

Online Supplementary Appendix

The Existence of Universal Qualitative Belief Spaces

Satoshi Fukuda[†]

December 17, 2023

The Supplementary Appendix provides additional discussions and results. It is organized as follows. Appendix B supplements Section 3, providing graphical illustrations of the six steps of the proof of Theorem 1.

Appendix C supplements the discussions in Section 4.2 about Bayesian equilibria. Appendix D supplements Section 5, demonstrating the existence of a terminal space for richer settings. Appendix D.1 establishes a terminal conditional-belief space, and Appendix D.2 a terminal dynamic knowledge-belief space. Appendix D.3 discusses further applications: knowledge and unawareness, preferences, and expectations.

Appendix E supplements Section 6, discussing the solution concepts of Iterated Elimination of Strictly Dominated Actions (IESDA), Iterated Elimination of Börgers Dominated Actions (IEBA), Iterated Elimination of Inferior Action Profiles (IEIP), pure-strategy Nash equilibria, and correlated equilibria. Appendix F supplements Section 7. It characterizes minimality. The proofs are relegated to Appendix G.

B Section 3

This section supplements the illustration of each step of the proof of Theorem 1 through figures. The first step starts with Definition 5, which inductively defines κ -expressions: logical formulas defined solely in terms of (S, \mathcal{S}) and players' interactive beliefs about (S, \mathcal{S}) .

The left panel of Figure S.1 illustrates Definition 5 and Remark 3. The bold rectangle defines the collection of expressions \mathcal{L} . As illustrated in the panel, I start with events E in $\mathcal{A}_\kappa(\mathcal{S})$ (the inner-most dashed rectangle). This corresponds to the base step of the inductive definition. The dashed arrow that connects the event E in the set $\mathcal{A}_\kappa(\mathcal{S})$ (the solid small rectangle) and the expression E in the set $\mathcal{A}_\kappa(\mathcal{S}) \subseteq \mathcal{L}$ (the inner-most dashed rectangle) shows that each event E in the set $\mathcal{A}_\kappa(\mathcal{S})$ is identified as a syntactic formula (i.e., an expression) in \mathcal{L} .

The set of expressions \mathcal{L} is closed, in particular, under the operation β_i : for each expression e , the expression $\beta_i(e)$ denotes “player i believes e .” The dashed arrow from E to $\beta_i(E)$ illustrates this point. In light of Remark 3, $\beta_i(E)$ resides in \mathcal{L}_1 . The dashed rectangle that encloses the inner-most dashed rectangle (that represents the set $\mathcal{L}_0 = \mathcal{A}_\kappa(\mathcal{S})$) represents \mathcal{L}_1 . Defining the sequence $(\mathcal{L}_\alpha)_{\alpha=0}^{\bar{\kappa}}$ of expressions that

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

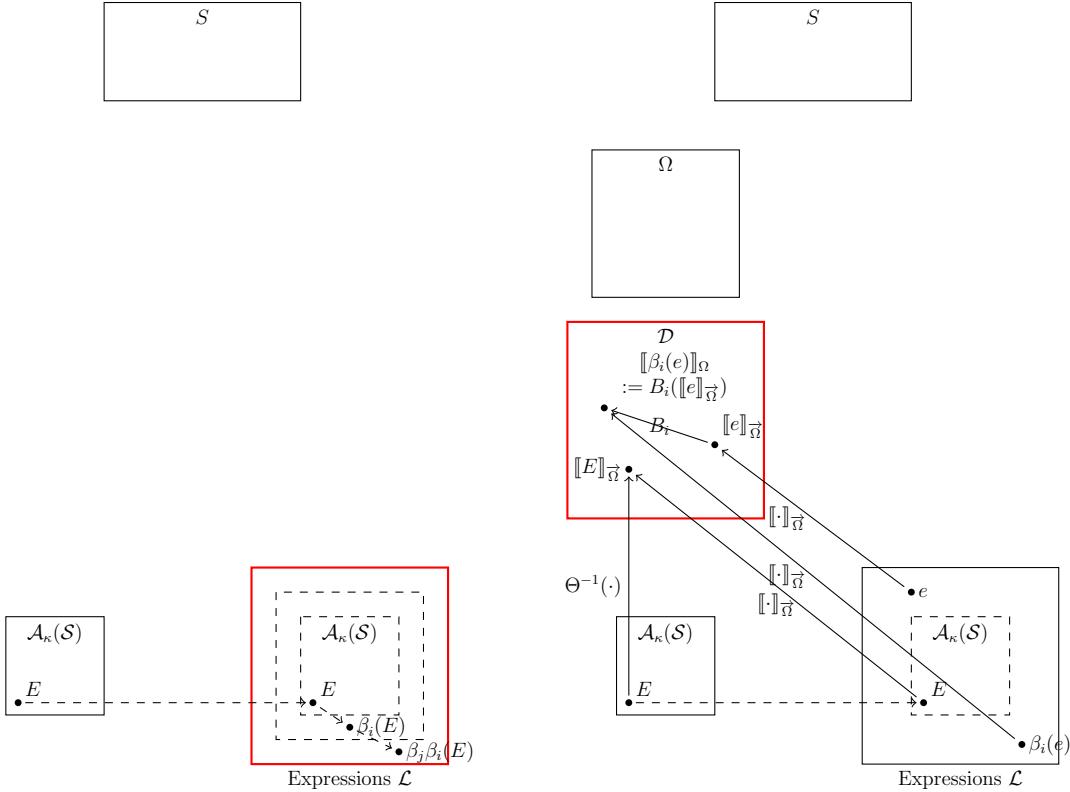


Figure S.1: Illustration of Step 1: Definition 5 and Remark 3 (Left) and Definition 6 (Right).

represents “ α -order beliefs” (formally, see Remark 2 for the depth α of an expression) inductively, one can construct the set of expressions \mathcal{L} as $\mathcal{L} = \mathcal{L}_{\vec{\alpha}}$, as in Remark 3. Note that the left panel only depicts $\mathcal{L}_0 = \mathcal{A}_\kappa(\mathcal{S})$, \mathcal{L}_1 , and $\mathcal{L} = \mathcal{L}_{\vec{\alpha}}$.

Once the expressions are defined, in each belief space $\vec{\Omega} = \langle (\Omega, \mathcal{D}), (B_i)_{i \in I}, \Theta \rangle$, Definition 6 identifies each expression $e \in \mathcal{L}$ with the corresponding event $\llbracket e \rrbracket_{\vec{\Omega}} \in \mathcal{D}$. The right panel of Figure S.1 illustrates Definition 6. Starting from an expression $E \in \mathcal{A}_\kappa(\mathcal{S})$ (in the dashed rectangle), $\llbracket E \rrbracket_{\vec{\Omega}} := \Theta^{-1}(E) \in \mathcal{D}$ is an event (in the bold rectangle) that corresponds to the event of nature $E \in \mathcal{A}_\kappa(\mathcal{S})$. If $e \in \mathcal{L}$ is an expression and if $\llbracket e \rrbracket_{\vec{\Omega}} \in \mathcal{D}$ is the corresponding event, then the expression $\beta_i(e) \in \mathcal{L}$ corresponds to the event $\llbracket \beta_i(e) \rrbracket_\Omega := B_i(\llbracket e \rrbracket_{\vec{\Omega}}) \in \mathcal{D}$ (in the bold rectangle).

The second step defines descriptions. For each belief space $\vec{\Omega}$, Definition 7 defines the description $D(\omega)$ of each state $\omega \in \Omega$. As illustrated in the bold rectangle in the left panel of Figure S.2, the collection of descriptions Ω^* , which is shown to be a set, turns out to be the underlying state space of the terminal belief space. For any given belief space $\vec{\Omega}$, the bold arrow depicts the description map D . It associates, with each state $\omega \in \Omega$, the description $D(\omega) \in \Omega^*$.

Each description $D(\omega)$ consists of the unique nature state $\Theta(\omega) \in S$ and the set

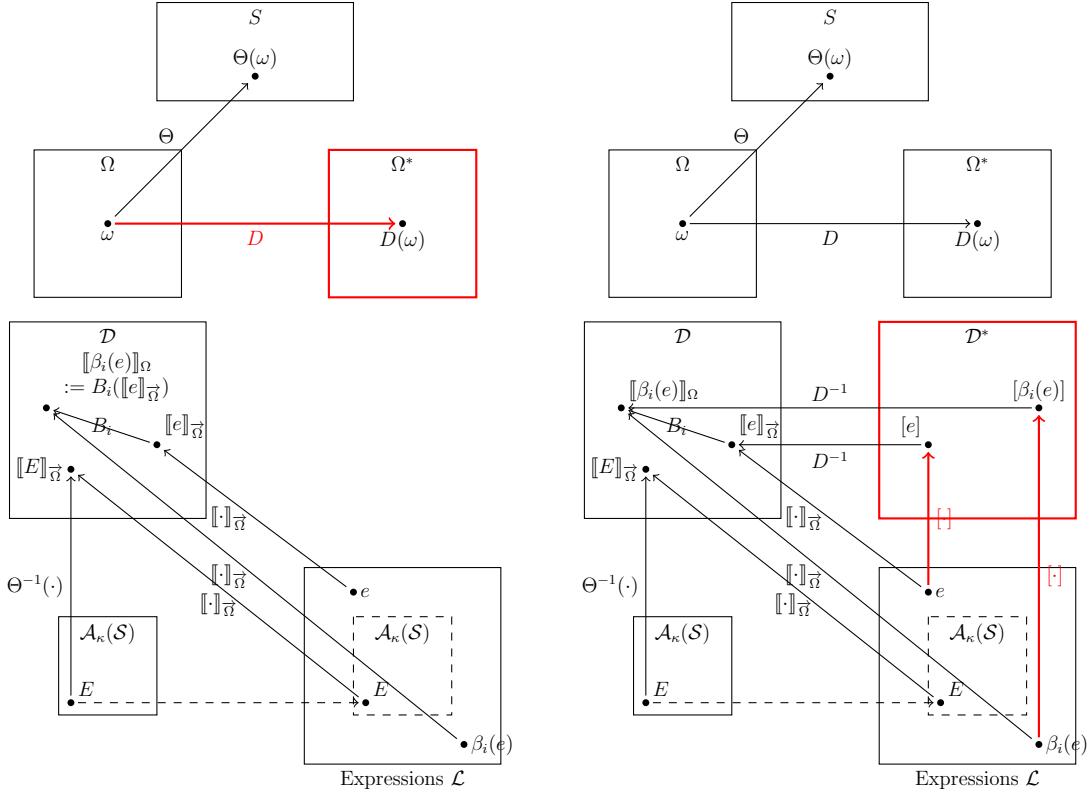


Figure S.2: Illustration of Step 2 (Left) and Step 3 (Right).

of expressions that hold at ω (formally, $\{e \in \mathcal{L} \mid \omega \in \llbracket e \rrbracket_{\vec{\Omega}}\}$). Figure S.3 depicts the subset of $\mathcal{L} = \mathcal{L}_\kappa(\mathcal{A}_\kappa(\mathcal{S}))$ associated with $D(\omega)$. The left-hand side depicts the state space Ω and the state ω . The arrow from $\omega \in \Omega$ to the shaded area, which depicts $D(\omega)$, shows that D is a mapping (the description map). The set $D(\omega)$ contains the expressions that hold at ω . As illustrated in the right-hand side of the figure, an expression f satisfies $f \in D(\omega)$ iff $(\neg f) \notin D(\omega)$. For instance, the figure illustrates this fact for $f \in \{E, \beta_i(E)\}$ (also, since $\beta_j \beta_i(E) \notin D(\omega)$, it follows that $(\neg \beta_j) \beta_i(E) \in D(\omega)$). In the state space Ω on the left-hand side, the set $\llbracket f \rrbracket_{\vec{\Omega}}$ of states at which an expression f holds is depicted as a circle, for each $f \in \{E, \beta_i(E), (\neg \beta_j) \beta_i(E)\}$. The state ω is in the intersection of the circles, and thus $D(\omega)$ contains the corresponding expressions. The figure also shows that one can interpret $D(\omega)$ as the players' belief hierarchies at ω (Remark 5).

The third step defines the domain \mathcal{D}^* , which consists of the sets $[e]$ for expressions e . In the right panel of Figure S.2, the bold rectangle depicts the domain \mathcal{D}^* . The bold arrow from $e \in \mathcal{L}$ to $[e] \in \mathcal{D}^*$ (also, the one from $\beta_i(e) \in \mathcal{L}$ to $[\beta_i(e)] \in \mathcal{D}^*$) depicts the definition. Lemma 1 shows that $\llbracket e \rrbracket_{\vec{\Omega}} = D^{-1}([e])$ for each expression $e \in \mathcal{L}$. In the right panel of Figure S.2, the triangle connecting $e \in \mathcal{L}$, $[e] \in \mathcal{D}^*$, and $\llbracket e \rrbracket_{\vec{\Omega}} \in \mathcal{D}$ depicts Lemma 1.

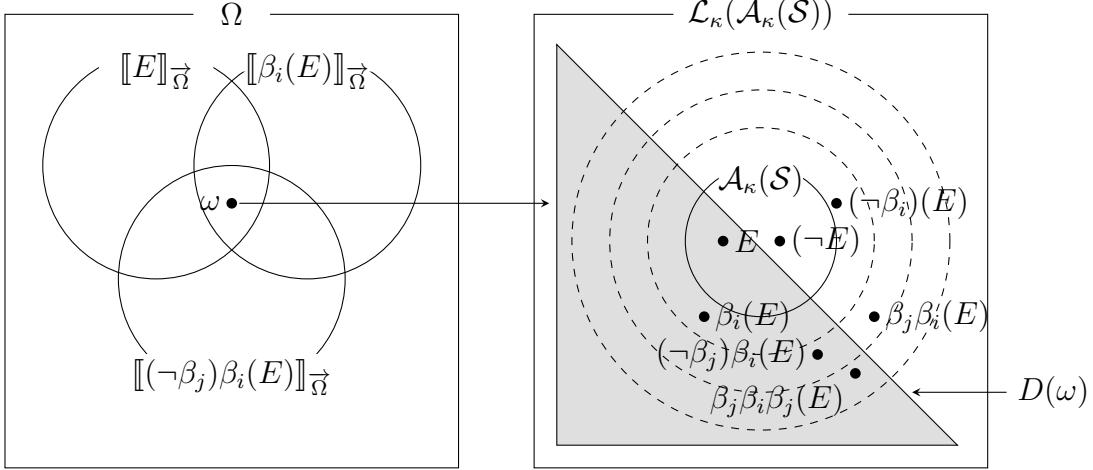


Figure S.3: Illustration of the Description $D(\omega)$ of ω (Step 2).

The fourth step defines $\Theta^* : \Omega^* \rightarrow S$. The bold arrow from $D(\omega) \in \Omega^*$ to $\Theta^*(D(\omega)) \in S$ in the left panel of Figure S.4 illustrates the mapping Θ^* . The diagram that connects ω , $D(\omega)$, and $\Theta(\omega) = \Theta^*(D(\omega))$ depicts the first part of Lemma 2.

The fifth step defines each player i 's belief operator $B_i^* : \mathcal{D}^* \rightarrow \mathcal{D}^*$ by $B_i^*([e]) := [\beta_i(e)]$. The bold arrow from $[e]$ to $B_i^*([e]) := [\beta_i(e)]$ in the right panel of Figure S.4 illustrates the definition of B_i^* in Lemma 3. The right panel also shows that, starting from $[e] \in \mathcal{D}^*$, one has $D^{-1}B_i^*([e]) = B_iD^{-1}([e])$. Thus, the quadrilateral connecting $[e]$, $B_i^*([e]) := [\beta_i(e)]$, $[\beta_i(e)]_{\vec{\Omega}} := B_i([\vec{e}]_{\vec{\Omega}})$, and $[\vec{e}]_{\vec{\Omega}}$ depicts the second part of Lemma 3.

