Epistemic Pseudo-Agreement and Hyper-Domain Semantics for Epistemic Modals

Yichi (Raven) Zhang[1] and Eric Pacuit[2]

[1] Heinrich Heine Universität, Düsseldorf, Germany
yichi.zhang@hhu.de
[2] University of Maryland, College Park, U.S.A.
epacuit@umd.edu

Abstract

We propose a hyper-domain semantics to analyze the group reading of epistemic modals according to which the relevant quantificational base for modals is not represented by a single aggregated information state but rather by a set of information states. The resulting framework allows us to capture what we call "epistemic pseudo-agreements". Additionally, to further investigate the corresponding logic of our hyper-domain semantics, we devise a translation of the consequence relation in our hyper-domain semantics into a new logic called "Logic of Epistemic Disagreement" interpreted over possibility models.

1 Introduction

A key issue for any semantics for epistemic modals [6, 14, 11, 13] is to identify the relevant body of information that supplies a quantificational base for the modal operator. Epistemic disagreements are vivid illustrations of this:

(1) John: Alice might be at the party.
Mary: No, that's not true. She's just tested positive for Covid this morning.

According to [12], there are three different readings of the *might*-claim uttered by John. On John's solipsistic reading, John's assertion of the *might*-claim is warranted since its prejacent is compatible with what John knows. On Mary's solipsistic reading, the quantificational base is supplied by what Mary knows, and since she knows that Alice is not at the party, this is, in part, what warrants her rejection of John's *might*-claim. Finally, there is the group reading in which the quantificational base for the *might*-claim is derived from the *distributed knowledge* of John and Mary. The distributed knowledge of John and Mary is what is known after they each share everything that they know (about who is at the party)—i.e., the distributed knowledge is derived from the intersection of the sets of worlds representing John's private information and the set of worlds representing Mary's private knowledge. Since Alice being at the party is not consistent with John and Mary's pooled information, using the distributed knowledge of John and Mary as the quantificational base of the *might*-claim explains Mary's rejection of the *might*-claim and predicts that John will retract his claim.

One problem with using distributed knowledge for the group reading of epistemic modals is that it incorrectly predicts disagreement about *might*-claims in what we call *epistemic pseudo-agreement* dialogues:

(2) (*Context: John knows that Alice is at the party but has no information on Bob; Mary on the other hand knows that either Alice or Bob is at the party but not both.*)
John: Bob might be at the party.
Mary: I agree.

John's knowledge can be represented by two possible worlds, one in which both Alice and Bob are at the party and one in which Alice is at the party but Bob is not. Mary's knowledge can be represented by two possible worlds, one in which Alice is at the party but Bob is not at the party, and one in which Bob is at the party but Alice is not at the party. The intersection of these two sets of possible worlds consists of a single possible world in which Alice is at the party but Bob is not at the party. Nevertheless, even though Bob being at the party is inconsistent with the distributed knowledge of John and Mary, Mary agrees with John's *might*-claim.

Von Fintel and Gillies [12] proposed the following additional constraint on the group information state to handle the above problem of pseudo-agreement:

> **Defeasible Closure:** If [the hearer] H knows that φ is compatible with what x knows, for each $x \in G$, then it is reasonable for H to defeasibly infer that φ is compatible with what [the group] G knows.

Assuming that Mary knows what information John has about who is at the party, she can defeasibly infer that Bob being at the party is compatible with what the group knows. So, Defeasible Closure does address the problems that arise from naively using distributed knowledge for group readings of epistemic modals. However, since the fix offered by defeasible closure requires the discourse participants to deliberately draw the inference that φ is compatible with what other people know, its application is limited. For instance, suppose that in (2), Mary has good reasons to suspect that John may not truthfully report what he really knows. In this case, she will not defeasibly infer that Bob's being at the party is compatible with what John knows, even if, in this particular instance, John is indeed being truthful and his information is compatible with Bob being at the party. Nonetheless, it seems she can still assent to John's *might*-claim in (2).

One response to the above discussion is to simply dispense with the group reading of epistemic modals. However, we argue that there are utterances that only make sense with a group reading of epistemic modals. Consider the following sentence adapted from DeRose [2]:

(3) I don't know whether John might have cancer; only the doctor knows.

This sentence is perfectly natural in a context where the doctor has conducted a screening test on John that will conclusively rule out cancer if the test is negative, but will not conclusively affirm that he has cancer if the test is positive. However, this sentence would not make sense if the relevant body of information for the *might*-claim is based on the speaker's knowledge. One plausible interpretation of (3) is that the *might*-claim is based on the doctor's information as illustrated by the following paraphrase:

(4) I don't know whether John's having cancer is compatible with what the doctor knows; only the doctor knows.

The problem with this reading, however, is that it predicts that (5) and (6) have the same meaning, and thus incorrectly predicts that (5) is felicitous.

(5) # I know John doesn't have diabetes, but I don't know whether John might have cancer and diabetes.

(6) I know John doesn't have diabetes, but I don't know whether John's having cancer and diabetes is compatible with what the doctor knows.

Correctly predicting that (5) is infelicitous seems to require a group reading for the *might*-claim taking into account both the doctor's and the speaker's information. Thus, a group reading of the epistemic modals is needed for (4) to be a correct paraphrase of (3) and for (5) to be infelicitous. In this paper, rather than using a single information state to represent group knowledge (e.g., the distributed knowledge of the group), we propose to represent the group information state as a set of information states each corresponding to the private information of a member of the group.

