Discourse consistency and dynamic modals in commitment space semantics[1]

Yichi (Raven) ZHANG — *Heinrich Heine University Düsseldorf*

Abstract. This paper examines puzzles surrounding epistemic contradiction as well as cases known as "standoff." To resolve these puzzles, I incorporate Veltman's (1996) test conception of epistemic modals into Krifka's (2021) commitment space semantics. The resulting framework is a two-pronged update semantics in which an epistemic possibility claim $\Diamond p$ carries the discourse effect of delimiting future developments of the discourse to those states where p remains as an open possibility.

Keywords: discourse consistency, update semantics, commitment space semantics, epistemic modals, epistemic contradiction, standoffs.

1. Introduction

Puzzles surrounding epistemic contradiction have long been a focal point in semantic analyses of epistemic modals. In its simplest form, epistemic contradiction concerns sentences of the form $\Diamond p \wedge \neg p$ and $\Diamond \neg p \wedge p$ as well as their reversed orders, as exemplified by (1).

(1) a. #It might be raining and it is not raining. $\Diamond p \wedge \neg p$
    b. #It is raining and it might not be raining. $p \wedge \Diamond \neg p$

As Yalcin (2007) argues, unlike the Moorean sentence "*p, but I don't know p*," the oddness of epistemic contradiction does not disappear when embedded, as the contrast in (2) demonstrates.

(2) a. Suppose it is raining and I don't know that it is raining.
    b. #Suppose it is raining and it might not be raining.

Hence, it appears a simple contextualist account where "*might p*" just means "*p* is compatible with what the speaker knows," supplemented with a pragmatic story that requires knowledge for assertion, is inadequate to explain epistemic contradiction. A semantic analysis seems needed.

The pursuit of an explanation for epistemic contradiction has sparked a slew of analyses of epistemic modals, and one testing ground for these different theories is in their treatment of epistemic contradiction of other varieties. This paper aims to contribute to the literature in this respect. In particular, I will focus on cases known as *standoff* (Bennett, 2003; Goldstein, 2022) as well as their variants. To illustrate a typical case of standoff, consider the following example adapted from Bennett (2003). Suppose there are two levers which together control a water gate: when Lever 1 but not 2 is down, Top Gate is open and water flows left; when Lever 2 but not 1 is down, top Gate is open and water flows right; when Lever 1 and 2 are both down, Top Gate remains closed. Suppose Ann and Bob each can only see the position of one lever: Ann sees that Lever 1 is down, and Bob sees that Lever 2 is down. They then each pass a note to Carl:

(3) Ann's note: Top Gate might be open, and if it is open then the water is flowing left.
    Bob's note: Top Gate might be open, and if it is open then the water is flowing right.

[1] I would like to thank the semantics and pragmatics group at Heinrich Heine University, attendants of the Nihil seminar at ILLC, and audience at SuB28 for helpful comments and discussion.

©2024 Yichi (Raven) Zhang. In: Baumann, Geraldine, Daniel Gutzmann, Jonas Koopman, Kristina Liefke, Agata Renans, and Tatjana Scheffler (eds.) 2024. Proceedings of Sinn und Bedeutung 28. Bochum: Ruhr-University Bochum, 1060-1078.

These two notes do not strike us as necessarily contradictory. In fact, after having received both notes, Carl should be able to infer that Top Gate is closed. By contrast, if the two conditionals are uttered by a single speaker as in (4), then they do sound like a blatant contradiction.

(4) Ann: #Top Gate might be open, and if it is open then the water is flowing left, and if it is open then the water is flowing right.

To explain this contrast, I will invoke a notion of commitment and provide an analysis of epistemic modals in *commitment space semantics*. In addition, the proposed framework can also capture a large variety of epistemic contradictions that have been discussed in the literature. In §2, I will review various epistemic contradictions this paper sets out to capture. I will evaluate these data against a Veltman (1996) style update semantics, mainly for two reasons. First, as we shall see, a more sophisticated update semantics can in fact capture all the data; however, it has difficulty explaining the contrast between (3) and (4). Second, it helps set up the stage for my own implementation of dynamic modals in commitment space semantics in §3. Then in §4, I will refine this analysis by offering a two-pronged update framework which separates the traditional updates on an information state from updating with a sentence's discourse effect on a commitment space. In §5, I will discuss some ways to extend the framework to address a wider range of issues. §6 concludes.

2. Desiderata: epistemic contradiction and standoff

2.1. Update semantics and variants of epistemic contradiction

In update semantics and dynamic semantics more generally, the semantic contribution of a sentence is given by its update potential, construed as a function from input contexts or information states to output contexts. Following Veltman, I use $c[\varphi]$ to represent the output from updating the context c, usually modeled as a set of worlds, with φ. In update semantics, epistemic modals receive a test interpretation under which an update with $[\Diamond\varphi]$ tests whether its input context is compatible with φ:

(5) $c[\Diamond\varphi] = \{w \in c \mid c[\varphi] \neq \varnothing\}$

If c is compatible with φ, then the test is successful and the update returns its input state c without any change; if c is incompatible with φ so that $c[\varphi] = \varnothing$, then the update returns the empty set, thereby signaling discourse anomaly. There are two commonly employed notions of discourse consistency in update semantics (see, e.g., Groenendijk et al., 1996):

(6) **Consistence:** φ is consistent iff $\exists c : c[\varphi] \neq \varnothing$.
Coherence: φ is coherent iff $\exists c : c \neq \varnothing \;\&\; c[\varphi] = c$.

Consistence states that the update with φ does not necessarily lead to the absurd information state represented by the empty set, while coherence, which entails consistence, states that there is a non-absurd state, a fixed point, which already incorporates the information conveyed by φ so that the update with φ is idle. To put it another way, given a notion of support:

(7) **Support:** an information state/context c supports φ (notation $c \vDash \varphi$) iff $c[\varphi] = c$.

for φ to be coherent, it means there must be some non-trivial body of information which supports φ. This notion of coherence plays a crucial role in explaining the contradictoriness of epistemic contradiction. Consider the update with $[\Diamond p \wedge \neg p]$ as in (1a). This is predicted to be consistent under both a static interpretation of conjunction (i.e., $c[\varphi \wedge \psi] := c[\varphi] \cap c[\psi]$) as well as a dynamic one (i.e., $c[\varphi \wedge \psi] := s[\varphi][\psi]$). The update with $[\Diamond p \wedge \neg p]$ on any c that contains both some p-world and some $\neg p$-world will yield a non-empty set as its output. Hence, the first notion of consistency falls short. By contrast, (1a) fails to be coherent regardless of whether a static or a dynamic conjunction is used, because no context that supports $\Diamond p$ can simultaneously supports $\neg p$.

However, prior literature has noted the existence of intuitively contradictory sentences that are nevertheless coherent (Mandelkern, 2020a; Holliday and Mandelkern, 2022; Yalcin, 2015; Aloni, 2001). Consider a *Wittgenstein disjunction* of the form $(\Diamond p \wedge \neg p) \vee (\Diamond \neg p \wedge p)$:

(8) #It might be raining and it is not raining, or it might not be raining and it is raining.

Assume Veltman's update clause for disjunction with $c[\varphi \vee \psi] := c[\varphi] \cup c[\psi]$; any information state that contains both some p-world and some $\neg p$-world will be a fixed point for the update with $[(\Diamond p \wedge \neg p) \vee (\Diamond \neg p \wedge p)]$. The update with the first disjunct removes from c all p-worlds, while the update with the second disjunct removes all $\neg p$-worlds; taking their union then returns the original set c. Hence, contrary to our intuition, (8) is predicted to be coherent.

If we instead opt for a dynamic interpretation of disjunction à la Heim (1983) with $c[\varphi \vee \psi] := c[\varphi] \cup c[\neg\varphi][\psi]$, we can in fact correctly predict $(\Diamond p \wedge \neg p) \vee (\Diamond \neg p \wedge p)$ to be incoherent. Since updating c with the negation of the first disjunct $[\neg(\Diamond p \wedge \neg p)]$ eliminates every $\neg p$-world from c, upon which updating with the second disjunct $[\Diamond \neg p \wedge p]$ necessarily returns the empty set, the update with the whole disjunction can never reach a non-empty fixed point.

