic or qualitative (i.e., binary) including knowledge.

Unlike having been previously thought, the existence of the canonical space does not depend on a specific nature of beliefs. Rather, the key insight behind the construction of the canonical space is the need to specify players’ possible depth of reasoning. Roughly speaking, it is a pre-specified level of interactive beliefs of the form, Alice believes that Bob believes that Alice believes that In the previous literature on Harsanyi (1967-68) type spaces, the specification of possible depth of reasoning is given by the implicit assumption that players’ countably-additive beliefs are defined on a σ -algebra. This implicit assumption allows one to analyze any (at-most) countable-level interactive beliefs. Turning to game-theoretic applications, for any given strategic game where the players posses qualitative beliefs (including the case of knowledge), a canonical belief space in which players’ possible depth of reasoning is appropriately chosen characterizes the solution concept of iterated elimination of strictly dominated actions as an implication of common belief in rationality.

To see the importance of specifying players’ possible depth of reasoning, consider the following two-player strategic game $\langle (A_i)_{i \in I}, (u_i)_{i \in I} \rangle$.

Example 1. Each player i announces an element from an ordered set $A_i := \{0, 1, 2, 3, 4, \dots\} \cup \{a, b\}$, where $0 < 1 < 2 < 3 < 4 < \dots < a < b$. Announcing b always yields a payoff of 1 irrespective of the opponent’s announcement. For any other announcements, if i ’s announcement is (strictly) higher than the opponent’s, she obtains a payoff of 2; if not, she obtains a payoff of 0. Table 1 depicts player i ’s payoff $u_i(a_i, a_{-i})$ as a function of a_i (Row) and a_{-i} (Column).

Consider the solution concept of iterated elimination of strictly dominated actions (IESDA). At each round of elimination, the minimal element is always strictly dominated. First, $a_i = 0$ is a unique strictly dominated action in A_i ; next, $a_i = 1$ is a unique strictly dominated action in $A_i \setminus \{0\}$; and so on. Once $\{0, 1, 2, 3, 4, \dots\}$ have been deleted from each player’s action set, in the subgame in which each player’s action set is $\{a, b\}$, action a is strictly dominated by b . Thus, the action profile $(a_1, a_2) = (b, b)$ is the unique prediction under IESDA after one more elimination (i.e., a) after eliminating $0, 1, 2, 3, 4, \dots$. Hence, in order to reach the unique prediction under IESDA, the players need to engage in a transfinite process of IESDA.¹

¹Lipman (1994) is a pioneering work pointing out the need for a transfinite (yet countable) process of iterated elimination of never-best-replies.

	0	1	2	3	4	a	b
0	0	0	0	0		0	0
1	2	0	0	0		0	0
2	2	2	0	0	0	0	0
3	2	2	2	0	0	0	0
4	2	2	2	2	0	0	0
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
a	2	2	2	2	2	0	0
b	1	1	1	1	1	1	1

Table 1: Player i 's payoff $u_i(a_i, a_{-i})$ in Example 1 as a function of a_i (Row) and a_{-i} (Column).

While this example suggests that players may need to engage in a transfinite yet countable level of interactive reasoning, the next example suggests that players may need to engage in an *arbitrarily long* transfinite level of interactive reasoning. The example also suggests that, for such a strategic game, a traditional type space may not be able to properly formalize players' uncountably many iterations of interactive reasoning.

Example 2. Instead of actions $\{0, 1, 2, 3, 4, \dots\}$ in Example 1, consider an interval $[0, 1]$ with a total order \leq satisfying the following two properties: (i) $x < a < b$ for any $x \in [0, 1]$, where $<$ is the strict order associated with \leq (i.e., $x < a$ if $x \leq a$ and $x \neq a$); and (ii) any non-empty subset of $[0, 1]$ has the minimum element with respect to \leq .² Each player i 's action set is $A_i := [0, 1] \cup \{a, b\}$. Define her payoff function u_i similarly to Example 1 using the strict order $<$.

Consider IESDA. As in the previous example, at each round of elimination, the minimal element is always strictly dominated. First, the least element (denote by $a_i = a^0$) is a unique strictly dominated action in A_i ; next, the second least element (denote by a^1) is a unique strictly dominated action in $A_i \setminus \{a^0\}$; and so on. By the ordering \leq on $[0, 1]$, one can eliminate actions $[0, 1]$ one by one.³ In the subgame in

²It is well-known in set theory that, by the Axiom of Choice, for any set, a total order satisfying Property (ii) exists. Technically, one can identify $[0, 1] = \{a^\beta\}_{\beta < \alpha}$ where α is some ordinal number, so that the elements $\{a^\beta\}_{\beta < \alpha}$ are ordered in a desired way. For the concepts in set theory which are used in this paper (e.g., ordinal and cardinal numbers), see, for instance, Hrbacek and Jech (1999).

³Action a^0 is eliminated from A_i . For any natural number n , action a^n is eliminated from $A_i \setminus \{a^0, \dots, a^{n-1}\}$. Once actions $\{a^0, a^1, \dots, a^n, \dots\}$ have been eliminated, the next least element in $A_i \setminus \{a^0, a^1, \dots, a^n, \dots\}$ is eliminated (as action a is eliminated after $\{0, 1, \dots, n, \dots\}$ have been eliminated in Example 1). The process of elimination lasts indefinitely until all actions $[0, 1]$ will be eliminated (this can be rigorously formalized using ordinal numbers). Section 6.3.2 revisits Examples 1 and 2 using ordinal numbers.

which each player's action set is $\{a, b\}$, action a is strictly dominated by b . Thus, the action profile (b, b) is the unique prediction under IESDA after uncountably many eliminations.

This example can be generalized by replacing the set $[0, 1]$ with any set with larger cardinality. Hence, in order to reach the unique prediction under IESDA, the players may need to engage in an arbitrarily long transfinite process of IESDA.

Thus, when analysts provide an epistemic characterization of such solution concepts as IESDA for a given strategic game, it would be natural for them to consider a class of interactive belief models in which players can engage in interactive reasoning of an arbitrary but predetermined ordinal level. A pre-determined ordinal level is sufficient because the number of eliminations of strictly dominated actions does not exceed the number of action profiles in the original strategic game.

To see how the appropriately-chosen canonical interactive belief model captures predictions under IESDA, I define a model of beliefs (a belief space).

A Belief Space. A model of beliefs (a belief space) consists of the following four ingredients. The first ingredient is a set Ω . Each element $\omega \in \Omega$ is a list of possible specifications of the prevailing nature state $s \in S$ (e.g., an action profile played at state ω if S is the set of action profiles) and players' interactive beliefs regarding nature states S (i.e., their beliefs about nature states S , their beliefs about their beliefs about S , and so on). Call each ω a state (of the world).

The second ingredient is a mapping Θ , which associates, with each state of the world $\omega \in \Omega$, the corresponding state of nature $\Theta(\omega) \in S$. For instance, when players are reasoning about their actions in a given strategic game (i.e., when S is the set of action profiles of the given strategic game), the mapping Θ is a profile of players' strategies which associate, with each state of the world, the corresponding actions played at that state.

The third ingredient is the set of statements about which players can reason. These statements, specified as subsets of states of the world Ω , are referred to as events. For instance, when S is the set of action profiles of the given strategic game, an event corresponds to a set of states of the world at which the players take a certain action profile. A belief space has to specify the collection of events about which players can reason, which I call the domain.

The fourth ingredient is players' belief operators defined on the domain. For each event E , player i 's belief operator assigns the set of states at which she believes that E has occurred (simply referred to as "she believes E "), i.e., the event that i believes E . Specifying players' beliefs through belief operators is general enough to accommodate their probabilistic beliefs as well as qualitative beliefs. This is because probabilistic beliefs can reduce to whether a player believes an event with probability at least p (she p -believes the event) or not (Monderer and Samet, 1989).⁴ This paper

⁴In fact, Samet (2000) demonstrates the equivalence between the type space and the p -belief operator approaches when players possess countably-additive probabilistic beliefs.

demonstrates that one can represent various notions of qualitative or probabilistic beliefs by imposing properties of those beliefs on belief operators. Thus, the framework can accommodate qualitative beliefs (including knowledge) and probabilistic (countably-additive, finitely-additive, or non-additive) beliefs.

Iterative applications of belief operators generate higher-order interactive reasoning. As will be discussed in Section 1.1, certain logical (i.e., set-theoretic) assumptions on the domain determine players' possible depth of reasoning through their belief operators. For example, if the domain of a belief space is closed under countable conjunction (i.e., set-theoretic countable intersection), then the belief space can represent common belief in rationality held by Alice and Bob, who are reasoning about their actions in a strategic game: they believe that they are rational, they believe that they believe that they are rational, and so on, *ad infinitum*. In the strategic game in Example 1, one can formalize the countable iterations of mutual beliefs in rationality (which lead to the countable elimination of strictly dominated actions in the example) if the domain of a given belief space is a σ -algebra.⁵

Main Result and Its Application to IESDA (Informal). The main result of the paper (Theorem 1 in Section 3) demonstrates the existence of a universal (precisely, terminal) belief space into which any belief space is embedded in a unique manner that maintains players' interactive beliefs about nature states in that smaller space. As each belief space enables one to describe interactive beliefs without explicitly specifying the set of belief hierarchies that the given belief space induces, the existence of a universal belief space guarantees that the belief space approach can represent all possible hierarchies of beliefs.

To discuss its game-theoretic application to the solution concept of IESDA, consider the following epistemic characterization. In any belief space, if the players have (correct) common belief in rationality at a state, then their actions at that state survive IESDA; and for any action profile that survives IESDA, there exists a belief space in which there is a state at which the players have (correct) common belief in rationality and take the given actions.⁶

In this characterization, within a particular belief space, some action profiles that survive IESDA may not be predictions under common belief in rationality. Thus, as

⁵However, for the strategic game in Example 2, a belief space defined on a σ -algebra (including a traditional type space) may not be able to formalize an uncountably long iteration of mutual beliefs in rationality because the σ -algebra may not be closed under uncountable intersection.

⁶Two remarks are in order. First, see Brandenburger and Dekel (1987), Stalnaker (1994), and Tan and Werlang (1988) for pioneering epistemic characterizations of IESDA. Second, technically, correct common belief in rationality means that whenever the players commonly believe their rationality, they are rational. This is weaker than directly assuming correctness on beliefs (i.e., Truth Axiom, stating that if a player believes that an event has occurred then the event indeed has occurred). When each player's belief is assumed to satisfy certain logical properties, rationality and common belief in rationality (instead of assuming correct common belief in rationality) also characterize IESDA. See Section 6.3 for details.

Friedenberg and Keisler (2021) argue, the above epistemic characterization pertains to an analyst who does not know the players' beliefs. Such an analyst would seek the predictions under IESDA across all belief spaces associated with the given strategic game. In contrast, in order to provide an epistemic characterization of IESDA from the players' perspective, namely, in order to obtain an epistemic characterization of IESDA in a belief space in which the players' reasoning leads to the set of *all* action profiles that survive IESDA, one would need to have a belief space that induces all possible belief hierarchies about the play of the game. This leads to the epistemic characterization in a universal belief space.⁷

Indeed, as a game-theoretic application, Section 6.3 shows that any state of a belief space at which the players commonly believe their rationality is uniquely mapped to the corresponding state of the universal space at which the players commonly believe their rationality. Thus, in the universal belief space, if the players commonly believe their rationality at a state, then their actions at the state survive IESDA; and for any action profile that survives IESDA, there exists a state in the universal space at which the players commonly believe their rationality and take the given actions. For any given strategic game with ordinal payoffs, there exists a universal belief space that can represent the players' iterated reasoning about their actions up to a predetermined ordinal level, which suffices to pin down the predictions. Section 6.4 characterizes Börgers (1993) dominance as an implication of common belief in weak-dominance rationality for any strategic game with ordinal payoffs (Bonanno and Tsakas, 2018).

I establish the existence of a universal belief space as long as players' beliefs are represented by belief operators. My result is theoretically interesting in that the existence of a universal belief space is unrelated to assumptions on properties of beliefs. For example, my paper reconciles the previous existence results on canonical probabilistic belief structures and the previous non-existence results on canonical knowledge (or more general qualitative-belief) structures. It is substantively interesting because I establish the canonical representation of beliefs even when players' beliefs may not satisfy logical or introspective properties.

Applications to Qualitative Beliefs and Knowledge. My framework nests partitional (Aumann, 1976) and non-partitional possibility correspondence models of knowledge and qualitative beliefs by identifying the conditions on players' belief operators under which their beliefs are induced from a possibility correspondence. Each player's possibility correspondence associates, with each state, the set of states she considers possible at that state. While a player in a partitional model is fully introspective about what she knows and what she does not know, a player in a non-partitional

⁷Beyond the epistemic characterization of IESDA, the epistemic characterizations of extensive-form rationalizability and iterated admissibility (avoidance of weak dominance) call for a “rich” (precisely, belief-complete) belief space: see Battigalli and Siniscalchi (2002) and Brandenburger, Friedenberg, and Keisler (2008), respectively. Proposition 1 in Section 4.1 shows that the universal belief space in this paper is belief-complete.

model may, for example, fail Negative Introspection—she does not know a certain event, and she does not know that she does not know it.⁸ My framework also nests other forms of possibility correspondence models of qualitative beliefs which may fail to be truthful.⁹ I can further relax players’ logical reasoning abilities inherent in possibility correspondence models. For example, players may fail to believe logical consequences of their beliefs.

Applications to Probabilistic Beliefs. I construct a universal belief space in a way such that beliefs can be probabilistic as in type spaces (Harsanyi, 1967-68): each player has a type mapping that associates, with each state, a probability distribution on the underlying states.¹⁰ Since the framework of this paper can accommodate various forms of probabilistic beliefs through a collection of p -belief operators, the main result of this paper also asserts the existence of a universal probabilistic (e.g., countably-/finitely-/non-additive) belief space (Corollary 2). Also, this paper establishes such universal belief space with or without a common prior. Technically, my construction of a universal qualitative-belief space extends the topology-free construction of a universal (countably-additive) type space by Heifetz and Samet (1998b).

1.1 Specification of Possible Depth of Reasoning

I discuss the role that the key notion, depth of reasoning, plays on the existence of a universal belief space and its game-theoretic applications.

Previous Non-Existence Results. A *standard* belief space is a model in which any subset of underlying states Ω is an event. Heifetz and Samet (1998a) demonstrate that a universal standard partitional knowledge space generically does not exist.¹¹ Moreover, Meier (2005) shows that there is no universal standard qualitative-belief (including knowledge) space represented by a general non-partitional possibility cor-

⁸Non-partitional models are motivated in part by notions of unawareness (e.g., see Schipper (2015) for a survey). The study of non-partitional models ranges from implications of common knowledge and common belief (e.g., Agreement theorems (Aumann, 1976)) to solution concepts in game theory. See, for example, Bacharach (1985), Brandenburger, Dekel, and Geanakoplos (1992), Dekel and Gul (1997), Geanakoplos (2021), Morris (1996), Samet (1990), and Shin (1993).

⁹In the literature, knowledge is distinguished from belief in that a player can only know what is true (i.e., Truth Axiom) while she can believe something false. This paper also has a game-theoretical application of the difference between knowledge and belief. Section 6.4 shows that the epistemic characterization of iterated elimination of “Börgers dominated” actions may call for the failure of Truth Axiom. In contrast, Appendix E.3 shows that common knowledge (i.e., belief satisfies Truth Axiom) of weak-dominance rationality characterizes the iterated elimination procedure of “inferior” action profiles first introduced by Stalnaker (1994).

¹⁰The existence of a universal type/belief space is pioneered by Armbruster and Böge (1979), Böge and Eisele (1979), and Mertens and Zamir (1985).

¹¹The negative results are also obtained by Fagin (1994), Fagin et al. (1999), Fagin, Halpern, and Vardi (1991), and Heifetz and Samet (1999).

respondence.

How do my positive results reconcile with the negative results? What plays a crucial role in establishing a universal knowledge (i.e., truthful-belief) space is the specification of the domain of a knowledge space.

A κ -Algebra. To rigorously formalize the idea that the specification of the domain of a knowledge space determines players' possible depth of reasoning, let κ be an infinite cardinal number. Call a collection of subsets of underlying states Ω a κ -algebra (a shorthand for a κ -complete algebra) if it is closed under complementation and under union (and consequently intersection) of any sub-collection with cardinality less than κ . The power set of Ω is always a κ -algebra. For example, a κ -algebra subsumes an algebra of sets if κ is the least infinite cardinal \aleph_0 . A κ -algebra subsumes a σ -algebra if κ is the least uncountable cardinal \aleph_1 . Call a knowledge space (a belief space with players' beliefs truthful) a κ -knowledge space if its domain is a κ -algebra.

Specifying the domain of a knowledge space by a κ -algebra amounts to assuming that players are able to interactively reason (at least) up to the ordinal depth of κ .¹² That is, any κ -knowledge space can capture any knowledge hierarchy of the form, Alice knows that Bob knows that ..., up to the ordinal level of κ . Thus, I formally define depth of reasoning in a κ -belief space as the ordinal depth of κ (Remark 2). For example, any \aleph_0 -knowledge space can capture any finite-level interactive knowledge, because $\kappa = \aleph_0$ is the least infinite cardinal and a knowledge hierarchy up to the ordinal level of $\aleph_0 = |\{0, 1, 2, \dots\}|$ (where $|\cdot|$ denotes the cardinality of a set) consists of all finite levels of interactive knowledge. Likewise, any \aleph_1 -knowledge space can capture any countable-level knowledge hierarchy of the form, Alice knows that Bob knows that Alice knows that ..., because \aleph_1 is the least uncountable cardinal.

Main Result (Formal). The main result (Theorem 1 in Section 3) establishes that, for each fixed cardinal κ , there is a universal κ -belief space in each class of κ -belief spaces that respect some given assumptions on players' beliefs. In particular, a universal κ -knowledge space exists within a class of κ -knowledge spaces. The key idea behind the existence of the universal κ -belief space is to construct a belief space that explicitly represents all levels of interactive beliefs up to the ordinal level of κ (e.g., all countable levels of beliefs if κ is the least uncountable cardinal).

Role of κ . The construction circumvents the previous non-existence results by explicitly specifying the domain of a qualitative belief (or knowledge) space as a κ -algebra. Once I specify the domain as a κ -algebra, any κ -knowledge space can take into consideration players' interactive knowledge up to the ordinal level of κ . The universal κ -knowledge space contains all knowledge hierarchies up to the ordinal depth of κ .

For a given set S of nature states, one can arbitrarily choose κ . Thus, when S is the

¹²Technically, the ordinal depth of an infinite cardinal number κ refers to the least ordinal number which has cardinality κ . Section 2.1 provides the precise definition.

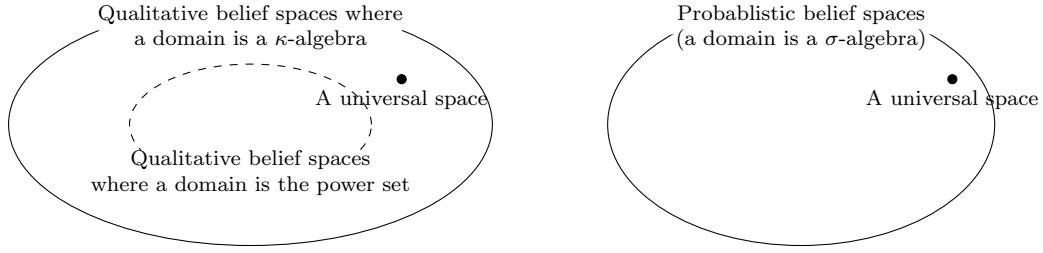


Figure 1: The Universal Belief Space within the Class of Belief Spaces: Qualitative Beliefs (Left) and Countably-Additive Probabilistic Beliefs (Right).

set of action profiles in a given strategic game, choose κ large enough: $\kappa > |S|$.¹³ The paper shows that the universal κ -belief space provides the epistemic characterization of IESDA. As the cardinality of a given strategic game increases, κ increases without end. Thus, I establish the existence of a universal κ -belief space for each κ .

Going back to the existence of a universal belief (especially, knowledge) space, I turn the previously mentioned negative results into the positive one in two ways. First, as shown in the left panel of Figure 1, I enlarge a class of knowledge spaces by allowing the domain of a knowledge space to be a κ -algebra. The left panel depicts the contrast between the previous non-existence results and the existence result (Theorem 1) by the fact that a universal space resides in the class of belief (knowledge) spaces in which a domain is a κ -algebra.

Second, I construct a universal κ -knowledge space by collecting all knowledge hierarchies of depth up to κ induced by the κ -knowledge spaces in the given class. In the left panel of Figure 2, each dot on the triangle represents a state of the universal κ -belief (knowledge) space with $\kappa = \aleph_1$. That is, each state (dot) represents a belief (knowledge) hierarchy of all countable levels. The one associated with the arrow is the belief hierarchy induced by some state of some belief (knowledge) space Ω .¹⁴ The set of belief (knowledge) hierarchies induced by the space Ω is embedded as a subset of the set of all belief (knowledge) hierarchies up to the depth of κ (the small ellipse). The set of all belief (knowledge) hierarchies up to the depth of κ is well-defined when κ is fixed (in contrast, the set of all belief (knowledge) hierarchies of an arbitrary depth is too big to be a set).

Comparison with the Probabilistic Case: Domain Specification. This paper shows

¹³In Example 1, take κ as the least uncountable cardinal number (i.e., the domain of a belief space is a σ -algebra). In Example 2, take κ as a cardinal number larger than the cardinality of $[0, 1]$ (especially, the least cardinal number larger than the cardinality of $[0, 1]$ exists). As seen in these two examples, a cardinal κ may vary with nature states S .

¹⁴Take any state of any κ -belief space. One can define the players' first-order beliefs about nature states S , their second-order beliefs, and so on, up to the ordinal level of κ . In the case of the left-panel of Figure 2, each such state induces a countable-level belief hierarchy because κ is the least uncountable cardinal number.

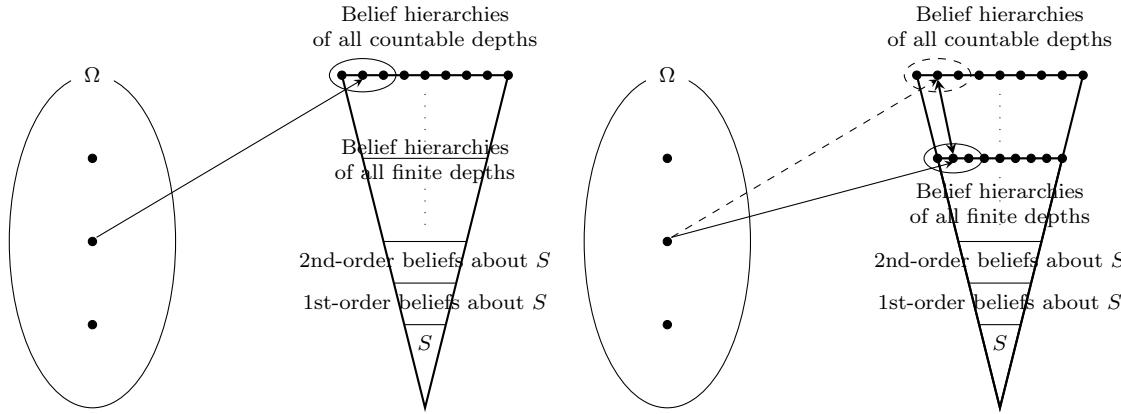


Figure 2: Construction of a Universal \aleph_1 -Belief Space: Qualitative Beliefs (Left) and Countably-Additive Probabilistic Beliefs (Right).

that the existence hinges on the specification of a domain (i.e., depth of reasoning) rather than on assumptions on beliefs themselves. Thus, the framework applies not only to qualitative belief or knowledge but also to various forms of probabilistic beliefs in a unified manner.¹⁵ Put differently, the facts that a universal probabilistic-belief space has been constructed and that a universal knowledge space has been shown not to exist reduce to whether the framework specifies depth of reasoning. As depicted in the right panel of Figure 1, the domain of a (countably-additive) type space is assumed to be a σ -algebra for the rather technical reason that a countably-additive probability measure may not necessarily be defined on the power set.¹⁶ The domain of a type space (σ -algebra) allows for capturing players' countable-level interactive probabilistic-beliefs. The domain of any \aleph_1 -qualitative-belief space (σ -algebra) also allows for capturing players' countable-level interactive qualitative-beliefs.

Comparison with the Probabilistic Case: Countable Additivity of Probabilistic Beliefs. Yet, I construct a universal qualitative \aleph_1 -belief space by collecting all belief hierarchies up to the ordinal depth of \aleph_1 (i.e., all countable-level belief hierarchies). This seems to be in contrast to the construction of a universal type space in the case of countably-additive probabilistic beliefs in the previous literature, as the universal type space consists only of finite-level belief hierarchies (beliefs about states of nature, beliefs about beliefs about states of nature, and so on, for only finitely many iterations). There, the continuity (i.e., countable additivity) of beliefs guarantees that finite-level beliefs can determine any subsequent countable level in a universal

¹⁵Appendix D.3 also discusses how the specification of a domain (i.e., depth of reasoning) sheds light on the constructions of universal unawareness, preference, and expectation spaces.

¹⁶Suppose that an outside analyst studies the strategic game in Example 1 using a (probabilistic) type space, as the set of action profiles is countable. Since the domain of the type space is a σ -algebra, the type space can coincidentally represent countable levels of interactive reasoning, which are sufficient for the unique prediction under IESDA.

type space (e.g., Fagin et al., 1999; Heifetz and Samet, 1998b). The two-way arrow in the right panel of Figure 2 illustrates that a belief hierarchy that contains all finite levels of interactive beliefs uniquely extends to the one that contains all countable levels as in the literature on type spaces (Proposition 4). In fact, in the context of finitely-additive beliefs, Meier (2006) shows, while a universal finitely-additive-belief space does not exist if all subsets are measurable (see also Fagin et al., 1999), it exists once players' beliefs are defined on a κ -algebra.

Comparison with Other Canonical Knowledge Structures. My existence result of a universal κ -knowledge space is related to the previous two positive results. First, Meier (2008) constructs a universal knowledge-belief space in which players' knowledge operators operate only on a σ -algebra on which players' probabilistic beliefs are defined. The construction in Theorem 1 of Section 3 nests Meier (2008) as a special class of \aleph_1 -knowledge(-belief) spaces under his assumptions on players' knowledge, which may not necessarily be induced from possibility correspondences. Technically, in addition to nesting his result, this paper shows the existence of a universal κ -belief space for such models as possibility correspondences of fully introspective or non-introspective knowledge and qualitative beliefs. Conceptually, this paper shows that the existence of a universal belief space is unrelated to a specific nature of beliefs and hinges rather on specifying depth of reasoning. The specification of depth of reasoning has game-theoretic applications as well. For the game in Example 1, a universal \aleph_1 -belief space can capture the players' countable levels of reasoning leading to the unique prediction under IESDA. For the game in Example 2, a universal κ -belief space can capture the players' interactive reasoning up to the ordinal level of κ which leads to the unique prediction under IESDA if a large enough $\kappa > |A_1 \times A_2|$ is chosen.

Second, Aumann (1999) constructs what he calls a canonical knowledge system, where each state of the world is a “complete and coherent” set of formulas describing finite levels of players' interactive knowledge.¹⁷

Theorem 2 in Section 4.4 reformulates a universal belief space by generalizing and modifying Aumann (1999)'s canonical knowledge system for any combination of assumptions on players' beliefs and for any domain (i.e., for any κ). In a particular case in which players with fully-introspective knowledge reason about finite levels of interactive knowledge, Theorem 2 formally proves that Aumann (1999)'s canonical space can be taken as a universal \aleph_0 -knowledge space, contrary to the conjecture of Heifetz and Samet (1998b, Section 6). Generally, Theorem 2 restates that the universal κ -belief space is the largest set (i) consisting of “complete and coherent” sets of formulas describing the players' belief hierarchies; (ii) satisfying the “(collective) coherency” condition on the entire space that induces the players' beliefs in a well-defined manner; and (iii) respecting given assumptions on their beliefs. For ex-

¹⁷Meier (2012) provides a logical axiomatization of (countably-additive probabilistic) type spaces and shows that its canonical system is a universal space, which is isomorphic to the universal type space constructed by Heifetz and Samet (1998b).

ample, when players can engage in countable-level interactive reasoning, i.e., $\kappa = \aleph_1$, a universal \aleph_1 -knowledge space, in contrast to Aumann (1999)'s canonical model, can accommodate countable levels of interactive knowledge including common knowledge, as in Example 1 (let alone Example 2 for an appropriately chosen κ). One can study an implication of common belief in rationality, instead of common knowledge of rationality, in my universal \aleph_1 -belief space.

The paper is organized as follows. Section 2 defines a belief space, properties of beliefs, and a universal belief space. Section 3 constructs a universal space (Theorem 1). Section 4 studies properties of the universal space. Especially, Section 4.4 characterizes the universal space as the “largest” set describing players' interactive beliefs in a complete and coherent manner (Theorem 2). Section 5 discusses applications to various forms of beliefs including probabilistic beliefs. Section 6 provides game-theoretic applications. Section 7 compares the existence result of a universal knowledge space with the previous non-existence results. Section 8 provides concluding remarks. Proofs are relegated to Appendix A. Supplementary Appendix (Appendices B to G), available online, provides supplementary discussions on the extensions of the existence result to richer forms of beliefs and on further game-theoretic applications.

2 Belief Spaces

Section 2.1 provides technical preliminaries. Section 2.2 defines a belief space and properties of beliefs. Section 2.3 defines a universal belief space.

2.1 Technical Preliminaries

This paper rigorously formalizes players' hierarchies of beliefs (beliefs, beliefs over beliefs, beliefs over beliefs over beliefs, etc) of transfinite levels by using ordinals (ordinal numbers) and cardinals (cardinal numbers). Ordinals are also used as a technical tool such as transfinite induction. Hence, this subsection provides technical definitions related to ordinals and cardinals.

To deal with ordinals and cardinals in the standard mathematical manner, throughout the paper, I assume the Axiom of Choice. The Axiom of Choice makes it possible to associate, with an (infinite) cardinal κ , the least ordinal $\bar{\kappa}$ with its cardinality $|\bar{\kappa}| = \kappa$. Thus, the cardinal κ is also identified with the ordinal $\bar{\kappa}$.¹⁸

Let κ be an infinite cardinal. For any given set X , call a collection \mathcal{X} of subsets of the set X a κ -complete algebra (κ -algebra, for short) (i) if $\{\emptyset, X\} \subseteq \mathcal{X}$ and (ii) if \mathcal{X} is closed under complementation and is closed under arbitrary union (and consequently intersection) of any non-empty sub-collection with cardinality less than κ . Formally, the second condition is: (i) $E \in \mathcal{X}$ implies $E^c \in \mathcal{X}$ (I denote the complement of E

¹⁸Technically, the ordinal $\bar{\kappa}$ is called the initial ordinal of κ .

by E^c or $\neg E$); and (ii) if $E \in \mathcal{X}$ for all $E \in \mathcal{E}$, where $0 < |\mathcal{E}| < \kappa$, then its union $\bigcup \mathcal{E} = \bigcup_{E \in \mathcal{E}} E$ and its intersection $\bigcap \mathcal{E} = \bigcap_{E \in \mathcal{E}} E$ belong to \mathcal{X} . When I stress the underlying space X , I also call the pair (X, \mathcal{X}) itself a κ -algebra.

For example, an \aleph_0 -algebra is an algebra of sets (i.e., a collection of subsets of X containing X itself and being closed under complementation and finite union), because \aleph_0 is the least infinite cardinal. An \aleph_1 -algebra is a σ -algebra (i.e., a collection of subsets of X containing X itself and being closed under complementation and countable union), because \aleph_1 is the least uncountable cardinal.

For an infinite cardinal κ , denote by $\mathcal{A}_\kappa(\cdot)$ the smallest κ -algebra (i.e., the intersection of all κ -algebras) including a given collection. For example, $\mathcal{A}_{\aleph_1}(\cdot) = \sigma(\cdot)$ generates the smallest \aleph_1 -algebra (i.e., σ -algebra).

I make a further technical remark on a κ -algebra. As mentioned in Meier (2006, Remark 1), when one considers a κ -algebra (X, \mathcal{X}) , it is without loss to assume the infinite cardinal κ to be regular. This is because, even if the infinite cardinal κ is not regular, (X, \mathcal{X}) is indeed a κ^+ -algebra, where the successor cardinal κ^+ (i.e., the least cardinal which is greater than κ) is known to be regular by the Axiom of Choice. Hence, if the analysts take a non-regular (i.e., singular) infinite cardinal κ , then they are implicitly taking an infinite regular cardinal κ^+ instead of κ . In this sense, while the statement and consequently the proof of the main result (i.e., Theorem 1) explicitly assume that a given infinite cardinal κ is regular, the reader (who is not familiar with the regularity of an infinite cardinal) could skip the precise definition of the regularity of an infinite regular cardinal. Finally, note that \aleph_0 and \aleph_1 are regular.

Throughout the paper, denote by I a non-empty set of players. Let S be a non-empty set of *states of nature*, endowed with a sub-collection \mathcal{S} of the power set $\mathcal{P}(S)$. An element of S is regarded as a specification of the exogenous values that are relevant to the strategic interactions among the players. For example, (S, \mathcal{S}) is the set of strategies or payoff functions endowed with a topological or measurable structure.

With an infinite regular cardinal κ fixed, I endow \mathcal{S} with a “logical” (precisely, set-algebraic) structure so that I call $E \in \mathcal{A}_\kappa(\mathcal{S})$ an *event of nature*. In the construction of a universal κ -belief space in Section 3, each $E \in \mathcal{A}_\kappa(\mathcal{S})$ plays a role of a “proposition” regarding nature states S about which players interactively reason. Hence, if E is a nature event, then so is its complement E^c ; if each $E \in \mathcal{E}$ with $0 < |\mathcal{E}| < \kappa$ is a nature event, then so are its union $\bigcup \mathcal{E}$ and its intersection $\bigcap \mathcal{E}$. For example, if (S, \mathcal{S}) is a topological space and if one studies (countably-additive) probabilistic interactive beliefs on (S, \mathcal{S}) so that $\kappa = \aleph_1$, then the collection $\mathcal{A}_{\aleph_1}(\mathcal{S})$ of nature events coincides with the generated Borel σ -algebra $\sigma(\mathcal{S})$, about which the players have beliefs.

2.2 Belief Spaces

I define a model of beliefs in which belief operators on a state space induce players’ interactive beliefs regarding nature states (S, \mathcal{S}) . I call the model a κ -belief space

when the underlying state space is a κ -algebra. Definition 1 formally defines a κ -belief space. Remark 1 discusses the role of specifying an infinite (regular) cardinal κ in a κ -belief space. Definition 2 specifies properties of beliefs.

Definition 1 (Belief Space). A κ -belief space of I on (S, \mathcal{S}) (a *belief space*, for short) is a tuple $\overrightarrow{\Omega} := \langle (\Omega, \mathcal{D}), (B_i)_{i \in I}, \Theta \rangle$ with the following three properties.

1. (Ω, \mathcal{D}) is a κ -algebra. Call Ω the set of *states* of the world (the *state space*). Call each $E \in \mathcal{D}$ an *event* (of the world). Call \mathcal{D} the *domain*.
2. For each $i \in I$, $B_i : \mathcal{D} \rightarrow \mathcal{D}$ is player i 's *belief operator*. For each $E \in \mathcal{D}$, $B_i(E)$ is the event that player i *believes* that E has occurred (player i *believes* E , for short). A player $i \in I$ *believes* an event $E \in \mathcal{D}$ at a state $\omega \in \Omega$ if $\omega \in B_i(E)$.
3. $\Theta : (\Omega, \mathcal{D}) \rightarrow (S, \mathcal{A}_\kappa(\mathcal{S}))$ is a measurable map: $\Theta^{-1}(E) \in \mathcal{D}$ for any $E \in \mathcal{S}$.

In the κ -belief space $\overrightarrow{\Omega}$, each state of the world $\omega \in \Omega$ is interpreted as a description of the corresponding nature state $\Theta(\omega) \in S$ and the players' interactive beliefs about the events at ω (e.g., player i believes an event E at ω if $\omega \in B_i(E)$, player j believes that i believes E at ω if $\omega \in B_j B_i(E)$, and so forth).

In Condition (3), since \mathcal{D} is a κ -algebra, Θ is measurable as long as $\Theta^{-1}(\mathcal{S}) \subseteq \mathcal{D}$. By this condition, any event of nature $E \in \mathcal{A}_\kappa(\mathcal{S})$ corresponds to the event $\Theta^{-1}(E) \in \mathcal{D}$ in the κ -belief space $\overrightarrow{\Omega}$.