The sixth step shows that the description map $D_{\vec{\Omega}^*}$ on Ω^* is the identity map (Lemma 4). The bold arrow from ω^* to itself in the left panel of Figure S.5 illustrates the lemma. Recall from the discussion in Remark 6 in the main text that there exists at most one morphism from a non-redundant belief space (recall also that a belief space is non-redundant if its description map is injective). Thus, Lemma 4 implies that the description map $D_{\vec{\Omega}}$ is the unique morphism from a given belief space $\vec{\Omega}$ to $\vec{\Omega}^*$, establishing Theorem 1.

In sum, the right panel of Figure S.5 illustrates the role of the definitions and lemmas in Section 3 in establishing Theorem 1 in a single panel.

C Section 4.2

As discussed in Section 4.2, I briefly and informally mention a possibility that one may be able to analyze Bayesian equilibria of a fixed underlying game in a terminal belief space (that depends on the fixed underlying game) by incorporating players' strategy

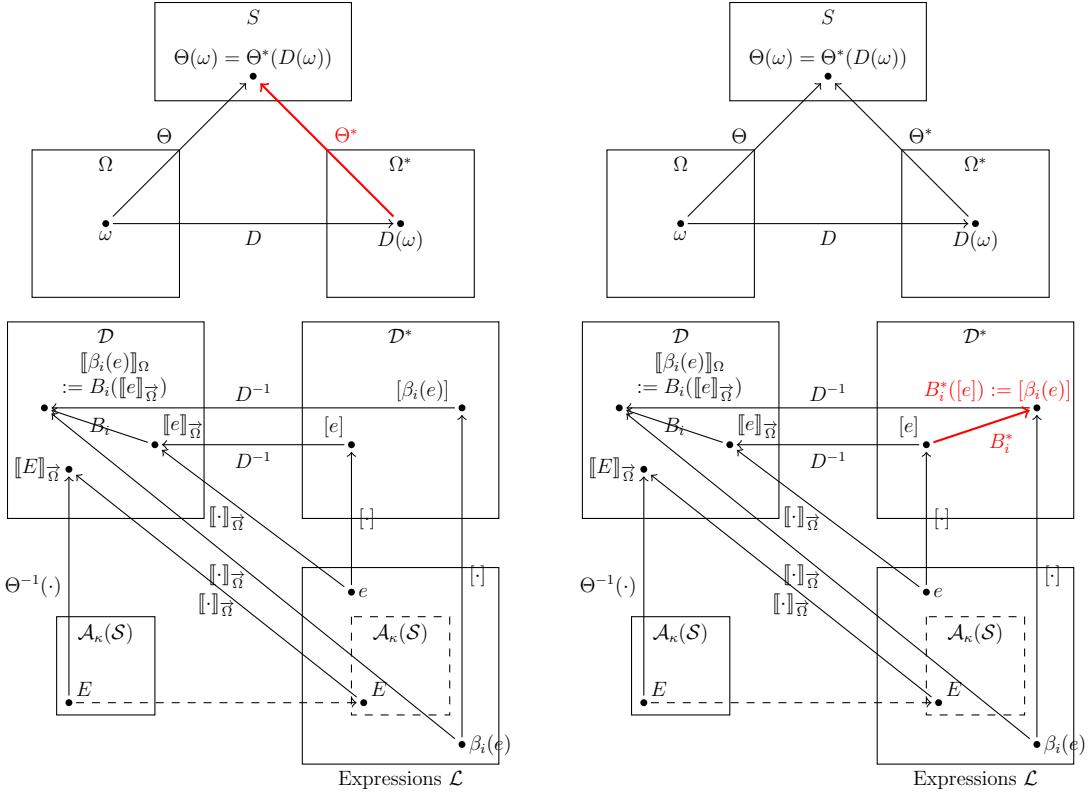


Figure S.4: Illustration of Step 4 (Left) and Step 5 (Right).

choices as a primitive of a belief space. Let (S, \mathcal{S}) be a measurable parameter space of payoff uncertainty. Let $\langle (A_i, \mathcal{A}_i)_{i \in I}, (u_i)_{i \in I} \rangle$ be an S -based game (i.e., an underlying game): each (A_i, \mathcal{A}_i) is a measurable space of player i 's actions; (A, \mathcal{A}) is the product measurable space of action profiles; and each $u_i : S \times A \rightarrow \mathbb{R}$ is a bounded Borel-measurable payoff function. Let “0” stand for nature.

A *Bayesian game* is a tuple $\vec{\Omega} = \langle (\Omega, \mathcal{D}), (m_i)_{i \in I}, \Theta \rangle$, which is a particular \aleph_1 -belief space on $S \times A$ in the following sense. The set of states of the world (Ω, \mathcal{D}) is a product measurable space consisting of the parameter space (S, \mathcal{S}) and each player's type set (T_i, \mathcal{T}_i) : $\Omega = S \times \prod_{i \in I} T_i$ and the domain \mathcal{D} is the corresponding product σ -algebra. Each player $i \in I$ has a type mapping $m_i : \Omega \rightarrow \Delta(\Omega)$, where $\Delta(\Omega)$ is the set of probability measures on (Ω, \mathcal{D}) . It associates, with each state of the world ω , her belief $m_i(\omega) \in \Delta(\Omega)$ held at ω with the requirement that the marginal of $m_i(\omega)$ on T_i is the Dirac measure concentrated at $\omega_i \in T_i$ (i.e., each player is certain of her own beliefs).¹ The measurable function $\Theta : (\Omega, \mathcal{D}) \rightarrow (S \times A, \mathcal{S} \times \mathcal{A})$ consists of the

¹Two technical remarks are in order. First, the type mapping m_i is required to be measurable (see Section 5.1 for the σ -algebra on $\Delta(\Omega)$ with respect to which m_i is required to be measurable). Second, Section 5.1 shows that the type mapping $m_i : \Omega \rightarrow \Delta(\Omega)$ is equivalently expressed as a collection of (p) -belief operators $(B_i^p)_{p \in [0,1]}$. There exist a list of properties on B_i^p under which B_i^p induces a type mapping m_i ; and m_i , in turn, induces the original B_i^p (see Definition 10). Technically,

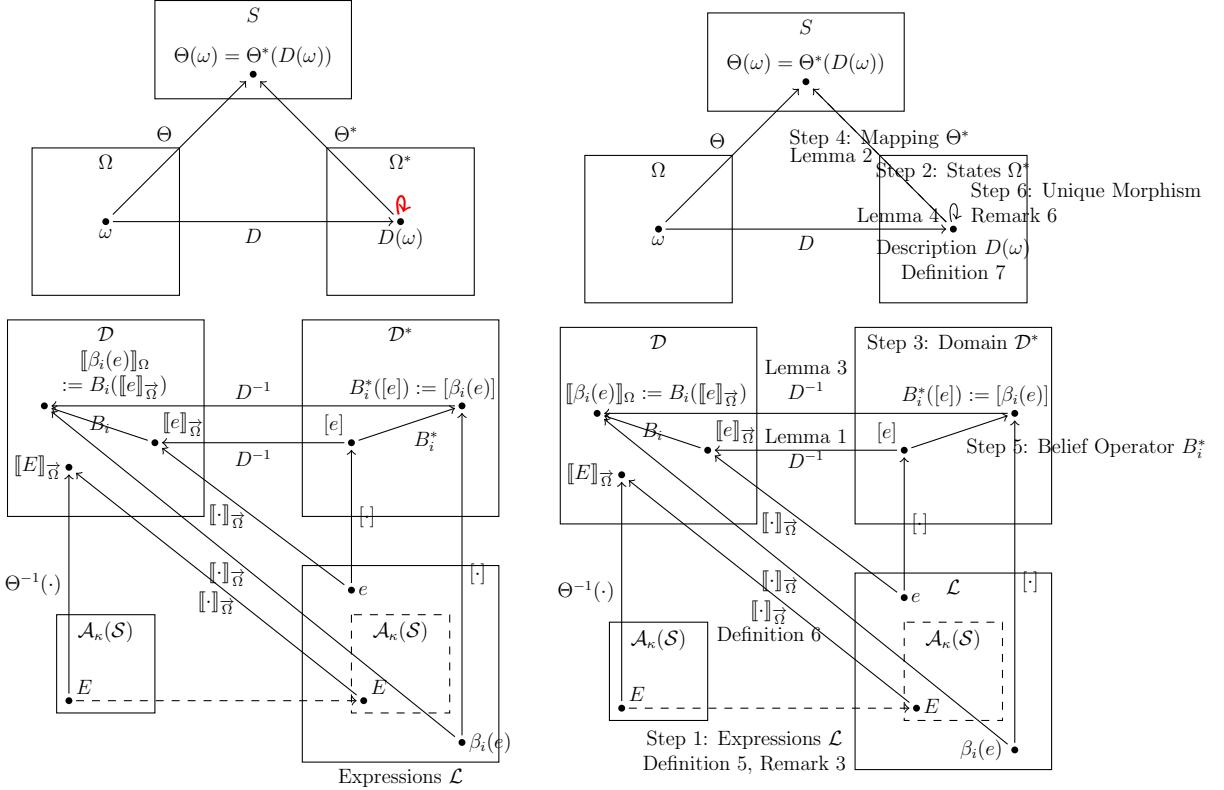


Figure S.5: Illustration of Step 6 (Left) and Overall Interrelations among the Definitions and Lemmas for Theorem 1 (Right).

identity map $\Theta_0 = \text{id}_S : S \rightarrow S$ and player i 's strategy $\Theta_i : (T_i, \mathcal{T}_i) \rightarrow (A_i, \mathcal{A}_i)$.² Thus, not only does Θ specify how each state of the world ω corresponds to the corresponding payoff parameter s but also it specifies each player's strategy Θ_i . Note that strategies are restricted to measurable pure strategies.³

With some abuse of terminology, a Bayesian equilibrium is a particular Bayesian game which respects the individual players' incentive-compatibility condition: for all $i \in I$, $\omega \in \Omega$, and $a_i \in A_i$,

$$\int_{\Omega} u_i(\Theta(\tilde{\omega})) m_i(\omega)(d\tilde{\omega}) \geq \int_{\Omega} u_i(\Theta_0(\tilde{\omega}), a_i, (\Theta_j(\tilde{\omega}))_{j \in I \setminus \{i\}}) m_i(\omega)(d\tilde{\omega}).$$

A *morphism* from one Bayesian game to another is a morphism in the category of belief spaces on $S \times A$. Thus, it preserves the states of nature and, by definition, the

the list of properties (including the property that player i is certain of her own beliefs) are defined, more generally, irrespective of whether an underlying state space admits a product structure.

²That is, for each $\omega = (\omega_i)_{i \in I \cup \{0\}}$, $\Theta(\omega) = (\Theta_i(\omega_i))_{i \in I \cup \{0\}} \in S \times A$.

³This is a restriction. While it is beyond the scope of this paper, see Friedenberg and Meier (2017), Hellman (2014), and Simon (2003) for the non-existence of a measurable Bayesian (behavior-strategy) equilibrium.

players' beliefs and strategy choices (in contrast to the case in which the players' strategy choices are separated from the representation of their beliefs).⁴ If the underlying S -based game has a measurable pure-strategy Bayesian equilibrium, I conjecture that there exists a terminal Bayesian equilibrium in which, for any Bayesian equilibrium $\langle(\Omega, \mathcal{D}), (m_i)_{i \in I}, \Theta\rangle$ (i.e., for any representation $\langle(\Omega, \mathcal{D}), (m_i)_{i \in I}, \Theta_0\rangle$ of the players' interactive beliefs about the payoff parameters S and for any strategy choices of the players $(\Theta_i)_{i \in I}$ which constitute a Bayesian equilibrium), there is a unique morphism that extends to the terminal belief space $\langle(\Omega^*, \mathcal{D}^*), (m_i^*)_{i \in I}, \Theta^*\rangle$ that consists of the representation $\langle(\Omega^*, \mathcal{D}^*), (m_i^*)_{i \in I}, \Theta_0^*\rangle$ of the players' interactive beliefs about S and the players' strategy choices $(\Theta_i^*)_{i \in I}$.

The space Ω^* would consist of all possible belief hierarchies over $S \times A$ ranged over all Bayesian equilibria. Thus, this terminal structure would be different from the terminal belief space that consists of all possible belief hierarchies over S , i.e., the terminal type space of Brandenburger and Dekel (1993), Heifetz and Samet (1998), and Mertens and Zamir (1985). In particular, this terminal structure on $S \times A$ would be redundant as a belief space on S . This terminal structure would also be different from the terminal belief space on $S \times A$ in that Ω^* would consist of those belief hierarchies that correspond to some Bayesian equilibrium.

D Section 5

The framework of this paper applies to various forms of beliefs as long as players' beliefs are represented by belief operators. Section 5 has established a terminal probabilistic-belief space. Appendix D discusses further applications. Appendix D.1 discusses a terminal space for conditional probability systems (CPSs). Appendix D.2 introduces players' knowledge and qualitative beliefs indexed by time. Appendix D.3 briefly discusses further possible applications, namely, terminal knowledge-unawareness, preference, and expectation spaces.

D.1 Terminal Conditional-Belief Space

I construct a terminal space for conditional belief systems (CPSs) (Battigalli and Siniscalchi, 1999; Guarino, 2017).⁵ While this subsection focuses on CPS-based conditional beliefs, it suggests that a terminal conditional-belief space exists for a wide

⁴Following the terminology of Friedenberg and Meier (2017), the “extension property” is built into the requirement of a morphism itself.

⁵As discussed in footnote 47 of the main text, a state space here is not restricted to a product space. The framework here does not presuppose any topological restriction on nature states or any cardinal restriction on conditioning events, either. Thus, the construction of the terminal conditional-belief space would be complementary to Battigalli and Siniscalchi (1999) and Guarino (2017). For instance, Di Tillio, Halpern, and Samet (2014) study conditional beliefs on a non-product space and also provide game-theoretical applications to dynamic games.

variety of qualitative or probabilistic beliefs, which can be used for epistemic analyses of dynamic games including dynamic psychological games.

Call a triple $(\Omega, \mathcal{D}, \mathcal{C})$ a *conditional space* if: (i) (Ω, \mathcal{D}) is an \aleph_1 -algebra; (ii) \mathcal{C} is a non-empty sub-collection of \mathcal{D} with $\emptyset \notin \mathcal{C}$; and (iii) there exists a *conditional probability system* (CPS) μ on $(\Omega, \mathcal{D}, \mathcal{C})$. A function $\mu(\cdot| \cdot) : \mathcal{D} \times \mathcal{C} \rightarrow [0, 1]$ is a CPS if: (i) each $\mu(\cdot|C)$ is a countably-additive probability measure; (ii) Normality: $\mu(C|C) = 1$ for each $C \in \mathcal{C}$; and (iii) Chain Rule: $\mu(E|C) = \mu(E|D)\mu(D|C)$ for any $E \in \mathcal{D}$ and $C, D \in \mathcal{C}$ with $E \subseteq D \subseteq C$. Call each $C \in \mathcal{C}$ a *conditioning event* (or a *condition*, for short). Fix a conditional space $(S, \mathcal{A}_{\aleph_1}(\mathcal{S}), \mathcal{C}_S)$, where (S, \mathcal{S}) is the set of nature states and $S \in \mathcal{C}_S$.