Despite its initial success, a Heimian disjunction leads to an additional problem—some classical validities no longer hold. As Mandelkern (2020a) noticed, the law of excluded middle (LEM) fails with a Heimian disjunction. Take $\neg(\Diamond p \wedge \neg p) \vee \neg\neg(\Diamond p \wedge \neg p)$ as an example. Since updating with the negation of the first disjunct $[\neg\neg(\Diamond p \wedge \neg p)]$, which amounts to updating with $[\Diamond p \wedge \neg p]$, eliminates all p-worlds, the subsequent update with the second disjunct $[\neg\neg(\Diamond p \wedge \neg p)]$ will yield the empty set. As a result, updating with the whole disjunction just becomes updating with the first disjunct $[\neg(\Diamond p \wedge \neg p)]$, which will not always be idle. Thus, $\neg(\Diamond p \wedge \neg p) \vee \neg\neg(\Diamond p \wedge \neg p)$ fails to be valid under a dynamic interpretation of disjunction.

Relatedly, another type of cases where intuitively contradictory sentences are nonetheless deemed coherent concerns disjunctions of the form $(\Diamond p \wedge \Diamond q \wedge \neg p) \vee (\Diamond p \wedge \Diamond q \wedge \neg q)$, as shown in (9):

(9) #Either Paul might be at the party and so might Quinn but Paul isn't at the party, or Paul might be at the party and so might Quinn but Quinn isn't at the party.

Let $c = \{p\bar{q}, \bar{p}q\}$ be our input context, with $p\bar{q}$ being a world where p is true but q is false and $\bar{p}q$ being a world where the opposite holds. Then under a static interpretation of disjunction, the first update $c[(\Diamond p \wedge \Diamond q \wedge \neg p)]$ returns $\{\bar{p}q\}$; the second update $c[(\Diamond p \wedge \Diamond q \wedge \neg q)]$ returns $\{p\bar{q}\}$; the update with the whole disjunction thus returns its input context, so the sentence is incorrectly predicted to be coherent. Although adopting a dynamic disjunction helps once again, given the aforementioned problem, one may hope for a better solution.

Additionally whereas (9) sounds infelicitous, the following sentence which is of the form $(\Diamond p \wedge \Diamond q) \wedge \neg(p \wedge q)$ sounds completely natural:

(10) Paul might be at the party and so might Quinn, but it is not the case they are both at the party.

However, in classical logic (9) and (10) are logically equivalent because of DeMorgan's laws and the law of distributivity.[2] To account for this apparent contrast, Holliday and Mandelkern (2022) provide a possibility semantics based on ortholattices according to which the law of distributivity in (11) fails when φ, ψ, or χ contains any epistemic modal:

(11) $\varphi \wedge (\psi \vee \chi) \vDash (\varphi \wedge \psi) \vee (\varphi \wedge \chi)$

While Holliday and Mandelkern's orthologic can in fact capture all the data presented so far, one may worry whether a ban on distributivity with embedded modals is an overkill. For one thing, the inference from (12a) to (12b) does feel rather uncontroversial.

(12) a. Paul might be at the party, and either Quinn or Rey is at the party.
b. Either Paul might be at the party and Quinn is at the party, or Paul might be at the party and Rey is at the party.

Lastly, coherence also fails to capture the oddness of embedded epistemic contradictions as shown in (13), and appealing to a dynamic disjunction is of no help here:

(13) #It might be the case that Paul might be at the party but he isn't. $\Diamond(\Diamond p \wedge \neg p)$

Since the update with the embedded $[\Diamond p \wedge \neg p]$ on any state that contains some p and some $\neg p$-worlds will not be empty, the test imposed by the matrix $\Diamond$ is satisfied. Consequently, the update with $[\Diamond(\Diamond p \wedge \neg p)]$ will simply return its input state.

2.2. Fixed-point updates and standoff

To address these inadequacies of the standard update semantics, Klinedinst and Rothschild (2014) proposed an ingenious fix which essentially imposes a coherence check locally by requiring every update to always be repeated until it reaches a fixed point. They define a fixed-point update $c[\varphi]^* := c[\varphi]...[\varphi]$ where the output c' from repeatedly updating c with $[\varphi]$ must always reach a fixed point such that $c'[\varphi] = c'$. Normal updates $[\cdot]$ are recursively defined in terms of fixed-point updates $[\cdot]^*$ as follows:

(14) $c[p] = c[p]^*$, where p is atomic
$c[\neg\varphi] = c - c[\varphi]^*$
$c[\varphi \wedge \psi] = c[\varphi]^* \cap c[\psi]^*$

[2] It is interesting to note that the following sentence, which is of the form $(\Diamond p \wedge \Diamond q) \wedge (\neg p \vee \neg q)$, sounds slightly degraded compared to (10):
(i) Paul might be at the party and so might Quinn, but either Paul isn't at the party or Quinn isn't at the party.
While a detailed exploration of this issue will have to wait for another time, it may be suggested that in (i) the disjunction $\neg p \vee \neg q$ receives an inquisitive interpretation which then interacts with the modals and results in markedness. With an inquisitive disjunction, the DeMorgan's inference from $\neg(p \wedge q)$ to $\neg p \vee \neg q$ is indeed blocked (see, e.g., Ciardelli et al., 2018).

$c[\varphi \vee \psi] = c[\varphi]^* \cup c[\psi]^*$
$c[\Diamond\varphi] = \{w \in c \mid c[\varphi]^* \neq \varnothing\}$

The update semantics augmented by fixed-point updates successfully captures all the data discussed so far. The simple epistemic contradiction is predicted to be inconsistent as the update $c[\Diamond p \wedge \neg p]^*$ which becomes $c[\Diamond p \wedge \neg p]...[\Diamond p \wedge \neg p]$ will always yield the empty set. The Wittgenstein disjunction $(\Diamond p \wedge \neg p) \vee (\Diamond \neg p \wedge p)$ is inconsistent as well given that with fixed-point updates both $c[\Diamond p \wedge \neg p]^*$ and $c[\Diamond \neg p \wedge p]^*$ become empty. Likewise for $\Diamond(\Diamond p \wedge \neg p)$, since $c[\Diamond p \wedge \neg p]^* = \varnothing$, the update $c[\Diamond(\Diamond p \wedge \neg p)]^*$ will also return the empty set.

However, I believe the fixed-point update is too strong as a general notion of updates. Consider again a case of standoff as in (3), repeated below. There does seem to be consistent way to update with the information contained in the two notes obtained from the two speakers.

(15) Ann's note: Top Gate might be open, and if it is open then the water is flowing left.
Bob's note: Top Gate might be open, and if it is open then the water is flowing right.

Granted, one may conclude that one of the notes must contain false information because the hearer may take one of the speakers to be more trustworthy than the other, but it is also fully reasonable for someone who takes both speakers to be equally trustworthy to draw the conclusion that Top Gate is in fact closed. Nevertheless, with fixed-point updates, the update with Bob's note must be repeated until it reaches a fixed point. Consequently, even if we can predict the inference that Top Gate is closed after the first update with Bob's note, the subsequent updates with "Top Gate might be open" will necessarily lead to the absurd state.

On the other hand, as (4)—repeated below as (16)—illustrates, the standoff conditionals are indeed contradictory when they are uttered by a single speaker.

(16) Ann: #Top Gate might be open, and if it is open then the water is flowing left, and if it is open then the water is flowing right.

Likewise, if the two conditionals are produced by two different speakers but in a single conversation as in (17), the discourse is also perceived to be odd.

(17) Ann: Top Gate might be open, and if it is open then the water is flowing left.
Bob: #Yes, and if it is open then the water is flowing right.

To take stock, we want a notion of updates that is able to predict inconsistency for various epistemic contradictions from §2.1, but is also flexible enough to capture the contrast between a standard case of standoff in (15) and its infelicitous cousins in (16) and (17). To accomplish these two tasks, I will utilize a notion of commitment. I will make a first pass at cashing out this view in the next section and further spell out and refine the analysis in §4 and §5.

