Possibilities for Awareness

Wes Holliday

ESSLLI 2023

Outline

Basic distinctions for awareness

Fully rational but partially unaware agents

From Burkhard Schipper's survey article on awareness:

> Formalized notions of awareness have been studied both in computer science and economics. In computer science the original motivation was mainly the modeling of agents who suffer from different forms of logical non-omniscience. The aim was to introduce logics that are more suitable than traditional logics for modeling beliefs of humans or machines with limited reasoning capabilities. In economics the motivation is similar but perhaps less ambitious. The goal is to model agents who may not only lack information but also conception. Intuitively, there is a fundamental difference between not knowing that an event obtained and not being able to conceive of that event. Despite such a lack of conception, agents in economics are still assumed to be fully rational in the sense of not making any errors in information processing. . . . (p. 1)

Two notions of awareness

Coarse-grained propositional awareness:

An agent is aware of a set-of-worlds proposition (a.k.a. event) or, more generally, a proposition in some Boolean algebra of propositions.

Cf. coarse-grained propositional knowledge/belief à la Lewis, Stalnaker, Aumann, etc.

Hyperintensional sentential awareness:

An agent is aware of a sentence φ in some language.

Hyperintensionality: where $[\![\varphi]\!]$ is the proposition expressed by φ, we may have $[\![\varphi_1]\!] = [\![\varphi_2]\!]$ yet the agent is aware of φ_1 and unaware of φ_2.

Plausible bridge principle: an agent is aware of a proposition E if and only if she is aware of some sentence φ such that $[\![\varphi]\!] = E$.

Two notions of awareness

Which principles we accept for awareness depends crucially on whether we have in mind the coarse-grained propositional notion or the hyperintensional sentential notion.

An unacceptable principle for coarse-grained propositional awareness

Adopting the coarse-grained view, let B be a Boolean algebra of propositions (e.g., the powerset of a state of states) and A_i a unary awareness operation on B.

First, observe that we should *not* require A_i to be monotonic with respect to the entailment order $\leq$ on propositions (i.e., $\subseteq$ in the case of the powerset algebra):

- for all $a, b \in B$, if $a \leq b$, then $A_i a \leq A_i b$.

Example: if an agent is aware of the proposition expressed by

'Ann and Bob will play a Nash equilibrium',

it does not follow that the student is aware of the proposition expressed by

'Ann and Bob will play a correlated equilibrium',

even though the first proposition entails the second.

An unacceptable principle for coarse-grained propositional awareness

Since $a \leq b$ is equivalent to $a = a \sqcap b$, monotonicity is equivalent to:

- for all $a, b \in B$, $A_i(a \sqcap b) \leq A_i b$.

A bad argument for this principle: "if you're aware of a conjunction, surely you're aware of the conjuncts!"

This is a bad argument because it switches to the hyperintensional sentential notion of awareness. The meet $a \sqcap b$ of a and b in the Boolean algebra is not somehow intrinsically a conjunction. We may have $a = a \sqcap b = c \sqcup d$, etc.

The upshot is that we must reject any treatment of coarse-grained propositional awareness that imposes either of the following equivalent principles:

- for all $a, b \in B$, if $a \leq b$, then $A_i a \leq A_i b$.
- for all $a, b \in B$, $A_i(a \sqcap b) \leq A_i b$.

Awareness in possibility frames

Key idea

To represent awareness using our possibility frames, for each agent $i \in I$, we will add a binary relation A_i (or, equivalently, a correspondence as is popular in econ).

We interpret $\omega A_i \nu$ as meaning "in possibility ω, agent i is aware of possibility ν."

We then say that i is aware in ω of a proposition E when the following condition holds at ω and its refinements:

- if i is aware of a possibility ν, then i is aware of any coarsest refinements of ν making E true and any coarsest refinements of ν making $\neg E$ true.

In other words, if you are aware of the proposition E, then you should be able to apply the E vs. $\neg E$ distinction starting from any possibility of which you are aware.

Review: possibility frames

For a poset $(\Omega, \sqsubseteq)$, an $E \subseteq \Omega$ is *regular open* (belongs to $\mathcal{RO}(\Omega, \sqsubseteq)$) iff it satisfies:

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

Definition

A (general) *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 $\neg$ in $\mathcal{RO}(\Omega, \sqsubseteq)$:

$$\neg E = \{\omega \in \Omega \mid \forall \omega' \sqsubseteq \omega\ \omega' \notin E\}.$$

Quasi-principal frames

Given a poset $(\Omega, \sqsubseteq)$ and $E \in \mathcal{RO}(\Omega, \sqsubseteq)$, 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.

Definition

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

In other words, any possibility that settles that E holds is a refinement of some coarsest possibility that settles that E holds, of which there are only finitely many.