2 Preliminary: Domain Semantics for Epistemic Modals

Our formal model builds on Yalcin's domain semantics [14]. One of the main motivations of domain semantics is to capture epistemic contradictions:

(7) # John doesn't have cancer but he might have cancer.

(8) # Mary supposes that John doesn't have cancer but he might have cancer.

Sentences of the form $\neg p \wedge \Diamond p$ and $p \wedge \Diamond \neg p$ (as well as the variants that swap the order of the conjuncts) intuitively feel contradictory even in an embedded environment (as in (8)). Domain semantics predicts that both (7) and (8) are contradictory. Formulas are interpreted at pairs $\langle w, s\rangle$ where w is a possible world and s is an information state (a set of possible worlds). Given a world-information-state pair $\langle w, s\rangle$, the atomic propositions and the Boolean connectives are interpreted using the possible world w and a valuation function as usual, and modalities quantify over the worlds in the information state s: for instance, $w, s \Vdash \Diamond p$ if, and only if, there exists $w' \in s$ such that $w', s \Vdash p$. That is, $\Diamond p$ tests that the prejacent is compatible with the given body of information specified by s. Say that an information state s *accepts* (or *supports*) a formula φ, denoted $s \vDash \varphi$, if, and only if, $w, s \Vdash \varphi$ for all $w \in s$. Then, sentences such as (7) are unassertable since any information state that accepts $\neg p$ will fail to accept $\Diamond p$, and vice versa.

To see why sentences such as (8) are unassertable, note that the attitude verb 'i suppose' shifts the information state to an information state representing what the agent i supposes: $w, s \Vdash i \text{ supposes } \varphi$ if, and only if, $\forall w' \in S_i^w : w', S_i^w \Vdash \varphi$, where S_i^w is the body of information that corresponds to what i supposes at w. It is not difficult to see that this will predict that sentences such as (8) are unassertable. However, as Yalcin noted [14], the semantics sketched above has difficulty explaining why sentences such as (3) of the form 'I don't know whether might p' are *not* contradictory. Furthermore, as pointed by Dorr and Hawthorne [3], shifting to the knowledge state of someone other than the speaker does not help since for all information states s, one of the following must be true: (i) for all $w' \in s$, $w', s \Vdash \Diamond\varphi$ or (ii) for all $w' \in s$, $w', s \Vdash \neg\Diamond\varphi$. Hence, the speaker cannot both fail to know that might p and know that not might p. We address this problem with domain semantics by moving to hyper-domains (i.e., sets of information states) as explained in the next section.

3 Hyper-Domain Semantics for Epistemic Modals

Suppose that $W \neq \varnothing$ and $w \in W$ and $\Sigma \subseteq \wp(W)$. Formulas are interpreted at pairs $\langle w, \Sigma\rangle$, where w is a possible world and Σ is a set of information states (called a hyper-domain). Each element of Σ represents the information possessed by a relevant party in a discourse. We assume that for all $\langle w, \Sigma\rangle$, we have that $w \in \bigcap \Sigma$, so that all discourse participants consider the actual world possible. For an atomic formula p, we have that $w, \Sigma \Vdash p$ iff p is true at w (according to some fixed valuation function mapping atomic propositions to sets of worlds). Compound formulas are evaluated as follows:

$w, \Sigma \Vdash \neg\varphi$ iff $\forall s \in \Sigma : w, \{s\} \nVdash \varphi$; $w, \Sigma \Vdash \varphi \to \psi$ iff $\forall s \in \Sigma : w, \{s\} \nVdash \varphi$ or $w, \{s\} \Vdash \psi$;
$w, \Sigma \Vdash \varphi \wedge \psi$ iff $\forall s \in \Sigma : w, \{s\} \Vdash \varphi$ & $w, \{s\} \Vdash \psi$; $w, \Sigma \Vdash \Diamond\varphi$ iff $\forall s \in \Sigma\ \exists w' \in s$: $w', \{s\} \Vdash \varphi$;
$w, \Sigma \Vdash \varphi \vee \psi$ iff $\forall s \in \Sigma : w, \{s\} \Vdash \varphi$ or $w, \{s\} \Vdash \psi$; $w, \Sigma \Vdash \Box\varphi$ iff $\forall s \in \Sigma\ \forall w' \in s$: $w', \{s\} \Vdash \varphi$.

