

AWARENESS*

This is a slightly extended version of a chapter prepared for the *Handbook of Logics for Knowledge and Belief* edited by Hans van Ditmarsch, Joseph Y. Halpern, Wiebe van der Hoek, and Barteld Kooi, College Publications, London.

Burkhard C. Schipper[†]

May 6, 2014

Abstract

Unawareness refers to the lack of conception rather than the lack of information. This chapter discusses various epistemic approaches to modeling (un)awareness from computer science and economics that have been developed over the last 25 years. While the focus is on axiomatizations of structures capable of modeling knowledge and propositionally determined awareness, we also discuss structures for modeling probabilistic beliefs and awareness as well as structures for awareness of unawareness. Further topics, such as dynamic awareness, games with unawareness, decision theory under unawareness, and applications are just briefly reviewed.

Contents

1	Introduction	1
2	Preliminary Discussion	3
2.1	Some Properties of Awareness in Natural Languages	3
2.2	Some Properties of Awareness in a Formal Language	3
3	Awareness and Knowledge	5
3.1	Awareness Structures by Fagin and Halpern [1988]	5
3.2	An Impossibility Result on Unawareness in Kripke Frames	12
3.3	Unawareness Frames	15
3.4	Unawareness Structures	20

*I thank Hans van Ditmarsch, Joe Halpern, and Fernando R. Velázquez Quesada for comments on an earlier draft. Moreover, I thank Aviad Heifetz, Sander Heinsalu, Li Jing, Salvatore Modica, and Giacomo Sillari for helpful discussions. Financial support through NSF SES-0647811 is gratefully acknowledged. A part of the paper has been written when the author visited New York University.

[†]Department of Economics, University of California, Davis. Email: bcschipper@ucdavis.edu.

3.5	Generalized Standard Models by Modica and Rustichini [1999]	23
3.6	Product Models by Li [2009]	26
3.7	Synopsis of Other Approaches: Pires [1994], Ewerhart [2001], and Feinberg [2004, 2005, 2012]	30
4	Awareness and Probabilistic Beliefs	31
4.1	Type Spaces with Unawareness	31
4.2	Universal Type Space with Unawareness	33
4.3	Speculative Trade and Agreement	34
5	Awareness of Unawareness	36
5.1	Propositional Quantifiers and Extended Awareness Structures	37
5.2	First-Order Logic with Unawareness of Objects	39
5.3	Neighborhood Semantics and First-Order Logic with Awareness of Objects and Properties	45
5.4	Awareness of Unawareness without Quantifiers	46
6	Synopsis of Extensions	50
6.1	Dynamic Awareness	50
6.2	Games with Unawareness	51
7	Summary	54

1 Introduction

Formalized notions of awareness have been studied both in computer science and economics.¹ In computer science the original motivation was mainly the modeling of agents who suffer from different forms of logical non-omniscience. The aim was to introduce logics that are more suitable than traditional logics for modeling beliefs of humans or machines with limited reasoning capabilities. In economics the motivation is similar but perhaps less ambitious. The goal is to model agents who may not only lack information but also conception. Intuitively, there is a fundamental difference between not knowing that an event obtained and not being able to conceive of that event. Despite such a lack of conception, agents in economics are still assumed to be fully rational in the sense of not making any errors in information processing such as violating introspection of beliefs for events they can conceive. By letting unawareness stand for lack of conception, economists seem to have aimed for a narrow notion of awareness, while computer scientists are more agnostic about the appropriate notion of awareness. Economists and computer scientists seem to have slightly different tastes over formalisms too. While computer scientists are clearly inspired by Kripke structures but formalize aware-

¹For a bibliography, see <http://www.econ.ucdavis.edu/faculty/schipper/unaw.htm>.

ness syntactically, economists seem to prefer purely event-based approaches similar to Aumann structures or Kripke frames as well as Harsanyi type-spaces. This may be due to different uses of those models. Central to economists are applications to game theory in which players in principle may use the model to reason about other players, to reason about other players' reasoning, reasoning about that etc. Consequently, states can be interpreted as "subjective" descriptions by players. In contrast, states in awareness structures by computer scientists are best understood as an outside analyst's description of agents' reasoning. They are typically not "accessible" to the agent themselves. Because of the different emphasis it is perhaps somewhat surprising that some approaches to awareness in computer science turn out to be equivalent in terms of expressivity to the approach taken in economics, especially with the focus on the notion of awareness that has been called "awareness generated by primitive propositions". At a first glance, the name of this notion suggests that it is essentially syntactic because it refers to primitive propositions or atomic formulae, thus presupposing a syntactic formalism. Yet, it also makes clear that syntactic constructions (such as the order of formulae in a conjunction of formulae) do not play a role in determining this notion of awareness. Hence, this notion should be well-suited to be captured with event-based approaches that economists have focused on. Consequently, the literature on awareness can be viewed from at least two angles. On one hand, the structures proposed by economists are equivalent to a subclass of structures proposed by computer scientists (in particular, the ones proposed in the seminal paper by Fagin and Halpern [1988]). On the other hand, the different modeling approaches pursued by economists make some of their structures directly more amenable to applications in game theory and allow us to isolate without further assumptions the effect of unawareness from other forms of logical non-omniscience. Throughout the chapter, I comment on the slightly different perspectives.

We find it useful to repeatedly refer as a backdrop to the following simple example that has been considered previously by Heifetz, Meier, and Schipper [2006]:

Speculative Trade Example. There are two agents, an owner o of a firm, and a potential buyer b of the firm. The status quo value of the firm is \$100 per share. The owner of the firm is aware of a potential lawsuit that reduces the value of the firm by \$20 per share. The owner does not know whether the lawsuit will occur. The buyer is unaware of the lawsuit and the owner knows that the buyer is unaware of the lawsuit. The buyer, however, is aware of a potential innovation that increases the value of the firm by \$20 per share. The buyer does not know whether the innovation will occur. The owner is unaware of the innovation and the buyer knows that the owner is unaware of the innovation.

A question of interest to economists is whether speculative trade between the owner and buyer based on their differences in awareness is possible. Speculative trade is trade purely due to differences in information/awareness. In this example, we may phrase the question as follows: Suppose that the buyer offers to purchase the firm from the owner for \$100 per share. Is the owner going to sell to her?

The purpose of the verbal description of the example is threefold: First, the "real-life" application should serve as a motivation to studying unawareness in multi-agent settings. We will provide an answer to what extent speculative trade is possible under unawareness in Section 4.3. Second, we will formalize the example using different approaches to awareness. This will allow us to illustrate the main features of those approaches and make them easily accessible

with the help of simple graphs. Third, an answer to this question is relevant for contrasting results under unawareness with results in standard structures without unawareness (e.g., the “No-agreeing-to-disagree” theorem of Aumann [1976] or the “No-speculative-trade” theorem of Milgrom and Stokey [1982]).

2 Preliminary Discussion

2.1 Some Properties of Awareness in Natural Languages

The words “aware” and “unaware” are used in many contexts with many different connotations. Sometimes “aware” is used in place of “knowing” like in the sentence “I was aware of the red traffic light.” On the other hand we interpret “aware” to mean “generally taking into account”, “being present in mind” (Modica and Rustichini [1999], p. 274), “thinking about” (Dekel, Lipman, and Rustichini [1998b]) or “paying attention to” like in the sentence “Be aware of sexually transmitted diseases!” In fact, the last sentence resonates closely with the etymology of “aware” since it has its roots in the old English “gewær” (which itself has roots in the German “gewahr”) emphasizing to be “wary”.² In psychiatry (see for instance, Green et al., 1993), lack of self-awareness means that a patient is oblivious to aspects of an illness that is obvious to social contacts. This is arguably closer to “not knowing”. But it also implies that the patient lacks introspection of her/his lack of knowledge of the illness. It turns out that lack of negative introspection will play a crucial role in modeling unawareness (see Section 2.2). In neuroscience, being aware is taken as making/having/enjoying some experience and being able to specify the content of consciousness (Zeman [2002], pp. 16). While the precise connotations of all those uses of awareness are different, they have in common that the agent is able to conceive something. Being unaware means then that he lacks conception of something.

Describing properties of awareness and unawareness informally with words like “knowing”, “not knowing”, “lack of conception”, “not thinking about it” etc. does not make awareness amenable to formal analysis. We turn now to formal approaches.

2.2 Some Properties of Awareness in a Formal Language

One attempt to avoid ambiguities of natural languages is to use a formal language. Given a nonempty set of agents $\mathbf{Ag} = \{1, \dots, n\}$ indexed by a and a nonempty set of atomic formulae $\mathbf{At}$ (also called primitive propositions) as well as the special formula $\top$, the formulae φ of the language $\mathbf{L}_n^{K,A}(\mathbf{At})$ are defined by the following grammar

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_a\varphi \mid A_a\varphi,$$

where $p \in \mathbf{At}$ and $a \in \mathbf{Ag}$. As usual, the propositional connectives $\neg$ and $\wedge$ denote negation and conjunction, respectively. The epistemic modal operators K_a and A_a are named knowledge and awareness, respectively. For instance, $A_a\varphi$ is read as “agent a is aware of φ ”. Atomic formulae represent propositions such as “penicillium rubens has antibiotic properties” that are

²<http://www.etymonline.com/index.php?term=aware>

not themselves formed of other propositions. As usual, disjunction $\vee$, implication $\rightarrow$, and bi-implication $\leftrightarrow$ are abbreviations, defined by $\varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi)$, $(\varphi \rightarrow \psi) := (\neg\varphi \vee \psi)$, and $(\varphi \leftrightarrow \psi) := (\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$, respectively.

One immediate property of awareness is implicitly assumed with the introduction of just one modal operator A_a : there are no “degrees” of awareness. That is, one does not have statements like “agent a is more aware of φ than she is of ψ but less aware of φ than she is of χ ”. I believe this is justified. The notion of awareness is essentially dichotomous. An agent is either aware of φ or unaware of φ .

With some abuse of notation, let $\text{At}(\varphi)$ denote the set of primitive propositions that appear in φ defined inductively as follows:

- $\text{At}(\top) := \emptyset$,
- $\text{At}(p) := p$, for $p \in \text{At}$,
- $\text{At}(\neg\varphi) := \text{At}(\varphi)$,
- $\text{At}(\varphi \wedge \psi) := \text{At}(\varphi) \cup \text{At}(\psi)$,
- $\text{At}(K_a\varphi) := \text{At}(\varphi) =: \text{At}(A_a\varphi)$.

With this notation on hand, we can formalize the property of “awareness generated by primitive propositions”:

$$\text{AGPP. } A_a\varphi \leftrightarrow \bigwedge_{p \in \text{At}(\varphi)} A_ap$$

An agent a is aware of φ if and only if she is aware of every primitive proposition that appears in φ . This property is not completely innocent. If an agent can think about every primitive proposition $p \in \text{At}(\varphi)$, can she also think about all those primitive propositions joint together in some potentially very complicated formula φ ? Isn’t one feature of unawareness that an agent is sometimes unable to “connect various thoughts”? This property differentiates unawareness generated by primitive propositions from other forms of logical non-omniscience.

Using this syntactic approach, we can state easily further properties of awareness that have been considered in the literature:

- KA. $K_a\varphi \rightarrow A_a\varphi$ (Knowledge implies Awareness)
- AS. $A_a\neg\varphi \leftrightarrow A_a\varphi$ (Symmetry)
- AKR. $A_a\varphi \leftrightarrow A_aK_a\varphi$ (Awareness Knowledge Reflection)
- AR. $A_a\varphi \leftrightarrow A_aA_a\varphi$ (Awareness Reflection)
- AI. $A_a\varphi \rightarrow K_aA_a\varphi$ (Awareness Introspection)

KA relates awareness to knowledge. Quite naturally, knowing φ implies being aware of φ . Symmetry is natural if we take the idea of awareness of φ to mean “being able to think about φ ”. For example, if an agent can think about that penicillium rubens could have antibiotic properties then she can also think about that penicillium rubens does not have antibiotic properties. Symmetry makes clear that awareness is different from notions of knowledge as knowledge does not satisfy symmetry. Awareness Knowledge Reflection is also natural: if an agent can think about some particular proposition, then she can also think about her knowledge of that proposition and vice versa. Similarly, if an agent is aware of a proposition, then she

knows that she is aware of this proposition (Awareness Introspection). Awareness Reflection states that an agent can reflect on her awareness.³

Modica and Rustichini [1994] defined awareness in terms of knowledge by

$$A_a\varphi := K_a\varphi \vee (\neg K_a\varphi \wedge K_a\neg K_a\varphi),$$

which in propositional logic is equivalent to

$$A_a\varphi = K_a\varphi \vee K_a\neg K_a\varphi,$$

a definition of awareness used frequently in the literature especially in economics. The Modica-Rustichini definition of awareness and some of the properties discussed above make use of the knowledge modality. This just begs the question about the properties of knowledge. One property often assumed is

$$5. \neg K_a\varphi \rightarrow K_a\neg K_a\varphi \text{ (Negative Introspection)}$$

This property is implicitly used (together with others) in most economic applications to model agents who are free from “mistakes in information processing.” It is immediate that the Modica-Rustichini definition of awareness is equivalent to Negative Introspection. Thus, if Negative Introspection is a valid formula, then awareness as defined by Modica-Rustichini must be trivial in the sense that the agent is aware of every formula. The discussion then begs the questions about which formulae are valid. Answers to this question are given with various structures discussed in sequel.

3 Awareness and Knowledge

3.1 Awareness Structures by Fagin and Halpern [1988]

Fagin and Halpern [1988] were the first to present a formal approach of modeling awareness. They augment Kripke structures with a syntactic awareness correspondence in order to provide a flexible approach for modeling logical non-omniscience. Their starting point was Levesque’s Logic of Implicit and Explicit Belief (Levesque [1984]). Perhaps that’s why they considered also another epistemic modal operator L_a interpreted as “implicit knowledge”. We denote the resulting language $\mathcal{L}_n^{L,K,A}(\text{At})$.

An *awareness structure* is a tuple $M = (S, (R_a)_{a \in \text{Ag}}, (\mathcal{A}_a)_{a \in \text{Ag}}, V)$ where $(S, (R_a)_{a \in \text{Ag}}, V)$ is a Kripke structure and $\mathcal{A}_a : \Omega \longrightarrow 2^{\mathcal{L}_n^{L,K,A}(\text{At})}$ is agent a ’s awareness correspondence⁴ that associates with each state $s \in S$ the set of formulae $\mathcal{A}_a(s) \subseteq \mathcal{L}_n^{L,K,A}(\text{At})$ of which agent a is aware at state s . A Kripke structure $(S, (R_a)_{a \in \text{Ag}}, V)$ consists of a nonempty set of states S , and for each agent $a \in \text{Ag}$ a binary relation $R_a \subseteq S \times S$, the accessibility relation.⁵ Intuitively,

³We use the term “introspection” when the agent reasons about knowledge of knowledge or awareness. We use the term “reflection” when the agent reasons about awareness of knowledge or awareness.

⁴The name “correspondence” refers to a set-valued function.

⁵See Fagin, Halpern, Moses, and Vardi [1995] or Chellas [1980].

$(s, t) \in R_a$ is interpreted as ‘at state s agent a considers the state t possible’. Of particular interest is the case when R_a is an equivalence relation, i.e., when R_a is a relation that is reflexive (i.e., $(s, s) \in R_a$, for all $s \in S$), transitive (i.e., for all $s_1, s_2, s_3 \in S$, if $(s_1, s_2) \in R_a$, and $(s_2, s_3) \in R_a$, then $(s_1, s_3) \in R_a$) and Euclidean (i.e., for all $s_1, s_2, s_3 \in S$, if $(s_1, s_2) \in R_a$ and $(s_1, s_3) \in R_a$, then $(s_2, s_3) \in R_a$). In this case (and only in this case), R_a forms a partition of the state space S , i.e., a collection of disjoint subsets such that the union covers S . We call an awareness structure in which each agent’s accessibility relation forms a partition a *partitional awareness structure*. Our exposition focuses on partitional awareness structures.⁶ The valuation function $V : S \times \text{At} \rightarrow \{\text{true}, \text{false}\}$ assigns to each state and atomic formula a truth value.

The awareness correspondence offers a very flexible approach to modeling various notions of awareness, some of which have yet to be explored. The authors stress that restrictions depend on which interpretation of awareness is desired in applications. For instance, economists became interested in framing affects including presentation order effects. Presentation order effects may be relevant when information is acquired with the help of online search engines, which typically present lists of search results. Rather than going through all of them, we usually stop when we found a satisfactory result. There may be search results further down the list of which we remain unaware but which we would find much more relevant if we were aware of them. Our awareness of search results depends crucially on the order in which they are presented and on our search aim. If we consider lists as conjunctions of propositions in which the order matters (i.e., for which we do not assume commutativity), then we may be able to model presentation order effects with awareness correspondences because at some state $s \in S$ we may have $\varphi \wedge \psi \in \mathcal{A}_a(s)$ but $\psi \wedge \varphi \notin \mathcal{A}_a(s)$ while in another state $s' \neq s$, $\psi \wedge \varphi \in \mathcal{A}_a(s')$. Such “realistic” approaches have not been explored in the awareness literature. But especially for such kind of applications, the syntactic awareness correspondence may have an advantage over event-based approaches that we will introduce later.

Note that by definition, the awareness correspondence imposes a dichotomous notion of awareness because at a state s a formula can be either in $\mathcal{A}_a(s)$ or not in $\mathcal{A}_a(s)$. Thus, at state s , agent a is either aware or unaware of that formula.

Given how the literature has developed so far, two restrictions on the awareness correspondence that are jointly called *propositionally determined awareness* are of particular interest:

- *Awareness is generated by primitive propositions* if for all $a \in \text{Ag}$ and $s \in S$, $\varphi \in \mathcal{A}_a(s)$ if and only if $\text{At}(\varphi) \subseteq \mathcal{A}_a(s)$.
- *Agents know what they are aware of* if for all $a \in \text{Ag}$ and $s, s' \in S$, $(s, s') \in R_a$ implies $\mathcal{A}_a(s') = \mathcal{A}_a(s)$.

A satisfaction relation specifies for each awareness structure and state which formulae are true. We denote by $M, s \models \varphi$ that φ is true (or satisfied) at state s in the awareness structure M , and define inductively on the structure of formulae in $\mathbb{L}_n^{L, K, A}(\text{At})$,

⁶The aim of our exposition is not necessarily to present the most general setting. Nor do we believe that partitional information structures describe best how humans reason in real life. Rather, partitional information structures serve as a “rational” benchmark. It allows us to demonstrate that one can model both, a “strong” notion of knowledge *and* unawareness. Thus, we can isolate unawareness from errors of information processing associated with lack of introspection. Moreover, as it turns out, the interpretation of propositionally generated unawareness as “lack of conception” is most transparent in partitional awareness structures.


$M, s \models \top$, for all $s \in S$,
 $M, s \models p$, for $p \in \text{At}$ if and only if $V(s, p) = \text{true}$,
 $M, s \models \varphi \wedge \psi$ if and only if both $M, s \models \varphi$ and $M, s \models \psi$,
 $M, s \models \neg\varphi$ if and only if $M, s \not\models \varphi$,
 $M, s \models L_a\varphi$ if and only if $M, t \models \varphi$ for all $t \in S$ such that $(s, t) \in R_a$,
 $M, s \models A_a\varphi$ if and only if $\varphi \in \mathcal{A}_a(s)$,
 $M, s \models K_a\varphi$ if and only if $M, s \models A_a\varphi$ and $M, s \models L_a\varphi$.

The satisfaction relation is standard except for awareness ($A_a\varphi$), which is new, and for knowledge ($K_a\varphi$), which is defined by $L_a\varphi \wedge A_a\varphi$. In the presence of an implicit knowledge modality L_a , the knowledge modality K_a is called explicit knowledge. Agent a explicitly knows φ if she is aware of φ and implicitly knows φ .

At this point, it may be helpful to illustrate awareness structures with our example.

Speculative Trade Example (continued). Denote by ℓ the atomic formula “the lawsuit is brought against the firm” and by n “the novel use of the firm’s product is discovered”. Figure 1 depicts an awareness structure that models the speculative trade example from the Introduction.⁷ There are four states. For simplicity, we name each state by the atomic formulae that are true or false at that state. For instance the upper right state $(n, \neg\ell)$, n is true and ℓ is false. The awareness correspondences are indicated by clouds, one for each player. For each state, the blue solid cloud represents the awareness set of the owner while the red intermitted cloud represents the awareness set of the potential buyer.

Figure 1: An Awareness Structure for the Speculative Trade Example



For graphical simplicity, we represent the accessibility relations of agents by possibility sets rather than arrows, a practice common in game theory. The blue solid-lined possibility set belongs to the owner while the one with the red intermitted line is the buyer’s. Each agent’s information is trivial as neither can distinguish between any states.

⁷Neither Fagin and Halpern [1988], Halpern [2001] nor Halpern and Rêgo [2008] provide a graphical rendering of an example of an awareness structure. We choose “clouds” to depict the awareness sets so as to suggest the interpretation of “thinking about”.

This simple figure models the story that we outlined in the Introduction. In any state, the owner is unaware of the potential innovation but aware of the lawsuit because his awareness set never contains formulae involving n but only formulae involving ℓ . Similarly, at any state the buyer is unaware of the potential lawsuit but aware of the innovation because his awareness set never contains formulae involving ℓ but only formulae involving n . The accessibility relations show us that the owner does not know whether the lawsuit obtains because he cannot distinguish between states in which the lawsuit obtains and states where it doesn't. Analogously, the buyer does not know whether the innovation obtains because she cannot distinguish between states in which the innovation obtains from states where it doesn't obtain. But the owner knows implicitly that the buyer is unaware of the lawsuit because at every state of his possibility set, the buyer's awareness set does not contain formulae involving ℓ . Moreover, he also explicitly knows that the buyer is unaware of the lawsuit because he implicitly knows it and his own awareness set contains formulae involving ℓ . (An analogous statement holds for the buyer n .)

This example is very special because each agent's accessibility relation is trivial and each agent's awareness correspondence is constant across states. Nevertheless, the example illustrates some particular features and properties of awareness structures. One thing to note is that the accessibility relation models implicit knowledge and not necessarily explicit knowledge. For instance, at every state the owner implicitly knows that he is unaware of the innovation although he does not explicitly know it because his awareness set never contains formulae that involve n . This is actually a general property (see axiom UI_L below), which we find hard to interpret. $\square$

The discussion of the example just begs the question about what are the general properties of awareness and knowledge in awareness structures? What are *all* properties of awareness and knowledge in awareness structures? To answer these questions, we characterize the properties of awareness and knowledge in terms of formulae that are valid in awareness structures. To set the stage for such an axiomatization, we need to introduce the following standard notions (see Fagin, Halpern, Moses, and Vardi [1995] or Chellas [1980]): An *axiom* is a formula assumed. An *inference rule* infers a formula (i.e., a conclusion) from a collection of formulae (i.e., the hypothesis). An *axiom system* consists of a collection of axioms and inference rules. A *proof* in an axiom system consists of a sequence of formulae, where each formula is either an axiom in the axiom system or follows by an application of an inference rule. A proof is a proof of a formula φ if the last formula in the proof is φ . A formula φ is *provable* in an axiom system if there is a proof of φ in the axiom system. The set of *theorems* of an axiom system is the smallest set of formulae that contain all axioms and that is closed under inference rules of the axiom system.

Given a class $\mathcal{M}^{FH}$ of awareness structures, a formula φ is *valid* in $\mathcal{M}^{FH}$ if $M, s \models \varphi$ for every awareness structure $M \in \mathcal{M}^{FH}$ and state s in M . An axiom system is said to be *sound* for a language L with respect to a class $\mathcal{M}^{FH}$ of awareness structures if every formula in L that is provable in the axiom system is valid with respect to every awareness structure in $\mathcal{M}^{FH}$. An axiom system is *complete* for a language L with respect to a class of awareness structures $\mathcal{M}^{FH}$ if every formula in L that is valid in $\mathcal{M}^{FH}$ is provable in the axiom system. A sound and complete axiomatization for a class of awareness structures characterizes these awareness structures in terms of properties of knowledge and awareness as codified in the axiom system.

Consider the following axiom system:⁸

- Prop. All substitutions instances of tautologies of propositional logic, including the formula $\top$.
 KL. $K_a\varphi \leftrightarrow L_a\varphi \wedge A_a\varphi$ (Explicit Knowledge is Implicit Knowledge and Awareness)
 AS. $A_a\neg\varphi \leftrightarrow A_a\varphi$ (Symmetry)
 AC. $A_a(\varphi \wedge \psi) \leftrightarrow A_a\varphi \wedge A_a\psi$ (Awareness Conjunction)
 AKR. $A_a\varphi \leftrightarrow A_aK_a\varphi$ (Awareness Explicit Knowledge Reflection)
 ALR. $A_a\varphi \leftrightarrow A_aL_a\varphi$ (Awareness Implicit Knowledge Reflection)
 AR. $A_a\varphi \leftrightarrow A_aA_a\varphi$ (Awareness Reflection)
 AI_L. $A_a\varphi \rightarrow L_aA_a\varphi$ (Awareness Introspection)
 UI_L. $\neg A_a\varphi \rightarrow L_a\neg A_a\varphi$ ($\neg$)
 K. $(L_a\varphi \wedge L_a(\varphi \rightarrow \psi)) \rightarrow L_a\psi$ (Distribution Axiom)
 T. $L_a\varphi \rightarrow \varphi$ (Implicit Knowledge Truth Axiom)
 4. $L_a\varphi \rightarrow L_aL_a\varphi$ (Implicit Positive Introspection Axiom)
 5. $\neg L_a\varphi \rightarrow L_a\neg L_a\varphi$ (Implicit Negative Introspection Axiom)
 MP. From φ and $\varphi \rightarrow \psi$ infer ψ (modus ponens)
 Gen. From φ infer $L_a\varphi$ (Implicit Knowledge Generalization)

Note that each of the axioms and inference rules is an instance of a *scheme*; it defines an infinite collection of axioms (inference rules, respectively), one for each choice of formulae.