For instance, when S is the set of action profiles of a strategic game, the mapping Θ is a profile of the players' strategy choices: it associates, with each state of the world ω , the corresponding action profile $\Theta(\omega) \in S$. For a set of action profiles $E \in \mathcal{S}$, $\Theta^{-1}(E)$ corresponds to the event that the players' actions are in E . At a state ω , player i believes that the players' actions are in E if $\omega \in B_i(\Theta^{-1}(E))$, player j believes i believes the players' actions are in E if $\omega \in B_j B_i(\Theta^{-1}(E))$, and so forth.

Moreover, Θ^{-1} preserves any set-algebraic ("logical") operations such as complementation (i.e., negation), intersection (i.e., conjunction), and union (i.e., disjunction) in $\mathcal{A}_\kappa(\mathcal{S})$ to the domain \mathcal{D} . For instance, the event $\Theta^{-1}(E^c)$ that the players' actions are in E^c equals to the event $(\Theta^{-1}(E))^c$ that the players' actions are not in E .

A standard partitional model assumes any subset of underlying states Ω to be an event, i.e., $\mathcal{D} = \mathcal{P}(\Omega)$. In contrast, Definition 1 explicitly specifies the domain to be a κ -algebra. The following remark illustrates that, when the domain is a κ -algebra, one can always introduce players' interactive beliefs about nature states up to the ordinal level $\bar{\kappa}$.

Remark 1 (Role of κ in a κ -Belief Space). For simplicity, assume $I = \{1, 2\}$ and consider mutual beliefs held by the two players. Denote by $B_I(\cdot) := B_1(\cdot) \cap B_2(\cdot)$ the mutual belief operator. For any event $F \in \mathcal{D}$, $B_I(F)$ is the event that players 1 and 2 believe F . Especially, take $F = \Theta^{-1}(E)$, where $E \in \mathcal{S}$.

When κ is the least infinite cardinal \aleph_0 , one can always introduce any finite-level mutual beliefs such as: the event $B_I(F)$ that players 1 and 2 believe F (i.e., the players' first-order beliefs about nature event E), the event $B_I^2(F) = B_I B_I(F)$ that players 1 and 2 believe that they believe F (i.e., the second-order beliefs of the players about their (first-order) beliefs about nature event E), and so forth, for any finite level $n < \aleph_0$ where \aleph_0 corresponds to the least infinite ordinal.

When κ is the least uncountable cardinal, one can always introduce any countable-level mutual beliefs such as the event $\bigcap_{n \in \mathbb{N}} B_I^n(F)$, where $B_I^n(F) = B_I(B_I^{n-1}(F))$: players 1 and 2 believe they believe ... they believe F (for any n times).¹⁹

Within the class of κ -belief spaces for a fixed κ , one can always introduce mutual beliefs of the players up to the ordinal level $\bar{\kappa}$. More specifically, one can define a chain of mutual beliefs B_I^α for any non-zero ordinal α with $|\alpha| < \kappa$ as follows:

$$B_I^\alpha(\cdot) := \begin{cases} B_I(B_I^\beta(\cdot)) & \text{if } \alpha = \beta + 1 \text{ for some } \beta \\ \bigcap_{\beta < \alpha} B_I^\beta(\cdot) & \text{if } \alpha \text{ is a limit ordinal} \end{cases} \quad .^{20}$$

Section 3 constructs the universal κ -belief space in a way so that the domain of the universal κ -belief space turns out to be a κ -algebra generated by events corresponding to nature states and players' interactive beliefs up to the ordinal level $\bar{\kappa}$. Game-theoretic applications in Section 6 also suggests the importance of explicitly specifying the ordinal level $\bar{\kappa}$ up to which players can interactively reason.

Moreover, in the context of probabilistic beliefs, the domain is usually specified as a σ -algebra. In fact, the specification of the domain allows for treating both knowledge and probabilistic beliefs on a κ -algebra (primarily, σ -algebra) at the same time.

Next, I define properties of qualitative beliefs. Theorem 1 (in Section 3) constructs a universal κ -belief space within a given class of κ -belief spaces satisfying an arbitrary combination of properties specified below. As will be seen, the following list of properties covers various classes of possibility correspondence models of (introspective/non-introspective) knowledge and qualitative beliefs. For example, if a belief operator B_i in a κ -belief space is induced by a partition, then B_i satisfies all the properties below (the converse also holds with some redundancies).²¹

Definition 2 (Properties of Beliefs). Let $\langle(\Omega, \mathcal{D}), (B_i)_{i \in I}, \Theta\rangle$ be a κ -belief space. Fix $i \in I$.

1. Monotonicity: $B_i(E) \subseteq B_i(F)$ for any $E, F \in \mathcal{D}$ with $E \subseteq F$.
2. Necessitation: $B_i(\Omega) = \Omega$.

¹⁹Note that this corresponds to the iterative definition of common belief. See Section 6.1.

²⁰Any ordinal α is either a successor ordinal (i.e., $\alpha = \beta + 1$ for some ordinal β) or a limit ordinal.

²¹Theorem 1 extends to a class of κ -belief spaces in which belief operators satisfy general set-theoretic properties (given in Lemma A.1 in Appendix A.1) beyond Definition 2. Using this result, Section 5 constructs a universal probabilistic-belief space for various notions of probabilistic beliefs.

3. λ -Conjunction (λ is a fixed infinite cardinal with $\lambda \leq \kappa$): $\bigcap_{E \in \mathcal{E}} B_i(E) \subseteq B_i(\bigcap \mathcal{E})$ for any $\mathcal{E} \in \mathcal{P}(\mathcal{D})$ with $0 < |\mathcal{E}| < \lambda$.
4. The Kripke property: for each $(\omega, E) \in \Omega \times \mathcal{D}$, $\omega \in B_i(E)$ if(f) $E \supseteq b_{B_i}(\omega) := \bigcap\{E \in \mathcal{D} \mid \omega \in B_i(E)\}$.
5. Consistency: $B_i(E) \subseteq (\neg B_i)(E^c)$ for any $E \in \mathcal{D}$.
6. Truth Axiom: $B_i(E) \subseteq E$ for any $E \in \mathcal{D}$.
7. Positive Introspection: $B_i(\cdot) \subseteq B_i B_i(\cdot)$.
8. Negative Introspection: $(\neg B_i)(\cdot) \subseteq B_i(\neg B_i)(\cdot)$.

First, Monotonicity states that if a player believes some event then she believes any of its logical consequences. Second, Necessitation means that a player believes any form of tautology such as $E \cup E^c$ expressed as Ω . Third, λ -Conjunction says that a player believes any conjunction of events (with cardinality less than λ) if she believes each event. For example, if $B_i(E)$ denotes the event that player i believes E with probability one (assume countable additivity), then B_i satisfies Monotonicity, Necessitation, and λ -Conjunction for $\lambda = \aleph_1$ but not necessarily for $\lambda > \aleph_1$.

Fourth, the Kripke property is the condition under which B_i is induced from the possibility correspondence $b_{B_i} : \Omega \rightarrow \mathcal{P}(\Omega)$.²² The information (or possibility) set $b_{B_i}(\omega) = \{\omega' \in \Omega \mid \omega' \in E \text{ for all } E \in \mathcal{D} \text{ with } \omega \in B_i(E)\}$ consists of states i considers possible at ω . The Kripke property implies Monotonicity, Necessitation, and κ -Conjunction.

Fifth, Consistency means that, if a player believes an event E then she does not believe its negation E^c . Probability-one belief satisfies Consistency, assuming additivity. Sixth, Truth Axiom says that a player can only “know” what is true. That is, if player i “knows” an event E at state ω then E is true at ω (i.e., $\omega \in E$). Truth Axiom distinguishes belief and knowledge, because a player may believe an event E at ω even if E is not true at ω . Truth Axiom implies Consistency. Seventh, Positive Introspection states that if a player believes some event then she believes that she believes it. Eighth, Negative Introspection states that if a player does not believe some event then she believes that she does not believe it.

Three remarks are in order. First, one can assume different properties of beliefs for different players. Players may also have multiple kinds of “belief” operators.²³

²²If there is $b : \Omega \rightarrow \mathcal{P}(\Omega)$ with $B_i(E) = \{\omega \in \Omega \mid b(\omega) \subseteq E\}$ for all $E \in \mathcal{D}$, then B_i satisfies the Kripke property, i.e., $B_i(E) = \{\omega \in \Omega \mid b_{B_i}(\omega) \subseteq E\}$ for all E (the converse trivially holds). See Fukuda (2019, Remark 1).

²³For instance, one can introduce both knowledge and belief by extending the set of players to $\{0, 1\} \times I$, where player i ’s knowledge operator (which satisfies Truth Axiom) is $K_i := B_{(0,i)}$ while her qualitative-belief operator is $B_i := B_{(1,i)}$. See Appendix D.2.

Second, my framework nests possibility correspondence models. Assume the Kripke property. The other properties can be expressed in terms of the possibility correspondence. First, B_i satisfies Consistency if and only if (iff, for short) b_{B_i} is serial (i.e., $b_{B_i}(\cdot) \neq \emptyset$). Second, B_i satisfies Truth Axiom iff b_{B_i} is reflexive (i.e., $\omega \in b_{B_i}(\omega)$ for all $\omega \in \Omega$). Third, B_i satisfies Positive Introspection iff b_{B_i} is transitive (i.e., $\omega' \in b_{B_i}(\omega)$ implies $b_{B_i}(\omega') \subseteq b_{B_i}(\omega)$). Fourth, B_i satisfies Negative Introspection iff b_{B_i} is Euclidean (i.e., $\omega' \in b_{B_i}(\omega)$ implies $b_{B_i}(\omega) \subseteq b_{B_i}(\omega')$). Thus, b_{B_i} forms a partition iff B_i satisfies Truth Axiom, Positive Introspection, and Negative Introspection (note that Negative Introspection and Truth Axiom imply Positive Introspection). Likewise, one can capture non-partitional models (see footnote 8): b_{B_i} is reflexive and transitive iff B_i satisfies Truth Axiom and Positive Introspection. Also, one can capture qualitative beliefs: b_{B_i} is serial, transitive, and Euclidean iff B_i satisfies Consistency, Positive Introspection, and Negative Introspection. In this way, one can identify various classes of possibility correspondence models on a κ -algebra (Ω, \mathcal{D}) . For each of such classes of possibility correspondence models, Theorem 1 in Section 3 implies the existence of a universal possibility correspondence model in the class.

Third, in order to accommodate Truth Axiom, the state space (Ω, \mathcal{D}) may not necessarily be assumed to be the product κ -algebra of the nature states $(S, \mathcal{A}_\kappa(\mathcal{S}))$ and the players' type sets $((T_i, \mathcal{T}_i))_{i \in I}$ (all of which form a κ -algebra). If player i 's beliefs depend only on her own types, then B_i would violate Truth Axiom.

2.3 A Terminal Belief Space

The previous subsection has defined belief spaces (Definition 1) and properties of beliefs (Definition 2). Any possible combination of properties of beliefs in Definition 2 and an infinite regular cardinal κ determine the class of κ -belief spaces that respect the given properties of beliefs. Theorem 1 in Section 3 constructs a universal κ -belief space in such a class of κ -belief spaces.

To that end, Definition 4 defines a universal belief space in a given class of belief spaces as a terminal belief space in the class. It is a belief space to which every belief space in the given class is uniquely mapped in a belief-preserving manner. I start by formalizing the notion of a belief-preserving map, a belief morphism.

Definition 3 (Belief Morphism). Let $\vec{\Omega} = \langle (\Omega, \mathcal{D}), (B_i)_{i \in I}, \Theta \rangle$ and $\vec{\Omega}' = \langle (\Omega', \mathcal{D}'), (B'_i)_{i \in I}, \Theta' \rangle$ be belief spaces of a given class. A (belief) morphism $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}'$ is a measurable map $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$ satisfying: (i) $\Theta = \Theta' \circ \varphi$; and (ii) $B_i(\varphi^{-1}(E')) = \varphi^{-1}(B'_i(E'))$ for each $(i, E') \in I \times \mathcal{D}'$.

The morphism φ associates, with each state $\omega \in \Omega$, the corresponding state $\varphi(\omega) \in \Omega'$ with the two conditions. Condition (i) requires that the same nature state prevail between two states ω and $\varphi(\omega)$. Condition (ii) states that the players' beliefs are preserved from one space to another: player i believes an event E' at $\varphi(\omega)$ iff she believes $\varphi^{-1}(E')$ at ω .

For any belief space $\vec{\Omega}$, the identity map id_Ω on Ω is a morphism from $\vec{\Omega}$ into itself. Denote by $\text{id}_\Omega : \vec{\Omega} \rightarrow \vec{\Omega}$ the identity (belief) morphism. Next, call a belief morphism $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}'$ a *(belief) isomorphism*, if there is a morphism $\psi : \vec{\Omega}' \rightarrow \vec{\Omega}$ with $\psi \circ \varphi = \text{id}_\Omega$ and $\varphi \circ \psi = \text{id}_{\Omega'}$ (that is, φ is bijective and its inverse φ^{-1} is a morphism). If φ is an isomorphism then its inverse φ^{-1} is unique. Belief spaces $\vec{\Omega}$ and $\vec{\Omega}'$ are *isomorphic*, if there is an isomorphism $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}'$.

Now, I define a terminal belief space. It “includes” all belief spaces in that any belief space can be mapped to the terminal space by a unique morphism.

Definition 4 (Terminal Belief Space). Fix a class of κ -belief spaces of I on (S, \mathcal{S}) . A κ -belief space $\vec{\Omega}^*$ in the class is *terminal* if, for any κ -belief space $\vec{\Omega}$ in the class, there is a unique morphism $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}^*$.

Fix a non-empty set of players I , a space of nature states (S, \mathcal{S}) , an infinite cardinal κ , and assumptions on the players’ beliefs. Then, the given class of κ -belief spaces of I on (S, \mathcal{S}) forms a *category*, where a belief space $\vec{\Omega}$ is an *object* and a belief morphism is a *morphism*. In the language of category theory, a terminal κ -belief space in the class is a terminal object in the category of belief spaces. As is well known in category theory, a terminal belief space (in a given class) is unique up to belief isomorphism.

3 Construction of a Terminal Belief Space

Throughout this section, fix a category of κ -belief spaces of I on (S, \mathcal{S}) that satisfy some given properties of beliefs in Definition 2, where κ is an infinite regular cardinal. A belief space refers to a κ -belief space of I on (S, \mathcal{S}) in the given category.

I construct a terminal belief space by employing the “expressions-descriptions” approach (Heifetz and Samet, 1998b; Meier, 2006, 2008). The construction in this section demonstrates that the existence of the terminal κ -belief space hinges on the specification of the infinite regular cardinal κ , which determines possible “depth of reasoning,” rather than on properties of beliefs themselves.

The construction of a terminal belief space consists of six steps.²⁴ The first step is to inductively define *expressions*, syntactic formulas defined solely in terms of nature and interactive beliefs about nature.²⁵ Informally, in the terminal type space construction (e.g., Brandenburger and Dekel, 1993; Mertens and Zamir, 1985), the counterpart to this step is to inductively construct players’ belief hierarchies (their first-order beliefs about nature states, their second-order beliefs about nature states and the first-order beliefs, and so on) from the nature states $(S, \sigma(\mathcal{S}))$. Since any nature event is an object of beliefs, I treat any nature event $E \in \mathcal{A}_\kappa(\mathcal{S})$ as a proposition

²⁴While the explanations in the main text are self-contained, Appendix B provides graphical illustrations which intend to supplement the explanations of these six steps.

²⁵Technically, the inductive definition uses transfinite induction.

and call it an expression. Since objects of beliefs are closed under conjunction, negation, and the players' beliefs, I define the corresponding syntactic (not set-theoretic) operations for expressions. For a set of expressions \mathcal{E} , let $(\bigwedge \mathcal{E})$ be a (syntactic) expression denoting the conjunction of expressions \mathcal{E} . For an expression e , let $(\neg e)$ be the (syntactic) expression denoting the negation of e , and let $(\beta_i(e))$ be the (syntactic) expression denoting that player i believes e .²⁶ Formally:

Definition 5 (Expressions). Let λ be an infinite regular cardinal with $\lambda \leq \kappa$. The set of all λ -expressions $\mathcal{L}_\lambda^I(\mathcal{A}_\kappa(\mathcal{S}))$ is the smallest set satisfying the following.

1. Nature: Every nature event $E \in \mathcal{A}_\kappa(\mathcal{S})$ is a λ -expression.
2. Conjunction: If \mathcal{E} is a set of λ -expressions with $0 < |\mathcal{E}| < \lambda$, then so is $(\bigwedge \mathcal{E})$.
3. Negation: If e is a λ -expression then so is $(\neg e)$.
4. Beliefs: If e is a λ -expression, then so is $(\beta_i(e))$ for each $i \in I$.

For $\lambda = \kappa$, call each κ -expression an expression, and denote $\mathcal{L} = \mathcal{L}_\kappa^I(\mathcal{A}_\kappa(\mathcal{S}))$.

Before providing technical discussions about Definition 5, I define the depth of an expression.

Remark 2 (Depth of an Expression). The set \mathcal{L} is a *language* (as a terminology in logic) that represents any form of interactive beliefs regarding nature states (i.e., player i 's beliefs about nature $\mathcal{A}_\kappa(\mathcal{S})$, player i 's beliefs about player j 's beliefs about nature, and so on) up to the ordinal depth $\bar{\kappa}$.

Formally, one can define the *depth* of a λ -expression e , $\text{depth}(e)$, as an ordinal in the following way. First, any nature event $E \in \mathcal{A}_\kappa$ has depth 0: $\text{depth}(E) := 0$. Second, for any collection of λ -expressions \mathcal{E} with $0 < |\mathcal{E}| < \lambda$, the depth of $\bigwedge \mathcal{E}$ is defined as $\text{depth}(\bigwedge \mathcal{E}) := \sup_{e \in \mathcal{E}} \text{depth}(e)$. Third, the depth of $(\neg e)$ is equal to that of e : $\text{depth}(\neg e) := \text{depth}(e)$. Fourth, the depth of $\beta_i(e)$ is defined as $\text{depth}(\beta_i(e)) := \text{depth}(e) + 1$ for any $i \in I$. By construction, (i) any λ -expression e has $\text{depth}(e) < \bar{\lambda}$; and (ii) for any ordinal $\alpha < \bar{\lambda}$, there exists a λ -expression e with $\text{depth}(e) = \alpha$.²⁷

With this definition in mind, the infinite regular cardinal κ determines the players' possible *depth of reasoning* $\bar{\kappa}$.

If, for instance, κ is the least infinite cardinal \aleph_0 , then the set $\mathcal{L} = \mathcal{L}_\kappa^I(\mathcal{A}_\kappa(\mathcal{S}))$ can capture any finite-level interactive beliefs in the sense that the set contains finite-depth expressions. Thus, the set \mathcal{L} contains such expressions as $(\beta_{i_1}\beta_{i_2} \cdots \beta_{i_n})(e)$ (i.e., the expression that represents "player i_1 believes that i_2 believes that ... i_n believes

²⁶Thus, for instance, for a nature event $E \in \mathcal{A}_\kappa(\mathcal{S})$, I define player i 's belief in E (denoted $\beta_i(E)$) not as an element of the set $\mathcal{A}_\kappa(\mathcal{S})$ but as a syntactic formula.

²⁷For instance, when $\lambda = \aleph_0$, any λ -expression has finite depth; and for any finite n there exists a λ -expression with depth n .

$e \in \mathcal{A}_\kappa(\mathcal{S})$ ” for any finite number $n \in \mathbb{N}$. The set \mathcal{L} , however, does not contain the infinite conjunction of such expressions.

In contrast, if κ is the least uncountable cardinal \aleph_1 , then the set $\mathcal{L} = \mathcal{L}_\kappa^I(\mathcal{A}_\kappa(\mathcal{S}))$ can capture any countable-level interactive beliefs in the sense that the set contains at-most-countable-depth expressions. Thus, the set \mathcal{L} contains such expressions as $\bigwedge_{n \in \mathbb{N}} (\beta_{i_1} \beta_{i_2} \cdots \beta_{i_n})(e)$ (i.e., the expression that represents “player i_1 believes that i_2 believes that ... i_n believes $e \in \mathcal{A}_\kappa(\mathcal{S})$ for any $n \in \mathbb{N}$ ”).

Three technical remarks on Definition 5 are in order. First, since κ is fixed, for each λ , the (smallest) set $\mathcal{L}_\lambda^I(\mathcal{A}_\kappa(\mathcal{S}))$ is well-defined by induction. Remark 3 below also implies that \mathcal{L} and consequently $\mathcal{L}_\lambda^I(\mathcal{A}_\kappa(\mathcal{S}))$ are well-defined sets.

Second, I consider κ -expressions beyond $\lambda = \aleph_0$ because, when it comes to qualitative beliefs, a belief hierarchy of all finite-level beliefs may not necessarily uniquely pin down the corresponding belief hierarchy of all countable-level beliefs. This contrasts with the literature on probabilistic type spaces.²⁸

Third, for ease of notation, I often add or omit parentheses in denoting expressions (and in other occurrences). If \mathcal{E} is a set of expressions with $0 < |\mathcal{E}| < \kappa$, then let $(\bigvee \mathcal{E}) := \neg(\bigwedge \{(\neg e) \in \mathcal{L} \mid e \in \mathcal{E}\})$. Thus, $\bigvee \mathcal{E}$ denotes the disjunction of \mathcal{E} . Also, I interchangeably denote, for instance, $e_1 \wedge e_2 = \bigwedge \{e_1, e_2\}$ and $e_1 \vee e_2 = \bigvee \{e_1, e_2\}$. I interchangeably denote $\bigwedge_{j \in J} e_j = \bigwedge \{e_j \mid j \in J\}$ and $\bigvee_{j \in J} e_j = \bigvee \{e_j \mid j \in J\}$ when expressions are indexed by some set J . Denote the implication $(e \rightarrow f) := ((\neg e) \vee f)$ and the equivalence $(e \leftrightarrow f) := ((e \rightarrow f) \wedge (f \rightarrow e))$.

The following remark shows how the set of expressions \mathcal{L} is inductively generated from the nature states $(S, \mathcal{A}_\kappa(\mathcal{S}))$ in $\bar{\kappa}$ steps.

Remark 3 (Restatement of Expressions \mathcal{L}). Let λ be an infinite regular cardinal with $\lambda \leq \kappa$. The following auxiliary ordinal sequence $(\mathcal{L}_\alpha)_{\alpha=0}^{\bar{\kappa}}$ generates the set of expressions $\mathcal{L}_\lambda^I(\mathcal{A}_\kappa(\mathcal{S})) = \mathcal{L}_{\bar{\kappa}}$. In particular, if $\lambda = \kappa$ then $\mathcal{L} = \mathcal{L}_{\bar{\kappa}}$. Namely, let $\mathcal{L}_0 := \mathcal{A}_\kappa(\mathcal{S})$. For any ordinal α with $0 < \alpha \leq \bar{\kappa}$, define

$$\begin{aligned} \mathcal{L}_\alpha &:= \mathcal{L}'_\alpha \cup \{(\neg e) \mid e \in \mathcal{L}'_\alpha\} \cup \{\bigwedge \mathcal{F} \mid \mathcal{F} \subseteq \mathcal{L}'_\alpha \text{ and } 0 < |\mathcal{F}| < \lambda\}, \text{ where} \\ \mathcal{L}'_\alpha &:= (\bigcup_{\beta < \alpha} \mathcal{L}_\beta) \cup \bigcup_{i \in I} \{\beta_i(e) \mid e \in \bigcup_{\beta < \alpha} \mathcal{L}_\beta\}. \end{aligned}$$

In the remark, $(\mathcal{L}_\alpha)_{\alpha=0}^{\bar{\kappa}}$ is an increasing sequence of sets of expressions such that, starting from nature events $\mathcal{L}_0 = \mathcal{A}_\kappa(\mathcal{S})$, each set \mathcal{L}_α contains interactive beliefs about expressions $\bigcup_{\beta < \alpha} \mathcal{L}_\beta$. For instance, the set \mathcal{L}_1 contains the players’ first-order beliefs about nature events $\beta_i(e)$. The set \mathcal{L}_2 contains their second-order beliefs about nature events $\beta_i \beta_j(e)$. For the least infinite ordinal α , the set \mathcal{L}_α contains any finite-depth expression of the form $(\beta_{i_1} \beta_{i_2} \cdots \beta_{i_n})(e)$. For $\alpha = \bar{\kappa}$, $\mathcal{L} = \mathcal{L}_{\bar{\kappa}}$ consists of expressions

²⁸Proposition 4 in Section 5.1 shows that, for countably-additive probabilistic beliefs on \aleph_1 -algebras, \aleph_0 -expressions $\mathcal{L}_{\aleph_0}^I(\mathcal{A}_{\aleph_1}(\mathcal{S}))$ suffice to capture countable-level interactive beliefs, consistently with the construction of a terminal type space in the literature.

of depth less than $\bar{\kappa}$, i.e., logical formulas expressing interactive beliefs regarding $(S, \mathcal{A}_\kappa(\mathcal{S}))$ up to depth $\bar{\kappa}$.²⁹

While expressions themselves are defined independently of any particular belief space, for any belief space $\vec{\Omega}$, I inductively identify each expression e with the corresponding event $\llbracket e \rrbracket_{\vec{\Omega}} \in \mathcal{D}$ so that $\llbracket e \rrbracket_{\vec{\Omega}}$ is the set of states of the world in which the expression e holds. Recalling the inductive definition in Definition 5, I start with an expression $E \in \mathcal{A}_\kappa(\mathcal{S})$. Let $\llbracket E \rrbracket_{\vec{\Omega}} := \Theta^{-1}(E) \in \mathcal{D}$ be the set of states at which $E \in \mathcal{A}_\kappa(\mathcal{S})$ is true. If $\llbracket e \rrbracket_{\vec{\Omega}}$ is defined as an event in \mathcal{D} for some expression e , then for the expression $\beta_i(e)$ (“ i believes e ”), let $\llbracket \beta_i(e) \rrbracket_{\vec{\Omega}} := B_i(\llbracket e \rrbracket_{\vec{\Omega}})$ be the set of states (in $\vec{\Omega}$) at which player i believes expression e . Thus, I inductively define the mapping $\llbracket \cdot \rrbracket_{\vec{\Omega}}$, which associates, with each expression $e \in \mathcal{L}$, the event $\llbracket e \rrbracket_{\vec{\Omega}} \in \mathcal{D}$ in $\vec{\Omega}$ that expression e holds. Following the terminology in logic, I refer to the mapping $\llbracket \cdot \rrbracket_{\vec{\Omega}}$ as the *semantic interpretation function* of $\vec{\Omega}$. Formally:

Definition 6 (Expressions Identified as Events). Fix a κ -belief space $\vec{\Omega}$. Inductively define the map $\llbracket \cdot \rrbracket_{\vec{\Omega}} : \mathcal{L} \rightarrow \mathcal{D}$, the semantic interpretation function of $\vec{\Omega}$, as follows.

1. Nature: $\llbracket E \rrbracket_{\vec{\Omega}} := \Theta^{-1}(E)$ for every $E \in \mathcal{A}_\kappa(\mathcal{S})$.
2. Conjunction: $\llbracket \bigwedge \mathcal{E} \rrbracket_{\vec{\Omega}} := \bigcap_{e \in \mathcal{E}} \llbracket e \rrbracket_{\vec{\Omega}}$ for any $\mathcal{E} \in \mathcal{P}(\mathcal{L})$ with $0 < |\mathcal{E}| < \kappa$.
3. Negation: $\llbracket \neg e \rrbracket_{\vec{\Omega}} := (\llbracket e \rrbracket_{\vec{\Omega}})^c$ for each expression e .
4. Beliefs: $\llbracket \beta_i(e) \rrbracket_{\vec{\Omega}} := B_i(\llbracket e \rrbracket_{\vec{\Omega}})$ for each $i \in I$ and expression e .

The semantic interpretation function $\llbracket \cdot \rrbracket_{\vec{\Omega}}$ of a given belief space is, by induction, uniquely extended from Θ^{-1} . Remark 4 below states that, by induction, a morphism preserves the meaning of an expression. Namely, let $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}'$ be a morphism. Then, an expression e holds at ω (i.e., $\omega \in \llbracket e \rrbracket_{\vec{\Omega}}$) iff it holds at $\varphi(\omega)$ (i.e., $\varphi(\omega) \in \llbracket e \rrbracket_{\vec{\Omega}'}$). Formally:

Remark 4 (Morphism Preserves the Meaning of an Expression). If $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}'$ is a morphism, then $\llbracket \cdot \rrbracket_{\vec{\Omega}} = \varphi^{-1}(\llbracket \cdot \rrbracket_{\vec{\Omega}'})$.

The second step is to define *descriptions* by the set of expressions and the nature state that obtain at each state of each belief space. Each description turns out to be a state of the terminal space. That is, each state in the terminal space describes the nature state and the set of expressions that hold at some state of some belief space. In defining a description, observe that nature states and expressions reside in different spaces. Thus, I define a description to be a subset of the disjoint union

$$S \sqcup \mathcal{L} := \{(0, s) \in \{0\} \times S \mid s \in S\} \cup \{(1, e) \in \{1\} \times \mathcal{L} \mid e \in \mathcal{L}\}.$$

²⁹It can be seen that, for any successor ordinal α , any expression in \mathcal{L}_α has depth weakly less than α and that, for any non-zero limit ordinal α , any expression in \mathcal{L}_α has depth less than α .

While this definition of the description is different from the one in the previous literature, this definition uniquely identifies the corresponding nature state for each description without any condition on (S, \mathcal{S}) .³⁰ Formally:

Definition 7 (Descriptions). For any belief space $\vec{\Omega}$ and $\omega \in \Omega$, define the description $D(\omega)$ of ω by $D(\omega) := \{\Theta(\omega)\} \sqcup \{e \in \mathcal{L} \mid \omega \in \llbracket e \rrbracket_{\vec{\Omega}}\}$.

Each description $D(\omega) = \{(0, \Theta(\omega))\} \cup \{(1, e) \in \{1\} \times \mathcal{L} \mid \omega \in \llbracket e \rrbracket_{\vec{\Omega}}\}$ contains the unique nature state $\Theta(\omega) \in S$ associated with ω and the expressions e which are true at ω .³¹ For ease of notation, write $s \in_0 D(\omega)$ for $(0, s) \in D(\omega)$. Also, write $e \in_1 D(\omega)$ for $(1, e) \in D(\omega)$. The reader could even read “ $s \in D(\omega)$ ” and “ $e \in D(\omega)$ ” by disregarding the subscripts 0 and 1.

Remark 5 (Belief Hierarchy of Depth up to $\bar{\kappa}$). One can interpret each $D(\omega)$ as a set of expressions that represents each player’s *belief hierarchy* (her belief about nature $\mathcal{A}_\kappa(\mathcal{S})$, her belief about the players’ beliefs about nature, and so on) of depth up to $\bar{\kappa}$ that holds at ω (together with the state of nature $\Theta(\omega)$). That is, for any expression $e \in \mathcal{L}$ (which, by construction, has depth less than $\bar{\kappa}$) and for any player $i \in I$, exactly one of the following holds: $\beta_i(e) \in_1 D(\omega)$ (if player i believes $\llbracket e \rrbracket_{\vec{\Omega}}$ at ω) or $(\neg\beta_i)(e) \in_1 D(\omega)$ (if she does not believe $\llbracket e \rrbracket_{\vec{\Omega}}$ at ω).³² In this sense, $D(\omega)$ represents the belief hierarchy (precisely, the players’ belief hierarchies) at ω , and this second step is analogous to defining belief hierarchies induced by some type of some player in the construction of a terminal type space.³³

Descriptions have two roles in constructing a terminal belief space. First, I will construct the terminal belief space so that the underlying states Ω^* consist of all descriptions of states of the world ranged over all belief spaces in the given category:

$$\Omega^* := \{\omega^* \in \mathcal{P}(S \sqcup \mathcal{L}) \mid \omega^* = D(\omega) \text{ for some } \vec{\Omega} \text{ and } \omega \in \Omega\}. \quad (1)$$

³⁰Meier (2006, 2008) assumes the following “separative” condition on (S, \mathcal{S}) : for any distinct $s, s' \in S$, there is $E \in \mathcal{S}$ with $(s \in E \text{ and } s' \notin E) \text{ or } (s' \in E \text{ and } s \notin E)$ (i.e., there is $E \in \mathcal{A}_\kappa(\mathcal{S})$ with $s \in E$ and $s' \notin E$). Then, $\{s\} = \bigcap\{E \in \mathcal{A}_\kappa(\mathcal{S}) \mid s \in E\}$ for each $s \in S$, though it may be the case that $\{s\} \notin \mathcal{A}_\kappa(\mathcal{S})$. Moss and Viglizzo (2004, 2006) also impose the separative condition.

³¹Note that $(0, s) \in D(\omega)$ indicates which nature event belongs to $D(\omega)$ in the following sense: for any $E \in \mathcal{A}_\kappa(\mathcal{S})$, $(1, E) \in D(\omega)$ iff $s \in E$. For ease of exposition, I simply include all expressions including nature events which are true at ω in the description $D(\omega)$.

³²More specifically, the players’ first-order beliefs are incorporated in $D(\omega)$ as follows: for any nature event $E \in \mathcal{A}_\kappa(\mathcal{S})$, player i believes $\llbracket E \rrbracket_{\vec{\Omega}}$ at ω (i.e., $\omega \in B_i \llbracket E \rrbracket_{\vec{\Omega}}$) iff $\beta_i(E) \in_1 D(\omega)$. For their second-order beliefs, player i believes j believes $\llbracket E \rrbracket_{\vec{\Omega}}$ at ω (i.e., $\omega \in B_i B_j \llbracket E \rrbracket_{\vec{\Omega}}$) iff $\beta_i \beta_j(E) \in_1 D(\omega)$. For any $n \in \mathbb{N}$, player i_1 believes i_2 believes \dots i_n believes $\llbracket E \rrbracket_{\vec{\Omega}}$ at ω (i.e., $\omega \in (B_{i_1} B_{i_2} \dots B_{i_n}) \llbracket E \rrbracket_{\vec{\Omega}}$) iff $(\beta_{i_1} \beta_{i_2} \dots \beta_{i_n})(E) \in_1 D(\omega)$. Thus, one can interpret the description $D(\omega)$ as the players’ belief hierarchies of depth up to $\bar{\kappa}$ at ω .

³³In fact, since any two terminal spaces are isomorphic, when the construction of a terminal space in this section is applied to the case with probabilistic beliefs (see Section 5), the description $D(\omega)$ is formally in a one-to-one relation with the profile of the corresponding nature state and the players’ belief hierarchies (in the form of a sequence of probability distributions) at ω .

Second, I regard D as a mapping $D : \Omega \rightarrow \Omega^*$ (or $D_{\vec{\Omega}}$ to stress the domain of D) for any belief space $\vec{\Omega}$, and I call D the *description map*. The description map D turns out to be a unique morphism.

Two technical remarks are in order. First, Ω^* is not empty because there is a belief space $\vec{\Omega}$ with $\Omega \neq \emptyset$ in the given category.³⁴

Second, Ω^* depends on the choice of a category of belief spaces. The more assumptions on beliefs one imposes, the smaller Ω^* becomes. Formally, consider two categories of belief spaces where assumptions on players' beliefs in the first are also imposed in the second. Denoting by Ω^{1*} and Ω^{2*} the spaces constructed according to Equation (1), $\Omega^{2*} \subseteq \Omega^{1*}$ holds by construction.

To show that the description map D is a unique morphism (in the sixth step), here I remark that a morphism preserves the descriptions in the following sense.

Remark 6 (Morphism Preserves Descriptions). If $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}'$ is a morphism, then $D_{\vec{\Omega}} = D_{\vec{\Omega}'} \circ \varphi$.

To see this, fix belief spaces $\vec{\Omega}$ and $\vec{\Omega}'$ and $(\omega, \omega') \in \Omega \times \Omega'$. Then, $D_{\vec{\Omega}}(\omega) = D_{\vec{\Omega}'}(\omega')$ iff (i) $\Theta(\omega) = \Theta'(\omega')$ and (ii) $\omega \in \llbracket e \rrbracket_{\vec{\Omega}}$ iff $\omega' \in \llbracket e \rrbracket_{\vec{\Omega}'}$ for all $e \in \mathcal{L}$. Thus, $D_{\vec{\Omega}}(\omega) = D_{\vec{\Omega}'}(\omega')$ means that states ω and ω' are equivalent in terms of a prevailing nature state and prevailing expressions, abstracting away from physical representations of $\vec{\Omega}$ and $\vec{\Omega}'$. By Remark 4, both conditions are met for any $(\omega, \varphi(\omega)) \in \Omega \times \Omega'$ where $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}'$ is a morphism. Thus, states ω and $\varphi(\omega)$ induce the same profile of the players' belief hierarchies together with the corresponding nature state.