3. Invoking commitment

With respect to the aforementioned two tasks, the benefit of invoking commitment is apparent in the case of capturing the contrast between the standoff in (15) and its infelicitous cousins in (16) and (17). In (16), by asserting this sentence, the speaker becomes committed to (i) "Top Gate might be open," (ii) "if Top Gate is open then the water is flowing left," and (iii) "if Top Gate is open then water is flowing right." However, given a suitable analysis of conditionals which enables the inference from (ii) and (iii) to "Top Gate is not open," these three commitments jointly constitute an instance of epistemic contradiction. Similarly in (17), since Bob affirms Ann's assertion, it commits himself to both "Top Gate might be open," and "if Top Gate is open then the water is flowing left." But because Bob's own assertion also commits himself to "if Top Gate is open then the water is flowing right," again it means that Bob's commitments are jointly incompatible. By contrast, in (15), since Bob is ignorant of what Ann knows, he is not committed to "if Top Gate is open then the water is flowing left." Hence, the same contradictoriness does not arise.

What is less obvious perhaps is how to define epistemic modals in a semantics that incorporates a notion of commitment so that epistemic contradictions can be properly accounted for. The analysis I will put forth is couched in Krifka's Commitment Space Semantics (CSS). In CSS, updates are performed on *commitment spaces* (notated by uppercase C). A commitment space is a set of *commitment states* (notated by lowercase c), modeled as sets of possible worlds.[3] A commitment state is like a Stalnakerian context set but it also embodies information about each participant's discourse commitment.

A commitment space C then contains the commitment state at the current stage of discourse—called the *root* $\surd C$—which functions similar to Stalnaker's (1999) common ground, along with all possible continuations of the common ground. These notions are formally defined as follows (cf. Krifka 2021: 68):

(18) a. A commitment space C is a set of non-empty commitment states.
b. A commitment state c' is a continuation of c (notated $c' \sqsubset c$) iff $c' \subset c$ and $c' \neq \varnothing$.[4]
c. $\surd C$, the root of C, is defined as $\{c \in C \mid \neg\exists c' \in C[c \sqsubset c']\}$.

The root of a commitment space thus consists of all maximal elements of C, the commitment states that are the least specific. In this paper, I will focus on cases where the root $\surd C$ contains only one member, so to simplify, whenever $\surd C$ is a singleton that contains only one commitment state, I will also call this commitment state the root of C and use $\surd C$ to represent it.

Given that in CSS assertions express discourse commitments, an update on a commitment space C with an assertion of φ always first involves an update of the form $[s \vdash \varphi]$, which reads "s is committed to φ" with s usually being the speaker. For the time being, let us set aside the detail of this update to which I will return in §5. If other discourse participants do not object to the asserted content of φ, then φ becomes part of the common ground and C will be updated with $[\varphi]$. This update is defined as follows where "$\sqsubseteq$" is the improper version of "$\sqsubset$":

(19) $C[\varphi] := \{c \in C \mid c \sqsubseteq \surd C[\varphi]\}$

Essentially, this update will first update the current common ground (i.e., the root) by removing all worlds incompatible with φ and then lift the result to commitment space by collecting all possible continuations of the updated common ground. What distinguishes CSS is its ability to define *meta speech act* operators directly at the level of commitment spaces. These operators

[3]Here, I follow Krifka (2021) in construing commitment states as sets of worlds. In Krifka (2015), commitment states are modeled as sets of propositions instead.
[4]Note that this is slightly different from how "continuation" is defined in Krifka 2021. This change is made to facilitate my exposition of the two-pronged update framework later in §4.

are capable of changing a commitment space without altering its root by pruning certain commitment states from its input, thereby delimiting possible future developments of the discourse. Most notably is what Cohen and Krifka (2014) call *denegations* as shown in (20):

(20) a. I don't promise to come.
b. I don't claim that Paul is at the party.

(21) a. I promise not to come.
b. I claim that Paul isn't at the party.

Different from the sentences in (21) where negation applies to the propositional content of the speech act, the denegations in (20) are characterized as refusals to perform certain speech acts. In CSS, the denegation "$\sim$" is defined as follows:

(22) Denegation: $C[\sim\varphi] := C - C[\varphi]$

Figure 1 illustrates the difference between updating a commitment space with a denegation $\sim p$ and with a propositional negation $\neg p$. The update $C[\sim p]$ retains the root $\sqrt{C}$ of its input commitment space and only constrains future legal continuations by removing all states where p is settled true. By contrast, the update $C[\neg p]$ alters the root of its input commitment space.

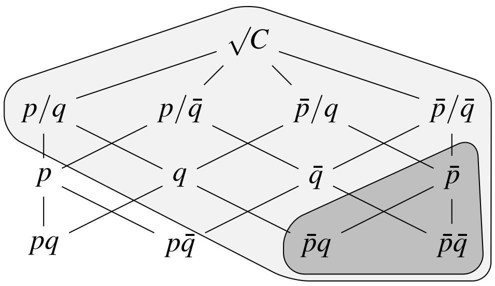

Figure 1: Updating C, which includes all the commitment states as shown above, with $\sim p$ (represented by the light grey area) and with $\neg p$ (represented by the dark grey area), respectively. In the figure, for example, p represents a state where p is settled true and q is unsettled, $p\bar{q}$ a state where p is settled true and q is settled false, and p/q a state where only the classical disjunction $p \vee q$ is settled true; with only two atomic propositions p and q, the root $\sqrt{C}$ above represents the minimal commitment state where nothing has been settled.

Moving on to the epistemic modal, I will also construe "$\Diamond$" as a meta speech act operator, at least for now. More specifically, $\Diamond\varphi$ will be viewed as the denegation of $\neg\varphi$:

(23) $C[\Diamond\varphi] := C[\sim\neg\varphi]$

An update with $[\Diamond p]$ for example constrains the commitment space by eliminating all commitment states where $\neg p$ is settled true, i.e., states that do not contain any p-worlds. Intuitively, an utterance of "*might p*" commits the speaker to keeping p as an open possibility, or at least until the *might*-claim is retracted.[5]

[5]The is perhaps reminiscent of the link between "$\Diamond$" and the *weak rejection* operator "$\ominus$" as proposed by Incurvati and Schlöder (2019, 2022). On their view, a strong assertion of $\Diamond p$ is inferentially equivalent to a weak assertion of p, which is in turn equivalent to a weak rejection of $\neg p$, i.e., $\ominus\neg p$. Since "$\ominus$" is an illucutionary force operator, a similar worry about embedded modals which I will discuss later also applies to their analysis. In fact, embedded epistemic contradictions are not immediately captured on their account and need to be derived via additional

Before elucidating how this analysis accounts for various epistemic contradictions, I will follow Cohen and Krifka (2014) and define speech act conjunction and disjunction as follows:

(24) a. Speech act conjunction: $C[\varphi \wedge \psi] := C[\varphi] \cap C[\psi]$
b. Speech act disjunction: $C[\varphi \vee \psi] := C[\varphi] \cup C[\psi]$

Lastly, we define consistency, the previous notion of support, validity, as well as *update-to-test* consequence (Veltman, 1996) at the level of commitment spaces:

(25) a. φ is consistent iff $\exists C : C[\varphi] \neq \varnothing$
b. C supports φ ($C \vDash \varphi$) iff $C[\varphi] = C$
c. φ is valid iff $\forall C : C \vDash \varphi$
d. $\varphi_1, \varphi_2, ..., \varphi_n$ entails ψ iff $\forall C : C[\varphi_1][\varphi_2]...[\varphi_n] \vDash \psi$

Let us now first consider how an update on a commitment space with the simple epistemic contradiction $\Diamond p \wedge \neg p$ looks like. Given that the update $C[\Diamond p]$ is defined as $C - C[\neg p]$, its intersection with the update $C[\neg p]$ will always be empty, as Figure 2 depicts. Hence, updates with $[\Diamond p \wedge \neg p]$ and $[\Diamond \neg p \wedge p]$ as well as their reversed orders are all predicted to be inconsistent.

