The Logic of Knightian Games

Shiri Alon* Aviad Heifetz†

January 24, 2014

Abstract

We devise a sound and complete epistemic-logic axiomatization for *Knightian type spaces* – the generalization of Harsanyi type spaces employed in strategic games with asymmetric *uncertainty/ambiguity.* In a Knightian type space, each type's epistemic attitude is represented by a *set* of probability measures. The axiomatization unravels how each such epistemic attitude embodies a (potentially) *partial* likelihood relation over formulas, and conversely, each partial likelihood relation over formulas is representable by a Knightian type.

Keywords: Knightian uncertainty, type spaces, epistemic logic, ambiguity
JEL classification: D81, D82, D83

*The department of Economics, Bar Ilan University. email: Shiri.Alon-Eron@biu.ac.il
†The Economics and Management department, the Open University of Israel. email: aviadhe@openu.ac.il

1 Introduction

Classical games of incomplete information are modeled by Harsanyi (1967-68) type spaces, where the belief of each type is a probability measure on nature and the players' types. Such beliefs are precise and complete, and cannot model Knightian uncertainty, in the form of ambiguity or incompleteness. To express the latter, a recent and growing literature on games with uncertainty generalizes Harsanyi type spaces to *Knightian* type spaces, where the belief of each type is expressed by a *set* of probability measures on nature and the players' types. A type then deems one event more likely than another if *all* of the type's probability measures assign a higher probability to the first event. In particular, some types may be unable to compare the likelihoods of some pairs of events, in which case these types' likelihood relation over events is incomplete. Ahn (2007) constructed the universal space in the class of Knightian type spaces, generalizing the classical Mertens-Zamir (1985) construction of the universal space for Harsanyi type spaces.[1]

Partial likelihood assessments allow for a variety of ways to eventually determine the player's choice of a strategy – the act which defines, given the other players' strategies, the mapping from states of the world to outcomes. The type's strategy choice could be determined by the Gilboa and Schmeidler (1989) Maxmin criterion across the type's set of probabilities, as in Kajii and Ui (2005, 2009) and Ganguli (2007) who study agreement and no-trade theorems, Ui (2009) who studies global games, Bose et al. (2006), Bose and Renou (2011) and Bodoh-Creed (2012) who study auctions and mechanism design, and de Castro, Pesce and Yannalis (2011) and de Castro and Chateauneuf (2011), who study competitive economies with ambiguity. Alternatively, the type's choice could be determined by the Bewley (2002) unanimity of probabilities criterion (and inertia in case the preferences over strategies are incomplete), as in Lopomo et al. (2009, 2011) who study mechanism design and moral hazard, and Stauber (2011) who studies equilibrium robustness. [2]

A state of the world in any kind of type space is aimed at describing all aspects of the situation relevant for the strategic interaction – the physical, 'natural' aspects like the monetary payoffs, as well as the state of mind of each player, i.e. the player's beliefs about nature and about the state of mind of the other players. In a Knightian type space, the probability measures associated with a type should therefore faithfully

[1]This construction was suggested by Armbruster and Boge (1979) and Boge and Eisele (1979), and Mertens and Zamir (1985) proved its universality, i.e. that it embeds any other type space by a unique belief morphism.

[2]The literature on Knightian games still awaits the applications of other multiple-prior formulations of ambiguity which were offered in the literature, such as the smooth model of Klibanoff et al. (2005) or the Maccheroni et al. (2006) variational preferences rendering of Hansen and Sargent's (2001) multiplier preferences – see the extensive survey by Gilboa and Marinacci (2011).

and succinctly represent the list of assertions which describe its state of mind, e.g. 'I believe it's likely to rain'; 'I can't tell whether it's more likely that Bob thinks it's likely to rain, or rather that it's more likely he thinks it's unlikely to rain'; etc. Obviously, these assertions should be consistent. But what exactly are the properties that these assertions should satisfy if they are to be represented by a set of probability measures of some type in a Knightian type space?

The way to address this question is to phrase these assertions as formulas in a formal syntax, and to look for an axiom system for this syntax which will be sound and complete with respect to the class of Knightian type spaces. Soundness means that every formula which can be proved from the axioms indeed obtains in each state of every Knightian type space; completeness means that every formula which obtains in each state of every Knightian type space can be proved from the axioms.

Our main contribution in this paper is the formulation of such a syntax, along with a sound and complete axiomatization for Knightian type spaces. Formulas in the syntax are built up inductively from atomic formulas using the propositional connectives of negation and conjunction, as well as, for each player, the operator 'at least as likely as' connecting pairs of formulas. The axiom system augments the standard axiomatization of propositional logic with axioms and inference rules regarding the properties of the 'at least as likely as' operator. The axioms employed are reflexivity, positivity and non-triviality, as well as a more involved axiom, which is a generalized version of a cancellation-type condition. The latter asserts that whenever there are two provably equivalent ways to represent the same disjunction of conjunctions, likelihood assessments of formulas across these representations are not contradictory in the aggregate. Put intuitively, if a player is certain that two finite sequences of formulas are 'equivalent', then these should be considered equally likely. Two sequences of formulas are considered equivalent whenever exactly m of the formulas from one sequence hold true, if, and only if, exactly m of the formulas from the other sequence hold true, for every number of formulas m. A similar axiom featured also in Heifetz and Mongin's (2001) sound and complete axiomatization of standard Harsanyi type spaces, where it was vital for securing the existence of a probability measure congruent with the formulas, based on a theorem of the alternative; indeed, an appropriate probability measure might fail to exist in the absence of such a cancellation property, as shown by Kraft et al. (1959). A *generalized* duality argument was used in Alon and Lehrer (2011) for characterizing incomplete likelihood relations by a *set* of probability measures, and we rely on their result here.

Our completeness proof relies on the method of finite filtrations of the syntax, restricting attention each time to the sub-language generated by the finite set of atomic formulas which appear in the formula of interest, and whose depth is equal or smaller than that of that formula. Indeed, example 1 shows that in the full language there are maximally consistent sets of formulas which cannot be represented by a set of probability measures in Knightian type space, and example 2 shows how this may happen even if the language has only one atomic formula. This phenomenon – the lack of *strong* completeness – is well known already from the axiomatization for standard Harsanyi type spaces (Heifetz and Mongin 2001); it stems from the fact that the syntax can express only finitary conjunctions and disjunctions, while the probability measures take values in the *complete* field of real numbers. Strong completeness for Harsanyi type spaces can be retained with an infinitary logic (Meier 2012, Zhou 2009). It remains open whether an analogous, strongly complete infinitary axiomatization can be devised also for Knightian type spaces.

The paper is organized as follows. Basic semantic and syntactic definitions are presented in the next section. The axiom system is presented in section 3. Section 4 states the characterization result, and elaborates the counterexamples to strong completeness. Section 5 discusses introspection. All proofs appear in the last section.

2 Setup

The paper is concerned with Knightian type spaces *à*-la Harsanyi (1967-68). A Knightian type space is a tuple $\tau = (I, \Omega, \mathcal{A}, S, \mathcal{B}, T_0, (T_i)_{i \in I})$, where I is a nonempty set of individuals, Ω is the set of states of the world endowed with the sigma algebra $\mathcal{A}$, S is the set of states of nature endowed with the sigma algebra $\mathcal{B}$, and $T_0 : \Omega \to S$ is a measurable mapping which determines which state of nature obtains in each state of the world.

Denote by $\Delta_\sigma(\Omega, \mathcal{A})$ the set of all countably additive probability measures over $(\Omega, \mathcal{A})$. For each player i, $T_i : \Omega \to 2^{\Delta_\sigma(\Omega, \mathcal{A})}$ is a measurable mapping which associates each state of the world with a *set* of countably additive probabilities over $(\Omega, \mathcal{A})$, where $2^{\Delta_\sigma(\Omega, \mathcal{A})}$ is endowed with the sigma algebra generated by the sets

$$\{\mathcal{C} \subseteq \Delta_\sigma(\Omega, \mathcal{A}) \mid \mu(E) \geq \mu(F) \ \forall \mu \in \mathcal{C}\}$$

for events $E, F \in \mathcal{A}$.

Note that T_i being measurable w.r.t. the above sigma algebra implies that for any two events E and F from $\mathcal{A}$, the set

$$\{\omega \in \Omega \mid \mu(E) \geq \mu(F) \ \forall \mu \in T_i(\omega)\}$$

is also an event (it is the inverse image under T_i of all subsets of probabilities which rank E above F).

The aim of this paper is to formulate a syntax and an axiom system which characterizes Knightian type spaces. To this end, let $\mathcal{L}$ be the language built from a set of atomic formulas $\mathcal{P}$. Formulas in $\mathcal{L}$ are constructed inductively using the operators $\neg$

(‘not’) and $\wedge$ (‘and’). For two formulas φ and ψ, The operator $\varphi \vee \psi$ is shorthand for $\neg(\neg\varphi \wedge \neg\psi)$, $\varphi \to \psi$ is shorthand for $\neg\varphi \vee \psi$ and $\leftrightarrow$ is shorthand for $(\varphi \to \psi) \wedge (\psi \to \varphi)$. ‘True’ is denoted as $\top$, and equals $\varphi \vee \neg\varphi$, and ‘False’ is denoted by $\perp$, formally defined as $\varphi \wedge \neg\varphi$ (for any formula φ).

A binary operator $\succsim_i$ over formulas is defined for each individual $i \in I$, where $\varphi \succsim_i \psi$ is interpreted as ‘individual i judges φ to be at least as likely as ψ’. The relation $\sim_i$ is the symmetric part of $\succsim_i$, so that $\varphi \sim_i \psi$ if and only if $(\varphi \succsim_i \psi) \wedge (\psi \succsim_i \varphi)$. The propositional connectives are regarded as having precedence over the operators $\succsim$ and $\sim$.

Given a formula φ, its depth, $dp(\varphi)$, is defined inductively by:

(1) $dp(\varphi) = 0$ if φ is atomic.

(2) $dp(\neg\varphi) = dp(\varphi)$.