I discuss two implications of Remark 6. First, call a belief space $\vec{\Omega}$ *non-redundant* (Mertens and Zamir, 1985, Definition 2.4) if its description map D is injective. In other words, for any distinct ω and ω' , either $\Theta(\omega) \neq \Theta(\omega')$ or they are separated by $\mathcal{D}_{\bar{\kappa}} := \{\llbracket e \rrbracket_{\vec{\Omega}} \in \mathcal{D} \mid e \in \mathcal{L}\}$ (i.e., there exists $\llbracket e \rrbracket_{\vec{\Omega}} \in \mathcal{D}_{\bar{\kappa}}$ such that $(\omega \in \llbracket e \rrbracket_{\vec{\Omega}} \text{ and } \omega' \notin \llbracket e \rrbracket_{\vec{\Omega}}) \text{ or } (\omega' \in \llbracket e \rrbracket_{\vec{\Omega}} \text{ and } \omega \notin \llbracket e \rrbracket_{\vec{\Omega}})$).³⁵

Second, Remark 6 implies that if $\vec{\Omega}'$ is non-redundant then there is at most one morphism from a given space $\vec{\Omega}$ into $\vec{\Omega}'$.³⁶ The sixth step shows that $D_{\vec{\Omega}} : \vec{\Omega} \rightarrow \vec{\Omega}^*$ is a unique morphism by demonstrating that $D_{\vec{\Omega}^*}$ is the identity.

The third step is to define the domain \mathcal{D}^* of the candidate terminal belief space Ω^* . Since each expression e corresponds to an object of interactive beliefs, define the set $[e]$ of descriptions that make e true (i.e., contain e) to be an event in Ω^* .

³⁴Consider a belief space $\vec{\{s\}} := \langle (\{s\}, \mathcal{P}(\{s\})), (\text{id}_{\mathcal{P}(\{s\})})_{i \in I}, \Theta \rangle$ where $s \in S$ and $\Theta : \{s\} \ni s \mapsto s \in S$. Each $B_i = \text{id}_{\mathcal{P}(\{s\})}$ satisfies all the properties of beliefs in Definition 2.

³⁵The collection $\mathcal{D}_{\bar{\kappa}}$, which consists of events $\llbracket e \rrbracket_{\vec{\Omega}}$ that correspond to expressions $e \in \mathcal{L}$ (i.e., events that are generated by states of nature and belief hierarchies), can be expressed solely by the primitives of the belief space by Definition 17 in Section 7.

³⁶*Proof.* If $\varphi, \psi : \vec{\Omega} \rightarrow \vec{\Omega}'$ are morphisms then $D_{\vec{\Omega}'} \circ \varphi = D_{\vec{\Omega}} = D_{\vec{\Omega}'} \circ \psi$. Since $D_{\vec{\Omega}}$ is injective, $\varphi = \psi$.

Formally, for each $e \in \mathcal{L}$, define the set of descriptions $[e] := \{\omega^* \in \Omega^* \mid e \in_1 \omega^*\}$. I show that $\mathcal{D}^* := \{[e] \in \mathcal{P}(\Omega^*) \mid e \in \mathcal{L}\}$ is a legitimate domain.

Lemma 1 (Domain \mathcal{D}^*). 1. $(\Omega^*, \mathcal{D}^*)$ is a κ -algebra.

2. For any belief space $\vec{\Omega}$, the description map $D : (\Omega, \mathcal{D}) \rightarrow (\Omega^*, \mathcal{D}^*)$ is a measurable map with $D^{-1}([e]) = \llbracket e \rrbracket_{\vec{\Omega}}$ for any $e \in \mathcal{L}$.

The property, $D_{\vec{\Omega}}^{-1}([\cdot]) = \llbracket \cdot \rrbracket_{\vec{\Omega}}$, exhibits duality between the semantic interpretation function $\llbracket \cdot \rrbracket_{\vec{\Omega}} : \mathcal{L} \rightarrow \mathcal{D}$ (which, by induction, is unique) and the description map $D_{\vec{\Omega}} : \vec{\Omega} \rightarrow \vec{\Omega}^*$ (which turns out to be a unique morphism) in the following sense. Through $\llbracket \cdot \rrbracket_{\vec{\Omega}}$, each expression $e \in \mathcal{L}$ that represents nature states and interactive beliefs is interpreted as an event $\llbracket e \rrbracket_{\vec{\Omega}}$ in the given space $\vec{\Omega}$. In contrast, through $D_{\vec{\Omega}}$, each state ω in the given space $\vec{\Omega}$ is re-formulated as the corresponding nature state and expressions (i.e., the description) $D_{\vec{\Omega}}(\omega)$. The property $D_{\vec{\Omega}}^{-1}([\cdot]) = \llbracket \cdot \rrbracket_{\vec{\Omega}}$ also implies that $\mathcal{D}_{\vec{\Omega}}$ is the sub- κ -algebra induced by $D_{\vec{\Omega}}$: $\mathcal{D}_{\vec{\Omega}} = \{D_{\vec{\Omega}}^{-1}([e]) \in \mathcal{D} \mid e \in \mathcal{L}\} = D_{\vec{\Omega}}^{-1}(\mathcal{D}^*)$.

Call $\vec{\Omega}$ *minimal* (Di Tillio, 2008) if $\mathcal{D}_{\vec{\Omega}} = \mathcal{D}$.³⁷ It will be shown that if a given belief space is non-redundant and minimal, then the belief space is isomorphic to a subspace of the terminal space.

The fourth step is to construct the mapping $\Theta^* : \Omega^* \rightarrow S$ that associates with each state $\omega^* \in \Omega^*$ the unique nature state s contained in ω^* (i.e., $s \in_0 \omega^*$).

Lemma 2 (Mapping Θ^*). There is a measurable map $\Theta^* : (\Omega^*, \mathcal{D}^*) \rightarrow (S, \mathcal{A}_{\kappa}(S))$ with the following properties:

1. $\Theta^*(D(\omega)) = \Theta(\omega)$ for any belief space $\vec{\Omega}$ and $\omega \in \Omega$; and
2. $(\Theta^*)^{-1}(E) = [E] \in \mathcal{D}^*$ for all $E \in \mathcal{A}_{\kappa}(S)$.

The fifth step is to introduce the players' belief operators on \mathcal{D}^* in a way such that player i believes an event $[e]$ at a state ω^* iff ω^* contains $\beta_i(e)$ (i.e., $\beta_i(e) \in_1 \omega^*$). That is, I define $B_i^* : \mathcal{D}^* \rightarrow \mathcal{D}^*$ by $B_i^*([e]) := [\beta_i(e)]$ for each $e \in \mathcal{L}$. I show that this is well-defined: if expressions e and f satisfy $[e] = [f]$, then $[\beta_i(e)] = [\beta_i(f)]$.

Lemma 3 (Belief Operators B_i^*). Fix $i \in I$.

1. B_i^* is a well-defined belief operator which inherits the properties of beliefs imposed in the given category.

³⁷Thus, the given belief space is minimal if the domain \mathcal{D} consists solely of events $\mathcal{D}_{\vec{\Omega}}$ that are generated by nature states and belief hierarchies (recall footnote 35). The notion of minimality turns out to be equivalent to what Friedenberg and Meier (2011) call *strong measurability*. See Appendix F for the characterization of minimality (strong measurability).

2. Moreover, for any belief space $\vec{\Omega}$, $D^{-1}(B_i^*([e])) = B_i(D^{-1}([e]))$ for all $[e] \in \mathcal{D}^*$.

Lemma A.1 in Appendix A.1 shows that each B_i^* inherits properties of beliefs beyond Definition 2. In other words, this paper shows the existence of a terminal belief space as long as players' beliefs are represented by their belief operators and properties of the belief operators satisfy the set-theoretic conditions of Lemma A.1 in Appendix A.1. In fact, Section 5 extends the construction of a terminal belief space to such cases as probabilistic beliefs.

Lemma A.1 implies that if there is a belief space $\vec{\Omega}$ such that B_i fails a given property with respect to some $\llbracket e \rrbracket_{\vec{\Omega}}$, then B_i^* fails that property with respect to $[e]$. Thus, B_i^* satisfies the properties of beliefs for player i that are common among all the belief spaces in the given category.³⁸

So far, $\vec{\Omega}^* = \langle (\Omega^*, \mathcal{D}^*), (B_i^*)_{i \in I}, \Theta^* \rangle$ is a belief space of the given category such that, for any belief space $\vec{\Omega}$, the description map $D : \Omega \rightarrow \Omega^*$ is a morphism.

The sixth step finally establishes that the description map D is a unique morphism. To that end, I show that the description map from $\vec{\Omega}^*$ into itself is the identity map.

Lemma 4 (Description Map $D_{\vec{\Omega}^*}$). *The description map $D_{\vec{\Omega}^*} : \vec{\Omega}^* \rightarrow \vec{\Omega}^*$ is the identity morphism.*

I prove Lemma 4 by showing $[\cdot] = \llbracket \cdot \rrbracket_{\vec{\Omega}^*}$, which means whether an expression e holds at ω^* (the right-hand side) is determined solely by whether $e \in_1 \omega^*$ (the left-hand side). The lemma implies that the belief space $\vec{\Omega}^*$ is non-redundant and minimal. Together with Remark 6, the lemma implies that $D_{\vec{\Omega}}$ is a unique morphism.³⁹ Thus:

Theorem 1 ($\vec{\Omega}^*$ is Terminal). *The space $\vec{\Omega}^* = \langle (\Omega^*, \mathcal{D}^*), (B_i^*)_{i \in I}, \Theta^* \rangle$ is a terminal κ -belief space of I on (S, \mathcal{S}) for a given category of κ -belief spaces.*

As discussed in Section 2.3, a terminal belief space exists uniquely up to isomorphism. Since terminality requires a unique morphism from a given space, a belief space $\vec{\Omega}$ is terminal iff the description map $D_{\vec{\Omega}}$ is an isomorphism. Especially, any terminal space is non-redundant and minimal.

I discuss how the belief space Ω^* "includes" all belief spaces. First, for any state ω of any particular space $\vec{\Omega}$, states $\omega \in \Omega$ and $D(\omega) \in \Omega^*$ are equivalent in that the same state of nature $\Theta(\omega) = \Theta^*(D(\omega))$ prevails and the same set of expressions regarding nature and interactive beliefs (i.e., the same profile of the players' belief hierarchies) obtains. This is because $D(\omega) = D_{\vec{\Omega}^*}(D(\omega))$ (recall discussions after Remark 6). To restate, for any representation $\vec{\Omega}$ of interactive beliefs regarding (S, \mathcal{S}) and for any

³⁸Consider, for instance, the category of possibility correspondence models of knowledge without Negative Introspection. Among the category, B_i^* fails to satisfy Negative Introspection (even though there are some spaces in the category in which the belief operator B_i satisfies Negative Introspection).

³⁹Let $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}^*$ be a morphism. Then $D_{\vec{\Omega}}(\cdot) = D_{\vec{\Omega}^*}(\varphi(\cdot)) = \varphi(\cdot)$, where the first equality follows from Remark 6 and the second from Lemma 4.

state $\omega \in \Omega$, the prevailing nature state and the prevailing set of expressions at ω are encoded in the state $D(\omega)$ in $\overrightarrow{\Omega}^*$. Especially, the expressions $\{e \in \mathcal{L} \mid \omega \in B_i(\llbracket e \rrbracket_{\overrightarrow{\Omega}})\}$ that player i believes at ω coincide with those $\{e \in \mathcal{L} \mid \omega^* \in B_i^*([e])\}$ that she believes at ω^* .

Second, a non-redundant belief space $\overrightarrow{\Omega}$ is, by definition, embedded into $\overrightarrow{\Omega}^*$: there is a belief (sub-)space $\overrightarrow{D(\Omega)} := \langle (D(\Omega), \mathcal{D}^* \cap D(\Omega)), (B'_i)_{i \in I}, \Theta^*|_{D(\Omega)} \rangle$ such that $D : \overrightarrow{\Omega} \rightarrow \overrightarrow{D(\Omega)}$ is a bijective morphism, where $\mathcal{D}^* \cap D(\Omega) := \{[e] \cap D(\Omega) \mid [e] \in \mathcal{D}^*\}$ and $B'_i([e] \cap D(\Omega)) := B_i^*([e]) \cap D(\Omega)$. If, in addition, $\overrightarrow{\Omega}$ is minimal, then $\overrightarrow{\Omega}$ and $\overrightarrow{D(\Omega)}$ are isomorphic (i.e., the inverse D^{-1} is also a morphism) because any $E \in \mathcal{D}$ is associated with some $[e] \in \mathcal{D}^*$ through $[e] = D^{-1}(E)$. Indeed, $\overrightarrow{\Omega}$ and $\overrightarrow{D(\Omega)}$ are isomorphic iff $\overrightarrow{\Omega}$ is non-redundant and minimal. Observe that if this is the case, then a κ -algebra \mathcal{D} is typically not the power set.

4 Characterization of the Terminal Space

The previous section has constructed the terminal space $\overrightarrow{\Omega}^*$, which is also non-redundant and minimal. This section studies properties of the terminal space $\overrightarrow{\Omega}^*$ further. Section 4.1 shows that the terminal space is (belief-)complete (Brandenburger, 2003): the space contains all possible beliefs about its states. Section 4.2 shows: any set of expressions that hold at some state of some belief space also hold at the corresponding state of the terminal space. Section 4.3 demonstrates how the terminal space resolves the issue of self-reference (e.g., Aumann, 1976): while each state is supposed to represent players' interactive beliefs, their beliefs are defined on the states. Section 4.4 introduces the notion of a belief-closed set and restates the terminal belief space as the largest belief-closed set (Theorem 2).

4.1 Belief-Completeness

In the literature on type spaces, epistemic characterizations of solution concepts of games often refer to (belief-)completeness of a type space.⁴⁰ Informally, call a belief space $\overrightarrow{\Omega}$ (belief-)complete if, for any profile of sets of expressions that individual players can possibly believe within the framework of belief spaces, there exists a state $\omega \in \Omega$ at which each player believes the corresponding set of expressions. This

⁴⁰A type space (i.e., a set of states of the world is the product space consisting of the set of nature states and each player's type set, and a type of each player induces a probabilistic belief over the types of the opponents) is (belief-)complete if, for any belief over the types of the players other than i , there exists a type of player i which induces the given belief. A terminal type space is (belief-)complete (Moss and Viglizzo, 2004, 2006). A (belief-)complete type space is terminal given topological conditions (Friedenberg, 2010), and Friedenberg and Keisler (2021) show that the topological conditions cannot be dropped (these two papers also study iterated elimination of strictly dominated actions in the type-space framework). See also Appendix D.1.

subsection shows that a non-redundant and minimal belief space is terminal iff it is (belief-)complete. Especially, the terminal space $\vec{\Omega}^*$ is (belief-)complete.⁴¹

To define (belief-)completeness, fix a category of κ -belief spaces (that satisfy some properties of beliefs specified in Definition 2). I define a new set Ω^{**} as follows:

$$\Omega^{**} := \{(s, \Psi) \in S \times \mathcal{P}(\mathcal{L})^I \mid (s, \Psi) = (\Theta(\omega), (\{e \in \mathcal{L} \mid \omega \in B_i[\![e]\!]_{\vec{\Omega}}\})_{i \in I}) \text{ for some belief space } \vec{\Omega} \text{ and } \omega \in \Omega\}.$$

The set Ω^{**} consists of a nature state s and sets of expressions Ψ that individual players believe in the category of belief spaces. Since assumptions on beliefs are arbitrary and thus it is cumbersome to explicitly define the conditions on sets of expressions Ψ that reflect given assumptions of beliefs (see Section 4.4 for an explicit characterization), here I consider sets of expressions that the players believe at some state of some belief space.

Call a belief space $\vec{\Omega}$ (in the given category) *(belief-)complete* if the mapping

$$\chi_{\vec{\Omega}} : \Omega \ni \omega \mapsto (\Theta(\omega), (\{e \in \mathcal{L} \mid \omega \in B_i(\![e]\!]_{\vec{\Omega}}\})_{i \in I}) \in \Omega^{**}$$

is surjective. Each mapping $\omega \mapsto \{e \in \mathcal{L} \mid \omega \in B_i(\![e]\!]_{\vec{\Omega}}\}$ that constitutes part of $\chi_{\vec{\Omega}}$ defines player i 's “type mapping” that associates, with each state ω , her beliefs on $\vec{\Omega}$ about the expressions \mathcal{L} (or $\mathcal{D}_{\vec{\Omega}}$).⁴² With this definition:

Proposition 1 (Belief-Completeness). 1. A belief space $\vec{\Omega}$ is terminal (in the given category) iff it is minimal, non-redundant, and (belief-)complete.

2. A belief space $\vec{\Omega}$ is non-redundant iff $\chi_{\vec{\Omega}}$ is injective.

3. The mapping $\chi_{\vec{\Omega}^*} : \Omega^* \rightarrow \Omega^{**}$ is bijective. In particular,

$$\Omega^{**} = \{(\Theta^*(\omega^*), (\{e \in \mathcal{L} \mid \omega^* \in B_i^*[e]\})_{i \in I}) \in S \times \mathcal{P}(\mathcal{L})^I \mid \omega^* \in \Omega^*\}. \quad (2)$$

The main part of the proposition is Part (1). It implies that a non-redundant and minimal belief space is (belief-)complete iff it is terminal. Especially, $\vec{\Omega}^*$ is (belief-)complete. The proof shows that (belief-)completeness of $\vec{\Omega}^*$ follows from

$$(\Theta(\omega), (\{e \in \mathcal{L} \mid \omega \in B_i(\![e]\!]_{\vec{\Omega}}\})_{i \in I}) = (\Theta^*(D(\omega)), (\{e \in \mathcal{L} \mid D(\omega) \in B_i^*[e]\})_{i \in I}). \quad (3)$$

⁴¹To define the notion of (belief-)completeness in a way free from the structure of an underlying set of states of the world (it does not need to be a product space) and properties of beliefs (i.e., beliefs can be qualitative and are independent of such properties as Truth Axiom, Positive Introspection and Negative Introspection), I define (belief-)completeness in terms of “language” (sets of expressions that individual players believe), following the idea of Brandenburger and Keisler (2006).

⁴²Section 5.1 applies the construction in Section 3 to probabilistic beliefs. There, player i 's p -belief operators $(B_i^p)_{p \in [0,1]}$ yield her type $m_{B_i}(\omega)$ at a state ω through the maximum probability p with which she believes an event at the state: $m_{B_i}(\omega)(E) := \sup\{p \in [0, 1] \mid \omega \in B_i^p(E)\}$ for each $E \in \mathcal{D}$. Hence, the set of expressions that a player believes at a state recovers her beliefs at that state.

For any state of nature and any profile of sets of expressions that individual players believe at some state ω of some belief space $\vec{\Omega}$ (the left-hand side), there exists a state $D(\omega)$ in the terminal belief space at which the corresponding state of nature prevails and the players believe the given sets of expressions (the right-hand side).

Part (3) implies that the mapping $\chi_{\vec{\Omega}^*}$, which associates, with each state $\omega^* \in \Omega^*$, the corresponding nature state and profile of sets of expressions that the individual players believe at ω^* , is a bijection. This gives a sense in which each state $\omega^* \in \Omega^*$ represents what each player believes at the state ω^* itself. This is because, for any expression e , whether player i believes $[e]$ or not at ω^* (i.e., whether $\omega^* \in B_i^*[e]$ or not) is incorporated within ω^* itself in the sense of whether $\beta_i(e) \in_1 \omega^*$ or not.

Proposition 1 demonstrates the sense in which $\vec{\Omega}^*$ contains all possible beliefs about its states, as a (belief-)complete type space contains all possible beliefs about its types (see footnote 40). Namely, the set Ω^* is in a bijective relation with the set

$$\{(s, \Psi) \in S \times \mathcal{P}(\mathcal{D}^*)^I \mid (s, \Psi) = (\Theta(\omega^*), (\{[e] \in \mathcal{D}^* \mid \omega^* \in B_i([e])\})_{i \in I}) \\ \text{for some belief space } \langle (\Omega^*, \mathcal{D}^*), (B_i)_{i \in I}, \Theta \rangle \text{ and } \omega^* \in \Omega^*\}$$

through a bijection $\omega^* \mapsto (\Theta(\omega^*), (\{[e] \in \mathcal{D}^* \mid \omega^* \in B_i([e])\})_{i \in I})$. For any profile of nature state that holds and sets of events that the individual players believe at some state of some belief structure on $(\Omega^*, \mathcal{D}^*)$, there is a state in $\vec{\Omega}^*$ that generates the given profile. Thus, $\vec{\Omega}^*$ contains all possible beliefs about its states.

4.2 Informational Robustness

This subsection formalizes a sense in which interactive beliefs about nature states S in a particular belief space can be extended to the terminal space. Recall that the set \mathcal{L} of expressions (Definition 5) specifies, without reference to any specific belief space, the nature states S and players' interactive beliefs about S . Concretely, let S be the set of action profiles of a strategic game. Let $e \in \mathcal{S}$ be a subset of action profiles. The expression e denotes "the players' actions are in e ." Then, the expression $f := \bigwedge_{i \in I} \beta_i(e)$ denotes "every player believes that the players' actions are in e ." In this way, one can capture a certain form of players' interactive beliefs by using the set of expressions.

Now, one can ask whether a set of expressions holds at a particular state of a particular belief space. For instance, the set of expressions $\{e, f\}$ holds true at some state ω (i.e., at ω , the players' actions are in e and every player believes that their actions are in e) in some belief space $\vec{\Omega}$ if $\omega \in \llbracket e \rrbracket_{\vec{\Omega}}$ and $\omega \in \llbracket f \rrbracket_{\vec{\Omega}}$. It may be the case that, in the particular space $\vec{\Omega}$, the players' actions at ω are in e and the players take actions in e iff they believe that their actions are in e (for instance, this holds when their actions are in e at any state and their beliefs satisfy Necessitation). Then, the set of expressions $\{e, f\}$ indeed holds true at ω . Moreover, e and f happen to be equivalent in the sense that $\llbracket e \rrbracket_{\vec{\Omega}} = \llbracket f \rrbracket_{\vec{\Omega}}$.

Proposition 2 below shows that, for the set of expressions $\{e, f\}$ that holds at some state ω of some belief space $\vec{\Omega}$, the set of expressions $\{e, f\}$ holds at the corresponding state $D(\omega)$ of the terminal space $\vec{\Omega}^*$. Thus, at $D(\omega)$, the players' actions are in e and they believe their actions are in e .

In contrast, suppose that there exists a belief space $\vec{\Omega}'$ with $\llbracket e \rrbracket_{\vec{\Omega}'} \neq \llbracket f \rrbracket_{\vec{\Omega}'}$. In this case, for instance, there may exist a state $\omega' \in \Omega'$ at which $\{e, (\neg f)\}$ holds. Proposition 2 below shows that if some belief space $\vec{\Omega}'$ distinguishes expressions e and f in that $\llbracket e \rrbracket_{\vec{\Omega}'} \neq \llbracket f \rrbracket_{\vec{\Omega}'}$, then so does the terminal space $\vec{\Omega}^*$. At $D_{\vec{\Omega}'}(\omega')$, while the players' actions are in e , they do not believe that their actions are in e . This suggests that the equivalence of e and f , i.e., $\llbracket e \rrbracket_{\vec{\Omega}} = \llbracket f \rrbracket_{\vec{\Omega}}$, is an additional assumption imposed by the structure of the particular belief space $\vec{\Omega}$. In fact, the expressions e and f are not equivalent in the terminal space, i.e., $\llbracket e \rrbracket_{\vec{\Omega}^*} \neq \llbracket f \rrbracket_{\vec{\Omega}^*}$.

To make these claims formal, I start with introducing:

Definition 8 (Semantic Properties). 1. (a) An expression $e \in \mathcal{L}$ is *valid in a belief space* $\vec{\Omega}$, written $\models_{\vec{\Omega}} e$, if $\llbracket e \rrbracket_{\vec{\Omega}} = \Omega$.

(b) If e is valid in any belief space of the given category, then e is *valid* (in the given category), written $\models e$.

2. (a) A set of expressions $\Phi \in \mathcal{P}(\mathcal{L})$ is *satisfiable* in $\vec{\Omega}$ if there is $\omega \in \Omega$ with $\omega \in \llbracket f \rrbracket_{\vec{\Omega}}$ for all $f \in \Phi$.

(b) Call Φ *satisfiable* if Φ is satisfiable in some belief space $\vec{\Omega}$.

3. A set of expressions $\Phi \in \mathcal{P}(\mathcal{E})$ is *maximally satisfiable* if it is satisfiable and if $\Phi = \Psi$ for any satisfiable set Ψ with $\Phi \subseteq \Psi$.

4. (a) An expression $e \in \mathcal{L}$ is a *semantic consequence* of $\Phi \in \mathcal{P}(\mathcal{L})$ in $\vec{\Omega}$, written $\Phi \models_{\vec{\Omega}} e$, if: $\omega \in \llbracket e \rrbracket_{\vec{\Omega}}$ holds whenever $\omega \in \llbracket f \rrbracket_{\vec{\Omega}}$ for all $f \in \Phi$.

(b) If $\Phi \models_{\vec{\Omega}} e$ for any belief space $\vec{\Omega}$, then $e \in \mathcal{L}$ is a *semantic consequence* of Φ , written $\Phi \models e$.

A morphism preserves the notion of validity, satisfiability, and semantic consequences in the following senses. Let $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}'$ be a morphism. By Remark 4, any valid expression e in $\vec{\Omega}'$ is also valid in $\vec{\Omega}$. If Φ is satisfiable in $\vec{\Omega}$, then so is it in $\vec{\Omega}'$. Suppose further that φ is surjective. If e is a semantic consequence of Φ in $\vec{\Omega}$, then so is it in $\vec{\Omega}'$.

The following proposition formalizes a sense in which interactive beliefs about nature states S in a particular belief space can be extended to the terminal space, using the semantic notions of satisfiability, semantic consequence, and validity in $\vec{\Omega}^*$.

Proposition 2 (Informational Robustness). *Let $e \in \mathcal{L}$ and $\Phi \in \mathcal{P}(\mathcal{L})$.*

1. Φ is satisfiable iff Φ is satisfiable in $\vec{\Omega}^*$.
2. e is a semantic consequence of $\vec{\Omega}$ iff e is a semantic consequence of Φ in $\vec{\Omega}^*$.
Also, e is valid iff e is valid in $\vec{\Omega}^*$.
3. $\Omega^* = \{\{s\} \sqcup \Phi \in \mathcal{P}(S \sqcup \mathcal{L}) \mid \Phi \text{ is maximally satisfiable and, for any } E \in \mathcal{A}_\kappa(\mathcal{S}), s \in E \text{ iff } E \in \Phi\}.$

Part (1) means that if expressions Φ hold at some state ω in some belief space $\vec{\Omega}$, then the expressions Φ hold at $D(\omega)$ in $\vec{\Omega}^*$. The set of expressions that hold at some state of some belief space is maximally satisfiable, and any maximally satisfiable set, in turn, can be associated with the set of expressions that hold at some state of some belief space. Thus, Part (3) demonstrates that each state of $\vec{\Omega}^*$ is a maximally satisfiable set of expressions and a corresponding nature state.

Part (2) implies that if expressions e and f satisfy $\llbracket e \rrbracket_{\vec{\Omega}} \neq \llbracket f \rrbracket_{\vec{\Omega}}$ for some belief space $\vec{\Omega}$, then $\llbracket e \rrbracket_{\vec{\Omega}^*} \neq \llbracket f \rrbracket_{\vec{\Omega}^*}$. Put differently, if $\llbracket e \rrbracket_{\vec{\Omega}^*} = \llbracket f \rrbracket_{\vec{\Omega}^*}$, then it is always the case in any belief space $\vec{\Omega}$ that $\llbracket e \rrbracket_{\vec{\Omega}} = \llbracket f \rrbracket_{\vec{\Omega}}$. Thus, the space $\vec{\Omega}^*$ makes the minimum possible assumptions on how the players' interactive beliefs about nature states are modeled.⁴³

Proposition 2 provides a formal sense in which reasoning about the states of nature S in a smaller belief space can be extended to the terminal space. Section 6.3 shows that players' rationality and common belief in rationality are well-defined in every qualitative-belief space. This is because, in the context of Section 6.3 in which S is the set of action profiles of a strategic game, players' rationality and common belief in rationality are expressed solely from players' interactive reasoning about S . If the players are rational at some state of some belief space, then they are rational at the corresponding state of the terminal space. For probabilistic beliefs, Appendix E.5 studies correlated equilibria. One epistemic characterization is Bayes rationality: a probabilistic-belief space in which each player is Bayes rational at every state (and thus players commonly believe their Bayes rationality) is a correlated equilibrium (Aumann, 1987). In the category of probabilistic-belief spaces in which Bayes rationality is valid, a terminal space exists. At every state of the terminal space, the players are Bayes rational (and hence they commonly believe their Bayes rationality).

To sum up this subsection, I discuss two papers. First, Bjorndahl, Halpern, and Pass (2013) study psychological games (Battigalli and Dufwenberg, 2009; Geanakoplos, Pearce, and Stacchetti, 1989) in which players' interactive qualitative beliefs themselves enter into their preferences. This paper implies that one can introduce

⁴³For instance, consider the category of belief spaces in which each player's belief operator B_i satisfies the Kripke property, Consistency, Positive Introspection, and Negative Introspection. Then, in the terminal space, B_i^* satisfies these four properties, and B_i^* violates any other properties of beliefs that are not derived from these four properties (e.g., Truth Axiom).

an arbitrary notion of beliefs including qualitative beliefs in psychological games.⁴⁴ BJORNDALH, HALPERN, and PASS (2013) introduce the language (as a formal terminology in logic) which describes players' strategy choices and their interactive beliefs (in the framework of this paper, $\mathcal{L}_{\aleph_0}(\mathcal{A}_{\aleph_0}(\mathcal{S}))$), and a player's utility function is defined on the set, each element of which is a maximally satisfiable set of formulas from the language. By Proposition 2 (3), a maximally satisfiable set of expressions, together with a corresponding nature state (i.e., the players' strategy choices), constitute a state of the terminal space $\vec{\Omega}^*$. Hence, each player i 's payoff function is defined over the terminal space Ω^* .⁴⁵ For solution concepts, BJORNDALH, HALPERN, and PASS (2013) define rationalizability by its epistemic characterization. As discussed above, Section 6.3 studies an implication of common belief in rationality.

Second, in the context of Bayesian games, FRIEDENBERG and MEIER (2017) show that a (smaller) type space may have a Bayesian equilibrium that cannot be extended to a larger type space (e.g., a terminal type space of BRANDENBURGER and DEKEL (1993), HEIFETZ and SAMET (1998b), and MERTENS and ZAMIR (1985)). On the one hand, recall that, in the above applications when the players interactively reason about their actions of an underlying strategic game, a belief space $\vec{\Omega} = \langle (\Omega, \mathcal{D}), (B_i)_{i \in I}, \Theta \rangle$ encodes the players' strategy choices Θ as part of the primitives. Also, the strategies Θ are always assumed to be measurable. Thus, strategy choices and rationality are well-defined within the belief space $\vec{\Omega}$. In fact, Appendix E.4 provides an epistemic characterization of pure-strategy Nash equilibria on the terminal belief space. Appendix E.5 shows that if a correlated equilibrium as a probabilistic-belief space is non-redundant and minimal then the underlying state space Ω of the correlated equilibrium can be replaced by a subspace $D(\Omega)$ of the terminal space which consists of players' belief hierarchies.

On the other hand, in a Bayesian game in which the players have uncertainty about payoff parameters S , a Bayesian equilibrium consists of (i) a belief space $\vec{\Omega}$ (in which the players reason about payoff parameters S) and, separately from the belief space, (ii) the players' (equilibrium) strategy choices, which are functions from the states of the world Ω to the action profiles A . In the belief space $\vec{\Omega}$, the function $\Theta : \Omega \rightarrow S$ associates, with each state of the world, the corresponding payoff parameter $\Theta(\omega) \in S$,

⁴⁴For probabilistic beliefs, this paper accommodates the existence of a universal probabilistic-belief space or that of a universal conditional-probabilistic-belief space. In fact, this paper establishes such existence result without imposing certain properties of beliefs such as countable additivity. Hence, it is possible to extend the scope of psychological games to non-standard probabilistic beliefs. As to qualitative beliefs, this paper extends the framework of BJORNDALH, HALPERN, and PASS (2013) in which interactive beliefs are represented through a logical system. Also, as an example of the use of qualitative belief in psychological games, suppose Alice can send Bob either flowers or chocolates (GEANAKOPLOS, PEARCE, and STACCHETTI, 1989). Alice enjoys surprising Bob. So, if Alice believes that Bob expects (i.e., believes that he receives) flowers, she sends chocolates, and vice versa.

⁴⁵As standard psychological games consider probabilistic beliefs or probabilistic conditional beliefs, Section 5 establishes a terminal space for probabilistic beliefs and Appendix D.1 a terminal space for probabilistic conditional beliefs.

not their action profile. As shown by Friedenberg and Meier (2017) (see also the references therein for the broader literature), Bayesian equilibria may depend on a specific belief space within which belief hierarchies are modeled and some Bayesian equilibria in a given belief space may fail to extend to a larger belief space.

The analyses of this subsection are different from Friedenberg and Meier (2017) in the sense that the players' strategy choices are incorporated within a belief space itself for the game-theoretic applications in Section 6, and this paper does not claim that players' reasoning about Bayesian equilibria in a smaller belief space is preserved to a larger space. In fact, in the context of Bayesian games, Proposition 2 would imply only that the players' reasoning about payoff uncertainty S in a smaller space is preserved in the terminal space. This paper would clarify that a terminal space includes reasoning in a smaller space only about a given set S of nature states. Only to clarify this issue further, Appendix C briefly mentions a possibility that one may be able to analyze Bayesian equilibria of a fixed underlying game in a terminal belief space by incorporating players' strategy choices as a primitive of a belief space. There, the set of states of nature is the product space $S \times A$, where S is the set of payoff parameters and A is the set of action profiles of the underlying game.

4.3 Self-Reference

A conceptual issue in modeling players' beliefs on a state space since Aumann (1976) is the interpretation of each state as a “full” description of players' beliefs at the state, as this interpretation leads to self-reference: each state refers to players' beliefs at that state. For instance, suppose player i believes some event E at a state ω . If each state is a full description, then the state ω has to include the description that “player i believes E at ω .” This subsection formalizes the sense in which each state ω^* of the terminal space $\overrightarrow{\Omega}^*$ is a full description of players' beliefs at the state.

Observe that each state ω^* of the terminal space $\overrightarrow{\Omega}^*$ represents each player's belief at that state by using expressions that hold at ω^* . That is, the state ω^* represents the fact that player i believes an event $[e]$ at ω^* , i.e., $\omega^* \in B_i^*([e])$, by $\beta_i(e) \in_1 \omega^*$.

Generally, each state ω^* of the terminal space $\overrightarrow{\Omega}^*$ contains expressions that hold at ω^* (i.e., $e \in_1 \omega^*$ iff $\omega^* \in [e]$) as well as the corresponding nature state $\Theta^*(\omega^*) \in S$. Thus, I define two conditions for each state ω^* to be a full description of the corresponding nature state and players' beliefs at that state.

The first is its logical structure, following Aumann (1999). Each ω^* is *coherent*: if ω^* contains an expression e then it does not contain $(\neg e)$. Each ω^* is *complete*: if ω^* does not contain e then it contains $(\neg e)$. As shown below, every ω^* is logically closed under implication and conjunction.

Proposition 3 (Logical Properties of Each State). *Fix $\omega^* \in \Omega^*$, $e \in \mathcal{L}$, $f \in \mathcal{L}$, and $\mathcal{E} \in \mathcal{P}(\mathcal{L})$ with $0 < |\mathcal{E}| < \kappa$.*

1. *Coherency and Completeness: $e \notin_1 \omega^*$ iff $(\neg e) \in_1 \omega^*$.*

2. *Closure under Implication:* If $e \in_1 \omega$ and $(e \rightarrow f) \in_1 \omega$, then $f \in_1 \omega$.

3. *Closure under Conjunction:* $\bigwedge \mathcal{E} \in_1 \omega^*$ iff $e \in_1 \omega^*$ for all $e \in \mathcal{E}$.

Second, the terminal space $\overrightarrow{\Omega}^*$ resolves self-reference as each player's beliefs at each state ω^* are written within ω^* itself in the sense that $\omega^* \in B_i^*([e])$ iff $\beta_i(e) \in_1 \omega^*$. Since each ω^* satisfies the logical requirement postulated in Proposition 3 and since B_i^* inherits the properties of beliefs assumed in a given category, the following corollary shows that the players' beliefs at each state are encoded within the state itself.