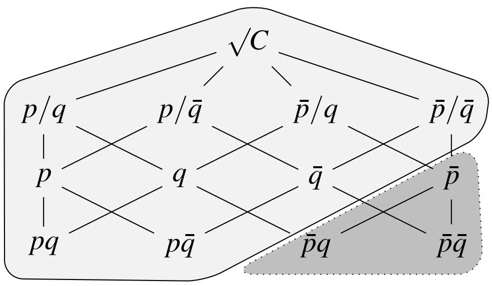

Figure 2: Epistemic contradiction $C[\Diamond p \wedge \neg p]$. The update $C[\Diamond p]$ is represented by the light grey area, and the update $C[\neg p]$ is represented by the dark grey area.

Analogously, for Wittgenstein disjunctions, since both $C[\Diamond p \wedge \neg p]$ and $C[\Diamond \neg p \wedge p]$ return the empty set. Updating with the disjunction $(\Diamond p \wedge \neg p) \vee (\Diamond \neg p \wedge p)$ which takes the union of the two sets will also be empty. Inconsistency is also predicted for disjunctions of the form $(\Diamond p \wedge \Diamond q \wedge \neg p) \vee (\Diamond p \wedge \Diamond q \wedge \neg q)$ for the same reason.

As for whether LEM is still preserved in the face of sentences like $\neg(\Diamond p \wedge \neg p) \vee \neg\neg(\Diamond p \wedge \neg p)$, we need to first modify this sentence a bit. Given that "$\Diamond$" is treated as a meta speech act operator for the moment, it follows that a propositional negation "$\neg$" cannot outscope it because otherwise the update $c[\neg(\Diamond p \wedge \neg p)]$ on a commitment state cannot be executed. Hence, the sentence needs to be reinterpreted as $\sim(\Diamond p \wedge \neg p) \vee \sim\sim(\Diamond p \wedge \neg p)$. For this sentence, LEM is indeed preserved, since updating any commitment space with it will always be idle.

This reinterpretation strategy fails to work, however, when we consider embedded epistemic contradictions such as $\Diamond(\Diamond p \wedge \neg p)$. By taking both occurrences of "$\Diamond$" to be "$\sim\neg$," again we have a case where "$\neg$" outscopes "$\sim$," and the sentence cannot be interpreted. On the other hand, if we treat the embedded "$\Diamond$" not as a meta speech act operator but simply as a Veltman test modality, then although the update can be executed, it will not give the correct result.

pragmatic means (see Incurvati and Schlöder, 2022).

The update $C[\Diamond(\Diamond p \wedge \neg p)]$ will become $C - \{c \in C \mid c \sqsubseteq \surd C[\neg(\Diamond p \wedge \neg p)]\}$. Since the update $\surd C[\neg(\Diamond p \wedge \neg p)]$ will not always return the root $\surd C$ as the update $\surd C[\Diamond p \wedge \neg p]$ will not always be empty, the update $C[\Diamond(\Diamond p \wedge \neg p)]$ also will not always be empty, which means $\Diamond(\Diamond p \wedge \neg p)$ is still predicted to be consistent.

I will address the challenge posed by embedded modals in the next section, but before moving on, let me highlight some additional features of the current analysis. One feature of this meta speech act analysis of epistemic modals is that some updates could end up eliminating all possible ways to settle an atomic sentence. For example, consider the update with the conjunction $\Diamond p \wedge \Diamond \neg p$. Since $\Diamond p$ requires all future developments to be compatible with p while $\Diamond \neg p$ requires all future developments to be compatible with $\neg p$, together they rule out all possible ways for p to be settled, yielding the commitment space as depicted in Figure 3.

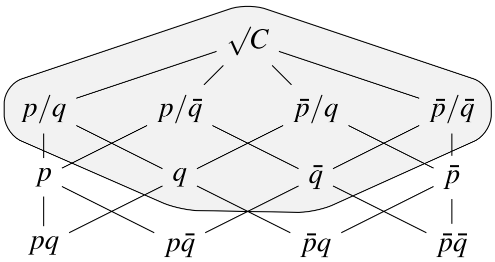

Figure 3: $C[\Diamond p \wedge \Diamond \neg p]$

To argue in favor of this result, let us compare the discourse effect of asserting $\Diamond p \wedge \Diamond \neg p$ and that of asserting one of its conjuncts. Consider the contrast between (26) on one hand, and (27) and (28) on the other.

(26) a. John might still be in France, and he might not.
b. #Yes, he is still in France/he is now in Germany.
c. No/Well actually, he is still in France/he is now in Germany.

(27) a. John might still be in France.
b. Yes, he is still in France.

(28) a. John might no longer be in France.
b. Yes, he is now in Germany.

In (26), after the update with $[\Diamond p \wedge \Diamond \neg p]$, the discourse cannot be further developed to settle John's whereabouts in either direction. To either affirm or reject p, the commitment carried by $\Diamond p \wedge \Diamond \neg p$ must be resisted or first retracted as in (26c). By contrast, in (27) the discourse can be further developed in the direction of affirming p, and in (28) in the direction of affirming $\neg p$. The current view predicts this since after updating on a commitment space with $[\Diamond p \wedge \Diamond \neg p]$, updating with either $[p]$ or $[\neg p]$ will become inconsistent.

Relatedly, recall (10), repeated below as (29), which is of the form $(\Diamond p \wedge \Diamond q) \wedge \neg(p \wedge q)$. Different from $(\Diamond p \wedge \Diamond q \wedge \neg p) \vee (\Diamond p \wedge \Diamond q \wedge \neg q)$, this sentence is intuitively unmarked.

(29) Paul might be at the party and so might Quinn, but it is not the case they are both at the party.

As Figure 4 shows, although (29) is in fact consistent, by delimiting future continuations so that Paul's being at the party and Quinn's being at the party must remain as open possibilities, we end up with a commitment space where the presence of Paul and Quinn at the party cannot be settled without the discourse participants first retracting $\Diamond p \wedge \Diamond q$. Similar to the update with $[\Diamond p \wedge \Diamond \neg p]$ in Figure 3, this update leads to a commitment space where both the issues of whether p and of whether q cannot be settled in the negative way unless one of the *might*-claims gets first retracted. Compare the following two conversations:

(30) a. Paul might be at the party and so might Quinn, but it's not the case that they are both at the party.
b. #That's right. In fact, Paul/Quinn isn't at the party.

(31) a. It's not the case that both Paul and Quinn are at the party.
b. That's right. In fact, Paul/Quinn isn't at the party.

(30b) feels odd for the very same reason that renders (26b) odd. In contrast to (31a), the addition of the might claims in (30a) rules out the possibility of smoothly developing the discourse by affirming either Paul's or Quinn's absence at the party.

Another feature of the current analysis is in its treatment of necessity modals as well as negated possibility modals. In standard update semantics, since $\Box\varphi$ continues to be defined as $\neg\Diamond\neg\varphi$, the update with $[\Box\varphi]$ amounts to the following update:

(32) $c[\Box\varphi] = \{w \in c \mid c[\varphi] = c\}$

This result is counter-intuitive, because it follows that in a context where the truth of φ has yet to be settled, an utterance of "*must* φ" will always be regarded as inconsistent, given that $c[\Box\varphi]$ will always be empty unless c already supports φ. In CSS, this problem can be avoided. The negation of a possibility modal will be represented as $\sim\Diamond\varphi$, which is equivalent to $\neg\varphi$, and the necessity modal $\Box\varphi$ can be reinterpreted as $\sim\Diamond\neg\varphi$, which is equivalent to φ—given that both "$\sim$" and "$\neg$" satisfy double negation elimination. As a result, updates on commitment spaces with $[\neg\Diamond\varphi]$ and $[\Box\varphi]$ are tantamount to updates with $[\neg\varphi]$ and $[\varphi]$. That being said, this is not the only way to construe necessity modals in CSS, and I will entertain a slightly weaker interpretation of necessity modals in §6.1.