It is easy to see that the usual inter-definability between the Boolean connectives and the duality between $\Diamond$ and $\Box$ are satisfied. Note that a *might*-claim ($\Diamond\varphi$) is satisfied when every discourse participant deems the prejacent possible, and that the negation of a *might*-claim is satisfied when every discourse participant deems its prejacent impossible. The analogue of Yalcin's notion of *acceptance* is defined as follows: Σ accepts φ, denoted $\Sigma \vDash \varphi$, iff $w, \Sigma \Vdash \varphi$ for all $w \in \bigcup \Sigma$. Now, there is an epistemic disagreement in a hyperdomain Σ when $\Sigma \nvDash \Diamond\varphi$ and $\Sigma \nvDash \neg\Diamond\varphi$. We thus capture the epistemic disagreement in (1) since the prejacent is compatible with John's private information s_j but not with Mary's s_m, and so the the hyper-dommain $\{s_j, s_m\}$ accepts neither the *might*-claim nor its negation. We also correctly predict the epistemic pseudo-agreement in (2): since both John's and Mary's private information states contain a world where Bob is at the party, the *might*-claim is accepted. Finally, semantic consequence, denoted $\vDash^{hd}$, is defined as preservation of acceptance at all hyper-domains: $\Gamma \vDash^{hd} \varphi$ iff for all W, $V : \mathsf{At} \to \wp(W)$ and $\Sigma \subseteq \wp(W)$, if $\Sigma \vDash \gamma$ for all $\gamma \in \Gamma$, then $\Sigma \vDash \varphi$. Hence, we have that $\neg p \wedge \Diamond p \vDash^{hd} \bot$ as every hyper-domain that accepts $\neg p$ will fail to accept $\Diamond p$ (similarly for the other variants of epistemic contradictions). In addition, as with domain semantics, our hyper-domain semantics collapses strings of modals to the innermost modality: for instance, we have that $\Diamond \Box \Diamond p \vDash^{hd} \Diamond p$.

4 Logic of Epistemic Disagreement (LED)

To investigate the logic of our hyper-domain semantics, we take inspiration from Schulz's [9] translation of the consequence relation of domain semantics into the modal logic **S5**. The target logical system for our translation of the semantic consequence for hyper-domain semantics is a bimodal logic interpreted in models based on *partial states*, also called *possibilities* or *situations*, where the truth value of a formula may be unsettled [5, 10, 4]. The two modalities included in our language are the usual epistemic modality $\Box$ representing "must" and a "consensus" modality $[C]$ indicating that the group collectively agrees on the prejacent. For simplicity, we assume that $[C]$ is always the outermost modality and does not iterate—i.e., it does not embed under either $[C]$ or $\Box$. More formally, the language $\mathcal{L}_{LED}$ is defined recursively as follows: $\varphi ::= \psi \mid [C]\psi \mid -\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \supset \varphi$, where $\psi ::= p \mid -\psi \mid \psi \wedge \psi \mid \psi \vee \psi \mid \psi \supset \psi \mid \Box\psi$ for $p \in Atoms$ (the set of atomic propositions). We define $\Diamond\psi$ as $-\Box-\psi$ and $\langle C\rangle\psi$ as $-[C]-\psi$.

Formulas of $\mathcal{L}_{LED}$ are interpreted in a possibility model $\mathcal{M} = \langle S, \sqsubseteq, C, V\rangle$ where S is a nonempty set, $\sqsubseteq$ is a partial order on S (reflexive, transitive and anti-symmetric relation on S), C is a binary relation on S, and V is a valuation that assigning to ever atomic formula a nonempty subset U of S satisfying persistence and refinability [4]:

Persistence: if $x \in U$ and $y \sqsubseteq x$, then $y \in U$;

Refinability: if $x \notin U$, then $\exists y \sqsubseteq x\ \forall z \sqsubseteq y : z \notin U$.

Intuitively, y is a refinement of x ($y \sqsubseteq x$) when y settles all the the formulas that x does and possibly more. Persistence guarantees that every atomic proposition that is settled true/false in x is also settled true/false in y. Refinability guarantees that if an atomic proposition not settled true at x, then there is a refinement that settles it false. The intended interpretation of $x\ C\ y$ is that y is a relevant body of information that corresponds to some discourse participant's private information in the *context* represented by x. We impose the following constraint on $\sqsubseteq$ and C mirroring the constraint on hyper-domains requiring that every information state in the hyper-domain contains the actual world: for all $x, y \in S$, if $x\ C\ y$, then $x \sqsubseteq y$. That is, if x represents the context of the conversation, then every participant's information at x must be able to be refinable into x.

Truth of a formula $\varphi \in \mathcal{L}_{LED}$ at a state x in a model $\mathcal{M}$ is defined as follows: For any atomic formula p, $\mathcal{M}, x \Vdash p$ iff $x \in V(p)$.

$\mathcal{M}, x \Vdash -\varphi$ iff $x \nVdash \varphi$
$\mathcal{M}, x \Vdash \varphi \wedge \psi$ iff $\mathcal{M}, x \Vdash \varphi$ and $\mathcal{M}, x \Vdash \psi$
$\mathcal{M}, x \Vdash \varphi \vee \psi$ iff $\mathcal{M}, x \Vdash \varphi$ or $\mathcal{M}, x \Vdash \psi$
$\mathcal{M}, x \Vdash \varphi \supset \psi$ iff either $\mathcal{M}, x \nVdash \varphi$ or $\mathcal{M}, x \Vdash \psi$
$\mathcal{M}, x \Vdash \Box\varphi$ iff $\forall y \sqsubseteq x$: $\mathcal{M}, y \Vdash \varphi$
$\mathcal{M}, x \Vdash [C]\varphi$ iff $\forall y$: if $x\ C\ y$ then $\mathcal{M}, y \Vdash \varphi$

Note that, unlike in standard possibility semantics where persistence holds for all formulas, $\Diamond$-formulas may fail to be persistent: a *might*-claim may hold at a situation but fail to hold at one of its refinements. The definitions of validity on a frame $\langle S, \sqsubseteq, C\rangle$ and a semantic consequence $\vDash^{led}$ are defined as usual.