Corollary 1 (Beliefs within Each State). *Fix $\omega^* \in \Omega^*$, $i \in I$, and $e \in \mathcal{L}$.*

1. *Exactly one of $\beta_i(e) \in_1 \omega^*$ or $(\neg\beta_i)(e) \in_1 \omega^*$ holds.*

2. *At least one of $\beta_i(e) \in_1 \omega^*$, $\beta_i(\neg e) \in_1 \omega^*$, or $(\neg\beta_i)(e) \wedge (\neg\beta_i)(\neg e) \in_1 \omega^*$ holds. Exactly one of them always holds (for any (e, ω^*)) iff Consistency holds for i .*

3. *Exactly one of $\beta_i(e) \in_1 \omega^*$, $(\neg\beta_i)(e) \wedge \beta_i(\neg\beta_i)(e) \in_1 \omega^*$, or $(\neg\beta_i)(e) \wedge (\neg\beta_i)(\neg\beta_i)(e) \in_1 \omega^*$ holds. The third condition never occurs (for any (e, ω^*)) iff Negative Introspection holds for i .*

Part (1) of Corollary 1 provides the sense in which each ω^* fully describes i 's beliefs: for any $e \in \mathcal{L}$, the state ω^* contains exactly one of the two expressions denoting “ i believes e ” or “ i does not believe e .” Parts (2) and (3) characterize how the space Ω^* encodes whether such assumptions as Consistency and Negative Introspection are made within Ω^* . For instance, if player i 's beliefs violate Consistency, then at some state ω^* at which Consistency is violated for some $[e]$ (i.e., $\omega^* \in B_i^*([e]) \cap B_i^*(\neg[e])$), the state ω^* contains the expressions $\beta_i(e)$ and $\beta_i(\neg e)$, explicitly indicating that player i 's beliefs violate Consistency. Corollary 1 is related to some consistency conditions of Gilboa (1988) for a state to be a full description of players' beliefs. The next subsection characterizes how states in Ω^* encode all the properties of beliefs specified in Definition 2 by demonstrating that properties imposed on player i 's beliefs in a given category are expressed within Ω^* .

4.4 Largest Belief-Closed Set Ω^*

In the terminal belief space $\overrightarrow{\Omega}^*$ constructed in Section 3, the state space Ω^* is implicitly characterized as the set of the players' belief hierarchies $D_{\overrightarrow{\Omega}}(\omega)$ at some state ω of some belief space $\overrightarrow{\Omega}$. This subsection, instead, explicitly characterizes the space $\overrightarrow{\Omega}^*$ as the largest *belief-closed* set.

Call a subset Ω of $\mathcal{P}(S \sqcup \mathcal{L})$ belief-closed if: (i) each $\omega \in \Omega$ is a complete and coherent set of expressions together with a unique nature state; (ii) the set Ω as a whole induces the players' beliefs in a well-defined manner; and (iii) each ω reflects assumptions on players' beliefs. This characterization holds for any infinite regular cardinal κ and assumptions on beliefs. Formally:

Definition 9 (Belief-Closed Set of Descriptions). Call a subset Ω of $\mathcal{P}(S \sqcup \mathcal{L})$ a *belief-closed* set of descriptions if it satisfies the following conditions.

1. Each $\omega \in \Omega$ satisfies the following.
 - (a) Nature State: There is a unique $s \in S$ with $s \in_0 \omega$. Moreover, for all $E \in \mathcal{A}_\kappa(\mathcal{S})$, $s \in E$ iff $E \in_1 \omega$.
 - (b) Coherency and Completeness: For each $e \in \mathcal{L}$, $e \notin_1 \omega$ iff $(\neg e) \in_1 \omega$.
 - (c) Closure under Implication: If $e \in_1 \omega$ and $(e \rightarrow f) \in_1 \omega$ then $f \in_1 \omega$.
 - (d) Closure under Conjunction: For any $\mathcal{E} \in \mathcal{P}(\mathcal{L})$ with $0 < |\mathcal{E}| < \kappa$, $\bigwedge \mathcal{E} \in_1 \omega$ iff $e \in_1 \omega$ for all $e \in \mathcal{E}$.
2. The set Ω satisfies the following ((2b) and (2c) depend on assumptions on beliefs).
 - (a) Equivalence: If $e, f \in \mathcal{L}$ satisfy $(e \leftrightarrow f) \in_1 \omega$ for all $\omega \in \Omega$, then $(\beta_i(e) \leftrightarrow \beta_i(f)) \in_1 \omega$ for all $\omega \in \Omega$.
 - (b) Let Monotonicity be assumed for player i . If $e, f \in \mathcal{L}$ satisfy $(e \rightarrow f) \in_1 \omega$ for all $\omega \in \Omega$, then $(\beta_i(e) \rightarrow \beta_i(f)) \in_1 \omega$ for all $\omega \in \Omega$.
 - (c) Let the Kripke property be assumed for player i . Then, $\beta_i(e) \in_1 \omega$ for any $(e, \omega) \in \mathcal{L} \times \Omega$ with the following condition: if $\omega' \in \Omega$ satisfies $f \in_1 \omega'$ for all $f \in \mathcal{L}$ with $\beta_i(f) \in_1 \omega$, then $e \in_1 \omega'$.
3. Depending on assumptions on beliefs, each $\omega \in \Omega$ contains any instance of the following expressions.
 - (a) Necessitation: $\beta_i(S)$.
 - (b) λ -Conjunction: $((\bigwedge_{e \in \mathcal{E}} \beta_i(e)) \rightarrow \beta_i(\bigwedge \mathcal{E}))$ with $0 < |\mathcal{E}| < \lambda$.
 - (c) Consistency: $(\beta_i(e) \rightarrow (\neg \beta_i)(\neg e))$.
 - (d) Truth Axiom: $(\beta_i(e) \rightarrow e)$.
 - (e) Positive Introspection: $(\beta_i(e) \rightarrow \beta_i \beta_i(e))$.
 - (f) Negative Introspection: $((\neg \beta_i)(e) \rightarrow \beta_i(\neg \beta_i)(e))$.

First, Condition (1a) states that each $\omega \in \Omega$ describes a corresponding nature state s in a well-defined manner. Also, ω contains those nature events $E \in \mathcal{A}_\kappa(\mathcal{S})$ that are true at s (i.e., $s \in E$). Conditions (1b) through (1d) are logical requirements on each $\omega \in \Omega$ (recall Proposition 3 for each $\omega^* \in \Omega^*$).

Second, Condition (2a) requires that if two expressions e and f are equivalent in the sense that every ω contains $(e \leftrightarrow f)$, then expressions $\beta_i(e)$ and $\beta_i(f)$ are equivalent in the same sense. This condition allows one to define the players' belief operators in a way such that if two expressions e and f correspond to the same event then the events associated with the beliefs in e and f are the same.

Third, Conditions (2b), (2c), and (3) describe how states Ω encode properties of beliefs. Corollary 1 has provided related characterizations for Consistency and Negative Introspection for Ω^* .

Now, I restate the terminal space Ω^* as the largest belief-closed set.

Theorem 2 (Ω^* as the Largest Belief-Closed Set). *The set Ω^* constructed in Section 3 is the largest belief-closed set of descriptions. Namely, for any belief-closed set Ω , there is a (non-redundant and minimal) belief space $\overrightarrow{\Omega}$ such that its description map $D_{\overrightarrow{\Omega}} : \overrightarrow{\Omega} \rightarrow \overrightarrow{\Omega}^*$ is an inclusion map and thus $\Omega \subseteq \Omega^*$.*

Two remarks on Theorem 2 are in order. First, Theorem 2 restates the terminal space $\overrightarrow{\Omega}^*$ as the largest belief-closed set irrespective of properties of beliefs. In particular, in the category of \aleph_0 -knowledge spaces in which each player's knowledge operator is induced by a partitional possibility correspondence, Theorem 2 formally proves that Aumann (1999)'s canonical space can be taken as a terminal \aleph_0 -knowledge space in this category, contrary to the conjecture of Heifetz and Samet (1998b, Section 6).

Second, the two constructions of a terminal belief space in Theorems 1 and 2 have some analogies with the following two constructions of a terminal type space. The first is Heifetz and Samet (1998b), in which a terminal type space consists of belief hierarchies induced by some type profile of some type space. This corresponds to the construction of Theorem 1, where each state in $\overrightarrow{\Omega}^*$ is a profile of the players' belief hierarchies (precisely, a description) induced by some state of some belief space. The second is Brandenburger and Dekel (1993), in which, using the terminology of this paper, a terminal type space is the largest "belief-closed" set of coherent belief hierarchies. In Theorem 2, each state in $\overrightarrow{\Omega}^*$ represents the players' coherent belief hierarchies, and Ω^* is the largest belief-closed set in the sense of Definition 9.⁴⁶

5 Applications to Other Forms of Beliefs

Section 3 has demonstrated that a terminal belief space exists when the players' beliefs are represented by their belief operators and when an infinite regular cardinal κ determines the depth of reasoning through the specification of a domain as a κ -algebra. The framework of this paper (especially, the existence and characterizations of a terminal belief space) applies to richer forms of beliefs such as probabilistic beliefs because probabilistic beliefs are represented by p -belief operators on a κ -algebra.

⁴⁶Two remarks for the expert reader familiar with the technical details of Brandenburger and Dekel (1993): First, the definition of a coherent belief hierarchy here is stronger than the one in the type space literature. In Brandenburger and Dekel (1993), a belief hierarchy (consisting of all finite-level beliefs) is coherent if no finite levels of beliefs contradict with each other (by their topological assumption, such coherent belief hierarchies extend to countable levels). In contrast, here, each belief hierarchy is of depth less than $\bar{\kappa}$, and no different levels of the belief hierarchy contradicts with each other. Second, in Brandenburger and Dekel (1993), belief closure is defined by the iterations of mutual certainty (known as "common certainty of coherency" in the literature).

Section 5.1 constructs a terminal space for countably-additive, finitely-additive, and non-additive (not-necessarily-additive) beliefs. The framework nests, for example, Heifetz and Samet (1998b), Meier (2006), and Pintér (2012), and establishes the existence of a terminal non-additive belief space irrespective of any continuity property on beliefs.⁴⁷ The terminal countably-additive belief space can also be reconstructed as belief hierarchies of all finite levels, consistently with the previous literature (Proposition 4).

The Supplementary Appendix discusses further applications. Appendix D.1 discusses a terminal space for conditional probability systems (CPSs) (Battigalli and Siniscalchi, 1999; Guarino, 2017). In Appendix D.2, players' knowledge and qualitative beliefs are indexed by time (Battigalli and Bonanno, 1997). One can also combine knowledge and probabilistic beliefs (Meier, 2008). Appendix D.3 briefly discusses further possible applications, namely, terminal knowledge-unawareness, preference, and expectation spaces.

5.1 Terminal Probabilistic-Belief Space

I formulate a probabilistic-belief space in terms of p -belief operators using the equivalence between a type mapping and p -belief operators established by Samet (2000). To study probabilistic beliefs, throughout this subsection, let $\kappa = \aleph_1$ so that the underlying state space (Ω, \mathcal{D}) of a belief space is a measurable space. Denote by $\Delta(\Omega)$ the set of countably-additive probability measures on (Ω, \mathcal{D}) . Let \mathcal{D}_Δ be the \aleph_1 -algebra on $\Delta(\Omega)$ generated by $\{\{\mu \in \Delta(\Omega) \mid \mu(E) \geq p\} \in \mathcal{P}(\Delta(\Omega)) \mid (E, p) \in \mathcal{D} \times [0, 1]\}$ as in Heifetz and Samet (1998b). I define a probabilistic-belief space as follows.

Definition 10 (Probabilistic-Belief Space). A *probabilistic-belief space* of I on (S, \mathcal{S}) is a tuple $\overrightarrow{\Omega} := \langle (\Omega, \mathcal{D}), (B_i^p)_{(i,p) \in I \times [0,1]}, \Theta \rangle$ with the following properties.

1. (Ω, \mathcal{D}) is an \aleph_1 -algebra and the map $\Theta : (\Omega, \mathcal{D}) \rightarrow (S, \mathcal{A}_{\aleph_1}(\mathcal{S}))$ is measurable.
2. Player i 's p -belief operators $B_i^p : \mathcal{D} \rightarrow \mathcal{D}$ satisfy the properties below. For each $E \in \mathcal{D}$, $B_i^p(E)$ is the event that player i p -believes E , i.e., she believes that the event E has occurred with probability at least p .
 - (a) Non-Negativity: $B_i^0(\cdot) = \Omega$.
 - (b) p -Regularity: If $p_n \uparrow p$ then $B_i^{p_n}(\cdot) \downarrow B_i^p(\cdot)$.
 - (c) Monotonicity: If $E \subseteq F$ then $B_i^p(E) \subseteq B_i^p(F)$.
 - (d) Normalization: $B_i^1(\Omega) = \Omega$.

⁴⁷Heifetz and Samet (1998b) employ a product of players' type sets while this paper does a single non-product state space Ω . If each player is always “certain” of her beliefs (see Definition 10), then the non-product terminal belief space is isomorphic to the product of nature states and players' type spaces (Mertens and Zamir, 1985). These remarks also apply to a terminal space for conditional probability systems (CPSs) in Appendix D.1.

- (e) Super-additivity: $B_i^p(E \cap F) \cap B_i^q(E \cap (\neg F)) \subseteq B_i^{p+q}(E)$ for $p + q \leq 1$.
- (f) Sub-additivity: $(\neg B_i^p)(E) \cap (\neg B_i^q)(F) \subseteq (\neg B_i^{p+q})(E \cup F)$ for $p + q \leq 1$.
- (g) Continuity-from-above: If $E_n \downarrow E$ and $E \in \mathcal{D}$ then $B_i^p(E_n) \downarrow B_i^p(E)$.
- (h) Continuity-from-below: If $E_n \uparrow E$ and $E \in \mathcal{D}$ then $B_i^p(E) = \bigcap_{r \in \mathbb{N}: p - \frac{1}{r} \geq 0} \bigcup_{n \in \mathbb{N}} B_i^{p - \frac{1}{r}}(E_n)$.
- (i) Certainty-of-Beliefs: If $[m_{B_i}(\omega)] \subseteq E$, then $\omega \in B_i^1(E)$, where

$$[m_{B_i}(\omega)] := \left(\bigcap_{\substack{(p, E) \in [0, 1] \times \mathcal{D} \\ \omega \in B_i^p(E)}} B_i^p(E) \right) \cap \left(\bigcap_{\substack{(p, E) \in [0, 1] \times \mathcal{D} \\ \omega \in (\neg B_i^p)(E)}} (\neg B_i^p)(E) \right).$$

In a probabilistic-belief space, player i 's p -belief operators induce her *type mapping* $m_{B_i} : (\Omega, \mathcal{D}) \rightarrow (\Delta(\Omega), \mathcal{D}_\Delta)$, a measurable map defined by

$$m_{B_i}(\omega)(E) := \sup\{p \in [0, 1] \mid \omega \in B_i^p(E)\} \text{ for each } (\omega, E) \in \Omega \times \mathcal{D}.$$

At each state ω , a countably-additive probability measure $m_{B_i}(\omega)$ represents i 's beliefs at that state. Since m_{B_i} is measurable, it reproduces the original p -belief operators:

$$B_i^p(E) = \{\omega \in \Omega \mid m_{B_i}(\omega)(E) \geq p\} \in \mathcal{D} \text{ for each } E \in \mathcal{D}.$$

Conditions (2), slightly different from Samet (2000), axiomatize type mappings. Conditions (2a) and (2b) guarantee that the map m_{B_i} is well-defined, irrespective of properties of probabilistic beliefs. By (2c), each $m_{B_i}(\omega)$ is *monotonic*: $E \subseteq F$ implies $m_{B_i}(\omega)(E) \leq m_{B_i}(\omega)(F)$. Condition (2d) is a normalization: $m_{B_i}(\cdot)(\Omega) = 1$. Thus, (2a)-(2d) yield non-additive beliefs (or capacities).

By (2e), each $m_{B_i}(\omega)$ is *super-additive*: $m_{B_i}(\omega)(E \cap F) + m_{B_i}(\omega)(E \cap (\neg F)) \leq m_{B_i}(\omega)(E)$. Note that (2a) and (2e) imply (2c). By (2f), each $m_{B_i}(\omega)$ is *sub-additive*: $m_{B_i}(\omega)(E \cup F) \leq m_{B_i}(\omega)(E) + m_{B_i}(\omega)(F)$. Thus, (2a)-(2f) yield *finitely-additive* (both super-additive and sub-additive) beliefs.

Condition (2g) guarantees that each $m_{B_i}(\omega)$ is *continuous from above*: $E_n \downarrow E$ implies $m_{B_i}(\omega)(E_n) \rightarrow m_{B_i}(\omega)(E)$. Condition (2h) guarantees that each $m_{B_i}(\omega)$ is *continuous from below*: $E_n \uparrow E$ implies $m_{B_i}(\omega)(E_n) \rightarrow m_{B_i}(\omega)(E)$. By (2g) or (2h), a finitely-additive probability measure $m_{B_i}(\omega)$ becomes *countably additive*. I have presented both (2g) and (2h) to accommodate non-additive beliefs.

Condition (2i) is the introspective property that player i is certain of her own beliefs. The set $[m_{B_i}(\omega)]$ consists of states ω' that player i cannot distinguish from ω in that $[m_{B_i}(\omega)] = \{\omega' \in \Omega \mid m_{B_i}(\omega') = m_{B_i}(\omega)\}$. Thus, player i is certain of her beliefs in that she believes E with probability one if $[m_{B_i}(\omega)]$ implies (i.e., is a subset of) E . Especially, $B_i^p(E) \subseteq B_i^1 B_i^p(E)$ and $(\neg B_i^p)(E) \subseteq B_i^1 (\neg B_i^p)(E)$ hold: if player i p -believes (does not p -believe) an event E , then she believes with probability one that she p -believes (does not p -believe) the event E .

A (*probabilistic belief*) *morphism* from $\overrightarrow{\Omega}$ to $\overrightarrow{\Omega}'$ is a measurable map $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$ satisfying: (i) $\Theta = \Theta' \circ \varphi$; and (ii) $B_i^p(\varphi^{-1}(\cdot)) = \varphi^{-1}(B_i'^p(\cdot))$ for all $(i, p) \in I \times$

$[0, 1]$. A probabilistic-belief space $\overrightarrow{\Omega}^*$ of I on (S, \mathcal{S}) is *terminal* if, for any probabilistic-belief space $\overrightarrow{\Omega}$ of I on (S, \mathcal{S}) , there is a unique morphism $\varphi : \overrightarrow{\Omega} \rightarrow \overrightarrow{\Omega}^*$.

I show that a terminal probabilistic-belief space exists. One can also extend it to various notions of beliefs (e.g., non-additive beliefs and finitely-additive beliefs) by dropping corresponding conditions in Definition 10 (2).

Corollary 2 (Terminal Probabilistic-Belief Space). *There exists a terminal probabilistic-belief space $\overrightarrow{\Omega}^*$ of I on (S, \mathcal{S}) .*

In the construction of the terminal belief space for a generic class of belief spaces in Section 3, the set $\mathcal{L} = \mathcal{L}_\kappa^I(\mathcal{A}_\kappa(\mathcal{S}))$ of expressions is infinitary (i.e., involves infinite conjunctions) when an infinite regular cardinal κ satisfies $\kappa > \aleph_0$. Each belief hierarchy in the terminal space $\overrightarrow{\Omega}^*$ in Corollary 2 consists of all countable-level beliefs. In contrast, each type of a terminal type space in the literature (e.g., Brandenburger and Dekel, 1993; Heifetz and Samet, 1998b; Mertens and Zamir, 1985) is a belief hierarchy consisting of all finite-level beliefs.

I show that the terminal probabilistic-belief space can be constructed when one restricts attention to all finite-level belief hierarchies. The intuition is that, by the continuity (countable-additivity) of beliefs, all finite-level belief hierarchies (that can be attained by some state of some belief space) uniquely extend to countable levels.

To that end, for the rest of this subsection, let $(\kappa, \lambda) = (\aleph_1, \aleph_0)$. Recalling Definition 5 and Remark 3 in Section 3, this means that I consider λ -expressions $\mathcal{L}_\lambda^I(\mathcal{A}_\kappa(\mathcal{S}))$: syntactic formulas that express nature and *finite-level* interactive beliefs. Applying the construction of a terminal space in Section 3 while the language is restricted to $\mathcal{L}_\lambda^I(\mathcal{A}_\kappa(\mathcal{S}))$, the proof of Lemma 1 implies that $\{[e] \in \mathcal{D}^* \mid e \in \mathcal{L}_\lambda^I(\mathcal{A}_\kappa(\mathcal{S}))\}$ is an algebra on Ω^* . I show below that the smallest σ -algebra including this set is \mathcal{D}^* .

Proposition 4 (Extension of Finitary Language). $\mathcal{D}^* = \sigma(\{[e] \in \mathcal{D}^* \mid e \in \mathcal{L}_\lambda^I(\mathcal{A}_\kappa(\mathcal{S}))\})$.

Proposition 4 implies that the terminal probabilistic-belief space can be constructed when the language is restricted to finite-level beliefs $\mathcal{L}_\lambda^I(\mathcal{A}_\kappa(\mathcal{S}))$. The players' countably-additive beliefs are well-defined on \mathcal{D}^* because a countably-additive probability measure defined on the algebra (the generator of the right-hand side) admits a unique extension to the generated σ -algebra (the left-hand side, which is \mathcal{D}^*).

Below, I provide an informal discussion of the fact that the terminal probabilistic-belief space can be constructed when the language is restricted to finite-level beliefs $\mathcal{L}_\lambda^I(\mathcal{A}_\kappa(\mathcal{S}))$. Take any probabilistic-belief space $\overrightarrow{\Omega}$. On the one hand, Section 3 shows that a unique morphism $D : \overrightarrow{\Omega} \rightarrow \overrightarrow{\Omega}^*$ associates, with each state ω , the corresponding nature state and the players' belief hierarchies that describe all countable-level interactive beliefs at ω . On the other hand, the hierarchical construction of a terminal type space in the literature (denote it by $\overrightarrow{\Omega}^{**}$) implies that there exists a unique morphism $h : \overrightarrow{\Omega} \rightarrow \overrightarrow{\Omega}^{**}$ which associates, with each state $\omega \in \Omega$, the corresponding nature state and the profile of the players' belief hierarchies (each player i 's first-order belief

$m_{B_i}(\omega) \circ \Theta^{-1}$ over S , each player's second-order belief, and so on) of all finite levels. Since both spaces are terminal, there exists a unique isomorphism $\varphi : \overrightarrow{\Omega^*} \rightarrow \overrightarrow{\Omega^{**}}$ such that $h = \varphi \circ D$. Hence, φ strips away transfinite levels of beliefs from $D(\omega)$ to generate the profile of the players' finite-level belief hierarchies $h(\omega) = \varphi(D(\omega))$ associated with $D(\omega)$.

Corollary 2 and Proposition 4 imply that one can separately ask the question of whether there exists a terminal space (Corollary 2) and that of whether the terminal space consists of all finite-level beliefs (Proposition 4). The continuity (countable-additivity) of probabilistic beliefs has to do with the latter question, not the former.⁴⁸

6 Game-Theoretic Applications

This section shows that the specification of an infinite (regular) cardinal κ , which determines depth of reasoning $\bar{\kappa}$, is also important in game-theoretic applications. Section 6.1 introduces the notion of common belief into a belief space, and demonstrates that, in a category of belief spaces in which common belief is well-defined, a terminal space exists. Section 6.2 defines strategic games with ordinal payoffs.

Section 6.3 characterizes iterated elimination of strictly dominated actions (IESDA) as an implication of common belief in rationality in a terminal qualitative-belief space for an arbitrary strategic game with ordinal payoffs. The subsection also provides an example of a game in which a unique prediction under IESDA requires an arbitrarily long process of IESDA.⁴⁹ Hence, in a situation in which the players reason about their actions in a given game, one would need to fix an appropriate level $\bar{\kappa}$ to accommodate all possible reasoning about the given game.

Section 6.4 shows that common belief in weak-dominance rationality characterizes the iterated elimination procedure first studied by Börgers (1993) for an arbitrary strategic game with ordinal payoffs.⁵⁰

6.1 Common Belief

I show that a terminal space exists in a class of belief spaces with common belief. To that end, I incorporate the notion of common belief, irrespective of a choice of κ

⁴⁸Corollary 2 and Proposition 4 imply that the previous non-existence result on a terminal qualitative belief space does not have to do with the continuity of beliefs, contrary to the claims of the previous literature (e.g., Fagin, 1994; Fagin, Halpern, and Vardi, 1991; Fagin et al., 1999; Heifetz and Samet, 1998a,b, 1999).

⁴⁹The example in Section 6.3 involves infinite action sets. Appendix E.1 provides an example of a strategic game with finitely many actions in which infinitely many players interactively reason so that a unique prediction under IESDA requires an arbitrarily long elimination process.

⁵⁰First, Appendix E.2 studies a strategic game in which a unique prediction requires an arbitrarily long elimination process. Second, Appendix E.3 contrasts belief and knowledge by showing that common knowledge of weak-dominance rationality characterizes the pure-strategy version of the iterated elimination procedure first studied by Stalnaker (1994).

and properties of beliefs, following Fukuda (2020). The definition of common belief does not resort to the chain of mutual beliefs. Thus, one can analyze players who fail logical reasoning (e.g., Monotonicity or λ -Conjunction) or players who reason only about finite levels of interactive beliefs (i.e., $\kappa = \aleph_0$).

Fix a non-empty set I of players, and let κ be an infinite regular cardinal with $\kappa > |I|$. By this assumption, in any κ -belief space $\langle(\Omega, \mathcal{D}), (B_i)_{i \in I}, \Theta\rangle$, one can define the *mutual belief operator* $B_I : \mathcal{D} \rightarrow \mathcal{D}$ by $B_I(\cdot) := \bigcap_{i \in I} B_i(\cdot)$.

An event E is a *common basis* if everybody believes any logical implication of E whenever E is true: $E \subseteq F$ implies $E \subseteq B_I(F)$ (Fukuda, 2020). If the mutual belief operator B_I satisfies Monotonicity, then E is a common basis if(f) it is publicly-evident: $E \subseteq B_I(E)$. Denote by \mathcal{J}_I the collection of common bases.

An event E is *common belief* at a state ω if there is a common basis $F \in \mathcal{J}_I$ which is true at ω and which implies the mutual belief in E : $\omega \in F \subseteq B_I(E)$. If B_I satisfies Monotonicity, then this definition of common belief reduces to Monderer and Samet (1989).

The common belief in E at a state ω implies the chain of mutual beliefs in E at that state: at ω , everybody believes E (i.e., $\omega \in B_I(E)$), everybody believes that everybody believes E (i.e., $\omega \in B_I B_I(E)$), and so on *ad infinitum*. The converse holds (formally, the common belief in E reduces to $\bigcap_{n \in \mathbb{N}} B_I^n(E)$) when, for example, B_I satisfies Monotonicity and \aleph_1 -Conjunction as in a possibility correspondence model (Fukuda, 2020).⁵¹

A κ -belief space with a common belief operator is a tuple $\overrightarrow{\Omega} := \langle(\Omega, \mathcal{D}), (B_i)_{i \in I}, C, \Theta\rangle$ such (i) that $\langle(\Omega, \mathcal{D}), (B_i)_{i \in I}, \Theta\rangle$ is a κ -belief space (satisfying given properties of beliefs) and (ii) that $C : \mathcal{D} \rightarrow \mathcal{D}$ satisfies $C(E) = \max\{F \in \mathcal{J}_I \mid F \subseteq B_I(E)\}$ for each $E \in \mathcal{D}$, where “max” is taken with respect to the set inclusion.⁵² I show that a terminal space exists within the class of κ -belief spaces with a common belief operator irrespective of κ and properties of beliefs.

Corollary 3 (Terminal Belief Space with a Common Belief Operator). *There exists a terminal κ -belief space $\overrightarrow{\Omega}^*$ of I on (S, \mathcal{S}) with a common belief operator.*

6.2 Strategic Games with Ordinal Payoffs

For the rest of Section 6, I consider strategic games with ordinal payoffs. A strategic game with ordinal payoffs is a tuple $\Gamma := \langle(A_i)_{i \in I}, (\succ_i)_{i \in I}\rangle$ satisfying the following: each A_i is player i ’s (non-empty) action set, and each \succ_i is player i ’s (complete and transitive) preference relation on $A := \prod_{i \in I} A_i$. Denote by \sim_i and \succ_i the indifference and strict preference relations, respectively. For ease of exposition, sometimes

⁵¹Also, in a probabilistic-belief space (recall Definition 10), let $B_I^p(\cdot) := \bigcap_{i \in I} B_i^p(\cdot)$ be the mutual p -belief operator. It can be seen that the common p -belief operator C^p coincides with the iteration of mutual p -beliefs: $C^p(\cdot) = \bigcap_{n \in \mathbb{N}} (B_I^p)^n(\cdot)$.

⁵²Then, we have $C(E) = \{\omega \in \Omega \mid \omega \in F \subseteq B_I(E) \text{ for some } F \in \mathcal{J}_I\}$ for each $E \in \mathcal{D}$.

I introduce a strategic game by specifying the players' payoff functions, that is, as a tuple $\Gamma := \langle (A_i)_{i \in I}, (u_i)_{i \in I} \rangle$, where each player's payoff function u_i represents her underlying preferences \succsim_i .

Since the players reason about their actions, for the rest of Section 6, I consider $S = A$ and an infinite regular cardinal κ with $\max(|I|, |A|) < \kappa$ for a given strategic game Γ .⁵³ In order for each player to be able to reason about her own actions, I assume $\{a_i\} \times A_{-i} \in \mathcal{S}$ for each $i \in I$ and $a_i \in A_i$. Since $\kappa > |A|$, the assumption amounts to $\mathcal{A}_\kappa(\mathcal{S}) = \mathcal{P}(S)$.

In a κ -belief space (with a common belief operator) $\langle (\Omega, \mathcal{D}), (B_i)_{i \in I}, C, \Theta \rangle$, the measurable map Θ is decomposed into $\Theta = (\Theta_i)_{i \in I}$, where $\Theta_i : \Omega \rightarrow A_i$ is interpreted as player i 's (behavioral) strategy: it associates, with each state ω , the corresponding action $\Theta_i(\omega) \in A_i$.

For each $a_i \in A_i$, $\Theta_i^{-1}(\{a_i\}) = \llbracket \{a_i\} \times A_{-i} \rrbracket_{\vec{\Omega}} \in \mathcal{D}$ is the event that player i plays a_i . Player i is *certain of her own strategy* if $\Theta_i^{-1}(\{a_i\}) \subseteq B_i(\Theta_i^{-1}(\{a_i\}))$ for all $a_i \in A_i$.

For each $(a'_i, a_i) \in A_i^2$, define $[a'_i \succ_i a_i] := \{a_{-i} \in A_{-i} \mid (a'_i, a_{-i}) \succ_i (a_i, a_{-i})\}$ and $\llbracket a'_i \succ_i a_i \rrbracket_{\vec{\Omega}} := (\Theta_{-i})^{-1}([a'_i \succ_i a_i]) \in \mathcal{D}$. The set $\llbracket a'_i \succ_i a_i \rrbracket_{\vec{\Omega}}$ is the event that player i strictly prefers a'_i to a_i given the other players' strategies. The sets $[a'_i \succsim_i a_i]$ and $\llbracket a'_i \succsim_i a_i \rrbracket_{\vec{\Omega}}$ are similarly defined.

6.3 Iterated Elimination of Strictly Dominated Actions

This subsection studies the solution concept of iterated elimination of strictly dominated actions (IESDA) in a terminal belief space as an implication of common belief in rationality for a game with ordinal payoffs and qualitative beliefs.

6.3.1 Common Belief in Rationality

I start with defining the notion of rationality, and show that: within a class of belief spaces with common belief, (i) players' rationality is well-defined in the terminal space, and (ii) common belief in rationality in any given belief space is preserved in the terminal space. Especially, the terminal space itself characterizes the solution concept of IESDA as an implication of common belief in rationality.

Definition 11 (Rationality). Let $\vec{\Omega}$ be a belief space. Player i is *rational* (e.g., Bonanno, 2008; Chen, Long, and Luo, 2007) at a state $\omega \in \Omega$ if, for no action a'_i , she believes at ω playing a'_i is strictly better than $\Theta_i(\omega)$ given the opponents' strategies:

$$\omega \in B_i(\llbracket a'_i \succ_i \Theta_i(\omega) \rrbracket_{\vec{\Omega}}).$$

Denote by RAT_i (or $\text{RAT}_i^{\vec{\Omega}}$) the set of states at which player i is rational. Likewise, let $\text{RAT}_I := \bigcap_{i \in I} \text{RAT}_i$.

⁵³Recalling the technical preliminaries in Section 2.1, one can always take the successor cardinal $\kappa = (\max(|I|, |A|))^+$.

It can be seen that each player's rationality $\text{RAT}_i^{\vec{\Omega}}$ is a well-defined event, as

$$\text{RAT}_i^{\vec{\Omega}} = \bigcap_{a_i \in A_i} ((\Theta_i^{-1}(\{a_i\}))^c \cup \bigcap_{a'_i \in A_i} (\neg B_i)(\llbracket a'_i \succ_i a_i \rrbracket_{\vec{\Omega}})) \in \mathcal{D}.$$

Moreover, a morphism preserves the notion of rationality. If $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}'$ is a morphism, then $\varphi^{-1}(\llbracket a'_i \succ_i a_i \rrbracket_{\vec{\Omega}'}) = \llbracket a'_i \succ_i a_i \rrbracket_{\vec{\Omega}}$ and $\varphi^{-1}(\text{RAT}_i^{\vec{\Omega}'}) = \text{RAT}_i^{\vec{\Omega}}$. Especially, $D^{-1}(C^*(\text{RAT}_I^{\vec{\Omega}^*})) = C(\text{RAT}_I^{\vec{\Omega}})$. Thus, for any state ω^* in the terminal space $\vec{\Omega}^*$ and for any state $\omega \in \Omega$ of a belief space $\vec{\Omega}$ with $\omega^* = D(\omega)$, the players commonly believe their rationality at ω in the belief space $\vec{\Omega}$ iff they commonly believe their rationality at $\omega^* = D(\omega)$ in the terminal space.

Next, I define a process of IESDA, following Chen, Long, and Luo (2007). To that end, I start with the notion of strictly dominated actions. Since this subsection studies an *arbitrary* strategic game with *ordinal* payoffs, it only allows for elimination of actions that are dominated by pure actions (see also Chen, Long, and Luo, 2007; Dufwenberg and Stegeman, 2002).

Definition 12 (Strict Dominance). Let $\langle (A_i)_{i \in I}, (\succ_i)_{i \in I} \rangle$ be a strategic game. Let X_{-i} be a non-empty subset of the set A_{-i} of the actions other than player i . An action $a_i \in A_i$ is *strictly dominated* given X_{-i} if there exists an action $\hat{a}_i \in A_i$ such that $(\hat{a}_i, x_{-i}) \succ_i (a_i, x_{-i})$ for all $x_{-i} \in X_{-i}$.

Definition 13 (IESDA). A process of *iterated elimination of strictly dominated actions* (IESDA) is an ordinal sequence of A^α (with $|\alpha| \leq |A|$) defined as follows: (i) $A^0 = A$; (ii) for a successor ordinal $\alpha = \beta + 1$, A^α is obtained by eliminating *at least one* action $a_i \in A_i^\beta$ which is strictly dominated given A_{-i}^β ; and (iii) for a non-zero limit ordinal α , $A^\alpha = \bigcap_{\beta < \alpha} A^\beta$. Since $(A^\alpha)_\alpha$ is weakly decreasing, let α (with $|\alpha| \leq |A|$) be the smallest ordinal with $A^\alpha = A^{\alpha+1}$. An action profile $a \in A$ *survives* the process of IESDA if $a \in A^{\text{IESDA}} := A^\alpha$. Call A^{IESDA} the *terminal set* of the process of IESDA.

A process of IESDA is order-independent, that is, the terminal set A^{IESDA} is uniquely determined (Chen, Long, and Luo, 2007).⁵⁴

Below, for a given strategic game, I characterize IESDA as an implication of common belief in rationality in κ -belief spaces in which $\kappa > \max(|I|, |A|)$ and in which the players have *correct common belief in rationality*: $C(\text{RAT}_I) \subseteq \text{RAT}_I$.⁵⁵

⁵⁴I adopt this order-independent elimination procedure also to compare it with other order-independent elimination procedures in Section 6.4 and Appendix E.3.

⁵⁵Proposition 5 also holds for each of the following three cases. First, if some player's belief operator B_i satisfies Truth Axiom, then common belief satisfies Truth Axiom. Second, if each B_i satisfies the Kripke property and Consistency and if each player is certain of her own strategy (i.e., $\Theta_i^{-1}(\{a_i\}) \subseteq B_i(\Theta_i^{-1}(\{a_i\}))$), then (each player correctly believes her own rationality and consequently) the players have correct common belief in rationality. Third, an alternative characterization in terms of rationality and common belief in rationality (formally, $\text{RAT}_I \cap C(\text{RAT}_I)$), which does not require correct common belief in rationality, also holds when each B_i satisfies Monotonicity and \aleph_0 -Conjunction.