Axioms AS, AKR, and AR were motivated in Section 2.2. Awareness Conjunction (AC) has a similar flavor as the property “awareness generated by primitive propositions” (AGPP) introduced in Section 2.2. Axioms ALR and AI_L are similar to axioms AKR and AI, respectively, but with explicit knowledge replaced by implicit knowledge. Axioms and inference rules Prop., K, T, 4, 5, MP, and Gen. together make up the well-known axiom system **S5** (see Fagin, Halpern, Moses, and Vardi [1995] or Chellas [1980]) but are stated here with implicit knowledge modalities. Axiom KL links explicit knowledge and implicit knowledge via awareness. Explicit knowledge is implicit knowledge and awareness. This resonates well with the interpretation of awareness as “being present in mind”. Explicit knowledge, i.e., knowledge that one is aware of, is knowledge that is “present in mind”. The notion of knowledge usually considered in economics corresponds to explicit knowledge despite the fact that standard properties on knowledge are now imposed in implicit knowledge! The axiom Unawareness Introspection (UI_L) is hard to interpret since it has no analog in which implicit knowledge is replaced by explicit knowledge. We denote the above axiom system by $\mathbf{S5}_n^{L,K,A}$ because it is analogous to **S5**.

Theorem 1 (Halpern [2001]) *For the language $\mathbf{L}_n^{L,K,A}(\text{At})$, the axiom system $\mathbf{S5}_n^{L,K,A}$ is a sound and complete axiomatization with respect to partitional awareness structures in which awareness is determined by propositions.*

The theorem is proved by modifying the proof for the well-known result that **S5** is a sound and complete axiomatization with respect to partitional Kripke structures for the language $\mathbf{L}^L(\text{At})$.

While we focus our exposition on the strong notion of knowledge as encapsulated in axiom systems analogous to **S5**, the literature considers weaker notions of knowledge as well.

⁸Some of the axioms involving awareness have been introduced in the literature under various different names. Here we attempt to assign them intuitive acronyms. “R” stands for “reflection” and “I” for introspection.

Halpern [2001] proves also axiomatizations for non-partitional awareness structures. Huang and Kwast [1991] study variants of awareness structures and discuss various notions of negation and implications in the context of logical non-omniscience. Sillari [2008a,b] uses a version of awareness structures in which he replaces the Kripke relations by a neighborhood correspondences $\mathcal{N}_a : S \rightarrow 2^{2^S}$ in the spirit of neighborhood semantics originally introduced by Scott [1970] and Montague [1970] (see Section 5.3). Intuitively, $\mathcal{N}_a(s)$ is the list of events that agent a knows at state s . The satisfaction relation for the case of implicit knowledge of a formula is then modified accordingly to $M, s \models L_a\varphi$ if and only if $\{t \in S : M, t \models \varphi\} \in \mathcal{N}_a(s)$. Modelers interested in studying various forms of logical non-omniscience will welcome the additional flexibility that awareness neighborhood structures provide over awareness (Kripke) structures. But this additional generality is less helpful when being interested in isolating the effect of unawareness per se in the presence of a strong notion of knowledge as the one encapsulated in S5.⁹

The axiomatization of Theorem 1 is somewhat dissatisfactory as most of the properties are stated in terms of implicit knowledge, a notion that we find very hard to interpret in the context of propositionally determined awareness. (In fact, since explicit knowledge is defined in terms of implicit knowledge and awareness, the expressivity of language $\mathbf{L}_n^{L,A}(\text{At})$ is equal to $\mathbf{L}_n^{L,K,A}(\text{At})$.) Implicit knowledge that is not explicit knowledge is as if Isaac Newton would say “I know the theory of relativity but unfortunately I am not aware of it”. In economics, we are only interested in knowledge that the agent is aware of, that can guide her decisions, and that in principle could be tested with choice experiments (see Morris [1996, 1997], Schipper [2014]). While an outsider may be able to reason about the implicit knowledge of an agent, it is hard to see how the agent herself could reason about her implicit knowledge that is not explicit knowledge as well. Some authors in the awareness literature interpret implicit knowledge as “knowledge that the agent would have if she were aware of it”. But this interpretation is flawed because if she really becomes aware of it, then maybe her explicit knowledge would not correspond anymore to her earlier implicit knowledge because her state of mind would have changed from the one she was in when she had this implicit knowledge and was unaware.

Fortunately, it is possible to axiomatize awareness structures without an implicit knowledge modality using the language $\mathbf{L}_n^{K,A}(\text{At})$. The following axiom system that we may denote by $\mathbf{S5}_n^{K,A}$ has been in part already motivated in Section 2.2:

- Prop. All substitutions instances of tautologies of propositional logic, including the formula $\top$.
- KA. $K_a\varphi \rightarrow A_a\varphi$ (Knowledge implies Awareness)
- AS. $A_a\neg\varphi \leftrightarrow A_a\varphi$ (Symmetry)
- AC. $A_a(\varphi \wedge \psi) \leftrightarrow A_a\varphi \wedge A_a\psi$ (Awareness Conjunction)
- AKR. $A_a\varphi \leftrightarrow A_aK_a\varphi$ (Awareness Knowledge Reflection)
- AR. $A_a\varphi \leftrightarrow A_aA_a\varphi$ (Awareness Reflection)
- AI. $A_a\varphi \rightarrow K_aA_a\varphi$ (Awareness Introspection)
- K. $(K_a\varphi \wedge K_a(\varphi \rightarrow \psi)) \rightarrow K_a\psi$ (Distribution Axiom)
- T. $K_a\varphi \rightarrow \varphi$ (Axiom of Truth)

⁹It should be feasible to model awareness using a hybrid of a Kripke structure and a neighborhood structure where, in place of a syntactic awareness correspondence of awareness structures, the neighborhood correspondence lists for each state the set of events that the agent is aware of at that state while knowledge continues to be modeled by the accessibility relation as in Kripke structures. To our knowledge, nobody has explored such a syntax-free approach.

- 4. $K_a\varphi \rightarrow K_aK_a\varphi$ (Positive Introspection Axiom)
- 5_A. $\neg K_a\varphi \wedge A_a\varphi \rightarrow K_a\neg K_a\varphi$ (Weak Negative Introspection Axiom)
- MP. From φ and $\varphi \rightarrow \psi$ infer ψ (modus ponens)
- Gen_A. From φ and $A_a\varphi$ infer $K_a\varphi$ (Modified Knowledge Generalization)

Note that we now require awareness for the formalization of the negative introspection axiom 5_A and knowledge generalization Gen_A. This is quite intuitive. If an agent does not know φ and is also not aware of φ , how could she know that she doesn't know φ ? Moreover, how could she infer knowledge of φ from φ if she isn't aware of φ ?

The definition of the satisfaction relation must be modified with respect to knowledge as follows:

$$M, s \models K_a\varphi \text{ if and only if } M, s \models A_a\varphi \text{ and } M, t \models \varphi \text{ for all } t \in S \text{ such that } (s, t) \in R_a.$$

Theorem 2 (Halpern [2001], Halpern and Rêgo [2008]) *For the language $L_n^{K,A}(At)$, the axiom system $S5_n^{K,A}$ is a sound and complete axiomatization with respect to partitional awareness structures in which awareness is determined by propositions.*

Halpern [2001] proves the theorem with an additional inference rule that Halpern and Rêgo [2008] show indirectly to be unnecessary. We like to mention that Halpern [2001] and Ditmarsch, French, Velázquez-Quesada, and Wáng [2013] prove also axiomatizations for awareness structures that are not necessarily partitional. The latter paper also introduces another notion of knowledge that they dub “speculative knowledge”. This notion is similar to explicit knowledge except that an agent always speculatively knows tautologies even though these tautologies may involve primitive propositions that she is unaware of.

Axiomatizing awareness structures for the language $L_n^{K,A}(At)$ does not completely avoid the issue of implicit knowledge because implicit knowledge and not explicit knowledge is represented by the relation R_a . This is quite in contrast to the interpretation of information partitions in for instance Aumann [1976] despite the fact that R_a is here assumed to be partitional. Note further that the issue with implicit knowledge is potentially much less severe if we are not interested in propositionally determined awareness. With other forms of logical non-omniscience, implicit knowledge may be to a certain extent present in the mind and the objections raised here against the notion of implicit knowledge may be misguided when different notions of awareness are considered.

Recall from Section 2.2 that Modica and Rustichini [1994, 1999] defined awareness in terms of knowledge by $A_a\varphi = K_a\varphi \vee K_a\neg K_a\varphi$.

Lemma 1 (Halpern [2001]) *For any partitional awareness structure in which awareness is generated by primitive propositions, the formula $A_a\varphi \leftrightarrow K_a\varphi \vee K_a\neg K_a\varphi$ is valid.*

The proof relies mainly on weak negative introspection and awareness generated by primitive propositions.

That is, in the class of partitional awareness structures we can define awareness in terms of knowledge as in Modica and Rustichini [1994, 1999] and the expressive power of the language

$L_n^{K,A}(At)$ is equivalent to $L_n^K(At)$. Note that the lemma does not hold in general for awareness structures which are not necessarily partitional.

In conclusion, we like to emphasize that the strength of awareness structures is their flexibility. Potentially they can be used to model many interesting notions of awareness and logical non-omniscience. But this flexibility comes also at a cost. Because of the syntactic awareness correspondence, the semantics of awareness structures is not completely syntax free. All applications that appeared so far in the literature avoided specifying any syntax and just work with states instead. Generally, modeling approaches in economics using syntax are extremely rare. Another issue is that awareness structures model reasoning about awareness and knowledge of agents from an outside analyst's point of view. While an outsider can easily use an awareness structure to analyze agents' interactive reasoning, it is hard to imagine that agents themselves could also use the awareness structure to analyze their reasoning. For instance, if the buyer in the speculative trade example is presented with the awareness structure in Figure 1, then presumably she must become aware of the lawsuit. Writing down an awareness structure from the buyer's point of view, would require us to erase everything involving ℓ (and analogously for the owner). The states in an awareness structure are "objective" descriptions of situations in the eyes of an analyst but not necessarily the agents themselves. This may be problematic in game theoretic models where we are interested in modeling strategic situations from each player's perspective. In the following sections we will introduce structures for modeling awareness that feature "subjective" states, have a syntax-free semantics, and in some cases avoid the notion of implicit knowledge altogether.

3.2 An Impossibility Result on Unawareness in Kripke Frames

We have seen that awareness structures are not syntax-free since they involve a syntactic awareness correspondence. This has been criticized early on by Konolige [1986]. Others conjectured the impossibility of a purely semantic approach to awareness. For instance, Thijsse [1991] writes "... a purely semantic and fully recursive approach would be preferable, but I believe it is intrinsically impossible, due to the psychological nature of awareness." Why would it be difficult to devise a syntax-free semantics for modeling awareness? Suppose we would "erase" all syntactic components of awareness structures; could we still model non-trivial awareness? When erasing each agent's awareness correspondence and the valuation from awareness structures, we are left with what is known as an Aumann structure (Aumann [1976]) or Kripke frame.

Let S be a nonempty space of states and consider a set of events which we may take for simplicity to be the set of all subsets 2^S . A natural occurrence like "penicillium rubens has antibiotic properties" is represented by an event, which is simply a subset of states $E \in 2^S$. That is, E is the set of states in which "penicillium rubens has antibiotic properties". Instead modal operators on propositions, we model epistemic notions of knowledge and unawareness by operators on events. The knowledge operator of agent a is denoted by $\mathbf{K}_a : 2^S \rightarrow 2^S$. For the event $E \in 2^S$, the set $\mathbf{K}_a(E)$ represents the event that agent a knows the event E . Yet, to interpret the operator $\mathbf{K}_a$ as knowledge, we should impose properties that reasonable notions knowledge should satisfy. Here we just require the knowledge operator to satisfy one extremely basic property: Agent a always knows the state space, i.e.,

$$\mathbf{K}_a(S) = S \text{ (Necessitation).}$$

Note that all notions of knowledge or belief in the literature satisfy this property, including the knowledge operator defined from an accessibility relation¹⁰ or possibility correspondence, the probability p -belief operator¹¹ (Monderer and Samet [1989]) as well as belief operators for ambiguous beliefs¹² (Morris [1997]). Yet, this property is not as innocent as it may appear at the first glance. The state space is the universal event; it always obtains. We may interpret S as a tautology. Thus, necessitation can be interpreted as knowing tautologies. Such a basic property may be violated by agents who lack logical omniscience and face potentially very complicated tautologies. For instance, mathematicians work hard to “discover new” theorems. They obviously don’t know all theorems beforehand. There may be at least two reason for why it is hard to know all tautologies. It may be due to logical non-omniscience in the sense that the agent does not realize all implications of her knowledge. Alternatively, it may be because the agent is unaware of some concepts referred to in the tautology.

To see how Necessitation may conflict with lack of awareness, we need a formal notion of unawareness in this purely event-based setting. Denote the unawareness operator of agent a by $\mathbf{U}_a : 2^S \rightarrow 2^S$. For the event $E \in 2^S$, the set $\mathbf{U}_a(E)$ represents the event that agent a is unaware of the event E . In our verbal discussion of the use of the term awareness in psychiatry in Section 2.1, we noted already that lack of awareness may imply lack of negative introspection. Formally we can state

$$\mathbf{U}_a(E) \subseteq \neg \mathbf{K}_a(E) \cap \neg \mathbf{K}_a \neg \mathbf{K}_a(E) \text{ (Plausibility)}$$

This property is implied by the Modica-Rustichini definition of awareness discussed earlier. Being unaware of an event implies not knowing the event and not knowing that you don’t know the event. The negation of an event is here defined by the relative complement of that event with respect to S , i.e., $\neg E := S \setminus E$ is the event that the event E does not occur. Conjunction of events is given by the intersection of events. Thus, $E \cap F$ denotes the event that the event E and the event F occurs. Implication of events is given by the subset relation; $E \subseteq F$ denotes that the event E implies the event F .

The next property states that an agent lacks positive knowledge of her unawareness. That is, she never knows that she is unaware of an event.

$$\mathbf{K}_a \mathbf{U}_a(E) \subseteq \emptyset \text{ (KU Introspection)}$$

While we may know that, in principle, there could exist some events that we are unaware of

¹⁰I.e., $\mathbf{K}_a(E) := \{s \in S : (s, s') \in R_a \text{ implies } s' \in E\}$.

¹¹Let $t_a : S \rightarrow \Delta(S)$ be a type mapping that assigns to each state in s a probability measure on S , where $\Delta(S)$ denotes the set of all probability measures on S . For $p \in [0, 1]$, the probability-of-at-least- p -belief operator is defined on events by $\mathbf{B}_a^p(E) := \{s \in S : t_a(\omega)(E) \geq p\}$.

¹²Let $C(S)$ denote the set of capacities on S , i.e., the set of set functions $\nu : 2^S \rightarrow [0, 1]$ satisfying monotonicity (for all $E, F \subseteq S$, if $E \subseteq F$ then $\nu(E) \leq \nu(F)$) and normalization ($\nu(\emptyset) = 0$, $\nu(S) = 1$). Capacities are like probability “measures” except that they not necessarily satisfy additivity. Capacities have been used extensively in decision theory to model Knightian uncertainty, ambiguous beliefs, or lack of confidence in one’s probabilistic beliefs (see for instance the seminal paper by Schmeidler [1989]). Typically, Knightian uncertainty is distinguished from risk in economics. Risk refers to situations in which the agent reasons probabilistically while Knightian uncertainty refer to a situation in which the agent is unable to form probability judgements. Let $t_a : S \rightarrow C(S)$ be a type mapping that assigns to each state a capacity on S . The capacity-of-at-least- p -belief operator is now defined analogously to the probability-of-at-least- p -belief operator.

(an issue we will turn to in Section 5), we cannot know that we are unaware of a specific event E . In the same vein we may also require that if an agent is aware that she is unaware of an event, then she should be aware of the event. Stated in the contrapositive,

$$\mathbf{U}_a(E) \subseteq \mathbf{U}_a \mathbf{U}_a(E) \text{ (AU Reflection).}$$

In order to interpret $\mathbf{K}_a$ as knowledge we should certainly impose further properties but we can already show a very simple but conceptually important impossibility result according to which the above notion of unawareness is inconsistent with any notion of knowledge satisfying Necessitation. Since we did not even assume that the knowledge operator is derived from an accessibility relation (or a possibility correspondence), Theorem 3 below will apply more generally to any state-space model with knowledge satisfying necessitation, and not just to Kripke frames or Aumann structures.

Theorem 3 (Dekel, Lipman, and Rustichini [1998a]) *If a state-space model satisfies Plausibility, KU-introspection, AU-reflection, and Necessitation, then $\mathbf{U}_a(E) = \emptyset$, for any event $E \in 2^S$.*

$$\text{PROOF. } \mathbf{U}_a(E) \stackrel{AU-Ref.}{\subseteq} \mathbf{U}_a(\mathbf{U}_a(E)) \stackrel{Plaus.}{\subseteq} \neg \mathbf{K}_a(\neg \mathbf{K}_a(\mathbf{U}_a(E))) \stackrel{KU-Intro.}{=} \neg \mathbf{K}_a(S) \stackrel{Nec.}{=} \emptyset. \quad \square$$

This shows that the (“standard”) state-space approach is incapable of modeling unawareness. Thus, we need more structure for modeling non-trivial unawareness than what Kripke structures have to offer.¹³

Our brief discussion of Necessitation already suggests that more careful descriptions of states (i.e., syntactic approaches) are useful for modeling awareness. Tautologies are descriptions that are true in every state. Knowing tautologies seems to imply that at every state the agent is able to reason with a language that is as expressive as the most complicated tautology. But if she can use this rich language at every state, then she should be able to describe and reason about any event expressible in this language and thus it may not come as a surprise that she must be aware of all events. The syntactic approach is nicely fine-grained as the “internal structure” of states can be made explicit. This allows us to write formally properties like “awareness generated by primitive propositions” (AGPP), $A_a\varphi \leftrightarrow \bigwedge_{p \in \text{At}(\varphi)} A_ap$, a property that we discussed already in Section 2.2.

It may be worthwhile to ask how awareness structures circumvent impossibility results like the one discussed in this section. Dekel, Lipman, and Rustichini [1998a] identify two assumptions that are implicitly satisfied in every event-based approach like Aumann structures or Kripke frames. One of them is *event-sufficiency*. It says that if two formulae are true in exactly the same subset of states, then (1) the subset of states in which the agent knows one formula must coincide with the subset of states in which the agent knows the other formula and (2) the subset of states in which the agent is unaware of one formula must coincide with the subset of states in which the agent is unaware of the other formula. Clearly, awareness

¹³Dekel, Lipman, and Rustichini [1998a] present also impossibility results in which necessitation is weakened or replaced by monotonicity. Chen, Ely, and Luo [2012] and Montiel Olea [2012] provide further elaborations of those impossibility results.

structures do not satisfy event-sufficiency since two formulae may be true exactly at the same subset of states but the awareness correspondence may be such that the agent is aware of one but not the other in some states. It is the awareness correspondence that allow awareness structures to overcome the impossibility result.

3.3 Unawareness Frames

Inspired by Aumann structures, Heifetz, Meier, and Schipper [2006] introduced an event-based approach to unawareness, that is, a syntax-free semantics for multi-agent unawareness. To circumvent the impossibility results by Modica and Rustichini [1994] and Dekel, Lipman, and Rustichini [1998a], they work with a lattice of state spaces rather than a single state-space.

Let $\mathcal{S} = (\{S_\alpha\}_{\alpha \in \mathcal{A}}, \succeq)$ be a complete lattice of disjoint *state-spaces*, with the partial order $\succeq$ on $\mathcal{S}$.¹⁴ A complete lattice is a partially ordered set in which each subset has a least upper bound (i.e., supremum) and a greatest lower bound (i.e., infimum). If S_α and S_β are such that $S_\alpha \succeq S_\beta$ we say that “ S_α is more expressive than S_β .” Intuitively, states of S_α “describe situations with a richer vocabulary” than states of S_β . Denote by $\Omega = \bigcup_{\alpha \in \mathcal{A}} S_\alpha$ the union of these spaces. This is by definition of $\{S_\alpha\}_{\alpha \in \mathcal{A}}$ a disjoint union.

For every S and S' such that $S' \succeq S$, there is a surjective projection $r_S^{S'} : S' \rightarrow S$, where r_S^S is the identity.¹⁵ (“ $r_S^{S'}(\omega)$ is the restriction of the description ω to the more limited vocabulary of S .”) Note that the cardinality of S is smaller than or equal to the cardinality of S' . We require the projections to commute: If $S'' \succeq S' \succeq S$ then $r_S^{S''} = r_S^{S'} \circ r_{S'}^{S''}$. If $\omega \in S'$, denote $\omega_S = r_S^{S'}(\omega)$. If $D \subseteq S'$, denote $D_S = \{\omega_S : \omega \in D\}$. Intuitively, projections “translate” states from “more expressive” spaces to states in “less expressive” spaces by “erasing” facts that can not be expressed in a lower space.

For $D \subseteq S$, denote $D^\uparrow = \bigcup_{S' \in \{S' : S' \succeq S\}} (r_S^{S'})^{-1}(D)$. (“All the extensions of descriptions in D to at least as expressive vocabularies.”) This is the union of inverse images of D in weakly higher spaces.

An *event* is a pair (E, S) , where $E = D^\uparrow$ with $D \subseteq S$, where $S \in \mathcal{S}$. D is called the *base* and S the *base-space* of (E, S) , denoted by $S(E)$. If $E \neq \emptyset$, then S is uniquely determined by E and, abusing notation, we write E for (E, S) . Otherwise, we write $\emptyset^S$ for $(\emptyset, S)$. Note that not every subset of Ω is an event. Intuitively, some fact may obtain in a subset of a space. Then this fact should be also “expressible” in “more expressive” spaces. Therefore the event contains not only the particular subset but also its inverse images in “more expressive” spaces.

Let Σ be the set of *events* of Ω , i.e., sets $D^\uparrow$ such that $D \subseteq S$, for some state space $S \in \mathcal{S}$. Note that unless $\mathcal{S}$ is a singleton, Σ is not an algebra because it contains distinct $\emptyset^S$ for all $S \in \mathcal{S}$. The event $\emptyset^S$ should be interpreted as a “logical contradiction phrased with the expressive power available in S ”. It is quite natural to have distinct vacuous events since “contradictions can be phrased with differing expressive powers”.

If $(D^\uparrow, S)$ is an event where $D \subseteq S$, the negation $\neg(D^\uparrow, S)$ of $(D^\uparrow, S)$ is defined by

¹⁴Recall that a binary relation is a partial order if it is reflexive, antisymmetric, and transitive.

¹⁵Recall that a function $f : X \rightarrow Y$ is surjective (or called onto) if for every $y \in Y$ there is some $x \in X$ such that $f(x) = y$.

$\neg(D^\uparrow, S) := ((S \setminus D)^\uparrow, S)$. Note that, by this definition, the negation of an event is an event. Abusing notation, we write $\neg D^\uparrow := \neg(D^\uparrow, S)$. By our notational convention, we have $\neg S^\uparrow = \emptyset^S$ and $\neg \emptyset^S = S^\uparrow$, for each space $S \in \mathcal{S}$. $\neg D^\uparrow$ is typically a proper subset of the complement $\Omega \setminus D^\uparrow$, that is, $(S \setminus D)^\uparrow \subsetneq \Omega \setminus D^\uparrow$. Intuitively, there may be states in which the description of an event $D^\uparrow$ is both expressible and valid – these are the states in $D^\uparrow$; there may be states in which its description is expressible but invalid – these are the states in $\neg D^\uparrow$; and there may be states in which neither its description nor its negation are expressible – these are the states in $\Omega \setminus (D^\uparrow \cup \neg D^\uparrow) = \Omega \setminus S(D^\uparrow)^\uparrow$. Thus unawareness structures are not standard state-space models in the sense of Dekel, Lipman, and Rustichini [1998a] because the definition of negation prevents them from satisfying what they call *real states*.

If $\left\{ (D_\lambda^\uparrow, S_\lambda) \right\}_{\lambda \in L}$ is a collection of events (with $D_\lambda \subseteq S_\lambda$, for $\lambda \in L$), their conjunction $\bigwedge_{\lambda \in L} (D_\lambda^\uparrow, S_\lambda)$ is defined by $\bigwedge_{\lambda \in L} (D_\lambda^\uparrow, S_\lambda) := \left(\left(\bigcap_{\lambda \in L} D_\lambda^\uparrow \right), \sup_{\lambda \in L} S_\lambda \right)$. Note, that since $\mathcal{S}$ is a *complete* lattice, $\sup_{\lambda \in L} S_\lambda$ exists. If $S = \sup_{\lambda \in L} S_\lambda$, then we have $\left(\bigcap_{\lambda \in L} D_\lambda^\uparrow \right) = \left(\bigcap_{\lambda \in L} \left((r_{S_\lambda}^S)^{-1}(D_\lambda) \right) \right)^\uparrow$. Again, abusing notation, we write $\bigwedge_{\lambda \in L} D_\lambda^\uparrow := \bigcap_{\lambda \in L} D_\lambda^\uparrow$ (we will therefore use the conjunction symbol $\wedge$ and the intersection symbol $\cap$ interchangeably).