We provide a translation scheme for a portion of the language of hyper-domain semantics in which (i) there are no iterated modalities (i.e., no formulas of the form $\Box\varphi$ where φ contains a modality) and (ii) there are no modalities in the scope of a disjunction or implication (e.g., no formula of the form $\Diamond\varphi \vee \psi$ or $\Box\varphi \supset \psi$). Suppose that φ is such a formula of hyper-domain semantics, the formula $\tau(\varphi) \in \mathcal{L}_{LED}$ is defined in two steps: (1) transform φ into a logically equivalent formula φ' where all negations scope below any modal; and (2) define $\tau'(\varphi')$ recursively: (i) $\tau'(p) = p$; (ii) $\tau'(\neg\varphi') = \Box{-}\tau'(\varphi')$; (iii) $\tau'(\varphi' \wedge \psi') = \tau'(\varphi') \wedge \tau'(\psi')$; (iv) $\tau'(\varphi' \vee \psi') = \Box\Diamond(\tau'(\varphi') \vee\!\!\vee\, \tau'(\psi'))$; (v) $\tau'(\varphi' \to \psi') = \Box(\tau'(\varphi') \supset \tau'(\psi'))$; (vi) $\tau'(\Box\varphi') = \Box\tau'(\varphi')$; and (vii) $\tau'(\Diamond\varphi') = \Diamond\tau'(\varphi')$. For example, $\tau(\neg\Box p) = \tau'(\Diamond\neg p) = \Diamond\Box{-}p$. Our first main result is τ preserves validity: for all formulas φ and sets of formulas Γ (in the restricted language), $\Gamma \vDash^{hd} \varphi$ iff $[C]\tau(\Gamma) \vDash^{led} [C]\tau(\varphi)$, where $[C]\tau(\Gamma)$ means $[C]\tau(\gamma)$ for all $\gamma \in \Gamma$ (see the Appendix).

The axiomatization of LED for language $\mathcal{L}_{LED}$ includes the following: (i) all axiom schemas from classical logic; (i) the **S4** axioms for $\Box$; (iii) the axiom scheme corresponding to the refinability $\Box\Diamond\varphi \supset \varphi$; (iv) a limited version of persistence $\alpha \supset \Box\alpha$, where α is a *positive formula*; (vi) the **KD** axioms for $[C]$; and (vii) $\varphi \supset [C]\Diamond\varphi$. The last axiom, called *entertainability*, says that if φ is true, then everyone should agree that φ is at least possible. This axiom corresponds to the aforementioned condition on C stating that if $x\ C\ y$, then $x \sqsubseteq y$. Importantly, the T axiom is not satisfied for $[C]$, in particular, sentences of the form $[C]\Diamond p \supset \Diamond p$ may not be true at a partial state. This is what happens in pseudo-agreement situations, such as (2), in which it is possible for everybody at a situation x to agree that $\Diamond p$ while p is actually false at x, and so is $\Diamond p$ is false at x. The logic is closed under Modus Ponens and necessitation for both $\Box$ and $[C]$. Some details of the proof of completeness can be found in the Appendix.

We conclude this section by discussing the consequences of imposing additional constraints concerning the interaction of C and $\sqsubseteq$. A full discussion of these constraints is left for the full paper. Given the restriction on the language that $[C]$ cannot be embed under $\Box$, the formulas corresponding to the constraints discussed below cannot be stated in the language $\mathcal{L}_{LED}$, but they do help to refine the notion of *context* we employ. We start with the following two constraints from [4]:

R-rule: if $x' \sqsubseteq x$, $x'\ C\ y'$, and $y' \between z$, then $\exists y : xCy$ and $y \between z$.
R⇒win: if $x\ C\ y$, then $\forall y' \sqsubseteq y \exists x' \sqsubseteq x \forall x'' \sqsubseteq x' \exists y'' : x''\ C\ y''$ and $y'' \between y'$.

For any two possibilities x and y, $x \between y$ means that x and y have a common refinement. The above two conditions guarantee that sets of partial states associated with $[C]$-formulas satisfy persistence and refinability. Persistence of $[C]$-formulas means that if every participant at a context x agrees on some formula φ (e.g., $\Diamond p$), then at all refinements of x, everybody should still agree on φ. This requires that the group of relevant discourse participants remain fixed across refinements. It also means that a refinement of a context is not accompanied by the discourse participants' gaining new information. Otherwise, there could be some discourse participant's information that is compatible with p, but no longer compatible after the discourse participant acquires new information (e.g., that p is false). Thus, our notion of context refinement is different from the standard notion of context *update*. Our proposal is to view a refinement of a context as an update with new relevant issues or questions under discussion [8, 7]. That is, if y is a refinement of x, then y may settle additional issues that are deemed irrelevant at x. For example, suppose that Alice, Bob, and Carla are at the party. If x represents a context in which the question under discussion only concerns the whereabouts of Alice and Bob, then a refinement of x may represent a context where the question under discussion concerns the whereabouts of all three individuals.

Finally, we consider three additional constraints from [10]. The first two constraints that one might impose are:

C(ontinuous): If $x' \sqsubseteq x$ and $x' \mathrel{C} y'$, then $\exists y : x \mathrel{C} y$ and $y' \sqsubseteq y$.