Proposition 5 (IESDA). *Fix a strategic game and $\kappa > \max(|I|, |A|)$.*

1. *Take any κ -belief space $\vec{\Omega}$ in which the players have correct common belief in rationality. If $\omega \in C(\text{RAT}_I)$ then $\Theta(\omega) \in A^{\text{IESDA}}$.*
2. *For any $a \in A^{\text{IESDA}}$, there exist a κ -belief space $\vec{\Omega}$ and a state $\omega \in \Omega$ such that: the players have correct common belief in rationality; $\Theta(\omega) = a$; and that $\omega \in C(\text{RAT}_I)$.*

The proof is similar to that of Fukuda (2020, Theorem 3) and is omitted. Part (1) states that, for any given κ -belief space, the players' actions at states at which they have common belief in rationality survive any process of IESDA but may not necessarily exhaust the entire predictions A^{IESDA} . Part (2) states that the entire predictions A^{IESDA} can be obtained if all belief spaces in the given category are considered.

Next, I restate the epistemic characterization of Proposition 5 on the terminal κ -belief space $\vec{\Omega}^*$ (in which the players have correct common belief in rationality) instead of considering all belief spaces in the given category. In the terminal κ -belief space $\vec{\Omega}^*$, (i) for any κ -belief space $\vec{\Omega}$ with $\omega \in C(\text{RAT}_I^{\vec{\Omega}})$, $D_{\vec{\Omega}}(\omega) \in C^*(\text{RAT}_I^{\vec{\Omega}^*})$ and $\Theta^*(D_{\vec{\Omega}}(\omega)) = \Theta(\omega) \in A^{\text{IESDA}}$; and (ii) conversely, for any $a \in A^{\text{IESDA}}$, there exists a state $\omega^* \in \Omega^*$ with $a = \Theta^*(\omega^*)$ and $\omega^* \in C^*(\text{RAT}_I^{\vec{\Omega}^*})$. The first part states that for any state of a belief space in which the players have (correct) common belief in rationality, they have (correct) common belief in rationality and their actions survive any process of IESDA in the corresponding state of the terminal space. The second part says that, for any action profile $a \in A^{\text{IESDA}}$, there exists a state in the terminal space at which the players, who commonly believe their rationality, take the given actions. Thus, it is sufficient to consider the terminal space.

Corollary 4 (IESDA on the Terminal Space). *Fix any strategic game and $\kappa > \max(|I|, |A|)$. There exists a terminal κ -belief space $\vec{\Omega}^*$ in which the players have correct common belief in rationality. On the terminal space,*

$$A^{\text{IESDA}} = \{a \in A \mid a = \Theta^*(\omega^*) \text{ for some } \omega^* \in C^*(\text{RAT}_I)\}.$$

On the one hand, the terminal κ -belief space can provide the exact characterization of IESDA as an implication of common belief in rationality. On the other hand, the second part of the corollary is indeed equivalent to Proposition 5.

Corollary 4 also clarifies properties of players' beliefs under which common belief in rationality characterizes IESDA, the direction which has not been explored before. First, one assumption is correct common belief in rationality, as Proposition 5 or Corollary 4 may not necessarily hold without the restriction that the players have correct common belief in rationality (Fukuda, 2020). Second, another epistemic characterization of IESDA is rationality and common belief in rationality when the players' belief operators satisfy Monotonicity and \aleph_0 -Conjunction (see footnote 55).

Proposition 5 and Corollary 4 hold in the category of belief spaces in which the players' beliefs satisfy these logical properties. Without these logical properties of beliefs, the characterization may not necessarily hold.

6.3.2 Example of a Game with a Transfinite Process of IESDA

To conclude the subsection, for an arbitrary limit ordinal α , I provide an example of a strategic game in which a unique prediction under IESDA is obtained after $\alpha + 1$ iterations. Hence, an epistemic analysis of a strategic game may necessitate appropriately choosing an infinite regular cardinal κ with $\kappa > \max(|I|, |A|)$.

Let α be a non-zero limit ordinal. Define a strategic game $\langle (A_i)_{i \in I}, (u_i)_{i \in I} \rangle$ (for ease of exposition, in terms of payoff functions) as follows. Let $A_i := \alpha + 2$ (i.e., $A_i = \{0, 1, \dots, \alpha, \alpha + 1\}$) be the set of actions available to $i \in I := \{1, 2\}$. Define i 's payoff function $u_i : A_i \times A_{-i} \rightarrow \mathbb{R}$ as

$$u_i(a_i, a_{-i}) := \begin{cases} 0 & \text{if } a_i < a_{-i} \text{ or } a_i = a_{-i} \neq \alpha + 1 \\ 2 & \text{if } a_{-i} < a_i \neq \alpha + 1 \\ 1 & \text{if } a_i = \alpha + 1 \end{cases}.$$

Table 1 in the Introduction depicts $u_i(a_i, a_{-i})$ when α is the least infinite ordinal (i.e., the set of non-negative integers). Action $\alpha + 1$ always yields a payoff of 1 irrespective of the opponent's action. For any other action, player i obtains a payoff of 2 if her action a_i is (strictly) higher than the opponent's, and she obtains a payoff of 0 otherwise.

Any process of IESDA yields a unique action profile $(a_1, a_2) = (\alpha + 1, \alpha + 1)$. For instance, the process of eliminating all strictly dominated actions at each step leads to the unique prediction $(\alpha + 1, \alpha + 1)$ after the $\alpha + 1$ round of elimination.

The rest of this subsection constructs a belief space $\langle (\Omega, \mathcal{D}), (B_i)_{i \in I}, C, \Theta \rangle$ in which common belief in rationality is attained exactly as the $\alpha + 1$ iterations of mutual beliefs in rationality: $C(\text{RAT}_I) = \bigcap_{\beta: 1 \leq \beta \leq \alpha+1} B_I^\beta(\text{RAT}_I)$ while $C(\text{RAT}_I) \subsetneq \bigcap_{\beta: 1 < \beta \leq \alpha} B_I^\beta(\text{RAT}_I)$.⁵⁶

First, let $(\Omega, \mathcal{D}) := (\alpha + 3, \mathcal{P}(\Omega))$. Second, define each B_i as (i) $B_i(E) := E$ if $E \in \{\emptyset, \{\alpha + 2\}, \Omega\}$; and (ii) $B_i(E) := E \setminus \{\min(E)\}$ otherwise.⁵⁷ Each player does not believe a contradiction of the form \emptyset , always believes a tautology of the form Ω , and believes $\{\alpha + 2\}$ at $\alpha + 2$. She believes any other event E at any state in E except at $\min E$. Note that B_i satisfies Truth Axiom and Monotonicity but

⁵⁶If each player's belief operator satisfies Monotonicity and \aleph_1 -Conjunction, then the least infinite chain of mutual beliefs $\bigcap_{n \in \mathbb{N}} B_I^n(\text{RAT}_I)$ already converges to common belief $C(\text{RAT}_I)$. Two remarks are in order. First, even in such a model, the model itself may need to be defined on a κ -algebra (with $\kappa > |A|$) to accommodate the $\alpha + 1$ iterations of mutual beliefs to reflect transfinite processes of IESDA. Second, in this example, common belief is not characterized by all the finite iterations of mutual beliefs as the players' belief operators violate \aleph_1 -Conjunction.

⁵⁷As Ω is an ordinal, $\min(E)$ is well-defined for any non-empty subset E of Ω .

fails \aleph_0 -Conjunction. Third, the common belief (or, common knowledge) operator $C : \mathcal{D} \rightarrow \mathcal{D}$, by definition, satisfies: (i) $C(\Omega) = \Omega$; (ii) $C(E) = \{\alpha + 2\}$ if $\alpha + 2 \in E$ and $E \neq \Omega$; and (iii) $C(E) = \emptyset$ if $\alpha + 2 \notin E$. Fourth, define $\Theta = (\Theta_i)_{i \in I}$ as: (i) $\Theta_i(\omega) := \alpha + 1$ if $\omega \neq 0$; and (ii) $\Theta_i(0) := 0$.

Each player is rational at each state except at 0: $\text{RAT}_i = \Omega \setminus \{0\}$. To find mutual beliefs in rationality, for any ordinal $\beta \leq \alpha + 2$, denote by $[\beta, \alpha + 2]$ the set of ordinals γ with $\beta \leq \gamma \leq \alpha + 2$. Then, for any ordinal β with $1 \leq |\beta| < \kappa$, $B_I^\beta(\text{RAT}_I) = [\beta, \alpha + 2]$ if $\beta < \alpha + 1$, and $B_I^\beta(\text{RAT}_I) = \{\alpha + 2\}$ if $\beta \geq \alpha + 1$. Hence, $C(\text{RAT}_I) = \bigcap_{\beta: 1 \leq \beta \leq \alpha + 1} B_I^\beta(\text{RAT}_I)$, while $C(\text{RAT}_I) \subsetneq \bigcap_{\beta: 1 \leq \beta \leq \alpha} B_I^\beta(\text{RAT}_I)$. That is, the chain of mutual beliefs in rationality converges to common belief in rationality at the $\alpha + 1$ -th round. Common belief in rationality $C(\text{RAT}_I)$ captures IESDA: $\Theta(\alpha + 2) = (\alpha + 1, \alpha + 1) \in A^{\text{IESDA}}$. The terminal κ -belief space of a category to which the example belief space belongs contains the belief hierarchies that the example belief space induces.

6.4 Iterated Elimination of Börgers Dominated Actions

This subsection provides an epistemic characterization of a pure-strategy dominance first studied by Börgers (1993): an action is Börgers dominated if it is weakly dominated by some pure action when opponents' play is restricted to arbitrary non-empty subsets of their actions. The idea behind Börgers dominance is that players' preferences (that represent their von Neumann-Morgenstern utility functions) over game outcomes are ordinal. He assumes that players' ordinal preferences are "commonly certain" (in the informal sense), each player forms probabilistic beliefs about the opponents' choices, and that each player chooses a pure action that maximizes her expected utility. He shows that, in this setting, a pure action is rationalized iff it is not dominated. He then studies action profiles that are consistent with common belief in rationality.

The subsection starts with the definition of the solution concept.

Definition 14 (Börgers Dominance). Let $\langle (A_i)_{i \in I}, (\succ_i)_{i \in I} \rangle$ be a strategic game.

1. Let X_{-i} be a non-empty subset of the set A_{-i} of the actions other than player i . An action $a_i \in A_i$ is *weakly dominated* given X_{-i} if there exists $\hat{a}_i \in A_i$ such that (i) $(\hat{a}_i, x_{-i}) \succ_i (a_i, x_{-i})$ for all $x_{-i} \in X_{-i}$; and that (ii) $(\hat{a}_i, \hat{x}_{-i}) \succ_i (a_i, \hat{x}_{-i})$ for some $\hat{x}_{-i} \in X_{-i}$.
2. An action $a_i \in A_i$ is *Börgers-dominated* (henceforth, *B-dominated*) if, for every non-empty subset $X_{-i} \subseteq A_{-i}$, the action a_i is weakly dominated given X_{-i} .

Definition 15 (IEBA). A process of *iterated elimination of B-dominated actions* (IEBA) is an ordinal sequence of A^α (with $|\alpha| \leq |A|$) defined as follows: (i) $A^0 = A$; (ii) for a successor ordinal $\alpha = \beta + 1$, A^α is obtained by eliminating *at least one* action $a_i \in A_i^\beta$ which is B-dominated given A_{-i}^β ; and (iii) for a non-zero limit ordinal

α , $A^\alpha = \bigcap_{\beta < \alpha} A^\beta$. Since $(A^\alpha)_\alpha$ is weakly decreasing, let α (with $|\alpha| \leq |A|$) be the smallest ordinal with $A^\alpha = A^{\alpha+1}$. An action profile $a \in A$ *survives* the process of IEBA if $a \in A^{\text{IEBA}} := A^\alpha$. Call A^{IEBA} the *terminal set* of the process of IEBA.

Note that, at the $(\beta + 1)$ -th step of a process of IEBA, dominating strategies are not necessarily restricted to A_i^β . With this in mind:

Lemma 5 (Order-Independence of IEBA). *A process of IEBA is order-independent, that is, A^{IEBA} is uniquely determined.*

I characterize the solution concept of IEBA as an implication of common belief in weak-dominance rationality when players' qualitative beliefs satisfy Consistency and the Kripke property.

Definition 16 (Weak-Dominance Rationality). Let $\vec{\Omega}$ be a belief space. Player i is *weak-dominance rational* at a state $\omega \in \Omega$ if, for no action $a_i \in A_i$,

$$\omega \in B_i(\llbracket a_i \succ_i \Theta_i(\omega) \rrbracket_{\vec{\Omega}}) \cap (\neg B_i)(\neg \llbracket a_i \succ_i \Theta_i(\omega) \rrbracket_{\vec{\Omega}}).$$

Denote by WDRAT_i (or $\text{WDRAT}_i^{\vec{\Omega}}$) the set of states at which player i is weak-dominance rational. Likewise, let $\text{WDRAT}_I := \bigcap_{i \in I} \text{WDRAT}_i$.

The intended meaning of weak-dominance rationality is as follows. Player i is weak-dominance rational at ω if, for no action $a_i \in A_i$, player i believes that playing a_i is as good as playing $\Theta_i(\omega)$ and she considers it possible that playing a_i is strictly better than playing $\Theta_i(\omega)$.

The intended meaning is best interpreted under Consistency and the Kripke property. Recall that $b_{B_i}(\omega)$ is the set of states that player i considers possible (the discussions after Definition 1). Under Consistency and the Kripke property, player i considers an event E *possible* if she does not believe its negation E^c (as $b_{B_i}(\omega) \cap E \neq \emptyset$). Also, if player i believes an event then she considers the event possible. Thus, the intended meaning of weak-dominance rationality is obtained. Moreover, player i is weak-dominance rational at ω if her action $\Theta_i(\omega)$ is not weakly dominated given the set of actions $\Theta_{-i}(b_{B_i}(\omega))$ she considers possible at ω .

Weak-dominance rationality is a stronger notion of rationality, i.e., $\text{WDRAT}_i \subseteq \text{RAT}_i$, under Consistency and Monotonicity (which is implied by the Kripke property).

In any belief space, the set $\text{WDRAT}_i^{\vec{\Omega}}$ is an event because

$$\text{WDRAT}_i^{\vec{\Omega}} = \bigcap_{a_i \in A_i} \left((\Theta_i^{-1}(\{a_i\}))^c \cup \bigcap_{a'_i \in A_i} ((\neg B_i)(\llbracket a'_i \succ_i a_i \rrbracket_{\vec{\Omega}}) \cup B_i(\neg \llbracket a'_i \succ_i a_i \rrbracket_{\vec{\Omega}})) \right) \in \mathcal{D}.$$

Moreover, a morphism preserves weak-dominance rationality. If $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}'$ is a morphism, then $\varphi^{-1}(\text{WDRAT}_i^{\vec{\Omega}'}) = \text{WDRAT}_i^{\vec{\Omega}}$. Especially, $D^{-1}(C^*(\text{WDRAT}_I^{\vec{\Omega}})) =$

$C(WDRAT_I^{\vec{\Omega}})$: if the players commonly believe their weak-dominance rationality at ω in a given belief space $\vec{\Omega}$, then so do they at $\omega^* = D(\omega)$ in the terminal space.

The proposition below provides an epistemic characterization of IEBA in terms of common belief in weak-dominance rationality when the players' beliefs satisfy the Kripke property and Consistency. The proposition generalizes the epistemic characterization by Bonanno and Tsakas (2018, Theorem 1) (of a finite strategic game) to an arbitrary strategic game.

Proposition 6 (IEBA). *Fix a strategic game and $\kappa > \max(|I|, |A|)$.*

1. *Take any κ -belief space $\vec{\Omega}$ in which each B_i satisfies the Kripke property and Consistency and in which each player is certain of her own strategy: $\Theta_i^{-1}(\{a_i\}) \subseteq B_i(\Theta_i^{-1}(\{a_i\}))$. If $\omega \in C(WDRAT_I)$ then $\Theta(\omega) \in A^{IEBA}$.*
2. *For any $a \in A^{IEBA}$, there exist a κ -belief space $\vec{\Omega}$ and a state $\omega \in \Omega$ such that: each B_i satisfies the Kripke property and Consistency; each player is certain of her own strategy; $\Theta(\omega) = a$; and that $\omega \in C(WDRAT_I)$.*

The proof of Proposition 6 is similar to that of Bonanno and Tsakas (2018, Theorem 1) and is omitted.⁵⁸

Three remarks are in order. First, since one can prove Proposition 6 (2) by constructing a possibility correspondence model in which every player's belief satisfies Consistency, Positive Introspection and Negative Introspection, Proposition 6 holds with or without Positive Introspection and Negative Introspection, as long as Consistency and the Kripke property are imposed.

Second, the above observation implies that $A^{IEBA} \subseteq A^{IESDA}$. For any $a \in A^{IEBA}$ there exists a possibility correspondence model $\vec{\Omega}$ such that: every player's belief operator satisfies Consistency, Positive Introspection, and Negative Introspection; and that there exists a state $\omega \in C(WDRAT_I)$ at which the players commonly believe their weak-dominance rationality and their action profile satisfies $\Theta(\omega) = a$. Since each player's belief satisfies Monotonicity and Consistency, $WDRAT_i \subseteq RAT_i$ for each $i \in I$. Also, in this model, it can be seen that each player has correct belief in her rationality: $B_i(RAT_i) \subseteq RAT_i$ and consequently they have correct common belief in their rationality. Thus, it follows from Proposition 5 that $a = \Theta(\omega) \in A^{IESDA}$. Appendix E.2 provides an example of a strategic game for which (i) the converse set inclusion does not hold and for which (ii) a unique prediction under IEBA is obtained after α iterations for any given non-zero limit ordinal α .

Third, the epistemic characterization of IEBA may call for the failure of Truth Axiom. Bonanno and Tsakas (2018, Section 4.2) provide such an example.

Proposition 6 is restated as:

⁵⁸For the first part, one needs to extend the induction proof of Bonanno and Tsakas (2018) to transfinite induction. For the second part, their proof works without any substantial modification.

Corollary 5 (IEBA on the Terminal Space). *Fix any strategic game and $\kappa > \max(|I|, |A|)$. There exists a terminal κ -belief space $\vec{\Omega}^*$ in which each B_i satisfies the Kripke property and Consistency and in which each player is certain of her own strategy. Then,*

$$A^{\text{IEBA}} = \{a \in A \mid a = \Theta^*(\omega^*) \text{ for some } \omega^* \in C^*(\text{WDRAT}_I)\}.$$

The first part of the corollary is an implication of Theorem 1. The second part is equivalent to the following: for any $\omega^* \in C^*(\text{WDRAT}_I)$, $\Theta^*(\omega^*) \in A^{\text{IEBA}}$; and conversely, for any action profile $a \in A^{\text{IEBA}}$, there exists $\omega^* \in C^*(\text{WDRAT}_I)$ such that $a = \Theta^*(\omega^*)$. In words, an action profile a survives a process of IEBA iff a is played at some state in the terminal belief space at which the players commonly believe their weak-dominance rationality.

7 Comparison with the Previous Negative Results

Throughout this section, unless otherwise stated, fix an infinite regular cardinal κ . Recall that the framework of this paper admits the category of κ -knowledge spaces $\vec{\Omega}$ in which (Ω, \mathcal{D}) is a κ -algebra and in which each B_i is induced by a partitional possibility correspondence (recall Definition 2 and its discussions). Theorem 1 constructs a terminal κ -knowledge space $\vec{\Omega}^*$ of this category (i.e., $(\Omega^*, \mathcal{D}^*)$ is a κ -algebra and each B_i^* is induced by a partitional possibility correspondence): for any κ -knowledge space $\vec{\Omega}$ in this category, there is a unique morphism $D : \vec{\Omega} \rightarrow \vec{\Omega}^*$.⁵⁹

This section compares the existence of a terminal κ -belief space with the previous non-existence results (e.g., Fagin et al., 1999; Heifetz and Samet, 1998a; Meier, 2005). Section 7.1 introduces the notion of a complete algebra. The previous non-existence results hold in the category of belief spaces where the domain of each belief space is a complete algebra.

Section 7.2 examines whether the domain specification of a κ -belief space $\vec{\Omega}$ as a κ -algebra may neglect any reasoning regarding the underlying states Ω . It shows that the κ -algebra \mathcal{D} always contains any event that corresponds to a belief hierarchy of depth up to $\bar{\kappa}$ that the underlying states Ω intend to represent.

Section 7.3 shows that the terminal space $\vec{\Omega}^*$ contains all belief hierarchies of depth up to $\bar{\kappa}$. By defining the notion of a κ -rank, it shows that the terminal space attains the highest κ -rank (which is indeed $\bar{\kappa}$ in a category of κ -qualitative-belief spaces).

Section 7.4 shows that, while the cardinality of a terminal λ -belief space with $\lambda > \kappa$ has at least as high as that of the terminal κ -belief space, the terminal λ -belief space is redundant and is not minimal as a κ -belief space. Thus, the terminal κ -belief space consists of all possible belief hierarchies of depth up to $\bar{\kappa}$, and only up to $\bar{\kappa}$.

⁵⁹Also, in the category of κ -belief spaces $\vec{\Omega}$ in which each B_i is induced by a serial, transitive, and Euclidean possibility correspondence (i.e., B_i satisfies Consistency, Positive Introspection, Negative Introspection, and the Kripke property), Theorem 1 establishes a terminal κ -belief space.

7.1 Complete Algebras

For any given set X , call the sub-collection \mathcal{X} of $\mathcal{P}(X)$ (or the pair (X, \mathcal{X}) itself) a *complete algebra* if (i) $\{\emptyset, X\} \subseteq \mathcal{X}$ and if (ii) \mathcal{X} is closed under complementation and is closed under arbitrary union and intersection. In other words, (X, \mathcal{X}) is a complete algebra if it is a λ -algebra for any infinite (regular) cardinal λ . For any set X , $(X, \mathcal{P}(X))$ is a complete algebra.

In order to conveniently refer to both a κ -algebra (where κ is an infinite cardinal) and a complete algebra, I call a complete algebra an \aleph_∞ -algebra. The symbol \aleph_∞ (which is not a cardinal) is used only for indicating a complete algebra.⁶⁰ With this notation in mind, denote by $\mathcal{A}_{\aleph_\infty}(\cdot)$ the smallest complete algebra (i.e., the intersection of all complete algebras) including a given collection. I also explicitly call a belief space $\vec{\Omega}$ an \aleph_∞ -belief space if the domain \mathcal{D} is a complete algebra.

7.2 Informational Content of the Domain

For the terminal κ -belief space constructed in Section 3, the domain \mathcal{D}^* is generically not the power set of the underlying states Ω^* . Does the domain \mathcal{D}^* have some limitation on the representation of players' interactive beliefs? This subsection shows that any κ -belief space $\vec{\Omega}$ can capture the belief hierarchies of depth up to $\bar{\kappa}$ that the underlying states Ω intend to represent in the sense that the domain \mathcal{D} contains the events $\mathcal{D}_{\bar{\kappa}}$ generated by the expressions $\mathcal{L} = \mathcal{L}_{\bar{\kappa}}$.

On the one hand, the set $\mathcal{L} = \mathcal{L}_{\bar{\kappa}}$ represents nature states (S, \mathcal{S}) and players' belief hierarchies of depth up to $\bar{\kappa}$ in a way free from particular belief spaces. On the other hand, in an arbitrary κ -belief space $\vec{\Omega}$, the semantic interpretation function $[\![\cdot]\!]_{\vec{\Omega}}$ associates, with each expression $e \in \mathcal{L}$, the corresponding event $[\![e]\!]_{\vec{\Omega}} \in \mathcal{D}$ in the space $\vec{\Omega}$. Thus, $\mathcal{D}_{\bar{\kappa}} = \{[\![e]\!]_{\vec{\Omega}} \in \mathcal{D} \mid e \in \mathcal{L}\}$ is the collection of events in $\vec{\Omega}$ that correspond to the nature states (S, \mathcal{S}) and players' belief hierarchies of depth up to $\bar{\kappa}$.⁶¹ Since $\mathcal{D}_{\bar{\kappa}} \subseteq \mathcal{D}$, the given belief space $\vec{\Omega}$ can capture the belief hierarchies of depth up to $\bar{\kappa}$ that the underlying state space intends to represent.

This leads to the observation that any belief hierarchy of depth up to $\bar{\kappa}$ generated by a given κ -belief space $\vec{\Omega}$ is also generated by the κ -belief space whose domain is restricted to $\mathcal{D}_{\bar{\kappa}}$. By construction, such belief space is minimal. Formally:

Remark 7 (Minimal Belief Space). If $\vec{\Omega}$ is a κ -belief space in a given category, then $\vec{\Omega}_{\bar{\kappa}} := \langle (\Omega, \mathcal{D}_{\bar{\kappa}}), (B_i|_{\mathcal{D}_{\bar{\kappa}}})_{i \in I}, \Theta \rangle$ is a κ -belief space in the same category with the following properties: (i) the identity map $\text{id}_\Omega : \vec{\Omega} \rightarrow \vec{\Omega}_{\bar{\kappa}}$ is a morphism (an

⁶⁰However, one can informally interpret the symbol \aleph_∞ as satisfying $\lambda < \aleph_\infty$ for any cardinal λ , because a complete algebra \mathcal{X} is closed under the union (and consequently the intersection) of any non-empty sub-collection of \mathcal{X} with cardinality less than λ , for any cardinal λ .

⁶¹In contrast, an event $E \in \mathcal{D} \setminus \mathcal{D}_{\bar{\kappa}}$, if there is any, cannot be captured by the given language \mathcal{L} . Such event E might be captured by a richer language, i.e., $E \in \mathcal{D}_{\bar{\lambda}}$ with $\bar{\lambda} > \bar{\kappa}$.

isomorphism iff $\vec{\Omega}$ is minimal); and (ii) $D_{\vec{\Omega}} = D_{\vec{\Omega}_{\bar{\kappa}}} \circ \text{id}_{\Omega}$. By construction, $\vec{\Omega}_{\bar{\kappa}}$ is minimal: $\mathcal{D}_{\bar{\kappa}} = D_{\vec{\Omega}_{\bar{\kappa}}}^{-1}(\mathcal{D}^*)$.

Using Remark 7, Appendix F characterizes minimality as in Friedenberg and Meier (2011).

7.3 Belief Hierarchies of Depth up to $\bar{\kappa}$

The previous subsection has shown that any κ -belief space $\vec{\Omega}$ contains the belief hierarchies of depth up to $\bar{\kappa}$ that the space intends to represent. For instance, starting from a depth-zero event $\llbracket E \rrbracket_{\vec{\Omega}} \in \mathcal{D}$ with $E \in \mathcal{A}_\kappa(\mathcal{S})$, one can consider a depth-one event $\llbracket \beta_i(E) \rrbracket \in \mathcal{D}$ (“player i believes E ”), a depth-two event $\llbracket \beta_j \beta_i(E) \rrbracket_{\vec{\Omega}} \in \mathcal{D}$ (i.e., “player j believes that i believes E ”), a depth-three event $\llbracket \beta_i \beta_j \beta_i(E) \rrbracket_{\vec{\Omega}} \in \mathcal{D}$ (i.e., “player i believes that j believes that i believes E ”), and so forth.

The particular belief space $\vec{\Omega}$, however, may not distinguish all different depths in the following sense. For instance, if Ω is a finite set, then there exists a finite number n such that any $(n+1)$ -th depth event (e.g., $\llbracket (\beta_{i_1} \beta_{i_2} \cdots \beta_{i_{n+1}})(E) \rrbracket_{\vec{\Omega}} \in \mathcal{D}$) has already appeared as some n -th depth event.

Throughout this subsection, let λ be an infinite regular cardinal with $\lambda \geq \kappa$. This subsection defines the κ -rank of a λ -belief space $\vec{\Omega}$. Roughly, it is the least ordinal α such that any higher-depth event appears already as an α -depth event. The following defines the collection \mathcal{C}_α of α -order events and the κ -rank of the λ -belief space $\vec{\Omega}$.

Definition 17 (κ -Rank). The κ -rank of a λ -belief space $\vec{\Omega}$ of I on (S, \mathcal{S}) is the least ordinal α with $\mathcal{C}_\alpha = \mathcal{C}_{\alpha+1}$, where the sequence $(\mathcal{C}_\alpha)_\alpha$ is defined as follows:

$$\mathcal{C}_\alpha := \begin{cases} \mathcal{A}_\kappa(\{\Theta^{-1}(E) \in \mathcal{D} \mid E \in \mathcal{S}\}) = \Theta^{-1}(\mathcal{A}_\kappa(\mathcal{S})) & \text{if } \alpha = 0 \\ \mathcal{A}_\kappa((\bigcup_{\beta < \alpha} \mathcal{C}_\beta) \cup \bigcup_{i \in I} \{B_i(E) \in \mathcal{D} \mid E \in \bigcup_{\beta < \alpha} \mathcal{C}_\beta\}) & \text{if } \alpha > 0 \end{cases}.$$

The κ -rank of a λ -belief space does not depend on a particular choice of $\lambda \geq \kappa$. The notion of κ -rank generalizes that of rank which Heifetz and Samet (1998a) define for an \aleph_∞ -knowledge space on $(S, \mathcal{P}(S))$. For an \aleph_∞ -belief space, define the *HS-rank* of the belief space as in Definition 17 by substituting $\kappa = \lambda = \aleph_\infty$.

With the notion of HS-rank, Heifetz and Samet (1998a) demonstrate that there is no terminal standard partitional knowledge space from the following two assertions. First, a morphism preserves the HS-ranks. Second, there is a standard partitional knowledge space with arbitrarily high HS-rank. Then, for any candidate terminal standard partitional knowledge space, there exists a standard partitional knowledge space with a higher HS-rank, and thus the candidate space must not be terminal.

The next proposition shows: first, a morphism preserves the κ -ranks; and second, the κ -rank of any λ -belief space (with $\lambda \geq \kappa$) is at most $\bar{\kappa}$. Formally:

Proposition 7 (κ -Rank). 1. If $\varphi : \overrightarrow{\Omega} \rightarrow \overrightarrow{\Omega}'$ is a morphism between λ -belief spaces $\overrightarrow{\Omega}$ and $\overrightarrow{\Omega}'$, then the κ -rank of $\overrightarrow{\Omega}'$ is at least as high as that of $\overrightarrow{\Omega}$.

2. The κ -rank of any λ -belief space $\overrightarrow{\Omega}$ is at most $\bar{\kappa}$.

To prove Part (2), I define $\mathcal{D}_\alpha := \mathcal{A}_\kappa(\{\llbracket e \rrbracket_{\overrightarrow{\Omega}} \in \mathcal{D} \mid e \in \mathcal{L}_\alpha\})$ for each ordinal $\alpha \leq \bar{\kappa}$, where $(\mathcal{L}_\alpha)_{\alpha=0}^{\bar{\kappa}}$ is defined as in Remark 3 so that $\mathcal{L} = \mathcal{L}_{\bar{\kappa}}$. Note that $\mathcal{D}_{\bar{\kappa}}$ coincides with the original definition $\mathcal{D}_{\bar{\kappa}} = \{\llbracket e \rrbracket_{\overrightarrow{\Omega}} \in \mathcal{D} \mid e \in \mathcal{L}\}$. The proof shows that $\mathcal{D}_\alpha = \mathcal{C}_\alpha$ for each ordinal $\alpha \leq \bar{\kappa}$. Then, $\mathcal{D}_{\bar{\kappa}} = \mathcal{C}_{\bar{\kappa}} = \mathcal{C}_{\bar{\kappa}+1}$, i.e., the κ -rank of $\overrightarrow{\Omega}$ is at most $\bar{\kappa}$.

I discuss the role of the infinite regular cardinal κ . By Heifetz and Samet (1998a) and Proposition 7, the κ -rank of the terminal κ -qualitative-belief space $\overrightarrow{\Omega}^*$ is generically $\bar{\kappa}$. By the construction of Heifetz and Samet (1998a), there exists a κ -belief space $\overrightarrow{\Omega}$ with κ -rank $\bar{\kappa}$, but the κ -rank of such space never exceeds $\bar{\kappa}$. Such a particular space (as a κ -belief space) does not contain all possible belief hierarchies of depth up to $\bar{\kappa}$. Thus, for the infinite regular cardinal κ , Heifetz and Samet (1998a)'s non-existence argument does not apply to a given class of κ -belief spaces.⁶²

Two further remarks on Proposition 7 are in order. First, since $\{\llbracket e \rrbracket_{\overrightarrow{\Omega}} \in \mathcal{D} \mid e \in \mathcal{L}\} = \mathcal{D}_{\bar{\kappa}} = \mathcal{C}_{\bar{\kappa}}$, one can examine whether a given κ -belief space $\overrightarrow{\Omega}$ is non-redundant through its primitives alone (i.e., $\mathcal{C}_{\bar{\kappa}}$): recall footnote 35.

Second, in the construction of a terminal κ -qualitative-belief space, each state is a belief hierarchy of depth up to $\bar{\kappa}$. This paper rather shows that the non-existence hinges on the fact that an infinite regular cardinal κ , which determines depth of reasoning, is not specified.⁶³

Within a category of κ -belief spaces (with κ an infinite regular cardinal), the description map preserves interactive beliefs of depth up to $\bar{\kappa}$ in a given κ -belief space to the terminal κ -belief space. At the same time, such preservation concerns only to the extent that belief hierarchies of depth up to $\bar{\kappa}$ are preserved.

7.4 Comparison of Terminal κ -Belief and λ -Belief Spaces

Throughout the subsection, let κ and λ be infinite regular cardinals with $\kappa < \lambda$. Fix I , (S, \mathcal{S}) , and some properties in Definition 2 for the players' beliefs. Denote by $\overrightarrow{\Omega}_\kappa^*$ and $\overrightarrow{\Omega}_\lambda^*$ the terminal κ -belief space and the terminal λ -belief space, respectively. I compare these terminal belief spaces:

⁶²On the contrary, if a terminal partitional \aleph_∞ -knowledge space existed in the category of \aleph_∞ -knowledge spaces, then, for any (infinite regular) cardinal κ , its HS-rank is at least $\bar{\kappa}$, which is a contradiction because the HS-rank of the terminal space is a fixed ordinal.

⁶³As discussed in Section 5, countable additivity of probabilistic beliefs makes it possible to restrict attention to belief hierarchies of depth up to \aleph_0 (i.e., finite-level beliefs) on an \aleph_1 -algebra. In the context of this subsection, Proposition 4 implies that, within the class of probabilistic-belief spaces, the terminal probabilistic-belief space has the \aleph_1 -rank \aleph_0 .

Proposition 8 (Cardinality of a Terminal Space). 1. $D_{\overrightarrow{\Omega_\lambda^*}} : \overrightarrow{\Omega_\lambda^*} \rightarrow \overrightarrow{\Omega_\kappa^*}$ is a surjective morphism so that $|\Omega_\kappa^*| \leq |\Omega_\lambda^*|$.

2. Let $|I| \geq 2$ and $|S| \geq 2$, and suppose \mathcal{S} contains E with $\emptyset \subsetneq E \subsetneq S$. Then, $\max(2^{\aleph_0}, \kappa) \leq |\Omega_\kappa^*|$.

Proposition 8 (1) implies that the terminal λ -belief space $\overrightarrow{\Omega_\lambda^*}$ is at least as rich as $\overrightarrow{\Omega_\kappa^*}$ in cardinality: $|\Omega_\kappa^*| \leq |\Omega_\lambda^*|$. This is because the description map $D_{\overrightarrow{\Omega_\lambda^*}}$ from $\overrightarrow{\Omega_\lambda^*}$ into the terminal κ -space $\overrightarrow{\Omega_\kappa^*}$ is surjective (note that both spaces reside in the same given category of κ -belief spaces). When it comes to belief hierarchies of depth up to $\bar{\kappa}$, however, the space $\overrightarrow{\Omega_\lambda^*}$ is redundant as a κ -belief space (because there would be two states which induce the same belief hierarchy of depth up to $\bar{\kappa}$) and is not minimal (because there is an expression $e \in \mathcal{L}_{\bar{\lambda}} \setminus \mathcal{L}_{\bar{\kappa}}$).

The rest of this subsection discusses two implications of Proposition 8. The first is the sense in which the space $\overrightarrow{\Omega_\lambda^*}$ contains all possible belief hierarchies of depth up to $\bar{\kappa}$. To that end, I introduce a weaker notion of terminality following Friedenberg (2010). Let ν be an infinite regular cardinal with $\nu \leq \kappa$. A κ -belief space $\overrightarrow{\Omega'}$ is ν -terminal (in the category of κ -belief spaces) if, for any state $\omega \in \Omega$ of any κ -belief space $\overrightarrow{\Omega}$, there is a state $\omega' \in \Omega'$ such that ω and ω' induce the same belief hierarchy of depth up to $\bar{\nu}$ for every player: formally, $D_{\overrightarrow{\Omega}}(\omega) \cap (S \sqcup \mathcal{L}_{\bar{\nu}}) = D_{\overrightarrow{\Omega'}}(\omega') \cap (S \sqcup \mathcal{L}_{\bar{\nu}})$. Then, a κ -belief space $\overrightarrow{\Omega}$ is terminal (in the sense of Definition 4) iff it is κ -terminal, non-redundant, and minimal. This is because, given the existence of a terminal κ -belief space $\overrightarrow{\Omega^*}$, a κ -belief space $\overrightarrow{\Omega'}$ is κ -terminal iff $D_{\overrightarrow{\Omega'}}(\omega)$ is surjective. With these in mind, the terminal λ -belief space (in the category of λ -belief spaces) is κ -terminal (as a κ -belief space) but is not minimal and is redundant.