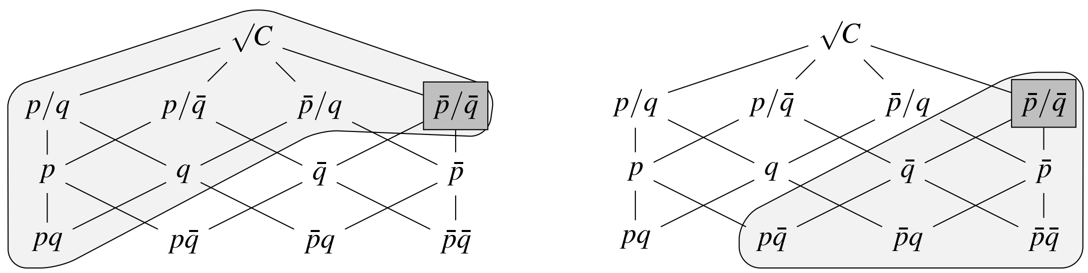

(a) $C[\Diamond p \wedge \Diamond q]$ (b) $C[\neg(p \wedge q)]$

Figure 4: $C[(\Diamond p \wedge \Diamond q) \wedge \neg(p \wedge q)]$. The final output, represented by the dark grey area, is the intersection of the outputs from the updates with the two conjuncts.

4. Refinement: a two-pronged update approach

The proposal laid out in §3 has difficulty dealing with embedded modals. As we have seen, it fails to predict the oddness of sentences like $\Diamond(\Diamond p \wedge \neg p)$. Also, treating epistemic modals as meta speech act operators across the broad does not seem to mesh with their distribution data. For instance, epistemic "*might*" can be freely embedded in downward entailing environments, which as Cohen and Krifka (2014) argue is uncharacteristic of meta speech acts (see also Krifka, 2023). Moreover, given that in standard CSS, the objects of one's commitment are propositional, we face difficulty formulating the very claim "$s \vdash \Diamond p$"—that is, s is committed to *might p*. We may relax the requirement for propositionality by allowing objects of one's commitment to be some dynamic content such as update potentials (cf. Kamp, 1990), but give that epistemic modals influence commitment spaces, it seems we still need to complicate the picture by saying objects of one's commitment are now functions over commitment spaces.

In what follows, I will argue that we can in fact retain a simple update semantics for dynamic modals and construe the objects of one's commitments, at most, as functions over commitment/information states. The meta speech act effects carried by epistemic modals will be derived as a result of some default conversational principle governing discourse developments. This results in a two-pronged update framework which imparts similar meta speech act effects to epistemic modals without directly treating them as meta speech act operators.

The default conversational principle that I was speaking of is the following:

(33) **Persistence of commitment**:
 a. if $c \vDash \varphi$ then $c[\psi] \vDash \varphi$, or equivalently
 b. if $c \vDash \varphi$ and $c' \sqsubset c$, then $c' \vDash \varphi$

Persistence states that if a commitment state already supports φ, then after updating it with any new information the output should still support φ, or equivalently, in CSS, it means that if c supports φ then it should continue to support it in all possible continuations of the discourse. When φ is a modal-free formula, persistence is always satisfied. By contrast, *test* modalities fail to satisfy persistence: a commitment state that supports $\Diamond p$ will no longer support it after an update with $\neg p$. This is in fact a desirable result for Veltman because processing a sequence like (34) does not cause any problems (Veltman, 1996: 223).

(34) Somebody is knocking at the door... Maybe it's John... It's Mary.

Here, I want to make a distinction between persistence with respect to acquiring and processing information and persistence of discourse commitment. While (34) may indicate that information processing does not have to be persistent, I argue that when it comes to making discourse commitment, persistence should still hold. If a speaker is already committed to something, then she should continue to uphold her commitment unless the commitment is retracted in light of new information. This should apply to modal commitment as well. Persistence of modal commitment highlights the forward-looking aspect in our daily usage of epistemic modals (cf. Mandelkern, 2020b). In uttering "$\Diamond p$," not only does the speaker express that p is compatible with what she knows or with the common ground information, but she also draws attention to p as an open possibility, and by doing so—that is, by proposing to treat p as an open possibility until new information becomes available—the speaker manages to orient future discourse in a given direction. To see this, let us convert (34) into a conversation.

(35) Ann: Somebody is knocking at the door.
Bob: Maybe it's John.
Ann: (*opens the door*) No, it's Mary./#Yes, it's Mary.

As (35) shows, the common ground cannot be updated with "it's Mary" straightaway without the commitment carried by "maybe it's John" having been first retracted.

At the same time, it is worth noting that persistence of commitment does not always result in the aforementioned meta speech effects of epistemic modals. For instance, prior work has pointed out a distinction between subjective and objective epistemic modalities as demonstrated by the following pair (Krifka, 2023: 131; see also Nilsen, 2004):

(36) a. It is possible that Le Pen will win, even though she certainly won't.
b. #Le Pen will possibly win, even though she certainly won't.

The objective modality in (36a) only claims that Le Pen's winning is possible according to some source of information, and the speaker does not have to be committed to this very possibility. Hence, persistence of commitment will only require in subsequent discourse that "Le Pen's winning is possible according to some source of information" still holds; it does not require Le Pen's winning to remain as an open possibility. Moreover, there are also special cases where persistence of commitment, as a default conversational principle, is overridden. In particular, this is what happens in a typical case of standoff as in (3), which I will return to in §5.

Now, given that we want the meta speech act effect of modals triggered by persistence to be a detachable enrichment, I will offer a two-pronged update framework. Updating a commitment state c with φ is accompanied by two separate updates: first, there is the usual update on a commitment state as per Veltman's semantics which takes in a commitment state and returns a commitment state—call this the *root update* $c[\varphi]$; second, there is the update with any discourse effects carried by φ on a commitment space which takes in a commitment state but returns a commitment space—call this the *discourse effect update* $c[\varphi]^+$—defined recursively as follows:

(37) **Discourse effect updates – DE updates**:
$c[p]^+ := \{c' \mid c' \sqsubseteq c[p]\}$
$c[\varphi \wedge \psi]^+ := c[\varphi]^+ \cap c[\psi]^+$
$c[\varphi \vee \psi]^+ := c[\varphi]^+ \cup c[\psi]^+$
$c[\neg\varphi]^+ := c^{\sqsubseteq} - c[\varphi]^+$, where $c^{\sqsubseteq} = \{c' \mid c' \sqsubseteq c\}$
$c[\Diamond\varphi]^+ := \{c' \sqsubseteq c \mid c'[\varphi]^+ \neq \varnothing\}$

For an atomic sentence p, its DE update outputs the set that contains all continuations of $c[p]$, i.e., all non-empty subsets of $c[p]$. If $c[p]$ itself is empty, $c[p]^+$ is also empty. Conjunction and disjunction are self-explanatory. The DE update of a negation yields a similar effect as that of updating with a denegation: $c[\neg\varphi]^+$ subtracts from all continuations of c those that are continuations of $c[\varphi]$. For $c[\Diamond\varphi]^+$, it collects only those continuations c' such that the discourse effect carried by φ can also be satisfied at c', viz., $c'[\varphi]^+ \neq \varnothing$. $\Box\varphi$ is defined as $\neg\Diamond\neg\varphi$ as usual.

To see these definitions at work, let us first consider the update $c[\Diamond p \wedge \neg p]^+$. It yields the set $\{c' \sqsubseteq c \mid c'[p]^+ \neq \varnothing\} \cap (c^{\sqsubseteq} - c[p]^+)$, which is $\{c' \sqsubseteq c \mid c'[p] \neq \varnothing\} \cap \{c' \sqsubseteq c \mid c'[p] = \varnothing\}$, which is empty. Hence, the discourse effect carried by $\Diamond p \wedge \neg p$ is inconsistent. Analogous results also hold for Wittgenstein disjunctions. Unlike treating "$\Diamond$" as a meta speech act operator outright, we can now also handle embedded epistemic contradictions. Consider the update $c[\Diamond(\Diamond p \wedge \neg p)]^+$. It yields the set $\{c' \sqsubseteq c \mid c'[\Diamond p \wedge \neg p]^+ \neq \varnothing\}$. As we just saw, updating any state with $[\Diamond p \wedge \neg p]^+$ produces the empty set, the set $\{c' \sqsubseteq c \mid c'[\Diamond p \wedge \neg p]^+ \neq \varnothing\}$ will be empty as well, which thereby makes the meta speech act effect carried by $\Diamond(\Diamond p \wedge \neg p)$ inconsistent.