Definition

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;
2. $\mathcal{A}_i : \Omega \to \wp(\Omega)$ is a correspondence satisfying the following conditions:
   2.1 awareness nonvacuity: $m \in \mathcal{A}_i(\omega)$;
   2.2 awareness expressibility: if $\nu \in \mathcal{A}_i(\omega)$, then $\downarrow\nu \in \mathcal{E}$;
   2.3 awareness persistence: if $\omega' \sqsubseteq \omega$, then $\mathcal{A}_i(\omega) \subseteq \mathcal{A}_i(\omega')$;
   2.4 awareness refinability: if $\nu \notin \mathcal{A}_i(\omega)$, then $\exists \omega' \sqsubseteq \omega \; \forall \omega'' \sqsubseteq \omega' \; \nu \notin \mathcal{A}_i(\omega'')$;
   2.5 awareness joinability: if $\nu \in \mathcal{A}_i(\omega)$ and $\nu_1, \ldots, \nu_n \in \mathcal{A}_i(\omega) \cap \downarrow\nu$, then $\max(\rho(\{\nu_1, \ldots, \nu_n\}) \cap \downarrow\nu) \subseteq \mathcal{A}_i(\omega)$.
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 $\forall \omega' \sqsubseteq \omega \; \forall \nu \in \mathcal{A}_i(\omega')$

$$\max(E \cap \downarrow\nu) \cup \max(\neg E \cap \downarrow\nu) \subseteq \mathcal{A}_i(\omega').$$

Closure of regular opens under awareness operations

Lemma

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* $\mathbf{A}_i(E) \in \mathcal{RO}(\Omega, \sqsubseteq)$*.*

Example 1

Example from Geanakoplos 1989 concerning 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.

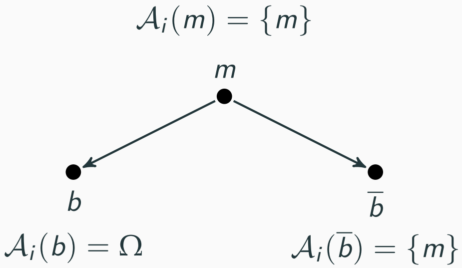

Where $Barks = \{b\}$, we have $\overline{b} \in \neg \mathbf{A}_i Barks$ and $\overline{b} \in \neg \mathbf{A}_i \neg \mathbf{A}_i Barks$.

Example 2

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 following game matrix with the middle row is greyed out:

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

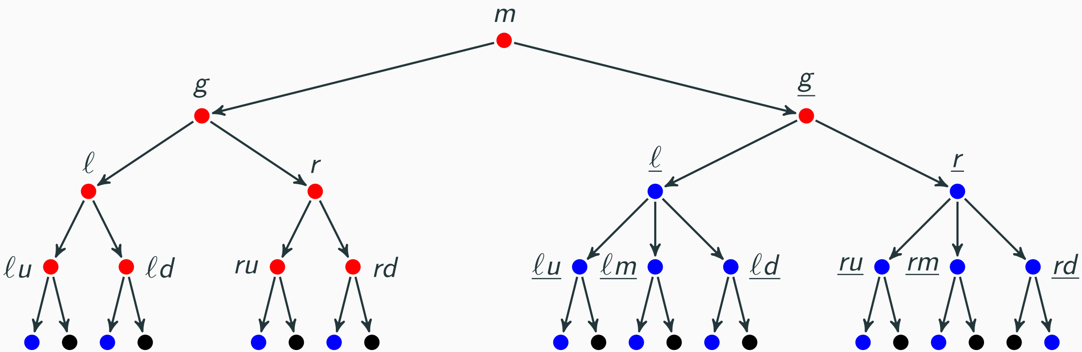

for each colored state ω, $\mathcal{A}_i(\omega) = \{\nu \in \Omega \mid \nu \text{ a red state}\}$

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

Where $\underline{Middle} = {\downarrow}\{\underline{\ell m}, \underline{rm}\}$, at each colored leaf ω, we have $\omega \in \neg \mathbf{A}_i \underline{Middle}$ and $\omega \in \neg \mathbf{A}_i \neg \mathbf{A}_i \underline{Middle}$. Yet where $Up = {\downarrow}\{\ell u, ru\}$, we have $\omega \in \mathbf{A}_i Up$.