(3) $dp(\varphi \wedge \psi) = \max(dp(\varphi), dp(\psi))$.

(4) $dp(\varphi \succsim_i \psi) = \max(dp(\varphi), dp(\psi)) + 1$.

A typical model is a Knightian type space $M = (I, \Omega, \mathcal{A}, 2^{\mathcal{P}}, \mathcal{B}, T_0, (T_i)_{i \in I})$, where the states of nature are all subsets of atomic formulas, $\mathcal{B}$ is the product algebra of $2^{\mathcal{P}}$, and the mapping T_0 is an interpretation mapping, specifying which atomic formulas belong to each state of the world. Whenever $p \in T_0(\omega)$ for a state $\omega \in \Omega$ in a model M we say that p obtains, or holds true in ω, and denote it by $M, \omega \models p$. The relation $\models$ is then extended inductively to all the formulas according to the rules:

$$\begin{aligned} M, \omega &\models \neg\varphi \quad \text{if} \quad M, \omega \not\models \varphi \\ M, \omega &\models (\psi \wedge \varphi) \quad \text{if} \quad M, \omega \models \varphi \text{ and } M, \omega \models \psi \\ M, \omega &\models \varphi \succsim_i \psi \quad \text{if} \quad \mu([\varphi]) \geq \mu([\psi]) \ \forall \mu \in T_i(\omega) \end{aligned}$$

where

$$[\varphi] = \{\omega \in \Omega \mid M, \omega \models \varphi\} .$$

The class of all models is denoted $\mathcal{M}$.

For a model M and a formula φ, we write $M \models \varphi$ whenever $M, \omega \models \varphi$ for every $\omega \in \Omega$. We write $\mathcal{M} \models \varphi$ whenever $M \models \varphi$ for every model $M \in \mathcal{M}$. For a set of formulas Γ, $\Gamma \models \varphi$ denotes that for any state ω of any model $M \in \mathcal{M}$, φ obtains in ω whenever all the formulas in Γ obtain in ω.

3 The axiom system

The axiom system suggested in this paper is an incomplete variant of Gärdenfors (1975), which in turn builds on a construction due to Segerberg (1971). Gärdenfors formulated a sound and complete epistemic-logic axiomatization of additive probability measures, over a binary 'at least as likely as' relation. One of his axioms, which is discarded here, is completeness: for an 'at least as likely as' relation $\succsim$, every pair of formulas φ and ψ satisfy either $\varphi \succsim \psi$ or $\psi \succsim \varphi$. This axiom requires that each pair of formulas be comparable. Where Knightian type spaces are concerned, the belief of an individual is represented by a *set* of probability measures, therefore cannot be expected to be complete (unless the set is a singleton). Completeness (axiom A2 of Gärdenfors) is thus relaxed here, and replaced by the weaker assumption of reflexivity. Reflexivity (our A2) simply asserts that the binary relation we consider is a weak relation rather than a strict one.

Other than completeness, the axioms employed here are similar to those in Gärdenfors, consisting of a positivity assumption (A1 here and in Gärdenfors), a non-triviality assumption stating that 'False' cannot be considered at least as likely as 'True' (A3 here and in Gärdenfors), and a cancellation-type condition. The cancellation condition employed here is stronger than that of Gärdenfors (see a comparison of the two after the presentation of our cancellation axiom).

Both here and in Gärdenfors a basic axiomatization of propositional calculus is assumed, as well as an additional inference rule, the Rule of Necessitation. The first axiom is a propositional calculus assumption (for further details on such axiomatization see, for example, Chellas (1980)).

(PropCalc) An axiomatization of propositional calculus, including the inference rule Modus Ponens.

Four additional axiom schemata are introduced next. The axioms express certain regularity conditions on the relation $\succsim_i$, which will be shown to be necessary and sufficient for a representation of the relation $\succsim_i$ by a consensus rule over a set of prior probabilities.

(A1) For all formulas φ and individuals $i \in I$, $\varphi \succsim_i \bot$.

(A2) For all formulas φ and individuals $i \in I$, $\varphi \succsim_i \varphi$.

(A3) For all individuals i, $\neg(\bot \succsim_i \top)$.

In order to introduce the last axiom, two definitions are required.

Definition 1. *Let* $(\varphi_1, \ldots, \varphi_n)$ *be a finite sequence of formulas. For* $m \le n$*,* $\varphi^{(m)}$ *denotes the formula,*

$$\bigvee_{1 \le \ell_1 < \ldots < \ell_m \le n} (\varphi_{\ell_1} \wedge \ldots \wedge \varphi_{\ell_m}) \ .$$

In words, $\varphi^{(m)}$ 'goes over' all possible choices of m formulas from the sequence, and calculates the conjunction of each such set of chosen formulas. All these conjunctions are then combined by disjunction. Hence $\varphi^{(m)}$ holds true whenever there are at least m formulas from the sequence that hold true. In particular, $\varphi^{(1)}$ denotes the disjunction of all the formulas in the sequence and $\varphi^{(n)}$ is the conjunction of all the formulas in the sequence.

Definition 2. *For two sequences of formulas,* $(\varphi_1, \ldots, \varphi_n)$ *and* $(\psi_1, \ldots, \psi_n)$*, the notation*

$$(\varphi_1, \ldots, \varphi_n) \leftrightarrow (\psi_1, \ldots, \psi_n)$$

stands for the formula

$$\bigwedge_{m=1}^{n} (\varphi^{(m)} \leftrightarrow \psi^{(m)}) \ .$$

For two sequences $(\varphi_1, \ldots, \varphi_n)$ and $(\psi_1, \ldots, \psi_n)$, if $\varphi^{(m)} \leftrightarrow \psi^{(m)}$, then there are at least m formulas in the φ-sequence that hold true if and only if there are at least m formulas in the ψ-sequence that hold true. Specifically, $\varphi^{(1)} \leftrightarrow \psi^{(1)}$ denotes that the disjunction of all the formulas in the φ-sequence is equivalent to the disjunction of all the formulas in the ψ-sequence, whereas $\varphi^{(n)} \leftrightarrow \psi^{(n)}$ means that the conjunction of all the φ-formulas is equivalent to the conjunction of all the ψ-formulas.

By running over all possible values of m, $\bigwedge_{m=1}^{n}(\varphi^{(m)} \leftrightarrow \psi^{(m)})$ holds true, if and only if, for every $m = 1, \ldots, n$, *exactly* m of the φ-formulas hold true iff *exactly* m of the ψ-formulas hold true. The notation $(\varphi_1, \ldots, \varphi_n) \leftrightarrow (\psi_1, \ldots, \psi_n)$ may therefore be viewed as a generalization of equivalence of formulas to the case of sequences of formulas.[3]

Introducing a likelihood relation over formulas, if two (single) formulas φ and ψ are believed by an individual to be equivalent, then it is natural to require that the two formulas be deemed by the individual equally likely. Similarly, generalizing the equivalence relation from single formulas to sequences of formulas, it would make sense to require that if an individual is certain that two sequences of formulas are equivalent, one cannot be considered more likely than the other. This intuition is expressed in the axiom below, which is a cancellation-type condition, generalized to the incomplete case.

(A4') Let $(\varphi_1, \ldots, \varphi_n, \underbrace{\varphi', \ldots, \varphi'}_{\text{k times}})$ and $(\psi_1, \ldots, \psi_n, \underbrace{\psi' \ldots, \psi'}_{\text{k times}})$ be two sequences of formulas, and $i \in I$ an individual. Then

[3] Gärdenfors (1975) defines the same relation over formulas with an 'E' symbol, and suggests to read it as 'are generalized equivalent to'. The 'E' notation originates in Segerberg (1971).

$$\left(\left((\varphi_1,\ldots,\varphi_n,\underbrace{\varphi',\ldots,\varphi'}_{\text{k times}}) \leftrightarrow (\psi_1,\ldots,\psi_n,\underbrace{\psi'\ldots,\psi'}_{\text{k times}})\right) \sim_i \top\right) \wedge$$
$$(\varphi_1 \succsim_i \psi_1) \wedge (\varphi_2 \succsim_i \psi_2) \wedge \ldots \wedge (\varphi_n \succsim_i \psi_n) \longrightarrow (\psi' \succsim_i \varphi')$$

A4' formalizes the idea expressed above, stating that if two sequences, the φ-sequence and the ψ-sequence, are deemed by an individual to be equivalent, and the individual further judges every formula in the φ-sequence to be weakly more likely than every corresponding formula in the ψ-sequence, none of these "more likely" relations can be strict. To wit, if an individual deems all the φ-formulas except one as at least as likely as their corresponding ψ-formulas, then the weak likelihood judgement between the last formulas should hold, that is the last formulas should also be comparable, and specifically the likelihood judgement should be reversed. That is, the last ψ-formula should be considered at least as likely as the last φ-formula, thus balancing the likelihood of the two equivalent sequences.

Gärdenfors (1975) Applied a similar, somewhat weaker condition, in which in the sequences above the last formula repeats only once. We refer to this condition as A4. For complete relations our condition is implied by that of Gärdenfors: If, under the conditions of the axiom, $\neg(\psi' \succsim_i \varphi')$, then completeness necessitates that $\varphi' \succ_i \psi'$, delivering a contradiction of A4. But without completeness the two formulas can be incomparable, so that both $\neg(\psi' \succsim_i \varphi')$ and $\neg(\varphi' \succ_i \psi')$ hold. To exclude this possibility, A4 is strengthened to A4'.