Denote by $\Delta^{\mathcal{C}}(\Omega)$ the set of CPSs on $(\Omega, \mathcal{D}, \mathcal{C})$. I endow it with the \aleph_1 -algebra

$$\mathcal{D}_{\Delta}^{\mathcal{C}} := \mathcal{A}_{\aleph_1}(\{\{\mu \in \Delta^{\mathcal{C}}(\Omega) \mid \mu(E|C) \geq p\} \in \mathcal{P}(\Delta^{\mathcal{C}}(\Omega)) \mid (E, C, p) \in \mathcal{D} \times \mathcal{C} \times [0, 1]\}).$$

A player i 's *conditional-type mapping* is a measurable map $m_i : (\Omega, \mathcal{D}) \rightarrow (\Delta^{\mathcal{C}}(\Omega), \mathcal{D}_{\Delta}^{\mathcal{C}})$. I formulate a conditional-belief space using conditional p -belief operators $B_i^p(\cdot|C)$ for each player i and each condition C , so that i 's conditional p -belief operators induce her conditional-type mapping.

Definition S.1 (Conditional-Belief Space). A conditional-belief space of I on $(S, \mathcal{S}, \mathcal{C}_S)$ is a tuple $\vec{\Omega} := \langle (\Omega, \mathcal{D}, \mathcal{C}), (B_i^p(\cdot|C))_{(i, p, C) \in I \times [0, 1] \times \mathcal{C}}, \Theta \rangle$ with the following properties.

1. $(\Omega, \mathcal{D}, \mathcal{C})$ is a conditional space and $\Theta : (\Omega, \mathcal{D}) \rightarrow (S, \mathcal{A}_{\aleph_1}(\mathcal{S}))$ is a measurable map with $\mathcal{C} = \Theta^{-1}(\mathcal{C}_S)$.
2. For each $i \in I$, player i 's conditional p -belief operators $B_i^p(\cdot| \cdot) : \mathcal{D} \times \mathcal{C} \rightarrow \mathcal{D}$ satisfy the following.
 - (a) For each $C \in \mathcal{C}$, $(B_i^p(\cdot|C))_{p \in [0, 1]}$ satisfies Definition 10 (2a)-(2h).
 - (b) Certainty-of-Conditional-Beliefs: If $[m_{B_i}(\omega)] \subseteq E$ then $\omega \in B_i^1(E|\Omega)$, where

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

- (c) Normality: $B_i^1(C|C) = \Omega$ for all $C \in \mathcal{C}$.
- (d) Chain Rule: $B_i^p(E|D) \cap B_i^q(D|C) \subseteq B_i^{pq}(E|C)$ for any $E \in \mathcal{D}$ and $C, D \in \mathcal{C}$ with $E \subseteq D \subseteq C$.

By the assumption $\mathcal{C} = \Theta^{-1}(\mathcal{C}_S)$ in (1), denote $B_{i, C_S}^p(\cdot) := B_i^p(\cdot|\Theta^{-1}(C_S))$ for each $(i, p, C_S) \in I \times [0, 1] \times \mathcal{C}_S$. This means that conditions in each conditional-belief space are exogenously given as in Battigalli and Siniscalchi (1999) and Guarino (2017). Thus, one's conditional belief may fail to be a condition (i.e., $B_i^p(E|C) \notin \mathcal{C}$). Also, since $S \in \mathcal{C}_S$, “unconditional” beliefs $B_{i, S}^p(\cdot) = B_i^p(\cdot|\Omega)$ are also considered.

Conditions (2) characterize each player's conditional-type mapping as in Di Tillio, Halpern, and Samet (2014, Theorem 1). First, by (2a), slightly abusing the notation, a measurable map $m_{B_i} : (\Omega, \mathcal{D}) \rightarrow (\Delta^{\mathcal{C}}(\Omega), \mathcal{D}_{\Delta}^{\mathcal{C}})$ is well-defined as in Section 5.1. For conditional beliefs, Condition (2b) is the introspective property stating that player i is certain of her own conditional beliefs. The set $[m_{B_i}(\omega)]$ satisfies $[m_{B_i}(\omega)] = \{\omega' \in \Omega \mid m_{B_i}(\omega') = m_{B_i}(\omega)\}$: it consists of states ω' that player i cannot distinguish from ω based on her conditional beliefs. That is, she unconditionally believes E with probability one when $[m_{B_i}(\omega)]$ implies (i.e., is a subset of) E . Especially, $B_{i,C_S}^p(E) \subseteq B_{i,S}^1 B_{i,C_S}^p(E)$ and $(\neg B_{i,C_S}^p)(E) \subseteq B_{i,S}^1 (\neg B_{i,C_S}^p)(E)$ hold: if player i p -believes (does not p -believe) E conditional on $\Theta^{-1}(C_S)$, then she unconditionally 1-believes that she p -believes (does not p -believe) E conditional on $\Theta^{-1}(C_S)$.

By (2c), each $m_{B_i}(\omega)(\cdot|\cdot)$ satisfies Normality. Under (2a) and (2c), it can be seen that (2d) characterizes the Chain Rule.

A (*conditional-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) $B_{i,C_S}^p(\varphi^{-1}(\cdot)) = \varphi^{-1}(B_{i,C_S}^p(\cdot))$ for all $(i, p, C_S) \in I \times [0, 1] \times \mathcal{C}_S$. A conditional-belief space $\overrightarrow{\Omega}^*$ of I on $(S, \mathcal{S}, \mathcal{C}_S)$ is *terminal* if, for any conditional-belief space $\overrightarrow{\Omega}$ of I on $(S, \mathcal{S}, \mathcal{C}_S)$, there is a unique morphism $\varphi : \overrightarrow{\Omega} \rightarrow \overrightarrow{\Omega}^*$. By considering the probabilistic-beliefs of each player for each conditioning event, as in Section 5.1, Theorem 1 implies:

Corollary S.1 (Terminal Conditional-Belief Space). *There exists a terminal conditional-belief space $\overrightarrow{\Omega}^*$ of I on $(S, \mathcal{S}, \mathcal{C}_S)$.*

Corollary S.1 establishes the existence of a terminal conditional belief space. I briefly mention some related models of interest, namely, “lexicographic probability systems (LPSs)” or “hypothetical reasoning.” For LPSs, Tsakas (2014) defines a formal equivalence between conditional and lexicographic belief hierarchies in respective type spaces (under some topological assumptions on nature states and beliefs), and establishes the existence of a terminal lexicographic-belief space from a terminal conditional-belief space.⁶ Brandenburger, Friedenberg, and Keisler (2008) provide an epistemic characterization of iterated admissibility (avoidance of weak dominance) using a belief-complete type space (recall footnote 40 in the main text for belief-completeness) in which each type is associated with an LPS. For hypothetical reasoning, Di Tillio, Halpern, and Samet (2014), for instance, discuss the connection of conditional beliefs with hypothetical knowledge and counterfactual reasoning.

While detailed discussions on models of counterfactual reasoning are beyond the scope of this paper, Arieli and Aumann (2015) study a logical system where each player has a probability-one belief operator (thus, the notion of belief is rather qualitative) and show that an outcome of a perfect-information game where each player moves just once is consistent with rationality and common strong belief in rationality

⁶Technically, the notion of terminality in his paper is slightly weaker than the one adopted in this paper, namely, “ \aleph_1 -terminality” in Section 7.4.

iff it is a backward induction outcome. Roughly, a player strongly believes a statement if she believes it unless it is logically inconsistent with her node being reached.⁷ The reason that Arieli and Aumann (2015) use the syntactic approach is that the formalization of the concept of strong belief calls for the operation that refers to a valid expression: in the semantic framework, an expression is valid (in all belief spaces) iff it corresponds to the entire state space in every belief space (recall Definition 8). Since an expression is valid in the terminal space iff it is valid in all belief spaces (recall Proposition 2), the analysis of this paper would show that Arieli and Aumann (2015)'s epistemic characterization of backward induction outcomes for their class of perfect-information games would be possible within the semantic framework, namely, in the terminal space.

D.2 Terminal Dynamic Knowledge-Belief Space

Epistemic analyses of dynamic games may call for players' knowledge and beliefs.⁸ As in Battigalli and Bonanno (1997), I consider players' knowledge and beliefs indexed by time. While a knowledge operator $K_{i,t} : \mathcal{D} \rightarrow \mathcal{D}$ represents player i 's knowledge at time $t \in \mathbb{N}$, a belief operator $B_{i,t} : \mathcal{D} \rightarrow \mathcal{D}$ does her qualitative belief at time t .

Definition S.2 (Dynamic Knowledge-Belief Space). A dynamic κ -knowledge-belief space of I on (S, \mathcal{S}) is a tuple $\overrightarrow{\Omega} := \langle (\Omega, \mathcal{D}), (K_{i,t}, B_{i,t})_{(i,t) \in I \times \mathbb{N}}, \Theta \rangle$ with the following properties.

1. (Ω, \mathcal{D}) is a κ -algebra and the map $\Theta : (\Omega, \mathcal{D}) \rightarrow (S, \mathcal{A}_\kappa(\mathcal{S}))$ is measurable.
2. Knowledge operators $K_{i,t} : \mathcal{D} \rightarrow \mathcal{D}$ satisfy Truth Axiom, (Positive Introspection,) Negative Introspection, and the Kripke property. Belief operators $B_{i,t} : \mathcal{D} \rightarrow \mathcal{D}$ satisfy Consistency, Positive Introspection, Negative Introspection, and the Kripke property.
3. Knowledge and belief operators jointly satisfy: (i) $K_{i,t}(\cdot) \subseteq B_{i,t}(\cdot)$; (ii) $B_{i,t}(\cdot) \subseteq K_{i,t}B_{i,t}(\cdot)$; and (iii) $B_{i,t}(\cdot) = B_{i,t}B_{i,t+1}(\cdot)$.

In Condition (2), for ease of illustration, I have assumed (i) both knowledge and belief are fully introspective, (ii) knowledge is truthful while belief is consistent, and (iii) both knowledge and belief are represented by a possibility correspondence.

In Condition (3), the first condition means that knowledge implies belief at each time. The second states that each player knows her own belief at each time. Note that

⁷Note that Battigalli and Siniscalchi (2002) provide the notion of strong belief and an epistemic characterization of extensive-form rationalizability (and sufficient epistemic conditions for a backward induction outcome) using a belief-complete conditional type space.

⁸For instance, one may analyze players' knowledge about their past-observed moves and their beliefs about past-unobserved and future moves in an extensive-form game (e.g., Battigalli and Bonanno, 1997).

$(\neg B_{i,t})(\cdot) \subseteq K_{i,t}(\neg B_{i,t})(\cdot)$ follows from Truth Axiom and Negative Introspection of knowledge. The third captures the idea of belief persistence (Battigalli and Bonanno, 1997): player i believes E at time t iff she believes at t that she (will) believe E at $t + 1$. Player i 's knowledge satisfies *perfect recall* if $K_{i,t}(\cdot) \subseteq K_{i,t+1}(\cdot)$ for all $t \in \mathbb{N}$. A dynamic knowledge-belief space *with perfect recall* is a dynamic knowledge-belief space such that each player's knowledge satisfies perfect recall.

A dynamic knowledge-belief space is mathematically a belief space of $I \times \mathbb{N} \times \{0, 1\}$, where “player $(i, t, 0)$'s belief operator” is $K_{i,t}$ while “player $(i, t, 1)$'s belief operator” is $B_{i,t}$, with the specified conditions. Thus:

Corollary S.2 (Terminal Dynamic Knowledge-Belief Space). *There exists a terminal dynamic κ -knowledge-belief space (with/without perfect recall) $\vec{\Omega}^*$ of I on (S, \mathcal{S}) .*

D.3 Futher Possible Extensions

This subsection briefly discusses three possible extensions: a terminal knowledge-unawareness space, a terminal preference space, and a terminal expectation space.

D.3.1 Terminal Knowledge-Unawareness Space

A *knowledge-unawareness space* is a tuple $\vec{\Omega} := \langle (\Omega, \mathcal{D}), (K_i, U_i)_{i \in I}, \Theta \rangle$ where $K_i : \mathcal{D} \rightarrow \mathcal{D}$ is player i 's knowledge operator and $U_i : \mathcal{D} \rightarrow \mathcal{D}$ is i 's unawareness operator. By Theorem 1, a terminal knowledge-unawareness space $\vec{\Omega}^* := \langle (\Omega^*, \mathcal{D}^*), (K_i^*, U_i^*)_{i \in I}, \Theta^* \rangle$ exists under various assumptions on properties of knowledge and unawareness.

Call a knowledge-unawareness space $\vec{\Omega}$ *non-trivial* if $U_i(\llbracket e \rrbracket_{\vec{\Omega}}) \neq \emptyset$ for some $(i, e) \in I \times \mathcal{L}$. That is, some player i is unaware of some $\llbracket e \rrbracket_{\vec{\Omega}}$ at some state. Then, there exists a non-trivial knowledge-unawareness space within a given category of knowledge-unawareness spaces iff the terminal knowledge-unawareness space $\vec{\Omega}^*$ is non-trivial: $U_i^*([e]) \neq \emptyset$ for some $(i, e) \in I \times \mathcal{L}$.

Since the literature on unawareness has demonstrated some limitation of possibility correspondence models on standard state spaces (e.g., Chen, Ely, and Luo, 2012; Dekel, Lipman, and Rustichini, 1998; Fukuda, 2021; Modica and Rustichini, 1994), an interesting research avenue is to extend the construction of a terminal knowledge-unawareness space to a generalized state space consisting of multiple state spaces as in Heifetz, Meier, and Schipper (2006, 2008). The key insight of the specification of a domain (i.e., the collection of events) on a knowledge-unawareness space carries over to the generalized state space model.

Specifically, while the framework of this paper requires the domain \mathcal{D} to be a κ -algebra of sets, the domain in the generalized state space has a more general lattice structure. A (generalized) knowledge-unawareness space would refer to a tuple $\vec{\Omega} := \langle (\Omega, \mathcal{D}), (K_i, U_i)_{i \in I}, \Theta \rangle$ with the following properties: $\Omega = \bigcup_{\alpha \in A} \Omega_\alpha$ is a state space that consists of multiple subspaces $(\Omega_\alpha)_{\alpha \in A}$; \mathcal{D} forms a κ -complete lattice; $K_i, U_i : \Omega \rightarrow \mathcal{D}$ are κ -continuous functions.

$\mathcal{D} \rightarrow \mathcal{D}$ are player i 's knowledge and unawareness operators; and $\Theta : \Omega \rightarrow S$ is a map that associates, with each state of the world, the corresponding nature state. I conjecture that the idea of this paper can be applied to knowledge and unawareness operators on a κ -complete lattice.⁹

D.3.2 Terminal Preference Space

A preference space refers to a model in which players reason about their interactive preferences instead of beliefs (e.g., Chen, 2010; Di Tillio, 2008; Epstein and Wang, 1996; Ganguli, Hiefetz, and Lee, 2016). In a preference space $\langle(\Omega, \mathcal{D}), (m_i)_{i \in I}, \Theta\rangle$, each player i 's preference-type mapping m_i associates, with each state of the world $\omega \in \Omega$, her preference relation over the set of acts (i.e., bounded measurable functions) on Ω , where Ω is endowed with a κ -algebra. This contrasts with a belief space $\langle(\Omega, \mathcal{D}), (m_{B_i})_{i \in I}, \Theta\rangle$, where each player i 's type mapping m_{B_i} associates, with each state of the world $\omega \in \Omega$, her belief over the set of events. Call a preference space a κ -preference space if it is defined on a κ -algebra.