O(pen): If $x \mathrel{C} y$ and $y' \sqsubseteq y$, then $\exists x' \sqsubseteq x : x' \mathrel{C} y'$.

While it is clear that **C** should be satisfied given our intended interpretation of C and our understanding of context discussed above, imposing both **O** and our **E**(ntertainability) condition (if $y \mathrel{C} x$, then $x \sqsubseteq y$) leads to a counterintuitive constraint on our models. Given **O** and **E**, it follows that if $x \sqsubseteq y$ and $y' \sqsubseteq y$, then $\exists x' : x' \sqsubseteq x$ and $x' \sqsubseteq y'$. That is, every two refinements of a single situation have a common refinement. This means that we cannot impose both persistence and refinability. For this reason, we are tempted to drop the constraint **O**. However, for now we just highlight the tension between **O** and **E** and leave additional investigation to future work.

The third constraint that might be imposed is:

Conv(ergence): If $x' \sqsubseteq x$ and $x \mathrel{C} y$, then $\exists y' \sqsubseteq y : x' \mathrel{C} y'$.

Conv guarantees the persistence of formulas of the form $\langle C \rangle \varphi$ (provided that φ is persistent). This squares with the current interpretation of context refinement: if someone endorses that φ at x then the same person should still endorse φ at a refinement of x' of x. In our ongoing work, we will explore other interactions between C and $\sqsubseteq$ and lift the syntactic restriction currently imposed on $\mathcal{L}_{LED}$ leading to a translation of the full language of hyper-domain semantics.

5 Adding Knowledge Operator

In this section, we define a knowledge operator in hyper-domain semantics. Rather than representing an agent's knowledge as a set of worlds, we use a hyper-domain. Consider (3) again: in order to make "I don't know that John might have cancer" and "I don't know that it's not the case that John might have cancer" both true (i.e., $w, \Sigma \Vdash \neg K_i \Diamond p$ and $w, \Sigma \Vdash \neg K_i \neg \Diamond p$), we assume that the speaker's knowledge $\mathcal{K}_i^w$ contains two information states s_1 and s_2 where s_1 is compatible with John's having cancer whereas s_2 is not. In this case, the speaker internalizes two different epistemic perspectives and her knowledge state is indeterminate between them (cf. [13]). The knowledge operator K_i is defined as follows:

- $w, \Sigma \Vdash K_i \varphi$ iff $\forall w' \in \bigcup \mathcal{K}_i^w : w', \Sigma \cup \mathcal{K}_i^w \Vdash \varphi$, where $\mathcal{K}_i^w$ is the hyper-domain that corresponds to what i knows at w.

Note that K_i shifts the hyper-domain for evaluating its prejacent to the combined state $\Sigma \cup \mathcal{K}_i^w$ rather than agent i's knowledge state. This ensures that the truth axiom $K_i \varphi \supset \varphi$ is accepted at any hyper-domain. In particular, we have that $K_i \Diamond p \supset \Diamond p$ is accepted at any hyper-domain. This is not true in standard domain semantics since p may be compatible with what the knowledge ascribee knows, but incompatible with what the interlocutor knows (see [1] for discussion of a related problem).

Now, we can unpack $w, \Sigma \Vdash \neg K_i \Diamond p$ and $w, \Sigma \Vdash \neg K_i \neg \Diamond p$ into the following conditions:

- $w, \Sigma \Vdash \neg K_i \Diamond p$ iff $\forall s \in \Sigma . \exists s' \in \{s\} \cup \mathcal{K}_i^w . \forall w' \in s' : w', \{s'\} \nVdash p$
- $w, \Sigma \Vdash \neg K_i \neg \Diamond p$ iff $\forall s \in \Sigma . \exists s' \in \{s\} \cup \mathcal{K}_i^w . \exists w' \in s' : w', \{s'\} \Vdash p$

Since $\mathcal{K}_i^w$ contains two information states s_1 and s_2 each of which witnesses one of the two conditions, we correctly predict $w, \Sigma \Vdash \neg K_i \Diamond p$ and $w, \Sigma \Vdash \neg K_i \neg \Diamond p$ to be jointly satisfiable.

6 Conclusion

In this abstract, we presented a hyper-domain semantics that offers a straightforward account of the group reading of epistemic modals. Combining the idea of representing the group information directly using a set of information states with a Yalcin-style domain semantics, we are also able to define a knowledge operator that resolves certain problems associated with the standard domain semantics. We then show that our hyper-domain semantics can be axiomatized via the translation into a new bimodal possibility semantics.

References

[1] Bob Beddor and Simon Goldstein. Mighty knowledge. *The Journal of Philosophy*, 118:229 – 269, 2021.

[2] Keith DeRose. Epistemic possibility. *Philosophical Review*, 100:581 – 605, 1991.

[3] Cian Dorr and John Hawthorne. Embedding epistemic modals. *Mind*, 122:867 – 913, 2013.

[4] Wesley H. Holliday. Possibility semantics. In Melvin Fitting, editor, *Selected Topics from Contemporary Logics*, pages 108–130. College Publications, 2021.

[5] Lloyd Humberstone. From worlds to possibilities. *Journal of Philosophical Logic*, 10:313 – 339, 1981.