At the level of commitment spaces, the update on a commitment space $C[\varphi]$ will be jointly decided by its root update $\sqrt{C}[\varphi]$ together with its DE update $\sqrt{C}[\varphi]^+$. They work in tandem to constrain the shape of the output commitment space. I assume the following two constraints in calculating the final output from a commitment space update:

(38) Let C' be the output of an update on a commitment space $C[\varphi]$, then
  a. root restriction: if $\sqrt{C}[\varphi] \neq \varnothing$ and $\sqrt{C}[\varphi]^+ \neq \varnothing$, then $\sqrt{C'} = \sqrt{C}[\varphi]$;
  b. DE restriction: $C' \subseteq \sqrt{C}[\varphi]^+ \cup \{\sqrt{C}[\varphi]\}$.

The root restriction states that whenever possible, the root of the output commitment space should be given by the root update; the DE restriction states that every commitment state in the final output must either come from the DE update $\sqrt{c}[\varphi]^+$ or be identical to the new root $\sqrt{c}[\varphi]$.[6] From these two constraints, we derive the following update clauses.

(39) **Updates on commitment spaces:** $C[\varphi] =$
  a. $\varnothing$, if $\sqrt{C}[\varphi]^+ = \varnothing$;
  b. $\{c \in \sqrt{C}[\varphi]^+ \mid c \sqsubseteq \sqrt{C}[\varphi]\}$, if $\sqrt{C}[\varphi] \in \sqrt{C}[\varphi]^+$;
  c. $\sqrt{C}[\varphi]^+ \cup \{\sqrt{C}[\varphi]\} - \{\varnothing\}$, if $\sqrt{C}[\varphi] \notin \sqrt{C}[\varphi]^+$ but $\sqrt{C}[\varphi]^+ \neq \varnothing$.

The first clause encompasses cases like epistemic contradictions where the DE update returns the empty set. Thus, the whole update is empty, and the sentence is deemed inconsistent.

The second clause covers two types of cases. First, the outcome of the root update $\sqrt{C}[\varphi]$ can be identical to the root of the commitment space generated by $\sqrt{C}[\varphi]^+$. This is the case, for instance, when φ is atomic. In such cases, since $\sqrt{C}[\varphi]^+$ is already rooted in $\sqrt{C}[\varphi]$, the update $C[\varphi]$ simply returns the whole commitment space generated by $\sqrt{C}[\varphi]^+$. On the other hand, it can also happen that $\sqrt{C}[\varphi]^+$ contains the commitment state $\sqrt{C}[\varphi]$ but not as its root. In particular, given the update clauses in (37), negation is treated like delegation when calculating the discourse effect it produces. As a result, the DE update $\sqrt{C}[\neg p]^+$, which amounts to the update with a denegation $C[\sim p]$ as shown in Figure 1, will have $\sqrt{C}$ as its root instead of the output from the root update $\sqrt{C}[\neg p]$. To satisfy the root restriction in (38), the new commitment space will be formed by collecting all states from $\sqrt{C}[\neg p]^+$ that are rooted in $\sqrt{C}[\neg p]$.

As for the last clause, there are also two types of cases falling under it. Firstly, given how disjunction is defined for the DE update, calculating the discourse effect associated with a disjunction will often return a commitment space that contains multiple roots. For instance, let $\sqrt{C}$ be the minimal state as shown in Figure 1; then the update $\sqrt{C}[p \vee q]^+$ will yield a commitment space that contains two roots, namely $\sqrt{C}[p]$ and $\sqrt{C}[q]$, represented by the two state p and q in Figure 1. The output of the root update $\sqrt{C}[p \vee q]$, represented by the state p/q in Figure 1, is not contained in $c[p \vee q]^+$. To amend this so as to satisfy the root restriction, we will add the new root $c[p \vee q]$ to the commitment space generated by $c[p \vee q]^+$.

[6] I have been assuming that the root of a commitment space will always be a single commitment state. That said, the framework can be easily generalized to accommodate roots that contain multiple commitment states.

Secondly, we have cases where the root update returns the empty set, in particular cases involving updating on a commitment state where p has yet to be settled true with the necessity claim $\Box p$. There are two possible outcomes: if p has already been settled false at $\sqrt{C}$, then the DE update $\sqrt{C}[\Box p]^+$, i.e., $\sqrt{C}[\neg\Diamond\neg p]^+$, will be empty; consequently, the update on the commitment space C will also be empty as per clause (39a). However, if p remains undetermined at $\sqrt{C}$, then the update $\sqrt{C}[\Box p]^+$ will restrict future developments of the discourse to states where p is settled true.[7] It follows that the commitment space update $C[\Box p]$ will, in effect, be tantamount to the update $C[p]$. The refined system thus inherits the desirable feature of the previous analysis concerning updating with necessity modals.

5. Standoff conditionals in discourse

In this section, I will elucidate how the update system delineated above works in a multi-agent scenario and explain how standoffs in a private information gathering setting as in (15) differ from those in a single speaker or a public discourse setting as in (16) and (17). As mentioned before, in CSS, updating with an assertion always involves first updating with a commitment clause $s \vdash \varphi$, which reads "*s is committed to* φ." Formally, $s \vdash \varphi$ will be treated as an atomic sentence and will have the same root and DE updates as other atomic sentences:

(40) a. Root update: $c[s \vdash \varphi] = \{w \in c \mid V(s \vdash \varphi, w) = 1\}$[8]
b. DE update: $c[s \vdash \varphi]^+ := \{c' \mid c' \sqsubseteq c[s \vdash \varphi]\}$

Given that updating with a claim about a commitment is different from updating with the content of the said commitment, to explain how an update with $[s \vdash \varphi]$ is transformed into an update with $[\varphi]$—in other words, how the content of a commitment becomes grounded in a commitment state—Krifka (2023) proposed the following closure condition (slightly modified):

(41) Commitment Closure of c: If s is a participant in the conversation that is trustworthy, and $s \vdash \varphi$ holds at every w in c (i.e., c supports $s \vdash \varphi$), and the other participants in conversation do not object, then updating c with $[\varphi]$.

In order to distinguish a public discourse setting where discourse participants are mutually visible to one another from a private information gathering scenario, I will make one more modification by relativizing commitment states to discourse groups, encoded via subscripts. For example, in a case like (15) where Carl receives separate notes from Ann and Bob, we can designate the relevant commitment space with C_{Carl}, for Ann and Bob are not part of the discourse. By contrast, in a public setting like (17), we can designate the relevant commitment space with $C_{\mathfrak{A}}$ with Ann and Bob being members of the discourse group $\mathfrak{A}$.

One last thing we need is a working semantics for indicative conditionals. I will assume a dynamic strict analysis of conditionals (see, e.g., Willer, 2017): an indicative conditional $\varphi \to \psi$ presupposes $\Diamond\varphi$ and asserts the strict material conditional $\Box(\neg\varphi \vee \psi)$. To simplify, we can set aside the presuppositional aspect of $\Diamond\varphi$ and take $\varphi \to \psi$ to bring about the following update:

[7] Formally, $\sqrt{C}[\neg\Diamond\neg p]^+ = \sqrt{C}^{\sqsubseteq} - \{c' \sqsubseteq \sqrt{C} \mid c'[\neg p]^+ \neq \varnothing\} = \{c' \sqsubseteq \sqrt{C} \mid c'[p]^+ = c'^{\sqsubseteq}\} = \{c' \sqsubseteq \sqrt{C} \mid c'[p] = c'\}$

[8] For Krifka (2014, 2023), updating with a commitment is more complex as it involves a performative aspect which changes the current world-time index where the discourse takes place.

