A partial-state space model of unawareness

Wesley H. Holliday
University of California, Berkeley

November 9, 2022

Abstract

We propose a model of unawareness that remains close to the paradigm of Aumann's model for knowledge [R. J. Aumann, International Journal of Game Theory 28 (1999) 263-300]: just as Aumann uses a correspondence on a state space to define an agent's knowledge operator on events, we use a correspondence on a state space to define an agent's awareness operator on events. This is made possible by three ideas. First, like the model of [A. Heifetz, M. Meier, and B. Schipper, Journal of Economic Theory 130 (2006) 78-94], ours is based on a space of partial specifications of the world, partially ordered by a relation of further specification or refinement, and the idea that agents may be aware of some coarser-grained specifications while unaware of some finer-grained specifications; however, our model is based on a different implementation of this idea, related to forcing in set theory. Second, we depart from a tradition in the literature, initiated by [S. Modica and A. Rustichini, Theory Decision 37 (1994) 107-124] and adopted by Heifetz et al. and [J. Li, Journal of Economic Theory 144 (2009) 977-993], of taking awareness to be definable in terms of knowledge. Third, we show that the negative conclusion of a well-known impossibility theorem concerning unawareness in [Dekel, Lipman, and Rustichini, Econometrica 66 (1998) 159-173] can be escaped by a slight weakening of a key axiom. Together these points demonstrate that a correspondence on a partial-state space is sufficient to model unawareness of events. Indeed, we prove a representation theorem showing that any abstract Boolean algebra equipped with awareness, knowledge, and belief operators satisfying some plausible axioms is representable as the algebra of events arising from a partial-state space with awareness, knowledge, and belief correspondences.

JEL classification: C70; C72; D80; D83

Keywords: Unawareness; Knowledge; Belief; Interactive epistemology

E-mail: wesholliday@berkeley.edu

1 Introduction

In recent decades, models of uncertainty in economics have been enriched so as to also represent *unawareness*. If an agent is uncertain about an event or proposition E, then she can conceive of E but does not know whether E obtains. By contrast, if an agent is unaware of E, then E is not even "present to mind" (Modica and Rustichini 1994, p. 107, Modica and Rustichini 1999, p. 274); the agent has a "lack of conception" of E (Heifetz et al. 2006, p. 90, Schipper 2015). While unawareness is thereby distinguished from uncertainty, one prominent tradition in the literature, initiated by Modica and Rustichini (1994) and adopted by Heifetz, Meier, and Schipper (2006) and Li (2009), takes unawareness to be definable in terms of knowledge: an agent is unaware of E if and only if she does not know E but also does not know that she does not know E. Dually, an agent is aware of E if and only if she knows E or knows that she does not know E.

While conceptually parsimonious, taking unawareness to be definable in terms of knowledge as above limits the relevant phenomena that can be modeled. Say that an agent is *overconfident* when he not only believes E but also believes that he knows E, though he does not know E, say because E is false.[1] For example, a potential investor in a firm might believe he knows that the firm is profitable, while in fact it is unprofitable and is committing fraud to suggest otherwise. But according to the Modica-Rustichini definition of unawareness, it is impossible for an agent to be aware of E and overconfident with respect to E. Since the overconfident agent believes that he knows E, he does not know that he does not know E.[2] Then since he also does not know E, according to Modica-Rustichini the overconfident agent must be unaware of E. But on the intuitive interpretation of awareness as being "present to mind" or of unawareness as "lack of conception," the overconfident agent is perfectly aware of E. The investor is perfectly aware of the idea of profitability of the firm; indeed, he believes (falsely) that the firm is profitable. All the behavioral predictions implied by such awareness and belief apply to the investor. Thus, if we wish to model overconfident agents, we cannot accept the Modica-Rustichini definition of unawareness.[3]

Moreover, according to notions of belief as subjective certainty according to which believing E entails believing that you know E (as in Stalnaker 2006, § 3), the Modica-Rustichini definition makes it impossible for an agent to be aware of E and falsely believe E; then assuming belief requires awareness (as in Fagin and Halpern 1988), their definition makes it impossible for an agent to falsely believe E.

Yet we should not give up on modeling unawareness in the face of false or overconfident beliefs. The investor's overconfidence may be due in part to his unawareness of the possibility of a sophisticated type of fraud by the firm; one might predict that if he were made aware—or if we consider a different investor who is aware—he would realize that none of his due diligence ruled out such fraud. Thus, while overconfidence with respect to E should be compatible with awareness of E, it may be related to unawareness of some other event F. We would like to have a model of unawareness that can capture these phenomena.

In this paper, we propose a new model of unawareness. Instead of defining awareness in terms of knowledge, we model awareness in a way that remains close to the paradigm of Aumann's (1999a) model for knowledge: just as Aumann uses a correspondence on a state space to define an agent's knowledge operator

[1] As in Modica and Rustichini 1994, 1999, Heifetz et al. 2006, and Li 2009 (and the classics Aumann 1999a,b, Fagin et al. 1995, and Hintikka 1962), knowing E requires that E is true. As for belief, for the rest of the paragraph in the main text one may replace belief with *p-belief* (Monderer and Samet 1989), i.e., subjective probability at least p, for any $p \in (0, 1]$ for which one agrees that p-believing E is inconsistent with knowing $\neg E$ (not E). For example, in Aumann's (1999b, (12.2)) framework, knowing $\neg E$ implies 1-believing $\neg E$, so we can use any $p > 0$.

[2] This simply applies the principle that if an agent believes F, then he does not know $\neg F$ (not F), which we assume in the next paragraph as well.

[3] In Section 2.2, we argue that changing the Modica-Rustichini definition to define unawareness in terms of belief instead of knowledge still does not provide a satisfactory definition of unawareness of events.

on events, we use a correspondence on a state space to define an agent's awareness operator on events. Like the model of Heifetz, Meier, and Schipper (2006), our model is based on the idea of *partial specifications* of the world, partially ordered by a relation of *further specification* or *refinement*, and the idea that agents may be aware of some coarser-grained specifications while unaware of some finer-grained specifications. However, our model is based on a different implementation of this idea with a long history in mathematics and mathematical logic. In particular, as exploited in forcing in set theory (e.g., Takeuti and Zaring 1973), partial specifications ordered by refinement give rise to a Boolean algebra of events via the regular open sets in the downset topology of the partial order (see Section 3); crucially, in this Boolean algebra, the negation operation $\neg$ is not set-theoretic complementation, so there may be partial specifications that belong to neither E nor $\neg E$.[4] Next, as exploited in so-called possibility semantics for modal logic (Humberstone 1981; Holliday 2014, 2015; van Benthem et al. 2017; Holliday 2021), relating these partial specifications via accessibility relations—or equivalently, possibility correspondences—for different agents provides a model of multi-agent knowledge and belief, which generalizes the Kripke frames (Kripke 1963) or Aumann structures (Aumann 1976, 1999a) that drop out as special cases when using only discrete partial orders. The final step we take here is to add correspondences representing agents' awareness of partial possibilities: $\nu \in \mathcal{A}_i(\omega)$ will mean that in possibility ω, agent i is aware of possibility ν.

Just as in Aumann structures one uses possibility correspondences to define knowledge *of events*, in our structures we will use the $\mathcal{A}_i$'s to define awareness of events, where an event is understood as a set of possibilities in a Boolean algebra of events. This project of modeling awareness of events in some Boolean algebra must be distinguished from the project of modeling awareness of *sentences* in some language. Awareness of sentences may be *hyperintensional* in the sense that where $[\![\varphi]\!]$ is the set of possibilities in which a sentence φ is true in a given model, we may have $[\![\varphi_1]\!] = [\![\varphi_2]\!]$ while the agent is aware of the sentence φ_1 and yet unaware of the sentence φ_2.[5] Hyperintensional models of awareness of sentences have been developed in the literature (e.g., Fagin and Halpern 1988, Modica and Rustichini 1999, Halpern 2001), but here we follow the event-based tradition of Aumann (1976; 1999a) with a non-hyperintensional model of awareness of events (see Schipper 2015 for comparison of event-based and sentence-based approaches). However, we believe these two projects should be related by the following bridge principle: an agent is aware of an event E if and only if she is aware of some sentence φ that she understands such that $[\![\varphi]\!] = E$.[6]

Distinguishing awareness of events vs. sentences is crucial for assessing axioms concerning awareness. For example, awareness of events should not be monotonic with respect to the entailment relation $\leq$ in the Boolean algebra of events (which is just the subset relation $\subseteq$ when events are sets of possibilities); that is, we should not require that if $E \leq F$ (E entails F) and the agent is aware of E, then she must be aware of F. For example, an agent may be aware of the event E expressed by 'Ann and Bob will play a Nash equilibrium' without being aware of the event F expressed by 'Ann and Bob will play a correlated equilibrium', due to not having the concept of correlated equilibrium, despite the fact that $E \leq F$. Since $E \leq F$ is equivalent to $E = E \sqcap F$, where $\sqcap$ is the meet operation in the Boolean algebra (corresponding to intersection of sets), monotonicity is equivalent to the principle that if an agent is aware of $E \sqcap F$, then she must be aware of F. A spurious argument for this principle is that "if an agent is aware of a conjunction, then she must be aware of

[4] This is analogous to rejecting the *real states* assumption of Dekel et al. 1998, except we consider only events and not formulas in a logical language.

[5] The denial of hyperintensionality for awareness and knowledge of sentences is what Dekel et al. (1998) call *event sufficiency*.

[6] This bridge principle has consequences for operationalizing the concept of awareness of events. Assuming one has a decision procedure for testing awareness of sentences, one obtains a semi-decision procedure for testing awareness of a given event E: enumerate sentences $\varphi_1, \varphi_2, \ldots$ that express E (assuming that for the given E, there are at most countably infinitely many sentences in the agent's language that express E) and check for each sentence the agent's awareness of that sentence.

each conjunct." The argument is spurious because it implicitly assumes awareness of a *sentence*. It is indeed plausible that if an agent is aware of a sentence $\varphi \wedge \psi$, where $\wedge$ is sentential conjunction, then the agent must be aware of ψ. But the event $E \sqcap F$ is not intrinsically a conjunction. We may have $E = E \sqcap F = C \sqcup D$, etc. Thus, in contrast with some other models of awareness of events in the literature (e.g., Heifetz et al. 2006, Li 2009), our model will not validate the axiom that awareness of $E \sqcap F$ implies awareness of F.

Our definition of awareness of events is informally as follows. In a possibility ω, an agent i is aware of an event E if the following condition holds at ω and its refinements: if i is aware of a possibility ν, then i is aware of any coarsest refinement of ν belonging to E and any coarsest refinement of ν belonging to $\neg E$. In other words, if you are aware of E, then you should be able to apply the E vs. $\neg E$ distinction starting from any possibility of which you are aware. For example, if an agent is aware of the event E expressed by 'The Centers for Medicare and Medicaid Services are investigating the firm', then for any possibility ν of which the agent is aware, if ν does not already belong to E or $\neg E$, then the agent should be aware of a coarsest further specification of ν belonging to E and a coarsest further specification of ν belong to $\neg E$. This model of awareness satisfies the *symmetry axiom* of Modica and Rustichini 1994, stating that an agent is aware of E if and only if she is aware of $\neg E$ (though again without accepting the Modica-Rutschini definition of awareness in terms of knowledge), which helps distinguish being unaware of an event from assigning it zero probability (cf. Schipper 2013, p. 727). Given symmetry, one may also think of the model as a model of awareness of *distinctions* in the space of possibilities or of awareness of *binary questions*.

Besides symmetry, another consequence of our model is that each agent i is aware of the trivial event Ω and the trivial distinction of Ω vs. $\varnothing$. This does not imply that i is aware of *each possibility in* Ω. Nor does it imply that i is aware of each *sentence* true throughout Ω; e.g., it does not imply that i is aware of the sentence 'every Nash equilibrium is a correlated equilibrium'. This point is related to Stalnaker's (1984, pp. 85-6) defense of the fact that in standard state-space models of knowledge, such as Aumann structures, each agent *knows* the trivial event Ω: this does not imply that i knows that the sentence 'every Nash equilibrium is a correlated equilibrium' is true, since i may fail to know the metalinguistic fact that that sentence is true throughout Ω. Returning to awareness, the bridge principle proposed above implies that i is aware of Ω if and only if i is aware of some sentence φ such that $[\![\varphi]\!] = \Omega$, e.g., a sentence φ such as 'It is raining or it is not raining'. We therefore call the axiom that i is aware of Ω the *tautology axiom*.

A third consequence of our model is the *agglomeration axiom*: if i is aware of events E and F, then i is also aware of $E \sqcap F$, or set theoretically, $E \cap F$. The corresponding assumption on sentential awareness is that if i is aware of φ and of ψ, then i is aware of $\varphi \wedge \psi$. This assumes some logical sophistication on the part of i. Indeed, we assume that while our agents may have unawareness, they are logically perfect within the domain of their awareness. As Schipper (2015, p. 78) puts it, "Despite such a lack of conception [i.e., unawareness], agents in economics are still assumed to be fully rational," in contrast to some models in computer science of agents who are both unaware and logically imperfect (e.g., Fagin and Halpern 1988).

After handling awareness, we add knowledge and belief to our model, as in possibility semantics for epistemic modal logic. This part of our model is analogous to the standard treatment of knowledge and belief in Kripke frames and Aumann structures, allowing adjustments to deal with the partiality of possibilities. We call the resulting structures *epistemic possibility frames*. We use these frames to show that an influential impossibility result concerning awareness due to Dekel, Lipman, and Rustichini (1998) becomes a possibility result (Fact 4.15) under a slight weakening of one of their axioms.

We can now informally state the main technical result about our model, which is a representation theorem. We define an *epistemic awareness algebra* to be a Boolean algebra equipped with awareness operators for each agent satisfying the axioms of symmetry, tautology, and agglomeration, plus knowledge and belief operators satisfying some minimal axioms. We then prove (Theorem 5.6) that any epistemic awareness algebra is representable as the algebra of events of an epistemic possibility frame with awareness operators represented using the awareness correspondences $\mathcal{A}_i$ as sketched above and with knowledge and belief operators represented using knowledge and belief correspondences. This theorem shows that our model of awareness given by a partially ordered set equipped with regular open events and the correspondences for awareness, knowledge, and belief is highly versatile: it can represent any situation involving multiple agents' event-based awareness, knowledge, and belief, provided some basic axioms are satisfied.

There is now a large literature on unawareness in economics and computer science, as surveyed up to around 2014 in Schipper 2014a, 2015. A non-exhaustive list of more recent contributions in economics includes Grant et al. 2015, Quiggin 2016, Karni and Vierø 2017, Galanis 2016, 2018, Piermont 2017, Dietrich 2018, Guarino 2020, Fukuda 2021, Schipper 2021b, Dominiak and Tserenjigmid 2022, and the special issue introduced by Schipper 2021a. We will discuss those works most relevant to the present paper in Section 2.

1.1 Organization

The rest of the paper is organized as follows. In Section 2, we review as background the impossibility theorem of Dekel et al. (1998) that threatens to preclude our development of a model of awareness of events (2.1). We argue that one of their axioms is too strong (2.2), so we may proceed with developing our model, after discussing some related models (2.3). In Section 3, we review mathematical preliminaries that underly our model. In Section 4, we introduce the model of awareness (4.1) and then add knowledge and belief (4.2), using the model to formalize several examples, after which we briefly discuss some decision-theoretic considerations related to the model (4.3). In Section 5, we state our main representation theorem. In Section 6, we conclude with directions for future work, including a sketch of how to add awareness of sentences and probability to our framework. All substantial proofs are collected in the Appendix.

A Jupyter notebook containing code to verify conditions and compute awareness, knowledge, and belief in all examples in this paper is available at https://github.com/wesholliday/awareness.

2 Background

2.1 The relation of unawareness and knowledge

Several previous works treat the foundational issue of how to model awareness of events as opposed to sentences. Famously, Dekel et al. (1998) prove theorems that are supposed to raise problems for such models, so we must explain why our project is not doomed by these results. Although Dekel et al. phrase their theorems in terms of state-space models, they are much more general. Let $\mathcal{E}$ be a nonempty set and $\leq$ a preorder on $\mathcal{E}$ with a minimum element 0. Think of elements of $\mathcal{E}$ as events and $\leq$ as the entailment relation between events. For example, given a nonempty set Ω, we could take $\mathcal{E}$ to be the powerset of Ω, $\leq\,=\,\subseteq$, and $0 = \varnothing$, but this is just one example. Next, we assume maps $U : \mathcal{E} \to \mathcal{E}$ for *unawareness*, $K : \mathcal{E} \to \mathcal{E}$ for *knowledge*, and $\neg : \mathcal{E} \to \mathcal{E}$ for *negation*.[7] Suppose these maps satisfy the following axioms for all $E \in \mathcal{E}$:

[7]Stipulating some operations U and K on $\mathcal{E}$ is of course not to provide any illuminating model of unawareness and knowledge, but this abstract setup will be useful for stating impossibility theorems. The actual models of unawareness discussed below (in Heifetz et al. 2006, Li 2009, Fritz and Lederman 2015, and the present paper) all attempt to *represent* unawareness operations using more concrete structures. To provide any additional insight or representational succinctness beyond stipulating a primitive $U : \mathcal{E} \to \mathcal{E}$, the model must derive such an operation from more concrete relations, correspondences, etc. on a set of states. For an example in which this requirement is not satisfied, note that in a state-space model based on a field of sets $(\Omega, \mathcal{E})$, stipulating

- $U(E) \leq U(U(E))$ (AU Introspection);
- $U(E) \leq \neg K(E)$ and $U(E) \leq \neg K \neg K(E)$ (Plausibility);
- $K(U(E)) = 0$ (KU Introspection);
- $K(\neg 0) = \neg 0$ (Necessitation);
- $\neg\neg 0 = 0$ (Double Negation).

AU Introspection says that if an agent is unaware of E, then she is unaware that she is unaware of E, while KU introspection says that it is impossible for an agent to know that she is unaware of a specific event E. Necessitation says that the agent knows the trivial event $\neg 0$, namely Ω in state-space models, and Double Negation says that negation is involutive as in classical logic and Boolean algebras.

Plausibility is one direction of the Modica-Rustichini (1994) definition of unawareness in terms of knowledge. In fact, only the second half of Plausibility ($U(E) \leq \neg K \neg K(E)$) is used for the following result.

Proposition 2.1 (Dekel et al. 1998, Theorem 1(i))**.** Assuming the axioms above, $U(E) = 0$ for all $E \in \mathcal{E}$.

Proof. $U(E) \leq U(U(E)) \leq \neg K \neg K(U(E)) \leq \neg K \neg 0 = \neg\neg 0 = 0$. □

Thus, unawareness is contradictory and hence impossible assuming the axioms above.[8]

Note, by contrast, that our argument concerning overconfidence at the beginning of Section 1 targeted instead the converse of Plausibility, namely $\neg K(E) \sqcap \neg K \neg K(E) \leq U(E)$ (Converse Plausibility), where we assume that for all $E, F \in \mathcal{E}$, there is a greatest lower bound $E \sqcap F$ of $\{E, F\}$ in $(\mathcal{E}, \leq)$.

Proposition 2.2. In addition to K, U, and $\neg$, assume a map $B : \mathcal{E} \to \mathcal{E}$ such that for all $E \in \mathcal{E}$, $B(E) \leq \neg K \neg E$ (Noncontradictory Belief and Knowledge). Further assume Converse Plausibility. Then for all $E \in \mathcal{E}$, we have (i) $B(K(E)) \sqcap \neg K(E) \leq U(E)$. Thus, if $B(E) \sqcap U(E) = 0$ (Belief Requires Awareness), then (ii) $B(E) \sqcap B(K(E)) \sqcap \neg K(E) = 0$.

Proof. For (i), $B(K(E)) \leq \neg K \neg K(E)$, so $B(K(E)) \sqcap \neg K(E) \leq \neg K(E) \sqcap \neg K \neg K(E) \leq U(E)$. □

Thus, Converse Plausibility essentially renders overconfidence impossible, given that the other axioms are uncontroversial. We will also give an argument against the second half of Plausibility in Section 2.2.