Intuitively, to take the intersection of events $(D_\lambda^\uparrow, S_\lambda)_{\lambda \in L}$, we express them “most economically in the smallest language” in which they are all expressible $S = \sup_{\lambda \in L} S_\lambda$, take the intersection, and then the union of inverse images obtaining the event $\left(\bigcap_{\lambda \in L} ((r_{S_\lambda}^S)^{-1}(D_\lambda)) \right)^\uparrow$ that is based in S .

We define the relation $\subseteq$ between events (E, S) and (F, S') , by $(E, S) \subseteq (F, S')$ if and only if $E \subseteq F$ as sets *and* $S' \preceq S$. If $E \neq \emptyset$, we have that $(E, S) \subseteq (F, S')$ if and only if $E \subseteq F$ as sets. Note however that for $E = \emptyset^S$ we have $(E, S) \subseteq (F, S')$ if and only if $S' \preceq S$. Hence we can write $E \subseteq F$ instead of $(E, S) \subseteq (F, S')$ as long as we keep in mind that in the case of $E = \emptyset^S$ we have $\emptyset^S \subseteq F$ if and only if $S \succeq S(F)$. It follows from these definitions that for events E and F , $E \subseteq F$ is equivalent to $\neg F \subseteq \neg E$ only when E and F have the same base, i.e., $S(E) = S(F)$.

Intuitively, in order to say “ E implies F ” we must be able to express F in the “language” used to express E . Hence, it must be that $S(F) \preceq S(E)$. The inclusion is then just $E \cap S(E) \subseteq F \cap S(E)$.

The disjunction of $\left\{ D_\lambda^\uparrow \right\}_{\lambda \in L}$ is defined by the de Morgan law $\bigvee_{\lambda \in L} D_\lambda^\uparrow = \neg \left(\bigwedge_{\lambda \in L} \neg (D_\lambda^\uparrow) \right)$. Typically $\bigvee_{\lambda \in L} D_\lambda^\uparrow \subsetneq \bigcup_{\lambda \in L} D_\lambda^\uparrow$, and if all D_λ are nonempty we have that $\bigvee_{\lambda \in L} D_\lambda^\uparrow = \bigcup_{\lambda \in L} D_\lambda^\uparrow$ holds if and only if all the $D_\lambda^\uparrow$ have the same base-space.

So far, we have just described an event structure. To formalize the state of mind of an agent, a possibility correspondence is introduced analogous to the one in standard game theory (see for instance, Osborne and Rubinstein [1994], Chapter 5). For each agent $a \in \mathbf{Ag}$ there is a *possibility correspondence* $\Pi_a : \Omega \rightarrow 2^\Omega$ with the following properties:

Confinement: If $\omega \in S$ then $\Pi_a(\omega) \subseteq S'$ for some $S' \preceq S$.

Generalized Reflexivity: $\omega \in \Pi_a^\uparrow(\omega)$ for every $\omega \in \Omega$.

Stationarity: $\omega' \in \Pi_a(\omega)$ implies $\Pi_a(\omega') = \Pi_a(\omega)$.

Projections Preserve Ignorance: If $\omega \in S'$ and $S \preceq S'$ then $\Pi_a^\uparrow(\omega) \subseteq \Pi_a^\uparrow(\omega_S)$.

Projections Preserve Knowledge: If $S \preceq S' \preceq S''$, $\omega \in S''$ and $\Pi_a(\omega) \subseteq S'$ then $(\Pi_a(\omega))_S = \Pi_a(\omega_S)$.

Note that Generalized Reflexivity implies that if $S' \preceq S$, $\omega \in S$ and $\Pi_a(\omega) \subseteq S'$, then $r_{S'}^S(\omega) \in \Pi_a(\omega)$. Additionally, we have the possibility correspondence is serial, i.e., $\Pi_a(\omega) \neq \emptyset$, for all $\omega \in \Omega$.

The possibility correspondence is the analogue to the accessibility relations in Kripke structures. Generalized Reflexivity and Stationarity are the analogues of the partitional properties of the possibility correspondence in partitional Aumann structures or Kripke structures. In particular, Generalized Reflexivity yields the truth property; Stationarity will guarantee the introspection properties (see Proposition 1). It captures both transitivity and Euclideaness.

The properties Projections Preserve Ignorance and Projections Preserve Knowledge guarantee the coherence of knowledge and awareness of individuals down the lattice structure. They compare the possibility sets of an individual in a state ω and its projection ω_S . The properties guarantee that, first, at the projected state ω_S the individual knows nothing she does not know at ω , and second, at the projected state ω_S the individual is not aware of anything she is unaware of at ω (Projections Preserve Ignorance). Third, at the projected state ω_S the individual knows every event she knows at ω , provided that this event is based in a space lower than or equal to S (Projections Preserve Knowledge). These properties also imply that at the projected state ω_S the individual is aware of every event she is aware of at ω , provided that this event is based in a space lower than or equal to S .¹⁶

The *knowledge operator* of agent a on events E is defined, as usual in Aumann structures, by

$$\mathbf{K}_a(E) := \{\omega \in \Omega : \Pi_a(\omega) \subseteq E\},$$

if there is a state ω such that $\Pi_a(\omega) \subseteq E$, and by $\mathbf{K}_a(E) := \emptyset^{S(E)}$ otherwise.

The *awareness operator* of agent a on events E can be defined by

$$\mathbf{A}_a(E) := \left\{ \omega \in \Omega : \Pi_a(\omega) \subseteq S(E)^\uparrow \right\},$$

if there is a state ω such that $\Pi_a(\omega) \subseteq S(E)^\uparrow$, and by $\mathbf{A}_a(E) := \emptyset^{S(E)}$ otherwise. Thus, an agent is aware of an event if she considers possible states in which this event is “expressible”.

Both, the knowledge and awareness operators are well-defined and easy to work with:

Lemma 2 *If E is an event, then $\mathbf{K}_a(E)$ and $\mathbf{A}_a(E)$ are $S(E)$ -based events.*

The proof of the lemma makes use of the properties imposed on the possibility correspondence as does the proof of the following proposition.

The *unawareness operator* on events is defined as the negation of awareness, $\mathbf{U}_a(E) := \neg \mathbf{A}_a(E)$.

Proposition 1 (Heifetz, Meier, and Schipper [2006, 2008]) *The Knowledge and Awareness operators satisfy following properties:*

¹⁶Heifetz, Meier, and Schipper [2006] also state another property, called Projections Preserve Awareness, and remark that it follows from other properties.

Necessitation: $\mathbf{K}_a(\Omega) = \Omega$,
Distribution: $\mathbf{K}_a(\bigcap_{\lambda \in L} E_\lambda) = \bigcap_{\lambda \in L} \mathbf{K}_a(E_\lambda)$,
Monotonicity: $E \subseteq F$ implies $\mathbf{K}_a(E) \subseteq \mathbf{K}_a(F)$,
Truth: $\mathbf{K}_a(E) \subseteq E$,
Positive Introspection: $\mathbf{K}_a(E) \subseteq \mathbf{K}_a \mathbf{K}_a(E)$,
Negative Non-Introspection: $\neg \mathbf{K}_a(E) \cap \neg \mathbf{K}_a \neg \mathbf{K}_a(E) \subseteq \neg \mathbf{K}_a \neg \mathbf{K}_a \neg \mathbf{K}_a(E)$,
Weak Negative Introspection: $\neg \mathbf{K}_a(E) \cap \mathbf{A}_a \neg \mathbf{K}_a(E) = \mathbf{K}_a \neg \mathbf{K}_a(E)$,
MR Awareness: $\mathbf{A}_a(E) = \mathbf{K}_a(E) \cup \mathbf{K}_a \neg \mathbf{K}_a(E)$,
Strong Plausibility: $\mathbf{U}_a(E) = \bigcap_{n=1}^{\infty} (\neg \mathbf{K}_a)^n(E)$,
KU Introspection: $\mathbf{K}_a \mathbf{U}_a(E) = \emptyset^{S(E)}$,
AU Reflection: $\mathbf{U}_a(E) = \mathbf{U}_a \mathbf{U}_a(E)$,
Symmetry: $\mathbf{A}_a(E) = \mathbf{A}_a(\neg E)$,
A-Conjunction: $\bigcap_{\lambda \in L} \mathbf{A}_a(E_\lambda) = \mathbf{A}_a(\bigcap_{\lambda \in L} E_\lambda)$,
AK-Reflection: $\mathbf{A}_a(E) = \mathbf{A}_a \mathbf{K}_a(E)$,
Awareness Reflection: $\mathbf{A}_a(E) = \mathbf{A}_a \mathbf{A}_a(E)$,
Awareness Introspection: $\mathbf{A}_a(E) = \mathbf{K}_a \mathbf{A}_a(E)$.

The event-based approach lends itself well to study *interactive* reasoning about knowledge and awareness. Common knowledge can be defined in the usual way (see Aumann [1999]). The mutual knowledge operator on events is defined by

$$\mathbf{K}(E) := \bigcap_{a \in \text{Ag}} \mathbf{K}_a(E).$$

The common knowledge operator on events is defined by

$$\mathbf{CK}(E) := \bigcap_{n=1}^{\infty} \mathbf{K}^n(E).$$

Analogously we can define mutual and common awareness. The mutual awareness operator on events is defined by

$$\mathbf{A}(E) = \bigcap_{a \in \text{Ag}} \mathbf{A}_a(E),$$

and the common awareness operator by

$$\mathbf{CA}(E) = \bigcap_{n=1}^{\infty} (\mathbf{A})^n(E).$$


Proposition 2 (Heifetz, Meier, and Schipper [2006, 2008]) *The following multi-agent properties obtain: For all events E and agents $a, b \in \text{Ag}$,*

- | | |
|---|--|
| 1. $\mathbf{A}_a(E) = \mathbf{A}_a \mathbf{A}_b(E)$, | 5. $\mathbf{A}(E) = \mathbf{CA}(E)$, |
| 2. $\mathbf{A}_a(E) = \mathbf{A}_a \mathbf{K}_b(E)$, | 6. $\mathbf{K}(E) \subseteq \mathbf{A}(E)$, |
| 3. $\mathbf{K}_a(E) \subseteq \mathbf{A}_a \mathbf{K}_b(E)$, | 7. $\mathbf{CK}(E) \subseteq \mathbf{CA}(E)$, |
| 4. $\mathbf{A}(E) = \mathbf{K}(S(E)^\uparrow)$, | 8. $\mathbf{CK}(S(E)^\uparrow) \subseteq \mathbf{CA}(E)$. |

At this point, it may be useful to illustrate unawareness frames with our speculative trade example:

Speculative Trade Example (continued). Consider the unawareness structure depicted in Figure 2. There are four spaces. Space $S_{\{n,\ell\}}$ is the richest space in which both the lawsuit and the innovation are expressible. Both spaces, $S_{\{n\}}$ and $S_{\{\ell\}}$, are less expressive than $S_{\{n,\ell\}}$. $S_{\{n\}}$ is the space in which only the innovation is expressible while $S_{\{\ell\}}$ is the space in which only the lawsuit is expressible. Finally, neither the innovation nor the lawsuit are expressible in the lowest space, $S_\emptyset$. We let $\succeq$ be defined by $S_{\{n,\ell\}} \succeq S_{\{n\}} \succeq S_\emptyset$ and $S_{\{n,\ell\}} \succeq S_{\{\ell\}} \succeq S_\emptyset$. Projections from higher to lower spaces are indicated by the grey dotted lines. For instance, state $(\neg n, \neg \ell) \in S_{\{n,\ell\}}$ projects to $(\neg n) \in S_{\{n\}}$. It also projects to $(\neg \ell) \in S_{\{\ell\}}$. Both $(\neg n) \in S_{\{n\}}$ and $(\neg \ell) \in S_{\{\ell\}}$ project to $(\top) \in S_\emptyset$. The possibility correspondence is given by the blue solid and red intermitted arrows and soft-edged rectangles for the owner and the potential buyer, respectively. At any state in $S_{\{n,\ell\}}$ the owner's possibility set is at $S_{\{\ell\}}$. Thus, he is unaware of the innovation but aware of the lawsuit. Further, the owner's possibility set includes all states in $S_{\{\ell\}}$, which means that he does not know whether the lawsuit obtains or not. Since at every state in $S_{\{\ell\}}$ the buyer's possibility set is on S , in any state in $S_{\{\ell\}}$ the buyer is unaware of the lawsuit and the owner knows that. At any state in $S_{\{n,\ell\}}$ the buyers's possibility set is at $S_{\{n\}}$. Thus, he is unaware of the lawsuit but aware of the innovation. Further, the buyers's possibility includes all states in $S_{\{n\}}$, which means that she does not know whether the lawsuit obtains or not. Since at every state in $S_{\{n\}}$ the owner's possibility set is on S , in any state in $S_{\{n\}}$ the owner is unaware of the innovation and the buyer knows that. Thus, the unawareness frame of Figure 2 models the speculative trade example.

Figure 2: An Unawareness Frame for the Speculative Trade Example



In comparison to awareness structures, we observe that the possibility correspondences model explicit knowledge.¹⁷ In fact, together with the lattice structure, they also model aware-

¹⁷This is not to say that one couldn't define implicit knowledge in unawareness frames. An "implicit" possibility

ness determined by the space in which the possibility set lies. Although we used suggestive labels such as (n, ℓ) etc., unawareness frames are syntax-free. Finally, while the entire unawareness frame is the analysts model of the situation, it contains directly the “submodels” of agents. For instance, the sublattice consisting of the two spaces, $S_{\{\ell\}}$ and $S_{\emptyset}$, corresponds to the model of the owner while the sublattice consisting of the two spaces, $S_{\{n\}}$ and $S_{\emptyset}$, is the buyer’s model. Moreover, the sublattice $S_{\emptyset}$ is the model that both agents attribute to each other. The states in all those spaces can be interpreted as subjective descriptions of situations in the respective agent’s mind. $\square$

Board, Chung, and Schipper [2011] study properties of unawareness frames in which the possibility correspondence does not necessarily satisfy generalized reflexivity and stationarity. In such a case of course, Truth, Positive Introspection, Negative Non-Introspection, and Weak Negative Introspection may fail. More interestingly, KU-Introspection fails as well. This suggests that unawareness may not only persist despite a strong notion of knowledge like embodied in the properties of S5 but that it may even be enhanced by it.

Galanis [2013a] studies a variant of unawareness frames in which he drops the property Projections Preserve Knowledge of the possibility correspondence. His motivation is to study to what extent unawareness can constrain an agent’s knowledge and can impair her reasoning about what other agents know.

Schipper [2014] complements unawareness frames with decision theoretic primitives like preference relation over acts, i.e., functions from states to real numbers. This allows him to characterize properties of the possibility correspondence by corresponding properties of a decision maker’s preference relation. This extends the approach by Morris [1996, 1997] for standard states-spaces to unawareness frames.

3.4 Unawareness Structures

While the event-based approach of unawareness frames is a tractable approach to modeling nontrivial reasoning about knowledge and awareness among multiple agents, it leaves many questions open. For instance, when introducing the event-based approach we often alluded to intuitive explanations typeset in quotation marks that referred to “expressibility” etc. What justifies such an interpretation? Is it possible to link formally the “expressibility” of state spaces to the expressivity of languages? What does the expressivity of languages has to do with the notion of awareness used in unawareness frames? We also saw that an event may obtain in some states, its negation may obtain in others, and yet in others this event or its negation may not even be defined. This suggests that implicitly a three-valued logic is lurking behind the approach. Again, can we make this explicit? Moreover, Proposition 1 presents properties that awareness and knowledge satisfy in the event-based approach. But are these *all* the properties? That is, can we axiomatize the event-based approach in terms of all the properties of awareness and knowledge? Moreover, can we guarantee that the event-based approach is comprehensive enough so that we can model with it all situations with such properties. These questions can be addressed by introducing a logical apparatus and constructing the canonical unawareness

correspondence could be defined from the possibility correspondence by taking the inverse images of possibility sets in the upmost space.

structure as done in Halpern and Rêgo [2008] and Heifetz, Meier, and Schipper [2008].

Consider the language $L_n^{K,A}(\text{At})$ and define, as in Modica and Rustichini [1994, 1999], awareness in terms of knowledge by

$$A_a\varphi := K_a\varphi \vee K_a\neg K_a\varphi.$$

With this definition, we consider the following axiom system that we call $\widetilde{S5}_n^{K,A}$:

- Prop. All substitutions instances of tautologies of propositional logic, including the formula $\top$.
- AS. $A_a\neg\varphi \leftrightarrow A_a\varphi$ (Symmetry)
- AC. $A_a(\varphi \wedge \psi) \leftrightarrow A_a\varphi \wedge A_a\psi$ (Awareness Conjunction)
- A_aK_bR . $A_a\varphi \leftrightarrow A_aK_b\varphi$, for all $b \in \text{Ag}$ (Awareness Knowledge Reflection)
- T. $K_a\varphi \rightarrow \varphi$ (Axiom of Truth)
- 4. $K_a\varphi \rightarrow K_aK_a\varphi$ (Positive Introspection Axiom)
- MP. From φ and $\varphi \rightarrow \psi$ infer ψ (modus ponens)
- RK. For all natural numbers $n \geq 1$, if $\varphi_1, \varphi_2, \dots, \varphi_n$ and φ are such that $\text{At}(\varphi) \subseteq \bigcup_{i=1}^n \text{At}(\varphi_i)$, then $\varphi_1 \wedge \varphi_2 \wedge \dots \wedge \varphi_n \rightarrow \varphi$ implies $K_a\varphi_1 \wedge K_a\varphi_2 \wedge \dots \wedge K_a\varphi_n \rightarrow K_a\varphi$. (RK-Inference)

We also define the modality U_a by $U_a\varphi := \neg A_a\varphi$ read as “agent a is unaware of φ ”.

Remark 3 (Heifetz, Meier, and Schipper [2008]) *The Modica and Rustichini definition of awareness and axiom system $\widetilde{S5}_n^{K,A}$ implies:*

- K. $K_a\varphi \wedge K_a(\varphi \rightarrow \psi) \rightarrow K_a\psi$
- $K_a\varphi \wedge K_a\psi \rightarrow K_a(\varphi \wedge \psi)$
- NNI. $U_a\varphi \rightarrow \neg K_a\neg K_a\neg K_a\varphi$
- AI. $A_a\varphi \rightarrow K_aA_a\varphi$
- AGPP. $A_a\varphi \leftrightarrow \bigwedge_{p \in \text{At}(\varphi)} A_ap$
- Gen_A . If φ is a theorem, then $A_a\varphi \rightarrow K_a\varphi$ is a theorem.

For every $\text{At}' \subseteq \text{At}$, let $S_{\text{At}'}$ be the set of maximally consistent sets $\omega_{\text{At}'}$ of formulae in the sublanguage $L_n^{K,A}(\text{At}')$. Given a language $L_n^{K,A}(\text{At}')$, a set of formulae Γ is consistent with respect to an axiom system if and only if there is no formula φ such that both φ and $\neg\varphi$ are provable from Γ . ω_{At} is maximally consistent if it is consistent and for any formula $\varphi \in L_n^{K,A}(\text{At}') \setminus \omega_{\text{At}'}$, the set $\omega_{\text{At}'} \cup \{\varphi\}$ is not consistent. By standard arguments (see Chellas [1980]) one can show that every consistent subset of $L_n^{K,A}(\text{At}')$ can be extended to a maximally consistent subset $\omega_{\text{At}'}$ of $L_n^{K,A}(\text{At}')$. Moreover, $\Gamma \subseteq L_n^{K,A}(\text{At}')$ is a maximally consistent subset of $L_n^{K,A}(\text{At}')$ if and only if Γ is consistent and for every $\varphi \in L_n^{K,A}(\text{At}')$, $\varphi \in \Gamma$ or $\neg\varphi \in \Gamma$.

Clearly, $\{S_{\text{At}'}\}_{\text{At}' \subseteq \text{At}}$ is a complete lattice of disjoint spaces by set inclusion defined on the set of atomic formulae. Define the partial order on $\{S_{\text{At}'}\}_{\text{At}' \subseteq \text{At}}$ by $S_{\text{At}_1} \succeq S_{\text{At}_2}$ if and only if $\text{At}_1 \supseteq \text{At}_2$. Let $\Omega := \bigcup_{\text{At}' \subseteq \text{At}} S_{\text{At}'}$. For any $S_{\text{At}_1} \succeq S_{\text{At}_2}$, surjective projections $r_{\text{At}_2}^{\text{At}_1} : S_{\text{At}_1} \rightarrow S_{\text{At}_2}$ are defined by $r_{\text{At}_2}^{\text{At}_1}(\omega) := \omega \cap L_n^{K,A}(\text{At}_2)$.

Theorem 4 (Heifetz, Meier, and Schipper [2008]) *For every ω and $a \in \text{Ag}$, the possibility correspondence defined by*

$$\Pi_a(\omega) := \left\{ \omega' \in \Omega : \text{For every formula } \varphi, \begin{array}{l} (i) K_a\varphi \text{ implies } \varphi \in \omega', \text{ and} \\ (ii) A_a\varphi \in \omega \text{ if and only if } \varphi \in \omega \text{ or } \neg\varphi \in \omega \end{array} \right\}$$

satisfies Confinement, Generalized Reflexivity, Projections Preserve Ignorance, and Projections Preserve Knowledge. Moreover, for every formula φ , the set of states $[\varphi] := \{\omega \in \Omega : \varphi \in \omega\}$ is a $S_{\text{At}(\varphi)}$ -based event, and $[\neg\varphi] = \neg[\varphi]$, $[\varphi \wedge \psi] = [\varphi] \cap [\psi]$, $[K_a\varphi] = \mathbf{K}_a[\varphi]$, $[A_a\varphi] = \mathbf{A}_a[\varphi]$, and $[U_a\varphi] = \mathbf{U}_a[\varphi]$.

The canonical unawareness structure is constructed such that states are consistent and comprehensive descriptions, and the “internal” descriptions of the states is reflected by operations on the events.

We can extend unawareness frames to unawareness structures by adding a valuation. An unawareness structure $M = (\mathcal{S}, (r_S^{S'})_{S' \succeq S; S, S' \in \mathcal{S}}, (\Pi_a)_{a \in \text{Ag}}, V)$ is an unawareness frame $(\mathcal{S}, (r_S^{S'})_{S' \succeq S; S, S' \in \mathcal{S}}, (\Pi_a)_{a \in \text{Ag}})$ and a valuation $V : \text{At} \longrightarrow \Sigma$ that assigns to each atomic formula in At an event in Σ . The set $V(p)$ is the event in which p obtains.

The satisfaction relation is defined inductively on the structure of formulae in $\mathbf{L}_n^{K,A}(\text{At})$

$$\begin{aligned} M, \omega &\models \top, \text{ for all } \omega \in \Omega, \\ M, \omega &\models p \text{ if and only if } \omega \in V(p), \\ M, \omega &\models \varphi \wedge \psi \text{ if and only if } [\varphi] \cap [\psi], \\ M, \omega &\models \neg\varphi \text{ if and only if } \omega \in [\neg\varphi], \\ M, \omega &\models K_a\varphi \text{ if and only if } \omega \in \mathbf{K}_a[\varphi], \end{aligned}$$

where $[\varphi] := \{\omega' \in \Omega : M, \omega' \models \varphi\}$, for every formula φ . Note that $[\varphi]$ is an event in the unawareness structure. Recall that $A_a\varphi := K_a\varphi \vee K_a\neg K_a\varphi$. Thus, indeed the satisfaction relation is defined for formulae in $\mathbf{L}_n^{K,A}(\text{At})$.

Our aim is to state a characterization of unawareness structures in terms of properties. More precisely, we seek a complete and sound axiomatization of unawareness structures. To this extent we need to define first the notion of validity. Recall that a formula is said to be valid in a Kripke structure if it is true in *every* state. Yet, in unawareness structures, a nontrivial formula is not even defined in all states. Thus, the definition of validity for Kripke structures is not directly applicable to unawareness structures. But the remedy is straightforward. We say that φ is *defined in state* ω in M if $\omega \in \bigcap_{p \in \text{At}(\varphi)} (V(p) \cup \neg V(p))$. Now, we say φ is valid in M if $M, \omega \models \varphi$ for all ω in which φ is defined. φ is valid if it is valid in all M . Note that this generalized definition of validity is identical to the notion of validity for Kripke structures if $\mathcal{S} = \{S\}$, i.e., if the lattice of spaces of the unawareness structure is a singleton and thus the unawareness structure is a Kripke structure. The notions of soundness and completeness are now defined analogously to Kripke structures but using the generalized definition of validity.

Theorem 5 (Heifetz, Meier, and Schipper [2008]) *For the language $\mathbf{L}_n^{K,A}(\text{At})$, the axiom system $\widetilde{\mathbf{S}}_n^{K,A}$ is a sound and complete axiomatization with respect to unawareness structures.*

Halpern and Rêgo [2008] present an alternative axiomatization of unawareness structures in which they extend the language by adding a non-standard implication operator. Recall that in an unawareness structure a formula may not be defined at every state. Implicitly, the non-standard implication operator combines standard implication with an “at least as defined” relation on formulae. That is, formula φ implies (non-standardly) ψ is valid only if

ψ is true whenever φ is true and ψ is “at least as defined as” φ . Then they are able to define a satisfaction relation such that every formula in the extended language is defined in every state *across all* spaces. In such a setting, one can apply directly the definition of validity as for Kripke structures. The authors axiomatize propositional logic with respect to the non-standard implication operator. Then they axiomatize unawareness structures by an axiom system that is similar to **S5** but makes use of the non-standard implication operator. Halpern and Rêgo [2008] prove also analogous axiomatizations of unawareness structures in which the possibility correspondence does not necessarily satisfy generalized reflexivity and stationarity.