To gain intuition when A4 is not sufficient for incomplete relations, consider the following example: let $\mathcal{P} = \{\varphi_1,\ldots,\varphi_6\}$, and $\mathcal{L}$ the implied language. Consider a relation $\succsim_i$ which satisfies $\varphi \succsim_i \bot$ and $\varphi \succsim \varphi$ for every φ in $\mathcal{L}$ (Positivity and Reflexivity), and other than that only the two relationships, $\varphi_2 \vee \varphi_3 \vee \varphi_5 \succsim_i \varphi_1 \vee \varphi_4 \vee \varphi_6$ and $\varphi_2 \vee \varphi_4 \vee \varphi_5 \succsim_i \varphi_1 \vee \varphi_3 \vee \varphi_6$. This relation satisfies also A4, since no sequences yield additional comparisons but the trivial ones (which are the result of Positivity and Reflexivity). However this relation cannot be represented by a consensus rule over a set of probability measures, since if there exists a set of probability measures that respects these relationships, then these must also entail, by summing over the implied inequalities, $\varphi_2 \vee \varphi_5 \succsim_i \varphi_1 \vee \varphi_6$ (which is not assumed to hold). Completing the relation with comparisons according to A4' delivers the last ranking by considering the sequences $(\varphi_2\vee\varphi_3\vee\varphi_5, \varphi_2\vee\varphi_4\vee\varphi_5, \varphi_1\vee\varphi_6, \varphi_1\vee\varphi_6)$ and $(\varphi_1\vee\varphi_4\vee\varphi_6, \varphi_1\vee\varphi_3\vee\varphi_6, \varphi_2\vee\varphi_5, \varphi_2\vee\varphi_5)$. Thus, A4 does not suffice to obtain a representation of the relation by a consensus rule over a set of probabilities, and A4', its strengthening, is required.

Next, a definition of a 'theorem' in our system is given. This is the following inductive definition:

(i) If φ is an axiom, then φ is a theorem.

(ii) If φ and $\varphi \to \psi$ are theorems, then ψ is a theorem. This is Modus Ponens, assumed as part of (A0).

(iii) If φ is a theorem, then, for every individual i, $\varphi \sim_i \top$ is a theorem.

To denote that φ is a theorem of our system we write $\vdash \varphi$. The rule in (iii) is named the *Rule of Necessitation*, applied also in Gärdenfors (1975):

(RN) If $\vdash \varphi$ then for every individual i, $\vdash \ (\varphi \sim_i \top)$.

This rule asserts that every individual is certain of any theorem of the system. It represents a choice made here, to exclude the possibility that the individual be unaware of formulas which are tautologies, even if those are very complicated. It thus imposes a form of unbounded rationality on the individuals in our framework.

Note that as a result of (RN), (A4') implies that the likelihood relation respects logical equivalences of single formulas, in that $\vdash \varphi \leftrightarrow \psi$ implies $\varphi \sim_i \psi$ for every individual i. This is achieved by examining the sequences (φ, ψ) and (φ, φ). These sequences satisfy $\vdash (\varphi, \psi) \leftrightarrow (\varphi, \varphi)$, and by Reflexivity (A2), $\varphi \succsim_i \varphi$, yielding, through (RN) and (A4'), both $\varphi \succsim_i \psi$ and $\psi \succsim_i \varphi$.

Definition 3. *The system (PropCalc),(A1)-(A3),(A4') and (RN) is denoted by* Σ.

The next remark completes the comparison of the system Σ with the previous literature.

Remark 1. Gärdenfors (1975) assumes in his system the following additional axiom (which he denotes by A0):

For any individual i and formulas $\varphi_1, \varphi_2, \psi_1, \psi_2$,

$$(\varphi_1 \leftrightarrow \varphi_2 \ \sim_i \top) \wedge (\psi_1 \leftrightarrow \psi_2 \ \sim_i \ \top) \ \to \ ((\varphi_1 \succsim_i \psi_1) \leftrightarrow (\varphi_2 \succsim_i \leftrightarrow \psi_2)) \ .$$

However, it turns out that assuming the rest of the axioms (here and in Gärdenfors' system as well), A0 is implied (a proof is given in the last section).

This section is concluded with a few definitions. For a set of formulas Γ, ψ is provable from Γ if:

(i) ψ is a formula in Γ.

(ii) ψ is a theorem of Σ.

(ii) φ and $\varphi \to \psi$ are provable from Γ (Modus Ponens).

A set of formulas Γ is consistent if one cannot prove from it a formula and its negation. Γ is said to be *maximally consistent* if: (1) it is consistent, and (2) For any formula φ, either φ or $\neg\varphi$ is in Γ.

4 The Characterization Theorem

The main result of this paper, stated in the next theorem, is that Σ is a sound and complete axiomatization of Knightian type spaces. In other words, every theorem of Σ holds true in every model in $\mathcal{M}$, and every formula that holds true in every model in $\mathcal{M}$ is a theorem of Σ.

Theorem 2. Σ *is a sound and complete axiomatization of* $\mathcal{M}$*, namely,*

$$\vdash \varphi \quad \Longleftrightarrow \quad \mathcal{M} \models \varphi \ .$$

As a corollary, using the completeness proof of the theorem, it is possible to conclude that Σ is *decidable*, namely that there is an effective finitary method for deciding for any formula whether or not it is a theorem of the system.[4]

Corollary 3. Σ *is decidable.*

All proofs appear in Section 6. While the soundness proof of the main theorem is straightforward, the proof of completeness, and the implied decidability, rely on two major components.

First, it relies on the method of formula-by-formula filtration. That is, when we want to prove that $\mathcal{M} \models \varphi$ implies $\vdash \varphi$, we consider the canonical collection Ω_φ of maximally consistent sets of formulas constructed only from atomic formulas appearing in φ, and whose depth does not exceed the depth of φ. The collection Ω_φ is finite, each of its subsets corresponds to some formula, and a formula belonging to all of the 'states' $\omega \in \Omega_\varphi$ is a theorem of Σ.

Second, we apply a result by Alon and Lehrer (2011) to show that the axioms imply that for every individual $i \in I$ and every state $\omega \in \Omega_\varphi$, there exists a set $T_i(\omega)$ of probabilities which represent the likelihood relation $\succsim_i$ as it appears in the formulas of ω. In the implied model, having Ω_φ as its state space, if φ holds true in all states of Ω_φ then $\vdash \varphi$, which yields the desired completeness result.

As the derivation of a set of probabilities is a central step in the proof, the representation result of Alon and Lehrer (for finite state spaces) is briefly stated below.

Let Ω be a finite state space endowed with the algebra of all its subsets, and $\succsim$ a binary relation over events in Ω. The following four axioms are necessary and sufficient for $\succsim$ to admit a representation through a consensus rule over a set of prior probabilities.

P1. For all events A, $A \succsim A$.

P2. For all events A, $A \succsim \emptyset$.

[4]For a more detailed discussion of decidability see Chellas (1980).

P3. $\neg(\emptyset \succsim \Omega)$.

P4. Let $(A_\ell)_{\ell=1}^n$ and $(B_\ell)_{\ell=1}^n$ be two sequences of events from Σ, and $k \in \mathbb{N}$ an integer. Then

$$\begin{aligned}
&\text{If} \quad \sum_{\ell=1}^{n-1} \mathbf{1}_{A_\ell}(\omega) + k\mathbf{1}_{A_n}(\omega) = \sum_{\ell=1}^{n-1} \mathbf{1}_{B_\ell}(\omega) + k\mathbf{1}_{B_n}(\omega) \quad \text{for all } \omega \in \Omega, \\
&\text{and} \quad A_\ell \succsim B_\ell \quad \text{for } \ell = 1, \ldots, n-1 \ , \\
&\text{then} \quad B_n \succsim A_n \quad .
\end{aligned}$$

The representation theorem of Alon and Lehrer (2011) states that the relation $\succsim$ satisfies the four axioms above if and only if there exists a nonempty set $\mathcal{C}$ of additive probability measures over events in Ω, such that for every two events A and B,

$$A \succsim B \quad \Leftrightarrow \quad \mu(A) \geq \mu(B) \ \text{ for every } \mu \in \mathcal{C} \ .$$

The proof employs a generalized separating hyperplane argument.

This method of proof generalizes the one used by Heifetz and Mongin (2001), who also relied on the formula-by-formula filtration technique and on a separation argument (a theorem of the alternative) to prove the completeness of their probability logic. Their axiom system contained an inference rule analogous to A4.

In both cases, *strong completeness* does *not* obtain, i.e. it is not the case that for *every* set of formulas Γ (in particular for infinite Γ), $\Gamma \models \varphi$ implies $\Gamma \vdash \varphi$. The reason is that probabilities obtain values in the $[0,1]$ interval of the complete field of real numbers, which contains the limits of all of its Cauchy sequences, while limits of such infinite sequences cannot be expressed by a finitary logic. The lack of strong completeness in our case is demonstrated in the following example, by constructing an infinite set of formulas Γ and a formula φ, such that $\Gamma \models \varphi$, while φ cannot be proved from Γ.

Example 1. Let Ω be any space of states of the world of some model in $\mathcal{M}$. For any $n \geq 1$, consider the atomic formulas φ_0 and $\varphi_1, \ldots, \varphi_n$, and define, for $L \subseteq \{1, \ldots, n\}$, the formula

$$\varphi_L = \left(\bigwedge_{\ell \in L} \varphi_\ell \right) \wedge \left(\bigwedge_{\ell \notin L} \neg(\varphi_\ell) \right) \ .$$

From the definition of φ_L it follows that

$$\vdash \left(\bigvee_{L \subseteq \{1,\ldots,n\}} \varphi_L \right) \leftrightarrow \top$$

and if L and K are two distinct subsets of indices in $\{1, \ldots, n\}$ then

$$\vdash (\varphi_L \wedge \varphi_K) \leftrightarrow \bot$$

consequently, the corresponding events $[\varphi_L] \subseteq \Omega$ satisfy $\bigcup_{L\subseteq\{1,\ldots,n\}}[\varphi_L] = \Omega$, and $\varphi_L \cap \varphi_K = \emptyset$ for every $L \neq K$. In other words, $\{[\varphi_L]\}_{L\subseteq\{1,\ldots,n\}}$ is a partition of Ω.