Although Dekel et al. originally framed their Theorem 1(i) as a result about standard state-space models, Proposition 2.1 applies to any model of awareness that provides a structure $(\mathcal{E}, \leq, 0, U, K, \neg)$. For example, the model of Heifetz, Meier, and Schipper (2006)[9] provides a set $\mathcal{E}$ of events constructed as certain pairs (E, S) of sets ordered by $(E, S) \leq (E', S')$ if $E \subseteq E'$ and $S \preceq S'$, where $\preceq$ is a complete lattice order (see Schipper 2013, §§ 2.1, 2.3, 2.5). Hence the minimum element is $0 = (\varnothing, S_\perp)$ where $S_\perp$ is the minimum element of $\preceq$. The model also provides maps U, K, and $\neg$ from $\mathcal{E}$ to $\mathcal{E}$. Since this model allows for unawareness, it must reject one of Dekel et al.'s axioms. Indeed, it rejects $K(U(E)) = 0$ (KU Introspection). Heifetz et al. instead set $K(U(E))$ to be a certain non-minimum element $(\varnothing, S(E))$ of $(\mathcal{E}, \leq)$; although they denote that non-minimum element by $\emptyset^{S(E)}$ and refer to the axiom $K(U(E)) = \emptyset^{S(E)}$ as "KU Introspection," the event $\emptyset^{S(E)}$ is not equal to the minimum element 0 in $(\mathcal{E}, \leq)$, so their "KU Introspection" is not Dekel et al.'s axiom. The model of Li 2009 similarly rejects Dekel et al.'s KU Introspection axiom.

an operation $U : \mathcal{E} \to \mathcal{E}$ is equivalent to stipulating a neighborhood function $N_U : \Omega \to \wp(\mathcal{E})$ via the definition: $E \in N_U(\omega)$ if and only if $\omega \in U(E)$. Thus, this repackaging with a neighborhood function (which merely lists the events of which an agent is supposed to be unaware at a state ω) offers no additional insight or representational succinctness.

[8] Dekel et al. prove another result (Theorem 1(ii)) that replaces Necessitation with the Monotonicity of K, i.e., if $E \leq F$, then $K(E) \leq K(F)$, but this principle is unacceptable assuming knowledge requires awareness, for reasons similar to those for rejecting monotonicity of awareness above (cf. Modica and Rustichini 1994, p. 123, Modica and Rustichini 1999).

[9] See Heifetz et al. 2008, Halpern and Rêgo 2008, and Belardinelli and Rendsvig 2020 on the relation between this model and syntactic, logical approaches.

Like Heifetz et al. and Li, Fritz and Lederman (2015, § 2.3) respond to Dekel et al.'s result by rejecting their package of axioms. Taking inspiration from a distinction Dekel et al. draw between "real" and "subjective" states in state-space models, Fritz and Lederman argue that AU Introspection, Plausibility, and KU Introspection need only hold with respect to some distinguished set S of states, rather than with respect to the set Ω of all states. This means restricting the axioms as follows: $U(E) \sqcap S \leq U(U(E))$; $U(E) \sqcap S \leq \neg K(E)$ and $U(E) \sqcap S \leq \neg K \neg K(E)$; and $K(U(E)) \sqcap S = 0$. Fritz and Lederman (2015, Theorem 2) show that there are state-space models of unawareness of events satisfying these restricted axioms, plus Necessitation and Double Negation, in which unawareness is possible. More recently, Fukuda (2021) has studied other models violating AU Introspection in particular.

Thus, there is precedent in the literature for retaining an event-based approach to awareness despite the results of Dekel et al., by rejecting one or more of their axioms. In the following, we will raise doubts about their axioms in light of the distinction between awareness of events vs. awareness of sentences.[10]

2.2 Plausibility

In particular, we will raise doubt about the second half of Plausibility: $U(E) \leq \neg K \neg K(E)$. Suppose F is some in principle *unknowable* event, meaning $K(F) = 0$, such that our agent is also unaware of F. According to KU Introspection, $U(E)$ is an unknowable event, but we need not use KU Introspection; under highly plausible assumptions, $E \sqcap \neg U(E) \sqcap \neg K(E)$ is unknowable[11] (see the proof of Proposition 2.3). Now since F is unknowable, $\neg K(F)$ is equal to Ω (or abstractly, $\neg 0$). But any agent knows Ω (again cf. Stalnaker 1984), being careful not to confuse Ω with any particular sentence true throughout Ω. Thus, the agent's unawareness of F contradicts Plausibility, which implausibly implies that an agent who knows Ω must be aware of every unknowable event. It is surely true that knowledge of a *sentence* $\neg \mathtt{K}\varphi$ implies awareness of the *embedded sentence* φ. But unlike sentences, events do not intrinsically "embed" other events; and there is no reason to suppose that an agent who knows Ω thinks about the event Ω as the event expressed by the sentence $\neg \mathtt{K}\varphi$.[12] We will return to these issues in Section 4.2. For now, let us prove the claim above.

Proposition 2.3. Assume Necessitation and Double Negation as in Section 2.1. Further suppose $E \sqcap \neg E = 0$ (Noncontradiction), $K(E) \leq E$ (Factivity), and that if $E \leq F$, then $K(E) \sqcap \neg U(F) \leq K(F)$ (Awareness-restricted Monotonicity, to which we return in Section 5). Finally, assume the second half of Plausibility. Then for every $E \in \mathcal{E}$, we have $U(E \sqcap \neg U(E) \sqcap \neg K(E)) = 0$.

Proof. Where $E' = E \sqcap \neg U(E) \sqcap \neg K(E)$, we have $K(E') \leq E' \leq \neg U(E)$ by Factivity, which implies $K(E') \leq K(E') \sqcap \neg U(E) \leq K(E)$ using Awareness-restricted Monotonicity with $E' \leq E$. But we also have $K(E') \leq E' \leq \neg K(E)$ using Factivity. Thus, $K(E') \leq K(E) \sqcap \neg K(E) = 0$, so $K(E') = 0$. Hence by Necessitation, $K \neg K(E') = \neg 0$. Then by the second half of Plausibility, $U(E') \leq \neg K \neg K(E') = \neg\neg 0 = 0$. □

For E' as in the proof, we should be able to have $E' \neq 0$ and $E' \neq \neg 0$, so the idea that it is impossible to be unaware of E' is highly counterintuitive; indeed, we will see how such unawareness is possible in Example 4.13 (see Footnote 23). We take Proposition 2.3 to cast doubt on the second half of Plausibility, rather than

[10] After writing this paper, I learned from Harvey Lederman that Elliot 2022 raises similar doubts about the axioms.

[11] $E \sqcap \neg K(E)$ is the classic example of an unknowable event from what is known as Fitch's paradox (Fitch 1963); for example, can Ann know the event expressed by "Bob played left but Ann doesn't know it"? We could use this simpler event and the axiom $K(E \sqcap \neg K(E)) = 0$, but we will instead derive $K(E \sqcap \neg U(E) \sqcap \neg K(E)) = 0$ from other axioms.

[12] Where $A(E) = \neg U(E)$, these points also show the problem with the principle $A(\neg K(E)) \leq A(E)$ when $\neg K(E)$ is Ω. Note that this principle follows from $A(K(E)) \leq A(E)$ (one direction of AK-Self Reflection in Heifetz et al. 2006, Prop. 3) and $A(F) = A(\neg F)$ (symmetry), casting doubt on the former when $K(E) = 0$. See Footnote 26 for a restricted principle.

any of the other axioms. Fortunately, we will see that a slight weakening of Plausibility delivers us away from Dekel et al.'s (1998) impossibility theorem to a possibility result (Fact 4.15).

Finally, we address whether the problems with the Modica-Rustichini definition of unawareness in terms of knowledge might be solved by changing the definition to use belief: $U(E) = \neg B(E) \sqcap \neg B \neg B(E)$. The answer is that under a standard assumption of models of rational belief, namely that $B(E \sqcap \neg B(E)) = 0$ (No Moorean Beliefs[13]), the same kind of argument given above against the second half of Plausibility can be applied to the second half of Belief Plausibility, i.e., to $U(E) \leq \neg B \neg B(E)$.[14]

Proposition 2.4. Assume Double Negation, Necessitation for B instead of K, and No Moorean Beliefs. In addition, assume the second half of Belief Plausibility. Then for all events $E \in \mathcal{E}$, we have $U(E \sqcap \neg B(E)) = 0$.

Proof. Where $E' = E \sqcap \neg B(E)$, we have $B(E') = 0$ by No Moorean Beliefs, so $B\neg B(E') = \neg 0$ by Necessitation, in which case the second half of Belief Plausibility yields $U(E') \leq \neg B \neg B(E') = \neg\neg 0 = 0$. □

In short, Belief Plausibility implausibly implies that an agent who believes Ω must be aware of every *unbelievable* event E' as in the proof. To see the implausibility, note that many people at the beginning of World War II were unaware of the unbelievable E': a nuclear weapon will end the war, but we do not believe that a nuclear weapon will end the war. Thus, changing the Modica-Rustichini definition of unawareness to use belief instead of knowledge still does not provide a satisfactory definition of unawareness of events.

2.3 Related models

In addition to rejecting Plausibility, Fritz and Lederman 2015 overlaps with the present paper in a second way, by providing a model of awareness of events that also validates the axioms of symmetry, tautology, and agglomeration, while also rejecting the principle that awareness of $E \sqcap F$ implies awareness of F (cf. Theorem 3 in their paper). They do so via a very different construction than ours, which we take to be further evidence of the naturality of the three axioms. In particular, they introduce structures $(\Omega, \{\approx_i\}_{i \in I})$ where Ω is a nonempty set of states and $\approx_i$ assigns to each $\omega \in \Omega$ an equivalence relation $\approx^i_\omega$ on Ω; then agent i is aware of $E \subseteq \Omega$ in state ω, denoted $\omega \in a^i(E)$, if and only if for all ρ and τ such that $\rho \approx^i_\omega \tau$, we have $\rho \in E$ if and only if $\tau \in E$. This approach and ours locate the complexity of modeling awareness in different places. Theirs is more complex in assigning to each state ω an equivalence relation $\approx^i_\omega$ on Ω, whereas ours simply assigns a subset $\mathcal{A}_i(\omega) \subseteq \Omega$; but ours is more complex in working with a partially ordered set $(\Omega, \sqsubseteq)$, whereas theirs simply works with a set Ω. There is also a fundamental difference in representational power. On the equivalence relation approach, the family of all $E \subseteq \Omega$ of which an agent is aware must be an *atomic* algebra of sets,[15] whose atoms are the equivalence classes. By contrast, on our approach, even if our ambient set of events is an atomic Boolean algebra (which it need not be), there is no requirement that the subalgebra of events of which the agent is aware is an atomic Boolean algebra (see Example 4.8); the agent may be

[13] The name is taken from what is known as Moore's paradox (see Hintikka 1962, § 4.5 and Holliday and Icard 2010). The standard axiomatization of the *implicit belief* operator (not requiring awareness), denoted by L in Fagin and Halpern 1988, entails No Moorean Belief for L in place of B. Then given that explicit belief (requiring awareness) entails implicit belief, we can derive $B(E \sqcap \neg L(E)) = 0$ and replace the conclusion of Proposition 2.4 with the equally unappealing $U(E \sqcap \neg L(E)) = 0$.

[14] Note that the belief modification of the Modica-Rustichini definition assumes an agent who is not mistaken about what she believes ($B(B(E)) \leq B(E)$). For otherwise we could have an agent who is aware of E, does not believe E, but believes that she does believe E, so she does not believe that she does not believe E, contradicting the modified definition. From here it is a short step to the assumption that the agent is not mistaken about what she does not believe and then to No Moorean Beliefs.

[15] An *atom* in a Boolean algebra $(B, \leq)$ (see Section 3 for the definition of Boolean algebras as special partially ordered sets) is an $a \in B$ such that $0 < a$ and there is no b with $0 < b < a$. A Boolean algebra is *atomic* if for each $b \in B$, there is an atom $a \leq b$. By contrast, it is *atomless* if it has no atoms. An algebra $\mathcal{E}$ of sets is atomic (resp. atomless) if it is atomic (resp. atomless) when regarded as a Boolean algebra $(\mathcal{E}, \subseteq)$.

aware of only the events in some standard non-atomic algebra used in measure theory. Finally, Fritz and Lederman (2015, Appendix B) prove the weak completeness,[16] with respect to a semantics based on their structures, of a modal logic of awareness with axioms of symmetry, tautology, and agglomeration. Although we do not introduce logical syntax in this paper, our representation theorem for arbitrary awareness algebras (Theorem 5.6) rather immediately yields strong completeness theorems for the logic of awareness—and extensions thereof—but now with respect to a semantics based on our epistemic possibility frames.

As far as we know, the only other work to consider both awareness in economics and possibility semantics from modal logic is Piermont 2021. Piermont relates both to his concept of a *relativized Boolean algebra*, which is an algebraic structure satisfying some of the laws of Boolean algebras but not $E \sqcup \neg E = 1$.[17] By contrast, we work with the Boolean algebra of regular open sets canonically associated with a partially ordered set (Theorem 3.3); and we explain away intuitions that due to unawareness $E \sqcup \neg E$ may fail to equal Ω by insisting on the distinction between events and sentences (see Example 4.6). Our classical approach makes available all standard results, e.g., from measure theory, that are applicable to Boolean algebras.

3 Preliminaries

Standard representations of uncertainty begin with a nonempty set Ω of *states of the world*, whose powerset we denote by $\wp(\Omega)$. The set of *events* is then some nonempty collection $\mathcal{E} \subseteq \wp(\Omega)$ closed under at least finite intersection and complement relative to Ω. Here we instead begin with a partially ordered set (poset) $(\Omega, \sqsubseteq)$. We think of elements of Ω as *partial possibilities* and take $\omega \sqsubseteq \nu$ to mean that ω is a *further specification* or *refinement* of ν. For example, the possibility ν may settle that *Ann plays up* in a game but leave undetermined what Bob plays; then a refinement ω of ν may settle not only that *Ann plays up* but also that *Bob plays left*, while another refinement ω' of ν may settle that *Ann plays up* and *Bob plays right*. With this picture, not every subset of Ω is eligible to count as an event. We will delimit the eligible events shortly.

Given a poset $(\Omega, \sqsubseteq)$, we define the *downward closure* operation $\downarrow$ on $\wp(\Omega)$ by

$$\downarrow E = \{\omega \in \Omega \mid \text{for some } \nu \in E, \omega \sqsubseteq \nu\}.$$

A set $E \subseteq \Omega$ is a *downset of* $(\Omega, \sqsubseteq)$ if $E = \downarrow E$. For $\omega \in \Omega$, we write $\downarrow\omega$ for $\downarrow\{\omega\}$, called a *principal downset*. Possibilities ω and ν are *compatible*, denoted $\omega \between \nu$, if $\downarrow\omega \cap \downarrow\nu \neq \varnothing$; otherwise they are *incompatible*, denoted $\omega \perp \nu$. The *downset topology* on Ω is the topology whose open sets are exactly the downsets of $(\Omega, \sqsubseteq)$. The interior and closure operations for this topology are given by

$$\begin{aligned} \mathsf{int}(E) &= \{\omega \in \Omega \mid \text{for all } \nu \sqsubseteq \omega,\ \nu \in E\} \\ \mathsf{cl}(E) &= \{\omega \in \Omega \mid \text{for some } \nu \sqsubseteq \omega,\ \nu \in E\}. \end{aligned}$$

[16] As usual in logic (see, e.g., Blackburn et al. 2001, p. 194), a *weak completeness* theorem states that for every formula φ, if φ is semantically valid, then φ is syntactically provable; a *strong completeness* theorem states that for every set Γ of formulas and formula φ, if φ is a semantic consequence of Γ, then φ is syntactically provable from assumptions in Γ.

[17] A standard construction (see, e.g., Givant and Halmos 2009, pp. 91-2) associates with a Boolean algebra $\mathbb{B}$ and element a in $\mathbb{B}$ a new Boolean algebra $\mathbb{B}(a)$, the *relativization of* $\mathbb{B}$ *to* a, whose bottom element and meet operation coincide with those of $\mathbb{B}$ but whose top element 1_a is a and whose complement and join operations are defined by $\neg_a c = \neg c \sqcap a$ and $c \sqcup_a d = (c \sqcup d) \sqcap a$. Piermont's relativized Boolean algebras can be obtained by lifting this construction to $\mathbb{B}_\star = \{(a, b) \mid a, b \in \mathbb{B}, a \leq b\}$ as follows: $\neg_\star(a, b) = (\neg_b a, b)$, $(a, b) \sqcap_\star (a', b') = (a \sqcap a', b \sqcap b')$, $(a, b) \sqcup_\star (a', b') = (a \sqcap_{b \sqcap b'} a', b \sqcap b')$, $0_\star = (0, 0)$, and $1_\star = (1, 1)$.

The *regularization* operation $\rho : \wp(\Omega) \to \wp(\Omega)$ is then defined by

$$\rho(E) = \mathsf{int}(\mathsf{cl}(\downarrow E)) = \{\omega \in \Omega \mid \forall \omega' \sqsubseteq \omega\, \exists \omega'' \sqsubseteq \omega' : \omega'' \in \downarrow E\}.$$

A set $E \subseteq \Omega$ is *regular open* if $\rho(E) = E$. Let $\mathcal{RO}(\Omega, \sqsubseteq)$ be the collection of all regular open sets. Note that if $\sqsubseteq$ is the discrete partial order (i.e., the identity relation) on Ω, then $\mathcal{RO}(\Omega, \sqsubseteq)$ is just $\wp(\Omega)$.

Regular open sets can be characterized by the following conditions. The proof is straightforward.

Lemma 3.1. Given a poset $(\Omega, \sqsubseteq)$ and $E \subseteq \Omega$, we have $E \in \mathcal{RO}(\Omega, \sqsubseteq)$ if and only if for all $\omega, \omega' \in \Omega$:

1. *persistence*: if $\omega \in E$ and $\omega' \sqsubseteq \omega$, then $\omega' \in E$;
2. *refinability*: if $\omega \notin E$, then $\exists \omega' \sqsubseteq \omega\ \forall \omega'' \sqsubseteq \omega'\ \omega'' \notin E$.

Only regular open sets will count as genuine *events* in our model. Persistence states that if a possibility ω settles that an event holds, so does any refinement of ω. Refinability states that if a possibility ω does not settle that an event holds, then there is a refinement ω' of ω that settles that the event does *not* hold, in the sense that no possible refinement of ω' settles that the event holds (cf. the definition of $\neg$ in Theorem 3.3).

A poset is *separative* if for any $\omega \in \Omega$, its principal downset $\downarrow\omega$ is a regular open set, which may be described as the event: *the possibility ω obtains*. One may assume without loss of generality in what follows that all posets are separative.[18] An example of a non-separative poset is a two-element linear order with $\omega \sqsubseteq \nu$; the principal downset $\{\omega\}$ does not satisfy refinability, since $\nu \notin \{\omega\}$ and yet there is no refinement of ν all of whose refinements are not in $\{\omega\}$. We can turn this into a separative poset by adding a third element ω' with $\omega' \sqsubseteq \nu$, $\omega' \not\sqsubseteq \omega$, and $\omega \not\sqsubseteq \omega'$, so we obtain the three-element tree; now $\{\omega\}$ satisfies refinability, since ν is refined by ω', all of whose refinements (namely just ω' itself) are not in $\{\omega\}$.

Given an arbitrary set E of possibilities, $\rho(E)$ is the event that may be described as: *one of the possibilities in E obtains*. In the language of lattice theory, ρ is not only a *closure operator* (satisfying conditions 1-3 in Lemma 3.2) but also a *nucleus* (satisfying 4 in addition to 1-3) on $\wp(\Omega)$. The proof is straightforward.

Lemma 3.2. The map ρ satisfies the following for all downsets $E, F \subseteq \Omega$: 1. if $E \subseteq F$, then $\rho(E) \subseteq \rho(F)$; 2. $E \subseteq \rho(E)$; 3. $\rho(\rho(E)) = \rho(E)$; 4. $\rho(E) \cap \rho(F) \subseteq \rho(E \cap F)$.