On the one hand, the framework of this paper may not directly apply to the existence of a terminal preference space, as its construction would call for the hierarchical construction (directly constructing a set of hierarchies of preferences) as opposed to the belief-operator construction of this paper.¹⁰

On the other hand, the analyses of this paper suggest two points on the existence and structure of a terminal preference space.¹¹ First, a terminal κ -preference space would exist, irrespective of such property of preferences as “continuity,” if one considers preference hierarchies up to the ordinal level $\bar{\kappa}$. For example, Di Tillio (2008) constructs a terminal \aleph_0 -preference space consisting of finite preference hierarchies in the category of \aleph_0 -preference-type spaces. I conjecture that, in the context of Di Tillio (2008) in which players' preference relations are merely complete and transitive, a terminal \aleph_1 -preference space would also exist if each preference hierarchy consists of all countable-level interactive preferences.

⁹In fact, this conjecture is consistent with the construction of a canonical knowledge-unawareness space by Heifetz, Meier, and Schipper (2008) based on a finitely language. As this paper shows that Aumann (1976)'s canonical knowledge space (which is constructed based on a finitely language) can be taken as a terminal \aleph_0 -knowledge space, the canonical space of Heifetz, Meier, and Schipper (2008) could be taken as a terminal \aleph_0 -knowledge-unawareness space within the (new) class of knowledge-unawareness spaces defined on an \aleph_0 -complete lattice.

¹⁰In fact, the aforementioned papers directly construct a terminal preference space as the space of preference hierarchies.

¹¹The conjectures are partly based on the fact that one can construct a terminal qualitative-belief space as a space of qualitative-belief hierarchies. In the context of κ -belief spaces, an earlier version of this paper (Fukuda, 2017, Sections 5 and 6) (i) provides a type-space reformulation of a κ -qualitative-belief (especially, κ -knowledge) space, where the type mapping of a player associates, with each state, a binary belief (that assigns either 0 or 1 to each event) instead of a probability measure (that assigns $p \in [0, 1]$ to each event); and (ii) constructs a terminal κ -qualitative-belief (especially κ -knowledge) space consisting of qualitative-belief hierarchies of depth up to $\bar{\kappa}$. Such terminal κ -qualitative-belief space exists regardless of properties of beliefs.

Second, under a regularity condition under which players' finite-level reasoning extends to countable levels, a terminal \aleph_1 -preference space would consist of finite-level preference hierarchies in the category of \aleph_1 -preference spaces.¹² As examples, in Epstein and Wang (1996), preferences satisfy some continuity properties (their P3 and P4). In Ganguli, Hiefetz, and Lee (2016), preferences are represented by a countable collection of continuous real-valued functionals over acts.

Moreover, in terms of game-theoretic applications, suppose that one would like to use a preference space to study implications of common belief in rationality, where beliefs and rationality are induced from preferences. Especially, for each event $E \in \mathcal{D}$, player i believes E in the sense of Savage-null: she believes E if her preferences never depend on any outcomes that happen when the event E does not occur.¹³ The resulting belief is qualitative belief for which Section 6 of the main text establishes the epistemic characterization for iterative elimination of strictly dominated actions as common belief in rationality. It is an interesting avenue for future research to identify properties of preferences and the corresponding properties of qualitative belief.

Instead, the rest of this subsection shows that a terminal expectation space exists. At each state of the world, each player has her (numerical) expectation of an act (or a random variable, i.e., a bounded measurable function).

D.3.3 Terminal Expectation Space

I construct a terminal expectation space where players interactively reason about their expectations of random variables by formalizing the correspondence between beliefs and expectations. While I focus on expectations that come from countably-additive beliefs, a terminal expectation space would exist for weaker notions of expectations when the objects of reasoning (e.g., a class of random variables) and the properties of expectations (i.e., additivity or continuity) are modified.¹⁴

To define the objects of players' expectations (i.e., random variables), denote by $\mathcal{B}(\Omega)$ the set of bounded Borel measurable functions $f : \Omega \rightarrow \mathbb{R}$ on a measurable space (Ω, \mathcal{D}) . Next, I define the properties of expectations (that come from countably-additive beliefs). Define the space $\Gamma(\Omega)$ of expectations on a measurable space (Ω, \mathcal{D}) as the subset of the space $\mathbb{R}^{\mathcal{B}(\Omega)}$ of the mappings from $\mathcal{B}(\Omega)$ into \mathbb{R} respecting the following five properties of expectations. Namely, any $J \in \Gamma(\Omega)$ satisfies: (a. Non-negativity) $f(\cdot) \geq 0$ implies $J[f] \geq 0$; (b. Additivity) $J[f + g] = J[f] + J[g]$;

¹²For the case of probabilistic beliefs, Proposition 4 in Section 5.1 corresponds to this assertion.

¹³Morris (1996, 1997) studies qualitative belief and knowledge from preferences. Brandenburger, Friedenberg, and Keisler (2008) provide a preference foundation for the notion of assumption in their lexicographic-probability-system framework. Trost (2013) studies an epistemic characterization of iterated elimination of inferior action profiles (see Appendix E.3) in preference spaces.

¹⁴The construction here solves an open question raised by Golub and Morris (2017, Section 3) on the construction of a terminal expectation space. Corollary S.3 constructs a terminal expectation space by transforming a terminal probabilistic-belief space into the terminal expectation space. Footnote 32 in Appendix G.1 also shows how one can explicitly construct a terminal expectation space by paralleling the construction of a terminal probabilistic-belief space in the literature.

(c. Homogeneity) $J[cf] = cJ[f]$ for all $c \in \mathbb{R}$; (d. Constancy) $J[\mathbb{I}_\Omega] = 1$; and (e. Continuity) $f_n \uparrow f$ (in $\mathcal{B}(\Omega)$) implies $J[f_n] \uparrow J[f]$. Next, let \mathcal{D}_Γ be the \aleph_1 -algebra on $\Gamma(\Omega)$ generated by $\{\{J \in \Gamma(\Omega) \mid J(f) \geq r\} \in \mathcal{P}(\Gamma(\Omega)) \mid (f, r) \in \mathcal{B}(\Omega) \times \mathbb{R}\}$.

An *expectation space* (of I on (S, \mathcal{S})) is a tuple $\overrightarrow{\Omega} := \langle (\Omega, \mathcal{D}), (\mathbb{E}_i)_{i \in I}, \Theta \rangle$ with the following properties: (i) (Ω, \mathcal{D}) is a measurable space of states of the world; (ii) $\Theta : (\Omega, \mathcal{D}) \rightarrow (S, \mathcal{A}_{\aleph_1}(\mathcal{S}))$ is a measurable map that associates, with each state of the world, the corresponding nature state; and (iii) each \mathbb{E}_i is player i 's expectation-type mapping $\mathbb{E}_i : (\Omega, \mathcal{D}) \rightarrow (\Gamma(\Omega), \mathcal{D}_\Gamma)$ satisfying the introspective property (certainty of expectation) below. For each $\omega \in \Omega$, define the set of states $[\mathbb{E}_i(\omega)] := \{\tilde{\omega} \in \Omega \mid \mathbb{E}_i(\tilde{\omega}) = \mathbb{E}_i(\omega)\}$ at which player i cannot distinguish from ω based on her own expectations. Then, assume that each player i is *certain of her expectation*: for any $(\omega, E) \in \Omega \times \mathcal{D}$, $[\mathbb{E}_i(\omega)] \subseteq E$ implies $\mathbb{E}_i(\omega)[\mathbb{I}_E] = 1$.

I discuss two ways in which an expectation space unpacks players' higher-order expectations. The first is analogous to a type mapping in a probabilistic-belief space as in the above definition. Player i 's expectation-type mapping \mathbb{E}_i associates, with each state $\omega \in \Omega$, the functional $\mathbb{E}_i(\omega)$ that maps each random variable f to the player's expectation of f at ω .

In contrast, the second views the expectation-type mapping as a mapping that associates, with each random variable f , another random variable $\mathbb{E}_i[f]$ that represents the player's expectation of the random variable at each state. Hence, by iterating players' expectations \mathbb{E}_i , one can represent higher-order expectations such as j 's expectation of i 's expectation of a random variable f by $\mathbb{E}_j \mathbb{E}_i[f]$.

Formally, $\mathbb{E}_i : \mathcal{B}(\Omega) \rightarrow \mathcal{B}(\Omega)$ associates, with each bounded Borel measurable function f , another bounded Borel measurable function $\mathbb{E}_i(\cdot)[f]$ that represents the player i 's expectation of f at each state. Denote by $\mathbb{E}_i[f \mid \omega] = \mathbb{E}_i(\omega)[f]$. The five properties on the expectation-type mapping \mathbb{E}_i are: (a. Non-negativity) $f(\cdot) \geq 0$ implies $\mathbb{E}_i[f \mid \cdot] \geq 0$; (b. Additivity) $\mathbb{E}_i[f + g] = \mathbb{E}_i[f] + \mathbb{E}_i[g]$; (c. Homogeneity) $\mathbb{E}_i[cf] = c\mathbb{E}_i[f]$ for all $c \in \mathbb{R}$; (d. Constancy) $\mathbb{E}_i[\mathbb{I}_\Omega] = \mathbb{I}_\Omega$; and (e. Continuity) $f_n \uparrow f$ (in $\mathcal{B}(\Omega)$) implies $\mathbb{E}_i[f_n] \uparrow \mathbb{E}_i[f]$.

To see interactive reasoning over random variables on $(S, \mathcal{A}_{\aleph_1}(\mathcal{S}))$, for any $f \in \mathcal{B}(S)$, denote $f_\Theta = f \circ \Theta \in \mathcal{B}(\Omega)$. Player i 's expectation of f on the state space (Ω, \mathcal{D}) is $\mathbb{E}_i[f_\Theta] \in \mathcal{B}(\Omega)$. Thus, one can analyze players' higher-order expectations by iterating their expectation-type mappings. For example, player i 's expectation of j 's expectation of $f \in \mathcal{B}(S)$ is a bounded Borel measurable function $\mathbb{E}_i \mathbb{E}_j[f_\Theta] \in \mathcal{B}(\Omega)$.

Next, I also remark that if player i is certain of her expectations then her expectations satisfy the law of iterated expectations. Again, denoting $\mathbb{E}_i[f \mid \omega] = \mathbb{E}_i(\omega)[f]$, it can be seen that $\mathbb{E}_i[\mathbb{E}_i[f \mid \tilde{\omega}] \mid \omega] = \mathbb{E}_i[f \mid \omega]$ for all $(\omega, f) \in \Omega \times \mathcal{B}(\Omega)$.

An (*expectation*) *morphism* φ between expectation spaces $\overrightarrow{\Omega}$ and $\overrightarrow{\Omega}'$ is a measurable map $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$ with the following two properties: (i) $\Theta = \Theta' \circ \varphi$; and (ii) $\mathbb{E}_i(\omega)[f' \circ \varphi] = \mathbb{E}'_i(\varphi(\omega))[f']$ for all $(\omega, f') \in \Omega \times \mathcal{B}(\Omega')$. Call a morphism φ an (*expectation*) *isomorphism* if φ is bijective and its inverse φ^{-1} is a morphism. The class of expectation spaces forms a category, where each object is an expectation

space and each arrow is an expectation morphism.

An expectation space $\overrightarrow{\Omega}^*$ (of I on (S, \mathcal{S})) is *terminal* (among the class of expectation spaces of I on (S, \mathcal{S})) if, for any expectation space $\overrightarrow{\Omega}$ (of I on (S, \mathcal{S})), there is a unique morphism φ from $\overrightarrow{\Omega}$ into $\overrightarrow{\Omega}^*$. A terminal expectation space is unique up to isomorphism. Then:

Corollary S.3 (Terminal Expectation Space). *There exists a terminal expectation space $\overrightarrow{\Omega}^*$ of I on (S, \mathcal{S}) .*

To construct a terminal expectation space, the proof in Appendix G.1 formulates the equivalence between expectation spaces and probabilistic-belief spaces that comes from the one-to-one correspondence between beliefs and expectations. Especially, the terminal expectation space is constructed from a terminal probabilistic-belief space.

I briefly remark on expectations that come from finitely-additive or non-additive beliefs. When players possess expectations that come from finitely-additive or non-additive beliefs on a κ -algebra, they interactively reason about their expectations of bounded measurable functions $f : (\Omega, \mathcal{D}) \rightarrow (\mathbb{R}, \mathcal{A}_\kappa(\{[a, b] \mid a < b\}))$, and their expectation-type mappings would satisfy the appropriate properties of the expectation functional.

Next, I remark on the average expectation operator $\overline{\mathbb{E}} : (\Omega, \mathcal{D}) \rightarrow (\Gamma(\Omega), \mathcal{D}_\Gamma)$, which associates, with each state ω , the weighted average of the players' expectations (for simplicity, I focus on symmetric weights on players).¹⁵ Consider the following two typical cases of a set I of players. The first case is when the set of players is finite: $I = \{1, \dots, n\}$. Then, define $\overline{\mathbb{E}}(\omega)[f] := \frac{1}{n} \sum_{i=1}^n \mathbb{E}_i(\omega)[f]$ for each $(\omega, f) \in \Omega \times \mathcal{B}(\Omega)$. The second case is when $I = [0, 1]$. Assume that, for each $(\omega, f) \in \Omega \times \mathcal{B}(\Omega)$, the mapping $I \ni i \mapsto \mathbb{E}_i(\omega)[f] \in \mathbb{R}$ is Borel-measurable (by construction, it is bounded: $\sup_{i \in I} |\mathbb{E}_i(\omega)[f]| \leq \sup_{\tilde{\omega} \in \Omega} |f(\tilde{\omega})| < \infty$). Hence, the following mapping $\overline{\mathbb{E}} : \Omega \times \mathcal{B}(\Omega) \rightarrow \mathbb{R}$ is well-defined: $\overline{\mathbb{E}}(\omega)[f] := \int_I \mathbb{E}_i(\omega)[f] di$ for each $(\omega, f) \in \Omega \times \mathcal{B}(\Omega)$. It can be seen that $\overline{\mathbb{E}} : (\Omega, \mathcal{D}) \rightarrow (\Gamma(\Omega), \mathcal{D}_\Gamma)$ is measurable. Especially, each agent can reason about the average expectations, because $\overline{\mathbb{E}}$ maps $f \in \mathcal{B}(\Omega)$ to $\overline{\mathbb{E}}[f] \in \mathcal{B}(\Omega)$. Since the mapping $I \ni i \mapsto \mathbb{E}_i^*(\omega^*)[f] \in \mathbb{R}$ is Borel-measurable and bounded, the average expectation operator $\overline{\mathbb{E}}^* : (\Omega^*, \mathcal{D}^*) \rightarrow (\Gamma(\Omega^*), \mathcal{D}_\Gamma^*)$ is well-defined in the terminal space.

E Section 6

Appendix E supplements Section 6. Appendix E.1 provides another example of a strategic game in which a unique prediction under IESDA is obtained by an arbitrarily long elimination process. Appendix E.2 provides an example of a strategic game in

¹⁵See Golub and Morris (2017) and the references therein for applications of average expectations to such literature as network games.

which a unique prediction under IEBDA is obtained by an arbitrarily long elimination process. Appendix E.3 studies common knowledge of weak-dominance rationality instead of common belief in weak-dominance rationality.¹⁶ Appendix E.4 studies an epistemic characterization of a (pure-strategy) Nash equilibrium for any strategic game with ordinal payoffs. Appendix E.5 focuses on probabilistic beliefs by studying correlated equilibria.