For an individual i, Consider the following finite subset Γ_n of formulas:

$$\begin{aligned} &\varphi_L \sim_i \varphi_K, \quad \text{for all } L, K \subseteq \{1,\ldots,n\}\ , \\ &\varphi_0 \succsim_i (\varphi_1 \wedge \neg(\varphi_{\{1,\ldots,n\}}))\ , \end{aligned} \tag{1}$$

and let

$$\Gamma = \bigcup_{n=1}^{\infty} \Gamma_n\ .$$

Then in each state ω in which Γ obtains,

$$\begin{aligned} & \mu([\varphi_L]) = \mu([\varphi_K]) \ \text{ for all } L, K \subseteq \{1,\ldots,n\} \text{ and all } \mu \in T_i(\omega) \\ \Rightarrow \quad & \mu([\varphi_L]) = \frac{1}{2^n} \ \text{ for all } L \subseteq \{1,\ldots,n\} \text{ and all } \mu \in T_i(\omega)\ , \\ & \text{and, as } [\varphi_{\{1,\ldots,n\}}] \subseteq [\varphi_1]\ , \\ & \mu([\varphi_0]) \geq \mu([\varphi_1]) - \mu([\varphi_{\{1,\ldots,n\}}]) = \mu([\varphi_1]) - \frac{1}{2^n} \ \text{ for all } \mu \in T_i(\omega)\ . \end{aligned}$$

Therefore, in each state ω in which Γ obtains, it must be that $\mu([\varphi_0]) \geq \mu([\varphi_1])$ for all $\mu \in T_i(\omega)$, which means that $\Gamma \models (\varphi_0 \succsim_i \varphi_1)$. Still, for any *finite* n, the formula $\neg(\varphi_0 \succsim_i \varphi_1)$ is consistent with the (finite) subset of formulas Γ_n. To see this, observe that a model for these formulas should satisfy:

$$\begin{aligned} & \mu([\varphi_0]) \geq \mu([\varphi_1]) - \frac{1}{2^n} \ \text{ for all } \mu \in T_i(\omega)\ , \\ & \lambda([\varphi_1]) > \lambda([\varphi_0]) \ \text{ for some } \lambda \in T_i(\omega)\ , \end{aligned}$$

and this may easily be accommodated by letting some $\lambda \in \mathcal{P}$ satisfy $\lambda([\varphi_1]) > \lambda([\varphi_0]) \geq \lambda([\varphi_1]) - \frac{1}{2^n}$ for the finite n in question. Consequently, $\varphi_0 \succsim_i \varphi_1$ cannot be proved from Γ in a *finite* number of steps.

The above example makes use of a countable number of atomic formulas. The next example shows that the same contradiction to strong completeness may be obtained even with a single atomic formula.

Example 2. Let ψ_0 be the only atomic formula and $I = \{1, 2\}$ the set of individuals. Define the following formulas, for individuals $i = 1, 2$ (so for individual i, $3 - i$ is the other individual):

$$\begin{aligned} \psi_1^i &= (\psi_0 \succsim_i \neg\psi_0) \vee (\neg\psi_0 \succsim_i \psi_0) \\ \psi_2^i &= (\psi_1^{3-i} \succsim_i \neg\psi_1^{3-i}) \vee (\neg\psi_1^{3-i} \succsim_i \psi_1^{3-i}) \\ &\ \ \cdots \\ \psi_n^i &= (\psi_{n-1}^{3-i} \succsim_i \neg\psi_{n-1}^{3-i}) \vee (\neg\psi_{n-1}^{3-i} \succsim_i \psi_{n-1}^{3-i}) \end{aligned}$$

In words, ψ_1^1 is the formula saying that individual 1 can compare the likelihoods of ψ_0 and its negation, ψ_2^1 says that 1 can determine which is more likely, that 2 can compare the likelihoods of ψ_0 and its negation or not, and so on. In order to build Example 1 out of these formulas it suffices to show the consistency of any subset of the formulas $\psi_0, \psi_1^i, \ldots, \psi_n^i$ together with the negation of the remaining formulas, so that an appropriate partition, like $\{\varphi_L\}_L$ of Example 1, can be constructed from them. In order to show this consistency consider the following model:

$$\begin{aligned}
&I = \{1, 2\} \\
&\Omega = \{(s, a_1 \cdots a_n, b_1 \cdots b_n) |\ s, a_\ell, b_\ell \in \{0, 1\},\ \ell = 1, \ldots, n\}\ ,
\end{aligned}$$

$\mathcal{A}$ is the algebra of all subsets of Ω. T_0 is defined by

$$\psi_0 \in T_0\left((s, a_1 \cdots a_n, b_1 \cdots b_n)\right) \text{ if and only if } s = 1$$

so s denotes whether ψ_0 or $\neg\psi_0$ obtains in $\omega = (s, a_1 \cdots a_n, b_1 \cdots b_n)$.
Next, we define the mappings T_i, $i = 1, 2$. To this end, we first define the auxiliary mappings T_i^ℓ, $\ell = 1, \ldots, n$, as follows. If $a_1 = 1$, then

$$T_1^1\left((s, a_1 \cdots a_n, b_1 \cdots b_n)\right) = \{\delta_s \times \delta_{a_1}\}$$

(where δ_x denotes the Dirac probability concentrated on x), signifying that individual 1 knows the actual state of nature s and his/her own first coordinate $a_1 = 1$. Otherwise, if $a_1 = 0$, then

$$T_1^1\left((s, a_1 \cdots a_n, b_1 \cdots b_n)\right) = \{\delta_{s=0} \times \delta_{a_1}, \delta_{s=1} \times \delta_{a_1}\}$$

signifying that individual 1 is totally ignorant which state of nature is more likely, but still knows that his/her own first coordinate is $a_1 = 0$. Similarly, if $b_1 = 1$ then

$$T_2^1\left((s, a_1 \cdots a_n, b_1 \cdots b_n)\right) = \{\delta_s \times \delta_{b_1}\}$$

while if $b_1 = 0$ then

$$T_2^1\left((s, a_1 \cdots a_n, b_1 \cdots b_n)\right) = \{\delta_{s=0} \times \delta_{b_1}, \delta_{s=1} \times \delta_{b_1}\}$$

Inductively, suppose T_i^ℓ has already been defined for $\ell < k$ and $i = 1, 2$. If $a_k = 1$ then

$$T_1^k\left((s, a_1 \cdots a_n, b_1 \cdots b_n)\right) = \left\{\delta_{b_{k-1}} \times \delta_{a_k}\right\}$$

signifying that individual 1 knows the actual value of b_{k-1} and his/her own k-th coordinate $a_k = 1$. Otherwise, if $a_k = 0$, then

$$T_1^k\left((s, a_1 \cdots a_n, b_1 \cdots b_n)\right) = \left\{\delta_{b_{k-1}=0} \times \delta_{a_k}, \delta_{b_{k-1}=1} \times \delta_{a_k}\right\}$$

signifying that individual 1 is totally ignorant about his/her opponent's $k-1$ coordinate, but still knows that his/her own k-th coordinate is $a_k = 0$. Similarly, if $b_k = 1$ then

$$T_2^k\left((s, a_1 \cdots a_n, b_1 \cdots b_n)\right) = \{\delta_s \times \delta_{b_k}\}$$

while if $b_k = 0$ then

$$T_2^k\left((s, a_1 \cdots a_n, b_1 \cdots b_n)\right) = \left\{\delta_{a_{k-1}=0} \times \delta_{b_k}, \delta_{a_{k-1}=1} \times \delta_{b_k}\right\}$$

Finally, consider each vector in

$$\prod_{\ell=1}^{n} T_i^\ell\left((s, a_1 \cdots a_n, b_1 \cdots b_n)\right)$$

and consider the product probability obtained by taking the product of the members of this vector (this is a product of Dirac measures). Define $T_i\left((s, a_1 \cdots a_n, b_1 \cdots b_n)\right)$ to be the set of all these product probabilities.

The result of the above definitions is that

$$[\psi_0] = \{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid s = 1\}$$

and in a state $(s, a_1 \cdots a_n, b_1 \cdots b_n)$, the value of a_ℓ determines whether the formula ψ_ℓ^1 obtains, and the value of b_ℓ determines whether the formula ψ_ℓ^2 obtains. This claim is summarized in the following lemma.

Lemma 4. *For* $\ell = 1, \ldots, n$

$$[\psi_\ell^1] = \{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid a_\ell = 1\}$$
$$[\psi_\ell^2] = \{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid b_\ell = 1\}$$

The proof appears in the appendix.

In order to reproduce Example 1, consider the formulas ψ_ℓ^1, $\ell = 1, \ldots, n$. According to the lemma above, for every $L \subseteq \{1, \ldots, n\}$

$$[\psi_L^1] \equiv \left(\bigcap_{\ell \in L} [\psi_\ell^1]\right) \cap \left(\bigcap_{\ell \notin L} \neg[\psi_\ell^1]\right)$$

is nonempty since

$$[\psi_L^1] = \{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid a_\ell = 1 \text{ if and only if } \ell \in L\}$$

Now, augment Ω with an additional state $\tilde{\omega}$, let $T_1(\tilde{\omega})$ be an arbitrary set of probabilities over $\Omega \cup \{\tilde{\omega}\}$, and let $T_2(\tilde{\omega})$ be some singleton probability satisfying

$$\begin{aligned} T_2(\tilde{\omega})[\psi_L^1] &= \frac{1}{2^n}, \quad L \subseteq \{1, \ldots, n\} \\ T_2(\tilde{\omega})[\psi_0] &= \frac{1}{2} - \frac{1}{2^{n+1}} \end{aligned}$$

Then

$$T_2(\tilde{\omega})[\psi_1^1] = \frac{1}{2}$$

so the formula $\neg(\psi_0 \succsim_2 \psi_1^1)$ obtains in $\tilde{\omega}$, and the following formulas Γ_n

$$\begin{aligned} &\psi_L^1 \sim_2 \psi_K^1, \quad \text{for all } L, K \subseteq \{1, \ldots, n\} , \\ &\psi_0 \succsim_2 (\psi_1^1 \wedge \neg(\psi_{\{1,\ldots,n\}}^1)) . \end{aligned}$$

obtain in $\tilde{\omega}$ as well. By the soundness part of theorem 2, $\Gamma_n \cup \{\neg(\psi_0 \succsim_2 \psi_1^1)\}$ is consistent. This means that $(\psi_0 \succsim_2 \psi_1^1)$ cannot be proved in a finite number of steps from

$$\Gamma = \bigcup_{n=1}^{\infty} \Gamma_n$$

despite of the fact that $\Gamma \models (\psi_0 \succsim_2 \psi_1^1)$.