Next we characterize the poset $(\mathcal{RO}(\Omega, \sqsubseteq), \subseteq)$. Recall that a poset $(L, \leq)$ is a *lattice* (resp. *complete lattice*) if every two-element subset $\{a, b\} \subseteq L$ (resp. every subset $\{a_i\}_{i\in I} \subseteq L$) has a least upper bound with respect to $\leq$, called the *join* and denoted $a \sqcup b$ (resp. $\bigsqcup_{i\in I} a_i$) and greatest lower bound with respect to $\leq$, called the *meet* and denoted $a \sqcap b$ (resp. $\bigsqcap_{i\in I} a_i$). A lattice is *bounded* if in addition it has a greatest element with respect to $\leq$, denoted 1, and a least element with respect to $\leq$, denoted 0.[19] A bounded lattice is *complemented* if for every $a \in L$, there is an $\neg a \in L$, called a *complement of* a, such that $a \sqcup \neg a = 1$ and $a \sqcap \neg a = 0$. A lattice is *distributive* if for all $a, b, c \in L$, we have $a \sqcap (b \sqcup c) = (a \sqcap b) \sqcup (a \sqcap c)$. A *Boolean algebra* is a complemented distributive lattice $\mathbb{B} = (B, \leq)$. We abuse notation and write $a \in \mathbb{B}$ for $a \in B$.

It is a classic result in lattice theory that the fixpoints of any closure operator on $\wp(\Omega)$, ordered by $\subseteq$, form a complete lattice with meet as intersection and join as closure of union (Burris and Sankappanavar 1981, Thm. 5.2); and as Tarski (1937) observed, the fixpoints of the regularization operation ρ form a complete Boolean algebra. For modern proofs of the following, see, e.g., Takeuti and Zaring 1973, Thms. 1.30, 1.40.

[18] One can always pass to a quotient poset by identifying ω and ω' when $\rho(\{\omega\}) = \rho(\{\omega'\})$, resulting in a separative poset with an isomorphic algebra of regular open sets.

[19] Every complete lattice is bounded, since the least upper bound of $\varnothing$ is 0 and the greatest lower bound of $\varnothing$ is 1.

Theorem 3.3. For any poset $(\Omega, \sqsubseteq)$, the poset $(\mathcal{RO}(\Omega, \sqsubseteq), \subseteq)$ is a complete Boolean algebra, called the *regular open algebra of* $(\Omega, \sqsubseteq)$, in which the Boolean complement, meet, and join are given by:

$$\neg E \quad = \quad \mathsf{int}(\Omega \setminus E) = \{\omega \in \Omega \mid \forall \omega' \sqsubseteq \omega \; \omega' \notin E\} \tag{1}$$

$$\bigsqcap_{i \in I} E_i \quad = \quad \bigcap_{i \in I} E_i \tag{2}$$

$$\bigsqcup_{i \in I} E_i \quad = \quad \rho\Big(\bigcup_{i \in I} E_i\Big) = \{\omega \in \Omega \mid \forall \omega' \sqsubseteq \omega \, \exists \omega'' \sqsubseteq \omega' \, \exists i \in I : \omega'' \in E_i\}. \tag{3}$$

Conversely, each complete Boolean algebra $(B, \leq)$ is isomorphic to $\mathcal{RO}(B_+, \leq_+)$, where B_+ is the set of nonzero elements of the algebra and $\leq_+$ is $\leq$ restricted to B_+, via the map $b \mapsto \{a \in B_+ \mid a \leq b\}$.

Remark 3.4. In the regular open algebra of a non-discrete poset, often $\bigsqcup_{i \in I} E_i \supsetneq \bigcup_{i \in I} E_i$. This represents a deep difference between our approach and that of Heifetz, Meier, and Schipper (2006), who define the join of events in such a way that typically $\bigsqcup_{i \in I} E_i \subsetneq \bigcup_{i \in I} E_i$.

In applications to reasoning under uncertainty (e.g., involving probability) a Boolean algebra of events is often not complete (though it may be countably complete), so we do not want to restrict attention to only representing complete Boolean algebras of events. To represent arbitrary Boolean algebras, we can equip a poset $(\Omega, \sqsubseteq)$ with a distinguished Boolean subalgebra of $\mathcal{RO}(\Omega, \sqsubseteq)$.

Definition 3.5. A *possibility frame* is a triple $(\Omega, \sqsubseteq, \mathcal{E})$ where $(\Omega, \sqsubseteq)$ is a poset and $\mathcal{E}$ is a nonempty subset of $\mathcal{RO}(\Omega, \sqsubseteq)$ closed under binary intersection and the operation $\neg$ from (1).

Compare the notion of a possibility frame to the more familiar notion of a *field of sets*, a pair $(\Omega, \mathcal{E})$ where Ω is a nonempty set and $\mathcal{E}$ is an algebra of subsets of Ω as at the beginning of this section. It is a classic result of Stone (1936) that for each Boolean algebra $\mathbb{B}$, there is a field of sets $(\Omega, \mathcal{E})$ such that $\mathbb{B}$ is isomorphic to $(\mathcal{E}, \subseteq)$. Since every field of sets may be regarded as a possibility frame $(\Omega, \sqsubseteq, \mathcal{E})$ in which $\sqsubseteq$ is the discrete partial order, Stone's theorem immediately implies that each Boolean algebra is representable by a possibility frame. But we will need non-discrete partial orders to model unawareness using possibility frames. As a consequence of Theorem 5.6, we will obtain for each Boolean algebra $\mathbb{B}$ a possibility frame $(\Omega, \sqsubseteq, \mathcal{E})$ with a non-discrete partial order such that $\mathbb{B}$ is isomorphic to $(\mathcal{E}, \subseteq)$.

Given a poset $(\Omega, \sqsubseteq)$ and $E \in \mathcal{RO}(\Omega, \sqsubseteq)$, we denote the set of maximal elements of E by

$$\max(E) = \{\omega \in E \mid \text{there is no } \nu \in E : \omega \sqsubset \nu\},$$

where $\omega \sqsubset \nu$ means that $\omega \sqsubseteq \nu$ and $\nu \not\sqsubseteq \omega$. Intuitively, $\max(E)$ contains the *coarsest* or *least refined* possibilities that settle that E holds. It is a natural thought that for any nonempty event E, there should be a *unique* coarsest possibility belonging to E, describable as "the possibility that E holds"; this is indeed the case for the possibility frame $(B_+, \leq_+, \mathcal{RO}(\Omega, \sqsubseteq))$ used in Theorem 3.3, and in fact we shall see that every Boolean algebra (not only complete ones) can be represented by a possibility frame satisfying this condition (Theorem 5.6). However, since in applications we typically wish to draw as few possibilities as possible to model a given situation (cf. Examples 4.7 and 4.13), we impose only the following less demanding condition.

Definition 3.6. A possibility frame $(\Omega, \sqsubseteq, \mathcal{E})$ is *quasi-principal* if for any $E \in \mathcal{E}$ and $\omega \in E$, we have $\omega \in {\downarrow}\max(E)$.

In other words, any possibility that settles that E holds is a refinement of some coarsest possibility that settles that E holds. Of course any possibility frame with Ω finite satisfies this condition.

4 Model

In this section, we introduce our model in two stages: first concentrating on awareness in Section 4.1 and then adding knowledge and belief alongside awareness in Section 4.2. We then briefly discuss decision-theoretic considerations related to our model in Section 4.3.

Just as the basic datum in Aumann's (1999a) model of knowledge is a correspondence $\mathcal{K}_i : \Omega \to \wp(\Omega)$, the basic datum in our model of awareness is a correspondence $\mathcal{A}_i : \Omega \to \wp(\Omega)$. First, it is important to recall what makes modeling knowledge with a correspondence possible. Suppose we begin with a knowledge operator $\mathbf{K}_i : \wp(\Omega) \to \wp(\Omega)$; in fact, it is convenient to start with the dual operator $\widehat{\mathbf{K}}_i(E) : \wp(\Omega) \to \wp(\Omega)$ such that $\widehat{\mathbf{K}}(E) = \neg\mathbf{K}_i\neg(E)$, so $\omega \in \widehat{\mathbf{K}}_i(E)$ means that E is consistent with i's knowledge in ω. The key idea of Aumann's model, like those of Hintikka (1962) and Kripke (1963), is to reduce $\widehat{\mathbf{K}}_i$ to its behavior on singleton events $\{\nu\}$, in the sense that we want:

$$\omega \in \widehat{\mathbf{K}}(E) \text{ if and only if for some } \nu \in E : \omega \in \widehat{\mathbf{K}}(\{\nu\}). \tag{4}$$

If this holds, and only if this holds, we can represent the operator $\widehat{\mathbf{K}}(E) : \wp(\Omega) \to \wp(\Omega)$ using a simpler correspondence $\mathcal{K}_i : \Omega \to \wp(\Omega)$ defined by

$$\nu \in \mathcal{K}_i(\omega) \text{ if and only if } \omega \in \widehat{\mathbf{K}}_i(\{\nu\}), \tag{5}$$

where the representation of $\widehat{\mathbf{K}}_i$ has the form:

$$\omega \in \widehat{\mathbf{K}}(E) \text{ if and only if for some } \nu \in E : \nu \in \mathcal{K}_i(\omega). \tag{6}$$

This is analogous to reducing a probability measure on $\wp(\Omega)$ to its values on singleton events, which is always possible in the finite case and is also possible in the countably infinite case assuming the probability measure is countably additive. Similarly, the reducibility of $\widehat{\mathbf{K}}_i$ to its behavior on singleton events as in (4)—and hence the representability of $\widehat{\mathbf{K}}_i$ using a correspondence as in (6)—is equivalent to $\widehat{\mathbf{K}}_i$ being *completely additive*, in the sense that for any family of events $\{E_j\}_{j\in J} \subseteq \Omega$:

$$\widehat{\mathbf{K}}_i(\bigcup_{j\in J} E_j) = \bigcup_{j\in J} \widehat{\mathbf{K}}_i(E).$$

Our approach to awareness is analogous to Aumann's approach to knowledge: we will reduce the behavior of an awareness operator $\mathbf{A}_i : \wp(\Omega) \to \wp(\Omega)$ to its behavior on special events, but since we work with a partially ordered set rather than a set, these special events will be *principal downsets* $\downarrow\omega$ rather than singletons $\{\omega\}$. Thus, our awareness correspondence $\mathcal{A}_i : \Omega \to \wp(\Omega)$ will be such that

$$\nu \in \mathcal{A}_i(\omega) \text{ if and only if } \omega \in \mathbf{A}_i(\downarrow\nu). \tag{7}$$

Unlike Aumann's representation of knowledge, however, the representation of $\mathbf{A}_i$ using $\mathcal{A}_i$ will not have a form like (6) above. It will have a different form, given in Definition 4.1.3, that reflects the distinction between an event being consistent with an agent's knowledge and the agent being aware of the event.

4.1 Awareness

As discussed above, to model awareness we equip possibility frames as in Definition 3.5 with awareness correspondences $\mathcal{A}_i$ for each agent. When $\nu \in \mathcal{A}_i(\omega)$, we say that in possibility ω, agent i is aware of possibility ν. In Section 4.3, we will discuss a possible decision-theoretic interpretation of this relation.

Definition 4.1. A *possibility frame with awareness* is a tuple $\mathscr{F} = (\Omega, \sqsubseteq, \mathcal{E}, \{\mathcal{A}_i\}_{i\in I})$ such that:

1. $(\Omega, \sqsubseteq, \mathcal{E})$ is a quasi-principal possibility frame with a maximum element m in the poset $(\Omega, \sqsubseteq)$;
2. $\mathcal{A}_i : \Omega \to \wp(\Omega)$ is a correspondence satisfying the following conditions for all $\omega, \omega', \nu \in \Omega$:
   (a) awareness nonvacuity: $m \in \mathcal{A}_i(\omega)$;
   (b) awareness expressibility: if $\nu \in \mathcal{A}_i(\omega)$, then $\downarrow\nu \in \mathcal{E}$;
   (c) awareness persistence: if $\omega' \sqsubseteq \omega$, then $\mathcal{A}_i(\omega) \subseteq \mathcal{A}_i(\omega')$;
   (d) awareness refinability: if $\nu \notin \mathcal{A}_i(\omega)$, then $\exists \omega' \sqsubseteq \omega \; \forall \omega'' \sqsubseteq \omega' \; \nu \notin \mathcal{A}_i(\omega'')$;
   (e) awareness joinability: if $\nu \in \mathcal{A}_i(\omega)$, $E, E' \in \mathcal{E}$, and $\max(E \cap \downarrow\nu) \cup \max(E' \cap \downarrow\nu) \subseteq \mathcal{A}_i(\omega)$, then $\max((E \sqcup E') \cap \downarrow\nu) \subseteq \mathcal{A}_i(\omega)$.
3. $\mathcal{E}$ is closed under the operations $E \mapsto \mathbf{A}_i(E)$ for $i \in I$ defined by
   - $\omega \in \mathbf{A}_i(E)$ if and only if $\forall \omega' \sqsubseteq \omega \; \forall \nu \in \mathcal{A}_i(\omega') \; \max(E \cap \downarrow\nu) \cup \max(\neg E \cap \downarrow\nu) \subseteq \mathcal{A}_i(\omega')$.

Finally, we call $\mathscr{F}$ *standard* if for all $\omega, \nu \in \Omega$, $\nu \in \mathcal{A}_i(w)$ implies $\omega \in \mathbf{A}_i(\downarrow\nu)$.

The interpretations of the conditions on $\mathcal{A}_i$ are as follows. Awareness nonvacuity says that each agent i is aware of at least the coarsest possibility of all. Awareness expressibility says that if i is aware of a possibility, then the principal downset generated by that possibility is a genuine event, eligible to be thought about.[20] Awareness persistence says that if ω settles that i is aware of a possibility ν, and ω' refines ω, then ω' still settles that i is aware of ν. Awareness refinability says that if ω does not settle that i is aware of ν, then there is a refinement ω' of ω that settles that i is definitely *not* aware of ν, so no refinement ω'' of ω' settles that i is aware of ν. Finally, awareness joinability says that if i is aware of ν and of the coarsest refinements of ν belonging to the event E, and similarly for E', then i must be aware of the coarsest refinements of ν belonging to the event E *or* E'. There is a convenient equivalent condition if Ω is finite, which quantifies over possibilities rather than arbitrary events: if i is aware of ν and some refinements $\nu_1, \dots, \nu_n$ of ν, then i must be aware of the coarsest refinements of ν belonging to the event that *one of the ν_i's obtains.*

Lemma 4.2. Suppose $\mathscr{F} = (\Omega, \sqsubseteq, \mathcal{E}, \{\mathcal{A}_i\}_{i\in I})$ satisfies part 1 of Definition 4.1 and awareness expressibility. If (i) $\mathscr{F}$ satisfies awareness joinability, then (ii) for all $\nu \in \mathcal{A}_i(\omega)$ and $\nu_1, \dots, \nu_n \in \mathcal{A}_i(\omega) \cap \downarrow\nu$, we have $\max((\downarrow\nu_1 \sqcup \dots \sqcup \downarrow\nu_n) \cap \downarrow\nu) \subseteq \mathcal{A}_i(\omega)$. Conversely, if $\max(E)$ is finite for each $E \in \mathcal{E}$, then (ii) implies (i).

Particular frames used in applications may of course satisfy additional conditions on $\mathcal{A}_i$ (cf. Remark 4.11).

The definition of the $\mathbf{A}_i$ operation in part 3 formalizes the account of awareness of events sketched in Section 1, namely that in a possibility ω, an agent i is aware of an event E if the following condition holds at ω and its refinements: if i is aware of a possibility ν, then i is aware of any coarsest refinement of ν belonging to E and any coarsest refinement of ν belonging to $\neg E$, where $\neg$ is the negation in $\mathcal{RO}(\Omega, \sqsubseteq)$ from (1).

[20] Recall from Section 3 that in a separative poset, every principal downset is a regular open set.

Remark 4.3. This definition of awareness of events—or equivalently binary questions—generalizes to a definition of awareness of arbitrary partitional questions, such as "In what month was Ann born?" Given a family $\{E_1, \dots, E_n\}$ of disjoint events such that $E_1 \sqcup \dots \sqcup E_n = \Omega$, we say that in ω, agent i is aware of $\{E_1, \dots, E_n\}$ if the following condition holds at ω and its refinements: for $1 \leq k \leq n$, if i is aware of a possibility ν, then i is aware of any coarsest refinement of ν belonging to E_k.

As for the requirement in part 3 that $\mathcal{E}$ be closed under $\mathbf{A}_i$, note this requires that $\mathbf{A}_i(E)$ is a regular open set, which is indeed the case. The quantification over $\omega' \sqsubseteq \omega$ in the definition of $\mathbf{A}_i$ guarantees persistence for $\mathbf{A}_i(E)$ (recall Lemma 3.1), while awareness persistence and refinability guarantee refinability for $\mathbf{A}_i(E)$.

Lemma 4.4. Let $(\Omega, \sqsubseteq, \mathcal{E})$ be a possibility frame and $\mathcal{A}_i : \Omega \to \wp(\Omega)$ satisfy awareness persistence and awareness refinability. Then for any $E \in \mathcal{RO}(\Omega, \sqsubseteq)$, we have $\mathbf{A}_i(E) \in \mathcal{RO}(\Omega, \sqsubseteq)$.

Then defining *unawareness* by $\mathbf{U}_i(E) = \neg\mathbf{A}_i(E)$, we have $\mathbf{U}_i(E) \in \mathcal{RO}(\Omega, \sqsubseteq)$ as well.

Finally, the condition that $\mathscr{F}$ is standard is simply the condition that awareness of a possibility ν implies awareness of the distinction between $\downarrow\nu$ and $\neg\downarrow\nu$, which reduces to the condition that if i is aware of ν and ν', then i is also aware of any coarsest refinements of ν' that are incompatible with ν: if $\nu, \nu' \in \mathcal{A}_i(\omega)$, then $\max(\{\nu^* \in \downarrow\nu' \mid \nu^* \perp \nu\}) \subseteq \mathcal{A}_i(\omega)$. Standardness implies the equivalence given in (7) above.

Lemma 4.5. For any standard possibility frame with awareness $\mathscr{F} = (\Omega, \sqsubseteq, \mathcal{E}, \{\mathcal{A}_i\}_{i \in I})$ and $\omega, \nu \in \Omega$, we have $\nu \in \mathcal{A}_i(\omega)$ if and only if $\omega \in \mathbf{A}_i(\downarrow\nu)$.

Proof. The left-to-right direction is just the definition of standardness. For the right-to-left direction, if $\omega \in \mathbf{A}_i(\downarrow\nu)$, then since $m \in \mathcal{A}_i(\omega)$ by awareness nonvacuity, it follows from the definition of $\mathbf{A}_i$ that in ω, i is aware of the coarsest refinement of m belonging to $\downarrow\nu$, which is ν itself, so $\nu \in \mathcal{A}_i(\omega)$. $\square$

We now give our first two examples of using possibility frames with awareness for modeling. For simplicity, we concentrate throughout on the awareness of a single agent; one can enrich each of our examples to a multi-agent example by adding additional possibilities representing differences in other agents' awareness.

Example 4.6. We begin with perhaps the simplest example of an interesting possibility frame with awareness, formalizing a story discussed in Geanakoplos 1989, Modica and Rustichini 1994, and Modica and Rustichini 1999. The story concerns Sherlock Holmes's assistant, Watson: if Watson hears a dog bark, then he will know—and hence be aware of—the event of the dog barking; but if he does not hear the dog bark, then he will not even be aware of the distinction between the dog barking vs. not barking. We formalize this using the frame in Figure 4.6: $\Omega = \{m, b, \overline{b}\}$; the partial order $\sqsubseteq$ is depicted by the arrows, so, e.g., we have $b \sqsubseteq m$ (arrows point toward more refined possibilities); the awareness correspondence for Watson, whom we label i, is given by $\mathcal{A}_i(m) = \{m\}$, $\mathcal{A}_i(b) = \Omega$, and $\mathcal{A}_i(\overline{b}) = \{m\}$; and $\mathcal{E} = \mathcal{RO}(\Omega, \sqsubseteq)$. It is easy to check that this is a standard possibility frame with awareness.