(42) $C[\varphi \rightarrow \psi] := C[\Diamond\varphi][\Box(\neg\varphi \vee \psi)]$

We can now represent the three scenarios in (15-17) as inducing the following updates:

(43) a. Ann's note: Top Gate might be open, and if it is open the water is flowing left.
Bob's note: Top Gate might be open, and if it is open the water is flowing right.
b. $C_{Carl}[Ann \vdash \Diamond O][Ann \vdash \Box(\neg O \vee L)][Bob \vdash \Diamond O][Bob \vdash \Box(\neg O \vee \neg L)]$

(44) a. Ann: #Top Gate might be open, and if it is open the water is flowing left, and if it is open the water is flowing right.
b. $C_{Ann}[Ann \vdash \Diamond O][Ann \vdash \Box(\neg O \vee L)][Ann \vdash \Diamond O][Ann \vdash \Box(\neg O \vee \neg L)]$.

(45) a. Ann: Top Gate might be open, and if it is open then the water is flowing left.
Bob: #Yes, and if it is open then the water is flowing right.
b. $C_{\mathfrak{A}}[Ann \vdash \Diamond O][Ann \vdash \Box(\neg O \vee L)][Bob \vdash \Diamond O][Bob \vdash \Box(\neg O \vee \neg L)]$

Consider (44) first where the standoff conditionals are produced by a single speaker. Since Ann is committed to both conditionals and nothing suggests that she has retracted any part of the assertion, applying commitment closure will update C_{Ann} with $[\Diamond O][\Box(\neg O \vee L)][\Diamond O][\Box(\neg O \vee \neg L)]$. It is easy to verify that under the definition given in (39), this update will return the empty set: on one hand, the DE update with $[\Diamond O]^{+}$ will remove from C all commitment states where O is settled false; on the other hand, the updates with $[\Box(\neg O \vee L)]$ and $[\Box(\neg O \vee \neg L)]$ together place the opposite requirement that O is settled true in all future continuations of the discourse.

Likewise in (45), since Bob, who is a member of the discourse group $\mathfrak{A}$, does not object to Ann's assertion, commitment closure will give rise to the update $C_{\mathfrak{A}}[\Diamond O][\Box(\neg O \vee L)][\Diamond O][\Box(\neg O \vee \neg L)]$, which will again return the empty set. Hence, we correctly predict infelicity in both (44) and (45).

Now, the key difference between (43) and (45) is that in (43), Ann and Bob are not part of the discourse group when we zone in on Carl's private information. This means Carl may opt for a reinterpretation of the standoff conditionals, because the default interpretation, with the meta speech act effect carried by $\Diamond O$ at full force, results in inconsistency.[9] Since Ann and Bob are not active participants of the discourse and thus cannot react to the other person's discourse moves, commitment closure does not have to ground every aspect of a sentence. In particular, as I have previously suggested, the triggering of the meta speech act effect associated with the DE update $[\Diamond O]^{+}$ is most natural when the speaker is part of the discourse, because by delimiting future continuations of the discourse in this way the speaker can facilitate communication by directing attention to certain possibilities. Thus, when the speaker is not part of the relevant discourse group, there is less reason to carry out the update $[\Diamond O]^{+}$. Moreover, in the present case, the minimum change required to render the update with the two conditionals consistent is to block the DE update $[\Diamond O]^{+}$. Consequently, a charitable hearer in this case can weaken the default update after applying commitment closure as follows:

(46) a. $C_{Carl}[\Diamond O]^{0}[\Box(\neg O \vee L)][\Diamond O]^{0}[\Box(\neg O \vee \neg L)]$, where
b. $C[\varphi]^{0} := \{c \in C \mid c \sqsubseteq \sqrt{C[\varphi]}\}$

[9] Again, this is not to say that the hearer will always choose to reinterpret. It is possible that after having received both notes, Carl decides that one of the speakers must be wrong.

That is, the update $C[\varphi]^0$ simply collects all possible developments of the root update $\sqrt{C[\varphi]}$ without considering any further restrictions on the commitment space that φ may carry. The update in (46) produces the desired result that the output commitment space now supports $\neg O$.

On the other hand, weakening in general is not possible in a public setting like (45). After hearing Ann's assertion, Bob should be able to draw the inference himself that Top Gate is not open. As a cooperative interlocutor, Bob should immediately declare this instead of taking up Ann's commitment and placing the burden of drawing this inference on his hearers. That being said, in a conversation where one of the participants has reason to not be fully cooperative, standoff conditionals can become admissible when consecutively uttered by two different speakers:

(47) Teacher: Can you tell me what you know?
Student: I know that if Top Gate is open, then the water is flowing left.
Teacher: From what I can see, if Top Gate is open then the water is flowing right, so based on what each of us knows, what can you infer?

6. Extensions and further applications

Before concluding, I will briefly highlight some ways to extend the proposed update framework so as to apply to a wider range of issues surrounding epistemic modals.

6.1. Weak necessity modal

Given the definitions in (37) and (39), the necessity modal currently employed is strong in the sense that $\Box\varphi$ entails φ, viz., $C[\Box\varphi] \vDash \varphi$ for all C. In the literature, the question of whether epistemic *must* is a strong modality is still an issue of contention (von Fintel and Gillies, 2021; Lassiter, 2016; Goodhue, 2017). Although I do not intend to engage with this debate here, I will suggest a way to extend the current framework to allow for a weaker necessity modal. Consider an example from Kibble (1994: 8).

(48) a. John has a guitar. It's a Fender Stratocaster.
b. John must have a guitar. #It's a Fender Stratocaster.

Suppose that *must* is in fact weak. Then we may suspect the reason why (48b) is infelicitous is that "John must have a guitar" does not make it common ground that John has a guitar; since no discourse referent is introduced into the common ground, the anaphora lacks an antecedent. To cash out this intuition in CSS, we can introduce a weak necessity modal "$\boxdot$" and define its two-pronged updates on a commitment state as follows:

(49) a. Root update: $c[\boxdot\varphi] := c$
b. DE update: $c[\boxdot\varphi]^+ := c[\varphi]^+$

As Figure 5a shows, the update $C[\boxdot p]$ does not alter the root commitment state but merely restricts future continuations by retaining only those states where p is settled true.[10] As a

[10]Having the root unchanged is an oversimplification. A more sophisticated root update can, for example, collect only those worlds w in c where p is true at all of the most prototypical worlds at w (see, e.g., Kratzer, 1991).

result, $\boxdot p$ no longer entails p, since an update with the latter will also shift the root state.

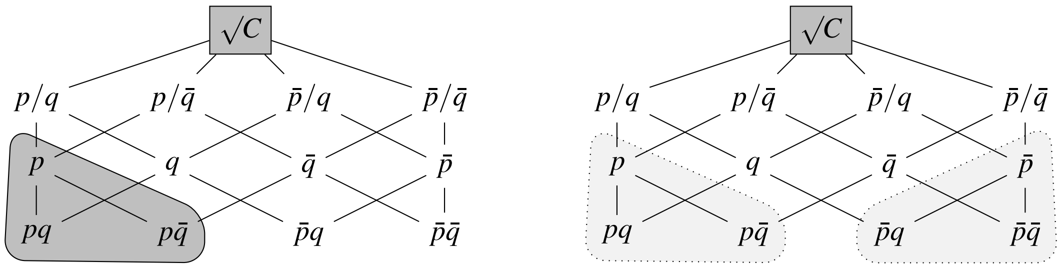

(a) $C[\boxdot p]$

(b) $C[\boxdot p \wedge \boxdot\neg p] = \{\sqrt{C}\}$

Figure 5: Weak necessity modal

One troublesome consequence arise from this interpretation of $\boxdot$ is that sentences of the form $\Box p \wedge \Box\neg p$ as well as $\Box p \wedge \neg\Diamond p$ are no longer inconsistent as updating with these sentences do not always result in the empty set as shown in Figure 4b. To explain the oddness of such sentences, we can resort to a stronger notion of consistency:

(50) **Strong Consistency:** φ is strongly consistent iff $\exists C : C[\varphi] \neq \varnothing$ & $C[\varphi] \neq \{\sqrt{C}\}$

For φ to be strongly consistent, it should be possible that an update with φ yields a commitment state which can be further developed. Since the update with $[\boxdot p \wedge \boxdot\neg p]$ necessarily results in a singleton set that contains only the current root, it is not strongly consistent. The two conjuncts impose incompatible restrictions on future continuations of the discourse so that there is no way for the current root to be further refined.