How are unawareness structures related to awareness structures introduced earlier? It turns out that despite the differences in motivation, their semantics are equivalent in terms of expressibility. That is, everything that can be described about awareness and knowledge in one structure can be described in the other structure and vice versa. More formally, let $\mathcal{M}^{HMS}(\text{At})$ be the class of unawareness structures over At .

Theorem 6 (Halpern and Rêgo [2008]) *For any partitional awareness structure $M = (\Omega, (R_a)_{a \in \text{Ag}}, (\mathcal{A}_a)_{a \in \text{Ag}}, V) \in \mathcal{M}^{FH}(\text{At})$ in which awareness is propositionally determined, there exists an unawareness structure $M' = (\Omega', (\Pi_a)_{a \in \text{Ag}}, (r_{\text{At}_1, \text{At}_2}^{\text{At}_1})_{\text{At}_1, \text{At}_2 \subseteq \text{At}}, V') \in \mathcal{M}^{HMS}(\text{At})$ such that $\Omega' := \Omega \times 2^{\text{At}}$, $S_{\text{At}_1} = \Omega \times \{\text{At}_1\}$ for all $\text{At}_1 \subseteq \text{At}$, and for all $\varphi \in \mathbf{L}_n^K(\text{At})$, if $\text{At}(\varphi) \subseteq \text{At}_1$, then $M, \omega \models \varphi$ if and only if $M', (\omega, \text{At}_1) \models \varphi$.*

Conversely, for every unawareness structure $M = (\Omega, (\Pi_a)_{a \in \text{Ag}}, (r_{\text{At}_1, \text{At}_2}^{\text{At}_1})_{\text{At}_1, \text{At}_2 \subseteq \text{At}}, V) \in \mathcal{M}^{HMS}(\text{At})$, there exists a partitional awareness structure $M' = (\Omega, (R_a)_{a \in \text{Ag}}, (\mathcal{A}_a)_{a \in \text{Ag}}, V') \in \mathcal{M}^{FH}(\text{At})$ in which awareness is propositionally determined such that for all $\varphi \in \mathbf{L}_n^K(\text{At})$, if $\omega \in S_{\text{At}_1}$ and $\text{At}(\varphi) \subseteq \text{At}_1$, then $(M, \omega) \models \varphi$ if and only if $(M', \omega) \models \varphi$.

The proof is by induction on the structure of φ . Halpern and Rêgo [2008] also prove versions of the result for unawareness structures with possibility correspondences (and awareness structures with accessibility relations, respectively) that not necessarily satisfy generalized reflexivity and stationary (reflexivity, transitivity, and Euclideaness, respectively).

The previous result implies alternative axiomatizations:

Corollary 1 *For the language $\mathbf{L}_n^{K,A}(\text{At})$, the axiom system $\mathbf{S5}_n^{K,A}$ is a sound and complete axiomatization with respect to unawareness structures. For the language $\mathbf{L}_n^{K,A}(\text{At})$, the axiom system $\widetilde{\mathbf{S5}}_n^{K,A}$ is a sound and complete axiomatization with respect to partitional awareness structures in which awareness is propositionally determined.*

As mentioned previously, Galanis [2013a] studies a variant of unawareness frames in which he drops the property Projections Preserve Knowledge of the possibility correspondence. Galanis [2011] axiomatizes his variant of unawareness structures with multiple knowledge modalities, one for each sub-language.

3.5 Generalized Standard Models by Modica and Rustichini [1999]

Modica and Rustichini [1999] were the first economists to present a semantics for propositionally determined awareness and partitional knowledge in the case of a single agent. In retrospect

we can understand unawareness structures introduced in the previous section as a multi-agent generalization of Modica and Rustichini [1999]. The exposition here follows mostly Halpern [2001].

A *generalized standard model* $M = (S, \Omega, \rho, \Pi, V)$ over At consists of a space of objective states S and a collection of nonempty disjoint subjective state spaces $\{\tilde{S}_{\text{At}'}\}_{\text{At}' \subseteq \text{At}}$ with $\Omega := \bigcup_{\text{At}' \subseteq \text{At}} \tilde{S}_{\text{At}'}$. Ω and S are disjoint. Further, there is a surjective projection $\rho : S \rightarrow \Omega$. Moreover, the agent has a generalized $\Pi : S \rightarrow 2^\Omega$ that satisfies:

Generalized Reflexivity: if $s \in S$, then $\Pi(s) \subseteq \tilde{S}_{\text{At}'}$ for some $\text{At}' \subseteq \text{At}$,
Stationarity: $\rho(s) = \rho(t)$, then $\Pi(s) = \Pi(t)$.

Finally, there is a valuation $V : \text{At} \rightarrow 2^S$ such that if $\rho(s) = \rho(t) \in \tilde{S}_{\text{At}'}$ then for all $p \in \text{At}'$ either $s, t \in V(p)$ or $s, t \notin V(p)$.

Intuitively, the states in the subjective state-space $\tilde{S}_{\text{At}'}$ describe situations conceivable by an agent who is aware of atomic formulae in At' only. Generalized reflexivity confines in each objective situation the perception of the agent to subjective situations that are all described with same “vocabulary”. Stationarity means that the agents’ perception depends only on her subjective states and summarizes transitivity and Euclideaness of the possibility correspondence.

Two caveats are to note: First, it is a single-agent structure. It is not immediate how to extent generalized standard models to a multi-agent setting. If we add additional possibility correspondences, one for each agent, then agents could reason about each other’s knowledge but presumably only at the same awareness level. At state $s \in S$, agent a may know that agent b does not know the event $E \subseteq \tilde{S}_{\text{At}'}$. But since at every state in $\tilde{S}_{\text{At}'}$, agent b ’s possibility set must be a subset of $\tilde{S}_{\text{At}'}$ as well, agent a is forced to know that b is aware of E . This is avoided in unawareness frames of Section 3.3 where at $\tilde{s} \in \tilde{S}_{\text{At}'}$, agent b ’s possibility set may be a subset of states in yet a lower space $\tilde{S}_{\text{At}''}$ with $\text{At}'' \subsetneq \text{At}'$. Generalized standard models are limited to the single-agent case. Yet, unawareness is especially interesting in interactive settings, where different agents may have different awareness and knowledge, and reason about each others awareness and knowledge.

Second, the condition on the valuation, if $\rho(s) = \rho(t) \in \tilde{S}_{\text{At}'}$ then for all $p \in \text{At}'$ either $s, t \in V(p)$ or $s, t \notin V(p)$, is also a condition on the projection ρ . Deleting the valuation does not yield straightforwardly an event-based approach or frame similar to Aumann structures, but one would need to add instead conditions on the projections.

We extend Π to a correspondence defined on the domain $S \cup \Omega$ by if $\tilde{s} \in \Omega$ and $\tilde{s} = \rho(s)$, then define $\Pi(\tilde{s}) = \Pi(s)$. This extension is well-defined by stationarity. A generalized standard model is said to be *partitional* if Π restricted to Ω is partitional. We also extend V to a valuation having the range $S \cup \Omega$ by defining $\tilde{V}(p) = V(p) \cup \bigcup_{\text{At}' \subseteq \text{At}} \left\{ \tilde{s} \in \tilde{S}_{\text{At}'} : p \in \text{At}', \rho^{-1}(\tilde{s}) \subseteq V(p) \right\}$.

For $\omega \in S \cup \Omega$, we define inductively on the structure of formulae in $\mathbf{L}_1^{K,A}(\text{At})$ the satisfaction relation

$M, \omega \models \top$, for all $\omega \in \Omega$,
 $M, \omega \models p$ if and only if $\omega \in \tilde{V}(p)$,
 $M, \omega \models \varphi \wedge \psi$ if and only if both $M, \omega \models \varphi$ and $M, \omega \models \psi$,

$M, \omega \models \neg\varphi$ if and only if $M, \omega \not\models \varphi$ and either $\omega \in S$ or $\omega \in \tilde{S}_{\text{At}'}$ and $\varphi \in \mathbb{L}_1^{K,A}(\text{At}')$,
 $M, \omega \models K\varphi$ if and only if $M, \omega' \models \varphi$ for all $\omega' \in \Pi(\omega)$.

Recall that Modica and Rustichini [1994, 1999] define $A\varphi := K\varphi \vee K\neg K\varphi$. Thus, indeed the satisfaction relation is defined for formulae in $\mathbb{L}_1^{K,A}(\text{At})$.

At this point, it may be useful to illustrate generalized standard models with our speculative trade example. Yet, since generalized standard models are defined for a single agent only, we cannot model the speculative trade example. While we could construct a separate generalized standard model for each of the agents, these models could not model the agent's reasoning about the other agent's awareness and knowledge. For instance, the sublattice consisting of the two spaces $S_{\{n,\ell\}}$ and $S_{\{\ell\}}$ in Figure 2 can be viewed as a generalized standard model of the owner.

To prove a characterization of generalized standard models in terms of properties of knowledge and awareness, we need to define validity. Recall that a formula is said to be valid in a Kripke structure if it is true in *every* state. Modica and Rustichini [1999] restrict the notion of validity to objective states in S only. We say φ is *objectively valid* in M if $M, \omega \models \varphi$ for all $\omega \in S$. The notions of soundness and completeness are now defined analogous to Kripke structures but using the notion of objective validity.

Modica and Rustichini [1999] consider the following axiom system that they call $\mathbf{U}$.

- Prop. All substitutions instances of tautologies of propositional logic, including the formula $\top$.
- AS. $A\neg\varphi \leftrightarrow A\varphi$ (Symmetry)
- AC. $A(\varphi \wedge \psi) \rightarrow A\varphi \wedge A\psi$
- T. $K\varphi \rightarrow \varphi$ (Axiom of Truth)
- 4. $K\varphi \rightarrow KK\varphi$ (Positive Introspection Axiom)
- MP. From φ and $\varphi \rightarrow \psi$ infer ψ (modus ponens)
- M, C. $K(\varphi \wedge \psi) \leftrightarrow K\varphi \wedge K\psi$ (Distribution)
- N. $K\top$
- RK_{sa}. From $\varphi \leftrightarrow \psi$ infer $K\varphi \leftrightarrow K\psi$, where φ and ψ are such that $\text{At}(\varphi) = \text{At}(\psi)$.

Theorem 7 (Modica and Rustichini [1999]) *The axiom system $\mathbf{U}$ is a complete and sound axiomatization of objective validity for the language $\mathbb{L}_1^{K,A}(\text{At})$ with respect to partitioned generalized standard models.*

When we restrict partitioned awareness structures that are propositionally determined to a single-agent, then those awareness structures and generalized standard structures are equally expressive. Everything that can be described about awareness and knowledge in a generalized standard model can be described in partitioned awareness structures in which awareness is propositionally determined. Let $\mathcal{M}^{MR}(\text{At})$ be the class of generalized standard models over At .

Theorem 8 (Halpern [2001]) *For any partitioned awareness structure $M = (S, R, \mathcal{A}, V) \in \mathcal{M}^{FH}(\text{At})$ in which awareness is propositionally determined, there exists a generalized standard model $M' = (S, \Omega, \Pi, \rho, V') \in \mathcal{M}^{MR}(\text{At})$ such that for all formulae $\varphi \in \mathbb{L}_1^K(\text{At})$, $M, s \models \varphi$ if and only if $M', s \models \varphi$.*

Conversely, for every generalized standard model $M = (S, \Omega, \Pi, \rho, V) \in \mathcal{M}^{MR}(\text{At})$, there exists a partitional awareness structure $M' = (S, R, \mathcal{A}, V') \in \mathcal{M}^{FH}(\text{At})$ in which awareness is propositionally determined such that for all $\varphi \in \mathcal{L}_1^K(\text{At})$, $M, s \models \varphi$ if and only if $M', s \models \varphi$.

Halpern [2001] proves also analogous results for generalized standard models (awareness structures, respectively) for which generalized reflexivity and stationarity (reflexivity, transitivity, and Euclideaness) may fail.

3.6 Product Models by Li [2009]

Li [2009] introduces what she calls product models by starting with a subset of questions about the relevant aspects of the world that can be answered either in the affirmative or negative. Awareness then differs by the subset of questions the agent has in mind. Such an approach to awareness is quite natural since lacking conception of some aspects of the world implies that one is not even able to ask questions about these aspects. In what follows, we slightly depart from her original exposition for better comparison.

Her original model was stated for a single agent only. The extension to the multi-agent setting is non-trivial (see Li [2008a]). We focus on the single-agent case but will consider the multi-agent case in the speculative trade example below. The primitives of the product model $M = (Q^*, \Omega^*, \mathcal{A}, P)$ are a set of questions Q^* and a space of objective states $\Omega^* := \prod_{q \in Q^*} \{1_q, 0_q\}$. Each state is a profile of zeros and ones; each component corresponding to a question in Q^* . The question is answered in the affirmative if the component corresponding to it is one. Further, an awareness correspondence $\mathcal{A} : \Omega^* \rightarrow 2^{Q^*}$ assigns to each state a subset of questions that the agent is aware of at that state. Finally, a possibility correspondence $P : \Omega^* \rightarrow 2^{\Omega^*} \setminus \{\emptyset\}$ assigns to each state a subset of states that the agent implicitly considers possible. We assume

Reflexivity: $\omega^* \in P(\omega^*)$ for all $\omega^* \in \Omega^*$.

Stationarity: For all $\omega_1^*, \omega_2^* \in \Omega^*$, $\omega_1^* \in P(\omega_2^*)$ implies both $P(\omega_1^*) = P(\omega_2^*)$ and $\mathcal{A}(\omega_1^*) = \mathcal{A}(\omega_2^*)$.

For $Q \subseteq Q^*$, let $\Omega_Q = \prod_{q \in Q} \{1_q, 0_q\}$ denote the state space in which states contain answers to questions in the subset Q only. Using the awareness correspondence, the subjective state-space of the agent at $\omega^* \in \Omega^*$ is $\Omega_{\mathcal{A}(\omega^*)}$.

Let $? : \bigcup_{Q \subseteq Q^*} 2^{\Omega_Q} \rightarrow 2^{Q^*}$ denote the correspondence that assigns to each subset of each state space the set of questions that define the space. That is, if $E \subseteq \Omega_Q$, then $?(E) = Q$.

For $Q \subseteq Q^*$, denote by $r_Q : \Omega^* \rightarrow \Omega_Q$ the surjective projection. If $E \subseteq \Omega_Q$, denote by $E^* = r_Q^{-1}(E)$ the set of objective states in Ω^* that project to E , i.e., the inverse image of E in Ω^* . An event is a pair $E = (E^*, ?(E))$. We define $\neg E := (\neg E^*, ?(E))$, $E_1 \wedge E_2 := (E_1^* \cap E_2^*, ?(E_1) \cup ?(E_2))$. Disjunction is defined by the De Morgan law using negation and conjunction as just defined. Note that there are many vacuous events $\emptyset_Q := \neg \Omega_Q$, $Q \subseteq Q^*$, one for each subset of questions.

The unawareness operator is defined by

$$\mathbf{U}(E) = \{\omega^* \in \Omega^* : ?(E) \not\subseteq \mathcal{A}(\omega^*)\}.$$

The awareness operator is defined by $\mathbf{A}(E) := \neg \mathbf{U}(E)$. Note that for every event E , the $\mathbf{U}(E)$ and $\mathbf{A}(E)$ are subsets of the objective space. Thus, they capture an agent’s reasoning about awareness from an outside modeler’s point of view.

There are two knowledge operators. As our notation suggests, “objective knowledge” in Li [2009] is best understood as implicit knowledge

$$\mathbf{L}(E) := \{\omega^* \in \Omega^* : P(\omega^*) \subseteq E^*\}.$$

The second knowledge operator refers to “subjective knowledge from the modeler’s perspective”. Although Li [2009] states it differently, it is equivalent to explicit knowledge

$$\mathbf{K}(E) := \mathbf{L}(E) \cap \mathbf{A}(E).$$

Note however, that both $\mathbf{L}(E)$ and $\mathbf{K}(E)$ are subsets of the objective state-space and therefore not necessarily “accessible” to the agent. As remedy, Li [2009] also defines a subjective possibility correspondence and uses it to define a knowledge operator reflecting “subjective knowledge from the agent’s perspective” (see also Heinsalu [2012]). Yet, since she states her results in terms of implicit and explicit knowledge, we focus on implicit and explicit knowledge only but will illustrate also the subjective versions in the speculative trade example below.

Proposition 4 (Li [2009]) *For the product model, the following properties obtain for any events E, E_1, E_2 and $\omega^* \in \Omega^*$,*


- Subjective Necessitation:* $\omega^* \in \mathbf{K}(\Omega(\omega^*))$,
- Distribution:* $\mathbf{K}(E_1) \cap \mathbf{K}(E_2) = \mathbf{K}(E_1 \wedge E_2)$,
- Monotonicity:* $E_1^* \subseteq E_2^*$ and $?(E_1) \supseteq ?(E_2)$ implies $\mathbf{K}(E_1) \subseteq \mathbf{K}(E_2)$,
- Truth:* $\mathbf{K}(E) \subseteq E^*$,
- Positive Introspection:* $\mathbf{K}(E) \subseteq \mathbf{K}\mathbf{K}(E)$,
- Weak Negative Introspection:* $\neg \mathbf{K}(E) \cap \mathbf{A}(E) = \mathbf{K}\neg \mathbf{K}(E)$,
- MR Awareness:* $\mathbf{A}(E) = \mathbf{K}(E) \cup \mathbf{K}\neg \mathbf{K}(E)$,
- Strong Plausibility:* $\mathbf{U}(E) = \bigcap_{n=1}^{\infty} (\neg \mathbf{K})^n(E)$,
- KU Introspection:* $\mathbf{K}\mathbf{U}(E) = \emptyset_{Q^*}$,
- AU Reflection:* $\mathbf{U}(E) = \mathbf{U}\mathbf{U}(E)$,
- Symmetry:* $\mathbf{A}(E) = \mathbf{A}(\neg E)$.

The proof follows from definitions and properties of the possibility correspondence.

At this point, it may be instructive to consider as an illustration the speculative trade example.

Speculative Trade Example (continued). Li [2009] introduces the product model for a single agent only. Thus, we cannot use it to model the speculative trade example. Yet, in an unpublished paper, Li [2008a] presents a multi-agent extension. We will illustrate the multi-agent extension with the speculative trade example. The set of questions is $\{n, \ell\}$, where we let n and ℓ stand for the questions “Is the innovation true?” and “Is the lawsuit true?”, respectively. The objective state space is $\Omega^* = \Omega_{\{n, \ell\}} = \{1_n, 0_n\} \times \{1_\ell, 0_\ell\}$; the upmost space in Figure 3. For instance, at the state $(1_n, 0_\ell)$ the question “Is the innovation true?” is answered in

Figure 3: A Product Model for the Speculative Trade Example



the affirmative, “The innovation is true.”, while the question “Is the lawsuit true?” is answered in the negative, “The lawsuit is not true.”

The awareness correspondences are indicated in Figure 3 by “speech bubbles” above each state. The solid blue speech bubbles belong to the owner, while the intermitted red speech bubbles are the buyer’s. Both awareness correspondences are very special as they are constant on Ω^* . At every state Ω^* , the owner is aware only of questions involving the lawsuit while the buyer is only aware of questions involving the innovation. The possibility correspondences are indicated in Figure 3 by the solid blue and intermitted red soft-edged rectangles for the owner and buyer, respectively. Again, the possibility correspondences are very special in this example as no agent can distinguish any objective states in Ω^* .

Given the awareness correspondences defined on the set of objective states Ω^* , we can construct the subjective state spaces of both agents by considering for each agent only the questions of which he is aware. At every state in Ω^* , the buyer’s subjective state space is the space to the left, $\Omega_b(\Omega^*)$, while the owner’s subjective state space is the space to the right, $\Omega_o(\Omega^*)$. So far, these are all the primitives of the product model. Li [2008a, 2009] also defines subjective versions of the awareness and possibility correspondences that do not play a role in her result (Proposition 4), but that are useful for modeling the example. To define the subjective awareness correspondence on $\Omega_b(\Omega^*)$, we extend the objective awareness correspondence on Ω^* to the subjective states in $\Omega_b(\Omega^*)$ and $\Omega_o(\Omega^*)$ by restricting the awareness sets at the subjective states to questions available at those subjective states, respectively.¹⁸ Similarly,

¹⁸Note that, for instance, the owner’s awareness correspondence on $\Omega_b(\Omega^*)$ cannot really be interpreted as the owner’s subjective awareness correspondence but rather as the owner’s awareness correspondence as perceived

we can extend the objective possibility correspondences defined on Ω^* to subjective states by taking the projections to $\Omega_b(\Omega^*)$ and $\Omega_o(\Omega^*)$, respectively.

The subjective awareness correspondences allow us to defined another subjective state space shown as the lowest space in Figure 3. At every state in $\Omega_b(\Omega^*)$, the owner is unaware of the innovation (and the lawsuit), hence his subjective state space (in the eyes of the buyer) is $\Omega_o(\Omega_b(\Omega^*))$. Similarly, at every state in the owner's subjective state space $\Omega_o(\Omega^*)$, the buyer is unaware of the lawsuit (and the innovation); thus his subjective state space (in the eyes of the owner) is $\Omega_b(\Omega_o(\Omega^*))$. Both spaces are defined from an empty set of questions. They are identical and singleton. Again, we can extend the awareness and possibility correspondences to the lowest space as outlined above.

This example illustrates the multi-agent product model of Li [2008a]. As it should be clear by now, it models the introductory example of speculative trade. Moreover, it suggests that Li's model bears features both of awareness structures by Fagin and Halpern [1988] and unawareness frames by Heifetz, Meier, and Schipper [2006]. First, with awareness structures it has in common that awareness is modeled separately with an awareness correspondence although questions are used as primitive instead atomic formulae. Since one can define a one-to-one relation between questions and atomic formulae, the upmost space is analogous to the awareness structure depicted in Figure 1. We will use this relationship more generally in the discussion below. Second, we see clearly that the possibility correspondence models implicit knowledge and not necessarily explicit knowledge. For instance, the possibility sets on the objective space Ω^* and the owner's (the buyer's, resp.) possibility sets on $\Omega_b(\Omega^*)$ (on $\Omega_o(\Omega^*)$, respectively) can be understood only in terms of implicit knowledge. With unawareness frames it has in common the idea of subjective states and the lattice structure. $\square$

The following discussion is confined to the single-agent product model of Li [2009]. The exposition follows mostly Heinsalu [2012]. Li [2009] does not present an axiomatization of product models. The product model is analogous to a frame but with an additional set of questions as a primitive. In order to extend it to a structure, we need to relate questions to formulae and introduce a valuation. Let $b : Q^* \rightarrow \mathbf{At}$ be a bijection. For every question $q \in Q^*$ there is exactly one primitive proposition p such that $b(q) = p$. The bijection is interpreted as assigning to each question $q \in Q^*$ exactly one primitive proposition $p \in \mathbf{At}$ that stands for “ q is answered affirmatively” and q stands for “Is p true?”. We can consider now the language $\mathcal{L}_1^{L,K,A}(b(Q^*))$.

A valuation $V : b(Q^*) \rightarrow 2^{\Omega^*}$ is defined by $V(p) = \{\omega^* \in \Omega^* : r_{\{b^{-1}(p)\}}(\omega^*) = 1\}$. The corresponding subjective event is $[p] = (V(p), \{b^{-1}(p)\})$. Note $[p]^* = V(p)$. The satisfaction relation is defined by induction on the structure of formulae:

$$\begin{aligned} M, \omega^* &\models p \text{ if and only if } \omega^* \in V(p), \\ M, \omega^* &\models \neg\varphi \text{ if and only if } \omega^* \in \neg[\varphi]^*, \\ M, \omega^* &\models \varphi \wedge \psi \text{ if and only if } \omega^* \in [\varphi]^* \wedge [\psi]^*, \\ M, \omega^* &\models A\varphi \text{ if and only if } \omega^* \in \mathbf{A}[\varphi], \\ M, \omega^* &\models L\varphi \text{ if and only if } \omega^* \in \mathbf{L}[\varphi], \\ M, \omega^* &\models K\varphi \text{ if and only if } \omega^* \in \mathbf{K}[\varphi]. \end{aligned}$$

by the buyer.

Note that the satisfaction relation is defined only for objective states in Ω^* . Thus, similar to Modica and Rustichini [1999] the setting so far allows for a notion of objective validity only. We can relate product models to structures discussed in the previous sections. Let $\mathcal{M}^{Li}$ denote the class of product models over the set of questions Q^* .

Theorem 9 (Heinsalu [2012]) *For any partitioned awareness structure $M = (S, R, \mathcal{A}, V) \in \mathcal{M}^{FH}(\text{At})$ in which awareness is generated by primitive propositions there is a product model $M = (b^{-1}(\text{At}), S, P, \mathcal{A}', V') \in \mathcal{M}^{Li}$ with $\mathcal{A}' := b^{-1} \circ \mathcal{A}$ such that for any $\varphi \in \mathcal{L}_1^{L,K,A}(\text{At})$, $M, s \models \varphi$ if and only if $M', s \models \varphi$.*

Conversely, for every product model $M = (Q^, \Omega^*, P, \mathcal{A}) \in \mathcal{M}^{Li}$ there is a partitioned awareness structure $M' = (\Omega^*, R, \mathcal{A}', V') \in \mathcal{M}^{FH}(b(Q^*))$ in which awareness is generated by primitive propositions such that for any $\varphi \in \mathcal{L}_1^{L,K,A}(b(Q^*))$, $M, \omega^* \models \varphi$ if and only if $M', \omega^* \models \varphi$.*

Heinsalu [2012] proves the theorem for product models (awareness structures, respectively) for which reflexivity and stationarity (reflexivity, transitivity, and Euclideaness) may fail. But it is straightforward to extend it to any corresponding subset of those properties. The relationship between the multi-agent extension of the product model by Li [2008a] and the rest of the literature is still open.