Let $Barks = \{b\}$; this set satisfies persistence and refinability, so by Lemma 3.1 it is an event in $\mathcal{RO}(\Omega, \sqsubseteq)$. Now it is immediate from the definition of $\mathbf{A}_i$ that when i is aware of all possibilities, i is also aware of all events. Hence in possibility b, i is aware of $Barks$ and $\neg Barks$. By contrast, in $\overline{b}$, i is unaware of these events, for the following reason: although in $\overline{b}$, i is aware of the coarsest possibility m, we have $m \notin Barks$ and $m \notin \neg Barks$ (the latter because $b \sqsubseteq m$ and $b \in Barks$); then since in $\overline{b}$, i is *only* aware of m, i is not aware of the coarsest refinement of m belonging to $Barks$ (namely, b) or of the coarsest refinement of m belonging to $\neg Barks$ (namely, $\overline{b}$). In short, in $\overline{b}$, i is not aware of the distinction $Barks$ vs. $\neg Barks$.

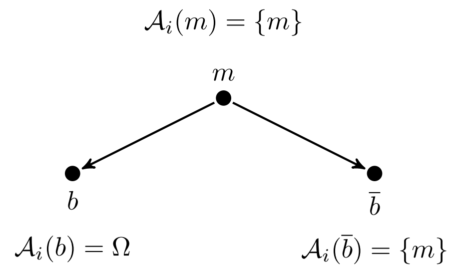

Figure 1: A possibility frame with awareness representing Watson in the story of Geanakoplos 1989. Refinement arrows implied by reflexivity are not drawn.

Moreover, in $\overline{b}$, i is unaware of his unawareness of $Barks$. The reason is similar to the above: although in $\overline{b}$, i is aware of m, we have $m \notin \mathbf{A}_i Barks = \{b\}$ and $m \notin \neg\mathbf{A}_i Barks = \{\overline{b}\}$ (the latter because $b \sqsubseteq m$ and $b \in \mathbf{A}_i Barks$); then since in $\overline{b}$, i is *only* aware of m, i is not aware of the coarsest refinement of m belonging to $\mathbf{A}_i Barks$ (namely, b) or of the coarsest refinement of m belonging to $\neg\mathbf{A}_i Barks$ (namely, $\overline{b}$). In short, in $\overline{b}$, i is not aware of the distinction $\mathbf{A}_i Barks$ vs. $\neg\mathbf{A}_i Barks$.

Finally, we must continue, as stressed in Section 1, to avoid conflating events with sentences. For example, since we can write $\Omega = Barks \sqcup \neg Barks$, where $\sqcup$ is the join in the Boolean algebra $\mathcal{RO}(\Omega, \sqsubseteq)$, does it follow from Watson's being aware of Ω that he is aware of $Barks$? As we have seen, it does not. Ω is the trivial event, not to be confused with the linguistic item '$Barks \sqcup \neg Barks$' of the analyst's language. Surely if Watson were aware of a sentence in a language that embeds '$Barks$', then he would be aware of '$Barks$'. But Watson's awareness of the trivial event Ω does not imply any such awareness of a sentence.

Example 4.7. Consider a game in which a column player is aware that the row player can move up or down but is unaware that the row player has a third move, middle. Informally, such a situation is represented by the game matrix at the top of Figure 2 in which the middle row is greyed out. Formally, we can represent the uawareness of the column player, whom we call i, using the frame in Figure 2; that this is a standard possibility frame with awareness can be checked by hand or more quickly with the notebook cited in Section 1.1. There are two games the players could play: game G in which the row player only has two moves, represented by the subtree with root g, and game $\underline{G}$ in which the row player has three moves, represented by the subtree with root $\underline{g}$. In each colored state, i is aware only of the red possibilities; note i's awareness of $\underline{g}$ is in effect just awareness of the possibility of *not playing* G, without any awareness of further refinements of that possibility. But before computing i's awareness of events, one should become comfortable with the treatment of 'not' and 'or' coming from Theorem 3.3. For example, although the partial possibility $\underline{\ell}$ does not belong to the event $\underline{Middle} = {\downarrow}\{\underline{\ell m}, \underline{rm}\}$ of the row player playing middle in $\underline{G}$, we have $\ell \notin \neg\underline{Middle}$, since $\underline{\ell}$ is refined by $\underline{\ell m}$ and $\underline{\ell m} \in \underline{Middle}$. Also note that where $\underline{Up} = {\downarrow}\{\underline{lu}, \underline{ru}\}$ and $\underline{Down} = {\downarrow}\{\underline{ld}, \underline{rd}\}$, we have $\underline{\ell} \in \underline{Up} \sqcup \underline{Middle} \sqcup \underline{Down}$, despite the fact that $\underline{\ell}$ does not belong to the union of these events; this is because every proper refinement of $\underline{\ell}$ does belong to the union. Thus, when dealing with partial possibilities, one must resist the temptation to interpret 'not' and 'or' using set-theoretic complement and union.

Turning to awareness, observe that at each of the colored states ω, it is not settled that i is aware of $\underline{Middle}$: $\omega \notin \mathbf{A}_i\underline{Middle}$. For in the colored states, i is aware of $\underline{g}$, $\underline{g} \notin \underline{Middle}$, and $\underline{g} \notin \neg\underline{Middle}$ (since $\underline{g}$ is refined by $\underline{\ell m}$, which belongs to $\underline{Middle}$), yet i is not aware of any coarsest refinement of $\underline{g}$ that belongs to $\underline{Middle}$, namely $\underline{\ell m}$ or $\underline{rm}$. Hence each colored *leaf* ω of the tree settles that i is *not* aware of $\underline{Middle}$: $\omega \in \neg\mathbf{A}_i\underline{Middle}$. By contrast, in the black leaves, i is aware of every event, in virtue of being aware of every

	ℓ	r
u	$3,3$	$0,4$
m	10,10	10,0
d	$4,0$	$1,1$

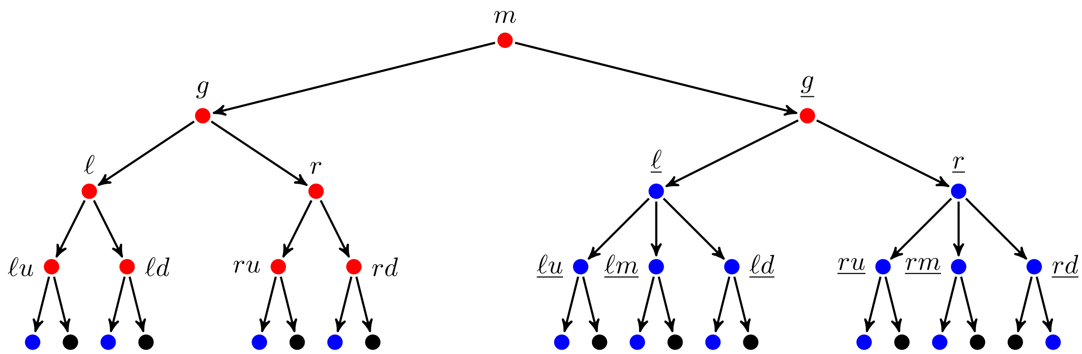

for each colored state ω, $\mathcal{A}_i(\omega) = \{\nu \in \Omega \mid \nu \text{ a red state}\}$

for each black state ω, $\mathcal{A}_i(\omega) = \Omega$ $\qquad\qquad \mathcal{E} = \mathcal{RO}(\Omega, \sqsubseteq)$

Figure 2: A possibility frame with awareness representing a column player's unawareness that the row player has an extra move m. Refinement arrows implied by reflexivity or transitivity are not drawn.

possibility. It follows that at each colored leaf ω of the tree, i is unaware of her unawareness of $\underline{Middle}$: $\omega \in \neg\mathbf{A}_i\neg\mathbf{A}_i\underline{Middle}$. For in the colored states, i is aware of possibilities that are refined by colored leaves and black leaves, yet i is unaware of all leaves. Thus, at each colored leaf, i is not only unaware of the distinction $\underline{Middle}$ vs. $\neg\underline{Middle}$ but also of the distinction $\mathbf{A}_i\underline{Middle}$ vs. $\neg\mathbf{A}_i\underline{Middle}$. Yet at all colored states, i is aware of the events $Up = \downarrow\{\ell u, ru\}$ and $Down = \downarrow\{\ell d, rd\}$ of playing up and down in G, since for every red state ω, every coarsest refinement of ω in Up, $\neg Up$, $Down$, and $\neg Down$ is itself a red state.

We close this section with a simple example of an infinite possibility frame with awareness.

Example 4.8. Let $2^{\leq\omega}$ be the set of all finite or countably infinite binary strings (such as 0110, etc.) ordered such that $\sigma \sqsubseteq \tau$ if τ is an initial segment of σ (so $0110 \sqsubseteq 011$, etc.). This could represent all possibilities for finitely many or countably infinitely many flips of a coin. Now consider the possibility frame $(2^{\leq\omega}, \sqsubseteq, \mathcal{RO}(2^{\leq\omega}, \sqsubseteq))$ equipped with the awareness correspondence $\mathcal{A}_i$ defined by

$$\mathcal{A}_i(\sigma) = \{\tau \in 2^{\leq\omega} \mid \tau \text{ a finite binary string}\}.$$

Then $\mathcal{A}_i$ satisfies awareness nonvacuity (as the empty string is the maximal element of $(2^{\leq\omega}, \sqsubseteq)$), expressibility (since $\mathcal{RO}(2^{\leq\omega}, \sqsubseteq)$ contains all principal downsets $\downarrow\tau$), persistence and refinability (since $\mathcal{A}_i(\sigma) = \mathcal{A}_i(\tau)$ for all $\sigma, \tau \in 2^{\leq\omega}$), and joinability (since if $\max(E) \cap \nu$ and $\max(E') \cap \nu$ contain only finite strings, then so does $\max((E \sqcup E') \cap \downarrow\nu)$). This frame represents an agent who can conceive of any finite sequence of coin flips but cannot conceive of infinite sequences. It is then easy to see that the set of events of which the agent is aware forms an *atomless* (recall Footnote 15) Boolean subalgebra of the atomic Boolean algebra $\mathcal{RO}(2^{\leq\omega}, \sqsubseteq)$.

4.2 Knowledge, belief, and awareness

We now add knowledge and belief correspondences to our possibility frames with awareness. The basic distinction is that what the agent knows depends on the true information she has received, whereas belief is subjective in the same sense as in subjective probability, which models belief quantitatively.[21] Take $\nu \in \mathcal{K}_i(\omega)$ (resp. $\nu \in \mathcal{B}_i(\omega)$) to mean that every event that i knows (resp. believes) in ω—or would know (resp. believe) if made aware of the event—holds true in ν. In this sense ν *conforms to what i knows* (resp. believes).

Definition 4.9. An *epistemic possibility frame* is a tuple $\mathscr{F} = (\Omega, \sqsubseteq, \mathcal{E}, \{\mathcal{A}_i\}_{i\in I}, \{\mathcal{K}_i\}_{i\in I}, \{\mathcal{B}_i\}_{i\in I})$ such that:

1. $(\Omega, \sqsubseteq, \mathcal{E}, \{\mathcal{A}_i\}_{i\in I})$ is a possibility frame with awareness;
2. each $\mathcal{R}_i \in \{\mathcal{K}_i, \mathcal{B}_i\}_{i\in I}$ is a correspondence $\mathcal{R}_i : \Omega \to \wp(\Omega)$ satisfying the following for all $\omega, \omega', \nu \in \Omega$:
   (a) $\mathcal{R}_i$-monotonicity: if $\omega' \sqsubseteq \omega$, then $\mathcal{R}_i(\omega') \subseteq \mathcal{R}_i(\omega)$;
   (b) $\mathcal{R}_i$-regularity: $\mathcal{R}_i(\omega) \in \mathcal{RO}(\Omega, \sqsubseteq)$;
   (c) $\mathcal{R}_i$-refinability: if $\nu \in \mathcal{R}_i(\omega)$, then $\exists \omega' \sqsubseteq \omega\ \forall \omega'' \sqsubseteq \omega'\ \exists \nu' \sqsubseteq \nu$: $\nu' \in \mathcal{R}_i(\omega'')$;
   (d) epistemic factivity: $\omega \in \mathcal{K}_i(\omega)$;
   (e) doxastic consistency: $\mathcal{B}_i(\omega) \neq \varnothing$;
   (f) doxastic inclusion: $\mathcal{B}_i(\omega) \subseteq \mathcal{K}_i(\omega)$.
3. for each $\mathcal{R}_i \in \{\mathcal{K}_i, \mathcal{B}_i\}_{i\in I}$ and $E \in \mathcal{E}$, we have $\{\omega \in \Omega \mid \mathcal{R}_i(\omega) \subseteq E\} \in \mathcal{E}$.

We call $\mathscr{F}$ *standard* if the underlying possibility frame with awareness is standard.

The interpretations of the first three conditions are as follows for knowledge; the belief interpretations are analogous. $\mathcal{R}_i$-monotonicity—in its equivalent formulation: if $\omega' \sqsubseteq \omega$, then $\nu \notin \mathcal{R}_i(\omega)$ implies $\nu \notin \mathcal{R}_i(\omega')$—says that if ν does not conform to i's knowledge in ω, then ν does not conform to i's knowledge in any refinement ω' of ω, since i retains in ω' whatever knowledge she had in ω. $\mathcal{R}_i$-regularity says that we can view the set of possibilities that conform to i's knowledge as a genuine event that i implicitly knows (in fact, the strongest such event, given the definition of implicit knowledge below). Finally, $\mathcal{R}_i$-refinability says that if ν conforms to i's knowledge in ω, then there is a refinement ω' of ω that settles that *some refinement of ν conforms to i's knowledge*, in the sense that at every refinement ω'' of ω', some refinement of ν conforms to i's knowledge in ω''.[22] Together these conditions imply the following closure property of $\mathcal{RO}(\Omega, \sqsubseteq)$, which shows that it is possible to satisfy the closure property in part 3 of Definition 4.9.

Lemma 4.10. Let $(\Omega, \sqsubseteq, \mathcal{E})$ be a possibility frame and $\mathcal{R}_i : \Omega \to \wp(\Omega)$ satisfy $\mathcal{R}_i$-monotonicity, $\mathcal{R}_i$-regularity, and $\mathcal{R}_i$-refinability. Then for any $E \in \mathcal{RO}(\Omega, \sqsubseteq)$, we have $\{\omega \in \Omega \mid \mathcal{R}_i(\omega) \subseteq E\} \in \mathcal{RO}(\Omega, \sqsubseteq)$.

As in Fagin and Halpern 1988, one may take $\mathbf{L}_i(E) = \{\omega \in \Omega \mid \mathcal{K}_i(\omega) \subseteq E\}$ (resp. $\{\omega \in \Omega \mid \mathcal{B}_i(\omega) \subseteq E\}$) to be the event of i *implicitly knowing* (resp. *implicitly believing*) E in the sense that i would know (resp. believe) E if i were aware of E. However, we will concentrate here on explicit knowledge $\mathbf{K}_i$ and belief $\mathbf{B}_i$:

$$\begin{aligned} \mathbf{K}_i(E) &= \{\omega \in \Omega \mid \mathcal{K}_i(\omega) \subseteq E \text{ and } \omega \in \mathbf{A}_i(E)\}; \\ \mathbf{B}_i(E) &= \{\omega \in \Omega \mid \mathcal{B}_i(\omega) \subseteq E \text{ and } \omega \in \mathbf{A}_i(E)\}. \end{aligned}$$

[21] As sketched in Section 6, we could use p-belief instead of belief, but for simplicity we use belief when introducing our model.

[22] This condition from Holliday 2015 is weaker than the refinability condition in Humberstone 1981. This weakening is useful for the representation theory of modal algebras by possibility frames (see Remark 2.39 of Holliday 2015).

By Lemma 4.10 and the closure of $\mathcal{E}$ under binary intersection, $\mathcal{E}$ is also closed under $\mathbf{K}_i$ and $\mathbf{B}_i$. If $\mathcal{E}$ is closed under countably infinite intersections (e.g., if $\mathcal{E} = \mathcal{RO}(\Omega, \sqsubseteq)$), then it is closed under the usual operations of *common knowledge* (Aumann 1999a, § 2) and *common belief* (see, e.g., Lismont and Mongin 1994) for $\Box_i \in \{\mathbf{K}_i, \mathbf{B}_i\}$:

$$\Box^1(E) = \bigcap_{i \in I} \Box_i(E) \qquad \Box^{n+1}(E) = \Box^1(\Box^n(E)) \qquad \Box^\infty(E) = \bigcap_{n \in \mathbb{N}_{>0}} \Box^n(E).$$

Conditions (d)–(f) of Definition 4.9 are the bare minimum constraints for knowledge and belief, familiar from the earliest formal models of knowledge and belief (Hintikka 1962). Epistemic factivity implies that if i knows E, then E is true; doxastic consistency implies that an agent cannot believe $\varnothing$; and doxastic inclusion (which implies that if $\mathcal{K}_i(\omega) \subseteq E$, then $\mathcal{B}_i(\omega) \subseteq E$) implies that if i knows E, then i believes E.

Remark 4.11. Particular frames used in applications may of course satisfy additional constraints, which may imply epistemic and doxastic introspection principles (see Ding et al. 2019 for a study of when such principles matter for multi-agent reasoning). As usual, introspection principles for $\mathbf{A}_i$, $\mathbf{K}_i$, and $\mathbf{B}_i$ that quantify over events—e.g., for all events E, $\mathbf{K}_i(E) \subseteq \mathbf{K}_i(\mathbf{K}_i(E))$—immediately correspond to conditions on $\mathcal{A}_i$, $\mathcal{K}_i$, and $\mathcal{B}_i$ that quantify over events, just by unpacking the definitions of $\mathbf{A}_i$, $\mathbf{K}_i$, and $\mathbf{B}_i$. There is then a technical question, studied in a branch of modal logic known as correspondence theory (van Benthem 2001), about whether the conditions on $\mathcal{A}_i$, $\mathcal{K}_i$, and $\mathcal{B}_i$ that quantify over events—called second-order conditions—are equivalent to conditions on $\mathcal{A}_i$, $\mathcal{K}_i$, and $\mathcal{B}_i$ that only quantify over possibilities in Ω—called first-order conditions. Correspondence theory for implicit knowledge and belief in possibility semantics is well understood (Holliday 2015, Yamamoto 2017, Zhao 2021). To take one example, in epistemic possibility frames, Positive Introspection for implicit knowledge—for all $E \in \mathcal{RO}(\Omega, \sqsubseteq)$, $\mathbf{L}_i(E) \subseteq \mathbf{L}_i(\mathbf{L}_i(E))$—holds if and only if $\mathcal{K}_i$ is transitive (if $\omega' \in \mathcal{K}_i(\omega)$ and $\omega'' \in \mathcal{K}_i(\omega')$, then $\omega'' \in \mathcal{K}_i(\omega)$), just as in Kripke frames. However, we will not delve into correspondence theory for awareness or explicit knowledge and belief here, as our focus is instead on representation theory in Section 5. The example frames in this paper satisfy first-order conditions that are sufficient but not necessary for introspection principles. For example, it is sufficient for $\mathbf{A}_i(E) \subseteq \mathbf{A}_i(\mathbf{A}_i(E))$ to hold for all $E \in \mathcal{E}$ that for all $\omega \in \Omega$, if $\mathcal{A}_i(\omega) \neq \Omega$, then for all $\nu \in \mathcal{A}_i(\omega)$ and $\nu' \sqsubseteq \nu$, if $\mathcal{A}_i(\nu') \neq \Omega$, then for some $\omega' \sqsubseteq \omega$, we have $\mathcal{A}_i(\nu') = \mathcal{A}_i(\omega')$; and if this also holds with 'if $\mathcal{A}_i(\omega) \neq \Omega$, then for all $\nu \in \mathcal{A}_i(\omega)$' replaced by 'for all $\nu \in \mathcal{K}_i(\omega)$', then $\mathbf{A}_i(E) \subseteq \mathbf{K}_i(\mathbf{A}_i(E))$ for all $E \in \mathcal{E}$.

We now present two examples of using epistemic possibility frames for modeling.