[6] Angelika Kratzer. *Modals and Conditionals: New and Revised Perspectives.* Oxford University Press, Oxford, 2012.

[7] Craige Roberts. Information structure in discourse. OSU Working Papers in Linguistics, volume 49, J. Yoon and A. Kathol (eds.), 1996.

[8] Craige Roberts. Information structure: Towards an integrated formal theory of pragmatics. *Semantics & Pragmatics*, 5:1 – 69, 2012.

[9] Moritz Schulz. Epistemic modals and informational consequence. *Synthese*, 174:385 – 395, 2010.

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

[11] Frank Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25:220 – 260, 1996.

[12] Kai von Fintel and Anthony Gillies. 'Might' made right. In Andy Egan and Brian Weatherson, editors, *Epistemic Modality*, pages 108–130. Oxford University Press, 2011.

[13] Malte Willer. Dynamics of epistemic modality. *Philosophical Review*, 122:45 – 92, 2013.

[14] Seth Yalcin. Epistemic modals. *Mind*, 116:983–1026, 2007.

A Formal Results

Theorem A.1. *The translation scheme τ preserves validity in the following way: $\Gamma \vDash^{hd} \varphi$ iff $[C]\tau(\Gamma) \vDash^{led} [C]\tau(\varphi)$ where $[C]\tau(\Gamma)$ abbreviates $[C]\tau(\gamma)$ for all $\gamma \in \Gamma$.*

Proof. We prove this by first proving **Proposition A.3**: given a pointed hyper-domain model $\langle W, V\rangle, \Sigma$ and its corresponding pointed possibility model $\mathcal{M}, x$, the following holds: $\Sigma \vDash \varphi$ iff $\mathcal{M} = \langle S, C, \sqsubseteq, V'\rangle, x \Vdash [C]\tau(\varphi)$. The correspondence is defined as follows: given a pair $\langle W, V\rangle, \Sigma$, we can generate a pair $\mathcal{M}, x$ by making (i) $x = \bigcap \Sigma$; (ii) $S = \wp(W) - \{\varnothing\}$; (iii) $y' \sqsubseteq y$ iff $y' \subseteq y$; (iv) xCy iff $y \in \Sigma$; (v) $V'(p) = \{s \in S \mid s \subseteq V(p)\}$. We can check that the resulting model is indeed of the right kind; for instance, for all $x, y \in S$, we have that if $x\ C\ y$, then $x \sqsubseteq y$ since $\bigcap \Sigma \subseteq y$ for any $y \in \Sigma$. Borrowing a term from [10], let us call the resulting $\mathcal{M}, x$ the posibilization of $\langle W, V\rangle, \Sigma$. As for the opposite direction going from $\mathcal{M}, x$ to $\langle W, V\rangle, \Sigma$, it suffices to show that $\mathcal{M}, x$ is isomorphic to the possibilization of $\langle W, V\rangle, \Sigma$. Given **Proposition A.3**, **Theorem A.1** immediately follows. Now to prove **Proposition A.3**, we will first prove **Proposition A.2**, the special case where Σ is a singleton hyper-domain. □

Proposition A.2. *Suppose that $W \neq \varnothing$, $V : Atoms \to \wp(W)$ and $s \subseteq W$. Let $\mathcal{M} = \langle S, C, \sqsubseteq, V\rangle$ be a model where (i) $x = \bigcap\{s\} = s$; (ii) $S = \wp(W) \setminus \{\varnothing\}$; (iii) $y' \sqsubseteq y$ iff $y' \subseteq y$; (iv) xCy iff $y \in \{s\}$; (v) $V'(p) = \{s \in S \mid s \subseteq V(p)\}$. Then, $\{s\} \vDash \varphi$ iff $\mathcal{M}, x \Vdash \tau(\varphi)$, where $x = s$.*

Proof. First, since $x = s$, *(iv)* implies that $x\ C\ x$. Let φ' be a formula equivalent to φ but with negations all scoping below modals, and all strings of modals collapsed. Then we can show $\{s\} \vDash \varphi$ iff $\mathcal{M}, x \Vdash \tau'(\varphi')$ by an induction on the structure of φ' by using induction hypotheses: for all $x' \in \wp(W) \setminus \{\varnothing\}, \{x'\} \vDash \alpha$ iff $\mathcal{M}, x' \Vdash \tau'(\alpha)$. Consider the left to right direction first.

1. Let φ' be of the form p: then $\tau'(p) = p$. Since $\{s\} \vDash \varphi'$, we have $\forall w \in \bigcup\{s\} : w, \{s\} \Vdash p$, which means $\forall w \in s : w \in V(p)$. Hence, $s \subseteq V(p)$ which means $\mathcal{M}, x \Vdash p$.