3.7 Synopsis of Other Approaches: Pires [1994], Ewerhart [2001], and Feinberg [2004, 2005, 2012]

Around the same time as Modica and Rustichini published the first paper on awareness in economics (Modica and Rustichini [1994]), Pires finished her doctoral dissertation in economics at MIT with an unpublished chapter on awareness that unfortunately has been ignored in the literature so far. Pires [1994] presents a model of non-trivial awareness for a single-agent that essentially captures awareness generated by primitive propositions. She already considers weak negative introspection, weak necessitation, and plausibility as properties of awareness and knowledge. Although she introduces both a logic and a state-space semantics, she has no soundness or completeness result. She also anticipates modeling awareness of unawareness very much in the spirit of later works by Ågotnes and N. [2007] and Walker [2014]. Finally, she studies updating of awareness as refinements of conceivable states.

Ewerhart [2001] introduces a state-space model in which at each state an agent may only be aware of a subset of states. But since no special event structure is assumed, an agent may be aware of an event but unaware of its complement or vice versa, thus violating symmetry. Ewerhart [2001] considers both implicit and explicit knowledge and the model satisfies, for instance, weak negative introspection with respect to explicit knowledge but not weak necessitation. Under an additional richness assumption, it satisfies KU-introspection and AU-introspection, strong plausibility with “ $\subseteq$ ” but not the Modica-Rustichini definition of awareness unless unawareness is trivial. He proves a generalization of Aumann’s “No-agreeing-to-disagree” theorem for his models with unawareness.

Feinberg [2004, 2005, 2012] provides different versions of an approach that models interactive awareness of components of games by explicit unbounded sequences of mutual views of players. Among the properties imposed on awareness is that (1) if a player is aware of what an opponent is aware of, then the player herself is also aware of it, and (2), if a player (“she”) is aware that an

opponent (“he”) is aware of something, then she is also aware that the opponent is aware that he is aware of it. These two properties are satisfied also by the notion of propositionally determined awareness discussed earlier. Yet, the precise connection between Feinberg’s approach and the rest of the literature is still open.

4 Awareness and Probabilistic Beliefs

4.1 Type Spaces with Unawareness

Models of unawareness are mostly applied in strategic contexts when agents are players in a game and have to take decisions that are rational/optimal with respect to their state of mind. In such situations, it is extremely helpful for players to judge uncertain events probabilistically. To model such agents, we need to replace the qualitative notions of knowledge or belief discussed until now by the quantitative notion of probabilistic beliefs. In standard game theory with incomplete information, this is done with Harsanyi type spaces (Harsanyi [1967/68]). Type spaces do not just model in a parsimonious way a player’s belief about some basic uncertain events but also their beliefs about other players’ beliefs, beliefs about that, and so on. That is, they model infinite hierarchies of beliefs. Under unawareness, the problem is complicated by the fact that agents may also be unaware of different events and may form beliefs about other players’ unawareness, their belief about other players’ beliefs about unawareness, etc. Combining ideas from unawareness frames and Harsanyi type spaces, Heifetz, Meier, and Schipper [2013a] define an unawareness type space $(\mathcal{S}, (r_S^{S'})_{S' \succeq S; S, S' \in \mathcal{S}}, (t_a)_{a \in \mathbf{Ag}})$ by a complete lattice of disjoint measurable spaces $\mathcal{S} = \{S_\alpha\}_{\alpha \in \mathcal{A}}$, each with a σ -field $\mathcal{F}_S$, and measurable surjective projections $(r_S^{S'})_{S' \succeq S; S, S' \in \mathcal{S}}$. Let $\Delta(S)$ be the set of probability measures on $(S, \mathcal{F}_S)$. We consider this set itself as a measurable space endowed with the σ -field $\mathcal{F}_{\Delta(S)}$ generated by the sets $\{\mu \in \Delta(S) : \mu(D) \geq p\}$, where $D \in \mathcal{F}_S$ and $p \in [0, 1]$.

For a probability measure $\mu \in \Delta(S')$, the marginal $\mu|_S$ of μ on $S \preceq S'$ is defined by

$$\mu|_S(D) := \mu \left(\left(r_S^{S'} \right)^{-1} (D) \right), \quad D \in \mathcal{F}_S.$$

Let S_μ be the space on which μ is a probability measure. Whenever $S_\mu \succeq S(E)$, we abuse notation slightly and write

$$\mu(E) = \mu(E \cap S_\mu).$$

If $S(E) \not\succeq S_\mu$, then we say that $\mu(E)$ is undefined.

For each agent $a \in \mathbf{Ag}$, there is a *type mapping* $t_a : \Omega \rightarrow \bigcup_{\alpha \in \mathcal{A}} \Delta(S_\alpha)$, which is measurable in the sense that for every $S \in \mathcal{S}$ and $Q \in \mathcal{F}_{\Delta(S)}$ we require $t_a^{-1}(Q) \cap S \in \mathcal{F}_S$. Analogous to properties of the possibility correspondence in unawareness frames, the type mapping t_a should satisfy the following properties:¹⁹

Confinement: If $\omega \in S'$ then $t_a(\omega) \in \Delta(S)$ for some $S \preceq S'$.

¹⁹Heifetz, Meier, and Schipper [2013a] introduce also an additional property and show that it is implied by the other properties.

- (2) If $S'' \succeq S' \succeq S$, $\omega \in S''$, and $t_a(\omega) \in \Delta(S')$ then $t_a(\omega_S) = t_a(\omega)|_S$.
- (3) If $S'' \succeq S' \succeq S$, $\omega \in S''$, and $t_a(\omega_{S'}) \in \Delta(S)$ then $S_{t_a(\omega)} \succeq S$.

$t_a(\omega)$ represents agent a 's belief at state ω . The properties guarantee the consistent fit of beliefs and awareness at different state spaces. *Confinement* means that at any given state $\omega \in \Omega$ an agent's belief is concentrated on states that are all described with the same "vocabulary" - the "vocabulary" available to the agent at ω . This "vocabulary" may be less expressive than the "vocabulary" used to describe statements in the state ω .

Properties (2) and (3) compare the types of an agent in a state $\omega \in S'$ and its projection to ω_S , for some $S \preceq S'$. Property (2) means that at the projected state ω_S the agent believes everything she believes at ω given that she is aware of it at ω_S . Property (3) means that at ω an agent cannot be unaware of an event that she is aware of at the projected state $\omega_{S'}$.

Define the set of states at which agent a 's type or the marginal thereof coincides with her type at ω by $Ben_a(\omega) := \{\omega' \in \Omega : t_a(\omega')|_{S_{t_a(\omega)}} = t_a(\omega)\}$. This is an event of the unawareness-belief frame although it may not be a measurable event (even in a standard type-space). It is assumed that if $Ben_a(\omega) \subseteq E$, for an event E , then $t_a(\omega)(E) = 1$. This assumption implies introspection with respect to beliefs.

For agent $a \in \mathbf{Ag}$ and an (not necessarily measurable) event E , define the *awareness operator* by

$$\mathbf{A}_a(E) := \{\omega \in \Omega : t_a(\omega) \in \Delta(S), S \succeq S(E)\}$$

if there is a state ω such that $t_a(\omega) \in \Delta(S)$ with $S \succeq S(E)$, and by $\mathbf{A}_a(E) := \emptyset^{S(E)}$ otherwise. This is analogous to awareness in unawareness frames.

For each agent $a \in \mathbf{Ag}$, $p \in [0, 1]$, and measurable event E , the *probability-of-at-least- p -belief operator* is defined as usual (see for instance Monderer and Samet [1989]) by

$$\mathbf{B}_a^p(E) := \{\omega \in \Omega : t_a(\omega)(E) \geq p\},$$

if there is a state ω such that $t_a(\omega)(E) \geq p$, and by $\mathbf{B}_a^p(E) := \emptyset^{S(E)}$ otherwise.

Lemma 3 *If E is an event, then both $\mathbf{A}_a(E)$ and $\mathbf{B}_a^p(E)$ are $S(E)$ -based events.*

The proof follows from the properties of the type mapping and the definitions.

The *unawareness operator* is defined by $\mathbf{U}_a(E) := \neg \mathbf{A}_a(E)$.

Let $\mathbf{Ag}$ be an at most countable set of agents. Interactive beliefs are defined as usual (e.g. Monderer and Samet [1989]). The *mutual p -belief operator* $\mathbf{B}^p$ is defined analogously to the mutual knowledge operator in Section 3.3 with $\mathbf{K}_a$ replaced by $\mathbf{B}_a$. The *common certainty operator* $\mathbf{CB}^1$ is defined analogously to the common knowledge operator but with $\mathbf{K}$ replaced by $\mathbf{B}^1$.

Proposition 5 (Heifetz, Meier, and Schipper [2013a]) *Let E and F be events, $\{E_l\}_{l=1,2,\dots}$ be an at most countable collection of events, and $p, q \in [0, 1]$. The following properties of belief obtain:*

- (o) $\mathbf{B}_a^p(E) \subseteq \mathbf{B}_a^q(E)$, for $q \leq p$,

- (i) *Necessitation*: $\mathbf{B}_a^1(\Omega) = \Omega$,
- (ii) *Additivity*: $\mathbf{B}_a^p(E) \subseteq \neg \mathbf{B}_a^q(\neg E)$, for $p + q > 1$,
- (iii a) $\mathbf{B}_a^p(\bigcap_{l=1}^{\infty} E_l) \subseteq \bigcap_{l=1}^{\infty} \mathbf{B}_a^p(E_l)$,
- (iii b) for any decreasing sequence of events $\{E_l\}_{l=1}^{\infty}$, $\mathbf{B}_a^p(\bigcap_{l=1}^{\infty} E_l) = \bigcap_{l=1}^{\infty} \mathbf{B}_a^p(E_l)$,
- (iii c) $\mathbf{B}_a^1(\bigcap_{l=1}^{\infty} E_l) = \bigcap_{l=1}^{\infty} \mathbf{B}_a^1(E_l)$,
- (iv) *Monotonicity*: $E \subseteq F$ implies $\mathbf{B}_a^p(E) \subseteq \mathbf{B}_a^p(F)$,
- (va) *Introspection*: $\mathbf{B}_a^p(E) \subseteq \mathbf{B}_a^1 \mathbf{B}_a^p(E)$,
- (vb) *Introspection II*: $\mathbf{B}_a^p \mathbf{B}_a^q(E) \subseteq \mathbf{B}_a^q(E)$, for $p > 0$.

Proposition 6 (Heifetz, Meier, and Schipper [2013a]) *Let E be an event and $p, q \in [0, 1]$. The following properties of awareness and belief obtain:*

1. *Plausibility*: $\mathbf{U}_a(E) \subseteq \neg \mathbf{B}_a^p(E) \cap \neg \mathbf{B}_a^p \neg \mathbf{B}_a^p(E)$,
2. *Strong Plausibility*: $\mathbf{U}_a(E) \subseteq \bigcap_{n=1}^{\infty} (\neg \mathbf{B}_a^p)^n(E)$,
3. *B^p U Introspection*: $\mathbf{B}_a^p \mathbf{U}_a(E) = \emptyset^{S(E)}$ for $p \in (0, 1]$ and $\mathbf{B}_a^0 \mathbf{U}_a(E) = \mathbf{A}_a(E)$,
4. *AU Reflection*: $\mathbf{U}_a(E) = \mathbf{U}_a \mathbf{U}_a(E)$,
5. *Weak Necessitation*: $\mathbf{A}_a(E) = \mathbf{B}_a^1(S(E)^\uparrow)$,
6. $\mathbf{B}_a^p(E) \subseteq \mathbf{A}_a(E)$ and $\mathbf{B}_a^0(E) = \mathbf{A}_a(E)$,
7. $\mathbf{B}_a^p(E) \subseteq \mathbf{A}_a \mathbf{B}_a^q(E)$,
8. *Symmetry*: $\mathbf{A}_a(E) = \mathbf{A}_a(\neg E)$,
9. *A Conjunction*: $\bigcap_{\lambda \in L} \mathbf{A}_a(E_\lambda) = \mathbf{A}_a(\bigcap_{\lambda \in L} E_\lambda)$,
10. *AB^p Reflection*: $\mathbf{A}_a \mathbf{B}_a^p(E) = \mathbf{A}_a(E)$,
11. *Awareness Reflection*: $\mathbf{A}_a \mathbf{A}_a(E) = \mathbf{A}_a(E)$, and
12. $\mathbf{B}_a^p \mathbf{A}_a(E) = \mathbf{A}_a(E)$.

Proposition 7 (Heifetz, Meier, and Schipper [2013a]) *Let E be an event and $p, q \in [0, 1]$. The following multi-person properties obtain:*

- | | |
|---|---|
| 1. $\mathbf{A}_a(E) = \mathbf{A}_a \mathbf{A}_b(E)$, | 7. $\mathbf{B}^p(E) \subseteq \mathbf{CA}(E)$,
$\mathbf{B}^0(E) = \mathbf{CA}(E)$, |
| 2. $\mathbf{A}_a(E) = \mathbf{A}_a \mathbf{B}_b^p(E)$, | 8. $\mathbf{B}^p(E) \subseteq \mathbf{A}(E)$,
$\mathbf{B}^0(E) = \mathbf{A}(E)$, |
| 3. $\mathbf{B}_a^p(E) \subseteq \mathbf{A}_a \mathbf{B}_b^q(E)$, | 9. $\mathbf{A}(E) = \mathbf{B}^1(S(E)^\uparrow)$, |
| 4. $\mathbf{B}_a^p(E) \subseteq \mathbf{A}_a \mathbf{A}_b(E)$, | 10. $\mathbf{CA}(E) = \mathbf{B}^1(S(E)^\uparrow)$, |
| 5. $\mathbf{CA}(E) = \mathbf{A}(E)$, | 11. $\mathbf{CB}^1(S(E)^\uparrow) \subseteq \mathbf{A}(E)$, |
| 6. $\mathbf{CB}^1(E) \subseteq \mathbf{CA}(E)$, | 12. $\mathbf{CB}^1(S(E)^\uparrow) \subseteq \mathbf{CA}(E)$, |

We conclude that unawareness type spaces are the probabilistic analogue to unawareness frames.

4.2 Universal Type Space with Unawareness

Unawareness type spaces capture unawareness and beliefs, beliefs about beliefs (including beliefs about unawareness), beliefs about that and so on in a parsimonious way familiar from standard type spaces. That is, hierarchies of beliefs are captured implicitly by states and type mappings. This begs two questions: First, can we construct unawareness type spaces from

explicit hierarchies of beliefs? Such a construction, if possible, is complicated by the multiple awareness levels involved. Player 1 with a certain awareness level may believe that player 2 has a lower awareness level. Moreover, he may believe that player 2 believes that player 1 has yet an even lower awareness level, etc. The second question that arises is whether there exists a universal unawareness type space in the sense that *every* hierarchy of beliefs is represented in it. In using such a universal type space for an application, the modeler ensures that she can analyze any hierarchy of beliefs. For standard type spaces these questions have been answered by Mertens and Zamir [1985] for the case when the space of underlying uncertainties is compact Hausdorff and beliefs are regular probability measures. Heifetz and Samet [1998] drop the topological assumptions and assume instead that the space of underlying uncertainties is measurable. This latter result has been generalized and reformulated in a category theoretic setting by Moss and Viglizzo [2004], who show a connection to coalgebraic modal logic.

Similar approaches can be taken for unawareness type spaces. Heifetz, Meier, and Schipper [2012] present the hierarchical construction and show the existence of a universal unawareness type space analogous to Mertens and Zamir [1985]. The advantage of the topological case over the measure theoretic case is that it is constructive. This is especially helpful for unawareness where hierarchies of beliefs are complicated by the presence of multiple awareness levels. Heinsalu proved independently the measurable case (see also Pinter and Udvari [2012]).


The presentation of unawareness type spaces is somewhat divorced from awareness structures and unawareness structures presented earlier. Those structures we could axiomatize. We could describe in minute detail knowledge and awareness of all agents in each state. While the hierarchical construction of unawareness type spaces by Heifetz, Meier, and Schipper [2012] retains the flavor of explicit descriptions of beliefs, it begs the question of whether unawareness type spaces could be axiomatized using a logic with modal operators p_a^μ interpreted as “agent a assigns probability at least μ ”, for rational numbers $\mu \in [0, 1]$. That is, can we axiomatize the probabilistic analogue of awareness structures or unawareness structures? Fagin, Halpern, and Meggido [1990] and Heifetz and Mongin [2001] axiomatized the class of standard type spaces without unawareness. But they do so in terms of a purely finitary logic that won’t allow for *strong* soundness and *strong* completeness. Meier [2012] circumvents the problem and devises an infinitary axiom system that he shows to be strongly sound and strongly complete for standard type spaces without unawareness. Sadzik [2007] presents extensions of both awareness and unawareness structures to the probabilistic cases and provides axiomatizations. In a recent paper, Cozic [2012] also extends Heifetz and Mongin [2001] to the case of unawareness of a single-agent analogous to generalized standard structures of Modica and Rustichini [1999]. An extension of Meier [2012] to unawareness and to multi-agent settings with unawareness is still open.

4.3 Speculative Trade and Agreement

In this section we revisit the speculative trade example discussed earlier. Unawareness type spaces allow us now to provide an answer to the question posed in the introduction, namely whether at a price of \$100 per share the owner is going to sell to the buyer. If this question is answered in the affirmative, then under unawareness we have a counterexample to the “No-speculative-trade” for standard structures (e.g. Milgrom and Stokey [1982]), thus illustrating that asymmetric awareness may have different implications from asymmetric (standard) infor-

mation. In standard structures, if there is a common prior probability (i.e., common among agents), then common certainty of willingness to trade implies that agents are indifferent to trade. To address how our example fits to the “No-speculative-trade” theorems, we need to recast it into an unawareness type space with a common prior. This is illustrated in Figure 4.

Figure 4: An Unawareness Type Space with a Common Prior for the Speculative Trade Example



The type-mappings are represented in Figure 4 as follows. At any state in the upmost space $S_{\{n, \ell\}}$, the buyer’s belief has full support on the left space $S_{\{n\}}$ given by the red intermitted soft-edged rectangle and the owner’s belief has full support on $S_{\{\ell\}}$ given by the solid blue soft-edge rectangle. At any state in $S_{\{n\}}$ the owner’s belief has full support on the lowest space $S_{\emptyset}$. Analogously, the owner is certain that the buyer is unaware of the law suit since at any state in $S_{\{\ell\}}$ the belief of the buyer has full support on the space $S_{\emptyset}$. This example is analogous to Figure 2 except that the supports of types are displayed rather the possibility sets and we write to the left of each state its common prior probability as well. For instance, state (n, ℓ) has common prior probability $\frac{1}{4}$. We see that, for instance, the common prior on S_n is the marginal of the common prior on $S_{\{n, \ell\}}$. Indeed, the common prior in unawareness type spaces generally constitutes a projective system of probability measures. Both agents’ beliefs are consistent with the common prior. Of course, referring to a “prior” is misleading terminology under unawareness as it is nonsensical to think of a prior stage at which all agents are aware of all states while at the interim stage, after they received their type, they are unaware of some events. Rather than understanding the prior as a primitive of the model, it should be considered as derived from the types of players. As in standard structures (see for instance Samet [1999]), it is a convex combination of types (for the definition and discussion of the common prior under unawareness, see Heifetz, Meier, and Schipper [2013a]).

Say that the buyer prefers to buy at price x if his expected value of the firm is at least x , while the owner prefers to sell at price x if her expected value is at most x . The buyer strictly prefers to buy at price x if his expected value of the firm is strictly above x , while the owner strictly prefers to sell at price x if her expected value is strictly below x . Note that at any state in $S_{\{n, \ell\}}$, the owner’s beliefs are concentrated on $S_{\{\ell\}}$ and thus his expected value of a firm’s

share is \$90. Similarly, at any state in $S_{\{n,\ell\}}$, the buyer's beliefs are concentrated on $S_{\{n\}}$ and thus her expected value of a firm's share is \$110. Thus the owner strictly prefers to sell at the price \$100 while the buyer strictly prefers to buy at the price of \$100. Moreover, at the price \$100 it is common certainty that each agent prefers to trade because each agent strictly prefers to trade at \$100 and is certain that the other agent is indifferent between trading or not at \$100. Hence, we have a common prior, common certainty of willingness to trade but each agent has a strict preference to trade. We conclude that speculative trade is possible under asymmetric awareness while it is ruled out under symmetric awareness by the standard “No-speculative-trade” theorems (i.e., see Milgrom and Stokey [1982]).

At a second glance, we realize that speculative trade is a knife-edge case in this example. Suppose that there are some transaction costs. For instance, the government may require the buyer to pay a tax of \$1 per share. Then the owner knows that the buyer is not just indifferent between buying or not but must have a strict preference to trade as well (similar for the buyer). This leads to the question of whether the common prior assumption rules out common certainty of strict preference to trade. Heifetz, Meier, and Schipper [2013a] (Theorem 1) prove for finite unawareness type spaces that a non-degenerate common prior rules out common certainty of strict preference to trade. This result has been extended to infinite unawareness type spaces by Meier and Schipper. Thus, arbitrary small transaction costs such as the famous Tobin tax on transactions rule out speculation under unawareness. The “No-speculative-trade” result under unawareness is also relevant for the following reason: one may casually conjecture that any behavior is possible when awareness is allowed to vary among agents and thus behavior under unawareness may have no testable predictions. The “No-speculative-trade” result under unawareness shows that this is not the case.

The common prior assumption is a sufficient condition for the “No-speculative-trade” result. This begs the question of whether it would also be necessary. In standard state-spaces, the absence of speculative trade implies a common prior (see for instance Feinberg [2000], Heifetz [2006]). This is the converse to the “No-speculative-trade” theorem. Such a converse is desirable because it provides a conceptual interpretation of the common prior assumption. Yet, Heifetz, Meier, and Schipper [2013a] show a counterexample to the converse of the “No-speculative-trade” theorem under unawareness.

The “No-speculative-trade” results for standard structures without unawareness extend Aumann's famous “No-agreeing-to-disagree” result according to which if agents share a common prior probability measure, then it cannot be common knowledge that their posteriors disagree (Aumann [1976]). Heifetz, Meier, and Schipper [2013a] prove also a generalization of Aumann [1976]'s “No-agreeing-to-disagree” result to unawareness structures.

5 Awareness of Unawareness

According to KU introspection, an agent never knows or believes that she is unaware of a specific event. This does not mean that she couldn't know that she is unaware of *something*. There is a difference between knowing (or not knowing) that you are unaware of *the proposition* φ and knowing (or not knowing) that *there exists some proposition* that you are unaware of. To borrow an example from Halpern and Rêgo [2009], a primary care physician may be unaware of a specific disease and may not even realize that she is unaware of this specific disease. But

she may refer a patient to a specialist because she believes that the specialist is aware of some diseases that she doesn't even think or have heard about. Grant and Quiggin [2013] argue that agents should generally induce from prior experience and experience with other agents that they may be unaware of something. All previous approaches outlined so far are silent on awareness of unawareness of something. As the discussion suggests, we could model awareness of unawareness with an existential quantifier like “I am uncertain about whether *there exists* a proposition that I am unaware of”. In this section, we will present several alternative approaches.

5.1 Propositional Quantifiers and Extended Awareness Structures

Halpern and Rêgo [2009] presented an extension of awareness structures with propositional quantifiers. They extend the syntax to allow for quantification over quantifier-free formulae. Unfortunately, formulae expressing that an agent considers it possible that she is aware of all formulae and also considers it possible that she is not aware of all formulae are not satisfiable in any of their extended awareness structures. This is a serious limitation for applications as it is very natural to consider agents who may be uncertain about whether they are aware of everything or not. As remedy, Halpern and Rêgo [2013] introduce extended awareness structures that allow different languages to be defined at different states, very much in the spirit of Modica and Rustichini [1999] and Heifetz, Meier, and Schipper [2008]. We will focus in this section on Halpern and Rêgo [2013].

Given a nonempty set of agents $\text{Ag} = \{1, \dots, n\}$ indexed by a , a countable infinite set of primitive propositions At as well as a countable infinite set of variables X , the languages are $L_n^{\forall, K, A}(\text{At}', X)$, $\emptyset \neq \text{At}' \subseteq \text{At}$. Different from $L_n^{K, A}(\text{At})$ introduced earlier, we allow for quantification with domain $L_n^{\bar{K}, A}(\text{At})$: If φ is a formula in $L_n^{\bar{K}, A}(\text{At}')$, then $\forall x \varphi$ is a formula in $L_n^{\forall, K, A}(\text{At}', X)$. That is, the domain of quantification are just quantifier-free formulae. As usual, we define $\exists x \varphi$ by $\neg \forall x \neg \varphi$.

An occurrence of a variable x is free in a formula φ if it is not bound by a quantifier. More formally, define inductively: If φ does not contain a quantifier, then every occurrence of x is free in φ . An occurrence of the variable x is free in $\neg \varphi$, (in $K_a \varphi$ and $A_a \varphi$, respectively) if and only if its corresponding occurrence is free in φ . An occurrence of the variable x is free in $\varphi \wedge \psi$ if and only if the corresponding occurrence of x in φ or ψ is free. An occurrence of x is free in $\forall y \varphi$ if and only if the corresponding occurrence of x is free in φ and x is different from y . A formula that contains no free variables is a *sentence*.

If ψ is a formula, we denote by $\varphi[x/\psi]$ the formula that results in replacing all free occurrences of the variable x in φ with ψ .