Example 4.12. Let us add knowledge and belief to Example 4.6. There are two pairs of knowledge and belief correspondences we might consider for Watson:

$$\mathcal{K}_i(b) = \mathcal{B}_i(b) = \{b\} \text{ and } \mathcal{K}_i(\bar{b}) = \mathcal{B}_i(\bar{b}) = \mathcal{K}_i(m) = \mathcal{B}_i(m) = \Omega;$$

$$\mathcal{K}'_i(b) = \mathcal{B}'_i(b) = \{b\} \text{ and } \mathcal{K}'_i(\bar{b}) = \mathcal{B}'_i(\bar{b}) = \{\bar{b}\} \text{ and } \mathcal{K}'_i(m) = \mathcal{B}'_i(m) = \Omega.$$

One can easily check that in both cases, all the conditions of Definition 4.9 are satisfied. Moreover, for any event E, we have $\mathbf{K}_i(E) = \mathbf{K}'_i(E)$ and $\mathbf{B}_i(E) = \mathbf{B}'_i(E)$. However, the primed pair of correspondences can be used to capture the idea that *if only Watson were aware in $\bar{b}$ of the distinction between Bark and $\neg$Bark*, then he would know and believe $\neg Bark$ in $\bar{b}$. In either case, the frame illustrates our reason for rejecting the Plausibility axiom of Dekel et al. 1998 discussed in Section 2. First observe that the event $\mathbf{U}_i(Bark)$ is unknowable. It cannot be known at b or m, because it is not true at these states (recall Example 4.6). It also cannot be known at $\overline{b}$, because although it is true at $\overline{b}$, Watson is not aware of $\mathbf{U}(Bark)$ in $\overline{b}$ (again recall Example 4.6). Thus, $\neg\mathbf{K}_i(\mathbf{U}_i(Bark)) = \Omega$. But of course $\mathbf{K}_i(\Omega) = \Omega$. So we have a violation of Plausibility at $\overline{b}$, as $\overline{b} \in \mathbf{U}_i(Bark)$ and $\overline{b} \notin \neg\mathbf{K}\neg\mathbf{K}_i(\mathbf{U}_i(Bark))$. But we have already explained in Section 2 the fallacy, based on conflating awareness of events with awareness of sentences, in thinking that if i is aware of Ω, and Ω can be obtained from E by applying some operations on sets, then i must be aware of E. It is similar to the fallacy in thinking that if a student is aware of a number n, and n can be obtained from a number m by some mathematical operations, then the student must also be aware of m.

Finally, let us return to the example of the *overconfident* agent with which we began in Section 1.

Example 4.13. A potential investor i in a firm believes that he knows the firm is profitable, while being unaware of a sophisticated type of fraud that the firm is in fact using to cover up unprofitability. This is the case in state r_1 in the possibility frame in Figure 4.13 with $\mathcal{E} = \mathcal{RO}(\Omega, \sqsubseteq)$ and the following awareness, belief, and knowledge correspondences:

- for each colored state ω, $\mathcal{A}_i(\omega) = \{\nu \in \Omega \mid \nu \text{ a red state}\}$;
- for each black or gray state ω, $\mathcal{A}_i(\omega) = \Omega$;
- for each square or blue state ω, $\mathcal{B}_i(\omega) = {\downarrow}pb$ and $\mathcal{K}_i(\omega) = {\downarrow}\{\omega, pb\}$;
- for each black state ω, $\mathcal{B}_i(\omega) = \{pb\overline{u}\}$ and $\mathcal{K}_i(\omega) = \{\omega, pb\overline{u}\}$;
- for each diamond or green state ω, $\mathcal{B}_i(\omega) = \mathcal{K}_i(\omega) = {\downarrow}\{\nu \in \Omega \mid \nu \text{ a diamond state}\}$;
- for each gray state ω, $\mathcal{B}_i(\omega) = \mathcal{K}_i(\omega) = \{\nu \in \Omega \mid \nu \text{ a gray state}\}$;
- $\mathcal{B}_i(p) = \rho(\bigcup_{\omega \sqsubset p} \mathcal{B}_i(\omega)) = {\downarrow}\{p, \overline{p}\overline{b}\}$ and $\mathcal{K}_i(p) = \rho(\bigcup_{\omega \sqsubset p} \mathcal{K}_i(\omega)) = {\downarrow}\{p, \overline{p}\overline{b}\}$;
- $\mathcal{B}_i(\overline{p}) = \rho(\bigcup_{\omega \sqsubset \overline{p}} \mathcal{B}_i(\omega)) = {\downarrow}\{p, \overline{p}\overline{b}\}$ and $\mathcal{K}_i(\overline{p}) = \rho(\bigcup_{\omega \sqsubset \overline{p}} \mathcal{K}_i(\omega)) = \Omega$;
- $\mathcal{B}_i(m) = {\downarrow}\{p, \overline{p}\overline{b}\}$ and $\mathcal{K}_i(m) = \Omega$.

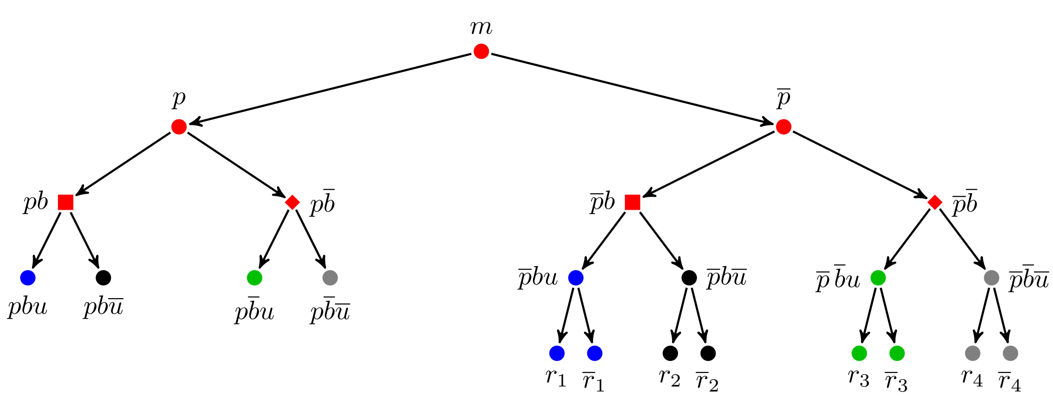

Figure 3: The refinement structure of a possibility frame for Example 4.13. Refinement arrows implied by reflexivity or transitivity are not drawn.

Let $Profit = {\downarrow}p$ and $Fraud = \{r_1, r_2, r_3, r_4\}$. Then we have the following:

- In the blue states, i is unaware of the distinction between $Fraud$ and $\neg Fraud$ and believes $Profit$ and believes that he knows $Profit$ (note that in state pb, i knows $Profit$).
- In the green states, i is unaware of the distinction between $Fraud$ and $\neg Fraud$ and is undecided (in belief) and uncertain (in knowledge) about $Profit$.
- In the black states, i is aware of the distinction between $Fraud$ and $\neg Fraud$ but still believes $Profit$ and believes he knows $Profit$ (note that in state $pb\overline{u}$, i knows $Profit$).
- In the gray states, i is aware of the distinction between $Fraud$ and $\neg Fraud$ and is undecided and uncertain about $Profit$.[23]

In r_1, $\overline{r}_1$, r_2, and $\overline{r}_2$, agent i is overconfident in $Profit$: i believes he knows $Profit$, but he does not, since $Profit$ does not obtain in these states. As analysts—or as other agents interacting with i—we might assign low probability to r_2 and $\overline{r}_2$, reflecting the view that i is unlikely to be overconfident when i is aware of the possibility of the sophisticated type of fraud. In general, say that a state ω determines that agent i has *information-based beliefs*[24] if for all $\omega' \sqsubseteq \omega$ and $\nu \in \Omega$, if in ω', i is aware of ν and i's information is in fact consistent with ν, then in ω', i does not mistakenly believe he can rule out ν:

$$\text{if } \nu \in \mathcal{A}_i(\omega') \text{ and } \nu \in \mathcal{K}_i(\omega'), \text{ then } \nu \in \mathcal{B}_i(\omega').$$

If ω determines that i has information-based beliefs, then ω determines that i can only have false beliefs if he is unaware of some possibilities that are in fact consistent with his information. In the above example, i has information-based beliefs at all leaves of the tree except for r_2 and $\overline{r}_2$. In principle, to what extent false beliefs in a population of agents are correlated with unawareness of possibilities consistent with their information, as distinguished from mistakes in reasoning, exposure to misleading evidence, etc., could be investigated experimentally, though we will not attempt to describe an experimental protocol here.

To sum up, in contrast to models that define awareness in terms of knowledge (recall Section 1), as in Modica and Rustichini 1994, 1999, Heifetz et al. 2006, and Li 2009, we are able to model an agent who is aware of $Profit$ while also being overconfident in $Profit$; and we are able to relate the agent's overconfidence in $Profit$ to his unawareness of $Fraud$.

Remark 4.14. Examples 4.12 and 4.13 satisfy all the principles of Stalnaker's (2006, p. 179) joint logic of knowledge and belief, with negative introspection for belief suitably modified to allow for unawareness: $\mathbf{K}_i(E) \subseteq \mathbf{K}_i(\mathbf{K}_i(E))$ and $\mathbf{B}_i(E) \subseteq \mathbf{K}_i(\mathbf{B}_i(E))$ (Positive Introspection); $\neg\mathbf{B}_i(E) \cap \mathbf{A}_i\neg\mathbf{B}_i(E) \subseteq \mathbf{K}_i\neg\mathbf{B}_i(E)$ (Weak Negative Introspection for Belief); and $\mathbf{B}_i(E) \subseteq \mathbf{B}_i(\mathbf{K}_i(E))$ (Strong Belief).

Finally, let us return to the impossibility theorem from Dekel et al. 1998 in Proposition 2.1. One can check[25] that the structures $(\mathcal{E}, \subseteq, \mathbf{U}_i, \mathbf{K}_i, \neg)$ arising from Examples 4.12 and 4.13 satisfy AU Introspection, KU Introspection, and the following weakening of Plausibility:

Nontrivial Plausibility: $U(E) \leq \neg K(E)$, and if $\neg K(E) \neq \Omega$, then $U(E) \leq \neg K \neg K(E)$.

Thus, a slight weakening of Plausibility leads us from impossibility to the following possibility result.

[23] Returning to Proposition 2.3, if we define $Fraud' = Fraud \cap \mathbf{A}_i(Fraud) \cap \neg\mathbf{K}_i(Fraud)$, then $Fraud' = \{r_2, r_4\}$ and $U(Fraud')$ is the set of blue and green states.

[24] Similarly, one could define *information-based p-beliefs* by replacing $\nu \in \mathcal{B}_i(\omega')$ in the displayed condition with i's assigning probability greater than $1 - p$ to the event $\downarrow\mu$, assuming an extension of our model with probability as sketched in § 6.

[25] By hand for Example 4.12 and using the Jupyter notebook cited in Section 1.1 for Example 4.13.

Fact 4.15. There are epistemic possibilities frames such that $(\mathcal{E}, \subseteq, \mathbf{U}_i, \mathbf{K}_i, \neg)$ satisfies the axioms of Dekel et al. 1998, Theorem 1(i) when Plausibility is replaced by Nontrivial Plausibility.[26]

4.3 Decision-theoretic considerations

We have shown how an awareness correspondence $\mathcal{A}_i$ on a partial-state space can be used to model an agent's awareness of events. Let us briefly consider a possible decision-theoretic interpretation of $\nu \in \mathcal{A}_i(\omega)$.

Given a poset $(\Omega, \sqsubseteq)$ and a set X of *consequences*, an *act* is a correspondence $f : \Omega \to \wp(X)$ satisfying at least the properties that

- if $\omega' \sqsubseteq \omega$, then $f(\omega') \subseteq f(\omega)$ (act persistence), and
- if $x \in f(\omega)$, then $\exists \omega' \sqsubseteq \omega\ \forall \omega'' \sqsubseteq \omega'\ x \in f(\omega'')$ (act refinability).

A natural further constraint is that if ω is a minimal element in the poset $(\Omega, \sqsubseteq)$, then $|f(\omega)| = 1$. The reason f is set-valued is that a partial possibility ω may rule out some possible consequences of an action without yet settling on a unique consequence. Act persistence and refinability guarantee that for any $x \in X$, the set $\{\omega \in \Omega \mid x \notin f(\omega)\}$ of possibilities that settle that x is not a possible consequence of the action belongs to $\mathcal{RO}(\Omega, \sqsubseteq)$, as does the set $\{\omega \in \Omega \mid f(\omega) = \{x\}\}$ of possibilities that settles that x is the only possible consequence of the action; thus, these sets are events that can be believed, assigned probabilities, etc. Now assume that ω settles that i has some strict preferences between acts, encoded by a relation $\succ_{i,\omega}$.[27] Then inspired by Schipper's (2013; 2014b) decision-theoretic characterization of awareness, we can say that in ω, agent i is aware of ν if there are acts f and g such that $f \succ_{i,\omega} g$ and either

(i) (a) for all $\nu_1 \sqsubseteq \nu$ and $\nu_2 \sqsubseteq \nu$, we have $f(\nu_1) = f(\nu_2)$ and $g(\nu_1) = g(\nu_2)$, and
   (b) for all $\nu' \not\sqsubseteq \nu$, there is a $\nu'' \sqsubseteq \nu'$ such that $f(\nu'') = g(\nu'')$,

   or

(ii) (a) for all $\nu_1 \perp \nu$ and $\nu_2 \perp \nu$, we have $f(\nu_1) = f(\nu_2)$ and $g(\nu_1) = g(\nu_2)$, and
   (b) for all $\nu' \between \nu$, there is a $\nu'' \sqsubseteq \nu'$ such that $f(\nu'') = g(\nu'')$.

That is, either (i) f and g are constant on all refinements of ν and are only determined to have different consequences in refinements of ν or (ii) f and g are constant on all possibilities incompatible with ν and are only determined to have different consequences in possibilities incompatible with ν. A strict preference between such acts *reveals* i's awareness of ν.[28] On this picture, which we leave as only a sketch, the conditions in Definition 4.1 are constraints on i's preference relations $\{\succ_{i,\omega}\}_{\omega \in \Omega}$.

[26] The frames from Examples 4.12 and 4.13 also satisfy the following, building on nomenclature of Heifetz et al. 2006, Prop. 3: if $\mathbf{K}_i(E) \neq \varnothing$, then $\mathbf{A}_i(\mathbf{K}_i(E)) = \mathbf{A}_i(E)$ (Nontrivial AK-Self Reflection); $\mathbf{A}_i(\mathbf{A}_i(E)) = \mathbf{A}_i(E)$ (AA-Self Reflection); and $\mathbf{K}_i(\mathbf{A}_i(E)) = \mathbf{A}_i(E)$ (A-Introspection) (recall Remark 4.11). However, we think it would be reasonable for at least the left-to-right inclusions in each of these principles to be violated in other examples, given the distinction between events and sentences we have stressed (e.g., from the facts that $\omega \in \mathbf{A}_i(F)$ and $F = \mathbf{A}_i(E)$, it should not be required that $\omega \in \mathbf{A}_i(E)$, since the agent need not conceive of the event F—which has no syntactic structure—in terms of anyone's awareness of E).

[27] This assignment $\omega \mapsto \succ_{i,\omega}$ should satisfy persistence and refinability conditions so that $\{\omega \in \Omega \mid f \succ_{i,\omega} g\} \in \mathcal{RO}(\Omega, \sqsubseteq)$.

[28] When attempting to infer such preferences from choice behavior, one confronts the problem that the agent may have false beliefs about the acts between which she is choosing. For example, Schipper (2013, p. 727) considers an agent faced with three contracts and argues that the agent's indifference between them reveals the agent's unawareness of a certain lawsuit that is addressed in different ways in the fine print of the contracts; but such indifference could be due to the agent being perfectly aware of the lawsuit but falsely believing that the contracts do not include any provisions related to the lawsuit. This is similar to the well-known problem that if an agent flips a coin to choose between a rotten apple and a ripe orange, this might not reveal that the agent is indifferent between rotten apples and ripe oranges, because the agent may falsely believe that the apple is ripe. The difference is that in the fruit case, we can fully inform the agent about the options, while in the contract case, doing so would make the agent aware of the very thing—the lawsuit—of which we are trying to check if the agent is aware. We leave the question of an experimental protocol for operationalizing unawareness for future work.

Example 4.16. In the context of Example 4.13, suppose f and g are functions from Ω to $\wp(X)$ such that if ω is a leaf of the tree, then (a) $|f(\omega)| = 1$, (b) if $\omega \neq r_1$, then $f(\omega) = g(\omega)$, (c) $f(r_1) \neq g(r_1)$, and (d) for any non-leaf node ν of the tree, $f(\nu) = \bigcup\{f(\omega) \mid \omega \text{ a leaf}\}$ and $g(\nu) = \bigcup\{g(\omega) \mid \omega \text{ a leaf}\}$. Then f and g satisfy act persistence and act refinability, so they are acts. Moreover, they satisfy condition (i) above, so the investor being unaware of r_1 implies that he cannot have a strict preference between f and g.

Heifetz, Meier, and Schipper (2006, § 3) argue that another consequence of unawareness for decision making is that agents with unawareness may engage in speculative trade of the kind forbidden in standard partition models of uncertainty (see Aumann 1976, Milgrom and Stokey 1982), wherein there is common knowledge that both parties are willing to trade, and each party strictly prefers to trade. However, their model of speculative trade does not satisfy a standard principle in decision theory accepted above (for simplicity, assume that singleton subsets of X are considered measurable):

- *event-measurability of acts*: for a given act f and $x \in X$, the set of states ω at which f produces the unique consequence x must be an *event* in $\mathcal{E}$.

In the model of Heifetz et al., as in ours, events must be *persistent*: if $\omega \in E \in \mathcal{E}$ and ω' refines ω, then $\omega' \in E$. Yet Heifetz et al. allow acts to have inconsistent values across refinements. In their speculative trade example, two relevant questions are whether a firm faces a lawsuit and whether it has a novel use for a product. The value of owning the firm is said to be \$100 in the coarsest state m in the state space, which settles nothing about the lawsuit or novelty; it is said to be \$90 in a refinement ω of m that settles there is a lawsuit but says nothing about the novelty; it is said to be \$110 in a refinement ω' of m that settles there is a novelty but says nothing about the lawsuit; and then there is a state refining both ω and ω', so there is a lawsuit and a novelty, at which the value of owning the firm cannot be both \$90 and \$110. Thus, the set of states in which the value of owning the firm is $\$x$ is not persistent and hence not an event (cf. Heifetz et al. 2013, p. 110). At a state in which owning the firm has a value of \$90, an agent may consider possible only states at which owning the firm has a value different than \$90; thus, one would say that the agent has a false belief. However, Heifetz et al. (2006, p. 90) say that the agent does not have a false belief in any *event*, because the set of states in which owning the firm has a particular value is not an event.

Our approach would instead be to accept the event-measurability of acts and make explicit the false beliefs of agents about the consequences of acts: the set of states in which owning the firm has a value of \$90 should be an event, and an agent may falsely believe that that event obtains. This does not undercut the interest in modeling unawareness, since as discussed in Example 4.13, false beliefs may be due to unawareness of possibilities rather than to mistakes in reasoning, misleading evidence, etc.; we can see this when the only possibilities in $\mathcal{K}_i(\omega) \setminus \mathcal{B}_i(\omega)$ are possibilities of which the agent is unaware.

5 Representation

Having seen how to model some concrete examples involving awareness, knowledge, and belief using possibility frames, we now identify exactly the class of examples that can be so represented. Such an example is specified abstractly by a Boolean algebra of events equipped with awareness, knowledge, and belief operators.

Definition 5.1. An *epistemic awareness algebra* is a tuple $\mathbb{A} = (\mathbb{B}, \{A_i\}_{i\in I}, \{K_i\}_{i\in I}, \{B_i\}_{i\in I})$ where $\mathbb{B}$ is a Boolean algebra and A_i, K_i, and B_i are unary operations on $\mathbb{B}$ such that for all $a, b \in \mathbb{B}$ and $\Box_i \in \{K_i, B_i\}$:

- $A_i 1 = 1$ (tautology), $A_i a = A_i \neg a$ (symmetry), and $A_i a \sqcap A_i b \leq A_i(a \sqcap b)$ (agglomeration);

- $K_i 1 = 1$ (knowledge necessitation) and $\Box_i a \sqcap \Box_i b \leq \Box_i(a \sqcap b)$ (knowledge/belief agglomeration);
- if $a \leq b$, then $\Box_i a \sqcap A_i b \leq \Box_i b$ (awareness-restricted monotonicity);
- $K_i a \leq a$ (factivity of knowledge) and $B_i 0 = 0$ (consistency of belief);
- $K_i a \leq B_i a$ (knowledge-belief entailment) and $B_i a \leq A_i a$ (belief-awareness entailment).