6.2. Free choice inference

Epistemic possibility modals give rise to so-called 'free-choice' inference:

(51) a. Paul might be at the party, or Quinn might be at the party. $\Diamond p \vee \Diamond q$
b. It is possible that either Paul or Quinn is at the party. $\Diamond(p \vee q)$
c. $\Rightarrow$ Paul might be at the party and Quinn might be at the party. $\Diamond p \wedge \Diamond q$

Both (51a) and (51b) give rise to the inference in (51c). The present account, however, fails to vindicate free choice. For example, a world where p is settled true and q false will survive the update with $[\Diamond p \vee \Diamond q]$ but will not survive the subsequent update with $[\Diamond p \wedge \Diamond q]$; hence, an update with $[\Diamond p \vee \Diamond q]$ will not always lead to a commitment space that supports $\Diamond p \wedge \Diamond q$.

There are various ways to extend the current framework to capture free choice. Most conveniently, we can modify the root update by adopting an update system capable of deriving free choice as an entailment at the level of commitment/information states (e.g., Goldstein, 2019; Aloni, 2022) or some more elaborate notions of information states (e.g., Ciardelli et al., 2009; Zhang, 2023). Then when calculating the meta speech effect carried by $\Diamond p \vee \Diamond q$, we also calculate those carried by its logical consequences, which in turn ensures that the output commitment space from this update will be delimited by $[\Diamond p \wedge \Diamond q]^{+}$.[11] Alternative, a more radical solution

[11] It is worth noting that this solution is incompatible with deriving free choice as a conversational implicature due

would be to modify the DE update so as to make it the case that $c[\Diamond p \vee \Diamond q]^+ = c[\Diamond p \wedge \Diamond q]^+$. I will leave further exploration of this issue to future work.

7. Conclusion

In this paper, I combined a commitment space semantics with a test semantics for epistemic modals. The resulting framework is a two-pronged update system wherein epistemic modals carry the discourse effects of delimiting future developments of the discourse: an update with $\Diamond p$ constrains future continuations to those where p remains as an open possibility. These discourse effects can be conceived of as resulting from a default conversational principle which requires discourse commitment to be persistent. This novel analysis captures a large variety of epistemic contradictions as well as discourse consistency in various standoff scenarios.

References

Aloni, M. (2001). *Quantification under conceptual covers*. Ph. D. thesis, University of Amsterdam.

Aloni, M. (2022). Neglect-zero effects in dynamic semantics. In *Tsinghua Interdisciplinary Workshop on Logic, Language, and Meaning*, pp. 1–24. Springer.

Bennett, J. (2003). *A Philosophical Guide to Conditionals*. Clarendon Press.

Ciardelli, I., J. Groenendijk, and F. Roelofsen (2009). Attention! ‘might’ in inquisitive semantics. *Semantics and Linguistic Theory (SALT) 19*, 91–108.

Ciardelli, I., J. Groenendijk, and F. Roelofsen (2018). *Inquisitive Semantics*. Oxford University Press.

Cohen, A. and M. Krifka (2014). Superlative quantifiers and meta-speech acts. *Linguistics and Philosophy 37*, 41–90.

Goldstein, S. (2019, November). Free choice and homogeneity. *Semantics and Pragmatics 12*(23), 1–53.

Goldstein, S. (2022). Sly pete in dynamic semantics. *Journal of Philosophical Logic 51*(5), 1103–1117.

Goodhue, D. (2017). Must φ is felicitous only if φ is not known. *Semantics and Pragmatics 10*, 14.

Groenendijk, J., M. Stokhof, and F. Veltman (1996). Coreference and modality in the context of multi-speaker discourse. In H. Kamp and B. Partee (Eds.), *Context Dependence in the Analysis of Linguistic Meaning*, pp. 195–215. Stuttgart: IMS.

Heim, I. (1983). On the projection problem for presuppositions. In P. H. Portner and B. Partee (Eds.), *Formal Semantics: The Essential Readings*, pp. 249–260. Blackwell.

Holliday, W. H. and M. Mandelkern (2022). The orthologic of epistemic modals. arXiv preprint arXiv:2203.02872.

to concerns about overgeneration. For example, it is plausible that $\Diamond p$ conversationally implicates $\Diamond\neg p$. However, if when calculating the meta speech act effect carried by $\Diamond p$ we also calculate the effect carried by its implicature and require the commitment space to be restricted by $[\Diamond\neg p]^+$, then we will incorrectly predict sentences like “*Paul might be at the party. In fact, he is at the party*” to be inconsistent.

Incurvati, L. and J. J. Schlöder (2019). Weak assertion. *The Philosophical Quarterly 69*(277), 741–770.

Incurvati, L. and J. J. Schlöder (2022). Epistemic multilateral logic. *The Review of Symbolic Logic 15*(2), 505–536.

Kamp, H. (1990). Prolegomena to a structural account of belief and other attitudes. In *Current Research in the Semantics/Pragmatics Interface*, pp. 513–583. Brill.

Kibble, R. (1994). Dynamics of epistemic modality and anaphora. In *International Workshop on Computational Semantics*, pp. 121–130. ITK, Tilburg.

Klinedinst, N. and D. Rothschild (2014). Epistemic contradictions: why idempotence is hygienic.

Kratzer, A. (1991). Modality. In A. von Stechow and D. Wunderlich (Eds.), *Semantik / Semantics: Ein internationales Handbuch zeitgenössischer Forschung*, pp. 639–650. Berlin and New York: De Gruyter Mouton.

Krifka, M. (2014). Embedding illocutionary acts. In T. Roeper and M. Speas (Eds.), *Recursion: Complexity in cognition*, pp. 59–87. Springer.

Krifka, M. (2015). Bias in commitment space semantics: Declarative questions, negated questions, and question tags. *Semantics and Linguistic Theory (SALT) 25*, 328–345.

Krifka, M. (2021). Modelling questions in commitment spaces. In M. Cordes (Ed.), *Asking and Answering: Rivalling approaches to interrogative methods*, pp. 63–95. Narr Francke Attempto Tübingen.

Krifka, M. (2023). Layers of assertive clauses: Propositions, judgements, commitments, acts. In J. M. Hartmann and A. Wollstein (Eds.), *Propositional Arguments in Cross-Linguistic Research: Theoretical and empirical issues*, pp. 115–182. Tübingen: Narr.

Lassiter, D. (2016). Must, knowledge, and (in)directness. *Natural Language Semantics 24*, 117–163.

Mandelkern, M. (2020a). Dynamic non-classicality. *Australasian Journal of Philosophy 98*(2), 382–392.

Mandelkern, M. (2020b). How to do things with modals. *Mind & Language 35*(1), 115–138.

Nilsen, Ø. (2004). Domains for adverbs. *Lingua 114*(6), 809–847.

Stalnaker, R. C. (1999). Pragmatic presuppositions. In *Context and Content: Essays on intentionality in speech and thought*. Oxford: Oxford University Press.

Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic 25*, 221–261.

von Fintel, K. and A. S. Gillies (2021). Still going strong. *Natural Language Semantics 29*, 91–113.

Willer, M. (2017). Lessons from Sobel sequences. *Semantics and Pragmatics 10*, 4–1.

Yalcin, S. (2007). Epistemic modals. *Mind 116*, 983–1026.

Yalcin, S. (2015). Epistemic modality *de re*. *Ergo, an Open Access Journal of Philosophy 2*(19).

Zhang, Y. (2023). On the modeling of live possibilities. *Semantics and Linguistic Theory (SALT) 33*, 686–706.