E.1 Another Example of a Game with a Transfinite Process of IESDA

I provide an example of a strategic game with finite action sets in which a unique prediction under IESDA involves an arbitrarily long process because there are infinitely many players.¹⁷

Let α be a non-zero limit ordinal. I define a strategic game $\langle (A_i)_{i \in I}, (u_i)_{i \in I} \rangle$ (in terms of payoff functions) as follows. First, let $I := \alpha + 1$ (i.e., $I = \{0, 1, \dots, \alpha\}$). For instance, when α is the smallest infinite ordinal (i.e., the set of non-negative integers), the set of players can be identified with $\{0, 1, 2, \dots\} \cup \{i^*\}$ with $0 < 1 < 2 < \dots < i^*$.¹⁸

Second, each player i 's action set is $A_i := \{x, y\}$. Third, for player $0 \in I$, her payoff is

$$u_0(a) := \begin{cases} 0 & \text{if } a_0 = x \\ 1 & \text{if } a_0 = y \end{cases}.$$

For any other player $i \in I \setminus \{0\}$, her payoff from taking $a_i = y$ is always 1 and her payoff from taking $a_i = x$ is

$$u_i(x, a_{-i}) := \begin{cases} 2 & \text{if } a_j = x \text{ for some } j < i \\ 0 & \text{if } a_j = y \text{ for all } j < i \end{cases}.$$

That is, for each player i , taking $a_i = y$ always yields a payoff of 1; in contrast, taking $a_i = x$ yields a payoff of 2 if some predecessor $j < i$ takes action x and otherwise it yields a payoff of 0. For player 0, taking $a_0 = 0$ thus always yields a payoff of 0.

For this game, the unique prediction under IESDA is the singleton set $A^{\text{IESDA}} = \{(y)_{i \in I}\}$ in which every player's action is y . In the original game, the unique strictly dominated action is $a_0 = x$. Once $a_0 = x$ is eliminated, the unique strictly dominated action is $a_1 = x$. At the $(n + 1)$ -th step of elimination, $a_n = x$ is the unique strictly

¹⁶Other possible applications for strategic games with qualitative beliefs are solution concepts by Guarino and Ziegler (2022) under which players follow maxmin (pessimism) or maxmax (optimism) decision criteria. For optimism (maxmax criteria), they study the role of Truth Axiom. For pessimism (maxmin criteria), they propose a new rationalizability solution concept which they call Wald rationalizability and compare it with Börgers rationalizability.

¹⁷This example is inspired by Dufwenberg and Stegeman (2002, Example 6).

¹⁸In fact, for any set I , by the Axiom of Choice, one can introduce a well-ordering on I so that I is order-isomorphic to some ordinal α . Then, add i^* to the set I as the largest element.

	0	1	2	3	4	α
0	0	0	0	0		0
1	1	0	0	0		0
2	1	1	0	0		0
3	1	1	1	0	0	0
4	1	1	1	1	0	0
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
α	0	0	0	0	0	1

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

dominated action and thus is eliminated. At the ϖ -th step (where ϖ denotes the least infinite ordinal), $a_n = x$ has been eliminated for all $n \in \varpi$. At the $(\varpi + 1)$ -th step, $a_\varpi = x$ is eliminated (in the special case in which $\alpha = \varpi$, $a_{i^*} = x$ is eliminated). In this way, $a_i = x$ is eliminated for all $i \in I$.

E.2 An Example of a Game with a Transfinite Process of IEBA

Similarly to the case of IESDA, I provide an example of a strategic game in which a unique prediction under IEBA is obtained after arbitrarily long iterations.

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 $A_i := \alpha + 1$ (i.e., $A_i = \{0, 1, \dots, \alpha\}$) 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

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

Table S.1 depicts $u_i(a_i, a_{-i})$ as a function of a_i (row) and a_{-i} (column). Action α yields a payoff of 1 if the opponent's action is also α , and it yields a payoff of 0 otherwise. Any other action a_i yields a payoff of 1 if it is larger than the opponent's action a_{-i} , and it yields a payoff of 0 otherwise. The process of eliminating all B-dominated actions at each step leads to the unique prediction (α, α) after the α round of elimination.¹⁹

¹⁹Two remarks are in order. First, since IEBA is order-independent, any process of IEBA yields a unique action profile $(a_1, a_2) = (\alpha, \alpha)$. Second, it can be seen that no action is strictly dominated in the original game, and thus $A^{\text{IESDA}} = A$.

E.3 Iterated Elimination of Inferior Action Profiles

I study a pure-strategy version of the iterated elimination procedure of “inferior” action profiles first introduced by Stalnaker (1994) and further studied, among others, by Bonanno (2008), Bonanno and Tsakas (2018), and Hillas and Samet (2020). An epistemic characterization of this solution concept calls for common knowledge of weak-dominance rationality instead of common belief in weak-dominance rationality. First, the following defines the notion of an inferior action profile.

Definition S.3 (Inferior Action Profiles). Let $\langle (A_i)_{i \in I}, (\succ_i)_{i \in I} \rangle$ be a strategic game. Let X be a subset of the action profiles A (X may not necessarily have a product structure). An action profile $x \in X$ is *inferior relative to X* if there exist a player $i \in I$ and an action $a_i \in A_i$ of player i (with a_i not necessarily belonging to the projection of X into A_i) with the following two properties: (i) $(a_i, x_{-i}) \succ_i (x_i, x_{-i})$; and (ii) $(a_i, a_{-i}) \succ_i (x_i, a_{-i})$ for any $a_{-i} \in A_{-i}$ with $(x_i, a_{-i}) \in X$.

The following defines the iterated elimination procedure of inferior action profiles.

Definition S.4 (IEIP). A process of *iterated elimination of inferior action profiles* (IEIP) 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* inferior action profile $a \in A^\beta$ relative to A^β ; and (iii) for a non-zero limit ordinal α , $A^\alpha = \bigcap_{\beta < \alpha} A^\beta$. Since $(A^\alpha)_\alpha$ is weakly decreasing, take the smallest ordinal α (with $|\alpha| \leq |A|$) with $A^\alpha = A^{\alpha+1}$. An action profile $a \in A$ *survives* the process of IEIP if $a \in A^{\text{IEIP}} := A^\alpha$. Call A^{IEIP} the *terminal set* of the process of IEIP.

A process of IEIP is order-independent, i.e., the terminal set A^{IEIP} is uniquely determined. The proof of this result is a minor modification to Hillas and Samet (2020, Proposition 1).

The following proposition demonstrates that the solution concepts of IEIP and IEBA shed light on the difference between knowledge and belief. Common knowledge of weak-dominance rationality characterizes IEIP, while common belief in weak-dominance rationality characterizes IEBA. Bonanno and Tsakas (2018) and Hillas and Samet (2020), among others, provide the epistemic characterization of IEIP as an implication of common knowledge of weak-dominance rationality when every player’s belief (in fact, knowledge) operator satisfies Truth Axiom and the Kripke property.²⁰ Technically, the following extends the epistemic characterization of IEIP by Bonanno and Tsakas (2018) and Hillas and Samet (2020) (for a finite strategic game) to any strategic game.

Proposition S.1 (IEIP). *Fix a strategic game and $\kappa > \max(|I|, |A|)$.*

1. *Take any κ -knowledge space $\overrightarrow{\Omega}$ in which each B_i satisfies the Kripke property and Truth Axiom and in which each player is certain of (in fact, knows) her*

²⁰Note that Truth Axiom implies Consistency.

own strategy: $\Theta_i^{-1}(\{a_i\}) \subseteq B_i(\Theta_i^{-1}(\{a_i\}))$.²¹ If $\omega \in C(\text{WDRAT}_I)$ then $\Theta(\omega) \in A^{\text{IEIP}}$.

2. For any $a \in A^{\text{IEIP}}$, there exist a κ -knowledge space $\vec{\Omega}$ and a state $\omega \in \Omega$ such that: each B_i satisfies the Kripke property and Truth Axiom; each player is certain of (in fact, knows) her own strategy; $\Theta(\omega) = a$; and that $\omega \in C(\text{WDRAT}_I)$.

The proof of Proposition S.1 is similar to that of Bonanno and Tsakas (2018, Proposition 1) or Hillas and Samet (2020, Theorem 1). Thus it is omitted.²²

Three remarks are in order. First, since one can prove Proposition S.1 (2) by constructing a partitional possibility correspondence model, Proposition S.1 holds with or without Positive Introspection and Negative Introspection, as long as Truth Axiom and the Kripke property are imposed.

Second, the above observation implies $A^{\text{IEIP}} \subseteq A^{\text{IEBA}}$. For any $a \in A^{\text{IEIP}}$, there exist a partitional model and its state ω with $a = \Theta(\omega)$ and $\omega \in C(\text{WDRAT}_I)$. Then, $a = \Theta(\omega) \in A^{\text{IEBA}}$. The converse set inclusion may not necessarily hold, because the epistemic characterization of IEBA may call for the failure of Truth Axiom.

Third, for the epistemic characterization of IEIP, Truth Axiom may be necessary. In fact, Bonanno (2008) provides a game and a belief space (see his Figures 6 and 7) such that the players have common belief in weak-dominance rationality and the resulting action profile does not survive any process of IEIP.²³

Proposition S.1 is restated as:

Corollary S.4 (IEIP on the Terminal Space). *Fix any strategic game and $\kappa > \max(|I|, |A|)$. There exists a terminal κ -knowledge space $\vec{\Omega}^*$ in which each B_i satisfies the Kripke property and Truth Axiom and in which every player is certain of (in fact, knows) her own strategy. Then,*

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

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

²¹I simply denote by B_i player i 's belief operator that satisfies Truth Axiom (i.e., knowledge operator).

²²For the first part, one needs to extend the induction proof of Bonanno and Tsakas (2018, Proposition 1) or Hillas and Samet (2020, Theorem 1) to transfinite induction. The second part of the proof is basically identical to that of Bonanno and Tsakas (2018, Proposition 1) or Hillas and Samet (2020, Theorem 1).

²³Moreover, in his example, common belief satisfies Truth Axiom because some player's belief satisfies Truth Axiom.

E.4 Pure-Strategy Nash Equilibria

It is well-known that if every player is rational and is certain of all the players' strategy choices then the resulting play constitutes a Nash equilibrium play (e.g., Aumann and Brandenburger, 1995). This subsection extends this epistemic characterization to an arbitrary strategic game with ordinal payoffs in a class of belief spaces in which each player's belief operator satisfies the Kripke property.

For a given strategic game $\langle (A_i)_{i \in I}, (\succ_i)_{i \in I} \rangle$, denote by A^{Nash} the set of Nash equilibria:

$$A^{\text{Nash}} := \{a \in A \mid \text{for all } i \in I, (a_i, a_{-i}) \succ_i (a'_i, a_{-i}) \text{ for all } a'_i \in A_i\}.$$

As in Section 6.3, for any belief space $\vec{\Omega}$, let RAT_i be the set of states at which player i is rational:

$$\text{RAT}_i := \{\omega \in \Omega \mid \omega \in B_i(\llbracket a_i \succ_i \Theta_i(\omega) \rrbracket_{\vec{\Omega}}) \text{ for no } a_i \in A_i\}.$$

Note that, in order to define RAT_i (precisely, in order for $\llbracket a_i \succ_i \Theta_i(\omega) \rrbracket_{\vec{\Omega}}$ to be an event), the given κ -belief space $\vec{\Omega}$ is assumed to satisfy $\kappa > \max(|A|, |I|)$ as in Section 6.2. Let $\text{RAT}_I := \bigcap_{i \in I} \text{RAT}_i$. Let STR_i be the set of states at which player i is certain of the players' strategies:

$$\text{STR}_i := \{\omega \in \Omega \mid \Theta^{-1}(\{\Theta(\omega)\}) \subseteq B_i(\Theta^{-1}(\{\Theta(\omega)\}))\}.$$

Let $\text{STR}_I := \bigcap_{i \in I} \text{STR}_i$. The next proposition shows that rationality and certainty of the players' strategies characterize Nash equilibrium plays.

Proposition S.2 (Nash Equilibria). *Fix any strategic game and $\kappa > \max(|I|, |A|)$.*

1. *Let $\langle (\Omega, \mathcal{D}), (B_i)_{i \in I}, \Theta \rangle$ be a κ -belief space such that each B_i satisfies the Kripke property. Then, $\omega \in \text{RAT}_I \cap \text{STR}_I$ implies $\Theta(\omega) \in A^{\text{Nash}}$.*
2. *For any $a \in A^{\text{Nash}}$, there exist a κ -belief space $\langle (\Omega, \mathcal{D}), (B_i)_{i \in I}, \Theta \rangle$ and $\omega \in \Omega$ such that: each B_i satisfies the Kripke property; $\omega \in \text{RAT}_I \cap \text{STR}_I$; and that $a = \Theta(\omega)$.*

Two remarks are in order. First, one can show $A^{\text{Nash}} \subseteq A^{\text{IESDA}}$ using the epistemic characterizations. For any $a \in A^{\text{Nash}}$, the proof of Proposition S.2 (2) constructs a belief space $\vec{\Omega}$ in which $\omega \in \Omega = \text{RAT}_I \cap \text{STR}_I$, $a = \Theta(\omega)$, and each B_i satisfies Truth Axiom. Then, by Proposition 5 (1), $\omega \in C(\text{RAT}_I) \subseteq \text{RAT}_I$ and $a = \Theta(\omega) \in A^{\text{IESDA}}$.

In Proposition S.2, the players have to possess the logical ability in the sense that their belief operators have to satisfy the Kripke property. For instance, one can construct a belief space in which the players cooperate in the prisoners' dilemma game even though they are rational and are certain of their strategies due to the failure of

Necessitation (e.g., they fail to believe, at some state, a tautology that cooperation is strictly dominated. At that state, taking cooperation is rational).²⁴

Proposition S.2 is re-stated as:

Corollary S.5 (Nash Equilibria 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. Then,*

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

E.5 Correlated Equilibria

Here, I discuss the role of the existence and structure of a terminal probabilistic-belief space on the solution concept of correlated equilibria. Section E.5.1 introduces, to a belief space, the players' common prior in addition to their posteriors (i.e., type mappings). It shows that a terminal space exists among such a class of belief spaces. Section E.5.2 shows that a correlated equilibrium is mapped to a subspace of the terminal probabilistic-belief space with a common prior if the correlated equilibrium as a belief space is non-redundant and minimal. Section E.5.3 then shows that, in particular, a correlated equilibrium in which the players are Bayes rational at every state (and thus they commonly believe Bayes rationality) can also be mapped to a subspace of the terminal probabilistic-belief space with a common prior and Bayes rationality under non-redundancy and minimality.

Thus, as Aumann (1987) argues, it is possible to take the underlying state space of the correlated equilibrium (in which the players are Bayes rational) as expressing the players' actions, their beliefs about their actions, their beliefs about their beliefs about their actions, and so on.

This subsection demonstrates the flexibility of the methodology of this paper. One can consider a certain class of belief spaces and then establish a terminal belief space among it (e.g., the class of belief spaces that admit a common prior and the class of correlated equilibria as belief spaces).²⁵ The existence of a terminal correlated equilibrium suggests that the underlying states in a correlated equilibrium can be replaced with the players' belief hierarchies about their play if the correlated equilibrium is non-redundant and minimal. Section E.5.4 briefly discusses the literature on the “intrinsic” view of correlation in games.