An *extended awareness structure* is a tuple $M = (S, (R_a)_{a \in \text{Ag}}, (\mathcal{A}_a)_{a \in \text{Ag}}, V, \mathcal{A}t)$, where $(S, (R_a)_{a \in \text{Ag}}, (\mathcal{A}_a)_{a \in \text{Ag}}, V)$ is an awareness structure as introduced in Section 3.1 and $\mathcal{A}t : S \rightarrow 2^{\text{At}} \setminus \{\emptyset\}$ is a correspondence that assigns to each state s in S a nonempty subset of primitive propositions in At . We require that at each state every agent is only aware of sentences that are in the language of this state. That is, $\mathcal{A}_a(s) \subseteq L_n^{\forall, K, A}(\mathcal{A}t(s), X)$. Moreover, the properties “awareness generated by primitive propositions” and “agents know what they are aware of” take the following form: for all $a \in \text{Ag}$ and $s, s' \in S$, if $(s, s') \in R_a$ then $\mathcal{A}_a(s) \subseteq L_n^{\forall, K, A}(\mathcal{A}t(s'), X)$.

The satisfaction relation is defined inductively on the structure of formulae in $L_n^{\forall, K, A}(\text{At}, X)$ as follows:

$M, s \models p$ if and only if $p \in \mathcal{At}(s)$ and $V(s, p) = \text{true}$,
 $M, s \models \varphi \wedge \psi$ if and only if $\varphi, \psi \in \mathbb{L}_n^{\forall, K, A}(\mathcal{At}(s), \mathbf{X})$ and both $M, s \models \varphi$ and $M, s \models \psi$,
 $M, s \models \neg\varphi$ if and only if $\varphi \in \mathbb{L}_n^{\forall, K, A}(\mathcal{At}(s), \mathbf{X})$ and $M, s \not\models \varphi$,
 $M, s \models A_a\varphi$ if and only if $\varphi \in \mathbb{L}_n^{\forall, K, A}(\mathcal{At}(s), \mathbf{X})$ and $\varphi \in \mathcal{A}_a(s)$,
 $M, s \models K_a\varphi$ if and only if $\varphi \in \mathbb{L}_n^{\forall, K, A}(\mathcal{At}(s), \mathbf{X})$ and both $M, s \models A_a\varphi$ and $M, t \models \varphi$ for all $t \in S$ such that $(s, t) \in R_a$,
 $M, s \models \forall x\varphi$ if and only if $\varphi \in \mathbb{L}_n^{\forall, K, A}(\mathcal{At}(s), \mathbf{X})$ and $M, s \models \varphi[x/\psi]$ for all $\psi \in \mathbb{L}_n^{K, A}(\mathbf{At})$.

Note that for a formula φ to be true at a world s , we require now also that $\varphi \in \mathbb{L}_n^{\forall, K, A}(\mathcal{At}(s), \mathbf{X})$. We say that a formula φ is defined at s if (with the obvious abuse of notation) $\mathbf{At}(\varphi) \subseteq \mathcal{At}(s)$. Otherwise, φ is undefined at s . Like in unawareness structures, a formula may not be defined in all states in an extended awareness structure. This requires us to define validity analogous to unawareness structures, that is, a formula φ is valid in an extended awareness structure M if $M, s \models \varphi$ for all s in which φ is defined. The notions of soundness and completeness are as defined previously.

Consider the following axiom system that we call $\mathbf{S5}_n^{\forall, K, A}$:

- Prop. All substitution instances of valid formulae of propositional logic.
- AGPP. $A_a\varphi \leftrightarrow \bigwedge_{p \in \mathbf{At}(\varphi)} A_ap$
- AI. $A_a\varphi \rightarrow K_aA_a\varphi$
- KA. $K_a\varphi \rightarrow A_i\varphi$
- K. $(K_a\varphi \wedge K_a(\varphi \rightarrow \psi)) \rightarrow K_a\psi$
- T. $K_a\varphi \rightarrow \varphi$
- 4. $K_a\varphi \rightarrow K_aK_a\varphi$
- 5_A. $\neg K_a\varphi \wedge A_a\varphi \rightarrow K_a\neg K_a\varphi$
- 1_∀. $\forall x\varphi \rightarrow \varphi[x/\psi]$ if ψ is a quantifier-free sentence
- 6_∀. $\forall x(\varphi \rightarrow \psi) \rightarrow (\forall x\varphi \rightarrow \forall x\psi)$
- N_∀. $\varphi \rightarrow \forall x\varphi$ if x is not free in φ
- FA. $\forall xU_ax \rightarrow K_a\forall xU_ax$
- Barcan_A^{*}. $(A_a(\forall x\varphi) \wedge \forall x(A_ax \rightarrow K_a\varphi)) \rightarrow K_a(\forall xA_ax \rightarrow \forall x\varphi)$
- MP. From φ and $\varphi \rightarrow \psi$ infer ψ .
- Gen_A. From φ infer $A_a\varphi \rightarrow K_a\varphi$.
- Gen_∀. If $q \in \mathbf{At}$, then from φ infer $\forall x\varphi[q/x]$.

All axioms and inference rules that do not involve quantification were discussed earlier. 1_∀ means that if a universally quantified formula is true then so is every instance of it. FA and Barcan_A^{*} are more difficult to interpret. FA says that if an agent is unaware of everything then she knows that she is unaware of everything. It is hard to judge the reasonableness of this axiom as the hypothesis of being unaware of everything is extreme. At a first glance, it may even appear paradoxical: If she knows that she is unaware of everything then by KA she is aware that she is unaware of everything. But if she is aware that she is unaware of everything, how can she be unaware of everything? Recall that quantification is just over quantifier-free sentences. Thus, the agent may be unaware of every quantifier-free sentence but still be aware that she is unaware of every quantifier-free sentence.

Barcan_A^{*} should be contrasted with the “standard” Barcan axiom $\forall xK_a\varphi \rightarrow K_a\forall x\varphi$: If the agent knows $\varphi[x/\psi]$ for every quantifier-free sentence ψ , then she knows $\forall x\varphi$. Barcan_A^{*} also

connects knowledge with quantification but it requires awareness. In the antecedent it requires that the agent is aware of the formula $\forall x\varphi$. Moreover, the agent is required to know $\varphi[x/\psi]$ only if she is aware of ψ . In the conclusions, $\forall x\varphi$ is true only if the agent is aware of all (quantifier-free) formulae.

Theorem 10 (Halpern and Rêgo [2013]) *For the language $L_n^{\forall,K,A}(\text{At}, X)$, the axiom system $S5_n^{\forall,K,A}$ is a sound and complete axiomatization with respect to extended awareness structures.*

The completeness part of the proof is by constructing a canonical model but dealing appropriately with complications arising from quantification as in Halpern and Rêgo [2009]. The soundness parts for modus ponens, Barcan *_A , and Gen $_{\forall}$ are nonstandard.

Extended awareness structures merge features of both awareness structures and unawareness structures introduced in Sections 3.1 and 3.4, respectively. Recall that in awareness structures the awareness correspondence associates potentially different subsets of formulae with different states but all formulae are defined at each state, while in unawareness structures potentially different subsets of formulae are defined at different states. This difference between the formalisms are immaterial as long as we are “just” interested in modeling reasoning about knowledge and propositionally determined awareness and do not care about the important conceptual issue of whether structures can be viewed from an agent’s subjective perspective. In extended awareness structures, the difference between formulae defined at a state and the formulae that an agent is aware of at that state is of conceptual significance for a second reason. Roughly these are the labels that the agent is aware that he is unaware of.

Halpern and Rêgo [2013] explore the connection between awareness and unawareness structures by showing that quantifier-free fragment of their logic is characterized by exactly the same axioms as the logic of Heifetz, Meier, and Schipper [2008]. Moreover, they show that under minimal assumptions they can dispense with Fagin and Halpern [1988]’s syntactic notion of awareness as this notion of awareness is essentially equivalent to the one used in Modica and Rustichini [1999] and Heifetz, Meier, and Schipper [2006, 2008].

5.2 First-Order Logic with Unawareness of Objects

In an unpublished paper, Board and Chung [2011a] proposed a first-order modal logic with unawareness in order to model awareness of unawareness. Different from extended awareness structures, the quantification is over objects rather than over quantifier-free formulae.

Given a nonempty set of agents $\text{Ag} = \{1, \dots, n\}$, a countable infinite set of variables X , and k -ary predicates P for every $k = 1, 2, \dots$, the set of atomic formulae At is generated by $P(x_1, \dots, x_k)$ where $x_1, \dots, x_k \in X$. We require that there is a unary predicate E . The intended interpretation of $E(x)$ is “ x is real”. The language we consider is $L_n^{\forall,L,K,A}(\text{At}, X)$. We allow for quantification: If φ is a formula in $L_n^{\forall,L,K,A}(\text{At}, X)$ and $x \in X$, then $\forall x\varphi \in L_n^{\forall,L,K,A}(\text{At}, X)$. We define a variable to be free in a formula as in Section 5.1. Moreover, if φ is a formula, we denote by $\varphi[x/y]$ the formula that results from replacing all free occurrences of x with y .

An *object-based unawareness structure* is a tuple $M = (S, D, \{D(s)\}_{s \in S}, (\Pi_a)_{a \in \text{Ag}}, (\mathcal{A}_a)_{a \in \text{Ag}}, \pi)$, where S is a nonempty set of states, D is a nonempty set of objects, $D(s)$ is a nonempty subset of D containing objects that are “real” in s , and $\Pi_a : S \rightarrow 2^S$ is a possibility correspondence

of agent $a \in \mathbf{Ag}$. We focus here on the case in which, for each agent $a \in \mathbf{Ag}$, the possibility correspondence forms a partition of the state-space. That is, we assume that it satisfies

Reflexivity: $s \in \Pi_a(s)$ for all $s \in S$,
Stationarity: $s' \in \Pi_a(s)$ implies $\Pi_a(s') = \Pi_a(s)$, for all $s, s' \in S$.

$\mathcal{A}_a : S \longrightarrow 2^D$ is the awareness correspondence of agent $a \in \mathbf{Ag}$. Different from awareness structures or the product model discussed earlier, the awareness correspondence in object-based unawareness structures assigns subsets of objects to states. We focus here on the case, in which for each agent $a \in \mathbf{Ag}$, the possibility correspondence and the awareness correspondence satisfy jointly

$s' \in \Pi_a(s)$ implies $\mathcal{A}_a(s') = \mathcal{A}_a(s)$.

Thus analogous to the corresponding property in awareness structures, agents know what they are aware of.

π is a state-dependent assignment of a k -ary relation $\pi(s)(P) \subseteq D^k$ to each k -ary predicate P . Intuitively, the assignment π ascribes in each state and to each property the subset of objects satisfying this property at that state. It is sometimes called a classical first-order interpretation function.

A valuation $V : \mathbf{X} \longrightarrow D$ assigns to each variable an object. Intuitively, $V(x)$ denotes the object referred to by variable x , provided that x is free in a given formula. Call V' is an x -alternative valuation of V if, for every variable y except possibly x , $V'(y) = V(y)$.

Since the truth value of a formula depends on the valuation, on the left-hand side of $\models$ we have a model, a state in the model, and a *valuation*. The satisfaction relation is defined inductively on the structure of formulae in $\mathbf{L}_n^{\forall, L, K, A}(\mathbf{At}, \mathbf{X})$ as follows:

$M, s, V \models E(x)$ if and only if $V(x) \in D(s)$,
 $M, s, V \models P(x_1, \dots, x_k)$ if and only if $(V(x_1), \dots, V(x_k)) \in \pi(s)(P)$,
 $M, s, V \models \neg\varphi$ if and only if $M, s, V \not\models \varphi$,
 $M, s, V \models \varphi \wedge \psi$ if and only if $M, s, V \models \varphi$ and $M, s, V \models \psi$,
 $M, s, V \models \forall x\varphi$ if and only if $M, s, V' \models \varphi$ and $V'(x) \in D(s)$ for every x -alternative valuation V' ,
 $M, s, V \models A_a\varphi$ if and only if $V(x) \in \mathcal{A}_a(s)$ for every x that is free in φ ,
 $M, s, V \models L_a\varphi$ if $M, s', V \models \varphi$ for all $s' \in \Pi_a(s)$,
 $M, s, V \models K_a\varphi$ if and only if $M, s, V \models A_a\varphi$ and $M, s, V \models L_a\varphi$.

A formula φ is valid in the object-based unawareness structure M under the valuation V if $M, s, V \models \varphi$ for all $s \in S$. The notions of soundness and completeness are now the standard notions.

The unawareness operator is defined, as usual, as the negation of awareness, that is, $U_a\varphi := \neg A_a\varphi$.

Consider the following axiom system that we call $\mathbf{S5}_n^{\forall, L, K, A}$:

Prop. All substitution instances of valid formulae of propositional logic.

- A. $A_a\varphi$ if there is no free variable in φ ;
- AC⁻. $A_a\varphi \wedge A_a\psi \rightarrow A_a(\varphi \wedge \psi)$
- A3. If every variable free in ψ is also free in φ , then $A_a\varphi \rightarrow A_a\psi$.
- K. $L_a(\varphi \rightarrow \psi) \rightarrow (L_a\varphi \rightarrow L_a\psi)$
- T. $L_a\varphi \rightarrow \varphi$
- 4. $L_a\varphi \rightarrow L_aL_a\varphi$
- 5. $\neg L_a\varphi \rightarrow L_a\neg L_a\varphi$
- KL. $K_a\varphi \leftrightarrow A_a\varphi \wedge L_a\varphi$.
- AI_L. $A_a\varphi \rightarrow L_aA_a\varphi$.
- UI_L. $U_a\varphi \rightarrow L_aU_a\varphi$.
- E. $\forall xE(x)$.
- 1 _{$\forall, E$} . $\forall x\varphi \rightarrow (E(y) \rightarrow \varphi[x/y])$.
- 6 _{$\forall$} . $\forall x(\varphi \rightarrow \psi) \rightarrow (\forall x\varphi \rightarrow \forall x\psi)$
- N _{$\forall$} ^{*}. $\varphi \leftrightarrow \forall x\varphi$ if x is not free in φ
- MP. From φ and $\varphi \rightarrow \psi$ infer ψ .
- Gen_L. From φ infer $L_a\varphi$.
- Gen _{$\forall, L$} . For all natural numbers $n \geq 1$, from $\varphi \rightarrow L_a(\varphi_1 \rightarrow \dots \rightarrow L_a(\varphi_n \rightarrow L_a\psi) \dots)$ infer $\varphi \rightarrow L_a(\varphi_1 \rightarrow \dots \rightarrow L_a(\varphi_n \rightarrow L_a\forall x\psi) \dots)$, provided that x is not free in $\varphi, \varphi_1, \dots, \varphi_n$.
- Gen _{$\forall$} ^{*}. From φ infer $\forall x\varphi$.

Theorem 11 (Board and Chung [2011a]) *For the language $\mathsf{L}_n^{\forall, L, K, A}(\mathsf{At}, \mathsf{X})$, the axiom system $\mathsf{S5}_n^{\forall, L, K, A}$ is a sound and complete axiomatization with respect to object-based unawareness structures.*

The proof uses standard methods. Board and Chung [2011a] present the proof of a version not imposing the assumption of a partitioned possibility correspondence.

Board and Chung [2011b] consider also “frames” of object-based unawareness structures for modeling reasoning about knowledge of unawareness. This approach is not purely event-based, though, as it requires the modeler to consider for each event also the set of objects referred to in the event.

Formally, an *event* in an object-based unawareness frame is a pair (E, O) , where $E \in 2^S$ is a subset of states and $O \in 2^D$ is a subset of objects. (E is now a set of states, not the existence predicate introduced earlier.) We let $\text{states}(E, O) := E$, and $\text{objects}(E, O) := O$. Negation and conjunction of events are defined by

$$\neg(E, O) := (S \setminus E, O),$$

$$\bigwedge_i (E_i, O_i) := \left(\bigcap_i E_i, \bigcup_i O_i \right).$$

The negation pertains to the set of states in which E does not obtain but refers to the same set of objects. The conjunction of events is the set of states in which all these events obtain and the union of objects referred to by those events. Conjunction is defined by the De-Morgan law by

$$\bigvee_i (E_i, O_i) = \neg \left(\bigwedge_i \neg(E_i, O_i) \right) = \left(\bigcup_i E_i, \bigcup_i O_i \right).$$

We let Σ denote the set of all events.

For each agent $a \in \mathbf{Ag}$, the awareness operator is defined on events by

$$A_a(E, O) := (\{s \in S : O \subseteq \mathcal{A}_a(s)\}, O).$$

As before, the unawareness operator is defined as the negation of the awareness, $U_a(E, O) := \neg A_a(E, O)$.

The implicit knowledge operator is defined on events by

$$L_a(E, O) := (\{s \in S : \Pi_a(s) \subseteq E\}, O).$$

Explicit knowledge is defined, as in awareness structures, by the conjunction of awareness and implicit knowledge, i.e.,

$$K_a(E, O) := A_a(E, O) \wedge L_a(E, O).$$

Awareness, implicit knowledge, and explicit knowledge of an event with a given subset of objects are events, respectively, with the same subset of objects.

Properties are defined as functions $p : D \rightarrow \Sigma$ such that $p(o) = (E_o^p, O^p \cup \{o\})$ for some $E_o^p \in 2^S$ and $O^p \in 2^D$. E_o^p is the set of states in which object o possesses property p and O^p is the set of objects referred to in that property. For instance, a property could be “... has as many legs as horses.” If object o is a unicorn, then E_o^p is the set of states in which this unicorn has as many legs as horses and O^p is the set of horses.

Object-based unawareness structures allow for quantification over objects. In the object-based unawareness frame, we will consider quantified events. Board and Chung [2011b] focus on an actualist quantifier that ranges over objects that “actually exist”. Formally, first define the property

$$e(o) = (\{s \in S : o \in D(s)\}, \{o\}).$$

That is, $e(o)$ is the event that object o exists. For any property p , the event that all (actually existing) objects satisfy property p is defined by

$$Allp = \left(\bigcap_{o \in D} E_o^{e \rightarrow p}, O^p \right).$$

$Allp$ obtains if all existing objects possess property p . Quantified events satisfy the following properties:


- (i) $All(\wedge_i p_i) = \wedge_i (Allp_i)$
- (ii) If $s \in E_o^p$ for every $o \in D$, then $s \in \text{states}(Allp)$.
- (iii) If $E_o^p = E_o^q$ for every $o \in D$, then $\text{states}(Allp) = \text{states}(Allq)$.

At this point, it may be helpful to consider our simple example.

Speculative Trade Example (continued). Let ℓ denote the object “lawsuit” and n the object “innovation”. Figure 5 presents a simply objective-based unawareness frame. There are

four states. Below each state we indicate which atomic formulas are true or false. The picture is analogous to the corresponding picture for awareness structures (Figure 1) except that the awareness correspondence is now indicated by rectangular text bubbles above states in which we indicate the set of objects the agent is aware of that state. The blue solid-lined rectangular text bubbles belong to the owner while the red intermitted-lined are the buyer's. Each agent's awareness correspondence is very special in this example because it is constant across states. As in Figure 1, the soft-edged rectangles indicate the possibility correspondences. The blue solid-lined possibility set belongs to the owner, while the red intermitted-lined is the buyer's. Both agents consider all states possible.

Figure 5: An Object-Based Unawareness Frame for the Speculative Trade Example



This simple figure models the story outlined in the introduction. The awareness correspondences shows that at any state, the owner is aware of the lawsuit and unaware of the innovation, while the buyer is aware of the innovation and unaware of the lawsuit. The possibility correspondences model implicit knowledge. The owner does not implicitly know whether the law suit obtains, and implicitly knows that the buyer is unaware of the lawsuit. But he also explicitly knows that because he is aware of the lawsuit. The buyer does not implicitly know whether the innovation obtains, and implicitly knows that the owner is unaware of the innovation. But she also explicitly knows that because she is aware of the innovation. Both agents also implicitly know what they are unaware of. This is hard to interpret. Similar to awareness structures discussed earlier, object-based unawareness structures are best understood from an outside modeler's point of view. The same structure can generally not be used as an analytical device by the agent herself to reason about her and other agents' knowledge and awareness.

We should mention that the introductory speculative trade example does not do full justice to objective-based unawareness structures as the example does not make use of quantification over objects. How would awareness of unawareness affect speculative trade? Board and Chung [2011b] show a “No-speculative-trade” theorem for the case that agents are paranoid in the sense that they always consider it possible that there is something they are unaware of. This echoes Grant and Quiggin [2013], who argue that agents' past experience lends support to the hypothesis that there are some contingencies that they are unaware of. This awareness of unawareness coupled with the precautionary principle may make them reluctant to engage in speculative trade. $\square$

The object-based unawareness frame allows us to easily compare this approach to unawareness frames introduced in Section 3.3. Already on an informal level, some differences become obvious. First, quantification is not explicitly considered in unawareness frames. Second, similar to awareness structures by Fagin and Halpern [1988], the set of states in object-based unawareness structures are “objective” descriptions that should be interpreted as given to an outside observer. This is different from “subjective” descriptions in unawareness structures. Third, Board and Chung [2011a] present no axiomatization for a language with explicit knowledge and awareness only. As in awareness structures of Fagin and Halpern [1988], their possibility correspondence models implicit knowledge. Arguably, explicit knowledge is what is of ultimate interest in applications and this is the notion focused on in unawareness structures. Finally, object-based unawareness frames model unawareness about propositions where the unawareness arises from unawareness of objects referred to in the propositions, and no unawareness of properties is considered. In contrast, unawareness structures model unawareness of abstract propositions. Yet, the more fine-grained distinction between objects and properties in object-based unawareness structures may yield an advantage in some applications where this distinction may be necessary.

More formally, common to both frames is that there is a set of events Σ , a negation operator $\neg$, a conjunction operator $\wedge$, and for each agent $a \in \mathbf{Ag}$ a knowledge operator K_a and an awareness operator A_a defined on events in Σ . We say that a frame $(\Sigma, \neg, \wedge, (K_a, A_a)_{a \in \mathbf{Ag}})$ can be embedded into a frame $(\Sigma', \neg, \wedge, (K'_a, A'_a)_{a \in \mathbf{Ag}})$ if there is an injective function $f : \Sigma \rightarrow \Sigma'$ with the following properties: For any events $E, F \in \Sigma$,

- Negation-Preserving: $f(\neg E) = \neg f(E)$
- Conjunction-Preserving: $f(E \wedge F) = f(E) \wedge f(F)$
- Knowledge-Preserving: $f(K_a(E)) = K'_a(f(E))$
- Awareness-Preserving: $f(A_a(E)) = A'_a(f(E))$

Theorem 12 (Board, Chung, and Schipper [2011]) *Every object-based unawareness frame can be embedded into some unawareness frame. Conversely, every unawareness frame can be embedded into an object-based unawareness frame that does not necessarily satisfy that agents know what they are aware of.*

The proof is constructive in that one can construct an embedding and show that it “works”. Moreover, one can show that the property of ‘agents know what they are aware of’ is required for any embedding of object-based unawareness frames into unawareness frames. Board, Chung, and Schipper [2011] also show that if generalized reflexivity (i.e., the truth axiom) and stationarity (i.e., the introspection properties of knowledge) are dropped from unawareness frames, then any object-based unawareness frame (not necessarily satisfying reflexivity, stationarity, or ‘agents know what they are aware of’) can be embedded into an unawareness frame and vice versa. While this result is mathematically more general and “cleaner” than Theorem 12 (because it has the full converse), it is of less interest because we want to know how even strong notions of knowledge are embedded into various frames. So far, it is open whether with strong notions of knowledge a full converse can be obtained with some different embedding than the one used in the proof.

5.3 Neighborhood Semantics and First-Order Logic with Awareness of Objects and Properties

Sillari [2008a,b] introduces a first-order logic of awareness but with a semantics based on awareness neighborhood frames mentioned already in Section 3.1. The language is as in Section 5.2. That is, $\mathcal{L}_n^{\forall, L, K, A}(\text{At}, \mathbf{X})$ consists of well-formed formulae having the following syntax

$$\varphi ::= P(x_1, \dots, x_k) \mid A!_a(x) \mid \neg\varphi \mid \varphi \wedge \psi \mid L_a\varphi \mid K_a\varphi \mid A_a\varphi \mid \forall x\varphi$$

$A!_a(x)$ denotes the awareness predicate of agent a , which is somewhat similar to the agent-independent existence predicate $E(x)$ in Section 5.2.

An awareness neighborhood structure is $M = (S, (\mathcal{N}_a)_{a \in \text{Ag}}, (\mathcal{A}_a)_{a \in \text{Ag}}, D, (D_a)_{a \in \text{Ag}}, \pi, (\pi_a)_{a \in \text{Ag}})$ in which S is a space of states and $\mathcal{N}_a : S \rightarrow 2^{2^S}$ is the neighborhood correspondence of agent a that assigns to each state the set of events that the agent knows at this state. Sillari [2008a,b] imposes no conditions on $\mathcal{N}_a$. As in awareness structures, $\mathcal{A}_a : S \rightarrow 2^{\mathcal{L}_n^{\forall, L, K, A}(\text{At}, \mathbf{X})}$ is the awareness correspondence of agent a . D is a nonempty set called the domain. $D_a : S \rightarrow 2^D$ is a correspondence of agent a assigning to each state s a subjective domain $D_a(s)$ of objects. Intuitively, $D_a(s)$ represents the objects that agent a is aware of at state s . As in the previous section, π is a state-dependent assignment of a k -ary relation $\pi(s)(P) \subseteq D^k$ to each k -ary predicate P . Finally, for each agent $a \in \text{Ag}$, π_a is a state-dependent assignment of a k -ary relation $\pi_a(s)(P) \subseteq D_a^k(s)$ to each k -ary predicate P that may possibly agree partially with π . This “agent-based” assignment is motivated by the author’s desire to model also awareness of properties.