As usual, an *isomorphism* between epistemic awareness algebras $\mathbb{A}$ and $\mathbb{A}'$ is a bijection from the carrier set of $\mathbb{A}$ to that of $\mathbb{A}'$ that respects the lattice orders and operations: $a \leq b$ if and only if $f(a) \leq' f(b)$; and for each $\triangle_i \in \{A_i, K_i, B_i\}_{i \in I}$, we have $f(\triangle_i(a)) = \triangle'_i(f(a))$.

Some immediate consequences are that $B_i 1 = 1$ (belief necessitation) and $K_i a \leq A_i a$ (knowledge-awareness entailment). It also follows that awareness is closed under not only meets but also joins.

Lemma 5.2. For any $a, b \in B$, $A_i a \sqcap A_i b \leq A_i(a \sqcup b)$.

Proof. We have $A_i a \sqcap A_i b \leq A_i \neg a \sqcap A_i \neg b \leq A_i(\neg a \sqcap \neg b) \leq A_i \neg(\neg a \sqcap \neg b) = A_i(a \sqcup b)$. □

Of course, we could impose additional axioms, but our representation theorems will be stronger if we can represent *any* epistemic awareness algebra, not just those with special additional properties.

Each epistemic possibility frame gives rise to an epistemic awareness algebra as follows.

Proposition 5.3. If $\mathscr{F} = (\Omega, \sqsubseteq, \mathcal{E}, \{\mathcal{A}_i\}_{i \in I}, \{\mathcal{K}_i\}_{i \in I}, \{\mathcal{B}_i\}_{i \in I})$ is an epistemic possibility frame, then $\mathscr{F}^+ = ((\mathcal{E}, \subseteq), \{\mathbf{A}_i\}_{i \in I}, \{\mathbf{K}_i\}_{i \in I}, \{\mathbf{B}_i\}_{i \in I})$ is an epistemic awareness algebra.

An immediate corollary of $\mathscr{F}^+$ satisfying tautology, symmetry, and agglomeration is that the family of events of which an agent is aware in a possibility forms a Boolean subalgebra of $\mathcal{RO}(\Omega, \sqsubseteq)$.

Corollary 5.4. If $\mathscr{F} = (\Omega, \sqsubseteq, \mathcal{E}, \{\mathcal{A}_i\}_{i \in I}, \{\mathcal{K}_i\}_{i \in I}, \{\mathcal{B}_i\}_{i \in I})$ is an epistemic possibility frame, then for any $\omega \in \Omega$ and $i \in I$, the family $\{E \in \mathcal{E} \mid \omega \in \mathbf{A}_i(E)\}$ contains Ω and is closed under $\neg$ from (1) and $\cap$.

We now proceed in the converse direction: given an epistemic awareness algebra, we wish to construct an epistemic possibility frame $\mathscr{F}$ that represents the algebra as $\mathscr{F}^+$. To this end, recall that a *filter* in a Boolean algebra $\mathbb{B}$ is a nonempty set F of elements of $\mathbb{B}$ that is upward closed under the lattice order of $\mathbb{B}$ (if $a \in F$ and $a \leq b$, then $b \in F$) and closed under the meet operation of $\mathbb{B}$ (if $a, b \in F$, then $a \sqcap b \in F$). A filter is *proper* if it does not contain all elements of $\mathbb{B}$. Given an element $a \in \mathbb{B}$, the set $\Uparrow a = \{b \in B \mid a \leq b\}$ is a filter, and a filter F is *principal* if $F = \Uparrow a$ for some $a \in \mathbb{B}$. The following type of construction is used in possibility semantics (Holliday 2015, § 5.5) but here we must add $\mathcal{A}_i$'s for awareness and modify the treatment of $\mathcal{K}_i$ and $\mathcal{B}_i$ due to the awareness-restricted monotonicity of knowledge and belief.

Definition 5.5. Given an awareness algebra $\mathbb{A} = (\mathbb{B}, \{A_i\}_{i \in I}, \{K_i\}_{i \in I}, \{B_i\}_{i \in I})$, define the frame $\mathbb{A}_+ = (\Omega, \sqsubseteq, \mathcal{E}, \{\mathcal{A}_i\}_{i \in I}, \{\mathcal{K}_i\}_{i \in I}, \{\mathcal{B}_i\}_{i \in I})$ as follows:

1. Ω is the set of all proper filters of $\mathbb{B}$, and $F \sqsubseteq G$ if $F \supseteq G$;
2. $\mathcal{E} = \{\widehat{a} \mid a \in B\}$ where $\widehat{a} = \{F \in \Omega \mid a \in F\}$;
3. $\mathcal{A}_i(F) = \{H \in \Omega \mid H$ is the principal filter of an element a such that $A_i a \in F\}$;
4. $\mathcal{K}_i(F) = \{H \in \Omega \mid$ for all $a_1, \dots, a_n \in \mathbb{B}$, if $K_i a_1 \sqcup \dots \sqcup K_i a_n \in F$, then $a_1 \sqcup \dots \sqcup a_n \in H\}$;

5. $\mathcal{B}_i(F) = \{H \in \Omega \mid \text{for all } a_1, \dots, a_n \in \mathbb{B}, \text{ if } B_i a_1 \sqcup \dots \sqcup B_i a_n \in F, \text{ then } a_1 \sqcup \dots \sqcup b_n \in H\}$.

In the Appendix, we prove the following main representation theorem.

Theorem 5.6. For any epistemic awareness algebra $\mathbb{A}$:

1. $\mathbb{A}_+$ is a standard epistemic possibility frame;
2. the map $a \mapsto \widehat{a}$ is an isomorphism from $\mathbb{A}$ to $(\mathbb{A}_+)^+$.

This result is analogous to the Stone Representation Theorem for Boolean algebras (Stone 1936) (or more precisely, the Stone-like representation without the Axiom of Choice[29] in Holliday 2015 and Bezhanishvili and Holliday 2020) but with awareness, knowledge, and belief added alongside Boolean operations.

Remark 5.7. In the case when $\mathbb{A}$ is finite, we can cut the set Ω in Definition 5.5 down to just the principal filters. In fact, one can typically use a much smaller quasi-principal possibility frame, as in Examples 4.7 and 4.13; one need only add enough partial possibilities to witness an agent's unawareness of events, rather than a partial possibility for each proposition from $\mathbb{A}$. For example, while the possibility frame in Example 4.7 has 37 possibilities, its associated algebra has $2^{20} = 1,048,576$ events. In practical modeling, we usually construct a concise possibility frame $\mathscr{F}$ and then calculate its algebra $\mathscr{F}^+$ of events, instead of starting with a very large algebra $\mathbb{A}$ of events and then applying Theorem 5.6 to obtain a possibility frame $\mathbb{A}_+$. The point of Theorem 5.6 is rather to show that possibility frames do not unreasonably limit what we can model.

6 Conclusion

Theorem 5.6 shows that epistemic possibility frames are capable of representing any scenario involving multi-agent awareness of events, plus knowledge and belief, provided some basic axioms are satisfied in the scenario. Thus, as far as event-based approaches to awareness are concerned, epistemic possibility frames provide a highly versatile modeling tool. We conclude by mentioning several avenues for further development.

First, we can immediately use our possibility frames to interpret a formal logical language for reasoning about unawareness and uncertainty, following standard practice in computer science (Fagin et al. 1995, Halpern 2003) and some work in economics (e.g., Board 2004, Heifetz et al. 2008, Alon and Heifetz 2014). We simply turn an epistemic possibility frame into an *epistemic possibility model* for a propositional language by equipping the frame with a valuation V of atomic formulas such that $V(\mathtt{p}) \in \mathcal{E}$ for each atomic formula $\mathtt{p}$ of the language. One can then recursively define the interpretation of complex formulas built up using negation and conjunction and modalities for awareness, knowledge, and belief for each agent, using the operations in the epistemic awareness algebra arising from the epistemic possibility frame. Finally, one can easily turn our main representation theorem (Theorem 5.6) into a strong completeness theorem for a modal logic of awareness, knowledge, and belief with axioms matching those of epistemic awareness algebras.

A logical approach can then be extended to more expressive languages than that of propositional modal logic, such as propositional modal logic with *propositional quantifiers*, as investigated in Halpern and Rêgo 2009, 2013. In such a language, one may express that agent i knows that there is some event of which she is unaware: $\mathtt{K}_i \exists \mathtt{p}\, \mathtt{U}_i \mathtt{p}$. There is a mathematically elegant semantics for modal logic with propositional quantifiers using complete Boolean algebras (see, e.g., Holliday 2019, Ding 2021) and hence a corresponding

[29] Using the Axiom of Choice, the set Ω in Definition 5.5 may be cut down to just the set of proper filters that are maximal (i.e., ultrafilters) or principal.

semantics using possibility frames in which $\mathcal{E} = \mathcal{RO}(\Omega, \sqsubseteq)$ (Holliday 2021, § 5.1), which one could consider for applications to unawareness. Ding (2021) solves a related problem: modeling a *modest* agent who believes that she must have some false belief, written as $\mathtt{B}_i\exists\mathtt{p}(\mathtt{B}_i\mathtt{p}\wedge\neg\mathtt{p})$, though she does not know which belief it is.

There is also the possibility of using a formal language not just to reason about *awareness of events* as modeled in this paper but also to model *awareness of sentences*, perhaps even developing a sentence-based approached on top of our event-based approach. For example, assume that for every event $E \in \mathcal{E}$ in our model, there is some atomic formula $\mathtt{p}$ for which $V(\mathtt{p}) = E$. For an arbitrary formula φ of the modal language, to formalize awareness of the formula φ itself, we could say that $\mathtt{A}_i\varphi$ is true at a possibility ω just in case for every subformula ψ of φ, we have $\omega \in \mathbf{A}_i[\![\psi]\!]$, where $[\![\psi]\!]$ is the set of possibilities at which ψ is true (an event). This would capture the idea that being aware of complex formulas such as $\varphi\wedge\psi$ or $\varphi\vee\neg\varphi$ requires being aware of φ and of ψ. However, this still has the consequence that if two atomic formulas $\mathtt{p}$ and $\mathtt{q}$ are true in exactly the same states, then i is aware of $\mathtt{p}$ if and only if i is aware of $\mathtt{q}$; to avoid this consequence, one must directly encode awareness of atomic formulas, as in Fagin and Halpern 1988.

Finally, though here we have focused on knowledge and qualitative belief, we can also add probability to our frames (cf. Aumann 1999b, Heifetz et al. 2013). This can be done by assigning to each possibility ω and agent i a set $\mathcal{P}_{\omega,i}$ of probability measures on the Boolean algebra $\{E \in \mathcal{E} \mid \omega \in \mathbf{A}_i(E)\}$ of events of which i is aware in ω. The reason for allowing a *set* of measures—besides wanting to allow multi-prior representations of uncertainty—is that a possibility ω may be *partial*, not settling exactly which probability measure captures the agent's subjective probabilities, leaving us with a set of measures to be narrowed down by further refinements of ω. Appropriate persistence and refinability conditions relating the sets $\mathcal{P}_{\omega,i}$ for different possibilities ω ensure that certain probabilistic events, such as i p-believing E (Monderer and Samet 1989) or i judging that E is at least as likely as F (cf. Alon and Lehrer 2014, Alon and Heifetz 2014), will themselves belong to $\mathcal{RO}(\Omega, \sqsubseteq)$, so we can require that they belong to $\mathcal{E}$. Thus, a full apparatus of multi-agent awareness, knowledge, and probability could be developed using possibility frames, enabling applications to decision theory and game theory with unawareness.

Acknowledgements

For helpful feedback, I thank Mathias Boehm, John Hawthorne, Peter Fritz, Harvey Lederman, Mikayla Kelley, Eric Pacuit, Matthew Mandelkern, and the audience in my course at NASSLLI 2022 at the University of Southern California, where this work was first presented.

A Proofs

In this appendix, we give proofs of results in the main text. To make clear where we use our various assumptions, we typeset them in bold.

Lemma 4.2. Suppose $\mathscr{F} = (\Omega, \sqsubseteq, \mathcal{E}, \{\mathcal{A}_i\}_{i\in I})$ satisfies part 1 of Definition 4.1 and awareness expressibility. If (i) $\mathscr{F}$ satisfies awareness joinability, then (ii) for all $\nu \in \mathcal{A}_i(\omega)$ and $\nu_1, \dots, \nu_n \in \mathcal{A}_i(\omega) \cap \downarrow\nu$, we have $\max((\downarrow\nu_1 \sqcup \dots \sqcup \downarrow\nu_n) \cap \downarrow\nu) \subseteq \mathcal{A}_i(\omega)$. Conversely, if $\max(E)$ is finite for each $E \in \mathcal{E}$, then (ii) implies (i).

Proof. First, assume (i). We prove (ii) by induction on n. For the base case of $n = 1$, if $\nu_1 \in \mathcal{A}_i(\omega) \cap \downarrow\nu$, so $\nu_1 \sqsubseteq \nu$, then $\max(\downarrow\nu_1 \cap \downarrow\nu) = \{\nu_1\} \subseteq \mathcal{A}_i(\omega)$. For the inductive step, suppose $\nu \in \mathcal{A}_i(\omega)$ and $\nu_1, \dots, \nu_{n+1} \in \mathcal{A}_i(\omega) \cap \downarrow\nu$. By the inductive hypothesis, we have $\max((\downarrow\nu_1 \sqcup \dots \sqcup \downarrow\nu_n) \cap \downarrow\nu) \subseteq \mathcal{A}_i(\omega)$; moreover, we have $\max(\downarrow\nu_{n+1}) = \{\nu_{n+1}\} \subseteq \mathcal{A}_i(\omega)$. By **awareness expressibility**, $\downarrow\nu_i \in \mathcal{E}$ for $1 \leq i \leq n+1$, and $\downarrow\nu \in \mathcal{E}$, so

$(\downarrow\nu_1 \sqcup \dots \sqcup \downarrow\nu_n) \cap \downarrow\nu \in \mathcal{E}$ by Definition 3.5 and $\downarrow\nu_{n+1} \in \mathcal{E}$. By the previous two sentences and **awareness joinability**, we have $\max((\downarrow\nu_1 \sqcup \dots \sqcup \downarrow\nu_n) \cap \downarrow\nu) \sqcup \downarrow\nu_{n+1}) \subseteq \mathcal{A}_i(\omega)$. Now we have

$$\begin{aligned}(\downarrow\nu_1 \sqcup \dots \sqcup \downarrow\nu_{n+1}) \cap \downarrow\nu &= ((\downarrow\nu_1 \sqcup \dots \sqcup \downarrow\nu_n) \sqcup \downarrow\nu_{n+1}) \cap \downarrow\nu \\ &= ((\downarrow\nu_1 \sqcup \dots \sqcup \downarrow\nu_n) \cap \downarrow\nu) \sqcup (\downarrow\nu_{n+1} \cap \downarrow\nu) \\ &= ((\downarrow\nu_1 \sqcup \dots \sqcup \downarrow\nu_n) \cap \downarrow\nu) \sqcup \downarrow\nu_{n+1}\end{aligned}$$

using the associative and distributive laws of Boolean algebras and the fact that $\nu_{n+1} \in \downarrow\nu$. From the previous two facts, it follows that $\max((\downarrow\nu_1 \sqcup \dots \sqcup \downarrow\nu_{n+1}) \cap \downarrow\nu) \subseteq \mathcal{A}_i(\omega)$, which establishes (ii).

Now assume (ii) and that $\max(E)$ is finite for each $E \in \mathcal{E}$. Toward proving (i), suppose $\nu \in \mathcal{A}_i(\omega)$, $E, E' \in \mathcal{E}$, and that $\max(E \cap \downarrow\nu) \cup \max(E' \cap \downarrow\nu) \subseteq \mathcal{A}_i(\omega) \cap \downarrow\nu$. By the finiteness assumption, we can write $\max(E \cap \downarrow\nu) \cup \max(E' \cap \downarrow\nu) = \{\nu_1, \dots, \nu_n\}$. Then by (ii), $\max((\downarrow\nu_1 \sqcup \dots \sqcup \downarrow\nu_n) \cap \downarrow\nu) \subseteq \mathcal{A}_i(\omega)$. Now we claim that $(\downarrow\nu_1 \sqcup \dots \sqcup \downarrow\nu_n) \cap \downarrow\nu = (E \sqcup E') \cap \downarrow\nu$. For the left-to-right inclusion, since $\nu_1, \dots, \nu_n \in E \cup E'$, we have $\downarrow\nu_1 \cup \dots \cup \downarrow\nu_n \subseteq E \cup E'$, so $\rho(\downarrow\nu_1 \cup \dots \cup \downarrow\nu_n) \subseteq \rho(E \cup E')$ by Lemma 3.2.1 and hence $\downarrow\nu_1 \sqcup \dots \sqcup \downarrow\nu_n \subseteq E \sqcup E'$ by Theorem 3.3. For the right-to-left inclusion, suppose $\mu \in (E \sqcup E') \cap \downarrow\nu$. To show $\mu \in \downarrow\nu_1 \sqcup \dots \sqcup \downarrow\nu_n$, it suffices to show that for every $\mu' \sqsubseteq \mu$, there is a $\mu'' \sqsubseteq \mu'$ with $\mu'' \in \downarrow\nu_1 \cup \dots \cup \downarrow\nu_n$. Given $\mu' \sqsubseteq \mu$ and $\mu \in E \sqcup E'$, there is a $\mu'' \sqsubseteq \mu'$ with $\mu'' \in E \cup E'$. Without loss of generality, suppose $\mu'' \in E$. Then since $\mu'' \sqsubseteq \mu' \sqsubseteq \mu \sqsubseteq \nu$, we have $\mu'' \in E \cap \downarrow\nu$. By **awareness expressibility**, $\downarrow\nu \in \mathcal{E}$, so we have $E \cap \downarrow\nu \in \mathcal{E}$. Then since $(\Omega, \sqsubseteq, \mathcal{E})$ is **quasi-principal** and $\mu'' \in E \cap \downarrow\nu$, there is a $\mu^* \in \max(E \cap \downarrow\nu)$ such that $\mu'' \sqsubseteq \mu^*$. It follows that $\mu'' \in \downarrow\nu_1 \cup \dots \cup \downarrow\nu_n$, which completes the proof of (i). □

Next we prove the two key lemmas that together show that the set of regular open sets of an epistemic possibility frame is closed under $\mathbf{A}_i$, $\mathbf{K}_i$, and $\mathbf{B}_i$.

Lemma 4.4. Let $(\Omega, \sqsubseteq, \mathcal{E})$ be a possibility frame and $\mathcal{A}_i : \Omega \to \wp(\Omega)$ satisfy awareness persistence and awareness refinability. Then for any $E \in \mathcal{RO}(\Omega, \sqsubseteq)$, we have $\mathbf{A}_i(E) \in \mathcal{RO}(\Omega, \sqsubseteq)$.

Proof. By Lemma 3.1, it suffices to verify persistence and refinability for $\mathbf{A}_i(E)$. Persistence is immediate from the $\forall\omega' \sqsubseteq \omega$ quantification in the definition of $\mathbf{A}_i$. As for refinability, suppose $\omega \notin \mathbf{A}_i(E)$, so there is some $\omega' \sqsubseteq \omega$, $\nu \in \mathcal{A}_i(\omega')$, and $\nu' \in \max(E \cap \downarrow\nu) \cup \max(\neg E \cap \downarrow\nu)$ such that $\nu' \notin \mathcal{A}_i(\omega')$. Given $\nu' \notin \mathcal{A}_i(\omega')$, by **awareness refinability** there is an $\omega'' \sqsubseteq \omega'$ such that for all $\omega''' \sqsubseteq \omega''$, we have $\nu' \notin \mathcal{A}_i(\omega''')$. We claim that for all $\omega''' \sqsubseteq \omega''$, we have $\omega''' \notin \mathbf{A}_i(E)$. Assume for contradiction that $\omega''' \in \mathbf{A}_i(E)$. Since $\omega''' \sqsubseteq \omega'' \sqsubseteq \omega'$, we have $\omega''' \sqsubseteq \omega'$, which with $\nu \in \mathcal{A}_i(\omega')$ implies $\nu \in \mathcal{A}_i(\omega''')$ by **awareness persistence**. Together $\omega''' \in \mathbf{A}_i(E)$ and $\nu \in \mathcal{A}_i(\omega''')$ imply $\max(E \cap \downarrow\nu) \cup \max(\neg E \cap \downarrow\nu) \subseteq \mathcal{A}_i(\omega''')$. Hence $\nu' \in \mathcal{A}_i(\omega''')$, contradicting what we derived above. Thus, $\omega''' \notin \mathbf{A}_i(E)$, which establishes refinability for $\mathbf{A}_i(E)$. □

Lemma 4.10. Let $(\Omega, \sqsubseteq, \mathcal{E})$ be a possibility frame and $\mathcal{R}_i : \Omega \to \wp(\Omega)$ satisfy $\mathcal{R}_i$-monotonicity, $\mathcal{R}_i$-regularity, and $\mathcal{R}_i$-refinability. Then for any $E \in \mathcal{RO}(\Omega, \sqsubseteq)$, we have $\{\omega \in \Omega \mid \mathcal{R}_i(\omega) \subseteq E\} \in \mathcal{RO}(\Omega, \sqsubseteq)$.