²⁴For the prisoners' dilemma, denote $A_i = \{c, d\}$ and $(d, c) \succ_i (c, c) \succ_i (d, d) \succ_i (c, d)$ for each $i \in I = \{1, 2\}$. Consider the following belief space: $(\Omega, \mathcal{D}) = (\{\omega_1, \omega_2, \omega_3\}, \mathcal{P}(\Omega))$; $B_i(E) = E$ for any $E \neq \Omega$ and $B_i(\Omega) = \{\omega_1, \omega_2\}$; and $(\Theta_i(\omega_\ell))_{\ell \in \{1, 2, 3\}} = (c, d, c)$. At ω_3 , each player does not believe a tautology. It can be seen that $\text{RAT}_i = \{\omega_2, \omega_3\}$ and $\text{STR}_i = \Omega$. Then, $\Theta(\omega_3) = (c, c)$, which is not a Nash equilibrium.

²⁵One can also establish a terminal space among a class of knowledge-belief spaces, e.g., a class of knowledge-belief spaces that admit a common prior or a class of knowledge-belief spaces in which knowledge is induced by a possibility correspondence. For the space consideration, I focus on probabilistic beliefs alone.

E.5.1 Terminal Belief Space with a Common Prior

I introduce probabilistic-belief spaces with a common prior, and show that a terminal space exists. A *probabilistic-belief space of I on (S, \mathcal{S}) with a common prior* is a tuple $\vec{\Omega} := \langle (\Omega, \mathcal{D}), (B_i^p)_{(i,p) \in I \times [0,1]}, \Theta, \mu \rangle$ such that: (i) $\langle (\Omega, \mathcal{D}), (B_i^p)_{(i,p) \in I \times [0,1]}, \Theta \rangle$ is a probabilistic-belief space in the sense of Definition 10; and that (ii) μ is a (countably-additive) probability measure, called the *common prior*, satisfying

$$\mu(E) = \int_{\Omega} m_{B_i}(\omega)(E) \mu(d\omega) \text{ for each } (i, E) \in I \times \mathcal{D}. \quad (\text{S.1})$$

Equation (S.1) says the prior probability of E is equal to the expectation of the posterior probabilities $m_{B_i}(\omega)(E)$ with respect to μ (e.g., Mertens and Zamir, 1985).²⁶

For probabilistic-belief spaces with a common prior $\vec{\Omega}$ and $\vec{\Omega}'$, $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}'$ is a *(probabilistic-belief) morphism* if φ is a morphism between probabilistic-belief spaces (i.e., $\Theta = \Theta' \circ \varphi$ and $B_i^p \varphi^{-1} = \varphi^{-1} B_i'^p$) and $\mu' = \mu \circ \varphi^{-1}$.²⁷ The last condition states that the prior probabilities are preserved. A probabilistic-belief space $\vec{\Omega}^*$ of I on (S, \mathcal{S}) with a common prior is *terminal* if, for any probabilistic-belief space $\vec{\Omega}$ of I on (S, \mathcal{S}) with a common prior, there is a unique morphism $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}^*$. As in Section 5.1:

Corollary S.6 (Terminal Probabilistic-Belief Space with a Common Prior). *There exists a terminal probabilistic-belief space $\vec{\Omega}^*$ of I on (S, \mathcal{S}) with a common prior.*

In the proof, I introduce a hypothetical player so that, in each probabilistic-belief space with a common prior, the beliefs of the hypothetical player are given by the common prior. Following the construction of a terminal space in Sections 3 and 5.1, one can introduce the beliefs of the hypothetical player on the candidate terminal space. Since it induces a common prior, the candidate space is indeed terminal.

E.5.2 Correlated Equilibrium and Belief Hierarchies

I introduce correlated equilibria. Since a correlated equilibrium is a particular belief space, I show that it is mapped to a subspace of the terminal probabilistic-belief space

²⁶In a probabilistic-belief space with a common prior $\vec{\Omega}$, Aumann (1976)'s Agreement theorem holds. If $\mu(C^p(\bigcap_{i \in I} \{\omega \in \Omega \mid m_{B_i}(\omega)(E) = r_i\})) > 0$, then $|r_i - r_j| \leq 1 - p$ for all $i, j \in I$: if the event that it is common p -belief that each player i 's belief in E is r_i has positive probability according to the common prior, then the difference between any two players' beliefs is at most $1 - p$ (see footnote 51 of the main text for common p -belief). See Fukuda (2019) for the Agreement theorem on an arbitrary measurable space.

²⁷Let $\vec{\Omega}$ be a probabilistic-belief space with a common prior, and let $\varphi : \langle (\Omega, \mathcal{D}), (B_i)_{i \in I}, \Theta \rangle \rightarrow \langle (\Omega', \mathcal{D}'), (B_i')_{i \in I}, \Theta' \rangle$ be a morphism in the sense of Section 5.1. Letting $\vec{\Omega}'$ with $\mu' := \mu \circ \varphi^{-1}$, $\vec{\Omega}'$ is a probabilistic-belief space with a common prior and $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}'$ is a morphism.

with a common prior. That is, the underlying state space of a correlated equilibrium can be replaced with the set of belief hierarchies that the state space induces.

To define a correlated equilibrium, let $\langle (A_i)_{i \in I}, (u_i)_{i \in I} \rangle$ be an underlying strategic game with the following properties. The set A_i of i 's actions is endowed with a σ -algebra \mathcal{S}_i containing singletons: $\{a_i\} \in \mathcal{S}_i$ for all $a_i \in A_i$.²⁸ Let the action profiles $A := \prod_{i \in I} A_i$ be endowed with the product σ -algebra \mathcal{S} . For ease of exposition, let $u_i : A \rightarrow \mathbb{R}$ be i 's bounded Borel measurable payoff function. Any measurable function $\Theta : (\Omega, \mathcal{D}) \rightarrow (A, \mathcal{S})$ can be decomposed into $\Theta = (\Theta_i)_{i \in I}$ such that each $\Theta_i : (\Omega, \mathcal{D}) \rightarrow (A_i, \mathcal{S}_i)$ is measurable.

A *correlated equilibrium* is an \aleph_1 -belief space $\vec{\Omega} := \langle (\Omega, \mathcal{D}), (B_i^p)_{(i,p) \in I \times [0,1]}, \Theta, \mu \rangle$ with a common prior satisfying the following two properties. The first is the certainty of actions: $\Theta_i^{-1}(\{a_i\}) \subseteq B_i^1(\Theta_i^{-1}(\{a_i\}))$ for each $a_i \in A_i$. Whenever player i takes action a_i , she believes with probability one that she takes a_i . The second is the (ex-ante) optimality condition. For any player i and for any measurable function $\tau_i : (\Omega, \mathcal{D}) \rightarrow (A_i, \mathcal{S}_i)$ with $\tau_i^{-1}(\{a_i\}) \subseteq B_i^1(\tau_i^{-1}(\{a_i\}))$ for each $a_i \in A_i$,

$$\int_{\Omega} u_i(\Theta_i(\omega), \Theta_{-i}(\omega)) \mu(d\omega) \geq \int_{\Omega} u_i(\tau_i(\omega), \Theta_{-i}(\omega)) \mu(d\omega).$$

For ease of exposition, I have defined a correlated equilibrium in terms of probability-one beliefs instead of knowledge induced by a partition or a σ -algebra. When each player's knowledge is induced by her partition, her type mapping (equivalently, her p -belief operators) is given as the Bayes conditional probability measure from the common prior conditional on the partition. Thus, instead of specifying the players' partitions, I specify their p -belief operators (equivalently, their type mappings). As discussed in footnote 25 in this subsection, one can establish a terminal correlated equilibrium with knowledge.

Next, I identify the set of belief hierarchies induced by some state of some correlated equilibrium. That is, I construct a terminal correlated equilibrium, i.e., a terminal \aleph_1 -belief space with a common prior satisfying the two requirements to be a correlated equilibrium: $\vec{\Omega}^* = \langle (\Omega^*, \mathcal{D}^*), (B_i^{*p})_{(i,p) \in I \times [0,1]}, \Theta^*, \mu^* \rangle$ such that, for any correlated equilibrium $\vec{\Omega} = \langle (\Omega, \mathcal{D}), (B_i^p)_{(i,p) \in I \times [0,1]}, \Theta, \mu \rangle$, there is a unique morphism $D : \vec{\Omega} \rightarrow \vec{\Omega}^*$. Indeed, as long as the underlying game admits a correlated equilibrium (e.g., when the underlying game is finite), the terminal correlated equilibrium exists, because the candidate terminal \aleph_1 -belief space with a common prior $\vec{\Omega}^*$ (constructed as in Sections 3 and 5.1) satisfies the two conditions to be a correlated equilibrium: (i) $(\Theta_i^*)^{-1}(\{a_i\}) \subseteq B_i^{*1}((\Theta_i^*)^{-1}(\{a_i\}))$ for each $a_i \in A_i$; and (ii) for any player i and for any measurable function $\tau_i^* : (\Omega^*, \mathcal{D}^*) \rightarrow (A_i, \mathcal{S}_i)$ with

²⁸If the players reason only about whether each player takes each action, then the σ -algebra \mathcal{S}_i is the one generated by the singleton actions. Generally, the action set A_i may have a natural measurable structure (e.g., the Borel σ -algebra) \mathcal{S}_i containing each singleton action.

$(\tau_i^*)^{-1}(\{a_i\}) \subseteq B_i^{*1}((\tau_i^*)^{-1}(\{a_i\}))$ for each $a_i \in A_i$,

$$\int_{\Omega^*} u_i(\Theta_i^*(\omega^*), \Theta_{-i}^*(\omega^*)) \mu^*(d\omega^*) \geq \int_{\Omega^*} u_i(\tau_i^*(\omega^*), \Theta_{-i}^*(\omega^*)) \mu^*(d\omega^*).$$

Each state of the terminal correlated equilibrium is associated with some correlated equilibrium distribution. For instance, take two correlated equilibria with different correlated equilibrium distributions and a state from each equilibrium. Then, the two states are mapped to different states in the terminal space, as different correlated equilibrium distributions induce different belief hierarchies about play.

I remark on two implications. First, whenever there exists a correlated equilibrium (and thus the terminal correlated equilibrium exists), any non-redundant and minimal correlated equilibrium $\vec{\Omega}$ can be embedded into the subspace $\overrightarrow{D(\Omega)}$ of the terminal space.²⁹ Put differently, the correlating device (i.e., the state space) of any non-redundant and minimal correlated equilibrium $\vec{\Omega}$ can be replicated by the space $\overrightarrow{D(\Omega)}$ of belief hierarchies: the correlated equilibria $\vec{\Omega}$ and $\overrightarrow{D(\Omega)}$ are isomorphic as a belief space (i.e., the players' beliefs and strategies and the common prior are preserved), and consequently induce the same correlated equilibrium distribution $\mu \circ \Theta^{-1} = \mu^* \circ (\Theta^*)^{-1}$.

Second, an action profile $a \in A$ is played under some correlated equilibrium, that is,

$$a = \Theta(\omega) \text{ for some state } \omega \in \Omega \text{ in some correlated equilibrium } \vec{\Omega},$$

if and only if it is played under the terminal correlated equilibrium $\vec{\Omega}^*$:

$$a = \Theta^*(\omega^*) \text{ for some state } \omega^* \in \Omega^*.$$

In the special case in which the set A of action profiles is finite, denote by A^{CE} the set of action profiles played with positive probability under some correlated equilibrium:

$$A^{\text{CE}} := \{a \in A \mid \mu(\Theta^{-1}(\{a\})) > 0 \text{ for some correlated equilibrium } \vec{\Omega}\}.$$

Then, A^{CE} is the set of action profiles played with positive probability under the terminal correlated equilibrium $\vec{\Omega}^*$:

$$A^{\text{CE}} = \{a \in A \mid \mu^*((\Theta^*)^{-1}(\{a\})) > 0\}.$$

E.5.3 Correlated Equilibria and Bayes Rationality

A similar analysis carries over to a belief space in which the players commonly believe Bayes rationality (Aumann, 1987).³⁰ A belief space with common belief in Bayes

²⁹Note that $\overrightarrow{D(\Omega)}$ is a belief space defined on the image $D(\Omega)$ of Ω under the description map D . As discussed at the end of Section 3, one can introduce a belief structure on $D(\Omega)$.

³⁰The similar analysis also holds for a refinement of a subjective correlated equilibrium called an a-posteriori equilibrium (Aumann, 1974; Brandenburger and Dekel, 1987).

rationality is an \aleph_1 -belief space $\overrightarrow{\Omega} := \langle (\Omega, \mathcal{D}), (B_i^p)_{(i,p) \in I \times [0,1]}, \Theta, \mu \rangle$ with a common prior satisfying the following two properties. The first is the certainty of actions: $\Theta_i^{-1}(\{a_i\}) \subseteq B_i^1(\Theta_i^{-1}(\{a_i\}))$ for each $a_i \in A_i$. The second is Bayes rationality (at every state): for all $\omega \in \Omega$ and for all $a_i \in A_i$,

$$\int_{\Omega} u_i(\Theta_i(\tilde{\omega}), \Theta_{-i}(\tilde{\omega})) m_{B_i}(\omega)(d\tilde{\omega}) \geq \int_{\Omega} u_i(a_i, \Theta_{-i}(\tilde{\omega})) m_{B_i}(\omega)(d\tilde{\omega}).$$

The players are Bayes rational at every state, and thus they commonly believe (at every state) that they are Bayes rational. It can be seen that a belief space with common belief in Bayes rationality is a correlated equilibrium.

Especially, whenever there exists a belief space with common belief in Bayes rationality, the class of belief spaces with common belief in Bayes rationality admits a terminal space. Now, the correlating device (i.e., the state space) of any non-redundant and minimal belief space with common belief in Bayes rationality $\overrightarrow{\Omega}$ can be replicated by the space $\overrightarrow{D(\Omega)}$ of belief hierarchies.

Thus, as Aumann (1987) argues, one can take the underlying state space of the correlated equilibrium (where the players are Bayes rational) as expressing the players' actions and belief hierarchies about their actions.

E.5.4 Intrinsic Correlation

The analysis in this subsection shows that any non-redundant and minimal correlated equilibrium is embedded into a subspace of the terminal correlated equilibrium, which consists of the players' belief hierarchies about play. Broadly, the analysis is somewhat related to an intrinsic view of correlation in non-cooperative games (Brandenburger and Friedenberg, 2008; Du, 2012). Brandenburger and Friedenberg (2008) formulate an *intrinsic* view of correlation in non-cooperative games, under which players' correlated assessments of their strategy choices come from their correlated assessments about their belief hierarchies (about their strategy choices). This is in contrast to the extrinsic view of correlation as external payoff-irrelevant signals (Aumann, 1974). While the analysis of intrinsic correlation is beyond the scope of this paper, I briefly discuss the literature.

The two aforementioned ingenious papers show that a certain correlated rationalizable (or correlated equilibrium) play cannot be played under intrinsic correlation. In contrast, Appendix E.5 asks a related but different question. Specifically, Brandenburger and Friedenberg (2008) study a refinement of correlated (and also independent) rationalizability by imposing epistemic assumptions on players' (i) beliefs about their belief hierarchies and (ii) rationality (and common belief in rationality) in a type space, and study the given strategic game as originally. In contrast, the analysis here takes an extended game (the given strategic game augmented with a correlation device—recall that a correlated equilibrium is a Nash equilibrium of the

extended game), and asks when the correlation device (in a correlated equilibrium) can be replaced with belief hierarchies about play.

Du (2012) characterizes a refinement of correlated equilibria in which each player's strategy in a type space is constant whenever types induce the same belief hierarchy about play. In contrast, this paper asks when the correlation device (the underlying state space) can be replaced with belief hierarchies without imposing the condition that each player's strategy is determined by her belief hierarchy alone.