Recall that in first-order logic an atomic formula takes the form $P(x_1, \dots, x_k)$ where P is a k -ary predicate and $x_1, \dots, x_k \in \mathbf{X}$. The notion of ‘awareness generated by atomic formulae’ is analogous to awareness structures, i.e., for all $s \in S$, $\varphi \in \mathcal{A}_a(s)$ if and only if $\text{At}(\varphi) \subseteq \mathcal{A}_a(s)$. We require that $P(x_1, \dots, x_k) \in \mathcal{A}_a(s)$ if and only if (i) $V(x_\ell) \in D_a(s)$ for $\ell = 1, \dots, k$, and (ii) $(V(x_1), \dots, V(x_k)) \in \pi_a(s)(P)$. As before, $V : \mathbf{X} \rightarrow D$ is a valuation or substitution. By Property (i), if at a state an agent is aware of an atomic formula then at that state she must be aware of any object referred to in the atomic formula. This formalizes the idea of awareness of objects. Property (ii) is interpreted as formalizing the idea of awareness of properties of objects. In order for an agent to be aware of a given atomic formula, she needs to be also aware of the property mentioned in the formula, i.e., she needs to be aware that the objects in the formula enjoy a given property. This is different from Section 5.2 where only awareness of objects is considered.

The satisfaction relation is defined by

$$\begin{aligned} M, s, V &\models A!_a(x) \text{ if and only if } V(x) \in D_a(s), \\ M, s, V &\models P(x_1, \dots, x_k) \text{ if and only if } (V(x_1), \dots, V(x_k)) \in \pi(s)(P) \\ M, s, V &\models \neg\varphi \text{ if and only if } M, s, V \not\models \varphi, \\ M, s, V &\models \varphi \wedge \psi \text{ if and only if } M, s, V \models \varphi \text{ and } M, s, V \models \psi, \\ M, s, V &\models \forall x\varphi(x) \text{ if and only if } M, s, V' \models \varphi \text{ for every } x\text{-alternative valuation } V' \text{ for} \\ &\text{which } V'(y) = V(y) \text{ for all } y \neq x, \\ M, s, V &\models L_a\varphi \text{ if and only if } \{t \in S : M, t, V \models \varphi\} \in \mathcal{N}_a(s), \\ M, s, V &\models A_a\varphi \text{ if and only if } \varphi \in \mathcal{A}_a(s), \\ M, s, V &\models K_a\varphi \text{ if and only if } M, s, V \models A_a\varphi \text{ and } M, s, V \models L_a\varphi. \end{aligned}$$

In addition to some axioms on awareness discussed already in Section 3.1, Sillari [2008a,b] discusses two axioms that relate awareness and quantifiers. Recall that $\exists x\varphi$ stands for $\neg\forall x\neg\varphi$. The first axiom is

$$A\exists. A_a\varphi[x/y] \rightarrow A_a\exists x\varphi(x).$$

If agent a is aware that y , which substitutes a free x , has property φ , then she is aware that there exists an x with property φ . Second,

$$A\forall. A_a\forall x\varphi(x) \rightarrow (A!_a(y) \rightarrow A_a\varphi[x/y]).$$

If agent a is aware that any x has property φ , then, provided that a is aware of y , she is aware that y , which substitutes a free x , has property φ . This axiom has a similar flavor as the axiom $1_{\forall,E}$ in the previous section except that it now involves awareness.

Sillari [2008a,b] does not present soundness and completeness proofs of first-order modal logic with awareness, but Sillari [2008a] suggests that results in Arló-Costa and Pacuit [2006] could be extended to awareness.

Sillari [2008b] proves two theorems. First, he shows that awareness neighborhood structures (without restrictions on the awareness correspondences and neighborhood correspondences) are equally expressive to impossible possible worlds structures introduced by Rantala [1982a,b] and Hintikka [1975]. This complements results on equal expressivity of awareness (Kripke) structures and impossible possible worlds structures by Wansing [1990]. Second, he shows an analogous result for quantified impossible possible worlds structures and quantified awareness neighborhood structures. This implies that one should be able to model awareness also with impossible possible worlds structures. Yet, without knowing how exactly various restrictions on awareness and belief translate into impossible possible worlds, it is not clear how tractable it would be to model awareness with these structures.

5.4 Awareness of Unawareness without Quantifiers

In this section we present the idea originally due to Ågotnes and N. [2007] of modeling awareness of unawareness by propositional constants such as “agent a is aware of everything” and “agent b is aware of everything that agent a is aware of” rather than with quantifiers. Unfortunately, their approach does not allow an agent to be uncertain about whether she is aware of everything or not. As a remedy, we will present the two-stage semantics by Walker [2014] in order to allow an agent also to be uncertain about her awareness of unawareness.

Let $L_n^{L,K,A,F,R}(\text{At})$ consists of well-formed formulae having the following syntactic forms

$$\varphi ::= p \mid F_a \mid R_{ab} \mid \varphi \mid \neg\varphi \mid \varphi \wedge \psi \mid L_a\varphi \mid K_a\varphi \mid A_a\varphi$$

The propositional constants, F_a and R_{ab} for $a, b \in \text{Ag}$, are new. Formula F_a stands for “agent a is aware of everything” (i.e., “full” awareness) while formula R_{ab} reads “agent b is aware of everything that agent a is aware of” (i.e., “relative” awareness). (Note that, different from previous sections, R_{ab} does not denote the accessibility relation but a propositional constant.)

A *modified awareness structure* $M = (S, (\Pi_a)_{a \in \text{Ag}}, (\mathcal{A}_a)_{a \in \text{Ag}}, (\triangleright_a^s)_{a \in \text{Ag}, s \in S}, V)$ consists of a space of states S and for each agent $a \in \text{Ag}$ a possibility correspondence $\Pi_a : S \rightarrow 2^S$. As in

the previous sections, we will focus on the case where Π_a forms a partition of S . That is, for all $a \in \mathbf{Ag}$ and $s \in S$, we require

Reflexivity: $s \in \Pi_a(s)$, and

Stationarity: $s' \in \Pi_a(s)$ implies $\Pi_a(s') = \Pi_a(s)$.

Modified awareness structures in which every Π_a forms a partition are called *partitional modified awareness structures*.

The awareness correspondence $\mathcal{A}_a : S \rightarrow 2^{\mathbf{At}}$ assigns to each state a subset of atomic formulae. Note that different from awareness structures introduced in Section 3.1, the codomain of the awareness correspondence is restricted to the set of all subsets of atomic formulae only (instead allowing for the entire language). We focus on the case in which *agents know what they are aware of*, that is, for all $a \in \mathbf{Ag}$ and $s \in S$,

$s' \in \Pi_a(s)$ implies $\mathcal{A}_a(s') = \mathcal{A}_a(s)$.

The next ingredient is new: For each state $s \in S$, $\succeq_a^s$ is a preorder (i.e., a reflexive and transitive binary relation) on $\mathbf{Ag} \cup \{\mathbf{At}\}$ with $\mathbf{At} \in \max_{\succeq_a^s} \{\mathbf{Ag} \cup \{\mathbf{At}\}\}$. The preorder $\succeq_a^s$ describes agent a 's conjecture about the relative extent of all agent's awareness at state s . $b \succeq_a^s c$ means that agent a conjectures in state s that agent b 's awareness is more extensive than agent c 's awareness. $b \succeq_a^s \mathbf{At}$ means that agent a conjectures agent b to be aware of everything. $\mathbf{At} \succeq_a^s b$ means that agent a conjectures agent b to be not more than aware of everything (which does not imply that agent a is aware of everything). We focus on the case in which for $s \in S$ and agents $a, b, c \in \mathbf{Ag}$, the awareness correspondences and the preorders jointly satisfy the condition that we may dub *coherent relative awareness*:

$\mathcal{A}_a(s) \cap \mathcal{A}_b(s) \not\supseteq \mathcal{A}_a(s) \cap \mathcal{A}_c(s)$ implies $b \not\succeq_a^s c$.

This condition may be interpreted as saying that agent a 's conjecture at state s about agent b 's awareness relative to agent c 's awareness is based on those agents' actual awareness conditional on agent a 's awareness at that state.

The last component of the modified awareness structure is the valuation function $V : S \times \mathbf{At} \rightarrow \{\text{true}, \text{false}\}$.

Walker [2014] introduces a two-stage semantics. At the first stage, an “individualized preliminary” truth value is assigned to every formula at every state. At the second stage, the final truth value is assigned. We denote the individualized preliminary satisfaction relation of agent a by $\models_a^1$ and the final satisfaction relation by $\models$. The individualized preliminary satisfaction relation is defined inductively on the structure of formulae in $\mathcal{L}_n^{L,K,A,F,R}(\mathbf{At})$ as follows:

$M, s \models_a^1 p$ if and only if $V(s, p) = \text{true}$,
 $M, s \models_a^1 F_b$ if and only if $b \succeq_a^s \mathbf{At}$,
 $M, s \models_a^1 R_{bc}$ if and only if $c \succeq_a^s b$,
 $M, s \models_a^1 \neg \varphi$ if and only if $M, s \not\models_a^1 \varphi$,
 $M, s \models_a^1 \varphi \wedge \psi$ if and only if both $M, s \models_a^1 \varphi$ and $M, s \models_a^1 \psi$,
 $M, s \models_a^1 L_b \varphi$ if and only if $M, s' \models_b^1 \varphi$ for all $s' \in \Pi_b(s)$,

$$\begin{aligned}
M, s \models_a^1 A_b \varphi & \text{ if and only if } \text{At}(\varphi) \subseteq \mathcal{A}_b(s), \\
M, s \models_a^1 K_b \varphi & \text{ if and only if both } M, s \models_a^1 A_b \varphi \text{ and } M, s \models_a^1 L_b \varphi.
\end{aligned}$$

All clauses with the exception of the second and third clause are familiar from the satisfaction relation defined for awareness structures. In the modified awareness structure M , formula F_b is preliminarily true for agent a at state s if and only if at state s agent a conjectures that agent b is aware of everything. Similarly, in the modified awareness structure M , formula R_{bc} is preliminarily true for agent a at state s if and only if at state s agent a conjectures that agent c 's awareness is more extensive than agent b 's awareness. Note that although the preliminary satisfaction relation is individualized, it is hard to interpret it as a subjective notion because states in an awareness structure should be interpreted as “objective” descriptions from an modeler's point of view and not necessarily from the agent's point of view.

The final satisfaction relation is defined inductively on the structure of formulae in $\mathbf{L}_n^{L,K,A,F,R}(\text{At})$ and makes use of the individualized preliminary satisfaction relation as follows:

$$\begin{aligned}
M, s \models p & \text{ if and only if } V(s, p) = \text{true}, \\
M, s \models F_a & \text{ if and only if } \mathcal{A}_a(s) = \text{At}, \\
M, s \models R_{ab} & \text{ if and only if } \mathcal{A}_b(s) \supseteq \mathcal{A}_a(s), \\
M, s \models \neg \varphi & \text{ if and only if } M, s \not\models \varphi, \\
M, s \models \varphi \wedge \psi & \text{ if and only if both } M, s \models \varphi \text{ and } M, s \models \psi, \\
M, s \models L_a \varphi & \text{ if and only if } M, s' \models_a^1 \varphi \text{ for all } s' \in \Pi_a(s), \\
M, s \models A_a \varphi & \text{ if and only if } \text{At}(\varphi) \subseteq \mathcal{A}_a(s), \\
M, s \models K_b \varphi & \text{ if and only if both } M, s \models A_b \varphi \text{ and } M, s \models L_b \varphi.
\end{aligned}$$

The second and third clauses use now the awareness correspondences instead individual conjectures captured by the preorders. In the modified awareness structure M , formula F_a is true at state s if and only if at state s agent a is aware of everything. Similarly, in the modified awareness structure M , formula R_{ab} is true at state s if and only if at state s agent b 's awareness as given by his awareness set is more extensive than agent a 's awareness. Most important is the clause giving semantics to implicit knowledge, which refers to the preliminary satisfaction relation of agent a . In the modified awareness structure M , agent a implicitly knows formula φ at state s if φ is *preliminarily true* for agent a at every state that he considers possible at s . Thus, whether or not an agent implicitly knows a formula depends on his preliminary satisfaction relation at states that he considers possible. This can be different from the final satisfaction relation for formulas involving F_a and R_{ab} .

The notion of validity is as in Kripke or awareness structures using the final satisfaction relation $\models$ just defined.

The aim is to characterize modified awareness structures in terms of properties of knowledge and awareness. To state the axiom system, it will be helpful to define the sublanguage $\mathbf{L}^- \subseteq \mathbf{L}_n^{L,K,A,F,R}(\text{At})$ that consists exactly of the set of formulae whose final truth values in any state and any modified awareness structure coincides with the individualized preliminary truth values. Define $\mathbf{L}^-$ inductively as follows:

$$\begin{aligned}
p \in \text{At} & \text{ implies } p \in \mathbf{L}^-, \\
A_a \varphi, L_a \varphi, K_a \varphi & \in \mathbf{L}^- \text{ for any } \varphi \in \mathbf{L}_n^{L,K,A,F,R}(\text{At}),
\end{aligned}$$

$\varphi \in \mathbf{L}^-$ implies $\neg\varphi \in \mathbf{L}^-$,

$\varphi, \phi \in \mathbf{L}^-$ implies $\varphi \wedge \phi \in \mathbf{L}^-$.

Note that $\varphi \in \mathbf{L}^-$ does not imply that F_a or R_{ab} for some $a, b \in \mathbf{Ag}$ could not be a subformula of φ . With this definition on hand, we consider the following axiom system that we call $\mathbf{S5}_n^{L,K,A,F,R}$:

Prop. All substitution instances of tautologies of propositional logic.

KL. $K_a\varphi \leftrightarrow L_a\varphi \wedge A_a\varphi$ (Explicit Knowledge is Implicit Knowledge and Awareness)

AS. $A_a\neg\varphi \leftrightarrow A_a\varphi$ (Symmetry)

AC. $A_a(\varphi \wedge \psi) \leftrightarrow A_a\varphi \wedge A_a\psi$ (Awareness Conjunction)

AKR. $A_a\varphi \leftrightarrow A_aK_a\varphi$ (Awareness Explicit Knowledge Reflection)

ALR. $A_a\varphi \leftrightarrow A_aL_a\varphi$ (Awareness Implicit Knowledge Reflection)

AR. $A_a\varphi \leftrightarrow A_aA_a\varphi$ (Awareness Reflection)

AI. $A_a\varphi \rightarrow K_aA_a\varphi$ (Awareness Introspection)

F0. A_aF_b

R0. A_aR_{bc}

F1. $F_a \rightarrow A_a\varphi$

F2. $F_a \rightarrow R_{ba}$

F3. $F_a \wedge R_{ab} \rightarrow F_b$

R1. $R_{ab} \rightarrow (A_a\varphi \rightarrow A_b\varphi)$

R2. $R_{ab} \wedge R_{bc} \rightarrow R_{ac}$ (Transitivity of Relative Awareness)

R3. R_{aa} (Reflexivity of Relative Awareness)

R4. $A_a\varphi \rightarrow K_a((A_b\varphi \wedge \neg A_c\varphi) \rightarrow \neg R_{bc})$

K. $(L_a\varphi \wedge L_a(\varphi \rightarrow \psi)) \rightarrow L_a\psi$ (Distribution Axiom)

T⁻. $L_a\varphi \rightarrow \varphi$ for any $\varphi \in \mathbf{L}^-$ (Modified Implicit Knowledge Truth Axiom)

4. $L_a\varphi \rightarrow L_aL_a\varphi$ (Implicit Positive Introspection Axiom)

5. $\neg L_a\varphi \rightarrow L_a\neg L_a\varphi$ (Implicit Negative Introspection Axiom)

MP. From φ and $\varphi \rightarrow \psi$ infer ψ (modus ponens)

Gen⁻. From φ proved without application of F1 or R1 infer $L_a\varphi$ (Modified Implicit Knowledge Generalization)

Most axiom schemes are familiar from previous sections. F0 means that the agent is always aware that an agent is aware of everything. R0 means that the agent is always aware that an agent's awareness is as extensive as another (or the same) agent's awareness. F1 states that full awareness implies awareness of any particular formula. Of course, these properties do not necessarily imply that an agent is aware of everything or is as aware as other agents. F2 means that full awareness implies relative awareness with respect to any agent.

F3 says that relative awareness of one agent with respect to a second agent implies full awareness of the first agent in the case that the second agent is fully aware. R1 states that if agent b 's awareness is as extensive as agent a 's awareness, then agent a being aware of a formula implies that agent b must be aware of it as well. R2 encapsulates the idea that relative awareness is transitive among agents. If agent b is as aware as agent a and agent c is as aware as agent b , then also agent c must be as aware as agent a . R3 states that relative awareness is reflexive in the sense that every agent is aware of everything that he is aware of. Finally, R4 says that when agent a is aware of a formula then he knows that if agent b is aware of it and agent c is not, then c 's awareness is not as extensive as b 's awareness. It is closely connected

to the condition “coherent relative awareness”. Axioms F1 to F3 and R1 to R3 appear also in Ågotnes and N. [2007].

The set of formulae for which the individualized preliminary truth values and the final truth values agree, L^- , play a role in the statement of the truth axiom T^- , which is restricted to just these formulae. This means that the agent may be delusional with respect to reasoning about full awareness of an agent or the awareness of an agent relative to another. Similarly, implicit knowledge generalization, Gen^- , has been weakened as it applies only to theorems that are deduced without use of axioms F1 or R1. We note that most of the axioms are stated in terms of implicit knowledge. As mentioned previously, axioms and inference rules that involve implicit knowledge are hard to interpret as it is not necessarily the knowledge that is “present in the agent’s mind”.

Theorem 13 (Walker [2014]) *For the language $L_n^{L,K,A,F,R}(At)$, the axiom system $S5_n^{L,K,A,F,R}$ is a sound and complete axiomatization with respect to partitioned modified awareness structures in which agents know what they are aware of and relative awareness conjectures are coherent.*

Modeling awareness of unawareness with propositional constants rather than quantification yields a language that is less expressive than the approaches introduced in Sections 5.1 to 5.3. For instance, we cannot express that an agent a knows that there is no more than one proposition that agent b is aware of but agent c is not. Nevertheless, the approach allows for modeling awareness of unawareness in relevant examples such as the doctor example mentioned earlier.

6 Synopsis of Extensions

6.1 Dynamic Awareness

All structures introduced in previous sections deal with agents in a static situation. Yet, many applications are likely to involve also changes of awareness and information. It is therefore desirable to devise structures capable of modeling changes of awareness and information. van Ditmarsch and French [2009, 2011a,b], van Ditmarsch, French, and Velázquez-Quesada [2012], as well as van Ditmarsch, French, Velázquez-Quesada, and Wáng [2013] study various logical semantics for changes of awareness. Central to their work are notions of bisimulation between awareness structures that may also be of independent interest. Intuitively, a bisimulation is a relation between Kripke models such that atomic formulas in related states have identical truth values and the information embedded in the accessibility relations is preserved. For awareness structures, we should also require from the notion of bisimulation that related states have identical awareness sets. Bisimulations are used to characterize modal equivalence, i.e., when states in distinct structures are indistinguishable by the formulae that are true in those states. Recall that accessibility relations in awareness structures model implicit knowledge. Thus, it is not surprising that the notion of bisimulation just mentioned yields modal equivalence with respect to a language with implicit knowledge. To characterize modal equivalence with respect to a language with explicit knowledge, the authors introduce a more appropriate notion of awareness bisimulation in which for each agent the information embedded in the accessibility relation is preserved when restricted to the fragment of the language of which the agent is aware. Using this notion of awareness bisimulation, the authors define a notion of speculative

knowledge. Speculative knowledge is similar to explicit knowledge in awareness structures; one notable difference is that an agent always knows tautologies even if those tautologies involve primitive propositions the agent may be unaware. van Ditmarsch, French, Velázquez-Quesada, and Wáng [2013] prove axiomatizations of awareness structures with respect to languages involving either implicit knowledge, explicit knowledge, or speculative knowledge. van Ditmarsch, French, and Velázquez-Quesada [2012], as well as van Ditmarsch, French, Velázquez-Quesada, and Wáng [2013] introduce epistemic awareness action models in which awareness and information of agents may change via certain actions (although agents cannot be differently aware of those actions) and prove an axiomatization. van Benthem and Velázquez-Quesada [2010] analyze changes of awareness by adding or dropping formulas from an agent’s awareness set, and prove completeness of their logic. They also discuss epistemic action models. Grossi and Velázquez-Quesada [2009] discuss changing awareness and additional inferences that may be induced by awareness changes in a nice example, “Twelve Angry Men”, using yet another framework. Hill [2010] introduces an algebraic semantic for modeling awareness of a single agent case and studies dynamic awareness logic.

Changes of awareness have also been studied within in stochastic processes (Modica [2008]), decision theory (Karni and Vierø [2013a,b], Li [2008b]), and in dynamic games discussed below.

6.2 Games with Unawareness

Situations in which agents with asymmetric awareness interact among each other are of particular interest. To what extent is an agent able to use her superior awareness to her advantage? Would agents want to make each other aware of some selected features of the situation but not on others? To study such issues, we need to complement epistemic structures capable of modeling awareness in multi-agent settings with actions that agents could take and incentives for agents. In other words, the analysis of such situations requires us to extend epistemic structures to games.

Strategic interaction among multiple players under incomplete information is usually modeled with Bayesian games. A standard Bayesian game (see for instance, Mertens and Zamir [1985], Section 5) consists of a type-space for n -players augmented with a set of actions and a utility function for each player.²⁰ The utility of a player depends both on the state of the world and the actions chosen by all players. A standard Bayesian game is implicitly common knowledge and common awareness among the players although they may have different beliefs about types of players. Feinberg [2012], Meier and Schipper [2014], and Sadzik [2007] extend Bayesian games to unawareness.

Meier and Schipper [2014] simply replace the type spaces in Bayesian games by unawareness type spaces as introduced in Section 4. They also model explicitly unawareness of actions and unawareness of players, define Bayesian equilibrium for games with unawareness, and prove existence of equilibrium. Meier and Schipper [2014] define a notion on unawareness perfection to capture the robustness of equilibria to uncertainty about opponents’ awareness of their actions. They show that an equilibrium is unawareness perfect if and only if it is an undominated

²⁰That is, a finite Bayesian game $\langle \mathbf{Ag}, \Omega, (t_a)_{a \in \mathbf{Ag}}, (M_a)_{a \in \mathbf{Ag}}, (u_a)_{a \in \mathbf{Ag}} \rangle$ consists of a nonempty finite set of players $\mathbf{Ag}$ and a nonempty finite space of states Ω . For each player $a \in \mathbf{Ag}$, there is a type mapping $t_a : \Omega \rightarrow \Delta(\Omega)$ satisfying introspection, $t_a(\{\omega' \in \Omega : t_a(\omega') = t_a(\omega)\})(\omega) = 1$ for all $\omega \in \Omega$. Further, each player $a \in \mathbf{Ag}$ has a nonempty finite set of actions M_a and a utility function $u_a : \Omega \times \prod_{b \in \mathbf{Ag}} M_b \rightarrow \mathbb{R}$.

equilibrium.

Sadzik [2007] develops a similar approach to Bayesian games with unawareness. Yet, different from Meier and Schipper [2014], he assumes a common prior probability distribution. Moreover, he imposes restrictions on strategies. In standard Bayesian games, a strategy is a mapping from types into probability distributions over actions.²¹ Consider now a type τ of a player (i.e., a set of states in which the player has the same beliefs and awareness) who is unaware of some parameter relevant to the strategic situation. Let there be also two other types, τ' and τ'' , of the this player who are aware of everything that τ is aware and something else, and hold the same beliefs as τ for everything that τ is aware of. Sadzik [2007] requires that the mixed action chosen by type τ is some average of the actions chosen by all such types τ' and τ'' . Meier and Schipper [2014] do not impose such a restriction because types τ' and τ'' are aware of some relevant parameter that τ is unaware of. Consequently, there is no way in which τ could take (even indirectly) some restrictions based on the corresponding more aware types into account.

Feinberg [2012] pursues a different approach. Rather than describing parsimoniously the players' beliefs and awareness by their types, analogously standard Bayesian games, he explicitly models each (mutual) view of the Bayesian game as a finite sequence of player names $i_1, \dots, i_n$ with the interpretation that this is how i_1 views how ... how i_n views the game. This is reminiscent of explicit syntactic descriptions or constructions of hierarchies of beliefs. He defines Bayesian Nash equilibrium in his setting and proves existence of equilibrium.