The second implication of Proposition 8 is the non-existence of a terminal \aleph_∞ -belief space. I informally argue that the cardinality of the \aleph_∞ -belief space, if it existed, would be as high as any cardinal, which is impossible.

Define the class $\mathcal{L}_{\aleph_\infty}^I(\mathcal{A}_{\aleph_\infty}(\mathcal{S}))$ of expressions as in Definition 5.⁶⁴ Define the class $\overrightarrow{\Omega_{\aleph_\infty}^*}$ as in Equation (1), and let $\overrightarrow{\Omega_{\aleph_\infty}^*}$ be the terminal space (as a class). Since $\overrightarrow{\Omega_{\aleph_\infty}^*}$ would be non-redundant, $\mathcal{D}_{\aleph_\infty}^*$ is a complete algebra that separates any two states, i.e., the power class $\mathcal{P}(\Omega_{\aleph_\infty}^*)$. Proposition 8 suggests that $\kappa \leq |\Omega_{\aleph_\infty}^*|$ for any (infinite regular) cardinal, meaning that $\Omega_{\aleph_\infty}^*$ is too big to be a set.

Alternatively, consider the terminal κ -belief space $\overrightarrow{\Omega_\kappa^*}$ in the category of κ -belief spaces satisfying at least the Kripke property (recall Definition 2). One can introduce the players' beliefs about any subset of $\overrightarrow{\Omega_\kappa^*}$: formally, see Remark A.1 in Appendix A.5. In the extended space $\langle (\Omega_\kappa^*, \mathcal{P}(\Omega_\kappa^*)), (\overrightarrow{B}_i^*)_{i \in I}, \Theta^* \rangle$, every state $\omega_\kappa^* \in \Omega_\kappa^*$ induces a belief hierarchy of an arbitrary depth, which is uniquely extended from the original

⁶⁴While $\mathcal{L}_\kappa^I(\mathcal{A}_\kappa(\mathcal{S}))$ (Definition 5) is a well-defined set for any infinite (regular) cardinal κ , it can be seen that the class $\mathcal{L}_{\aleph_\infty}^I(\mathcal{A}_{\aleph_\infty}(\mathcal{S}))$ is too big to be a set in the standard set theory.

belief hierarchy (of depth $\bar{\kappa}$). Thus, if the terminal \aleph_∞ -belief space $\overrightarrow{\Omega_{\aleph_\infty}^*}$ existed as a set, then one can construct an injection from Ω_κ^* into $\Omega_{\aleph_\infty}^*$, asserting again that $\kappa \leq |\Omega_{\aleph_\infty}^*|$ for any κ .

8 Concluding Remarks

The main result of this paper (Theorem 1) is the construction of the terminal belief space $\overrightarrow{\Omega^*}$ for varieties of assumptions on properties of beliefs. The space $\overrightarrow{\Omega^*}$ contains all belief hierarchies of depth up to $\bar{\kappa}$ at some state of some belief space. The space Ω^* is belief-complete, i.e., contains all possible beliefs about its states (Proposition 1). The space also exhausts any statement regarding interactive beliefs about the nature states that holds at some state of some belief space (Proposition 2). Each state in Ω^* coherently and completely describes the corresponding nature state and interactive beliefs (Proposition 3 and Corollary 1). Explicitly, Theorem 2 shows that $\overrightarrow{\Omega^*}$ is the largest belief-closed set of coherent descriptions that reflects assumptions on beliefs.

This paper shows that a terminal belief space exists regardless of whether beliefs are qualitative or probabilistic (Corollary 2). Proposition 4 shows, under the framework of this paper, finite-level belief hierarchies uniquely extend to countable levels for countably-additive probabilistic beliefs, as in the literature on type spaces. Appendix D shows that the framework of this paper also applies to richer forms of beliefs such as conditional beliefs. Appendix D also discusses the extensions to the existence of terminal knowledge-unawareness, preference, and expectation spaces.

This paper circumvents the previous non-existence of a terminal knowledge space by restricting attention to all knowledge (belief) hierarchies of depth up to $\bar{\kappa}$. The κ -algebra \mathcal{D}^* can capture all possible interactive beliefs of depth up to $\bar{\kappa}$ (Proposition 7 and Remark 7). While an infinite regular cardinal κ can be taken arbitrarily for given nature states (S, \mathcal{S}) , the terminal κ -belief space $\overrightarrow{\Omega^*}$ contains all belief hierarchies of depth up to $\bar{\kappa}$ and only up to $\bar{\kappa}$ (Proposition 8).

The paper demonstrates that the existence of a terminal belief space hinges on the specification of an infinite regular cardinal κ , which determines depth of reasoning $\bar{\kappa}$, rather than on properties of beliefs themselves (Remark 3). Section 6 shows that the specification of $\bar{\kappa}$ is crucial also for epistemic characterizations of game-theoretic solution concepts such as iterated elimination of strictly dominated (and also B-dominated) actions. Thus, the specification of depth of reasoning also plays an important role in such game-theoretic applications especially when players face a general infinite game or when their beliefs may not satisfy logical properties.

Given a strategic game, this paper shows that there exists a terminal κ -belief space with the following properties: (i) the players can engage in interactive reasoning up to predetermined depth $\bar{\kappa}$ which is sufficient to reason about their interactive beliefs about their play; (ii) rationality, beliefs in rationality, and common belief in rationality are expressible within the terminal space; (iii) common belief in rational-

ity characterizes iterated elimination of strictly dominated (or B-dominated) actions within the terminal space; and (iv) the above (i)-(iii) hold irrespective of assumptions on the players' beliefs. The paper also shows that, through examples, the choice of an infinite (regular) cardinal κ plays a crucial role in epistemic characterizations of solution concepts in games.

A Appendix

A.1 Section 3

Proof of Remark 3. First, it follows from induction that $\mathcal{L}_{\bar{\lambda}} \subseteq \mathcal{L}_{\lambda}^I(\mathcal{A}_{\kappa}(\mathcal{S}))$. Namely, (i) $\mathcal{L}_0 \subseteq \mathcal{L}_{\lambda}^I(\mathcal{A}_{\kappa}(\mathcal{S}))$; and (ii) if $\mathcal{L}_{\beta} \subseteq \mathcal{L}_{\lambda}^I(\mathcal{A}_{\kappa}(\mathcal{S}))$ for all $\beta < \alpha (\leq \bar{\lambda})$ then $\mathcal{L}_{\alpha} \subseteq \mathcal{L}_{\lambda}^I(\mathcal{A}_{\kappa}(\mathcal{S}))$. Conversely, it can be seen that if $e \in \mathcal{L}_{\bar{\lambda}}$ then $e \in \mathcal{L}_{\alpha}$ for some $\alpha < \bar{\lambda}$. I show, by induction on λ -expressions, that $\mathcal{L}_{\lambda}^I(\mathcal{A}_{\kappa}(\mathcal{S})) \subseteq \mathcal{L}_{\bar{\lambda}}$. First, $\mathcal{A}_{\kappa}(\mathcal{S}) \subseteq \mathcal{L}_{\bar{\lambda}}$. Second, if $e \in \mathcal{L}_{\bar{\lambda}}$ then $e \in \mathcal{L}_{\alpha}$ for some $\alpha < \bar{\lambda}$ and thus $(\neg e), (\beta_i(e)) \in \mathcal{L}_{\alpha+1} \subseteq \mathcal{L}_{\bar{\lambda}}$. Third, take $\mathcal{F} \subseteq \mathcal{L}_{\bar{\lambda}}$ with $0 < |\mathcal{F}| < \lambda$. For each $f \in \mathcal{F}$, there is $\alpha_f < \bar{\lambda}$ with $f \in \mathcal{L}_{\alpha_f}$. Since $|\mathcal{F}| < \bar{\lambda}$, the definition of an infinite regular cardinal yields $\gamma := \sup_{f \in \mathcal{F}} \alpha_f < \bar{\lambda}$. Since $\mathcal{F} \subseteq \mathcal{L}_{\gamma} \subseteq \mathcal{L}_{\bar{\lambda}}$, it follows that $\bigwedge \mathcal{F} \in \mathcal{L}_{\bar{\lambda}}$. Hence, $\mathcal{L}_{\lambda}^I(\mathcal{A}_{\kappa}(\mathcal{S})) \subseteq \mathcal{L}_{\bar{\lambda}}$. \square

Observe that the proof of Remark 3 used the fact that a fixed infinite cardinal κ is regular. When $\lambda = \kappa$, the proof states that the set $\mathcal{L}_{\bar{\kappa}}$ contains the conjunction of a set of expressions of cardinality less than κ , where each expression in the set captures interactive reasoning of depth less than $\bar{\kappa}$ (see Remark 2 for the formal definition of depth of an expression). Specifically, let $\kappa = \aleph_0$. Then, $\mathcal{L}_{\bar{\kappa}}$ contains, for instance, the conjunction of “player i believes E ,” “player j believes E ,” “player i believes that player j believes E ,” and “player j believes that player i believes E .” It is a conjunction of finitely many expressions, each of which has finite depth. Likewise, let $\kappa = \aleph_1$. Then, $\mathcal{L}_{\bar{\kappa}}$ contains, for instance, the countable conjunction of “player i_1 believes i_2 believes ... i_n believes E ” for any natural number n . It is a conjunction of countably many expressions, each of which is of at most countable depth.

Proof Sketch of Remark 4. The proof is similar to Heifetz and Samet (1998b, Proposition 4.1) and Meier (2006, Proposition 2). For nature events, since φ is a morphism, $[\cdot]_{\vec{\Omega}} = \Theta^{-1}(\cdot) = \varphi^{-1}(\Theta')^{-1}(\cdot) = \varphi^{-1}([\cdot]_{\vec{\Omega}'})$. Then, use the property that φ^{-1} commutes with set-algebraic operations and belief operators. \square

Proof of Lemma 1. 1. The proof consists of two steps. The first step establishes the following correspondence between syntactic and set-theoretic operations.

- (a) $[(\neg e)] = [e]^c$ for any $e \in \mathcal{L}$.
- (b) $[S] = \Omega^*$ (and $[\emptyset] = \emptyset$).
- (c) $[\bigwedge \mathcal{E}] = \bigcap_{e \in \mathcal{E}} [e]$ (and $[\bigvee \mathcal{E}] = \bigcup_{e \in \mathcal{E}} [e]$) for any $\mathcal{E} \in \mathcal{P}(\mathcal{L})$ with $0 < |\mathcal{E}| < \kappa$.

To prove (1a), fix $e \in \mathcal{L}$. Then, $\omega^* \in [(\neg e)]$ iff $(\neg e) \in_1 \omega^* = D(\omega)$ iff $\omega \in \llbracket \neg e \rrbracket_{\vec{\Omega}} = \neg \llbracket e \rrbracket_{\vec{\Omega}}$ iff $\omega \notin \llbracket e \rrbracket_{\vec{\Omega}}$ iff $e \notin_1 D(\omega) = \omega^*$ iff $\omega^* \notin [e]$ iff $\omega^* \in [e]^c$, where a belief space $\vec{\Omega}$ and a state $\omega \in \Omega$ satisfy $\omega^* = D(\omega)$. Thus, $[(\neg e)] = [e]^c$.

To prove (1b), if $\omega^* \in \Omega^*$ then $\omega^* = D(\omega)$ for some belief space $\vec{\Omega}$ and $\omega \in \Omega$. Since $\omega \in \Omega = \Theta^{-1}(S) = \llbracket S \rrbracket_{\vec{\Omega}}$, I get $S \in_1 D(\omega) = \omega^*$ and thus $\omega^* \in [S]$. This establishes $\Omega^* = [S]$. Also, by (1a), $[\emptyset] = [\neg S] = [S]^c = \emptyset$.

For (1c), $\omega^* \in [\bigwedge \mathcal{E}]$ iff $\bigwedge \mathcal{E} \in_1 \omega^* = D(\omega)$ iff $\omega \in \llbracket \bigwedge \mathcal{E} \rrbracket_{\vec{\Omega}} = \bigcap_{e \in \mathcal{E}} \llbracket e \rrbracket_{\vec{\Omega}}$ iff $\omega^* \in \bigcap_{e \in \mathcal{E}} [e]$, where a belief space $\vec{\Omega}$ and a state $\omega \in \Omega$ satisfy $\omega^* = D(\omega)$.

The second step establishes that \mathcal{D}^* is a κ -algebra on Ω^* . By the first step, \mathcal{D}^* contains $\Omega^* = [S]$ and $\emptyset = [\emptyset]$. Next, if $[e] \in \mathcal{D}^*$, then it follows from the first step and $(\neg e) \in \mathcal{L}$ that $[e]^c = [(\neg e)] \in \mathcal{D}^*$. Next, take $\mathcal{E} \in \mathcal{P}(\mathcal{L})$ with $0 < |\mathcal{E}| < \kappa$. It follows from the first step and $\bigwedge \mathcal{E} \in \mathcal{L}$ that $\bigcap_{e \in \mathcal{E}} [e] = [\bigwedge \mathcal{E}] \in \mathcal{D}^*$.

2. I show $D^{-1}([\cdot]) = \llbracket \cdot \rrbracket_{\vec{\Omega}}$ for any belief space $\vec{\Omega}$. For any $e \in \mathcal{L}$, $\omega \in D^{-1}([e])$ iff $D(\omega) \in [e]$ iff $e \in_1 D(\omega)$ iff $\omega \in \llbracket e \rrbracket_{\vec{\Omega}}$.

□

Proof of Lemma 2. 1. For any $\omega^* \in \Omega^*$, choose a belief space $\vec{\Omega}$ and $\omega \in \Omega$ with $\omega^* = D(\omega)$, and define $\Theta^*(\omega^*) := \Theta(\omega)$, where $\Theta(\omega) \in_0 D(\omega)$. I show $\Theta^*(\omega^*)$ does not depend on a particular choice of $\vec{\Omega}$ and ω (i.e., $\Theta^* : \Omega^* \rightarrow S$ is well-defined). Indeed, if $\omega^* = D_{\vec{\Omega}}(\omega) = D_{\vec{\Omega}'}(\omega')$ for some $\omega \in \Omega$ and $\omega' \in \Omega'$, then $(0, \Theta(\omega)) = (0, \Theta'(\omega'))$.

2. For each $E \in \mathcal{A}_\kappa(\mathcal{S})$, $\omega^* \in (\Theta^*)^{-1}(E)$ iff $\Theta^*(\omega^*) = \Theta(\omega) \in E$ iff $\omega \in \Theta^{-1}(E) = \llbracket E \rrbracket_{\vec{\Omega}}$ iff $E \in_1 D(\omega) = \omega^*$ iff $\omega^* \in [E]$, where $\vec{\Omega}$ and $\omega \in \Omega$ satisfy $\omega^* = D(\omega)$.

□

To establish Lemma 3, I provide Lemma A.1 below. Suppose that a certain property of beliefs is represented by operators $f_{\vec{\Omega}} : \mathcal{D} \rightarrow \mathcal{D}$ in each belief space $\vec{\Omega}$. Operators would be generated by composing belief operators $(B_i)_{i \in I}$ and set-algebraic as well as constant and identity operations. For example, let $f_{\vec{\Omega}}(\cdot) = B_i(\cdot)$ and $g_{\vec{\Omega}}(\cdot) = B_i B_i(\cdot)$. Positive Introspection is characterized by $f_{\vec{\Omega}}(\cdot) \subseteq g_{\vec{\Omega}}(\cdot)$. Truth Axiom is characterized by $f_{\vec{\Omega}}(\cdot) \subseteq \text{id}_{\mathcal{D}}(\cdot)$. Monotonicity is expressed as $f_{\vec{\Omega}}$ being *monotone*: $f_{\vec{\Omega}}(E) \subseteq f_{\vec{\Omega}}(F)$ for all $E, F \in \mathcal{D}$ with $E \subseteq F$. Likewise, λ -Conjunction is expressed as $f_{\vec{\Omega}}$ satisfying *λ -conjunction*: $\bigcap_{E \in \mathcal{E}} f_{\vec{\Omega}}(E) \subseteq f_{\vec{\Omega}}(\bigcap \mathcal{E})$ for all $\mathcal{E} \in \mathcal{P}(\mathcal{D})$ with $0 < |\mathcal{E}| < \lambda$. Abusing the notation, denote by $f_{\vec{\Omega}^*}$ the corresponding operation in $\vec{\Omega}^* = \langle (\Omega^*, \mathcal{D}^*), (B_i^*)_{i \in I}, \Theta^* \rangle$ (note that B_i^* is shown to be well-defined on $(\Omega^*, \mathcal{D}^*)$ irrespective of Lemma A.1 below).

Lemma A.1 (Preservation of Properties of Beliefs). *Suppose that $f_{\vec{\Omega}} : \mathcal{D} \rightarrow \mathcal{D}$ and $g_{\vec{\Omega}} : \mathcal{D} \rightarrow \mathcal{D}$ are defined in each κ -belief space $\vec{\Omega}$. Suppose further that if $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$ is measurable then $\varphi^{-1} f_{\vec{\Omega}'}(\cdot) = f_{\vec{\Omega}} \varphi^{-1}(\cdot)$ and $\varphi^{-1} g_{\vec{\Omega}'}(\cdot) = g_{\vec{\Omega}} \varphi^{-1}(\cdot)$. Then:*

1. (a) If $f_{\vec{\Omega}}(\cdot) \subseteq g_{\vec{\Omega}}(\cdot)$ holds for every belief space $\vec{\Omega}$, then $f_{\vec{\Omega}^*}(\cdot) \subseteq g_{\vec{\Omega}^*}(\cdot)$.
 - (b) If $f_{\vec{\Omega}}(\llbracket e \rrbracket_{\vec{\Omega}}) \not\subseteq g_{\vec{\Omega}}(\llbracket e \rrbracket_{\vec{\Omega}})$ for some belief space $\vec{\Omega}$ and some $e \in \mathcal{L}$, then $f_{\vec{\Omega}^*}([e]) \not\subseteq g_{\vec{\Omega}^*}([e])$.
 - (c) Let $\vec{\Omega}$ and $\vec{\Omega}'$ be a belief space. If there exists a surjective measurable map $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$, then $f_{\vec{\Omega}}(\cdot) \subseteq g_{\vec{\Omega}}(\cdot)$ implies $f_{\vec{\Omega}'}(\cdot) \subseteq g_{\vec{\Omega}'}(\cdot)$.
2. (a) If $f_{\vec{\Omega}}$ is monotone for every belief space $\vec{\Omega}$, then so is $f_{\vec{\Omega}^*}$.
 - (b) Suppose $f_{\vec{\Omega}}(\llbracket e \rrbracket_{\vec{\Omega}}) \not\subseteq f_{\vec{\Omega}}(\llbracket \hat{e} \rrbracket_{\vec{\Omega}})$ for some belief space $\vec{\Omega}$ and some $e, \hat{e} \in \mathcal{L}$. Then, $f_{\vec{\Omega}^*}([e]) \not\subseteq f_{\vec{\Omega}^*}([\hat{e}])$.
 - (c) Let $\vec{\Omega}$ and $\vec{\Omega}'$ be a belief space. Suppose there exists a surjective measurable map $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$. If $f_{\vec{\Omega}}$ is monotone, then so is $f_{\vec{\Omega}'}$.
3. (a) If $f_{\vec{\Omega}}$ satisfies λ -conjunction for every belief space $\vec{\Omega}$, then so does $f_{\vec{\Omega}^*}$.
 - (b) Suppose $\bigcap_{e \in \mathcal{E}} f_{\vec{\Omega}}(\llbracket e \rrbracket_{\vec{\Omega}}) \not\subseteq f_{\vec{\Omega}}(\bigcap_{e \in \mathcal{E}} \llbracket e \rrbracket_{\vec{\Omega}})$ for some belief space $\vec{\Omega}$ and some $\mathcal{E} \in \mathcal{P}(\mathcal{L})$ with $0 < |\mathcal{E}| < \lambda$. Then, $\bigcap_{e \in \mathcal{E}} f_{\vec{\Omega}^*}([e]) \not\subseteq f_{\vec{\Omega}^*}(\bigcap_{e \in \mathcal{E}} [e])$.
 - (c) Let $\vec{\Omega}$ and $\vec{\Omega}'$ be a belief space. Suppose there exists a surjective measurable map $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$. If $f_{\vec{\Omega}}$ satisfies λ -conjunction, then so does $f_{\vec{\Omega}'}$.

Proof of Lemma A.1. 1. The following intermediate result can be obtained. Let $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$ be measurable. Suppose that, for all $E \in \mathcal{D}$, if $\omega \in f_{\vec{\Omega}}(E)$ then $\omega \in g_{\vec{\Omega}}(E)$. Then, for any $E' \in \mathcal{D}'$, $\varphi(\omega) \in f_{\vec{\Omega}'}(E')$ implies $\varphi(\omega) \in g_{\vec{\Omega}'}(E')$. Then, it suffices to show the following:

- (a) If $\omega^* \in f_{\vec{\Omega}^*}([e])$, then there are a belief space $\vec{\Omega}$ and $\omega \in \Omega$ with $\omega^* = D(\omega)$. Now, $\omega^* = D(\omega) \in g_{\vec{\Omega}}([e])$.
- (b) By hypothesis, there is $\omega \in \Omega$ such that $\omega \in f_{\vec{\Omega}}(\llbracket e \rrbracket_{\vec{\Omega}}) = D^{-1}f_{\vec{\Omega}^*}([e])$ and $\omega \notin g_{\vec{\Omega}}(\llbracket e \rrbracket_{\vec{\Omega}}) = D^{-1}g_{\vec{\Omega}^*}([e])$.
- (c) If $\omega' \in f_{\vec{\Omega}'}(E')$, then $\omega' = \varphi(\omega)$ for some $\omega \in \Omega$. Now, $\omega' = \varphi(\omega) \in g_{\vec{\Omega}'}(E')$.

2. The following intermediate result can be obtained. Let $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$ be measurable, and let $f_{\vec{\Omega}}$ be monotone. For any $E', F' \in \mathcal{D}'$ with $E' \subseteq F'$, if $\varphi(\omega) \in f_{\vec{\Omega}'}(E')$, then $\varphi(\omega) \in f_{\vec{\Omega}'}(F')$. Then, it suffices to show the following:
 - (a) Let $[e] \subseteq [\hat{e}]$. If $\omega^* \in f_{\vec{\Omega}^*}([e])$, then there are a belief space $\vec{\Omega}$ and $\omega \in \Omega$ with $\omega^* = D(\omega)$. Now, $\omega^* = D(\omega) \in f_{\vec{\Omega}}([\hat{e}])$.
 - (b) By hypothesis, there is $\omega \in \Omega$ such that $\omega \in f_{\vec{\Omega}}(\llbracket e \rrbracket_{\vec{\Omega}}) = D^{-1}f_{\vec{\Omega}^*}([e])$ and $\omega \notin f_{\vec{\Omega}}(\llbracket \hat{e} \rrbracket_{\vec{\Omega}}) = D^{-1}f_{\vec{\Omega}^*}([\hat{e}])$.
 - (c) Let $E' \subseteq F'$. If $\omega' \in f_{\vec{\Omega}'}(E')$, then $\omega' = \varphi(\omega)$ for some $\omega \in \Omega$. Now, $\omega' = \varphi(\omega) \in f_{\vec{\Omega}'}(F')$.

3. The following intermediate result can be obtained. Let $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$ be measurable. If $f_{\vec{\Omega}}$ satisfies λ -conjunction, then, for any $\mathcal{E}' \in \mathcal{P}(\mathcal{D}')$ with $0 < |\mathcal{E}'| < \lambda$, $\varphi(\omega) \in \bigcap_{E' \in \mathcal{E}'} f_{\vec{\Omega}'}(E')$ implies $\varphi(\omega) \in f_{\vec{\Omega}'}(\bigcap \mathcal{E}')$. Then, it suffices to show the following:

- (a) Fix $\mathcal{E}^* \in \mathcal{P}(\mathcal{D}^*)$ with $0 < |\mathcal{E}^*| < \lambda$. If $\omega^* \in \bigcap_{[e] \in \mathcal{E}^*} f_{\vec{\Omega}^*}([e])$, then there are a belief space $\vec{\Omega}$ and $\omega \in \Omega$ with $\omega^* = D(\omega)$. Now, $\omega^* = D(\omega) \in f_{\vec{\Omega}^*}(\bigcap \mathcal{E}^*)$.
- (b) By hypothesis, there is $\omega \in \Omega$ with $\omega \in \bigcap_{e \in \mathcal{E}} f_{\vec{\Omega}}(\llbracket e \rrbracket_{\vec{\Omega}}) = D^{-1}(\bigcap_{e \in \mathcal{E}} f_{\vec{\Omega}^*}([e]))$ and $\omega \notin f_{\vec{\Omega}}(\bigcap_{e \in \mathcal{E}} \llbracket e \rrbracket_{\vec{\Omega}}) = D^{-1}(f_{\vec{\Omega}^*}(\bigcap_{e \in \mathcal{E}} [e]))$.
- (c) Fix $\mathcal{E}' \in \mathcal{P}(\mathcal{D}')$ with $0 < |\mathcal{E}'| < \lambda$. If $\omega' \in \bigcap_{E' \in \mathcal{E}'} f_{\vec{\Omega}'}(E')$, then there is $\omega \in \Omega$ with $\omega' = \varphi(\omega)$. Now, $\omega' = \varphi(\omega) \in f_{\vec{\Omega}'}(\bigcap \mathcal{E}')$.

□

Two remarks on Lemma A.1 are in order. First, B_i^* violates some property of beliefs if there exists a belief space $\vec{\Omega}$ in the given category such that B_i violates the corresponding property with respect to $\mathcal{D}_{\vec{\Omega}} = D^{-1}(\mathcal{D}^*)$. Second, for belief spaces $\vec{\Omega}$ and $\vec{\Omega}'$, if there is a surjective measurable map $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$ with $B_i \varphi^{-1}(\cdot) = \varphi^{-1} B_i'(\cdot)$, then B_i' inherits the properties of B_i . Now, I prove Lemma 3.

Proof of Lemma 3. 1. To show B_i^* is well-defined, take $e, f \in \mathcal{D}$ with $[e] = [f]$. If $\omega^* \in B_i^*([e]) = [\beta_i(e)]$, then there are a belief space $\vec{\Omega}$ and $\omega \in \Omega$ with $D(\omega) = \omega^* \in [\beta_i(e)]$, i.e., $\beta_i(e) \in_1 D(\omega)$. Thus, $\omega \in \llbracket \beta_i(e) \rrbracket_{\vec{\Omega}} = B_i(\llbracket e \rrbracket_{\vec{\Omega}}) = B_i(D^{-1}([e]))$. Since $[e] = [f]$, I get $\omega \in \llbracket \beta_i(f) \rrbracket_{\vec{\Omega}}$, and thus $\omega^* = D(\omega) \in [\beta_i(f)] = B_i^*([f])$. By changing the role of e and f , $B_i^*([e]) = B_i^*([f])$.

Next, I show B_i^* inherits properties specified in Definition 2. For Monotonicity, apply Lemma A.1 (2a) by taking $f_{\vec{\Omega}} = B_i$. For λ -Conjunction, apply Lemma A.1 (3a) by taking $f_{\vec{\Omega}} = B_i$. Next, apply Lemma A.1 (1a) to the following. For Necessitation, take $(f_{\vec{\Omega}}, g_{\vec{\Omega}}) = (\Omega, B_i(\Omega))$. For Consistency, take $(f_{\vec{\Omega}}, g_{\vec{\Omega}}) = (B_i(\cdot) \cap (\neg B_i)(\cdot), \emptyset)$. For Truth Axiom, take $(f_{\vec{\Omega}}, g_{\vec{\Omega}}) = (B_i, \text{id}_{\mathcal{D}})$. For Positive Introspection, take $(f_{\vec{\Omega}}, g_{\vec{\Omega}}) = (B_i, B_i B_i)$. For Negative Introspection, take $(f_{\vec{\Omega}}, g_{\vec{\Omega}}) = (B_i, B_i(\neg B_i))$.

Finally, consider the Kripke property. Suppose $b_{B_i^*}^*(\omega^*) \subseteq [e]$. There are a belief space $\vec{\Omega}$ and $\omega \in \Omega$ such that $\omega^* = D(\omega)$. Operating D^{-1} on both sides of the set inclusion, one can obtain $b_{B_i}(\omega) = \bigcap_{F \in \mathcal{D}: \omega \in B_i(F)} F \subseteq D^{-1} \bigcap_{[f] \in \mathcal{D}^*: \omega^* \in B_i^*([f])} [f] \subseteq D^{-1}[e]$. Since B_i satisfies the Kripke property, $\omega \in B_i D^{-1}[e] = D^{-1} B_i^*([e])$. Then, $\omega^* = D(\omega) \in B_i^*([e])$, as desired.

2. For any $e \in \mathcal{L}$, $B_i(D^{-1}([e])) = B_i(\llbracket e \rrbracket_{\vec{\Omega}}) = \llbracket \beta_i(e) \rrbracket_{\vec{\Omega}} = D^{-1}([\beta_i(e)]) = D^{-1}(B_i^*([e]))$.

□

Proof of Lemma 4. First, if $s \in_0 \omega^*$ and $s' \in_0 D(\omega^*)$, then there are a belief space $\vec{\Omega}$ and $\omega \in \Omega$ such that $s = \Theta(\omega) = \Theta^*(D(\omega)) = \Theta^*(\omega^*) = s'$. Note that the argument does not depend on a particular choice of belief spaces. Second, similarly to Heifetz and Samet (1998b, Lemma 4.6) and Meier (2006, Lemma 6), below I show by induction that $\llbracket \cdot \rrbracket_{\vec{\Omega}^*} = [\cdot]$. Then, $\omega^* = \{s\} \sqcup \{e \in \mathcal{L} \mid e \in_1 \omega^*\} = \{s\} \sqcup \{e \in \mathcal{L} \mid \omega^* \in [e]\} = \{s\} \sqcup \{e \in \mathcal{L} \mid \omega^* \in \llbracket e \rrbracket_{\vec{\Omega}^*}\} = D(\omega^*)$ for any $\omega^* \in \Omega^*$.

To establish $\llbracket \cdot \rrbracket_{\vec{\Omega}^*} = [\cdot]$, start from $E \in \mathcal{A}_\kappa(\mathcal{S})$. In fact, $\omega^* \in \llbracket E \rrbracket_{\vec{\Omega}^*} = (\Theta^*)^{-1}(E)$ iff $\Theta^*(\omega^*) = \Theta^*(D(\omega)) = \Theta(\omega) \in E$ iff $\omega \in \Theta^{-1}(E) = \llbracket E \rrbracket_{\vec{\Omega}}$ iff $E \in_1 D(\omega)$ iff $\omega^* = D(\omega) \in [E]$, where $\vec{\Omega}$ and $\omega \in \Omega$ satisfy $\omega^* = D(\omega)$.

Next, let $\mathcal{E} \in \mathcal{P}(\mathcal{L})$ with $0 < |\mathcal{E}| < \kappa$. Assume the induction hypothesis that $\llbracket e \rrbracket_{\vec{\Omega}^*} = [e]$ for all $e \in \mathcal{E}$. Then, $\llbracket \bigwedge \mathcal{E} \rrbracket_{\vec{\Omega}^*} = \bigcap_{e \in \mathcal{E}} \llbracket e \rrbracket_{\vec{\Omega}^*} = \bigcap_{e \in \mathcal{E}} [e] = [\bigwedge \mathcal{E}]$.

Next, assume the induction hypothesis that $\llbracket e \rrbracket_{\vec{\Omega}^*} = [e]$. By definition, $[\beta_i(e)] = B_i^*([e]) = B_i^*(\llbracket e \rrbracket_{\vec{\Omega}^*}) = \llbracket \beta_i(e) \rrbracket_{\vec{\Omega}^*}$. Also, $[(\neg e)] = \neg [e] = \neg \llbracket e \rrbracket_{\vec{\Omega}^*} = \llbracket \neg e \rrbracket_{\vec{\Omega}^*}$. \square

Proof of Theorem 1. I have already shown that $\vec{\Omega}^* = \langle (\Omega^*, \mathcal{D}^*), (B_i^*)_{i \in I}, \Theta^* \rangle$ is a belief space of I on (S, \mathcal{S}) of the given category such that, for any belief space $\vec{\Omega}$, the description map $D_{\vec{\Omega}}$ is a morphism. Remark 6 and Lemma 4 imply that $D_{\vec{\Omega}}$ is a unique morphism (see footnote 39). \square

A.2 Section 4

Proof of Proposition 1. I show Part (2) first and then Part (1). Part (3) immediately follows from these two parts.

Part (2). For the “only if” part, let $\vec{\Omega}$ be non-redundant. Assume $\chi_{\vec{\Omega}}(\omega) = \chi_{\vec{\Omega}}(\omega')$. Since $D_{\vec{\Omega}}$ is injective, I show, by induction, that $D_{\vec{\Omega}}(\omega) = D_{\vec{\Omega}}(\omega')$, which leads to $\omega = \omega'$. First, $\Theta(\omega) = \Theta(\omega')$ follows from $\chi_{\vec{\Omega}}(\omega) = \chi_{\vec{\Omega}}(\omega')$. Second, for any $E \in \mathcal{A}_\kappa(\mathcal{S})$, $E \in_1 D_{\vec{\Omega}}(\omega)$ iff $\Theta(\omega) = \Theta(\omega') \in E$ iff $E \in_1 D_{\vec{\Omega}}(\omega')$. Third, assume the induction hypothesis $e \in_1 D_{\vec{\Omega}}(\omega)$ iff $e \in_1 D_{\vec{\Omega}}(\omega')$. Then, $(\neg e) \in_1 D_{\vec{\Omega}}(\omega)$ iff $e \notin_1 D_{\vec{\Omega}}(\omega)$ iff $e \notin_1 D_{\vec{\Omega}}(\omega')$ iff $(\neg e) \in_1 D_{\vec{\Omega}}(\omega')$. Fourth, assume the induction hypothesis $e \in_1 D_{\vec{\Omega}}(\omega)$ iff $e \in_1 D_{\vec{\Omega}}(\omega')$ for all $e \in \mathcal{E}$, where $\mathcal{E} \in \mathcal{P}(\mathcal{L})$ and $0 < |\mathcal{E}| < \kappa$. Then, $\bigwedge \mathcal{E} \in_1 D_{\vec{\Omega}}(\omega)$ iff $e \in_1 D_{\vec{\Omega}}(\omega)$ for all $e \in \mathcal{E}$ iff $e \in_1 D_{\vec{\Omega}}(\omega')$ for all $e \in \mathcal{E}$ iff $\bigwedge \mathcal{E} \in_1 D_{\vec{\Omega}}(\omega')$. Fifth, $\beta_i(e) \in_1 D_{\vec{\Omega}}(\omega)$ iff $\omega \in B_i(\llbracket e \rrbracket_{\vec{\Omega}})$ iff $\omega' \in B_i(\llbracket e \rrbracket_{\vec{\Omega}})$ iff $\beta_i(e) \in_1 D_{\vec{\Omega}}(\omega')$, where the second equivalence follows from $\chi_{\vec{\Omega}}(\omega) = \chi_{\vec{\Omega}}(\omega')$. Hence, $D_{\vec{\Omega}}(\omega) = D_{\vec{\Omega}}(\omega')$.

For the “if” part, let $\chi_{\vec{\Omega}}$ be injective. Assume $D_{\vec{\Omega}}(\omega) = D_{\vec{\Omega}}(\omega')$. It suffices to show $\chi_{\vec{\Omega}}(\omega) = \chi_{\vec{\Omega}}(\omega')$. First, $D_{\vec{\Omega}}(\omega) = D_{\vec{\Omega}}(\omega')$ yields $\Theta(\omega) = \Theta(\omega')$. Second, if $\omega \in B_i(\llbracket e \rrbracket_{\vec{\Omega}}) = D_{\vec{\Omega}}^{-1}([\beta_i(e)])$, then $D(\omega') = D(\omega) \in [\beta_i(e)]$ and thus $\omega' \in B_i(\llbracket e \rrbracket_{\vec{\Omega}})$. Likewise, if $\omega' \in B_i(\llbracket e \rrbracket_{\vec{\Omega}})$ then $\omega \in B_i(\llbracket e \rrbracket_{\vec{\Omega}})$. Hence, $\chi_{\vec{\Omega}}(\omega) = \chi_{\vec{\Omega}}(\omega')$.

Part (1). For the “only if” part, a terminal space $\vec{\Omega}$ is non-redundant and minimal.