5 Introspection

In this section, the beliefs of an individual regarding his own beliefs are considered. Two axioms of introspection are suggested. While the standard introspection assumptions of epistemic logic apply to certainty judgements, $\varphi \sim_i \top$ in our notation, these are extended here to any judgement of the form $\varphi \succsim_i \psi$.

The first introspection assumption states that if an individual considers one formula to be at least as likely as another, then the individual is certain of this judgement.

(I1) For any individual i and formulas φ and ψ, if $\varphi \succsim_i \psi$ then $(\varphi \succsim_i \psi) \sim_i \top$.

The second introspection assumption asserts that if an individual cannot determine one formula to be at least as likely as another, then the individual is certain of this absence of ranking.

(I2) For any individual i and formulas φ and ψ, if $\neg(\varphi \succsim_i \psi)$ then $\neg(\varphi \succsim_i \psi) \sim_i \top$.

Put together, (I1) and (I2) yield that the individual is certain of his or her beliefs. Note that if the relation is complete, then the second axiom implies that the individual is certain of any *strict* relationship he or she establishes, while the first axiom concerns weak relationships. However, under incompleteness lack of ranking is not equivalent to strict ranking in the opposite direction, thus the two introspection axioms are conceptually different.

In order to characterize the beliefs of an individual in a Knightian type space under the introspection assumptions, a definition is required.

Consider the following equivalence relation over sets of probability measures over $(\Omega, \mathcal{A})$:

Definition 4. *A set of probabilities $\mathcal{C}$ is equivalent to a set $\mathcal{D}$, denoted $\mathcal{C} \sim \mathcal{D}$, if for every two events $A, B \in \mathcal{A}$,*

$$\begin{aligned} &\textit{there exists } \mu \in \mathcal{C} \textit{ such that } \mu(A) > \mu(B) \ \Leftrightarrow \\ &\textit{there exists } \mu' \in \mathcal{D} \textit{ such that } \mu'(A) > \mu'(B) \ . \end{aligned}$$

The definition identifies when two sets generate the same likelihood relation over events through a consensus rule. It is required since two *different* sets of probabilities may generate the same likelihood relation over events under a consensus rule (see Alon and Lehrer (2013) for more details). We use the above definition to denote the set of states in which an individual has identical beliefs.

Definition 5. *For the set of probabilities $T_i(\omega)$ of individual i in state ω,*

$$[T_i(\omega)] = \{\omega' \in \Omega \mid T_i(\omega') \sim T_i(\omega)\}.$$

The set $[T_i(\omega)]$ is not necessarily an event. It will be an event whenever the sigma-algebra $\mathcal{A}$ is separable, thus a probability measure on $\mathcal{A}$ can be uniquely identified using only a countable number of events. But generally the probability of $[T_i(\omega)]$ may not be defined. On that account, in order to characterize type spaces in which the introspection axioms (I1) and (I2) are satisfied, events that contain $[T_i(\omega)]$ are employed.

Proposition 5. *The axiom system Σ augmented by the introspection axioms (I1) and (I2) is a sound and complete axiomatization of type spaces in which for every individual i and every state ω, if E is an event such that $T_i(\omega) \subseteq E$, then $\mu(E) = 1$ for every $\mu \in T_i(\omega)$.*

6 Proofs

Let $\mathbf{1}_A$ denote the indicator function of event A, for any $A \in \mathcal{A}$.

6.1 Proof of Remark 1

Claim 6. *Under the axiom system Σ, for any individual i and formula φ, $\top \succsim_i \varphi$.*

Proof. Consider the two sequences $(\neg\varphi, \varphi)$ and $(\bot, \top)$. It holds that $\vdash (\neg\varphi \vee \varphi) \leftrightarrow \top$ and $\vdash (\neg\varphi \wedge \varphi) \leftrightarrow \bot$, therefore $\vdash (\neg\varphi, \varphi) \leftrightarrow (\bot, \top)$, and by (RN),

$(\neg\varphi, \varphi) \leftrightarrow (\bot, \top) \sim_i \top$. A1 implies that $\neg\varphi \precsim_i \bot$, therefore by A4' (and also A4),$\top \precsim_i \varphi$. ■

Claim 7. *The axiom system Σ implies transitivity: for any three formulas φ, ψ and χ, if $\varphi \precsim_i \psi$ and $\psi \precsim_i \chi$ then $\varphi \precsim_i \chi$.*

Proof. Follows from applying A4' (or A4) to the sequences (φ, ψ, χ) and (ψ, χ, φ). ■

Claim 8. *If $\vdash (\varphi \to \psi)$ then $\psi \precsim_i \varphi$.*

Proof. Consider the sequences $(\psi \wedge \neg\varphi, \varphi)$ and $(\bot, \psi)$. First, $\vdash ((\psi \wedge \neg\varphi) \vee \varphi) \leftrightarrow (\psi \vee \varphi)$, and as it is assumed that $\vdash (\neg\varphi \vee \psi)$, then $\vdash (\psi \vee \varphi) \leftrightarrow (\psi \vee \varphi) \wedge (\psi \vee \neg\varphi)$, which in turn satisfies, $\vdash ((\psi \vee \varphi) \wedge (\psi \vee \neg\varphi)) \leftrightarrow \psi$, and finally $\vdash \psi \leftrightarrow (\psi \vee \bot)$. In addition, $\vdash ((\psi \wedge \neg\varphi) \wedge \varphi) \leftrightarrow \bot$, and $\vdash (\psi \wedge \bot) \leftrightarrow \bot$. It therefore holds that $\vdash (\psi \wedge \neg\varphi, \varphi) \leftrightarrow (\bot, \psi)$, and by (RN), $(\psi \wedge \neg\varphi, \varphi) \leftrightarrow (\bot, \psi) \sim_i \top$. As $\psi \wedge \neg\varphi \precsim_i \bot$, it follows from A4' (or A4) that $\psi \precsim_i \varphi$. ■

Claim 9. *If $\varphi \sim_i \top$ and $\psi \sim_i \top$ then $\varphi \wedge \psi \sim_i \top$.*

Proof. This claim is proved using the sequences (φ, ψ) and $(\varphi \vee \psi, \varphi \wedge \psi)$. Obviously the disjunction and conjunction of the elements in each sequence are the same. Employing transitivity and Claim 6, $\varphi \precsim_i \top \precsim_i \varphi \vee \psi$, therefore A4' (or A4) and transitivity yield $\varphi \wedge \psi \precsim_i \psi \precsim_i \top$. ■

To prove that A0 of Gandenfors is implied by our axiom system, suppose first that $\varphi_1 \leftrightarrow \varphi_2 \sim_i \top$ for a pair of formulas φ_1 and φ_2, and let ψ be some formula. Consider the sequences (φ_1, ψ) and (ψ, φ_2).

Observe that,

$$\begin{aligned}
&\vdash \quad ((\varphi_1 \vee \psi) \leftrightarrow (\varphi_2 \vee \psi)) \leftrightarrow [(\neg(\varphi_1 \vee \psi) \vee (\varphi_2 \vee \psi)) \wedge (\neg(\varphi_2 \vee \psi) \vee (\varphi_1 \vee \psi))] \\
&\vdash \quad [(\neg(\varphi_1 \vee \psi) \vee (\varphi_2 \vee \psi)) \wedge (\neg(\varphi_2 \vee \psi) \vee (\varphi_1 \vee \psi))] \leftrightarrow [(\neg\varphi_1 \vee \varphi_2 \vee \psi) \wedge (\neg\varphi_2 \vee \varphi_1 \vee \psi)] \\
&\vdash \quad [(\neg\varphi_1 \vee \varphi_2 \vee \psi) \wedge (\neg\varphi_2 \vee \varphi_1 \vee \psi)] \leftrightarrow ((\varphi_1 \leftrightarrow \varphi_2) \vee \psi) \ .
\end{aligned}$$

Therefore, since $\vdash (\varphi_1 \leftrightarrow \varphi_2) \to ((\varphi_1 \leftrightarrow \varphi_2) \vee \psi)$, and by the assumption $(\varphi_1 \leftrightarrow \varphi_2) \sim_i \top$, Claim 8 implies that also $(\varphi_1 \vee \psi) \leftrightarrow (\varphi_2 \vee \psi) \sim_i \top$.

In addition,

$$\begin{aligned}
&\vdash \quad ((\varphi_1 \wedge \psi) \leftrightarrow (\varphi_2 \wedge \psi)) \leftrightarrow [(\neg(\varphi_1 \wedge \psi) \vee (\varphi_2 \wedge \psi)) \wedge (\neg(\varphi_2 \wedge \psi) \vee (\varphi_1 \wedge \psi))] \\
&\vdash \quad [(\neg(\varphi_1 \wedge \psi) \vee (\varphi_2 \wedge \psi)) \wedge (\neg(\varphi_2 \wedge \psi) \vee (\varphi_1 \wedge \psi))] \leftrightarrow [(\neg\varphi_1 \vee \varphi_2 \vee \neg\psi) \wedge (\neg\varphi_2 \vee \varphi_1 \vee \neg\psi)] \\
&\vdash \quad [(\neg\varphi_1 \vee \varphi_2 \vee \neg\psi) \wedge (\neg\varphi_2 \vee \varphi_1 \vee \neg\psi)] \leftrightarrow ((\varphi_1 \leftrightarrow \varphi_2) \vee \neg\psi) \ .
\end{aligned}$$

Hence in the same manner as above, $(\varphi_1 \wedge \psi) \leftrightarrow (\varphi_2 \wedge \psi) \sim_i \top$. Claim 9 yields that $(\varphi_1, \psi) \leftrightarrow (\psi, \varphi_2) \sim_i \top$, hence by A4' (or simply A4), $(\varphi_1 \succsim_i \psi) \leftrightarrow (\varphi_2 \succsim_i \psi)$.