Knowledge, belief, and awareness

Definition

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

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:
   2.1 $\mathcal{R}_i$-monotonicity: if $\omega' \sqsubseteq \omega$, then $\mathcal{R}_i(\omega') \subseteq \mathcal{R}_i(\omega)$;
   2.2 $\mathcal{R}_i$-regularity: $\mathcal{R}_i(\omega) \in \mathcal{RO}(\Omega, \sqsubseteq)$;
   2.3 $\mathcal{R}_i$-refinability: if $\nu \in \mathcal{R}_i(\omega)$, then $\exists \omega' \sqsubseteq \omega\ \forall \omega'' \sqsubseteq \omega'\ \exists \nu' \sqsubseteq \nu$: $\nu' \in \mathcal{R}_i(\omega'')$;
   2.4 epistemic factivity: $\omega \in \mathcal{K}_i(\omega)$;
   2.5 doxastic consistency: $\mathcal{B}_i(\omega) \neq \varnothing$;
   2.6 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 $\{\omega \in \Omega \mid \mathcal{R}_i(\omega) \subseteq E\} \in \mathcal{E}$.

Regular opens are closed under implicit knowledge and belief

Lemma

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*

$$\{\omega \in \Omega \mid \mathcal{R}_i(\omega) \subseteq E\} \in \mathcal{RO}(\Omega, \sqsubseteq).$$

Knowledge and belief

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 $\mathbf{B}_i$:

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

By the previous Lemma and the closure of $\mathcal{E}$ under binary intersection, $\mathcal{E}$ is also closed under $\mathbf{K}_i$ and $\mathbf{B}_i$.

Example 3

$$\mathcal{A}_i(m) = \{m\}$$

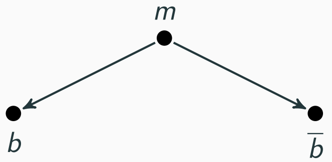

$$\mathcal{A}_i(b) = \Omega \qquad \mathcal{A}_i(\overline{b}) = \{m\}$$

Two pairs of knowledge and belief correspondences we might consider for Watson:

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

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

For any event E, $\mathbf{K}_i(E) = \mathbf{K}'_i(E)$ 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 $\overline{b}$ of the distinction between Bark and $\neg$Bark*, then he would know and believe $\neg$*Bark* in $\overline{b}$.

Example 4

A potential investor in a firm believes that he knows the firm is profitable, while being unaware of an unprecedented type of fraud that the firm is in fact using to cover up unprofitability (in f_1 below). $\mathcal{A}_i$, $\mathcal{B}_i$, $\mathcal{K}_i$ defined on handout...