2. Let φ' be of the form $\neg\alpha$: then $\tau'(\neg\alpha) = \Box - \tau'(\alpha)$. Since $\{s\} \vDash \neg\alpha$, it follows that $\forall w \in s : w, \{s\} \Vdash \neg\alpha$ and thus $\forall w \in s : w, \{s\} \nVdash \alpha$. Now, assume $\mathcal{M}, x \nVdash \Box - \tau'(\alpha)$ for contradiction. Then, $\exists x' \sqsubseteq x : \mathcal{M}, x' \Vdash \tau'(\alpha)$. Given that $x' \sqsubseteq x$ iff $x' \subseteq x$ and $x = s$, this just amounts to $\exists x' \subseteq s : \mathcal{M}, x' \Vdash \tau'(\alpha)$. By the induction hypothesis, we derive $\exists x' \subseteq s : \{x'\} \vDash \alpha$. Then, $\exists x' \subseteq s \forall w' \in x' : w', \{x'\} \Vdash \alpha$. Now recall that φ' is a formula where all negations scope below modals, this means that α does not contain any modal operators. Given this, $\exists x' \subseteq s \forall w' \in x' : w', \{x'\} \Vdash \alpha$ contradicts $\forall w \in s : w, \{s\} \nVdash \alpha$ since the truth of Boolean formulas do not depend on the hyper-domain.

3. Let φ' be of the form $\alpha \wedge \beta$: then $\tau'(\alpha \wedge \beta) = (\tau'(\alpha) \wedge \tau'(\beta))$. Given $\{s\} \vDash \alpha \wedge \beta$, $\mathcal{M}, x \Vdash (\tau'(\alpha) \wedge \tau'(\beta))$ follows straightforwardly from the induction hypothesis.

4. Let φ' be of the form $\alpha \vee \beta$: then $\tau'(\alpha \vee \beta) = \Box\Diamond(\tau'(\alpha) \vee\!\!\vee \tau'(\beta))$. Given $\{s\} \vDash \alpha \vee \beta$, it follows that $\forall w \in s : w, \{s\} \Vdash \alpha \vee \beta$, which means $\forall w \in s : w, \{s\} \Vdash \alpha$ or $w, \{s\} \Vdash \beta$. We also have $\forall w \in s$: if $w, \{s\} \nVdash \alpha$ then $w, \{s\} \Vdash \beta$, which entails if $\forall w \in s : w, \{s\} \nVdash \alpha$ then $\forall w \in s : w, \{s\} \Vdash \beta$, which amounts to if $\{s\} \vDash \neg\alpha$ then $\{s\} \vDash \beta$. From $\{s\} \vDash \beta$ by induction hypothesis, we have $\mathcal{M}, s \vDash \tau'(\beta)$; from $\{s\} \vDash \alpha$, since $\neg\alpha$ does not contain any modal operators, by what we have just shown in step (2) above, we have $\mathcal{M}, s \Vdash \Box - \tau'(\alpha)$. Hence, we can rewrite the conditional as: if $\mathcal{M}, s \Vdash \Box - \tau'(\alpha)$ then $\mathcal{M}, s \vDash \tau'(\beta)$. First assume $\mathcal{M}, s \Vdash \Box - \tau'(\alpha)$, then we immediately derive $\mathcal{M}, s \Vdash \tau'(\beta)$. Since β does not contain any modals, by persistence, $\forall s' \sqsubseteq s : \mathcal{M}, s' \Vdash \tau'(\beta)$. It then follows that $\mathcal{M}, s \Vdash \Box\Diamond(\tau'(\alpha) \vee\!\!\vee \tau'(\beta))$. Now assume $\mathcal{M}, s \nVdash \Box - \tau'(\alpha)$ instead. Then, $\exists s' \sqsubseteq s : \mathcal{M}, s' \Vdash \tau'(\alpha)$, which means $\exists s' \sqsubseteq s : \mathcal{M}, s' \Vdash \tau'(\alpha) \vee\!\!\vee \tau'(\beta)$. Since α does not contain modals, by persistence, we have $\mathcal{M}, s \Vdash \Box\Diamond(\tau'(\alpha) \vee\!\!\vee \tau'(\beta))$ again.

5. The proof for $\rightarrow$ is analogous to the proof for $\vee$.

6. Let φ' be of the form $\Diamond\alpha$: then $\tau'(\Diamond\alpha) = \Diamond\tau'(\alpha)$. Since $\{s\} \vDash \Diamond\alpha$, it follows that $\forall w \in s : w, \{s\} \Vdash \Diamond\alpha$ which just means $\exists w \in s : w, \{s\} \Vdash \alpha$. Assume $\mathcal{M}, x \nVdash \Diamond\tau'(\alpha)$ for contradiction, from which it follows that $\forall x' \subseteq x : \mathcal{M}, x' \nVdash \tau'(\alpha)$. By the induction hypothesis, we derive $\forall x' \subseteq x : \{x'\} \nvDash \alpha$. Then, $\forall x' \subseteq x \exists w' \in x' : w', \{x'\} \nVdash \alpha$. Since $x = s$, this contradicts $\exists w \in s : w, \{s\} \Vdash \alpha$. To see this, consider the singleton states $\{w'\} \subseteq x$.

7. Let φ' be of the form $\Box\alpha$: then $\tau'(\Box\alpha) = \Box\tau'(\alpha)$. Since $\{s\} \vDash \Box\alpha$, it follows that $\forall w \in s : w, \{s\} \Vdash \alpha$. Assume $\mathcal{M}, x \nVdash \Box\tau'(\alpha)$ for contradiction; it follows that $\exists x' \subseteq x : \mathcal{M}, x' \nVdash \tau'(\alpha)$. By the induction hypothesis, we derive $\exists x' \subseteq x : \{x'\} \nvDash \alpha$. Unpacking gives us $\exists x' \subseteq x \exists w' \in x' : w', \{x'\} \nVdash \alpha$, which $\forall w \in s : w, \{s\} \Vdash \alpha$ given that α is Boolean and the truth of Boolean formulas does not depend on the hyper-domain.