In the same manner it can be shown that if $\psi_1 \leftrightarrow \psi_2 \sim_i \top$ then $(\varphi_2 \succsim_i \psi_1) \leftrightarrow (\varphi_2 \succsim_i \psi_2)$, hence the conclusion of Gärdenfors' A0.

6.2 Proof of Theorem 2 – Soundness: $\vdash \psi \quad \Rightarrow \quad \mathcal{M} \models \psi$

Suppose that ψ is a theorem of Σ. It should be proved that ψ obtains in every state of any model in $\mathcal{M}$. Let $M = (I, \Omega, \mathcal{A}, 2^{\mathcal{P}}, \mathcal{B}, T_0, (T_i)_{i \in I})$ be an arbitrary model in $\mathcal{M}$. All axioms and rules of Σ should be shown to hold in M.

By definition of the interpretation function on formulae of the form $\psi \succsim_i \varphi$, in every state ω, Reflexivity, Positivity and Non Triviality are easily seen to hold.

For A4', let ω be a state, and suppose two sequences, $(\varphi_1, \dots, \varphi_n, \varphi', \dots, \varphi')$ and $(\psi_1, \dots, \psi_n, \psi' \dots, \psi')$, with the last formulas, φ' and ψ', repeating k times, so that $((\varphi_1, \dots, \varphi_n, \varphi', \dots, \varphi') \leftrightarrow (\psi_1, \dots, \psi_n, \psi' \dots, \psi')) \sim_i \top$ obtains in ω, as well as $\varphi_\ell \succsim_i \psi_\ell$, for every $\ell = 1, \dots, n$.

The first indifference implies that $\mu([((\varphi_1, \dots, \varphi_n, \varphi', \dots, \varphi') \leftrightarrow (\psi_1, \dots, \psi_n, \psi' \dots, \psi'))]) = 1$ for every $\mu \in T_i(\omega)$. By the definition of the relation $\models$, the event in question equals the set of states in which exactly m of the φ-formulas hold true if and only if exactly m of the ψ-formulas hold true. These are precisely the states ω for which

$$\sum_{\ell=1}^{n} \mathbf{1}_{[\varphi_\ell]}(\omega) + k\mathbf{1}_{[\varphi']}(\omega) = \sum_{\ell=1}^{n} \mathbf{1}_{[\psi_\ell]}(\omega) + k\mathbf{1}_{[\psi']}(\omega) \ . \tag{2}$$

Denote by Ω' the set of states for which (2) is satisfied. Let μ be a probability measure from $T_i(\omega)$. Then μ lends probability one to Ω'. By the assumption above, μ also satisfies $\mu([\varphi_\ell]) \geq \mu([\psi_\ell])$, and this inequality is maintained if we restrict attention only to states in Ω' (as all other states are of μ-probability zero). Taking expectation of (2) with respect to μ implies that $\mu([\psi']) \geq \mu([\varphi'])$ on Ω', and the same holds true if we add the states outside Ω' (again, as they add μ-probability zero). It is concluded that for every $\mu \in T_i(\omega)$, $\mu([\psi']) \geq \mu([\varphi'])$, hence $\psi' \succsim_i \varphi'$, as required.

The above shows that if $\vdash \psi$ then ψ obtains in every state of any model in $\mathcal{M}$. Thus, for any model, $[\psi] = \Omega$, implying that for every state ω, every individual i and every $\mu \in T_i(\omega)$, $\mu([\psi]) = 1$. Therefore, $\psi \sim_i \top$ obtains for every individual i, in every state of every model.

6.3 Proof of Theorem 2 – Completeness: $\mathcal{M} \models \psi \ \Rightarrow \vdash \psi$

Suppose that $\mathcal{M} \models \psi$, i.e., that ψ obtains in all states of all models in $\mathcal{M}$. Let $\mathcal{L}(\psi)$ be the restricted language which contains only the atomic formulas that appear in ψ, and formulas of depth no larger than the depth of ψ. Similarly define $\mathcal{L}(\psi)^{+1}$ to be the language which contains the same atomic formulas, and formulas of depth no larger than $dp(\psi)+1$. Denote by Ω all the maximally consistent subsets of $\mathcal{L}(\psi)$, and by Ω^{+1} all the maximally consistent subsets of $\mathcal{L}(\psi)^{+1}$. Typically, each subset $\omega \in \Omega$ is equivalent to a collection of subsets from Ω^{+1}.

By the finiteness of the languages $\mathcal{L}(\psi)$ and $\mathcal{L}(\psi)^{+1}$, the number of subsets in Ω and in Ω^{+1} is finite. Let $\mathcal{A} = 2^{\Omega}$ be the collection of all subsets of Ω, and let $\mathcal{A}^{+1} = 2^{\Omega^{+1}}$.

For $\varphi \in \mathcal{L}(\psi)^{+1}$, $|\varphi| = \{\omega \in \Omega^{+1} \mid \varphi \in \omega\}$. If φ is in $\mathcal{L}(\psi)$ then this may equivalently be written as $|\varphi| = \{\omega \in \Omega \mid \varphi \in \omega\}$, as φ either belongs to all state of Ω^{+1} which compose a state $\omega \in \Omega$, or does not belong to all of them. Two important properties of the sets $|\varphi|$ are summarized in the next lemma: [5]

Lemma 10. *(a) Every subset of Ω corresponds to a formula of $\mathcal{L}(\psi)$. That is, for every $E \in \mathcal{A}$ there is a formula $\varphi \in \mathcal{L}(\psi)$ such that $E = |\varphi|$.*

(b) For all $\varphi_1, \varphi_2 \in \mathcal{L}(\psi)^{+1}$, $|\varphi_1| \subseteq |\varphi_2|$ if and only if $\vdash \varphi_1 \to \varphi_2$.

Proof.

(a) Every state $\omega \in \Omega$ corresponds to a maximally consistent subset of $\mathcal{L}(\psi)$, denote it by $\{\varphi_1^{\omega}, \ldots, \varphi_{k_\omega}^{\omega}\}$. By the definition of maximally consistent subsets it follows that $\{\omega\} = |\varphi_1^{\omega} \wedge \ldots \wedge \varphi_{k_\omega}^{\omega}|$, hence $E = \bigcup_{\omega \in E} |\varphi_1^{\omega} \wedge \ldots \wedge \varphi_{k_\omega}^{\omega}| = |\bigvee_{\omega \in E} \left(\varphi_1^{\omega} \wedge \ldots \wedge \varphi_{k_\omega}^{\omega}\right)|$. The formula $\bigvee_{\omega \in E} \left(\varphi_1^{\omega} \wedge \ldots \wedge \varphi_{k_\omega}^{\omega}\right)$, its depth being no more than the maximal depth of the φ_ℓ^{ω}'s, is in $\mathcal{L}(\psi)$.

(b) First suppose that $\vdash \varphi_1 \to \varphi_2$, which means that φ_2 holds true whenever φ_1 does. As each state $\omega^{+1} \in \Omega^{+1}$ corresponds to a *maximally* consistent set of formulas of $\mathcal{L}(\psi)^{+1}$, it follows that if $\varphi_1 \in \omega$, then also $\varphi_2 \in \omega$.

For the other direction assume that $\nvdash \varphi_1 \to \varphi_2$, or in other words, that φ_1 and $\neg\varphi_2$ are consistent. Since Ω^{+1} is the set of *all* maximally consistent subsets of formulas of $\mathcal{L}(\psi)^{+1}$, there exists a state $\omega^{+1} \in \Omega^{+1}$ such that $\varphi_1 \in \omega^{+1}$ and $\neg\varphi_2 \in \omega^{+1}$, implying $\omega^{+1} \in |\varphi_1|$ but $\omega^{+1} \notin |\varphi_2|$. Therefore $|\varphi_1| \nsubseteq |\varphi_2|$.

■

Fix an individual $i \in I$ and a state $\omega^{+1} \in \Omega^{+1}$. Define a (partial) likelihood relation $\succsim_i^{\omega^{+1}}$ over events in $\mathcal{A}$ (i.e., over events that are composed of states from Ω) by $|\varphi| \succsim_i^{\omega^{+1}} |\varphi'|$ whenever $(\varphi \succsim_i \varphi') \in \omega^{+1}$. By the depth restriction on formulas in $\mathcal{L}(\psi)^{+1}$, φ is

[5] This lemma is analogous to Lemma A.1 from Heifetz and Mongin (2001).

necessarily in $\mathcal{L}(\psi)$. The relation is well defined over the algebra $\mathcal{A}$ according to part (a) of Lemma 10.

By definition of the maximally consistent subsets that constitute Ω^{+1}, for every $\varphi \in \mathcal{L}(\psi)$ and $i \in I$, $(\varphi \succsim_i \varphi) \in \omega^{+1}$ according to A1, $(\varphi \succsim_i \bot) \in \omega^{+1}$ by A2, and $\bot \succsim \top \notin \omega^{+1}$ by A3. Hence each $\succsim_i^{\omega^{+1}}$ satisfies P1, P2 and P3, respectively, from Alon and Lehrer (2013), over events in $\mathcal{A}$.