F Section 7: Minimality

As discussed in the main text, a belief space is minimal (or strongly measurable) if the domain \mathcal{D} consists solely of events that are generated by nature states and belief hierarchies. Proposition S.3 below characterizes minimality as in Friedenberg and Meier (2011, Theorem 5.1).

Recalling Remark 6 and its discussions, any belief morphism φ preserves descriptions: the states ω and $\varphi(\omega)$ induce the same nature state and players' belief hierarchies, i.e., $D_{\vec{\Omega}}(\omega) = D_{\vec{\Omega}'}(\varphi(\omega))$. By following the insightful notion of Friedenberg and Meier (2011), Proposition S.3 instead directly considers a map φ with the property that the description of (i.e., the nature state and players' belief hierarchies at) a state ω is associated with the description of $\varphi(\omega)$: $D_{\vec{\Omega}} = D_{\vec{\Omega}'} \circ \varphi$. Friedenberg and Meier (2011) call such φ a *hierarchy morphism*, as the analysts often directly work with a mapping that preserves players' belief hierarchies rather than with a belief morphism. Roughly, the proposition below states that $\vec{\Omega}'$ is minimal iff a hierarchy morphism φ is a belief morphism. Thus, this section shows that the insight of Friedenberg and Meier (2011) carries over to qualitative belief.

Proposition S.3 (Minimality). *Fix a category of κ -belief spaces of I on (S, \mathcal{S}) .*

1. *Let $\vec{\Omega}$ be a belief space, and let $\vec{\Omega}'$ be a minimal belief space. A measurable map $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$ is a belief morphism iff the map $\varphi : \Omega \rightarrow \Omega'$ satisfies $D_{\vec{\Omega}} = D_{\vec{\Omega}'} \circ \varphi$.*
2. *Let a belief space $\vec{\Omega}'$ be such that, for any belief space $\vec{\Omega}$, a measurable map $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$ is a belief morphism iff the map $\varphi : \Omega \rightarrow \Omega'$ satisfies $D_{\vec{\Omega}} = D_{\vec{\Omega}'} \circ \varphi$. Then, $\vec{\Omega}'$ is minimal.*

To shed more light on conditions under which a hierarchy morphism φ is a belief morphism, I introduce the following two definitions that generalize the notion of non-redundancy. First, following Friedenberg and Meier (2011, Definition 7.2), a κ -belief space $\vec{\Omega}$ of I on (S, \mathcal{S}) is *measurably non-redundant* if, for any $(\omega, E) \in \Omega \times \mathcal{D}$,

$$\text{either } D_{\vec{\Omega}}^{-1}(\{D_{\vec{\Omega}}(\omega)\}) \subseteq E \text{ or } D_{\vec{\Omega}}^{-1}(\{D_{\vec{\Omega}}(\omega)\}) \cap E = \emptyset.$$

A belief space is measurably non-redundant if two states which induce the same description (i.e., players' belief hierarchies) cannot be separated by events in \mathcal{D} . If the belief space $\vec{\Omega}$ is non-redundant (i.e., $D_{\vec{\Omega}}$ is injective), then it is measurably non-redundant.

Second, following Brandenburger and Friedenberg (2008, Definition 8.2) and Friedenberg and Meier (2011, Definition 7.2), a κ -belief space $\vec{\Omega}$ of I on (S, \mathcal{S}) is *bimeasurable* if the description map $D_{\vec{\Omega}} : (\Omega, \mathcal{D}) \rightarrow (\Omega^*, \mathcal{D}^*)$ satisfies

$$D_{\vec{\Omega}}(E) \in \mathcal{D}^* \text{ for all } E \in \mathcal{D}$$

in addition to the measurability condition of $D_{\vec{\Omega}}$ (i.e., $D_{\vec{\Omega}}^{-1}([e]) \in \mathcal{D}$ for all $[e] \in \mathcal{D}^*$). This condition generalizes non-redundancy under topological assumptions.³¹ With these two definitions in mind, the following remark provides a sufficient condition for minimality (Friedenberg and Meier, 2011, Lemma 7.1):

Remark S.1 (Sufficient Condition for Minimality). 1. A minimal κ -belief space is measurably non-redundant.

2. A measurably non-redundant and bimeasurable κ -belief space is minimal.

Proposition S.3 and Remark S.1 imply the following (Friedenberg and Meier, 2011, part of Corollary 7.2). Let $\vec{\Omega}$ and $\vec{\Omega}'$ be a κ -belief space in a given category. Under one of the following conditions, a measurable map $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$ is a belief morphism iff the map $\varphi : \Omega \rightarrow \Omega'$ satisfies $D_{\vec{\Omega}} = D_{\vec{\Omega}'} \circ \varphi$ (i.e., φ is a hierarchy morphism).

1. Proposition S.3 (1): $\vec{\Omega}'$ is minimal.
2. Remark S.1 (2): $\vec{\Omega}'$ is measurably non-redundant and bimeasurable.

G Proofs

G.1 Appendix D

Proof of Corollary S.1. Construct Ω^* as in the proof of Theorem 1, by viewing the set of players in each conditional-belief space as $\bar{I} := I \times [0, 1] \times \mathcal{C}_S$. To see that Ω^* is not empty, take a CPS μ on $(S, \mathcal{A}_{\aleph_1}(\mathcal{S}), \mathcal{C}_S)$. Consider $\langle (S, \mathcal{A}_{\aleph_1}(\mathcal{S}), \mathcal{C}_S), (B_{i,C}^p)_{(i,p,C) \in \bar{I}}, \text{id}_S \rangle$, where (i) $B_{i,C}^p(E) := \emptyset$ if $\mu(E|C) < p$; and (ii) $B_{i,C}^p(E) := S$ if $\mu(E|C) \geq p$.

³¹For the expert reader: In the context of type spaces in which players' probabilistic beliefs are countably additive, if the set of states of nature (S, \mathcal{S}) has a suitable topological structure, then any belief-closed subspace of a terminal type space is bimeasurable. Consistently with Remark S.1 below, any non-redundant type space is minimal and thus can be identified with a belief-closed subspace of the terminal type space (Mertens and Zamir, 1985).

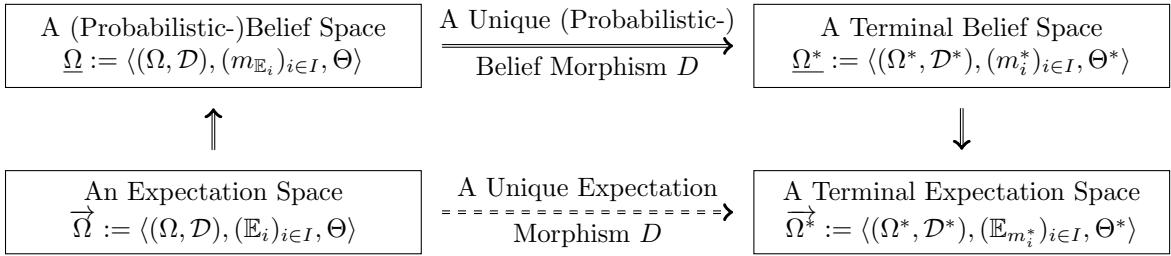


Figure S.6: The Third Step of the Proof of Corollary S.3

Next, as in the proof of Theorem 1, define \mathcal{D}^* , Θ^* , and an auxiliary collection of p -belief operators $(B_{i,C_S}^{*p})_{(i,p,C_S) \in \bar{I}}$ as $B_{i,C_S}^{*p}([e]) := [\beta_{i,C_S}^p(e)]$ for each $[e] \in \mathcal{D}^*$. By construction, $D^{-1}(B_{i,C_S}^{*p}([e])) = B_{i,C_S}^p(D^{-1}[e])$. Since $(\Theta^*)^{-1}(C_S) = [C_S] \in \mathcal{D}^*$, let $\mathcal{C}^* := \{[C_S] \in \mathcal{D}^* \mid C_S \in \mathcal{A}_{\aleph_1}(\mathcal{S})\}$. By construction, $\mathcal{C}^* \subseteq \mathcal{D}^*$, $(\Theta^*)^{-1}(\mathcal{C}_S) = \mathcal{C}^*$, and $\emptyset \notin \mathcal{C}^*$ (this is because Θ^* is surjective). Then, $(\Omega^*, \mathcal{D}^*, \mathcal{C}^*)$ is a conditional space, and $(B_{i,C_S}^{*p})_{(i,p,C_S) \in \bar{I}}$ is a well-defined collection of p -belief operators (observe $B_i^{*p}(\cdot | [C_S]) = B_{i,C_S}^{*p}$). As in the proof of Corollary 2, the p -belief operators satisfy the specified properties, i.e., $\overrightarrow{\Omega}^* := \langle (\Omega^*, \mathcal{D}^*, \mathcal{C}^*), (B_i^{*p}(\cdot | [C_S]))_{(i,p,[C_S]) \in I \times [0,1] \times \mathcal{C}^*}, \Theta^* \rangle$ is a conditional-belief space. By construction, $\overrightarrow{\Omega}^*$ is terminal. \square

Proof of Corollary S.3. The proof consists of three steps. The first step shows that any expectation space induces the corresponding (probabilistic-)belief space and that any belief space induces the corresponding expectation space. The second step shows that any expectation morphism induces the corresponding (probabilistic-)belief morphism and that any belief morphism induces the corresponding expectation morphism. These two steps establish the equivalence between an expectation space and a belief space. The third step then shows that the expectation space induced by a terminal belief space is a terminal expectation space. Figure S.6 illustrates the third step.

Step 1. The first step establishes the correspondence between expectation and belief spaces. For ease of notation, throughout the proof, I formulate any probabilistic-belief space using type mappings (recall the discussion after Definition 10 in the main text). First, let $\langle (\Omega, \mathcal{D}), (\mathbb{E}_i)_{i \in I}, \Theta \rangle$ be an expectation space. I define the corresponding probabilistic-belief space $\langle (\Omega, \mathcal{D}), (m_{\mathbb{E}_i})_{i \in I}, \Theta \rangle$: define a measurable map $m_{\mathbb{E}_i} : (\Omega, \mathcal{D}) \rightarrow (\Delta(\Omega), \mathcal{D}_\Delta)$ by $m_{\mathbb{E}_i}(\omega)(E) := \mathbb{E}_i(\omega)[\mathbb{I}_E]$.

Conversely, let $\langle (\Omega, \mathcal{D}), (m_i)_{i \in I}, \Theta \rangle$ be a belief space. I define the corresponding expectation space $\langle (\Omega, \mathcal{D}), (\mathbb{E}_{m_i})_{i \in I}, \Theta \rangle$. For any $f \in \mathcal{B}(\Omega)$, define $\mathbb{E}_{m_i}(\omega)[f] := \int_{\Omega} f(\tilde{\omega}) m_i(\omega)(d\tilde{\omega})$ for each $\omega \in \Omega$. It can be seen that $\mathbb{E}_{m_i} : \mathcal{B}(\Omega) \rightarrow \mathcal{B}(\Omega)$ is a well-defined map satisfying Non-negativity, Additivity, Homogeneity, Constancy, and Continuity. Moreover, $m_i = m_{\mathbb{E}_{m_i}}$ and $\mathbb{E}_i = \mathbb{E}_{m_{\mathbb{E}_i}}$ for every $i \in I$.³²

³²One can formalize this equivalence using category theory (namely, functors Δ and Γ are isomor-

Step 2. The second step establishes the correspondence between expectation and belief morphisms. Let $\varphi : \overrightarrow{\Omega} \rightarrow \overrightarrow{\Omega}'$ be an expectation morphism. Thus, $\mathbb{E}_i(\omega)[f' \circ \varphi] = \mathbb{E}'_i(\varphi(\omega))[f']$ for any $f' \in \mathcal{B}(\Omega')$. Take $\mathbb{I}_{E'}$ with $E' \in \mathcal{D}$. Then, $m_{\mathbb{E}_i}(\omega)(\varphi^{-1}(E')) = m'_{\mathbb{E}'_i}(\varphi(\omega))(E')$ for each $(\omega, E') \in \Omega \times \mathcal{D}'$. Since this condition is equivalent to the one in terms of p -belief operators, the measurable map $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$ is a belief morphism $\varphi : \langle(\Omega, \mathcal{D}), (m_{\mathbb{E}_i})_{i \in I}, \Theta\rangle \rightarrow \langle(\Omega', \mathcal{D}'), (m'_{\mathbb{E}'_i})_{i \in I}, \Theta'\rangle$.

Conversely, let $\varphi : \underline{\Omega} \rightarrow \underline{\Omega}'$ be a belief morphism (to distinguish expectation and belief spaces, I use the underline to indicate a belief space). For any $E' \in \mathcal{D}'$,

$$\mathbb{E}_{m_i}(\omega)[\mathbb{I}_{E'} \circ \varphi] = \mathbb{E}_{m_i}(\omega)[\mathbb{I}_{\varphi^{-1}(E')}] = m_i(\omega)(\varphi^{-1}(E')) = m'_i(\varphi(\omega))(E') = \mathbb{E}'_{m'_i}[\mathbb{I}_{E'}].$$

Using the properties of an expectation type-mapping, $\mathbb{E}_{m_i}(\omega)[f' \circ \varphi] = \mathbb{E}'_{m'_i}(\varphi(\omega))[f']$ for all $f' \in \mathcal{B}(\Omega')$. In other words, the measurable map $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$ is an expectation morphism $\varphi : \langle(\Omega, \mathcal{D}), (\mathbb{E}_{m_i})_{i \in I}, \Theta\rangle \rightarrow \langle(\Omega', \mathcal{D}'), (\mathbb{E}'_{m'_i})_{i \in I}, \Theta'\rangle$.

Step 3. The third step shows that the expectation space $\overrightarrow{\Omega^*} := \langle(\Omega^*, \mathcal{D}^*), (\mathbb{E}_{m_i^*})_{i \in I}, \Theta^*\rangle$ induced by a terminal probabilistic belief space $\underline{\Omega^*} := \langle(\Omega^*, \mathcal{D}^*), (m_i^*)_{i \in I}, \Theta^*\rangle$ is a terminal expectation space.

Take any expectation space $\overrightarrow{\Omega}$. Consider a corresponding belief space $\underline{\Omega} := \langle(\Omega, \mathcal{D}), (m_{\mathbb{E}_i})_{i \in I}, \Theta\rangle$. Since $\underline{\Omega}$ is terminal, there is a (unique) belief morphism $D : \underline{\Omega} \rightarrow \underline{\Omega^*}$. By Step 2, the measurable mapping D is an expectation morphism $\overrightarrow{\Omega}$ into $\overrightarrow{\Omega^*}$. To show that the mapping D is unique, let $\varphi : \overrightarrow{\Omega} \rightarrow \overrightarrow{\Omega^*}$ be an expectation morphism. By Step 2, $\varphi : \underline{\Omega} \rightarrow \underline{\Omega^*}$ is a belief morphism. Since $\underline{\Omega^*}$ is a terminal belief space, it follows that $\varphi = D$. Figure S.6 illustrates Step 3. \square

G.2 Appendix E

Proof of Proposition S.2. 1. It suffices to show that, for each $i \in I$,

$$(\Theta(\omega)) \succcurlyeq_i (a_i, \Theta_{-i}(\omega)) \text{ for all } a_i \in A_i.$$

Suppose to the contrary that there exist $i \in I$ and $a_i \in A_i$ such that

$$(\Theta(\omega)) \not\succcurlyeq_i (a_i, \Theta_{-i}(\omega)), \text{ i.e., } (a_i, \Theta_{-i}(\omega)) \succ_i (\Theta(\omega)).$$

Since $\omega \in \text{STR}_i$ and B_i satisfies the Kripke property, $\Theta(\omega) = \Theta(\tilde{\omega})$ for all $\tilde{\omega} \in b_{B_i}(\omega)$. Thus, $b_{B_i}(\omega) \subseteq \llbracket a_i \succ_i \Theta_i(\omega) \rrbracket_{\overrightarrow{\Omega}}$, i.e., $\omega \in B_i(\llbracket a_i \succ_i \Theta_i(\omega) \rrbracket_{\overrightarrow{\Omega}})$, which contradicts $\omega \in \text{RAT}_i$.

phic). Hence, one could also construct a terminal expectation space by replacing Δ with Γ in such constructions of a terminal probabilistic type space as Brandenburger and Dekel (1993), Heifetz and Samet (1998, Section 5), and Mertens and Zamir (1985).