Proof. By Lemma 3.1, it suffices to verify persistence and refinability for $\{\omega \in \Omega \mid \mathcal{R}_i(\omega) \subseteq E\}$. For persistence, suppose $\omega' \sqsubseteq \omega$. Then by **$\mathcal{R}_i$-monotonicity**, $\mathcal{R}_i(\omega') \subseteq \mathcal{R}_i(\omega)$, so $\mathcal{R}_i(\omega) \subseteq E$ implies $\mathcal{R}_i(\omega') \subseteq E$. For refinability, suppose $\mathcal{R}_i(\omega) \not\subseteq E$, so there is some $\nu \in \mathcal{R}_i(\omega) \setminus E$. Since $\nu \notin E$ and $E \in \mathcal{RO}(\Omega, \sqsubseteq)$, there is a $\nu' \sqsubseteq \nu$ such that for all $\nu'' \sqsubseteq \nu'$, $\nu'' \notin E$. Since $\nu \in \mathcal{R}_i(\omega)$ and $\nu' \sqsubseteq \nu$, we have $\nu' \in \mathcal{R}_i(\omega)$ by **$\mathcal{R}_i$-regularity**. Then by **$\mathcal{R}_i$-refinability**, there is an $\omega' \sqsubseteq \omega$ such that for all $\omega'' \sqsubseteq \omega'$ there is a $\nu'' \sqsubseteq \nu'$ with $\nu'' \in \mathcal{R}_i(\omega'')$. But as above, $\nu'' \sqsubseteq \nu'$ implies $\nu'' \notin E$, so $\mathcal{R}_i(\omega'') \not\subseteq E$. Thus, we have shown that for all $\omega' \sqsubseteq \omega$ there is an $\omega'' \sqsubseteq \omega'$ with $\mathcal{R}_i(\omega'') \not\subseteq E$, which completes the proof of refinability. □

By these lemmas and the fact that $\mathcal{RO}(\Omega, \sqsubseteq)$ is closed under intersection, it is also closed under $\mathbf{K}_i$ and $\mathbf{B}_i$.

Next we prove that epistemic possibility frames as in Definition 4.9 give rise to epistemic awareness algebras as in Definition 5.1.

Proposition 5.3. If $\mathscr{F} = (\Omega, \sqsubseteq, \mathcal{E}, \{\mathcal{A}_i\}_{i\in I}, \{\mathcal{K}_i\}_{i\in I}, \{\mathcal{B}_i\}_{i\in I})$ is an epistemic possibility frame, then $\mathscr{F}^+ = ((\mathcal{E}, \subseteq), \{\mathbf{A}_i\}_{i\in I}, \{\mathbf{K}_i\}_{i\in I}, \{\mathbf{B}_i\}_{i\in I})$ is an epistemic awareness algebra.

Proof. That $\mathbf{A}_i$ satisfies the tautology and symmetry axioms of Definition 5.1 is obvious. For the agglomeration axiom, suppose $\omega \in \mathbf{A}_i(E_1) \cap \mathbf{A}_i(E_2)$. Toward showing that $\omega \in \mathbf{A}_i(E_1 \cap E_2)$, suppose $\omega' \sqsubseteq \omega$ and $\omega' \mathcal{A}_i \nu$, so ${\downarrow}\nu \in \mathcal{E}$ by **awareness expressibility**. Then since $E_1, E_2 \in \mathcal{E}$, we have $E_1 \cap {\downarrow}\nu, E_2 \cap {\downarrow}\nu, E_1 \cap E_2 \cap {\downarrow}\nu \in \mathcal{E}$.

First suppose that $\nu' \in \max(E_1 \cap E_2 \cap {\downarrow}\nu)$. Then since $\nu' \in E_1$ and the frame is **quasi-principal** (recall Definition 3.6), there is some $\nu^* \in \max(E_1 \cap {\downarrow}\nu)$ with $\nu' \sqsubseteq \nu^*$. Then since $\omega \in \mathbf{A}_i(E_1)$, $\omega' \sqsubseteq \omega$, $\nu \in \mathcal{A}_i(\omega')$, and $\nu^* \in \max(E_1 \cap {\downarrow}\nu)$, we have $\nu^* \in \mathcal{A}_i(w')$. Moreover, we have $\nu' \in \max(E_2 \cap {\downarrow}\nu^*)$, for if there is some ν'' such that $\nu' \sqsubset \nu'' \in E_2 \cap {\downarrow}\nu^*$, then $\nu'' \in E_1 \cap E_2 \cap {\downarrow}\nu$, contradicting the fact that $\nu' \in \max(E_1 \cap E_2 \cap {\downarrow}\nu)$. Then since $\omega \in \mathbf{A}_i(E_2)$, $\omega' \sqsubseteq \omega$, $\nu^* \in \mathcal{A}_i(w')$, and $\nu' \in \max(E_2 \cap {\downarrow}\nu^*)$, we have $\nu' \in \mathcal{A}_i(\omega')$.

Now suppose $\nu' \in \max(\neg(E_1 \cap E_2) \cap {\downarrow}\nu) = \max((\neg E_1 \sqcup \neg E_2) \cap {\downarrow}\nu)$. Since $\omega \in \mathbf{A}_i(E_1) \cap \mathbf{A}_i(E_2)$, we have $\omega \in \mathbf{A}_i(\neg E_1) \cap \mathbf{A}_i(\neg E_2)$, which with $\omega' \sqsubseteq \omega$ and $\nu \in \mathcal{A}_i(\omega')$ implies $\max(\neg E_1 \cap {\downarrow}\nu) \cup \max(\neg E_2 \cap {\downarrow}\nu) \subseteq \mathcal{A}_i(\omega')$. Thus, by **awareness joinability**,

$$\max(\rho(\max(\neg E_1 \cap {\downarrow}\nu) \cup \max(\neg E_2 \cap {\downarrow}\nu))) \subseteq \mathcal{A}_i(\omega'). \tag{8}$$

Moreover, we claim that

$${\downarrow}(\max(\neg E_1 \cap {\downarrow}\nu) \cup \max(\neg E_2 \cap {\downarrow}\nu)) = (\neg E_1 \cup \neg E_2) \cap {\downarrow}\nu. \tag{9}$$

The left-to-right inclusion is obvious, since $\neg E_1$ and $\neg E_2$ are downsets. For the right-to-left inclusion, suppose $\nu' \in (\neg E_1 \cup \neg E_2) \cap {\downarrow}\nu$. Hence $\nu' \in \neg E_i \cap {\downarrow}\nu$ for some $i \in \{1, 2\}$. Since the poset is **quasi-principal**, we have $\nu' \in {\downarrow}\max(\neg E_i \cap {\downarrow}\nu)$ and hence $\nu' \in {\downarrow}(\max(\neg E_1 \cap {\downarrow}\nu) \cup \max(\neg E_2 \cap {\downarrow}\nu))$, which establishes (9). Thus, we have:

$$\begin{aligned}
& {\downarrow}(\max(\neg E_1 \cap {\downarrow}\nu) \cup \max(\neg E_2 \cap {\downarrow}\nu)) = (\neg E_1 \cup \neg E_2) \cap {\downarrow}\nu \\
\Rightarrow\ & \rho({\downarrow}(\max(\neg E_1 \cap {\downarrow}\nu) \cup \max(\neg E_2 \cap {\downarrow}\nu))) = \rho((\neg E_1 \cup \neg E_2) \cap {\downarrow}\nu) \\
\Rightarrow\ & \rho(\max(\neg E_1 \cap {\downarrow}\nu) \cup \max(\neg E_2 \cap {\downarrow}\nu)) = \rho((\neg E_1 \cup \neg E_2) \cap {\downarrow}\nu) \text{ by definition of } \rho \text{, idempotence of } {\downarrow} \\
\Rightarrow\ & \rho(\max(\neg E_1 \cap {\downarrow}\nu) \cup \max(\neg E_2 \cap {\downarrow}\nu)) = \rho(\neg E_1 \cup \neg E_2) \cap \rho({\downarrow}\nu) \text{ by Lemma 3.2.4} \\
\Rightarrow\ & \rho(\max(\neg E_1 \cap {\downarrow}\nu) \cup \max(\neg E_2 \cap {\downarrow}\nu)) = (\neg E_1 \sqcup \neg E_2) \cap {\downarrow}\nu \text{ by definition of } \sqcup \text{ and } {\downarrow}\nu \in \mathcal{E} \subseteq \mathcal{RO}(\Omega, \sqsubseteq) \\
\Rightarrow\ & \max(\rho(\max(\neg E_1 \cap {\downarrow}\nu) \cup \max(\neg E_2 \cap {\downarrow}\nu)) = \max((\neg E_1 \sqcup \neg E_2) \cap {\downarrow}\nu).
\end{aligned}$$

Then given (8), we have $\max((\neg E_1 \sqcup \neg E_2) \cap {\downarrow}\nu) \subseteq \mathcal{A}_i(\omega')$. This completes the proof of $\omega \in \mathbf{A}_i(E_1 \cap E_2)$.

Checking the axioms on knowledge and belief from Definition 5.1 is straightforward $\square$

Finally, we prove our main representation theorem.

Theorem 5.6. For any epistemic awareness algebra $\mathbb{A}$:

1. $\mathbb{A}_+$ is a standard epistemic possibility frame;
2. the map $a \mapsto \widehat{a}$ is an isomorphism from $\mathbb{A}$ to $(\mathbb{A}_+)^+$.

Proof. For part 1, let $\mathbb{B}$ be the underlying Boolean algebra of $\mathbb{A}$, and recall the construction of the frame $\mathbb{A}_+ = (\Omega, \sqsubseteq, \mathcal{E}, \{\mathcal{A}_i\}_{i\in I}, \{\mathcal{K}_i\}_{i\in I}, \{\mathcal{B}_i\}_{i\in I})$ from Definition 5.5. Also note that for a filter F in $\mathbb{B}$ and $b \in \mathbb{B}$, the smallest filter extending $F \cup \{b\}$ is $\{c \in \mathbb{B} \mid \text{for some } a \in F, a \sqcap b \leq c\}$. A basic fact about Boolean algebras is that if F is a proper filter and $b \notin F$, then the smallest filter extending $F \cup \{\neg b\}$ is proper.

Clearly $(\Omega, \sqsubseteq)$ is a poset. To see that $(\Omega, \sqsubseteq, \mathcal{E})$ is a possibility frame, we must show that $\mathcal{E}$ is a subalgebra of $\mathcal{RO}(\Omega, \sqsubseteq)$. First, we show that each $\widehat{a}$ is regular open. By Lemma 3.1, it suffices to show that $\widehat{a}$ satisfies persistence and refinability. For persistence, if $F \in \widehat{a}$, so $a \in F$, and $F' \sqsubseteq F$, so $F' \supseteq F$, then $a \in F'$ and hence $F' \in \widehat{a}$. For refinability, if $F \notin \widehat{a}$, so $a \notin F$, then the smallest filter F' extending $F \cup \{\neg a\}$ is proper, so $F' \sqsubseteq F$, and for all $F'' \sqsubseteq F'$, we have $\neg a \in F''$, so $a \notin F''$ since F'' is proper, so $F'' \notin \widehat{a}$. Finally, $\mathcal{E}$ is closed under $\cap$ and $\neg$, as (i) $\widehat{a} \cap \widehat{b} = \widehat{a \sqcap b}$, and (ii) $\neg\widehat{a} = \widehat{\neg a}$. Condition (i) follows from that fact that if F is a filter, then $a, b \in F$ if and only if $a \sqcap b \in F$. For condition (ii), if $F \in \widehat{\neg a}$, so $\neg a \in F$, then for any proper filter F' extending F, $a \notin F'$, so $F' \notin \widehat{a}$; this shows $F \in \neg\widehat{a}$. Conversely, if $F \notin \widehat{\neg a}$, so $\neg a \notin F$, then the smallest filter F' extending $F \cup \{a\}$ is proper, so $F' \sqsubseteq F$; this shows $F \notin \neg\widehat{a}$.

To show that $(\Omega, \sqsubseteq)$ is quasi-principal, we must show that for all $\widehat{a} \in \mathcal{E}$ and $F \in \widehat{a}$, we have that $F \in {\downarrow}\max(\widehat{a})$. Indeed, $\max(\widehat{a}) = \{{\Uparrow}a\}$ (recall that ${\Uparrow}a$ is the principal filter generated by a), and from $F \in \widehat{a}$ we have $a \in F$, so $F \supseteq {\Uparrow}a$ and hence $F \sqsubseteq {\Uparrow}a$, so $F \in {\downarrow}\max(\widehat{a})$. Finally, the maximum element m of the poset of proper filters is the principal filter $\{1\}$.

Next we verify the five conditions on $\mathcal{A}_i$:

- awareness nonvacuity: for all $F \in \Omega$, $m \in \mathcal{A}_i(F)$.

  We have $\{1\} \in \mathcal{A}_i(F)$ by the **tautology** axiom of Definition 5.1.

- awareness expressibility: for all $F \in \Omega$ and $H \in \mathcal{A}_i(F)$, ${\downarrow}H \in \mathcal{E}$.

  If $H \in \mathcal{A}_i(F)$, then $H = {\Uparrow}a$ for some $a \in \mathbb{B}$, in which case ${\downarrow}H = \widehat{a} \in \mathcal{E}$.

- awareness persistence: if $F' \sqsubseteq F$, then $\mathcal{A}_i(F) \subseteq \mathcal{A}_i(F')$.

  Immediate from the definitions of $\mathcal{A}_i$ and $\sqsubseteq$ in $\mathbb{A}_+$.

- awareness refinability: if $H \notin \mathcal{A}_i(F)$, then $\exists F' \sqsubseteq F\ \forall F'' \sqsubseteq F'\ H \notin \mathcal{A}_i(F'')$.

  Suppose not $H \notin \mathcal{A}_i(F)$. If H is not a principal filter, then set $F' = F$, and then for all $F'' \sqsubseteq F'$, we have $H \notin \mathcal{A}_i(F'')$. Now suppose H is the principal filter of a. Since $H \notin \mathcal{A}_i(F)$, it follows that $A_i a \notin F$. Hence the smallest filter F' extending $F \cup \{\neg A_i a\}$ is proper. Then for all $F'' \sqsubseteq F'$, i.e., $F'' \supseteq F'$, we have $\neg A_i a \in F''$ and hence $A_i a \notin F''$ since F'' is proper, so $H \notin \mathcal{A}_i(F'')$.

- awareness joinability: if $H \in \mathcal{A}_i(F)$ and $\max(\widehat{a} \cap {\downarrow}H) \cup \max(\widehat{b} \cap {\downarrow}H) \subseteq \mathcal{A}_i(F)$, then

  $$\max((\widehat{a} \sqcup \widehat{b}) \cap {\downarrow}H) \subseteq \mathcal{A}_i(F).$$

  Assuming $H \in \mathcal{A}_i(F)$, we have $H = {\Uparrow}c$ for some $c \in \mathbb{B}$. Then $\max((\widehat{a} \sqcup \widehat{b}) \cap {\downarrow}H) = \{{\Uparrow}((a \sqcup b) \sqcap c)\}$, $\max(\widehat{a} \cap {\downarrow}H) = \{{\Uparrow}(a \sqcap c)\}$, and $\max(\widehat{b} \cap {\downarrow}H) = \{{\Uparrow}(b \sqcap c)\}$. Assuming $\max(\widehat{a} \cap {\downarrow}H) \cup \max(\widehat{b} \cap {\downarrow}H) \subseteq \mathcal{A}_i(F)$, it follows that $A_i(a \sqcap c), A_i(b \sqcap c) \in F$. Hence $A_i((a \sqcap c) \sqcup (b \sqcap c)) \in F$ by Lemma 5.2, so $A_i((a \sqcup b) \sqcap c) \in F$ by the distributive law of Boolean algebras. Thus, ${\Uparrow}((a \sqcup b) \sqcap c) \in \mathcal{A}_i(F)$, so indeed $\max((\widehat{a} \sqcup \widehat{b}) \cap {\downarrow}H) \subseteq \mathcal{A}_i(F)$.

Finally, we verify the conditions on $\mathcal{R}_i \in \{\mathcal{K}_i, \mathcal{B}_i\}$, reasoning about $\Box_i \in \{K_i, B_i\}$:

- $\mathcal{R}_i$-monotonicity: if $F' \sqsubseteq F$, then $\mathcal{R}_i(F') \subseteq \mathcal{R}_i(F)$.

  Assume $F' \sqsubseteq F$ and $H \in \mathcal{R}_i(F')$. Toward showing $H \in \mathcal{R}_i(F)$, suppose $\Box_i a_1 \sqcup \cdots \sqcup \Box_i a_n \in F$. Then since $F' \sqsubseteq F$, $F' \supseteq F$, so $\Box_i a_1 \sqcup \cdots \sqcup \Box_i a_n \in F'$, which with $H \in \mathcal{R}_i(F')$ implies $a_1 \sqcup \cdots \sqcup a_n \in H$.

- $\mathcal{R}_i$-regularity: $\mathcal{R}_i(F) \in \mathcal{RO}(\Omega, \sqsubseteq)$.

  By Lemma 3.2, it suffices to show that $\mathcal{R}_i(F)$ satisfies persistence and refinability. That $\mathcal{R}_i(F)$ satisfies persistence is immediate from the definition of $\mathcal{R}_i$ and the definition of $\sqsubseteq$ as $\supseteq$. For refinability, suppose $H \notin \mathcal{R}_i(F)$. Hence there is $\Box_i a_1 \sqcup \cdots \sqcup \Box_i a_n \in F$ such that $a_1 \sqcup \cdots \sqcup a_n \notin F$. It follows that the smallest filter H' extending $H \cup \{\neg(a_1 \sqcup \cdots \sqcup a_n)\}$ is proper. Now suppose $H'' \sqsubseteq H'$, so H'' is a proper filter with $H'' \supseteq H'$ and hence $\neg(a_1 \sqcup \cdots \sqcup a_n) \in H''$. Then $a_1 \sqcup \cdots \sqcup a_n \notin H''$, so $H'' \notin \mathcal{R}_i(F)$.

- $\mathcal{R}_i$-refinability: if $H \in \mathcal{R}_i(F)$, then $\exists F' \sqsubseteq F\ \forall F'' \sqsubseteq F'\ \exists H' \sqsubseteq H$: $H' \in \mathcal{R}_i(F'')$.