To show that $\chi_{\vec{\Omega}}$ is surjective, observe that $\chi_{\vec{\Omega}} = \chi_{\vec{\Omega}^*} \circ D_{\vec{\Omega}}$ follows because

$$\begin{aligned}\chi_{\vec{\Omega}}(\omega) &= (\Theta(\omega), (\{e \in \mathcal{L} \mid \omega \in B_i(\llbracket e \rrbracket_{\vec{\Omega}})\})_{i \in I}) \\ &= (\Theta(D_{\vec{\Omega}}(\omega)), (\{e \in \mathcal{L} \mid D_{\vec{\Omega}}(\omega) \in B_i([e])\})_{i \in I}) = \chi_{\vec{\Omega}^*}(D_{\vec{\Omega}}(\omega)).\end{aligned}$$

Since $\vec{\Omega}$ is terminal, $D_{\vec{\Omega}}$ is bijective. By Expression (3) in the discussion of this proposition, $\chi_{\vec{\Omega}^*}$ is surjective. Thus, $\chi_{\vec{\Omega}}$ is surjective (note that, by Part (2), it is also injective because $\vec{\Omega}^*$ is non-redundant).

For the “if” part, I show that $D_{\vec{\Omega}}$ is an isomorphism. Since $\vec{\Omega}$ is minimal (i.e., $\mathcal{D} = D_{\vec{\Omega}}^{-1}(\mathcal{D}^*)$), if $D_{\vec{\Omega}}$ is bijective then $D_{\vec{\Omega}}^{-1} : \Omega^* \rightarrow \Omega$ is measurable. By Part (2), $D_{\vec{\Omega}}$ is injective. Thus, it suffices to show that $D_{\vec{\Omega}}$ is surjective. Take any $\omega^* \in \Omega^*$. Since $\chi_{\vec{\Omega}}$ is surjective, for $\chi_{\vec{\Omega}^*}(\omega^*) \in \Omega^{**}$, there is $\omega \in \Omega$ such that $\chi_{\vec{\Omega}^*}(\omega^*) = \chi_{\vec{\Omega}}(\omega)$: $\Theta^*(\omega^*) = \Theta(\omega)$ and $\{e \in \mathcal{L} \mid \omega^* \in B_i^*([e])\} = \{e \in \mathcal{L} \mid \omega \in B_i(\llbracket e \rrbracket_{\vec{\Omega}})\}$ for each $i \in I$. Then, it suffices to show that $\omega^* = D_{\vec{\Omega}}(\omega)$.

To that end, I show by induction that $e \in_1 \omega^*$ iff $e \in_1 D_{\vec{\Omega}}(\omega)$. For any $E \in \mathcal{A}_\kappa(\mathcal{S})$, $E \in_1 \omega^*$ iff $\Theta^*(\omega^*) \in E$ iff $\Theta(\omega) \in E$ iff $E \in_1 D_{\vec{\Omega}}(\omega)$. Next, assume $e \in_1 \omega^*$ iff $e \in_1 D_{\vec{\Omega}}(\omega)$. Then, $(\neg e) \in_1 \omega^*$ iff $e \notin_1 \omega^*$ iff $e \notin_1 D_{\vec{\Omega}}(\omega)$ iff $(\neg e) \in_1 D_{\vec{\Omega}}(\omega)$. Next, assume $e \in_1 \omega^*$ iff $e \in_1 D_{\vec{\Omega}}(\omega)$ for all $e \in \mathcal{E}$, where $\mathcal{E} \in \mathcal{P}(\mathcal{L})$ and $0 < |\mathcal{E}| < \kappa$. Then, $\bigwedge \mathcal{E} \in_1 \omega^*$ iff $e \in_1 \omega^*$ for all $e \in \mathcal{E}$ iff $e \in_1 D_{\vec{\Omega}}(\omega)$ for all $e \in \mathcal{E}$ iff $\bigwedge \mathcal{E} \in_1 D_{\vec{\Omega}}(\omega)$. Next, $\beta_i(e) \in_1 \omega^*$ iff $\omega^* \in B_i^*([e])$ iff $\omega \in B_i(\llbracket e \rrbracket_{\vec{\Omega}})$ iff $\beta_i(e) \in_1 D_{\vec{\Omega}}(\omega)$. \square

Proof of Proposition 2. For Part (1), it suffices to show that if Φ is satisfiable then it is satisfiable in $\vec{\Omega}^*$. If there are a belief space $\vec{\Omega}$ and $\omega \in \Omega$ with $\omega \in \llbracket f \rrbracket_{\vec{\Omega}} = D^{-1}([f])$ for all $f \in \Phi$, then $D(\omega) \in [f] = \llbracket f \rrbracket_{\vec{\Omega}^*}$ for all $f \in \Phi$.

For the first assertion of Part (2), it is enough to show that $\Phi \models_{\vec{\Omega}^*} e$ implies $\Phi \models e$. Let $\vec{\Omega}$ be a belief space. If $\omega \in \llbracket f \rrbracket_{\vec{\Omega}} = D^{-1}([f])$ for all $f \in \Phi$, then $D(\omega) \in [f] = \llbracket f \rrbracket_{\vec{\Omega}^*}$ for all $f \in \Phi$. By assumption, $D(\omega) \in \llbracket e \rrbracket_{\vec{\Omega}^*} = [e]$, i.e., $\omega \in D^{-1}([e]) = \llbracket e \rrbracket_{\vec{\Omega}}$. Thus, $\Phi \models e$. The second assertion can be seen as a special case of the first. Or, for any belief space $\vec{\Omega}$, $\llbracket e \rrbracket_{\vec{\Omega}} = D^{-1}([e]) = D^{-1}(\llbracket e \rrbracket_{\vec{\Omega}^*}) = D^{-1}(\Omega^*) = \Omega$.

For Part (3), let Ω consist of $\{s\} \sqcup \Phi \in \mathcal{P}(S \sqcup \mathcal{L})$ such that Φ is maximally satisfiable and that, for any $E \in \mathcal{A}_\kappa(\mathcal{S})$, $s \in E$ iff $E \in \Phi$. I show $\Omega^* = \Omega$ in two steps. The first step establishes $\Omega^* \subseteq \Omega$. The second step proves $\Omega \subseteq \Omega^*$ by showing that there exists a belief space defined on Ω such that its description map is an inclusion map.

Step 1. Take $\omega^* \in \Omega^*$. Denote $\omega^* = \{s\} \sqcup \Phi$, i.e., $s \in_0 \omega^*$ and $\Phi = \{e \in \mathcal{L} \mid e \in_1 \omega^*\}$. Since $\omega^* \in [e] = \llbracket e \rrbracket_{\vec{\Omega}^*}$ for all $e \in \Phi$, the set Φ is satisfiable. To show it is maximally satisfiable, take a satisfiable set of expressions Ψ with $\Phi \subseteq \Psi$. If there is $e \in \Psi \setminus \Phi$, then $(\neg e) \in \Phi \subseteq \Psi$. Then, Ψ is not satisfiable, a contradiction. Thus, $\Phi = \Psi$, i.e., Φ is maximally satisfiable. Next, $\Theta^*(\omega^*) = s$. For any $E \in \mathcal{A}_\kappa(\mathcal{S})$, $s \in E$ iff $\Theta^*(\omega^*) \in E$ iff $E \in_1 \omega^*$, i.e., $E \in \Phi$.

Step 2. I construct, in four substeps, a belief space $\vec{\Omega} := \langle (\Omega, \mathcal{D}), (B_i)_{i \in I}, \Theta \rangle$ defined on Ω such that the description map $D_{\vec{\Omega}}$ is an inclusion map. To that end, observe that, for any maximally satisfiable set Φ , there are a belief space $\vec{\Omega}'$ and $\omega' \in \Omega'$ with $\Phi = \{e \in \mathcal{L} \mid \omega' \in \llbracket e \rrbracket_{\vec{\Omega}'}\}$ (i.e., $e \in_1 \Phi$ iff $\omega' \in \llbracket e \rrbracket_{\vec{\Omega}'}$).

Step 2.1. Slightly abusing the notation, define $[e]_{\vec{\Omega}} := \{\omega \in \Omega \mid e \in_1 \omega\}$ for each $e \in \mathcal{L}$ (Step 2.4 establishes $[\cdot]_{\vec{\Omega}} = [\cdot]$). I show that $\mathcal{D} := \{[e]_{\vec{\Omega}} \in \mathcal{P}(\Omega) \mid e \in \mathcal{L}\}$ is a κ -algebra. First, $\Omega = [S]_{\vec{\Omega}} \in \mathcal{D}$ and $\emptyset = [\emptyset]_{\vec{\Omega}} \in \mathcal{D}$. Second, I show $[\neg e]_{\vec{\Omega}} = ([e]_{\vec{\Omega}})^c$. In fact, $\omega \in [\neg e]_{\vec{\Omega}}$ iff $(\neg e) \in_1 \omega$ iff there are a belief space $\vec{\Omega}'$ and $\omega' \in \Omega'$ with $\omega' \in \llbracket \neg e \rrbracket_{\vec{\Omega}'} = ([e]_{\vec{\Omega}'})^c$ iff $e \notin_1 \omega$ iff $\omega \notin [e]_{\vec{\Omega}}$ iff $\omega \in ([e]_{\vec{\Omega}})^c$. Now, $([e]_{\vec{\Omega}})^c = [\neg e]_{\vec{\Omega}} \in \mathcal{D}$. Third, I show $\bigcap_{e \in \mathcal{E}} [e]_{\vec{\Omega}} = [\bigwedge \mathcal{E}]_{\vec{\Omega}}$, where $\mathcal{E} \in \mathcal{P}(\mathcal{L})$ and $0 < |\mathcal{E}| < \kappa$. By definition, $\omega \in \bigcap_{e \in \mathcal{E}} [e]_{\vec{\Omega}}$ iff $\omega \in [e]_{\vec{\Omega}}$, i.e., $e \in_1 \omega$, for all $e \in \mathcal{E}$ iff there are a belief space $\vec{\Omega}'$ and $\omega' \in \Omega'$ with $\omega' \in \bigcap_{e \in \mathcal{E}} \llbracket e \rrbracket_{\vec{\Omega}'} = \llbracket \bigwedge \mathcal{E} \rrbracket_{\vec{\Omega}'}$ iff $\bigwedge \mathcal{E} \in_1 \omega$ iff $\omega \in [\bigwedge \mathcal{E}]_{\vec{\Omega}}$. Now, $\bigcap_{e \in \mathcal{E}} [e]_{\vec{\Omega}} = [\bigwedge \mathcal{E}]_{\vec{\Omega}} \in \mathcal{D}$.

Step 2.2. For each $[e]_{\vec{\Omega}} \in \mathcal{D}$, define $B_i([e]_{\vec{\Omega}}) := [\beta_i(e)]_{\vec{\Omega}}$. First, I show that B_i is well-defined: $[e]_{\vec{\Omega}} = [f]_{\vec{\Omega}}$ implies $[\beta_i(e)]_{\vec{\Omega}} = [\beta_i(f)]_{\vec{\Omega}}$. Assume $[e]_{\vec{\Omega}} = [f]_{\vec{\Omega}}$. Let $\omega \in [\beta_i(e)]_{\vec{\Omega}}$. There are a belief space $\vec{\Omega}'$ and $\omega' \in \Omega'$ such that $\omega' \in \llbracket \beta_i(e) \rrbracket_{\vec{\Omega}'} = B'_i(\llbracket e \rrbracket_{\vec{\Omega}'}) = B'_i(\llbracket f \rrbracket_{\vec{\Omega}'}) = \llbracket \beta_i(f) \rrbracket_{\vec{\Omega}'}$, where the second equality follows from $\llbracket e \rrbracket_{\vec{\Omega}'} = \llbracket f \rrbracket_{\vec{\Omega}'}$, which can be shown as follows. If $\tilde{\omega} \in \llbracket e \rrbracket_{\vec{\Omega}'}$, then $e \in \Phi := \{\hat{e} \in \mathcal{L} \mid \tilde{\omega} \in \llbracket \hat{e} \rrbracket_{\vec{\Omega}'}\}$. Since Φ is maximally satisfiable and $[e]_{\vec{\Omega}} = [f]_{\vec{\Omega}}$, it follows that $f \in \Phi$, i.e., $\tilde{\omega} \in \llbracket f \rrbracket_{\vec{\Omega}'}$. Similarly, $\tilde{\omega} \in \llbracket f \rrbracket_{\vec{\Omega}'} \text{ implies } \tilde{\omega} \in \llbracket e \rrbracket_{\vec{\Omega}'}$. Now, $\omega \in [\beta_i(f)]_{\vec{\Omega}}$, establishing $[\beta_i(e)]_{\vec{\Omega}} \subseteq [\beta_i(f)]_{\vec{\Omega}}$. Similarly, $[\beta_i(f)]_{\vec{\Omega}} \subseteq [\beta_i(e)]_{\vec{\Omega}}$.

Second, I show that each B_i inherits the properties of beliefs in Definition 2 imposed in a given category of belief spaces. For Monotonicity, let $\omega \in B_i([e]_{\vec{\Omega}})$ and $[e]_{\vec{\Omega}} \subseteq [f]_{\vec{\Omega}}$. There are a belief space $\vec{\Omega}'$ and $\omega' \in \Omega'$ with $\omega' \in \llbracket \beta_i(e) \rrbracket_{\vec{\Omega}'} = B'_i(\llbracket e \rrbracket_{\vec{\Omega}'})$. Since $[e]_{\vec{\Omega}'} \subseteq [f]_{\vec{\Omega}'}$ implies $\llbracket e \rrbracket_{\vec{\Omega}'} \subseteq \llbracket f \rrbracket_{\vec{\Omega}'}$, it follows from Monotonicity of B'_i that $\omega' \in B'_i(\llbracket f \rrbracket_{\vec{\Omega}'}) = \llbracket \beta_i(f) \rrbracket_{\vec{\Omega}'}$. Then, $\omega \in [\beta_i(f)]_{\vec{\Omega}} = B_i([f]_{\vec{\Omega}})$.

For Necessitation, if $\omega \in \Omega$ then there are a belief space $\vec{\Omega}'$ and $\omega' \in \Omega'$ such that $\omega' \in \llbracket e \rrbracket_{\vec{\Omega}'}$ for all $e \in_1 \omega$. Then, $\omega' \in \llbracket S \rrbracket_{\vec{\Omega}'} = B'_i(\llbracket S \rrbracket_{\vec{\Omega}'}) = \llbracket \beta_i(S) \rrbracket_{\vec{\Omega}'}$, where the first equality follows from Necessitation of B'_i . Hence, $(\beta_i(S)) \in_1 \omega$, i.e., $\omega \in [\beta_i(S)]_{\vec{\Omega}} = B_i([S]_{\vec{\Omega}})$. Thus, $\Omega = B_i([S]_{\vec{\Omega}}) = B_i(\Omega)$.

For λ -Conjunction, if $\omega \in \bigcap_{e \in \mathcal{E}} B_i([e]_{\vec{\Omega}})$ where $\mathcal{E} \in \mathcal{P}(\mathcal{L})$ and $0 < |\mathcal{E}| < \lambda$, then there are a belief space $\vec{\Omega}'$ and $\omega' \in \Omega'$ such that $\omega' \in B'_i(\llbracket e \rrbracket_{\vec{\Omega}'})$ for all $e \in \mathcal{E}$. Since B'_i satisfies λ -Conjunction, $\omega' \in \bigcap_{e \in \mathcal{E}} B'_i(\llbracket e \rrbracket_{\vec{\Omega}'}) \subseteq B'_i(\bigcap_{e \in \mathcal{E}} \llbracket e \rrbracket_{\vec{\Omega}'}) = B'_i(\llbracket \bigwedge \mathcal{E} \rrbracket_{\vec{\Omega}'})$. Then, $\omega \in B_i(\llbracket \bigwedge \mathcal{E} \rrbracket_{\vec{\Omega}'})$.

For the Kripke property, let $\bigcap_{[e]_{\vec{\Omega}} \in \mathcal{D}: \omega \in B_i([e]_{\vec{\Omega}})} [e]_{\vec{\Omega}} \subseteq [f]_{\vec{\Omega}}$. There are a belief space $\vec{\Omega}'$ and $\omega' \in \Omega'$ with $\{e \in \mathcal{L} \mid \omega \in B_i([e]_{\vec{\Omega}})\} = \{e \in \mathcal{L} \mid \omega' \in B'_i(\llbracket e \rrbracket_{\vec{\Omega}'})\}$. Then, since $\bigcap_{[e]_{\vec{\Omega}} \in \mathcal{D}: \omega' \in B'_i(\llbracket e \rrbracket_{\vec{\Omega}'})} [e]_{\vec{\Omega}} \subseteq [f]_{\vec{\Omega}'}$ and since B'_i satisfies the Kripke property,

$\omega' \in B'_i(\llbracket f \rrbracket_{\vec{\Omega}'})$. Then, $\omega \in B_i(\llbracket f \rrbracket_{\vec{\Omega}})$.

For Consistency, if $\omega \in B_i([e]_{\vec{\Omega}}) = [\beta_i(e)]_{\vec{\Omega}}$, then there are a belief space $\vec{\Omega}'$ and $\omega' \in \Omega'$ with $\omega' \in B'_i(\llbracket e \rrbracket_{\vec{\Omega}'}) \subseteq (\neg B'_i)(\llbracket \neg e \rrbracket_{\vec{\Omega}'})$, where the set inclusion follows from Consistency of B'_i . Then, $\omega' \in \llbracket (\neg \beta_i)(\neg e) \rrbracket_{\vec{\Omega}'}$, and thus $\omega \in \llbracket (\neg \beta_i)(\neg e) \rrbracket_{\vec{\Omega}} = (\neg B_i)(\neg [e]_{\vec{\Omega}})$. For Truth Axiom, if $\omega \in B_i([e]_{\vec{\Omega}})$ then there are a belief space $\vec{\Omega}'$ and $\omega' \in \Omega'$ such that $\omega' \in \llbracket \beta_i(e) \rrbracket_{\vec{\Omega}'} = B'_i(\llbracket e \rrbracket_{\vec{\Omega}'}) \subseteq \llbracket e \rrbracket_{\vec{\Omega}'}$. Thus, $\omega \in [e]_{\vec{\Omega}}$.

For Positive Introspection, if $\omega \in B_i([e]_{\vec{\Omega}})$, then there are a belief space $\vec{\Omega}'$ and $\omega' \in \Omega'$ with $\omega' \in \llbracket \beta_i(e) \rrbracket_{\vec{\Omega}'} = B'_i(\llbracket e \rrbracket_{\vec{\Omega}'}) \subseteq B'_i B'_i(\llbracket e \rrbracket_{\vec{\Omega}'}) = \llbracket \beta_i \beta_i(e) \rrbracket_{\vec{\Omega}'}$, where the set inclusion follows from Positive Introspection of B'_i . Then, $\omega \in \llbracket \beta_i \beta_i(e) \rrbracket_{\vec{\Omega}} = B_i B_i([e]_{\vec{\Omega}})$. The proof for Negative Introspection is similar.

Step 2.3. For each $\omega \in \Omega$, let $\Theta(\omega)$ be the unique $s \in S$ with $s \in_0 \omega$. Since $\Theta^{-1}(E) = [E]_{\vec{\Omega}} \in \mathcal{D}$ for all $E \in \mathcal{A}_\kappa(\mathcal{S})$, $\Theta : (\Omega, \mathcal{D}) \rightarrow (S, \mathcal{A}_\kappa(\mathcal{S}))$ is measurable.

Step 2.4. So far, I have constructed a belief space $\vec{\Omega}$. To show that the description map $D_{\vec{\Omega}}$ is an inclusion map, I start with showing that $[\cdot]_{\vec{\Omega}} = \llbracket \cdot \rrbracket_{\vec{\Omega}}$ (once $\vec{\Omega} = \vec{\Omega}^*$ is established, $[\cdot]_{\vec{\Omega}} = \llbracket \cdot \rrbracket_{\vec{\Omega}^*} = [\cdot]$). For each $E \in \mathcal{A}_\kappa(\mathcal{S})$, $[E]_{\vec{\Omega}} = \Theta^{-1}(E) = \llbracket E \rrbracket_{\vec{\Omega}}$. If $[e]_{\vec{\Omega}} = \llbracket e \rrbracket_{\vec{\Omega}}$, then $\llbracket \neg e \rrbracket_{\vec{\Omega}} = \neg [e]_{\vec{\Omega}} = \neg \llbracket e \rrbracket_{\vec{\Omega}} = \llbracket \neg e \rrbracket_{\vec{\Omega}}$ and $\llbracket \beta_i(e) \rrbracket_{\vec{\Omega}} = B_i([e]_{\vec{\Omega}}) = B_i(\llbracket e \rrbracket_{\vec{\Omega}}) = \llbracket \beta_i(e) \rrbracket_{\vec{\Omega}}$. Assume $[e]_{\vec{\Omega}} = \llbracket e \rrbracket_{\vec{\Omega}}$ for each $e \in \mathcal{E}$, where $\mathcal{E} \in \mathcal{P}(\mathcal{L})$ and $0 < |\mathcal{E}| < \kappa$. Then, $\llbracket \bigwedge \mathcal{E} \rrbracket_{\vec{\Omega}} = \bigcap_{e \in \mathcal{E}} [e]_{\vec{\Omega}} = \bigcap_{e \in \mathcal{E}} \llbracket e \rrbracket_{\vec{\Omega}} = \llbracket \bigwedge \mathcal{E} \rrbracket_{\vec{\Omega}}$.

Finally, I establish $D_{\vec{\Omega}}(\omega) = \omega$ for all $\omega \in \Omega$. First, for any ω , $\Theta(\omega) \in_0 D_{\vec{\Omega}}(\omega)$ and $s \in_0 \omega$ satisfy $s = \Theta(\omega)$. Second, $e \in_1 D_{\vec{\Omega}}(\omega)$ iff $D_{\vec{\Omega}}(\omega) \in [e]$ iff $\omega \in D_{\vec{\Omega}}^{-1}([e]) = \llbracket e \rrbracket_{\vec{\Omega}} = [e]_{\vec{\Omega}}$ iff $e \in_1 \omega$. \square

Proof of Proposition 3. For Part (1), $e \notin_1 \omega^*$ iff $\omega^* \notin [e]$ iff $\omega^* \in \neg [e] = \llbracket \neg e \rrbracket$ iff $\neg e \in_1 \omega^*$. For Part (2), let $e \in_1 \omega^*$ and $(e \rightarrow f) \in_1 \omega^*$. Then, $\omega^* \in [e]$ and $\omega^* \in [e \rightarrow f] = \llbracket (\neg e) \vee f \rrbracket = \neg [e] \cup [f]$. Thus, $\omega^* \in [f]$, i.e., $f \in_1 \omega^*$. For Part (3), $\bigwedge \mathcal{E} \in_1 \omega^*$ iff $\omega^* \in \llbracket \bigwedge \mathcal{E} \rrbracket = \bigcap_{e \in \mathcal{E}} [e]$ iff, for all $e \in \mathcal{E}$, $\omega^* \in [e]$, i.e., $e \in_1 \omega^*$. \square

Proof of Corollary 1. Part (1) follows from Proposition 3. Part (3) follows from Proposition 3 and the fact that $(\neg \beta_i)(e) \wedge (\neg \beta_i)(\neg e) \in_1 \omega^*$ iff $\omega^* \in (\neg B_i^*)([e]) \cap (\neg B_i^*)(\neg [e])$. Thus, I prove Part (2).

First, if $\beta_i(e) \notin_1 \omega^*$ and $\beta_i(\neg e) \notin_1 \omega^*$, then by Proposition 3, $(\neg \beta_i)(e) \wedge (\neg \beta_i)(\neg e) \in_1 \omega^*$. Second, assume Consistency on i 's beliefs. Since $B_i^*([e]) \cap B_i^*(\neg [e]) = \emptyset$, $\beta_i(e) \in_1 \omega^*$ and $\beta_i(\neg e) \in_1 \omega^*$ do not hold simultaneously. If $\beta_i(e) \in_1 \omega^*$ then $(\neg \beta_i)(e) \wedge (\neg \beta_i)(\neg e) \notin_1 \omega^*$. If $\beta_i(\neg e) \in_1 \omega^*$ then $(\neg \beta_i)(e) \wedge (\neg \beta_i)(\neg e) \notin_1 \omega^*$. Also, $(\neg \beta_i)(e) \wedge (\neg \beta_i)(\neg e) \in_1 \omega^*$ implies $\beta_i(e) \notin_1 \omega^*$ and $\beta_i(\neg e) \notin_1 \omega^*$. Conversely, suppose exactly one of the three conditions holds. If $\omega^* \in B_i^*([e])$ then $\beta_i(e) \in_1 \omega^*$. Then, $\beta_i(\neg e) \notin_1 \omega^*$, i.e., $\omega^* \in (\neg B_i^*)(\neg [e]) = (\neg B_i^*)([e]^c)$, establishing Consistency. \square

Proof of Theorem 2. *Step 1.* The proof consists of two steps. The first step shows that Ω^* is belief-closed. For (1a), take $\omega^* \in \Omega^*$. There is a unique nature state $s = \Theta^*(\omega^*)$

with $s \in_0 \omega^*$. For any $E \in \mathcal{A}_\kappa(\mathcal{S})$, $s \in E$ iff $\Theta^*(\omega^*) \in E$ iff $\omega^* \in (\Theta^*)^{-1}(E) = [E]$, i.e., $E \in_1 \omega^*$. Conditions (1b) to (1d) follow from Proposition 3.

Next, consider (2a). If $(e \leftrightarrow f)$ is valid in $\vec{\Omega}^*$, then $\llbracket e \rrbracket_{\vec{\Omega}^*} = \llbracket f \rrbracket_{\vec{\Omega}^*}$ and $B_i^*(\llbracket e \rrbracket_{\vec{\Omega}^*}) = B_i^*(\llbracket f \rrbracket_{\vec{\Omega}^*})$, i.e., $\llbracket \beta_i(e) \rrbracket_{\vec{\Omega}^*} = \llbracket \beta_i(f) \rrbracket_{\vec{\Omega}^*}$. It follows that $\llbracket \beta_i(e) \leftrightarrow \beta_i(f) \rrbracket_{\vec{\Omega}} = \Omega^*$, i.e., $(\beta_i(e) \leftrightarrow \beta_i(f))$ is valid in $\vec{\Omega}^*$.

For (2b), similarly to the above argument, if $(e \rightarrow f)$ is valid in $\vec{\Omega}^*$, then $\llbracket e \rrbracket_{\vec{\Omega}^*} \subseteq \llbracket f \rrbracket_{\vec{\Omega}^*}$ and thus $B_i^*(\llbracket e \rrbracket_{\vec{\Omega}^*}) \subseteq B_i^*(\llbracket f \rrbracket_{\vec{\Omega}^*})$, i.e., $\llbracket \beta_i(e) \rrbracket_{\vec{\Omega}^*} \subseteq \llbracket \beta_i(f) \rrbracket_{\vec{\Omega}^*}$. Then, $\llbracket \beta_i(e) \rightarrow \beta_i(f) \rrbracket_{\vec{\Omega}} = \Omega^*$, i.e., $(\beta_i(e) \rightarrow \beta_i(f))$ is valid in $\vec{\Omega}^*$.

For (2c), by supposition, $\bigcap \{\llbracket f \rrbracket_{\vec{\Omega}^*} \in \mathcal{D}^* \mid \omega^* \in B_i^*(\llbracket f \rrbracket_{\vec{\Omega}^*})\} \subseteq \llbracket e \rrbracket_{\vec{\Omega}^*}$. By the Kripke property of $\vec{\Omega}^*$, $\omega^* \in B_i^*(\llbracket e \rrbracket_{\vec{\Omega}^*}) = [\beta_i(e)]$, i.e., $\beta_i(e) \in_1 \omega^*$.

Next, I show (3). Fix $\omega^* \in \Omega^*$. It is enough to show that each of the following expressions is valid in $\vec{\Omega}^*$. For Necessitation, consider $\llbracket \beta_i(S) \rrbracket_{\vec{\Omega}^*} = B_i^*(\llbracket S \rrbracket_{\vec{\Omega}^*}) = B_i^*(\Omega^*) = \Omega^*$. For λ -Conjunction, consider: $\llbracket (\bigwedge_{e \in \mathcal{E}} \beta_i(e)) \rightarrow \beta_i(\bigwedge \mathcal{E}) \rrbracket_{\vec{\Omega}^*} = (\neg \bigcap_{e \in \mathcal{E}} B_i^*(\llbracket e \rrbracket_{\vec{\Omega}^*})) \cup B_i^*(\bigcap_{e \in \mathcal{E}} \llbracket e \rrbracket_{\vec{\Omega}^*}) = \Omega^*$. For Consistency, consider $\llbracket \beta_i(e) \rightarrow (\neg \beta_i)(\neg e) \rrbracket_{\vec{\Omega}^*} = (\neg B_i^*)(\llbracket e \rrbracket_{\vec{\Omega}^*}) \cup (\neg B_i^*)(\neg \llbracket e \rrbracket_{\vec{\Omega}^*}) = \Omega^*$. For Truth Axiom, consider $\llbracket \beta_i(e) \rightarrow e \rrbracket_{\vec{\Omega}^*} = (\neg B_i^*)(\llbracket e \rrbracket_{\vec{\Omega}^*}) \cup \llbracket e \rrbracket_{\vec{\Omega}^*} = \Omega^*$. For Positive Introspection, consider $\llbracket \beta_i(e) \rightarrow \beta_i \beta_i(e) \rrbracket_{\vec{\Omega}^*} = (\neg B_i^*)(\llbracket e \rrbracket_{\vec{\Omega}^*}) \cup B_i^* B_i^*(\llbracket e \rrbracket_{\vec{\Omega}^*}) = \Omega^*$. For Negative Introspection, consider $\llbracket (\neg \beta_i)(e) \rightarrow \beta_i(\neg \beta_i)(e) \rrbracket_{\vec{\Omega}^*} = B_i^*(\llbracket e \rrbracket_{\vec{\Omega}^*}) \cup B_i^*(\neg B_i^*)(\llbracket e \rrbracket_{\vec{\Omega}^*}) = \Omega^*$.

Step 2. The second step shows that $\Omega \subseteq \Omega^*$ for any belief-closed set Ω . To that end, I introduce a belief space $\vec{\Omega}$ defined on Ω , and show that the description map $D : \vec{\Omega} \rightarrow \vec{\Omega}^*$ is an inclusion map.

Step 2.1. By slightly abusing the notation, let $[e]_{\vec{\Omega}} := \{\omega \in \Omega \mid e \in_1 \omega\}$ for each $e \in \mathcal{L}$ (it turns out that $[.]_{\vec{\Omega}} = [.] \cap \Omega$). Let $\mathcal{D} := \{[e]_{\vec{\Omega}} \in \mathcal{P}(\Omega) \mid e \in \mathcal{L}\}$. I show (Ω, \mathcal{D}) is a κ -algebra. First, $\Omega = [S]_{\vec{\Omega}} \in \mathcal{D}$ and $\emptyset = [\emptyset]_{\vec{\Omega}} \in \mathcal{D}$. Second, $[\neg e]_{\vec{\Omega}} = ([e]_{\vec{\Omega}})^c$ follows because $\omega \in [\neg e]_{\vec{\Omega}}$ iff $(\neg e) \in_1 \omega$ iff $e \notin_1 \omega$ iff $\omega \in ([e]_{\vec{\Omega}})^c$. Third, I show $[\bigwedge \mathcal{E}]_{\vec{\Omega}} = \bigcap_{e \in \mathcal{E}} [e]_{\vec{\Omega}}$ for any $\mathcal{E} \in \mathcal{P}(\mathcal{L})$ with $0 < |\mathcal{E}| < \kappa$. Indeed, $\omega \in \bigcap_{e \in \mathcal{E}} [e]_{\vec{\Omega}}$ iff $e \in_1 \omega$ for all $e \in \mathcal{E}$ iff $\bigwedge \mathcal{E} \in_1 \omega$ iff $\omega \in [\bigwedge \mathcal{E}]_{\vec{\Omega}}$.

Step 2.2. Define $\Theta : \Omega \rightarrow S$ as follows: for each $\omega \in \Omega$, let $\Theta(\omega)$ be the unique $s \in S$ with $s \in_0 \omega$. By construction, it is a well-defined map. I show that the map Θ is a measurable map such that $(\Theta)^{-1}(E) = [E]_{\vec{\Omega}}$ for each $E \in \mathcal{A}_\kappa(\mathcal{S})$. If $\omega \in [E]_{\vec{\Omega}}$, then $E \in_1 \omega$. Hence, $\Theta(\omega) \in E$, i.e., $\omega \in \Theta^{-1}(E)$. Conversely, if $\omega \in \Theta^{-1}(E)$ then $\Theta(\omega) \in E$, and thus $E \in_1 \omega$. Hence, $\omega \in [E]_{\vec{\Omega}}$.

Step 2.3. Fix $i \in I$, and define i 's belief operator $B_i : \mathcal{D} \rightarrow \mathcal{D}$ by: $B_i([e]_{\vec{\Omega}}) := [\beta_i(e)]_{\vec{\Omega}}$ for each $[e]_{\vec{\Omega}} \in \mathcal{D}$. I show B_i is well-defined. If $[e]_{\vec{\Omega}} = [f]_{\vec{\Omega}}$, then $\llbracket (e \leftrightarrow f) \rrbracket_{\vec{\Omega}} = \Omega$. This implies $\llbracket (\beta_i(e) \leftrightarrow \beta_i(f)) \rrbracket_{\vec{\Omega}} = \Omega$. Thus, $[\beta_i(e)]_{\vec{\Omega}} = [\beta_i(f)]_{\vec{\Omega}}$.

Next, I show that B_i inherits assumptions on beliefs made in the given category. For Necessitation, $\Omega = [\beta_i(S)]_{\vec{\Omega}} = B_i([S]_{\vec{\Omega}}) = B_i(\Omega)$. For Monotonicity,

take $[e]_{\vec{\Omega}}, [f]_{\vec{\Omega}} \in \mathcal{D}$ with $[e]_{\vec{\Omega}} \subseteq [f]_{\vec{\Omega}}$. Then, $[e \rightarrow f]_{\vec{\Omega}} = \Omega$. It follows that $[\beta_i(e) \rightarrow \beta_i(f)]_{\vec{\Omega}} = \Omega$, i.e., $[\beta_i(e)]_{\vec{\Omega}} \subseteq [\beta_i(f)]_{\vec{\Omega}}$. Thus, $B_i([e]_{\vec{\Omega}}) \subseteq B_i([f]_{\vec{\Omega}})$.

For λ -Conjunction, take $\mathcal{E} \in \mathcal{P}(\mathcal{L})$ with $0 < |\mathcal{E}| < \lambda$. If $\omega \in \bigcap_{e \in \mathcal{E}} B_i([e]_{\vec{\Omega}}) = [\bigwedge_{e \in \mathcal{E}} \beta_i(e)]_{\vec{\Omega}}$ then $\bigwedge_{e \in \mathcal{E}} \beta_i(e) \in_1 \omega$. Since $(\bigwedge_{e \in \mathcal{E}} \beta_i(e) \rightarrow \beta_i(\bigwedge \mathcal{E})) \in_1 \omega$, it follows $\beta_i(\bigwedge \mathcal{E}) \in_1 \omega$, i.e., $\omega \in [\beta_i(\bigwedge \mathcal{E})]_{\vec{\Omega}} = B_i(\bigcap_{e \in \mathcal{E}} [e]_{\vec{\Omega}})$. For the Kripke property, $\omega \in B_i([e]_{\vec{\Omega}})$ for any $(\omega, [e]_{\vec{\Omega}}) \in \Omega \times \mathcal{D}$ with $\bigcap \{[f]_{\vec{\Omega}} \in \mathcal{D} \mid \omega \in B_i([f]_{\vec{\Omega}})\} \subseteq [e]_{\vec{\Omega}}$.

For Consistency, if $\omega \in B_i([e]_{\vec{\Omega}}) = [\beta_i(e)]_{\vec{\Omega}}$ then $\beta_i(e) \in_1 \omega$. Since $(\beta_i(e) \rightarrow (\neg \beta_i)(\neg e)) \in_1 \omega$, it follows $(\neg \beta_i)(\neg e) \in_1 \omega$, i.e., $\omega \in [(\neg \beta_i)(\neg e)]_{\vec{\Omega}} = (\neg B_i)(\neg [e]_{\vec{\Omega}})$. For Truth Axiom, if $\omega \in B_i([e]_{\vec{\Omega}}) = [\beta_i(e)]_{\vec{\Omega}}$ then $\beta_i(e) \in_1 \omega$. Since $(\beta_i(e) \rightarrow e) \in_1 \omega$, it follows $e \in_1 \omega$, i.e., $\omega \in [e]_{\vec{\Omega}}$. For Positive Introspection, if $\omega \in B_i([e]_{\vec{\Omega}}) = [\beta_i(e)]_{\vec{\Omega}}$ then $\beta_i(e) \in_1 \omega$. Since $(\beta_i(e) \rightarrow \beta_i \beta_i(e)) \in_1 \omega$, it follows $\beta_i \beta_i(e) \in_1 \omega$, i.e., $\omega \in [\beta_i \beta_i(e)]_{\vec{\Omega}} = B_i B_i([e]_{\vec{\Omega}})$. The proof for Negative Introspection is similar.