Furthermore, suppose that for two sequences of sets in $\mathcal{A}$, $(E_1, \ldots, E_n, E_{n+1}, \ldots, E_{n+k})$ and $(F_1, \ldots, F_n, F_{n+1}, \ldots, F_{n+k})$, it is the case that:

$$E_{n+1} = E_{n+\ell} \text{ and } F_{n+1} = F_{n+\ell} \quad \text{for } \ell = 1, \ldots, k \; , \tag{3}$$

$$\bigcup_{1 \leq \ell_1 < \ldots < \ell_m \leq n+k} (E_{\ell_1} \cap \ldots \cap E_{\ell_m}) = \bigcup_{1 \leq \ell_1 < \ldots < \ell_m \leq n+k} (F_{\ell_1} \cap \ldots \cap F_{\ell_m}) \; , \quad \text{for } m = 1, \ldots, n+k \tag{4}$$

$$\text{and } E_\ell \succsim_i^{\omega^{+1}} F_\ell \quad \text{for } \ell = 1, \ldots, n \; . \tag{5}$$

By part (a) of Lemma 10, every event E is in fact $|\varphi|$ for some $\varphi \in \mathcal{L}(\psi)$, so the sequences above translate to sequences $(\varphi) = (|\varphi_1|, \ldots, |\varphi_n|, \underbrace{|\varphi'|, \ldots, |\varphi'|}_{\text{k times}})$ and $(\psi) = (|\psi_1|, \ldots, |\psi_n|, \underbrace{|\psi'|, \ldots, |\psi'|}_{\text{k times}})$. Denote $\varphi_{n+1} = \ldots = \varphi_{n+k} = \varphi'$ and $\psi_{n+1} = \ldots = \psi_{n+k} = \psi'$. By consistency and maximality of the subsets of formulas which constitute Ω,

$$\bigcup_{1 \leq \ell_1 < \ldots < \ell_m \leq n+k} (|\varphi_{\ell_1}| \cap \ldots \cap |\varphi_{\ell_m}|) = \left| \bigvee_{1 \leq \ell_1 < \ldots < \ell_m \leq n} (\varphi_{\ell_1} \wedge \ldots \wedge \varphi_{\ell_m}) \right| = |\varphi^{(m)}| \; .$$

and the same for $\psi^{(m)}$. Therefore (4) implies that $|\varphi^{(m)}| = |\psi^{(m)}|$, which by (b) of Lemma 10 delivers $\vdash \varphi^{(m)} \leftrightarrow \psi^{(m)}$, yielding $\vdash (\varphi^{(m)} \leftrightarrow \psi^{(m)}) \sim_i \top$ by (RN). As this holds for every $m = 1, \ldots, n+k$, Claim 9 yields that $\vdash (\varphi) \leftrightarrow (\psi) \sim_i \top$. Under condition (5) and the definition of the relation $\succsim_i^{\omega^{+1}}$, maximality of the subset ω^{+1} and A4' entail that $(\psi_{n+1} \succsim_i \varphi_{n+1}) \in \omega^{+1}$, hence $|\psi_{n+1}| \succsim_i^{\omega^{+1}} |\varphi_{n+1}|$.

It is established that the relation $\succsim_i^{\omega^{+1}}$ satisfies assumptions P1-P4 of Alon and Lehrer (2013) over the finite state space Ω and the algebra $\mathcal{A}$. By their Theorem 1 (as sketched above in Section 4), for any state ω^{+1} and individual i there exists a set $\mathcal{C}_i^{\omega^{+1}}$ of probabilities over $(\Omega, \mathcal{A})$, such that for any two events $E, F \in \mathcal{A}$, $E \succsim_i^{\omega^{+1}} F$ if and only if $\mu(E) \geq \mu(F)$ for every $\mu \in \mathcal{C}_i^{\omega^{+1}}$.

Let $\mathcal{A}^{+1} = 2^{\Omega^{+1}}$ denote the (sigma) algebra of all subsets of Ω^{+1}. For each individual i and each state $\omega^{+1} \in \Omega^{+1}$ define a set of probabilities $T_i^{+1}(\omega^{+1})$ over $(\Omega^{+1}, \mathcal{A}^{+1})$ in the following manner: for each probability $\mu \in \mathcal{C}_i^{\omega^{+1}}$, the set $T_i^{+1}(\omega^{+1})$ contains all probability measures over $\mathcal{A}^{+1}$ which obtain the same values as μ on events in $\mathcal{A}$.

Consider the model $M(\psi) = (I, \Omega^{+1}, \mathcal{A}^{+1}, 2^{\mathcal{P}}, \mathcal{B}, T_0, (T_i)_{i\in I}^{+1})$, such that for every atomic formula $p \in \mathcal{P}$, $p \in T_0(\omega)$ if and only if $p \in \omega$. It is next shown that the resulting $[\cdot]$ notation is equivalent to the $|\cdot|$ notation.

Claim 11. *For any formula φ in $\mathcal{L}(\psi)^{+1}$, $[\varphi] = |\varphi|$.*

Proof. The proof is conducted using induction on the depth of the formula.

If φ is an atomic formula, the equivalence holds by definition. If for a formula φ', $|\varphi'| = [\varphi']$ and $\varphi = \neg\varphi'$, then by the definition of the states as maximal consistent subsets, for any state ω,

$$\begin{aligned} \neg\varphi' \in \omega &\Leftrightarrow \varphi' \notin \omega \\ &\Leftrightarrow \omega \notin |\varphi'| \\ &\Leftrightarrow \omega \notin [\varphi'] \\ &\Leftrightarrow M(\psi), \omega \not\models \varphi' \\ &\Leftrightarrow M(\psi), \omega \models \neg\varphi' \ . \end{aligned}$$

Next, let φ_1 and φ_2 be two formulas that satisfy $|\varphi_1| = [\varphi_1]$ and $|\varphi_2| = [\varphi_2]$, and consider their disjunction, $\varphi_1 \vee \varphi_2$. In that case,

$$\begin{aligned} \omega \in |\varphi_1 \vee \varphi_2| &\Leftrightarrow \varphi_1 \vee \varphi_2 \in \omega \\ &\Leftrightarrow \varphi_1 \in \omega \text{ or } \varphi_2 \in \omega \\ &\Leftrightarrow \omega \in |\varphi_1| = [\varphi_1] \text{ or } \omega \in |\varphi_2| = [\varphi_2] \\ &\Leftrightarrow \omega \in [\varphi_1] \cup [\varphi_2] = [\varphi_1 \vee \varphi_2] \ . \end{aligned}$$

Lastly, suppose that φ_1 and φ_2 are such that $|\varphi_1| = [\varphi_1]$ and $|\varphi_2| = [\varphi_2]$, and consider the formula $\varphi_1 \succsim_i \varphi_2$ for some i. Using the definition of the relation $\succsim_i^\omega$,

$$\begin{aligned} [\varphi_1 \succsim_i \varphi_2] &= \{\omega \in \Omega^{+1} \mid M(\psi), \omega \models \varphi_1 \succsim_i \varphi_2\} \\ &= \{\omega \in \Omega^{+1} \mid \mu([\varphi_1]) \geq \mu([\varphi_2]), \forall \mu \in T_i(\omega)\} \\ &= \{\omega \in \Omega^{+1} \mid \mu(|\varphi_1|) \geq \mu(|\varphi_2|), \forall \mu \in T_i(\omega)\} \\ &= \{\omega \in \Omega^{+1} \mid |\varphi_1| \succsim_i^\omega |\varphi_2|\} \\ &= \{\omega \in \Omega^{+1} \mid (\varphi_1 \succsim_i \varphi_2) \in \omega\} \\ &= |\varphi_1 \succsim_i \varphi_2| \ . \end{aligned}$$

Therefore, for any formula φ in $\mathcal{L}(\psi)^{+1}$, $[\varphi] = |\varphi|$. ■

If ψ is true under any model in $\mathcal{M}$, then it is true in particular under $M(\psi)$, that is, it obtains under any state ω in $M(\psi)$. Put differently, $[\top] = [\psi]$. From the claim and part (b) of Lemma 10 it follows that $\vdash \psi$, as required. ■

6.4 Proof of Corollary 3

The following proof imitates similar arguments in a proof by Gärdenfors (1975).

Any formula ψ belongs to the language $\mathcal{L}(\psi)^{+1}$ defined above, and by Lemma 10, ψ is a theorem of Σ if and only if it obtains in all states of $M(\psi)$. $M(\psi)$, in turn, is a finite model: it is constructed from a finite number of formulas, thus consisting of a finite number of states and events. Though the number of probabilities in a set $T_i(\omega)$ is infinite, there are only a finite number of partial orders on (the finite number of) events that admit such representation (which are those that satisfy assumptions P1-P4 of Alon and Lehrer (2012)), therefore in order to determine whether ψ obtains in all states of $M(\psi)$ it is required to check only this finite number of partial orders. Decidability is thus satisfied.

6.5 Proof of Lemma 4

First observe that,

$$\begin{aligned} [\psi_1^1] &= [(\psi_0 \succsim_1 \neg\psi_0) \vee (\neg\psi_0 \succsim_1 \psi_0)] \\ &= \{\omega \in \Omega \mid \mu([\psi_0]) \geq \mu([\neg\psi_0]) \text{ for all } \mu \in T_1(\omega)\} \cup \\ &\quad\ \{\omega \in \Omega \mid \mu([\neg\psi_0]) \geq \mu([\psi_0]) \text{ for all } \mu \in T_1(\omega)\} \\ &= \{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid a_1 = 1\}\ , \end{aligned}$$

and similarly $[\psi_1^2] = \{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid b_1 = 1\}$. It is next proved, by induction on ℓ, that for every $\ell = 1, \ldots, n$,

$$[\psi_\ell^1] = \{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid a_\ell = 1\}$$

and

$$[\psi_\ell^2] = \{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid b_\ell = 1\}\ .$$

Assume that $[\psi_{\ell-1}^1] = \{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid a_{\ell-1} = 1\}$ and $[\psi_{\ell-1}^2] = \{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid b_{\ell-1} = 1\}$. If $a_\ell = 1$, then individual 1 knows the value of $b_{\ell-1}$. In that case individual 1 can obviously compare the likelihoods of $\{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid b_{\ell-1} = 1\}$ and $\{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid b_{\ell-1} = 0\}$ (individual 1 attributes probability 1 to one of these events). Otherwise, if $a_\ell \neq 1$, then $a_\ell = 0$ and individual 1 is totally ignorant about the value of $b_{\ell-1}$, and hence individual 1 deems the events $\{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid b_{\ell-1} = 1\}$ and $\{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid b_{\ell-1} = 0\}$ incomparable. By the induction assumption, $\{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid b_{\ell-1} = 1\} = [\psi_{\ell-1}^2]$, thus if $\omega \in \{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid a_\ell = 1\}$, then either $\mu([\psi_{\ell-1}^2]) \geq \mu([\neg\psi_{\ell-1}^2])$ for every $\mu \in T_1(\omega)$ or $\mu([\neg\psi_{\ell-1}^2]) \geq \mu([\psi_{\ell-1}^2])$ for every $\mu \in T_1(\omega)$, resulting that either $\psi_{\ell-1}^2 \succsim_1 \neg\psi_{\ell-1}^2$ or