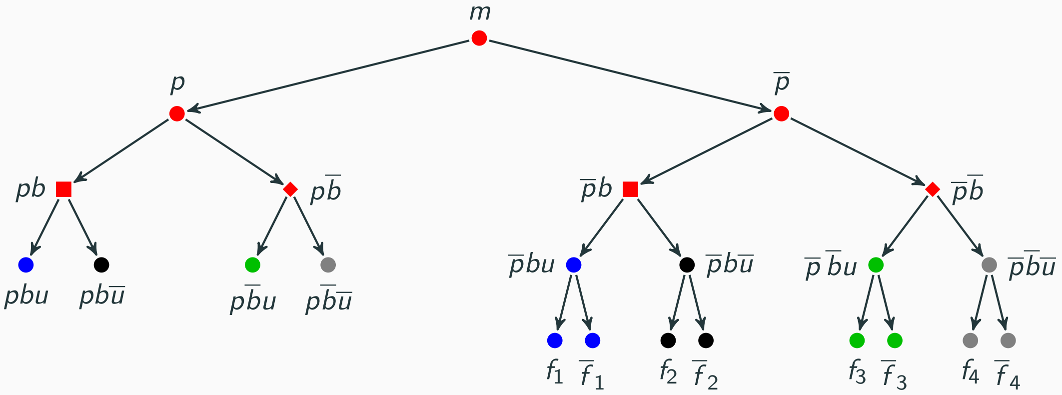

Example 4

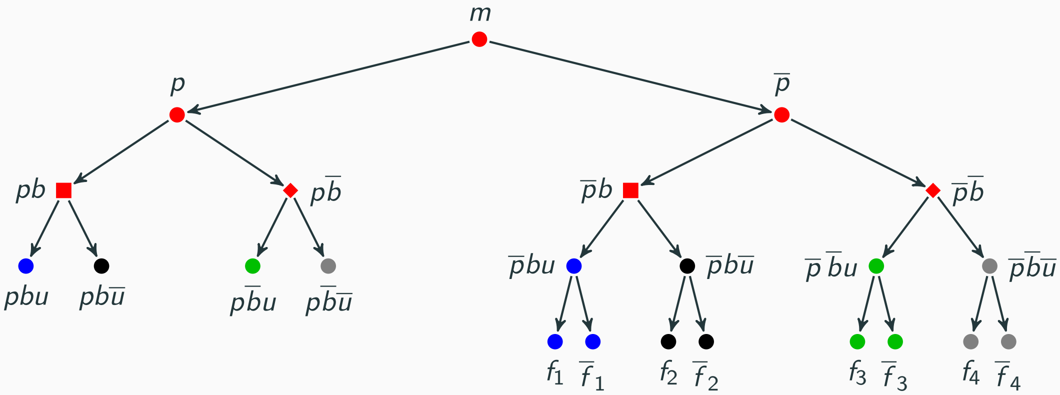

As analysts—or as other agents interacting with i—we might assign low probability to f_2 and $\overline{f}_2$, reflecting the view that i is unlikely to falsely believe he knows *Profit* when i is aware of the possibility of the sophisticated type of fraud.

Information-based beliefs

In general, say that a state ω determines that agent i has **information-based beliefs** 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 this is so, 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.

E.g., i has information-based beliefs at all leaves of the tree except for f_2 and $\overline{f}_2$.

In principle, to what extent false beliefs in a population of agents are correlated with unawareness of possibilities consistent with their information, versus mistakes in reasoning, exposure to misleading evidence, etc., could be investigated experimentally.

Representation of epistemic algebras

Definition

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 $\Box_i \in \{K_i, B_i\}$:

- $A_i 1 = 1$, $A_i a = A_i \neg a$, and $A_i a \sqcap A_i b \leq A_i(a \sqcap b)$;
- $K_i 1 = 1$ and $\Box_i a \sqcap \Box_i b \leq \Box_i(a \sqcap b)$;
- if $a \leq b$, then $\Box_i a \sqcap A_i b \leq \Box_i b$;
- $K_i a \leq a$ and $B_i 0 = 0$;
- $K_i a \leq B_i a$ and $B_i a \leq A_i a$.

The three axiom on awareness are also the key axioms in Fritz and Lederman 2015.

From frames to algebras

Proposition

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.*

From algebras to frames

Definition

Given an epistemic 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. $\mathcal{K}_i(F) = \{H \in \Omega \mid$ if $K_i a_1 \sqcup \cdots \sqcup K_i a_n \in F$, then $a_1 \sqcup \cdots \sqcup a_n \in H\}$;
5. $\mathcal{B}_i(F) = \{H \in \Omega \mid$ if $B_i a_1 \sqcup \cdots \sqcup B_i a_n \in F$, then $a_1 \sqcup \cdots \sqcup b_n \in H\}$.

From algebras to frames

Theorem

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

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

This 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.

Conclusion

Conclusion

I hope to have given some evidence that there is a mathematically elegant and philosophically reasonable way to add awareness to possibility frames.

But there are also a number of avenues for development...

Awareness of sentences

Though we have emphasized our event-based approach to unawareness, we could also develop a sentence-based approach on top of our event-based approach.

For example, we can turn a possibility frame into a possibility model for a propositional modal language by equipping the frame with a valuation V of atomic sentences such that $V(\mathrm{p}) \in \mathcal{E}$ for each atomic sentence p. One could even assume that for every $E \in \mathcal{E}$, there is some atomic sentence p for which $V(\mathrm{p}) = E$.

For an arbitrary sentence φ of the modal language, we could say that $\mathrm{A}_i\varphi$ is true at a possibility ω just in case for every subsentence ψ of φ, we have $\omega \in \mathbf{A}_i[\![\psi]\!]$, where $[\![\psi]\!]$ is the set of possibilities at which ψ is true. This would capture the idea that being aware of complex sentences such as $\varphi \wedge \psi$ or $\varphi \vee \neg\varphi$ requires being aware of φ and ψ.

Quantification

The literature on awareness also employs the languages of propositional modal logic with propositional quantifiers and first-order modal logic.

With propositional quantifiers, one may express that agent i knows that there is some event of which she is unaware: $\mathrm{K}_i \exists \mathrm{p}\, \mathrm{U}_i \mathrm{p}$

(cf. Ding 2020 on *modest* agents: $\mathrm{B}_i \exists \mathrm{p}(\mathrm{B}_i \mathrm{p} \wedge \neg \mathrm{p})$).

With first-order quantifiers, we can express unawareness of objects (unawareness of the proposition stating that the object exists? or something else?).

Possibility semantics for these extended languages works well and offers advantages over possible world semantics—see my "Possibility Semantics." Developing possibility semantics for such languages including awareness is a natural next step.

Probability

Probability can be added to our frames 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 (i.e., assigning subjective probability at least p) or i judging that E is at least as likely as F, will themselves belong to $\mathcal{RO}(\Omega, \sqsubseteq)$, so we can require that they belong to $\mathcal{E}$.

This would enable application in decision and game theory.

Next up

Time and the openness of the future...