Bayesian games with unawareness are most faithfully interpreted as one-shot situations in which players choose actions simultaneously. The beliefs and awareness of players refer to their state of mind at the moment of choosing their actions. Yet, most strategic situations involve some time dimension. Players may not move simultaneously. Some actions may be taken before others. Awareness and beliefs may change during the course actions. Such a dynamic interaction is usually modeled with dynamic games. A dynamic game (or extensive-form game) consists of a tree in which subsets of players (including nature) are associated with nodes and edges represent profiles of actions that players at the emanating node can take. Further, information is modeled with information sets of nodes (modeling which histories of play the players at those nodes cannot distinguish from each other). Finally, for each player there is a utility function that assigns a payoff to each terminal nodes which represent the outcomes of game. Standard extensive-form games are implicitly common knowledge and common awareness among players. I.e., all players are aware of all players, all actions, etc. To allow for different levels of awareness, the game tree has to be replaced with a partially ordered forest of game trees, which is very much in analogy to the lattice of state-spaces in unawareness structures. Intuitively, each tree of the forest corresponds to a more or less rich description of the strategic situation. This is essentially the approach taken by Halpern and Rêgo [2014], Heifetz, Meier, and Schipper [2013b], and Feinberg [2012], who present general frameworks for modeling dynamic strategic interaction. The frameworks differ in the details of how the forest of game trees is constructed, the modeling of the players subjective views of the game, and their changes thereof. Recall that awareness structures and unawareness structures take slightly different approach to modeling information. These differences also surface again in the proposals by Halpern and Rêgo [2014] and Heifetz,

²¹I.e., a strategy of player a is $\sigma_a : \Omega \rightarrow \Delta(M_a)$ such that for all $\omega \in \Omega$, $\sigma_a(\omega') = \sigma_a(\omega)$ for all $\omega' \in \{\omega'' \in \Omega : t_a(\omega'') = t_a(\omega)\}$.

Meier, and Schipper [2013b]. Whereas information sets in Halpern and Rêgo [2014] are best understood as modeling implicit knowledge, information sets in Heifetz, Meier, and Schipper [2013b] model explicit knowledge and thus also awareness. The information set associated with a node in a given tree may comprise of nodes in a less expressive tree. In contrast, Halpern and Rêgo [2014] specify information sets at nodes of a game tree even if a player is not aware of that game tree. Then they devise a mapping that associates with each game tree and node a subtree and an information set in this subtree. We view this mapping somewhat analogous to “awareness” correspondences in awareness structures. Feinberg [2012] extends his approach to Bayesian games discussed above to extensive-form games with unawareness. That is, he models explicitly each (mutual) view of the dynamic game as a finite sequence of player names $i_1, \dots, i_n$ with the interpretation that this is how i_1 views how $\dots$ how i_n views the dynamic game. Earlier, Feinberg [2004] discussed a nice example of a repeated prisoners’ dilemma game in which a small grain of uncertainty about the opponent’s unawareness of the defect-action induces cooperation even for finite repetitions. This echoes the literature on reputations in game theory that obtained an analogous result by adding an “irrational” type whose irrationality is suitably tailored to the solution. In contrast, Feinberg’s example demonstrates that cooperation in such games can be obtained with a rather natural assumption on players’ beliefs about opponents’ unawareness of actions.

In game theory, there is a clear “division of labor” between the game and the solution concept. While the game represents the players’ (change of) awareness and information of the strategic situation, the solution concept captures the behavioral assumptions about the players. Various solution concepts to standard dynamic games exist in the literature. Although their mathematical definitions can be somehow extended to dynamic games with unawareness, their application to strategic situations under unawareness may no longer be conceptually appropriate. Most commonly used solution concepts are refinements of Nash equilibrium. A profile of strategies, one for each player, is a Nash equilibrium if each player’s strategy is a best response to the opponents’ strategies. It presumes that strategies are mutual knowledge among players. This is often informally motivated with interactive learning of the equilibrium convention: If players interact in the game repeatedly, then eventually they will learn somehow about the strategies used by opponents. Such a motivation cannot apply to games with unawareness in general. Games with unawareness model situations where some players may be unaware of some actions; thus they couldn’t have learned previously about such actions. If such an action is played during the play of the game, then it is far from clear where the players’ knowledge of the new equilibrium convention should come from. Therefore, equilibrium notions in strategic situations with unawareness may make sense only in special situations such as when players’ awareness along the equilibrium path never changes, or when becoming aware also implies that by some kind of process the new equilibrium convention becomes mutual knowledge.

To avoid the conceptual problems of equilibrium under unawareness, Heifetz, Meier, and Schipper [2013b] extend extensive-form rationalizability to dynamic games with unawareness. Extensive-form rationalizability is an algorithmic solution concept that iteratively eliminates possible beliefs of players about opponents’ strategies. It does not presume equilibrium. Nevertheless it is a strong solution for standard dynamic games because it entails forward-induction (Pearce [1984], Battigalli [1997]). In contrast to backward induction, which assumes that players’ future behavior will be rational, forward induction also attributes rationality to players’ past behavior if possible. Rather than simply excusing unexpected behavior of opponents as mistakes, a player who uses forward-induction tries to rationalize opponents’ past behavior to

form predictions about their future behavior. This is important under unawareness because if a rational player “becomes aware”, she is by definition surprised. If becoming aware is a result of an opponent’s action, then she should consider the opponent’s intention for making her aware (rather than discounting it as a mistake) and should use this information to play optimally.

Heifetz, Meier, and Schipper [2011] introduce prudent rationalizability, an outcome refinement of extensive-form rationalizability and an extensive-form analogue to iterated admissibility, for dynamic games with unawareness. Meier and Schipper [2012] define the associated normal-form game to dynamic games with unawareness and characterize both extensive-form rationalizability and prudent rationalizability in dynamic games with unawareness by iterated conditional strict (weak, resp.) dominance in the associated normal-form. Halpern and Rêgo [2014] extend Nash equilibrium to dynamic games with unawareness. Rêgo and Halpern [2012] extend sequential equilibrium, a refinement of Nash equilibrium, to dynamic games with unawareness because it is known that Nash equilibrium is a quite weak solution concept even in standard extensive-form games: it does not eliminate, for instance, incredible threats. Feinberg [2012] extends assessments, the main “ingredient” of sequential equilibrium, to his framework of dynamic games with unawareness.

While the frameworks briefly discussed above are completely general, some authors consider certain special classes of games with unawareness. Li [2006] studies a class of dynamic games that are restricted to perfect information but allow for unawareness. Ozbay [2008] studies dynamic interaction among one fully aware first-mover and a potentially unaware second mover. He introduces a refinement of an analogue to Perfect Bayesian equilibrium that entails forward induction. Grant and Quiggin [2013] present a framework for dynamic games with unawareness and apply sequential equilibrium as solution concept. Their framework is somewhat special because it excludes situations in which the set of terminal nodes of which one player is aware may be disjoint from the set of terminal nodes of which another player is aware.

Nielsen and Sebald [2012] merge dynamic unawareness games with another conceptual innovation of recent game theory: dynamic psychological games (Battigalli and Dufwenberg [2009]). In standard dynamic games, preferences of players are defined over terminal nodes of the game. But players may also have a variety of psychological attitudes such as emotions and intentions like guilt and reciprocity. Thus, they may not just care about the material outcome of the game but also about other player’s beliefs about them. Preferences over opponent’s beliefs are excluded in standard dynamic games but explicitly allowed in dynamic psychological games. Nielsen and Sebald [2012] are motivated by the observation that psychological attitudes of a player like feeling guilty when taking a certain action depend very much on whether the opponent is aware that the player could have acted otherwise. They consider dynamic psychological games with unawareness and sequential equilibrium but restrict themselves to two players only.

7 Summary

An overview over approaches discussed in this chapter is given in Figure 6. For lack of space, we excluded in this picture probabilistic approaches to unawareness (Section 4) as well as extensions to dynamic awareness and games with unawareness (Section 6). While the upper part of the figure lists single-agent structures, the middle part shows multi-agent structures. Finally, the lower part presents structures with awareness of unawareness. Roughly, we indicate

generalizations by an arrow and equivalence by a bi-directional arrow. Beside the arrows, we sometimes list articles that show the connection between the approaches. Often these results imply further relationships. The interested reader should consult the original papers for the precise notions of equivalence.

Figure 6 also shows a connection to the impossible worlds approach by Rantala [1982a,b]. Wansing [1990] shows that it is equally expressive to awareness structures of Fagin and Halpern [1988]. Thijsse [1991], Thijsse and Wansing [1996], and Sillari [2008a,b] contain further results along those lines.


Our review leaves out many topics. For instance, some of the discussed papers also contain results on the complexity of deciding the satisfiability of formulas (e.g., Fagin and Halpern [1988], Ågotnes and N. [2007], Halpern and Rêgo [2009] and van Ditmarsch and French [2011a]).

Complexity may also be related to awareness on a conceptual level. Already Fagin and Halpern [1988] suggested that one may want to consider a computational-based notion of awareness of agents who may lack the computational ability to deduce all logical consequences of their knowledge. Fagin, Halpern, Moses, and Vardi [1995] (Chapter 10.2.2) discuss the connection between algorithmic knowledge and awareness. One may also conceive of a computational-based notion of awareness of an object that roughly corresponds to the amount of time needed to generate that object within a certain environment. Such an approach is pursued by Devanur and Fortnow [2009] using Kolmogorov complexity.

There is a growing literature to unforeseen contingencies in decision theory. Here is not the space to give an adequate review and the interested reader may want to consult Dekel, Lipman, and Rustichini [1998b] for an early review of some of the approaches. Most work on unforeseen contingencies is best understood in terms of awareness of unawareness discussed in Section 5. More recent work appears, among others, in Nehring [1999], Dekel, Lipman, and Rustichini [2001], Epstein, Marinacci, and Seo [2007], Krishna and Sandowski [2013], Ahn and Ergin [2010] and the literature cited therein. Decision theoretic approaches to unawareness are pursued in Schipper [2013, 2014], Li [2008b] and Karni and Vierø [2013a,b]. Schipper [2013] replaces the state-space in the Anscombe-Aumann approach to subjective expected utility by a lattice of spaces (like in unawareness frames) and axiomatizes subjective expected utility that depends on the decision maker’s awareness level. He then uses the approach to show that unawareness has behavioral implications distinct from zero-probability. Li (2008) also studies zero-probability versus unawareness. Karni and Vierø [2013a,b] study updating of beliefs and awareness. Expanding awareness is analogous to “reverse Bayesian updating”.

Although we expect many more applications of unawareness to emerge, this concept has been applied to various contexts already. In Section 4.3 we outlined the application of unawareness to speculative trade. Another application pertains to the disclosure of verifiably information (and awareness) (see Heifetz, Meier, and Schipper [2011], Schipper and Woo [2014] and Li, Peitz, and Zhao [2014]). For instance, Schipper and Woo [2014] study modern electoral campaigning in which candidates microtarget voters with limited political awareness by raising certain political issues and providing some information on their political preferences over those issues. Galanis [2013b] analyzes the value of information under unawareness. Unawareness has been naturally applied to incomplete contracting (Lee [2008], Filiz-Ozbay [2012], Auster [2013], Grant, Kline, and Quiggin [2012] and von Thadden and Zhao [2012a,b]). Board and Chung [2011b] discuss unawareness in the presence of some legal doctrines. Liu [2008] discusses an application to fair disclosure in financial markets.

Figure 6: Partial overview over the literature



References

- T. Ågotnes and Alechina. N. Full and relative awareness: A decidable logic for reasoning about knowledge of unawareness. In D. Samet, editor, *Theoretical Aspects of Rationality and Knowledge, Proceedings of the 11th Conference (TARK 2007)*, pages 6–14. Presses Universitaires De Louvain, 2007.
- D. Ahn and H. Ergin. Framing contingencies. *Econometrica*, 78:655–695, 2010.
- H. Arló-Costa and E. Pacuit. First-order classical modal logic. *Studia Logica*, 84:171–210, 2006.
- R. Aumann. Agreeing to disagree. *Annals of Statistics*, 4:1236–1239, 1976.
- R. Aumann. Interactive epistemology i: Knowledge. *International Journal of Game Theory*, 28:263–300, 1999.
- S. Auster. Asymmetric awareness and moral hazard. *Games and Economic Behavior*, 82:503–521, 2013.
- P. Battigalli. On rationalizability in extensive games. *Journal of Economic Theory*, 74:40–61, 1997.
- P. Battigalli and M. Dufwenberg. Dynamic psychological games. *Journal of Economic Theory*, 144:1–35, 2009.
- J. van Benthem and F.R. Velázquez-Quesada. The dynamic of awareness. *Synthese*, 177:5–27, 2010.
- O. Board and K.S. Chung. Object-based unawareness: Axioms. University of Minnesota, 2011a.
- O. Board and K.S. Chung. Object-based unawareness: Theory and applications. University of Minnesota, 2011b.
- O. Board, K.S. Chung, and B.C. Schipper. Two models of unawareness: Comparing the object-based and subjective-state-space approaches. *Synthese*, 179:13–34, 2011.
- B.F. Chellas. *Modal logic: An introduction*. Cambridge University Press, 1980.
- Y.C. Chen, J.C. Ely, and X. Luo. Note on unawareness: Negative introspection versus au introspection (and ku introspection). *International Journal of Game Theory*, 41:325–329, 2012.
- M. Cozic. Probabilistic unawareness. Université Paris-Est Créteil Va-de-Marne, 2012.
- E. Dekel, B. Lipman, and A. Rustichini. Standard state-space models preclude unawareness. *Econometrica*, 66:159–173, 1998a.
- E. Dekel, B. Lipman, and A. Rustichini. Recent developments in modeling unforeseen contingencies. *European Economic Review*, 42:523–542, 1998b.
- E. Dekel, B. Lipman, and A. Rustichini. Representing preferences with a unique subjective state space. *Econometrica*, 69:891–934, 2001.
- N. Devanur and L. Fortnow. A computational theory of awareness and decision making. In A. Heifetz, editor, *Theoretical Aspects of Rationality and Knowledge, Proceedings of the 12th Conference (TARK 2009)*, pages 99–107. ACM, 2009.
- H. van Ditmarsch and T. French. Awareness and forgetting of facts and agents. In P. Boldi, G. Vizzari, G. Pasi, and R. Baeza-Yates, editors, *Proceedings of WI-IAT Workshops, Workshop WLIAMAS*, pages 478–483. IEEE Press, 2009.
- H. van Ditmarsch and T. French. Becoming aware of propositional variables. In M. Banerjee and A. Seth, editors, *ICLA 2011*, LNAI 6521, pages 204–218. Springer, 2011a.

- H. van Ditmarsch and T. French. On the dynamics of awareness and certainty. In D. Wang and M. Reynolds, editors, *Proceedings of 24th Australasian Joint Conference on Artificial Intelligence (AI 11)*, LNAI 7106, pages 727–738. Springer, 2011b.
- H. van Ditmarsch, T. French, and F.R. Velázquez-Quesada. Action models for knowledge and awareness. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, 2012.
- H. van Ditmarsch, T. French, F.R. Velázquez-Quesada, and Y.N. Wáng. Knowledge, awareness, and bisimulation. In B.C. Schipper, editor, *Proceedings of the 14. Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2013)*, pages 61–70. University of California, Davis, 2013.
- L. Epstein, M. Marinacci, and K.W. Seo. Coarse contingencies. *Theoretical Economics*, 2:355–394, 2007.
- C. Ewerhart. Heterogeneous awareness and the possibility of agreement. Universität Mannheim, 2001.
- R. Fagin and J. Halpern. Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34:39–76, 1988.
- R. Fagin, J. Halpern, and N. Meggido. A logic for reasoning about probability. *Information and Computation*, 87:78–128, 1990.
- R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning about knowledge*. MIT Press, 1995.
- Y. Feinberg. Characterizing common priors in form of posteriors. *Journal of Economic Theory*, 91: 127–179, 2000.
- Y. Feinberg. Subjective reasoning - games with unawareness. Stanford University, 2004.
- Y. Feinberg. Games with incomplete awareness. Stanford University, 2005.
- Y. Feinberg. Games with unawareness. Stanford University, 2012.
- E. Filiz-Ozbay. Incorporating unawareness into contract theory. *Games and Economic Behavior*, 76: 181–194, 2012.
- S. Galanis. Syntactic foundation for unawareness of theorems. *Theory and Decision*, 71:593–614, 2011.
- S. Galanis. Unawareness of theorems. *Economic Theory*, 52:41–73, 2013a.
- S. Galanis. The value of information under unawareness. University of Southampton, 2013b.
- S. Grant and J. Quiggin. Inductive reasoning about unawareness. *Economic Theory*, 54:717–755, 2013.
- S. Grant, J.J. Kline, and J. Quiggin. Differential awareness and incomplete contracts: A model of contractual disputes. *Journal of Economic Behavior and Organization*, 82:494–504, 2012.
- D. Grossi and F.R. Velázquez-Quesada. Twelve angry men: A study on the fine-grain of announcements. In X. He, J. Horty, and E. Pacuit, editors, *LORI 2009*, LNAI 5834, pages 147–160. Springer, 2009.
- J. Halpern. Alternative semantics for unawareness. *Games and Economic Behavior*, 37:321–339, 2001.
- J. Halpern and L.C. Rêgo. Interactive unawareness revisited. *Games and Economic Behavior*, 61: 232–262, 2008.
- J. Halpern and L.C. Rêgo. Reasoning about knowledge of unawareness. *Games and Economic Behavior*, 67:503–525, 2009.
- J. Halpern and L.C. Rêgo. Reasoning about knowledge of unawareness revisited. *Mathematical Social Sciences*, 65:73–84, 2013.

- J. Halpern and L.C. Rêgo. Extensive games with possibly unaware players. *Mathematical Social Sciences*, forthcoming, 2014.
- J. Harsanyi. Games with incomplete information played ‘bayesian’ players, part i, ii, and iii. *Management Science*, 14:159–182, 320–334, 486–502, 1967/68.
- A. Heifetz. The positive foundation of the common prior assumption. *Games and Economic Behavior*, 56:105–120, 2006.
- A. Heifetz and P. Mongin. Probability logic for type spaces. *Games and Economic Behavior*, 35:31–53, 2001.
- A. Heifetz and D. Samet. Topology-free typology of beliefs. *Journal of Economic Theory*, 82:324–341, 1998.
- A. Heifetz, M. Meier, and B.C. Schipper. Interactive unawareness. *Journal of Economic Theory*, 130:78–94, 2006.
- A. Heifetz, M. Meier, and B.C. Schipper. A canonical model for interactive unawareness. *Games and Economic Behavior*, 62:304–324, 2008.
- A. Heifetz, M. Meier, and B.C. Schipper. Prudent rationalizability in generalized extensive-form games. University of California, Davis, 2011.
- A. Heifetz, M. Meier, and B.C. Schipper. Unawareness, beliefs, and speculative trade. University of California, Davis, 2012.
- A. Heifetz, M. Meier, and B.C. Schipper. Unawareness, beliefs, and speculative trade. *Games and Economic Behavior*, 77:100–121, 2013a.
- A. Heifetz, M. Meier, and B.C. Schipper. Dynamic unawareness and rationalizable behavior. *Games and Economic Behavior*, 81:50–68, 2013b.
- S. Heinsalu. Universal type spaces with unawareness. *Games and Economic Behavior*.
- S. Heinsalu. Equivalence of the information structure with unawareness to the logic of awareness. *Journal of Economic Theory*, 147:2453–2468, 2012.
- B. Hill. Awareness dynamics. *Journal of Philosophical Logic*, 39:113–137, 2010.
- J. Hintikka. Impossible possible worlds vindicated. *Journal of Philosophical Logic*, 4:475–484, 1975.
- Z.S. Huang and K. Kwast. Awareness, negation and logical omniscience. In J. van Eijck, editor, *Logics in AI*, volume 478 of *Lectures Notes in Computer Science*, pages 282–300. Springer, 1991.
- E. Karni and M.L. Vierø. “reverse bayesianism”: A choice-based theory of growing awareness. *American Economic Review*, forthcoming, 2013a.
- E. Karni and M.L. Vierø. Probabilistic sophistication and reverse bayesianism. Queen’s University, 2013b.
- K. Konolige. What awareness isn’t: A sentential view of implicit and explicit belief. In J. Halpern, editor, *Proceedings of the First Conference on Theoretical Aspects of Reasoning and Knowledge (TARK I)*, pages 241–250. Morgan Kaufmann, 1986.
- R.V. Krishna and P. Sandowski. Dynamic preference for flexibility. University of North Carolina and Duke University, 2013.

- J. Lee. Unforeseen contingencies and renegotiation with asymmetric information. *Economic Journal*, 118:678–694, 2008.
- H.J. Levesque. A logic of implicit and explicit belief. In *AAAI-84 Proceedings*, pages 198–220, 1984.
- J. Li. Dynamic games of complete information with unawareness. University of Pennsylvania, 2006.
- J. Li. Interactive knowledge with unawareness. University of Pennsylvania, 2008a.
- J. Li. A note on unawareness and zero probability. University of Pennsylvania, 2008b.
- J. Li. Information structures with unawareness. *Journal of Economic Theory*, 144:977–993, 2009.
- S.X. Li, M. Peitz, and X.J. Zhao. Vertically differentiated duopoly with unaware consumers. *Mathematical Social Sciences*, forthcoming, 2014.
- Z. Liu. Fair disclosure and investor asymmetric awareness in stock markets. University of Buffalo, 2008.
- M. Meier. An infinitary probability logic for type spaces. *Israel Journal of Mathematics*, 192:1–58, 2012.
- M. Meier and B.C. Schipper. Speculative trade under unawareness: The infinite case. *Economic Theory Bulletin*.
- M. Meier and B.C. Schipper. Conditional dominance in games with unawareness. University of California, Davis, 2012.
- M. Meier and B.C. Schipper. Bayesian games with unawareness and unawareness perfection. *Economic Theory*, forthcoming, 2014.
- J.F. Mertens and S. Zamir. Formulation of bayesian analysis for games with incomplete information. *International Journal of Game Theory*, 14:1–29, 1985.
- P. Milgrom and N. Stokey. Information, trade and common knowledge. *Journal of Economic Theory*, 26:17–27, 1982.
- S. Modica. Unawareness, priors and posteriors. *Decisions in Economics and Finance*, 31:81–94, 2008.
- S. Modica and A. Rustichini. Awareness and partitional information structures. *Theory and Decision*, 37:107–124, 1994.
- S. Modica and A. Rustichini. Unawareness and partitional information structures. *Games and Economic Behavior*, 27:265–298, 1999.
- D. Monderer and D. Samet. Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1:170–190, 1989.
- R. Montague. Universal grammar. *Theoria*, 36:373–398, 1970.
- J.L. Montiel Olea. A simple characterization of trivial unawareness. Harvard University, 2012.
- S. Morris. The logic of belief and belief change: A decision theoretic approach. *Journal of Economic Theory*, 69:1–23, 1996.
- S. Morris. Alternative notions of knowledge. In M. Bacharach, L.-A. Gérard-Varet, P. Mongin, and H.S. Shin, editors, *Epistemic logic and the theory of games and decisions*, pages 217–234. Kluwer Academic Press, 1997.
- L.S. Moss and I.D. Viglizzo. Harsanyi type spaces and final coalgebras constructed from satisfied theories. *Electronic Notes in Theoretical Computer Science*, 106:279–295, 2004.

- K. Nehring. Preference for flexibility in a savage framework. *Econometrica*, 67:101–119, 1999.
- C.S. Nielsen and A. Sebal. Unawareness in dynamic psychological games. University of Copenhagen, 2012.
- M. Osborne and A. Rubinstein. *A course in game theory*. MIT Press, 1994.
- E. Ozbay. Unawareness and strategic announcements in games with uncertainty. University of Maryland, 2008.
- D.G. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52:1029–1050, 1984.
- M. Pinter and Z. Udvari. Generalized type-spaces. Corvinus University Budapest, 2012.
- C.P. Pires. Do i know ω ? an axiomatic model of awareness and knowledge. Universidade Nova de Lisboa, 1994.
- V. Rantala. Impossible world semantics and logical omniscience. *Acta Philosophica Fennica*, 35:106–115, 1982a.
- V. Rantala. Quantified modal logic: non-normal worlds and propositional attitudes. *Studia Logica*, 41: 41–65, 1982b.
- L.C. Rêgo and J. Halpern. Generalized solution concepts in games with possibly unaware players. *International Journal of Game Theory*, 41:131–155, 2012.
- T. Sadzik. Knowledge, awareness and probabilistic beliefs. Stanford University, 2007.
- D. Samet. Bayesianism without learning. *Research in Economics*, 53:227–242, 1999.
- B.C. Schipper. Awareness-dependent subjective expected utility theory. *International Journal of Game Theory*, 42:725–753, 2013.
- B.C. Schipper. Preference-based unawareness. *Mathematical Social Sciences*, forthcoming, 2014.
- B.C. Schipper and H.Y. Woo. Political awareness, microtargeting of voters, and negative electoral campaigning. University of California, Davis, 2014.
- D. Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57:571–587, 1989.
- D. Scott. Advice in modal logic. In K. Lambert, editor, *Philosophical problems in logic*, pages 143–173. Reidel, 1970.
- G. Sillari. Models of awareness. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Logic and the foundations of games and decisions*, pages 209–240. University of Amsterdam, 2008a.
- G. Sillari. Quantified logic of awareness and impossible possible worlds. *Review of Symbolic Logic*, 1: 514–529, 2008b.
- E.L. von Thadden and X.J. Zhao. Incentives for unaware agents. *Review of Economic Studies*, 79: 1151–1174, 2012a.
- E.L. von Thadden and X.J. Zhao. Multitask agency with unawareness. Hong Kong University of Science and Technology, 2012b.

- E. Thijsse. On total awareness logic. with special attention to monotonicity constraints and flexibility. In M. De Rijke, editor, *Diamond and defaults: Studies in pure and applied intensional logic*, pages 309–347. Kluwer, 1991.
- E. Thijsse and H. Wansing. A fugue on the themes of awareness logic and correspondence. *Journal of Applied Non-Classical Logics*, 6:127–136, 1996.
- O. Walker. Unawareness with “possible” possible worlds. *Mathematical Social Sciences*, forthcoming, 2014.
- H. Wansing. A general possible worlds framework for reasoning about knowledge and belief. *Studia Logica*, 49:523–539, 1990.
- A. Zeman. *Consciousness. A user’s guide*. Yale University Press, 2002.