$\neg\psi^2_{\ell-1} \succsim_1 \psi^2_{\ell-1}$ obtain in ω, and this is precisely the definition of $[\psi^1_\ell]$. Similarly, if $\omega \in \{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid a_\ell = 0\}$, then neither $\mu([\psi^2_{\ell-1}]) \geq \mu([\neg\psi^2_{\ell-1}])$ for all $\mu \in T_1(\omega)$ nor $\mu([\neg\psi^2_{\ell-1}]) \geq \mu([\psi^2_{\ell-1}])$ for all $\mu \in T_1(\omega)$, hence $\neg\left((\psi^2_{\ell-1} \succsim_1 \neg\psi^2_{\ell-1}) \vee (\neg\psi^2_{\ell-1} \succsim_1 \psi^2_{\ell-1})\right)$ obtains in ω. It follows that $[\psi^1_\ell] = \{(s, a_1 \cdots a_n, b_1 \cdots b_n) \mid a_\ell = 1\}$. Analogous arguments prove the result for $[\psi^2_\ell]$. ■

6.6 Proof of Proposition 5

For soundness, let M be a model as in the above proof of the soundness part of Theorem 2. Observe first that if $\mu([\varphi]) \geq \mu([\psi])$ for every $\mu \in T_i(\omega)$, then for every $\omega' \in [T_i(\omega)]$, by definition of equivalence of probabilities sets, it also holds that $\mu'([\varphi]) \geq \mu'([\psi])$ for every $\mu' \in T_i(\omega')$. It follows that $[T_i(\omega)] \subseteq [\varphi \succsim_i \psi]$, hence for all $\mu \in T_i(\omega)$, $\mu([\varphi \succsim_i \psi]) = 1$. Therefore if $M, \omega \models \varphi \succsim_i \psi$ then $M, \omega \models (\varphi \succsim_i \psi) \sim_i \top$, for every state ω of every model M.

Similarly, if $M, \omega \models \neg(\varphi \succsim_i \psi)$ then there exists $\mu \in T_i(\omega)$ for which $\mu([\psi]) > \mu([\varphi])$, hence there is $\mu' \in T_i(\omega')$ for which $\mu'([\psi]) > \mu'([\varphi])$, for every $\omega' \in [T_i(\omega)]$. It follows again that $[T_i(\omega)] \subseteq [\neg(\varphi \succsim_i \psi)]$, implying that for all $\mu \in T_i(\omega)$, $\mu([\neg(\varphi \succsim_i \psi)]) = 1$, thus $M, \omega \models \neg(\varphi \succsim_i \psi) \sim_i \top$.

Completeness of the system Σ + (I1) + (I2) is proved in a similar manner to the completeness proof of Theorem 2. Similarly to the restricted languages defined there, the language $\mathcal{L}(\psi)^{+2}$ is now applied, having the same atomic formulas as in ψ, and depth no more than $dp(\psi)+2$. This way formulas that express introspection concerning likelihood judgements over formulas from $\mathcal{L}(\psi)$ may be added. Now the state space Ω^{+2}, generated by maximally consistent subsets of formulas from $\mathcal{L}(\psi)^{+2}$, may be constructed. Any such subset, whenever it contains a formula $\varphi \succsim_i \psi$ from $\mathcal{L}(\psi)^{+1}$, will also contain the formula $(\varphi \succsim_i \psi) \sim_i \top$, and similarly for $\neg(\varphi \succsim_i \psi)$. Put differently, $|\varphi \succsim_i \psi| = |(\varphi \succsim_i \psi) \sim_i \top|$, for every event $|\varphi \succsim_i \psi|$ from $\mathcal{A}^{+1}$.

As in the proof of Theorem 2, for every state in Ω^{+2} a set of probability measures that represents the relation $\succsim_i$ over events in $\mathcal{A}$ may be derived. Each such set may be extended to events in the finer, finite (sigma-)algebra of all subsets of Ω^{+2}: for each μ in the representing set and formula $(\varphi \succsim_i \psi) \sim_i \top$ that belongs to $\mathcal{L}(\psi)^{+2}$ but not to $\mathcal{L}(\psi)^{+1}$, set $\mu((|\varphi \succsim_i \psi) \sim_i \top|) = 1$ whenever $(\varphi \succsim_i \psi) \sim_i \top \in \omega$. Other than that restriction, the probabilities are extended to all subsets of Ω^{+2} in the same manner as in the proof of Theorem 2.

Suppose that $(\varphi \succsim_i \psi) \in \omega$. If $(\varphi \succsim_i \psi) \in \mathcal{L}(\psi)$ then $|\varphi \succsim_i \psi| \in \mathcal{A}$, and since it must be that $(\varphi \succsim_i \psi) \sim_i \top \in \omega$, and $(\varphi \succsim_i \psi) \sim_i \top \in \mathcal{L}(\psi)^{+1}$, then the set of probabilities represents that relation, and $\mu(|\varphi \succsim_i \psi|) = 1$ for all $\mu \in T_i(\omega)$. Otherwise $(\varphi \succsim_i \psi) \sim_i \top \in \mathcal{L}(\psi)^{+2}$, and $\mu(|\varphi \succsim_i \psi|) = \mu(|(\varphi \succsim_i \psi) \sim_i \top|) = 1$ by construction.

Finally, the states in $T_i(\omega)$ are precisely those states that consist of the same likelihood rankings as in ω, and as the languages considered are finite, there are only finitely many such rankings. That is, $T_i(\omega) = \bigcap_{(\varphi \succsim_i \psi)\in\omega} |\varphi \succsim_i \psi|$, where there are only finitely many events in this intersection. By the above, each $\mu \in T_i(\omega)$ lends probability one to each event in the intersection, thus to the entire intersection.

References

Ahn, D. (2007), Hierarchies of Ambiguous Beliefs, Journal of Economic Theory 136(1): 286-301

Alon S. and E. Lehrer (2013), "Subjective multi-prior probability: A representation of a partial likelihood relation", mimeo

Armbruster, W. and W. Böge (1979), "Bayesian Game Theory", in O. Moeschlin and D. Pallaschke (eds.), Game Theory and Related Topics, North-Holland, Amsterdam.

Bewley, T. (2002), "Knightian Decision Theory: Part I", Decisions in Economics and Finance, 25: 79-110. (Working paper, 1986)

Bodoh-Creed, A. (2012) "Ambiguous Beliefs and Mechanism Design", Games and Economic Behavior, 75, 518-537.

Böge, W. and T. Eisele (1979), "On Solutions of Bayesian Games", International Journal of Game Theory, 8 (4), 193-215.

Bose, S., Ozdenoren, E., and A. Pape (2006), "Optimal Auctions with Ambiguity", Theoretical Economics, 1, 411-438.

Bose, S. and L. Renou (2011), Mechanism Design With Ambiguous Communication Devices, mimeo

de Castro, L. I., M. Pesce, and N. C. Yannelis (2011), "Core and equilibria under ambiguity", Economic Theory, 48 (2-3), 519-548.

de Castro, L. I. and A. Chateauneuf (2011), "Ambiguity aversion and trade", Economic Theory, 48 (2-3), 243-273.

Chellas, B. (1980), Modal Logic, Cambridge: Cambridge University Press

Ganguli, J. (2007), Common p-belief and uncertainty, Proceedings of the XI conference on Theoretical Aspects of Rationality and Knowledge (TARK), Dov Samet (ed.), Presses Univesitaires De Louvain, 155-159

Gärdenfors, P. (1975), "Qualitative Probability as an Intensional Logic", J. Philosoph. Logic 4, 171185.

Gilboa, I. and M. Marinacci (2011),"Ambiguity and the Bayesian Paradigm", mimeo

Gilboa, I. and D. Schmeidler (1989), "Maxmin Expected Utility with a Non-Unique Prior", Journal of Mathematical Economics, 18: 141-153

Hansen, L. P., and T. J. Sargent (2001), Robust control and model uncertainty, American Economic Review, 91, 60-66.

Harsanyi, J. (1967-8), Games with incomplete information played by 'Bayesian' players, Parts I-III, Management Science 8,159-182, 320-334, 486-502.

Heifetz, A and P. Mongin (2001), Probability Logic for Type Spaces, Games and Economic Behavior, 35, 31-53.

Kajii, A., Ui, T. (2005), Incomplete information games with multiple priors. Japanese Econ. Rev. 56, 332-351.

Kajii, A. and T. Ui (2006), Interim Efficient Allocations Under Uncertainty, Journal of Economic Theory 144, 337–353.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005), "A Smooth Model of Decision Making under Ambiguity," Econometrica, 73: 1849-1892.

Lopomo, G., L. Rigotti, and C. Shannon, (2009), Uncertainty in Mechanism Design, mimeo.

Lopomo, G., L. Rigotti, and C. Shannon, (2011), Knightian uncertainty and moral hazard, Journal of Economic Theory 146 (3), 1148-1172.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006), "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences", Econometrica, 74: 1447-1498.

Meier, M. (2012), An infinitary probability logic for type spaces, Israel Journal of Mathematics, 192, 1-58.

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

Segerberg, K., Qualitative Probability in a Modal Setting, Proc. of the 2nd Scand. Log. Symp. (ed. Fenstad), Amsterdam 1971.

Stauber, R. (2011), Knightian games and robustness to ambiguity, Journal of Economic Theory 146, 248-274.

Ui, T. (2009), Ambiguity and Risk in Global Games, mimeo.

Zhou, C. (2009), A Complete Deductive System for Probability Logic, J. Log. Comput. 19, 1427-1454.