The right to left direction is analogous. □

Proposition A.3. *Suppose that $W \neq \varnothing$, $V : Atoms \rightarrow \wp(W)$ and $s \subseteq W$. Let $\mathcal{M} = \langle S, C, \sqsubseteq, V\rangle$ be the possibilization of $\langle W, V\rangle, \Sigma$. Then, $\Sigma \vDash \varphi$ iff $\mathcal{M}, x \Vdash [C]\tau(\varphi)$.*

Proof. We will illustrate with the left to right direction: if $\Sigma \vDash \varphi'$ then $\mathcal{M}, x \Vdash [C]\tau'(\varphi')$.

1. Let φ' be atomic; we want to show $\mathcal{M}, x \Vdash [C]p$. Since $\Sigma \vDash p$, we have $\forall w \in \bigcup\Sigma : w, \Sigma \Vdash p$, which means $\forall w \in \bigcup\Sigma : w \in V(p)$. Assume $\mathcal{M}, x \nVdash [C]p$ for contradiction. It follows that $\exists y : xCy$ & $\mathcal{M}, y \nVdash p$, which means $y \notin V(p)$. Since $y \in \Sigma$, this contradicts $\forall w \in \bigcup\Sigma : w \in V(p)$. Hence, $\mathcal{M}, x \Vdash [C]p$.

2. Let φ' be of the form $\neg\alpha$; we want to show $\mathcal{M}, x \Vdash [C]\Box{-}\tau'(\alpha)$. Since $\Sigma \vDash \neg\alpha$, we have $\forall w \in \bigcup\Sigma : w, \Sigma \Vdash \neg\alpha$, which means $\forall w \in \bigcup\Sigma \forall s \in \Sigma : w, \{s\} \nVdash \alpha$. By the definition of acceptance again, it follows that $\forall s \in \Sigma : \{s\} \vDash \neg\alpha$. Given Proposition A.2, we have $\forall s \in \Sigma : \mathcal{M}', s \Vdash [C]\Box{-}\tau'(\alpha)$, where $\mathcal{M}', s$ is the possibilization of $\langle W, V\rangle, \{s\}$. Now since $\{s\}$ only contains one state, it means if $s\,C\,y$ then $y = s$, and given $\mathcal{M}', s \Vdash [C]\Box{-}\tau'(\alpha)$, it follows that $\mathcal{M}', s \Vdash \Box - \tau'(\alpha)$. Since $\mathcal{M}'$ and $\mathcal{M}$ are both generated from $\langle W, V\rangle$, they can only differ possibly with respect to C. But since α does not contain any $[C]$, $\forall s \in \Sigma : \mathcal{M}, s \Vdash \Box - \tau'(\alpha)$. This implies $\mathcal{M}, x \Vdash [C]\Box{-}\tau'(\alpha)$.

3. Proofs for all the other operators can be given along the same lines, for example:

   Let φ' be of the form $\Diamond\alpha$; we want to show $\mathcal{M}, x \Vdash [C]\Diamond\tau'(\alpha)$. Since $\Sigma \vDash \Diamond\alpha$, we have $\forall w \in \bigcup\Sigma.\forall s \in \Sigma.\exists w' \in s : w', \{s\} \Vdash \alpha$, from which it follows that $\forall s \in \Sigma : \{s\} \vDash \Diamond\alpha$. Given Proposition A.2, we have $\forall s : \mathcal{M}', s \Vdash [C]\Diamond\tau'(\alpha)$. Since s only C-relates to itself in $\mathcal{M}'$, we have $\forall s : \mathcal{M}', s \Vdash \Diamond\tau'(\alpha)$. Because $\mathcal{M}'$ and $\mathcal{M}$ are again generated from the same $\langle W, V\rangle$, it follows that $\forall s \in \Sigma : \mathcal{M}, s \Vdash \Diamond\tau'(\alpha)$. Hence, $\mathcal{M}, x \Vdash [C]\Diamond\tau'(\alpha)$.

The right to left direction can be established analogously. □

Definition A.4. The canonical model for LED is defined as $\mathfrak{M}^\Lambda = \{S^\Lambda, \sqsubseteq^\Lambda, C^\Lambda, V^\Lambda\}$, where (i) S^Λ is the set of all consistent and deductively closed sets of formulas of LED; (ii) $\Gamma' \sqsubseteq^\Lambda \Gamma$ iff $\{\varphi \mid \Box\varphi \in \Gamma\} \subseteq \Gamma'$, that is, for all formulas φ, $\Box\varphi \in \Gamma$ implies $\varphi \in \Gamma'$; (iii) $\Gamma' \in C^\Lambda(\Gamma)$ iff $\{\varphi \mid [c]\varphi \in \Gamma\} \subseteq \Gamma'$; (iv) for any atomic formula p, $V^\Lambda(p) = \{\Gamma \in S^\Lambda \mid p \in \Gamma\}$.

Theorem A.5 (Soundness and Strong Completeness)**.** $\Gamma \vdash^{led} \varphi$ *if and only if* $\Gamma \vDash^{led} \varphi$.