  Assuming $H \in \mathcal{R}_i(F)$, we claim that the smallest filter F' extending

  $$F \cup \{\neg\Box_i c \mid \exists b \in H : c \leq \neg b\}$$

  is proper. If not, then there are $a \in F$, $b_1, \dots, b_n \in H$, and $c_1, \dots, c_n$ with $c_k \leq \neg b_k$ for $1 \leq k \leq n$ such that $a \sqcap \neg\Box_i c_1 \sqcap \cdots \sqcap \neg\Box_i c_n = 0$, which implies $a \leq \neg(\neg\Box_i c_1 \sqcap \cdots \sqcap \neg\Box_i c_n) = \Box_i c_1 \sqcup \cdots \sqcup \Box_i c_n$. Hence $\Box_i c_1 \sqcup \cdots \sqcup \Box_i c_n \in F$, which with $H \in \mathcal{R}_i(F)$ implies $c_1 \sqcup \cdots \sqcup c_n \in H$ and hence $\neg b_1 \sqcup \cdots \sqcup \neg b_n \in H$, contradicting the fact that $b_1, \dots, b_n \in H$ and H is a proper filter. Thus, F' is indeed proper, so $F' \sqsubseteq F$. Now consider any $F'' \sqsubseteq F'$, so $F'' \supseteq F'$. We claim that the smallest filter H' extending

  $$H \cup \{a_1 \sqcup \cdots \sqcup a_n \mid \Box_i a_1 \sqcup \cdots \sqcup \Box_i a_n \in F''\}$$

  is proper. If not, there is some $b \in H$ and a family $\{a_1^j \sqcup \cdots \sqcup a_{n_j}^j\}_{j \in J}$ for a nonempty finite J such that $\Box_i a_1^j \sqcup \cdots \sqcup \Box_i a_{n_j}^j \in F''$ and

  $$\bigsqcap_{j \in J}(a_1^j \sqcup \cdots \sqcup a_{n_j}^j) \leq \neg b.$$

  Let Λ be the set of all choice functions λ such that for $j \in J$, $1 \leq \lambda(j) \leq n_j$. Then for each $\lambda \in \Lambda$,

  $$\bigsqcap_{j \in J} a_{\lambda(j)}^j \leq \neg b.$$

  Hence for each $\lambda \in \Lambda$, we have $\neg\Box_i \bigsqcap_{j \in J} a_{\lambda(j)}^j \in F'$ by construction of F' and hence $\neg\Box_i \bigsqcap_{j \in J} a_{\lambda(j)}^j \in F''$. On the other hand,

  $$\bigsqcap_{j \in J}(\Box_i a_1^j \sqcup \cdots \sqcup \Box_i a_{n_j}^j) \leq \bigsqcup_{\lambda \in \Lambda} \bigsqcap_{j \in J} \Box_i a_{\lambda(j)}^j \leq \bigsqcup_{\lambda \in \Lambda} \Box_i \bigsqcap_{j \in J} a_{\lambda(j)}^j, \tag{10}$$

  using Boolean distributivity for the first inequality and **knowledge/belief agglomeration** for the second, since J is finite. Then since $\bigsqcap_{j \in J}(\Box_i a_1^j \sqcup \cdots \sqcup \Box_i a_{n_j}^j) \in F''$, we have $\bigsqcup_{\lambda \in \Lambda} \Box_i \bigsqcap_{j \in J} a_{\lambda(j)}^j \in F''$. But we concluded above that for each $\lambda \in \Lambda$ we have $\neg\Box_i \bigsqcap_{j \in J} a_{\lambda(j)}^j \in F''$, which implies $\bigsqcap_{\lambda \in \Lambda} \neg\Box_i \bigsqcap_{j \in J} a_{\lambda(j)}^j \in F''$ and hence $\neg \bigsqcup_{\lambda \in \Lambda} \Box_i \bigsqcap_{j \in J} a_{\lambda(j)}^j \in F''$, contradicting the fact that F'' is proper. Thus, we conclude that H' is proper. Hence $H' \sqsubseteq H$, and $H' \in \mathcal{R}_i(F'')$ by construction.

- epistemic factivity: $F \in \mathcal{K}_i(F)$.
  Immediate from **factivity for knowledge** and the definition of $\mathcal{K}_i$.

- doxastic consistency: $\mathcal{B}_i(F) \neq \varnothing$.

  We claim that the smallest filter H extending $\{a_1 \sqcup \dots \sqcup a_n \mid B_i a_1 \sqcup \dots \sqcup B_i a_n \in F\}$ is proper. If not, then there is a family $\{a_1^j \sqcup \dots \sqcup a_{n_j}^j\}_{j\in J}$ for a nonempty finite J such that $B_i a_1^j \sqcup \dots \sqcup B_i a_{n_j}^j \in F$ and $\bigsqcap_{j\in J}(a_1^j \sqcup \dots \sqcup a_{n_j}^j) = 0$. As above, let Λ be the set of all choice functions λ such that for $j \in J$, $1 \leq \lambda(j) \leq n_j$. It follows that for each $\lambda \in \Lambda$, $\bigsqcap_{j\in J} a_{\lambda(j)}^j = 0$. Then $\bigsqcap_{j\in J}(B_i a_1 \sqcup \dots \sqcup B_i a_n) \leq \bigsqcup_{\lambda\in\Lambda} B_i \bigsqcap_{j\in J} a_{\lambda(j)}^j \leq 0$ using the same reasoning as in (10) for the first inequality and **belief consistency** for the second. Thus, $0 \in F$, contradicting the fact that F is proper. Hence H is proper, and by construction, $H \in \mathcal{B}_i(F)$.

- doxastic inclusion: $\mathcal{B}_i(F) \subseteq \mathcal{K}_i(F)$.
  Immediate from **knowledge-belief entailment** and the definitions of $\mathcal{B}_i$ and $\mathcal{K}_i$.

This completes the proof of part 1.

For part 2, we already observed that $\widehat{\cdot}$ commutes with meet and complement in (i) and (ii) in the second paragraph of the proof, so next we prove that $\widehat{\cdot}$ commutes with the awareness operation. Suppose $F \in \widehat{A_i b}$, so $A_i b \in F$. Hence $A_i \neg b \in F$ by the **symmetry** axiom. Suppose $\Uparrow b \neq \Omega$ and $\Uparrow b \neq \varnothing$. It follows that neither b nor $\neg b$ is 0 in $\mathbb{B}$. Now consider any $F' \sqsubseteq F$, so $F' \supseteq F$, and suppose $H \in \mathcal{A}_i(F')$. Hence $H = \Uparrow c$ for some c such that $A_i c \in F'$. Then since $A_i b \in F'$ and $A_i \neg b \in F'$, we have $A_i(b \sqcap c) \in F'$ and $A_i(\neg b \sqcap c) \in F'$ by **agglomeration**, so $\Uparrow(b \sqcap c) \in \mathcal{A}_i(F')$ and $\Uparrow(\neg b \sqcap c) \in \mathcal{A}_i(F')$. Then since $\max(\widehat{b} \cap {\downarrow}H) = \{\Uparrow(b \sqcap c)\}$ and $\max(\neg\widehat{b} \cap {\downarrow}H) = \{\Uparrow(\neg b \sqcap c)\}$, we have $\max(\widehat{b} \cap {\downarrow}H) \cup \max(\neg\widehat{b} \cap {\downarrow}H) \subseteq \mathcal{A}_i(F')$. This shows that $F \in \mathbf{A}_i\widehat{b}$. Now suppose $F \notin \widehat{A_i b}$, so $A_i b \notin F$ and hence $\Uparrow b \notin \mathcal{A}_i(F)$. Then since $\{1\} \in \mathcal{A}_i(F)$ by the **tautology** axiom and $\max(\widehat{b} \cap {\downarrow}\{1\}) = \{\Uparrow b\}$, it follows that $F \notin \mathbf{A}_i\widehat{b}$. This completes the proof that $\widehat{A_i b} = \mathbf{A}_i\widehat{b}$.

It is now easy to see that the frame is standard: for if $H \in \mathcal{A}_i(F)$, then H is the principal filter of an element a such that $A_i a \in F$, in which case by the previous paragraph, $F \in \mathbf{A}_i\widehat{a}$, so $F \in \mathbf{A}_i{\downarrow}H$.

Finally, we show that $\widehat{\cdot}$ commutes with the knowledge operation (the proof for belief is analogous). Suppose $F \in \widehat{K_i b}$, so $K_i b \in F$. Then by **knowledge-awareness entailment**, $A_i b \in F$, so by the previous paragraph, $F \in \mathbf{A}_i\widehat{b}$. Moreover, $K_i b \in F$ implies that for each $H \in \mathcal{K}_i(F)$, we have $b \in H$ and hence $H \in \widehat{b}$, so $\mathcal{K}_i(F) \subseteq \widehat{b}$. Thus, by definition of $\mathbf{K}_i$, $F \in \mathbf{K}_i\widehat{b}$. Now suppose $F \notin \widehat{K_i b}$, so $K_i b \notin F$ and hence $b \neq 1$ by **necessitation**. Case 1: $A_i b \notin F$. Then by the previous paragraph, $F \notin \mathbf{A}_i\widehat{b}$, which implies $F \notin \mathbf{K}_i\widehat{b}$. Case 2: $A_i b \in F$. Then we claim that the smallest filter H extending

$$\{a_1 \sqcup \dots \sqcup a_n \mid K_i a_1 \sqcup \dots \sqcup K_i a_n \in F\} \cup \{\neg b\}$$

is proper. If not, then since $b \neq 1$, there is a family $\{a_1^j \sqcup \dots \sqcup a_{n_j}^j\}_{j\in J}$ for a nonempty finite J such that $K_i a_1^j \sqcup \dots \sqcup K_i a_n^j \in F$ and $\bigsqcap_{j\in J}(a_1^j \sqcup \dots \sqcup a_{n_j}^j) \leq b$. As before, let Λ be the set of all choice functions λ such that for $j \in J$, $1 \leq \lambda(j) \leq n_j$. Then for each $\lambda \in \Lambda$, $\bigsqcap_{j\in J} a_{\lambda(j)}^j \leq b$ and hence

$$A_i b \sqcap \bigsqcap_{j\in J} K_i a_{\lambda(j)}^j \leq A_i b \sqcap K_i \bigsqcap_{j\in J} a_{\lambda(j)}^j \leq K_i b, \tag{11}$$

using **knowledge agglomeration** for the first inequality and **awareness-restricted monotonicity** for the second. Then

$$A_i b \sqcap \prod_{j \in J} (K_i a_1^j \sqcup \cdots \sqcup K_i a_{n_j}^j) \leq A_i b \sqcap \bigsqcup_{\lambda \in \Lambda} \prod_{j \in J} K_i a_{\lambda(j)}^j \leq \bigsqcup_{\lambda \in \Lambda} (A_i b \sqcap \prod_{j \in J} K_i a_{\lambda(j)}^j) \leq K_i b, \tag{12}$$

using Boolean distributivity for the first two inequalities and then (11) for the last inequality. Then since the element on the far left belongs to F, so does $K_i b$, contradicting what we derived above. Thus, H is indeed a proper filter. By construction $H \in \mathcal{K}_i(F)$, and $b \notin H$ since H is proper, so $H \notin \widehat{b}$. Hence $F \notin \mathbf{K}_i \widehat{b}$. $\square$

References

Shiri Alon and Aviad Heifetz. The logic of Knightian games. *Economic Theory Bulletin*, 2(2):161–182, 2014.

Shiri Alon and Ehud Lehrer. Subjective multi-prior probability: A representation of a partial likelihood relation. *Journal of Economic Theory*, 151:476–492, 2014.

Robert J. Aumann. Agreeing to disagree. *Annals of Statistics*, 4:1236–1239, 1976.

Robert J. Aumann. Interactive epistemology I: Knowledge. *International Journal of Game Theory*, 28(3): 263–300, 1999a.

Robert J. Aumann. Interactive epistemology II: Probability. *International Journal of Game Theory*, 28(3): 301–314, 1999b.

Gaia Belardinelli and Rasmus K. Rendsvig. Awareness logic: A Kripke-based rendition of the Heifetz-Meier-Schipper model. In Manuel A. Martins and Igor Sedlár, editors, *Dynamic Logic. New Trends and Applications. DaLi 2020*, pages 22–50, Cham, 2020. Springer.

Johan van Benthem. Correspondence theory. In D. M. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume 3, pages 325–408. Springer, Dordrecht, 2nd edition, 2001.

Johan van Benthem, Nick Bezhanishvili, and Wesley H. Holliday. A bimodal perspective on possibility semantics. *Journal of Logic and Computation*, 27(5):1353–1389, 2017.

Nick Bezhanishvili and Wesley H. Holliday. Choice-free Stone duality. *The Journal of Symbolic Logic*, 85(1): 109–148, 2020.

Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge University Press, New York, 2001.

Oliver Board. Dynamic Interactive Epistemology. *Games and Economic Behavior*, 49:49–80, 2004.

Stanley Burris and H. P. Sankappanavar. *A Course in Universal Algebra*. Springer-Verlag, New York, 1981.

Eddie Dekel, Barton L. Lipman, and Aldo Rustichini. Standard state-space models preclude unawareness. *Econometrica*, 66(1):159–173, 1998.

Franz Dietrich. Savage's theorem under changing awareness. *Journal of Economic Theory*, 176:1–54, 2018.

Yifeng Ding. On the logic of belief and propositional quantification. *Journal of Philosophical Logic*, 50: 1143–1198, 2021.

Yifeng Ding, Wesley H. Holliday, and Cedegao Zhang. When do introspection axioms matter for multi-agent epistemic reasoning? In Lawrence S. Moss, editor, *Proceedings of TARK 2019*, volume 297, pages 121–139. Electronic Proceedings in Theoretical Computer Science, 2019.

Adam Dominiak and Gerelt Tserenjigmid. Ambiguity under growing awareness. *Journal of Economic Theory*, 199(105256), 2022.

Edward Elliot. Unawareness and implicit belief. Manuscript, 2022.

Ronald Fagin and Joseph Y. Halpern. Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34: 39–76, 1988.

Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge, Mass., 1995.

Frederic B. Fitch. A logical analysis of some value concepts. *The Journal of Symbolic Logic*, 28(2):135–142, 1963.

Peter Fritz and Harvey Lederman. Standard state space models of unawareness. In *Proceedings of TARK 2015*, pages 141–158. Electronic Proceedings in Theoretical Computer Science, 2015.

Satoshi Fukuda. Unawareness without AU Introspection. *Journal of Mathematical Economics*, 94(102456), 2021.

Spyros Galanis. The value of information in risk-sharing environments with unawareness. *Games and Economic Behavior*, 97:1–18, 2016.

Spyros Galanis. Speculation under unawareness. *Games and Economic Behavior*, 109:598–615, 2018.

John Geanakoplos. Game theory without partitions, and applications to speculation and consensus. Cowles Foundation Discussion Paper no. 914, 1989.

S. Givant and P. Halmos. *Introduction to Boolean Algebras*. Springer, New York, 2009.

Simon Grant, J. Jude Kline, Patrick O'Callaghan, and John Quiggin. Sub-models for interactive unawareness. *Theory and Decision*, 79:601–613, 2015.

Pierfrancesco Guarino. An epistemic analysis of dynamic games with unawareness. *Games and Economic Behavior*, 120:257–288, 2020.

Joseph Y. Halpern. Alternative semantics for unawareness. *Games and Economic Behavior*, 37:321–339, 2001.

Joseph Y. Halpern. *Reasoning about Uncertainty*. MIT Press, 2003.

Joseph Y. Halpern and Leandro C. Rêgo. Reasoning about knowledge of unawareness. *Games and Economic Behavior*, 67(2):503–525, 2009.

Joseph Y. Halpern and Leandro C. Rêgo. Reasoning about knowledge of unawareness revisited. *Mathematical Social Sciences*, 65(2):73–84, 2013.

Joseph Y. Halpern and Leandro Chaves Rêgo. Interactive unawareness revisisted. *Games and Economic Behavior*, 62:232–262, 2008.

Aviad Heifetz, Martin Meier, and Burkhard C. Schipper. Interactive unawareness. *Journal of Economic Theory*, 130:78–94, 2006.

Aviad Heifetz, Martin Meier, and Burkhard C. Schipper. A canonical model for interactive unawareness. *Games and Economic Behavior*, 62:304–324, 2008.

Aviad Heifetz, Martin Meier, and Burkhard C. Schipper. Unawareness, beliefs, and speculative trade. *Games and Economic Behavior*, 77:100–121, 2013.

Jaakko Hintikka. *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Cornell University Press, Ithaca, NY, 1962.

Wesley H. Holliday. Partiality and adjointness in modal logic. In R. Goré, B. Kooi, and A. Kurucz, editors, *Advances in Modal Logic, Vol. 10*, pages 313–332. College Publications, London, 2014.

Wesley H. Holliday. Possibility frames and forcing for modal logic. Forthcoming in *The Australasian Journal of Logic*, UC Berkeley Working Paper in Logic and the Methodology of Science, https://escholarship.org/uc/item/0tm6b30q, 2015.

Wesley H. Holliday. A note on algebraic semantics for S5 with propositional quantifiers. *Notre Dame Journal of Formal Logic*, 60(2):311–332, 2019.

Wesley H. Holliday. Possibility semantics. In Melvin Fitting, editor, *Selected Topics from Contemporary Logics*, Landscapes in Logic, pages 363–476. College Publications, 2021.

Wesley H. Holliday and Thomas F. Icard, III. Moorean Phenomena in Epistemic Logic. In Lev Beklemishev, Valentin Goranko, and Valentin Shehtman, editors, *Advances in Modal Logic*, volume 8, pages 178–199. College Publications, London, 2010.

I. L. Humberstone. From worlds to possibilities. *Journal of Philosophical Logic*, 10(3):313–339, 1981.

Edi Karni and Marie-Louise Vierø. Awareness of unawareness: A theory of decision making in the face of ignorance. *Journal of Economic Theory*, 168:301–328, 2017.

Saul A. Kripke. Semantical analysis of modal logic I. Normal modal propositional calculi. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 9:67–96, 1963.

Jing Li. Information structures with unawareness. *Journal of Economic Theory*, 144:977–993, 2009.

Luc Lismont and Phillippe Mongin. On the logic of common belief and common knowledge. *Theory and Decision*, 37:75–106, 1994.

Paul Milgrom and Nancy Stokey. Information, trade and common knowledge. *Journal of Economic Theory*, 26(1):17–27, 1982.

Salvatore Modica and Aldo Rustichini. Awareness and partitional information structures. *Theory and Decision*, 37:107–124, 1994.

Salvatore Modica and Aldo Rustichini. Unawareness and partitional information structures. *Games and Economic Behavior*, 27:265–298, 1999.

Dov Monderer and Dov Samet. Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1:170–190, 1989.

Evan Piermont. Introspective unawareness and observable choice. *Games and Economic Behavior*, 106: 134–152, 2017.

Evan Piermont. Algebraic semantics for relative truth, awareness, and possibility. arXiv:1910.07275, 2021.

John Quiggin. The value of information and the value of awareness. *Theory and Decision*, 80:167–185, 2016.

Burkhard C. Schipper. Awareness-dependent subjective expected utility. *International Journal of Game Theory*, 42:725–753, 2013.

Burkhard C. Schipper. Unawareness—A gentle introduction to both the literature and the special issue. *Mathematical Social Sciences*, 70:1–9, 2014a.

Burkhard C. Schipper. Preference-based unawareness. *Mathematical Social Sciences*, 70:34–31, 2014b.

Burkhard C. Schipper. Awareness. In Hans van Ditmarsch, Joseph Y. Halpern, Wiebe van der Hoek, and Barteld Kooi, editors, *Handbook of Epistemic Logic*, pages 77–146. College Publications, London, 2015.

Burkhard C. Schipper. Introduction to the special issue on unawareness. *The B.E. Journal of Theoretical Economics*, 21(2):355–360, 2021a.

Burkhard C. Schipper. Discovery and equilibrium in games with unawareness. *Journal of Economic Theory*, 198(105365), 2021b.

Robert Stalnaker. On logics of knowledge and belief. *Philosophical Studies*, 128:169–199, 2006.

Robert C. Stalnaker. *Inquiry*. MIT Press, Cambridge, Mass., 1984.

M. H. Stone. The theory of representation for Boolean algebras. *Transactions of the American Mathematical Society*, 40(1):37–111, 1936.

G. Takeuti and W. M. Zaring. *Axiomatic Set Theory*. Springer-Verlag, New York, 1973.

Alfred Tarski. Über additive und multiplikative Mengenkörper und Mengenfunktionen. *Sprawozdania z Posiedzeń Towarzystwa Naukowego Warszawskiego, Wydział III, Nauk Matematyczno-fizycznych*, 30:151–181, 1937.

Kentarô Yamamoto. Results in modal correspondence theory for possibility semantics. *Journal of Logic and Computation*, 27(8):2411–2430, 2017.

Zhiguang Zhao. Algorithmic correspondence and canonicity for possibility semantics. *Journal of Logic and Computation*, 31(2):523–572, 2021.