2. Take $a \in A^{\text{Nash}}$. Let $(\Omega, \mathcal{D}) = (\{a\}, \mathcal{P}(\{a\}))$, $\Theta = \text{id}_{\{a\}}$, and $B_i = \text{id}_{\mathcal{D}}$ for each $i \in I$. By construction, each B_i satisfies the Kripke property, $\text{RAT}_i = \text{STR}_i = \{a\}$, and $\Theta(a) = a$.

□

Proof of Corollary S.6. The proof consists of two steps. The first step introduces a hypothetical player “0” to each probabilistic-belief space $\vec{\Omega}$ with a common prior so that it can be identified as a probabilistic belief space $\langle (\Omega, \mathcal{D}), (B_i^p)_{(i,p) \in (I \cup \{0\}) \times [0,1]}, \Theta \rangle$ satisfying Expression (S.1). Namely, define B_0^p as follows: $B_0^p(E) = \Omega$ if $\mu(E) \geq p$; and $B_0^p(E) = \emptyset$ if $\mu(E) < p$. Hence, player 0 has a state-independent type mapping $m_{B_0}(\omega)(\cdot) = \mu(\cdot)$ for all $\omega \in \Omega$. For any spaces $\vec{\Omega}$ and $\vec{\Omega}'$, $B_0^p \varphi^{-1} = \varphi^{-1} B_0'^p$ is equivalent to $\mu \circ \varphi^{-1} = \mu'$.

The second step constructs a terminal space. Following the construction in Section 3, the p -belief operators $(B_0^{*p})_{p \in [0,1]}$ satisfy $D^{-1} B_0^{*p}([e]) = B_0^p D^{-1}([e])$ for all $[e] \in \mathcal{D}^*$. Define $\mu^* : \mathcal{D}^* \rightarrow [0, 1]$ by $\mu^*([e]) := \sup\{p \in [0, 1] \mid \Omega^* = B_0^{*p}([e])\}$ for each $[e] \in \mathcal{D}^*$. Since the p -belief operators $(B_0^{*p})_{p \in [0,1]}$ satisfy the properties in Definition 10 (2), μ^* is a countably-additive probability measure. Also, $\mu^* = \mu \circ D^{-1}$ follows from $D^{-1} B_0^{*p}(\cdot) = B_0^p D^{-1}(\cdot)$. Thus, μ^* satisfies Equation (S.1) (see footnote 27 in Appendix E.5.2). □

G.3 Appendix F

Proof of Proposition S.3. For both parts, by Remark 6, a belief morphism $\varphi : \vec{\Omega} \rightarrow \vec{\Omega}'$ satisfies $D_{\vec{\Omega}} = D_{\vec{\Omega}'} \circ \varphi$. For Part (1), it suffices to show the “if” part. First, I show that $\varphi : (\Omega, \mathcal{D}) \rightarrow (\Omega', \mathcal{D}')$ is measurable. Since $D_{\vec{\Omega}} = D_{\vec{\Omega}'} \circ \varphi$, $D_{\vec{\Omega}}^{-1} = \varphi^{-1} \circ D_{\vec{\Omega}'}^{-1}$. Since $\vec{\Omega}'$ is minimal, $D_{\vec{\Omega}'}^{-1}(\mathcal{D}') = \mathcal{D}'$. For any $E' \in \mathcal{D}'$, there is $[e] \in \mathcal{D}^*$ with $E' = D_{\vec{\Omega}'}^{-1}([e])$. Then, $\varphi^{-1}(E') = \varphi^{-1}(D_{\vec{\Omega}'}^{-1}([e])) = D_{\vec{\Omega}}^{-1}([e]) \in \mathcal{D}$. Second, $\Theta = \Theta' \circ \varphi$ follows from $D_{\vec{\Omega}} = D_{\vec{\Omega}'} \circ \varphi$. Third, I show $B_i(\varphi^{-1}(\cdot)) = \varphi^{-1}(B'_i(\cdot))$. For any $E' \in \mathcal{D}'$, take $[e] \in \mathcal{D}^*$ with $E' = D_{\vec{\Omega}'}^{-1}([e])$. Then, $B_i(\varphi^{-1}(E')) = B_i D_{\vec{\Omega}}^{-1}([e]) = D_{\vec{\Omega}}^{-1} B_i^*([e]) = \varphi^{-1} D_{\vec{\Omega}'}^{-1}(B_i^*([e])) = \varphi^{-1} B'_i D_{\vec{\Omega}'}^{-1}([e]) = \varphi^{-1} B'_i(E')$.

For Part (2), suppose that $\vec{\Omega}'$ is not minimal. By Remark 7, $\vec{\Omega}'_{\bar{\kappa}}$ is minimal and $\text{id}_{\Omega'}$ satisfies $D_{\vec{\Omega}'} = D_{\vec{\Omega}'_{\bar{\kappa}}} \circ \text{id}_{\Omega'}$. However, $\text{id}_{\Omega'} : (\Omega', \mathcal{D}'_{\bar{\kappa}}) \rightarrow (\Omega', \mathcal{D}')$ is not measurable because $\mathcal{D}'_{\bar{\kappa}} \subsetneq \mathcal{D}'$, and hence it is not a belief morphism. □

Proof of Remark S.1. 1. Take $(\omega, E) \in \Omega \times \mathcal{D}$. Suppose $D_{\vec{\Omega}}^{-1}(\{D_{\vec{\Omega}}(\omega)\}) \cap E \neq \emptyset$.

Since $\vec{\Omega}$ is minimal, there exists $[e] \in \mathcal{D}^*$ such that $E = D_{\vec{\Omega}}^{-1}([e])$. Thus, there is $\tilde{\omega} \in D_{\vec{\Omega}}^{-1}(\{D_{\vec{\Omega}}(\omega)\}) \cap E = D_{\vec{\Omega}}^{-1}(\{D_{\vec{\Omega}}(\omega)\}) \cap D_{\vec{\Omega}}^{-1}([e])$, and hence $D_{\vec{\Omega}}(\tilde{\omega}) = D_{\vec{\Omega}}(\omega) \in [e]$. Now, I show $D_{\vec{\Omega}}^{-1}(\{D_{\vec{\Omega}}(\omega)\}) \subseteq E$. If $\hat{\omega} \in D_{\vec{\Omega}}^{-1}(\{D_{\vec{\Omega}}(\omega)\})$ then $D_{\vec{\Omega}}(\hat{\omega}) = D_{\vec{\Omega}}(\omega) \in [e]$ and thus $\hat{\omega} \in D_{\vec{\Omega}}^{-1}([e]) = E$.

2. Since $D_{\vec{\Omega}}$ is measurable, $D_{\vec{\Omega}}^{-1}(\mathcal{D}^*) \subseteq \mathcal{D}$. Conversely, take $E \in \mathcal{D}$. It suffices to show $E \in D_{\vec{\Omega}}^{-1}(\mathcal{D}^*)$. Since $D_{\vec{\Omega}}$ is bimeasurable, $D_{\vec{\Omega}}(E) = [e]$ for some $[e] \in \mathcal{D}^*$. Operating $D_{\vec{\Omega}}^{-1}$, I have $D_{\vec{\Omega}}^{-1}(D_{\vec{\Omega}}(E)) = D_{\vec{\Omega}}^{-1}([e])$. Thus, to show $E \in D_{\vec{\Omega}}^{-1}(\mathcal{D}^*)$, it suffices to show $E = D_{\vec{\Omega}}^{-1}(D_{\vec{\Omega}}(E))$. Now, $E \subseteq D_{\vec{\Omega}}^{-1}(D_{\vec{\Omega}}(E))$ follows from the definition of the inverse map $D_{\vec{\Omega}}^{-1}$. If $\omega \in D_{\vec{\Omega}}^{-1}(D_{\vec{\Omega}}(E))$ then there is $\tilde{\omega} \in D_{\vec{\Omega}}^{-1}(\{D_{\vec{\Omega}}(\omega)\}) \cap E$. Since $\vec{\Omega}$ is measurably non-redundant, $\omega \in D_{\vec{\Omega}}^{-1}(\{D_{\vec{\Omega}}(\omega)\}) \subseteq E$.

□

References for the Supplementary Appendix

- [1] I. Arieli and R. J. Aumann. “The Logic of Backward Induction”. *J. Econ. Theory* 159 (2015), 443–464.
- [2] R. J. Aumann. “Subjectivity and Correlation in Randomized Strategies”. *J. Math. Econ.* 1 (1974), 67–96.
- [3] R. J. Aumann. “Agreeing to Disagree”. *Ann. Statist.* 4 (1976), 1236–1239.
- [4] R. J. Aumann. “Correlated Equilibrium as an Expression of Bayesian Rationality”. *Econometrica* 55 (1987), 1–18.
- [5] R. J. Aumann and A. Brandenburger. “Epistemic Conditions for Nash Equilibrium”. *Econometrica* 63 (1995), 1161–1180.
- [6] P. Battigalli and G. Bonanno. “The Logic of Belief Persistence”. *Econ. Philos.* 13 (1997), 39–59.
- [7] P. Battigalli and M. Siniscalchi. “Hierarchies of Conditional Beliefs and Interactive Epistemology in Dynamic Games”. *J. Econ. Theory* 88 (1999), 188–230.
- [8] P. Battigalli and M. Siniscalchi. “Strong Belief and Forward Induction Reasoning”. *J. Econ. Theory* 106 (2002), 356–391.
- [9] 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.
- [10] 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.
- [11] A. Brandenburger and E. Dekel. “Rationalizability and Correlated Equilibria”. *Econometrica* 55 (1987), 1391–1402.
- [12] A. Brandenburger and E. Dekel. “Hierarchies of Beliefs and Common Knowledge”. *J. Econ. Theory* 59 (1993), 189–198.
- [13] A. Brandenburger and A. Friedenberg. “Intrinsic Correlation in Games”. *J. Econ. Theory* 141 (2008), 28–67.
- [14] A. Brandenburger, A. Friedenberg, and H. J. Keisler. “Admissibility in Games”. *Econometrica* 76 (2008), 307–352.
- [15] Y.-C. Chen. “Universality of the Epstein-Wang Type Structure”. *Games Econ. Behav.* 68 (2010), 389–402.

- [16] Y.-C. Chen, J. C. Ely, and X. Luo. “Note on Unawareness: Negative Introspection versus AU Introspection (and KU Introspection)”. *Int. J. Game Theory* 41 (2012), 325–329.
- [17] E. Dekel, B. L. Lipman, and A. Rustichini. “Standard State-Space Models Preclude Unawareness”. *Econometrica* 66 (1998), 159–173.
- [18] A. Di Tillio. “Subjective Expected Utility in Games”. *Theor. Econ.* 3 (2008), 287–323.
- [19] A. Di Tillio, J. Halpern, and D. Samet. “Conditional Belief Types”. *Games Econ. Behav.* 87 (2014), 253–268.
- [20] S. Du. “Correlated Equilibrium and Higher Order Beliefs about Play”. *Games Econ. Behav.* 76 (2012), 74–87.
- [21] M. Dufwenberg and M. Stegeman. “Existence and Uniqueness of Maximal Reductions under Iterated Strict Dominance”. *Econometrica* 70 (2002), 2007–2023.
- [22] L. G. Epstein and T. Wang. “Beliefs about Beliefs” without Probabilities”. *Econometrica* 64 (1996), 1343–1373.
- [23] A. Friedenberg and M. Meier. “On the Relationship between Hierarchy and Type Morphisms”. *Econ. Theory* 46 (2011), 377–399.
- [24] A. Friedenberg and M. Meier. “The Context of the Game”. *Econ. Theory* 63 (2017), 347–386.
- [25] S. Fukuda. “The Existence of Universal Knowledge Spaces”. *Essays in the Economics of Information and Epistemology*. Ph.D. Dissertation, the University of California at Berkeley, 2017, 1–113.
- [26] S. Fukuda. “On the Consistency among Prior, Posteriors, and Information Sets (Extended Abstract)”. *Proc. 17th Conference on Theoretical Aspects of Rationality and Knowledge*. Ed. by L. S. Moss. 2019, 189–205.
- [27] S. Fukuda. “Unawareness without AU Introspection”. *J. Math. Econ.* 94 (2021), 102456.
- [28] J. Ganguli, A. Hefetz, and B. S. Lee. “Universal Interactive Preferences”. *J. Econ. Theory* 162 (2016), 237–260.
- [29] B. Golub and S. Morris. “Higher-Order Expectations”. Aug. 2017.
- [30] 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.
- [31] P. Guarino and G. Ziegler. “Optimism and Pessimism in Strategic Interactions under Ignorance”. *Games Econ. Behav.* 136 (2022), 559–585.
- [32] A. Heifetz, M. Meier, and B. C. Schipper. “Interactive Unawareness”. *J. Econ. Theory* 130 (2006), 78–94.
- [33] A. Heifetz, M. Meier, and B. C. Schipper. “A Canonical Model for Interactive Unawareness”. *Games Econ. Behav.* 62 (2008), 304–324.
- [34] A. Heifetz and D. Samet. “Topology-Free Typology of Beliefs”. *J. Econ. Theory* 82 (1998), 324–341.
- [35] Z. Hellman. “A Game with No Bayesian Approximate Equilibria”. *J. Econ. Theory* 153 (2014), 138–151.
- [36] J. Hillas and D. Samet. “Dominance Rationality: A Unified Approach”. *Games Econ. Behav.* 119 (2020), 189–196.
- [37] K. Hrbacek and T. Jech. Introduction to Set Theory. Third Edition. CRC Press, 1999.

- [38] J. F. Mertens and S. Zamir. “Formulation of Bayesian Analysis for Games with Incomplete Information”. *Int. J. Game Theory* 14 (1985), 1–29.
- [39] S. Modica and A. Rustichini. “Awareness and Partitional Information Structures”. *Theory Decis.* 37 (1994), 107–124.
- [40] S. Morris. “The Logic of Belief and Belief Change: A Decision Theoretic Approach”. *J. Econ. Theory* 69 (1996), 1–23.
- [41] S. Morris. “Alternative Definitions of Knowledge”. *Epistemic Logic and the Theory of Games and Decisions*. Ed. by M. Bacharach, L.-A. Gérard-Varet, P. Mongin, and H. S. Shin. Kluwer Academic Publishers, 1997, 217–233.
- [42] R. S. Simon. “Games of Incomplete Information, Ergodic Theory, and the Measurability of Equilibria”. *Isr. J. Math.* 138 (2003), 73–92.
- [43] R. Stalnaker. “On the Evaluation of Solution Concepts”. *Theory Decis.* 37 (1994), 49–73.
- [44] M. Trost. “Epistemic Characterization of Iterated Deletion of Inferior Strategy Profiles in Preference-Based Type Spaces”. *Int. J. Game Theory* 42 (2013), 755–776.
- [45] E. Tsakas. “Epistemic Equivalence of Extended Belief Hierarchies”. *Games Econ. Behav.* 86 (2014), 126–144.