Step 2.4. So far, $\vec{\Omega} := \langle (\Omega, \mathcal{D}), (B_i)_{i \in I}, \Theta \rangle$ is shown to be a belief space of the given category. Finally, I demonstrate that the description map $D : \vec{\Omega} \rightarrow \vec{\Omega}^*$ is an inclusion map (consequently, $\vec{\Omega}$ is non-redundant and $\Omega \subseteq \Omega^*$) by showing that $[\cdot]_{\vec{\Omega}} : \mathcal{L} \rightarrow \mathcal{D}$, viewed as a mapping, coincides with the semantic interpretation function $\llbracket \cdot \rrbracket_{\vec{\Omega}}$ (consequently, $\vec{\Omega}$ is minimal).

I show by induction that $[\cdot]_{\vec{\Omega}} = \llbracket \cdot \rrbracket_{\vec{\Omega}}$. First, fix $E \in \mathcal{A}_\kappa(\mathcal{S})$. Then, $\omega \in \llbracket E \rrbracket_{\vec{\Omega}} = \Theta^{-1}(E)$ iff $\Theta(\omega) \in E$ iff $\omega \in [E]_{\vec{\Omega}}$. Second, suppose $\llbracket e \rrbracket_{\vec{\Omega}} = [e]_{\vec{\Omega}}$. Then, $\omega \in \llbracket \neg e \rrbracket_{\vec{\Omega}} = \neg [e]_{\vec{\Omega}}$ iff $\omega \notin [e]_{\vec{\Omega}}$ iff $\omega \notin [e]_{\vec{\Omega}}$ iff $\omega \in [\neg e]_{\vec{\Omega}}$. Also, $\llbracket \beta_i(e) \rrbracket_{\vec{\Omega}} = B_i([e]_{\vec{\Omega}}) = B_i([e]_{\vec{\Omega}}) = [\beta_i(e)]_{\vec{\Omega}}$. Third, suppose $\llbracket e \rrbracket_{\vec{\Omega}} = [e]_{\vec{\Omega}}$ for all $e \in \mathcal{E}$ with $\mathcal{E} \in \mathcal{P}(\mathcal{L})$ and $0 < |\mathcal{E}| < \kappa$. Then, $\omega \in \llbracket \bigwedge \mathcal{E} \rrbracket_{\vec{\Omega}} = \bigcap_{e \in \mathcal{E}} [e]_{\vec{\Omega}}$ iff $\omega \in \bigcap_{e \in \mathcal{E}} [e]_{\vec{\Omega}} = [\bigwedge \mathcal{E}]_{\vec{\Omega}}$.

I show that $D(\omega) = \omega$ for all $\omega \in \Omega$. First, $e \in_1 D(\omega)$ iff $D(\omega) \in [e]$ iff $\omega \in D^{-1}([e]) = \llbracket e \rrbracket_{\vec{\Omega}} = [e]_{\vec{\Omega}}$ iff $e \in_1 \omega$. Second, $s \in_0 \omega$ iff $s = \Theta(\omega) = \Theta^*(D(\omega))$ iff $s \in_0 D(\omega)$. Hence, D is an inclusion map. \square

A.3 Section 5

Proof of Corollary 2. Construct $\vec{\Omega}^* := \langle (\Omega^*, \mathcal{D}^*), (B_i^{*p})_{(i,p) \in I \times [0,1]}, \Theta^* \rangle$, similarly to the proof of Theorem 1. The set Ω^* is not empty (consider $\overrightarrow{\{s\}}$). To see that $\vec{\Omega}^*$ is terminal, it suffices to show that the p -belief operators B_i^{*p} satisfy the properties specified in Definition 10 (2).

First, (2a), (2b), (2d), and (2h) follow from Lemma A.1 (1a). Next, (2c) follows from Lemma A.1 (2a). Next, (2g) follows from Lemma A.1 (2a) and (3a).

Next, (2e) and (2f) follow from the following variant of Lemma A.1 (1a). Let $f_{\vec{\Omega}} : \mathcal{D}^2 \rightarrow \mathcal{D}$ and $g_{\vec{\Omega}} : \mathcal{D}^2 \rightarrow \mathcal{D}$ be defined in each probabilistic-belief space $\vec{\Omega}$ and satisfy, for any measurable map $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$, $\varphi^{-1} f_{\vec{\Omega}}(E', F') = f_{\vec{\Omega}}(\varphi^{-1}(E'), \varphi^{-1}(F'))$ and $\varphi^{-1} g_{\vec{\Omega}}(E', F') = g_{\vec{\Omega}}(\varphi^{-1}(E'), \varphi^{-1}(F'))$ for all $E', F' \in \mathcal{D}'$. If $f_{\vec{\Omega}}(E, F) \subseteq g_{\vec{\Omega}}(E, F)$ (for all $E, F \in \mathcal{D}$) for every probabilistic-belief space $\vec{\Omega}$, then $f_{\vec{\Omega}^*}([e], [f]) \subseteq g_{\vec{\Omega}^*}([e], [f])$ for all $[e], [f] \in \mathcal{D}^*$.

For (2i), if $[m_{B_i^*}^*(\omega^*)] \subseteq [e]$ then $[m_{B_i}(\omega)] \subseteq D^{-1}[e]$ and thus $\omega \in B_i^1(D^{-1}[e]) = D^{-1}B_i^{*1}([e])$, where a probabilistic-belief space $\overrightarrow{\Omega}$ and $\omega \in \Omega$ satisfy $\omega^* = D(\omega)$. \square

To prove Proposition 4, I show the following preliminary result.

Lemma A.2 (Extension of p -Belief Operators). *Let (Ω, \mathcal{D}) be an \aleph_0 -algebra, and let $(B_i^p)_{p \in [0,1]}$ be a collection of player i 's p -belief operators $B_i^p : \mathcal{D} \rightarrow \mathcal{D}$ satisfying Definition 10 (2). Then, there is a unique collection $(\overline{B}_i^p)_{p \in [0,1]}$ of p -belief operators $\overline{B}_i^p : \sigma(\mathcal{D}) \rightarrow \sigma(\mathcal{D})$ satisfying Definition 10 (2) and $\overline{B}_i^p|_{\mathcal{D}} = B_i^p$.*

Proof of Lemma A.2. In this proof, denote explicitly by $\Delta(\Omega, \mathcal{D})$ the set of countably-additive probability measures on the \aleph_0 -algebra (Ω, \mathcal{D}) . Denote by $m_{B_i} : (\Omega, \mathcal{D}) \rightarrow (\Delta(\Omega, \mathcal{D}), \mathcal{A}_{\aleph_0}(\{\{\mu \in \Delta(\Omega, \mathcal{D}) \mid \mu(E) \geq p\} \mid (E, p) \in \mathcal{D} \times [0, 1]\}))$ the measurable type mapping associated with $(B_i^p)_{p \in [0,1]}$. That is, $m_{B_i}(\omega)(E) := \sup\{p \in [0, 1] \mid \omega \in B_i^p(E)\}$ for each $E \in \mathcal{D}$.

To construct \overline{B}_i^p , let $\overline{m}_{B_i}(\omega) \in \Delta(\Omega, \sigma(\mathcal{D}))$ be the unique Carathéodory extension of $m_{B_i}(\omega) \in \Delta(\Omega, \mathcal{D})$. Denote by $\Sigma := \sigma(\{\{\mu \in \Delta(\Omega, \sigma(\mathcal{D})) \mid \mu(E) \geq p\} \mid (E, p) \in \sigma(\mathcal{D}) \times [0, 1]\})$ the σ -algebra on $\Delta(\Omega, \sigma(\mathcal{D}))$. By Heifetz and Samet (1998b, Lemma 4.5), $\Sigma = \sigma(\{\{\mu \in \Delta(\Omega, \sigma(\mathcal{D})) \mid \mu(E) \geq p\} \mid (E, p) \in \mathcal{D} \times [0, 1]\})$.

I show $\overline{m}_{B_i} : (\Omega, \sigma(\mathcal{D})) \rightarrow (\Delta(\Omega, \sigma(\mathcal{D})), \Sigma)$ is measurable. For each $(E, p) \in \mathcal{D} \times [0, 1]$, it follows from $\{\omega \in \Omega \mid \overline{m}_{B_i}(\omega)(E) \geq p\} = \{\omega \in \Omega \mid m_{B_i}(\omega)(E) \geq p\}$ that $\overline{m}_{B_i}^{-1}(\{\mu \in \Delta(\Omega, \sigma(\mathcal{D})) \mid \mu(E) \geq p\}) \in \sigma(\mathcal{D})$. Hence, define $\overline{B}_i^p(E) := \{\omega \in \Omega \mid \overline{m}_{B_i}(\omega)(E) \geq p\} \in \sigma(\mathcal{D})$ for each $(E, p) \in \sigma(\mathcal{D}) \times [0, 1]$. One can show that the collection $(\overline{B}_i^p)_{p \in [0,1]}$ satisfies Definition 10 (2) and $\overline{B}_i^p|_{\mathcal{D}} = B_i^p$.

To show uniqueness, let \tilde{B}_i^p be an extension. If $\omega \in \tilde{B}_i^p(E)$, then $\overline{m}_{B_i}(\omega)(E) = m_{\tilde{B}_i}(\omega)(E) \geq p$. Thus, $\omega \in \overline{B}_i^p(E)$. The converse holds similarly. \square

Proof of Proposition 4. For ease of notation, denote $\mathcal{L}_\lambda^I = \mathcal{L}_\lambda^I(\mathcal{A}_\kappa(\mathcal{S}))$. Also, denote $\mathcal{L} = \mathcal{L}_\kappa^I$. Let $(\mathcal{L}_\alpha)_{\alpha=0}^{\bar{\kappa}}$ be the auxiliary sequence that generates \mathcal{L} as in Remark 3.

Define $[\mathcal{L}_\lambda^I] := \{[e] \in \mathcal{D}^* \mid e \in \mathcal{L}_\lambda^I\}$. As in Lemma 1, $[\mathcal{L}_\lambda^I]$ is an algebra on Ω^* . Since $\mathcal{L}_\lambda^I \subseteq \mathcal{L}$ and since \mathcal{D}^* is a σ -algebra, $\sigma([\mathcal{L}_\lambda^I]) \subseteq \mathcal{D}^*$. To prove the converse set inclusion, I show $[e] \in \sigma([\mathcal{L}_\lambda^I])$ for all $e \in \mathcal{L}$. To that end, I first establish $B_i^{*p}([e]) \in \sigma([\mathcal{L}_\lambda^I])$ for any $[e] \in \sigma([\mathcal{L}_\lambda^I])$, i.e., $B_i^{*p}|_{\sigma([\mathcal{L}_\lambda^I])} : \sigma([\mathcal{L}_\lambda^I]) \rightarrow \sigma([\mathcal{L}_\lambda^I])$. Applying Lemma 3 to $B_i^{*p}|_{[\mathcal{L}_\lambda^I]} : [\mathcal{L}_\lambda^I] \rightarrow [\mathcal{L}_\lambda^I]$, the operator $B_i^{*p}|_{[\mathcal{L}_\lambda^I]}$ satisfies Definition 10 (2) (with the slight modification that events are restricted to $[\mathcal{L}_\lambda^I]$). It follows from Lemma A.2 that $B_i^{*p}|_{[\mathcal{L}_\lambda^I]}$ uniquely extends to $\sigma([\mathcal{L}_\lambda^I])$, and it coincides with $B_i^{*p}|_{\sigma([\mathcal{L}_\lambda^I])}$. Then, $B_i^{*p}|_{\sigma([\mathcal{L}_\lambda^I])} : \sigma([\mathcal{L}_\lambda^I]) \rightarrow \sigma([\mathcal{L}_\lambda^I])$.

Now, I show $[e] \in \sigma([\mathcal{L}_\lambda^I])$ for all $e \in \mathcal{L}$ by induction on the construction of \mathcal{L} (recall Remark 3). For $\alpha = 0$, $[E] \in \sigma([\mathcal{L}_\lambda^I])$ for all $E \in \mathcal{L}_0 = \mathcal{A}_\kappa(\mathcal{S})$. Suppose $[e] \in \sigma([\mathcal{L}_\lambda^I])$ for all $e \in \bigcup_{\beta < \alpha} \mathcal{L}_\beta$. For any $e \in \bigcup_{\beta < \alpha} \mathcal{L}_\beta$, $[\beta_i^p(e)] = B_i^{*p}|_{\sigma([\mathcal{L}_\lambda^I])}([e]) \in \sigma([\mathcal{L}_\lambda^I])$ (note that $\beta_i^p(e)$ is the (syntactic) expression for “ i p -believes e ”). Thus, $[e] \in \sigma([\mathcal{L}_\lambda^I])$ for all $e \in \mathcal{L}'_\alpha$. Then, $[\neg e] = \neg[e] \in \sigma([\mathcal{L}_\lambda^I])$ for any $e \in \mathcal{L}'_\alpha$. Also, $[\bigwedge \mathcal{F}] = \bigcap_{e \in \mathcal{F}} [e] \in \sigma([\mathcal{L}_\lambda^I])$ for any $\mathcal{F} \subseteq \mathcal{L}'_\alpha$ with $0 < |\mathcal{F}| < \kappa (= \aleph_1)$. \square

A.4 Section 6

Proof of Corollary 3. In each belief space (with a common belief operator), identify the common belief operator as the belief operator of a hypothetical player who represents common belief. As in Section 3, construct a candidate terminal κ -belief space $\overrightarrow{\Omega^*} = \langle (\Omega^*, \mathcal{D}^*), (B_i^*)_{i \in I}, C^*, \Theta^* \rangle$. For any belief space $\overrightarrow{\Omega} = \langle (\Omega, \mathcal{D}), (B_i)_{i \in I}, C, \Theta \rangle$, $D^{-1}C^*[\cdot] = CD^{-1}[\cdot]$. To show that $\overrightarrow{\Omega^*}$ is terminal, it suffices to show that C^* is a common belief operator: $C^*[e] = \max\{[f] \in \mathcal{J}_I^* \mid [f] \subseteq B_I^*[e]\}$ for any $[e] \in \mathcal{D}^*$.

Let $\omega^* \in C^*[e]$. There are a belief space $\overrightarrow{\Omega}$ and a state $\omega \in \Omega$ with $\omega^* = D(\omega)$. Thus, $\omega \in D^{-1}C^*[e] = CD^{-1}[e]$. Since $CD^{-1}[e] \subseteq B_I D^{-1}[e] = D^{-1}B_I^*[e]$, $\omega^* = D(\omega) \in B_I^*[e]$. Thus, $C^*[e] \subseteq B_I^*[e]$. To show $C^*[e] \in \mathcal{J}_I^*$, take any $[f] \in \mathcal{D}^*$ with $C^*[e] \subseteq [f]$. Take $\omega^* \in C^*[e]$. There are a belief space $\overrightarrow{\Omega}$ and a state $\omega \in \Omega$ with $\omega^* = D(\omega)$. Since $CD^{-1}[e] = D^{-1}C^*[e] \subseteq D^{-1}[f]$, it follows $D^{-1}C^*[e] = CD^{-1}[e] \subseteq B_I D^{-1}[f] = D^{-1}B_I^*[f]$. Thus, $\omega^* = D(\omega) \in B_I^*[f]$. It follows that $C^*[e] \subseteq \max\{[f] \in \mathcal{J}_I^* \mid [f] \subseteq B_I^*[e]\}$.

To get the converse set inclusion, take any $[f] \in \mathcal{J}_I^*$ with $[f] \subseteq B_I^*[e]$. If $\omega^* \in [f]$, then there are a belief space $\overrightarrow{\Omega}$ and a state $\omega \in \Omega$ with $\omega \in D^{-1}[f] \subseteq D^{-1}B_I^*[e] = B_I D^{-1}[e]$. For the belief space $\overrightarrow{\Omega}$, consider $\overrightarrow{\Omega'} = \langle (\Omega, \mathcal{D}'), (B_i|_{\mathcal{D}'}), C|_{\mathcal{D}'}, \Theta \rangle$ with $\mathcal{D}' = D^{-1}(\mathcal{D}^*)$. One can show that $\overrightarrow{\Omega'}$ is a belief space and that the identify map $\text{id}_\Omega : \overrightarrow{\Omega} \rightarrow \overrightarrow{\Omega'}$ is a morphism (see also Remark 7 in Section 7.2). Then, $D_{\overrightarrow{\Omega}} = D_{\overrightarrow{\Omega'}} \circ \text{id}_\Omega$. Since one can retake $\omega^* = D_{\overrightarrow{\Omega}}(\omega)$, without loss, assume $\mathcal{D} = D^{-1}(\mathcal{D}^*)$.

To establish $\omega^* \in C^*[e]$, it suffices to show that $D^{-1}[f] \in \mathcal{J}_I$, because it implies $\omega \in D^{-1}[f] \subseteq C(D^{-1}[e]) = D^{-1}C^*[e]$. Take any $E' \in \mathcal{D} = D^{-1}(\mathcal{D}^*)$ with $D^{-1}[f] \subseteq E'$. Without loss, one can assume $E' = D^{-1}[e']$ for some $[e'] \in \mathcal{D}^*$ with $[f] \subseteq [e']$, because $D^{-1}[e' \vee f] = D^{-1}([e'] \cup [f]) = D^{-1}[e'] \cup D^{-1}[f] = D^{-1}[e'] = E'$. Since $[f] \in \mathcal{J}_I^*$, it follows $[f] \subseteq B_I^*[e']$. Thus, $D^{-1}[f] \subseteq D^{-1}B_I^*[e'] = B_I D^{-1}[e'] = B_I(E')$. \square

Proof of Lemma 5. Define an operator $g : \prod_{i \in I} \mathcal{P}(A_i) \rightarrow \prod_{i \in I} \mathcal{P}(A_i)$ as follows: for any $X = (X_i)_{i \in I} \in \prod_{i \in I} \mathcal{P}(A_i)$, let

$$g(X) := \prod_{i \in I} \{a_i \in X_i \mid a_i \text{ is not B-dominated given } X_{-i}\}.$$

Observe that g is monotone: for any $X = (X_i)_{i \in I}$ and $Y = (Y_i)_{i \in I}$, if $X_i \subseteq Y_i$ for all $i \in I$, then $g(X) \subseteq g(Y)$. This is because if $a_i \in X_i \subseteq Y_i$ is B-dominated given Y_{-i} then it is B-dominated given X_{-i} .

Let $(A^\beta)_\beta$ be a process of iterated elimination of B-dominated actions. I show by induction that, for any $T \in \prod_{i \in I} \mathcal{P}(A_i)$ with $g(T) = T$, $T \subseteq A^\beta$ for all β . Fix such T . For $\beta = 0$, $T \subseteq A = A^0$ follows by definition. For some $\beta = \gamma + 1$, suppose that $T \subseteq A^\gamma$. Since g is monotone, it follows that $T = g(T) \subseteq g(A^\gamma) \subseteq A^\beta$. For a non-zero limit ordinal β , if $T \subseteq A^\gamma$ for all $\gamma < \beta$ then $T \subseteq A^\beta = \bigcap_{\gamma: \gamma < \beta} A^\gamma$. Thus, terminal sets of different processes include each other so that they are all the same. \square

A.5 Section 7

Proof of Proposition 7. Part (1). Let $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}'$ be a morphism. In a similar way to Heifetz and Samet (1998a), I show by induction that $\mathcal{C}_\alpha = \varphi^{-1}(\mathcal{C}'_\alpha)$ for all α . For $\alpha = 0$, $\mathcal{C}_0 = \varphi^{-1}(\mathcal{C}'_0)$ follows because $(\Theta)^{-1}(E) = \varphi^{-1}((\Theta')^{-1}(E))$ for any $E \in \mathcal{A}_\kappa(\mathcal{S})$. Suppose $\mathcal{C}_\beta = \varphi^{-1}(\mathcal{C}'_\beta)$ for all $\beta < \alpha$. Then,

$$\begin{aligned}\mathcal{C}_\alpha &= \mathcal{A}_\kappa(\{\varphi^{-1}(E') \in \mathcal{D} \mid E' \in \bigcup_{\beta < \alpha} \mathcal{C}'_\beta\} \cup \bigcup_{i \in I} \{\varphi^{-1}(B'_i(E')) \in \mathcal{D} \mid E' \in \bigcup_{\beta < \alpha} \mathcal{C}'_\beta\}) \\ &= \varphi^{-1}(\mathcal{A}_\kappa(\{E' \in \mathcal{D} \mid E' \in \bigcup_{\beta < \alpha} \mathcal{C}'_\beta\} \cup \bigcup_{i \in I} \{B'_i(E') \in \mathcal{D} \mid E' \in \bigcup_{\beta < \alpha} \mathcal{C}'_\beta\})) = \varphi^{-1}(\mathcal{C}'_\alpha).\end{aligned}$$

The first and last equations follow from the definitions of \mathcal{C}_α and \mathcal{C}'_α , respectively. The second equation follows because φ^{-1} commutes with $\mathcal{A}_\kappa(\cdot)$. Then, $\mathcal{C}'_\alpha = \mathcal{C}'_{\alpha+1}$ implies $\mathcal{C}_\alpha = \mathcal{C}_{\alpha+1}$. Hence, if the κ -rank of $\vec{\Omega}'$ is α then that of $\vec{\Omega}$ is at most α .

Part (2). Fix a λ -belief space $\vec{\Omega}$, where $\lambda \geq \kappa$. Let $\mathcal{D}_\alpha := \mathcal{A}_\kappa(\{[e]_{\vec{\Omega}} \in \mathcal{D} \mid e \in \mathcal{L}_\alpha\})$ for each $\alpha \leq \bar{\kappa}$, where $(\mathcal{L}_\alpha)_{\alpha=0}^{\bar{\kappa}}$ is defined as in Remark 3 so that $\mathcal{L} = \mathcal{L}_{\bar{\kappa}}$. I show $\mathcal{D}_\alpha = \mathcal{C}_\alpha$ for all $\alpha \leq \bar{\kappa}$. For $\alpha = 0$, $\mathcal{D}_0 = \{\Theta^{-1}(E) \in \mathcal{D} \mid E \in \mathcal{A}_\kappa(\mathcal{S})\} = \mathcal{C}_0$. If $\mathcal{D}_\beta = \mathcal{C}_\beta$ for all $\beta < \alpha$, then $\mathcal{D}_\alpha = \mathcal{A}_\kappa(\{\bigcup_{\beta < \alpha} \mathcal{D}_\beta\} \cup \bigcup_{i \in I} \{B_i([e]) \in \mathcal{D} \mid [e] \in \bigcup_{\beta < \alpha} \mathcal{D}_\beta\}) = \mathcal{C}_\alpha$. Hence, $\mathcal{C}_{\bar{\kappa}} = \mathcal{D}_{\bar{\kappa}} = \{[e]_{\vec{\Omega}} \in \mathcal{D} \mid e \in \mathcal{L}\}$, implying $\mathcal{C}_{\bar{\kappa}} = \mathcal{C}_{\bar{\kappa}+1}$, i.e., the κ -rank of $\vec{\Omega}$ is at most $\bar{\kappa}$. \square

Proof of Proposition 8. Part (1). Since the λ -belief space $\vec{\Omega}_\lambda^*$ is also a κ -belief space, there is a unique morphism $D_{\vec{\Omega}_\lambda^*} : \vec{\Omega}_\lambda^* \rightarrow \vec{\Omega}_\kappa^*$, which takes the following form. While each $\omega^* = \{s \in S \mid s \in_0 \omega^*\} \sqcup \{e \in \mathcal{L}_\lambda^I(\mathcal{A}_\lambda(\mathcal{S})) \mid e \in_1 \omega^*\} \in \Omega_\lambda^*$ consists of the unique nature state s and expressions $e \in \mathcal{L}_\lambda^I(\mathcal{A}_\lambda(\mathcal{S}))$ that obtain, $D_{\vec{\Omega}_\lambda^*}(\omega^*) = \{s \in S \mid s \in_0 \omega^*\} \sqcup \{e \in \mathcal{L}_\kappa^I(\mathcal{A}_\kappa(\mathcal{S})) \mid e \in_1 \omega^*\}$ consists of the same unique nature state s and expressions $e \in \mathcal{L}_\kappa^I(\mathcal{A}_\kappa(\mathcal{S}))$ (observe $\mathcal{L}_\kappa^I(\mathcal{A}_\kappa(\mathcal{S})) \subseteq \mathcal{L}_\lambda^I(\mathcal{A}_\lambda(\mathcal{S}))$) that obtain. Thus, $D_{\vec{\Omega}_\lambda^*}$ is a surjective morphism. Hence, $|\Omega_\kappa^*| \leq |\Omega_\lambda^*|$.

Part (2). To simplify the proof, I make the following assumptions. Since the proof does not depend on the cardinality of I , let $I = \{1, 2\}$. Next, by the (second) remark following Definition 7, assume all the properties of beliefs in Definition 2. Next, assume $(S, \mathcal{S}) = (\{s_0, s_1\}, \mathcal{P}(S))$ (the proof goes through by taking $s_1 \in E$ and $s_0 \in E^c$ for E in the statement of the proposition). First, the knowledge space $\vec{\Omega}$ constructed by Hart, Heifetz, and Samet (1996) is a non-redundant κ -knowledge space with $|\Omega| \geq 2^{\aleph_0}$. Since the morphism $D_{\vec{\Omega}}$ is injective, $2^{\aleph_0} \leq |\Omega| \leq |\Omega_\kappa^*|$. Second, Heifetz and Samet (1998a, Theorem 2.5) construct a non-redundant κ -knowledge space $\vec{\Omega}'$ with $|\Omega'| = \kappa$. Since the morphism $D_{\vec{\Omega}'}$ is injective, $\kappa = |\Omega'| \leq |\Omega_\kappa^*|$. \square

Remark A.1 (Extension of the Domain). Let (Ω, \mathcal{D}) be a κ -algebra. Let $B_i : \mathcal{D} \rightarrow \mathcal{D}$ satisfy the Kripke property. Define $\bar{B}_i : \mathcal{P}(\Omega) \rightarrow \mathcal{P}(\Omega)$ by $\bar{B}_i(E) := \{\omega \in \Omega \mid b_{B_i}(\omega) \subseteq$

$E\}$ for each $E \in \mathcal{P}(\Omega)$. By construction, \overline{B}_i satisfy the Kripke property. Also, \overline{B}_i inherits Consistency, Truth Axiom, Positive Introspection, and Negative Introspection from B_i . Moreover, $B_i = \overline{B}_i|_{\mathcal{D}}$ and $b_{\overline{B}_i} = b_{B_i}$.

References

- [1] W. Armbruster and W. Böge. “Bayesian Game Theory”. *Game Theory and Related Topics*. Ed. by O. Moeschlin and D. Pallaschke. North-Holland, 1979, 17–28.
- [2] R. J. Aumann. “Agreeing to Disagree”. *Ann. Statist.* 4 (1976), 1236–1239.
- [3] R. J. Aumann. “Correlated Equilibrium as an Expression of Bayesian Rationality”. *Econometrica* 55 (1987), 1–18.
- [4] R. J. Aumann. “Interactive Epistemology I, II”. *Int. J. Game Theory* 28 (1999), 263–300, 301–314.
- [5] M. Bacharach. “Some Extensions of a Claim of Aumann in an Axiomatic Model of Knowledge”. *J. Econ. Theory* 37 (1985), 167–190.
- [6] P. Battigalli and G. Bonanno. “The Logic of Belief Persistence”. *Econ. Philos.* 13 (1997), 39–59.
- [7] P. Battigalli and M. Dufwenberg. “Dynamic Psychological Games”. *J. Econ. Theory* 144 (2009), 1–35.
- [8] P. Battigalli and M. Siniscalchi. “Hierarchies of Conditional Beliefs and Interactive Epistemology in Dynamic Games”. *J. Econ. Theory* 88 (1999), 188–230.
- [9] P. Battigalli and M. Siniscalchi. “Strong Belief and Forward Induction Reasoning”. *J. Econ. Theory* 106 (2002), 356–391.
- [10] A. Bjoern Dahl, J. Y. Halpern, and R. Pass. “Language-based Games”. *Proc. 14th Conference on Theoretical Aspects of Rationality and Knowledge*. Ed. by B. C. Schipper. 2013, 39–48.
- [11] W. Böge and T. Eisele. “On Solutions of Bayesian Games”. *Int. J. Game Theory* 8 (1979), 193–215.
- [12] G. Bonanno. “A Syntactic Approach to Rationality in Games with Ordinal Payoffs”. *Logic and the Foundations of Game and Decision Theory (LOFT 7)*. Ed. by G. Bonanno, W. van der Hoek, and M. Wooldridge. Amsterdam University Press, 2008, 59–86.
- [13] G. Bonanno and E. Tsakas. “Common Belief of Weak-dominance Rationality in Strategic-form Games: A Qualitative Analysis”. *Games Econ. Behav.* 112 (2018), 231–241.
- [14] T. Börgers. “Pure Strategy Dominance”. *Econometrica* 61 (1993), 423–430.
- [15] A. Brandenburger. “On the Existence of a “Complete” Possibility Structure”. *Cognitive Processes and Economic Behavior*. Ed. by N. Dimitri, M. Basili, and I. Gilboa. Routledge, 2003, 30–34.
- [16] A. Brandenburger and E. Dekel. “Rationalizability and Correlated Equilibria”. *Econometrica* 55 (1987), 1391–1402.
- [17] A. Brandenburger and E. Dekel. “Hierarchies of Beliefs and Common Knowledge”. *J. Econ. Theory* 59 (1993), 189–198.
- [18] A. Brandenburger, E. Dekel, and J. Geanakoplos. “Correlated Equilibrium with Generalized Information Structures”. *Games Econ. Behav.* 4 (1992), 182–201.

- [19] A. Brandenburger, A. Friedenberg, and H. J. Keisler. “Admissibility in Games”. *Econometrica* 76 (2008), 307–352.
- [20] A. Brandenburger and H. J. Keisler. “An Impossibility Theorem on Beliefs in Games”. *Studia Logica* 84 (2006), 211–240.
- [21] Y.-C. Chen, N. V. Long, and X. Luo. “Iterated Strict Dominance in General Games”. *Games Econ. Behav.* 61 (2007), 299–315.
- [22] E. Dekel and F. Gul. “Rationality and Knowledge in Game Theory”. *Advances in Economics and Econometrics: Theory and Applications, Seventh World Congress*. Ed. by D. M. Kreps and K. F. Wallis. Vol. 1. Cambridge University Press, 1997, 87–172.
- [23] A. Di Tillio. “Subjective Expected Utility in Games”. *Theor. Econ.* 3 (2008), 287–323.
- [24] M. Dufwenberg and M. Stegeman. “Existence and Uniqueness of Maximal Reductions under Iterated Strict Dominance”. *Econometrica* 70 (2002), 2007–2023.
- [25] R. Fagin. “A Quantitative Analysis of Modal Logic”. *J. Symb. Log.* 59 (1994), 209–252.
- [26] R. Fagin, J. Y. Halpern, and M. Y. Vardi. “A Model-theoretic Analysis of Knowledge”. *J. ACM* 38 (1991), 382–428.
- [27] R. Fagin, J. Geanakoplos, J. Y. Halpern, and M. Y. Vardi. “The Hierarchical Approach to Modeling Knowledge and Common Knowledge”. *Int. J. Game Theory* 28 (1999), 331–365.
- [28] A. Friedenberg. “When Do Type Structures Contain All Hierarchies of Beliefs?” *Games Econ. Behav.* 68 (2010), 108–129.
- [29] A. Friedenberg and H. J. Keisler. “Iterated Dominance Revisited”. *Econ. Theory* 72 (2021), 377–421.
- [30] A. Friedenberg and M. Meier. “On the Relationship between Hierarchy and Type Morphisms”. *Econ. Theory* 46 (2011), 377–399.
- [31] A. Friedenberg and M. Meier. “The Context of the Game”. *Econ. Theory* 63 (2017), 347–386.
- [32] S. Fukuda. “Epistemic Foundations for Set-algebraic Representations of Knowledge”. *J. Math. Econ.* 84 (2019), 73–82.
- [33] S. Fukuda. “Formalizing Common Belief with No Underlying Assumption on Individual Beliefs”. *Games Econ. Behav.* 121 (2020), 169–189.
- [34] J. Geanakoplos. “Game Theory without Partitions, and Applications to Speculation and Consensus”. *B. E. J. Theor. Econ.* 21 (2021), 361–394.
- [35] J. Geanakoplos, D. Pearce, and E. Stacchetti. “Psychological Games and Sequential Rationality”. *Games Econ. Behav.* 1 (1989), 60–79.
- [36] I. Gilboa. “Information and Meta-Information”. *Proc. 2nd Conference on Theoretical Aspects of Reasoning about Knowledge*. Morgan Kaufmann Publishers Inc., 1988, 227–243.
- [37] P. Guarino. “The Topology-Free Construction of the Universal Type Structure for Conditional Probability Systems”. *Proc. 16th Conference on Theoretical Aspects of Rationality and Knowledge*. Ed. by J. Lang. 2017.
- [38] J. C. Harsanyi. “Games with Incomplete Information Played by “Bayesian” Players, I-III”. *Manag. Sci.* 14 (1967-68), 159–182, 320–334, 486–502.
- [39] S. Hart, A. Heifetz, and D. Samet. ““Knowing Whether,” “Knowing That,” and The Cardinality of State Spaces”. *J. Econ. Theory* 70 (1996), 249–256.
- [40] A. Heifetz and D. Samet. “Knowledge Spaces with Arbitrarily High Rank”. *Games Econ. Behav.* 22 (1998), 260–273.

- [41] A. Heifetz and D. Samet. “Topology-Free Typology of Beliefs”. *J. Econ. Theory* 82 (1998), 324–341.
- [42] A. Heifetz and D. Samet. “Hierarchies of Knowledge: An Unbounded Stairway”. *Math. Soc. Sci.* 38 (1999), 157–170.
- [43] K. Hrbacek and T. Jech. Introduction to Set Theory. Third Edition. CRC Press, 1999.
- [44] B. L. Lipman. “A Note on the Implications of Common Knowledge of Rationality”. *Games Econ. Behav.* 6 (1994), 114–129.
- [45] M. Meier. “On the Nonexistence of Universal Information Structures”. *J. Econ. Theory* 122 (2005), 132–139.
- [46] M. Meier. “Finitely Additive Beliefs and Universal Type Spaces”. *Ann. Probab.* 34 (2006), 386–422.
- [47] M. Meier. “Universal Knowledge-Belief Structures”. *Games Econ. Behav.* 62 (2008), 53–66.
- [48] M. Meier. “An Infinitary Probability Logic for Type Spaces”. *Isr. J. Math.* 192 (2012), 1–58.
- [49] J. F. Mertens and S. Zamir. “Formulation of Bayesian Analysis for Games with Incomplete Information”. *Int. J. Game Theory* 14 (1985), 1–29.
- [50] D. Monderer and D. Samet. “Approximating Common Knowledge with Common Beliefs”. *Games Econ. Behav.* 1 (1989), 170–190.
- [51] S. Morris. “The Logic of Belief and Belief Change: A Decision Theoretic Approach”. *J. Econ. Theory* 69 (1996), 1–23.
- [52] L. S. Moss and I. D. Viglizzo. “Harsanyi Type Spaces and Final Coalgebras Constructed from Satisfied Theories”. *Electron. Notes Theor. Comput. Sci.* 106 (2004), 279–295.
- [53] L. S. Moss and I. D. Viglizzo. “Final Coalgebras for Functors on Measurable Spaces”. *Inf. Comput.* 204 (2006), 610–636.
- [54] M. Pintér. “Type Spaces with Non-Additive Beliefs”. 2012.
- [55] D. Samet. “Ignoring Ignorance and Agreeing to Disagree”. *J. Econ. Theory* 52 (1990), 190–207.
- [56] D. Samet. “Quantified Beliefs and Believed Quantities”. *J. Econ. Theory* 95 (2000), 169–185.
- [57] B. C. Schipper. “Awareness”. *Handbook of Epistemic Logic*. Ed. by H. van Ditmarsch, J. Y. Halpern, W. van der Hoek, and B. P. Kooi. College Publications, 2015, 77–146.
- [58] H. S. Shin. “Logical Structure of Common Knowledge”. *J. Econ. Theory* 60 (1993), 1–13.
- [59] R. Stalnaker. “On the Evaluation of Solution Concepts”. *Theory Decis.* 37 (1994), 49–73.
- [60] T. C.-C. Tan and S. R.d. C. Werlang. “The Bayesian Foundations of Solution Concepts of Games”. *J. Econ. Theory* 45 (1988), 370–391.