Alternatives and Local Contexts in Disjunctions

word count: 7176 words

September 25, 2026

Abstract

By appealing to presupposition filtering, pronominal anaphora, the *or*-to-*if* inference and modal domain restriction, dynamic theories of disjunctions establish a convincing case that the right disjunct in a disjunction is evaluated against a local context entailing the negation of the left disjunct. But this is mistaken. The local context of a right disjunct doesn't always entail the negation of the left disjunct. In particular, when the left disjunct is modalized, the right disjunct can sometimes have a local context only entailing the negation of the left disjunct's prejacent. We develop a new QUD-sensitive account of local contexts in disjunctions to explain the variability, and show that it explains inferences the standard account doesn't capture.

1 Introduction

As Karttunen (1974) famously observed, the presupposition of a right disjunct is "filtered" by the left disjunct's negation. For example, the right disjunct of (1) presupposes that John previously smoked, but the entire disjunction presupposes nothing, or, at most the trivial conditional (1b):

(1) Either John never smoked, or John stopped smoking.

a. $\not\rightsquigarrow$ John previously smoked.

b. $\rightsquigarrow$ If John previously smoked, he previously smoked.

A promising approach to presupposition says, based on this observation, that a sentence's presuppositions need to be satisfied (entailed) by its *local* context, which is the global context updated with some temporary suppositional information.[1] In particular, many contemporary theories of disjunctions posit that a right disjunct is evaluated in a local context updated with the negation of the left disjunct. This accounts for Karttunen's observation. Since the negation of "John never smoked" is "John previously smoked", if the second disjunct of (1) is evaluated in a local context updated with the negation of the first disjunct, its presupposition is *locally satisfied*. This approach is exemplified by Beaver's (2001) dynamic entry for disjunction:[2]

(2) $c[\varphi \text{ or } \psi] := c[\varphi] \cup c[\neg\varphi][\psi]$

which says the result of updating a context c with a disjunction is the union of c updated with the left disjunct and c first updated with the negation of the left disjunct and then the right disjunct.

Besides presupposition filtering, this dynamic approach can also be motivated by its application to pronominal anaphora, the *or*-to-*if* inference and modal domain restriction in information-sensitive modals. For instance, as we will explain in §2.3, the assumption that the right disjunct in a disjunction has a local context entailing the negation of the left disjunct help deliver an elegant account of how (3) seems to imply (4)–the *or*-to-*if* inference–

(3) John is in New York, or he must be in Boston.

(4) If John is not in New York, he must be in Boston.

and how the modal domain of *must* in both (3) and (4) is intuitively only a subset of worlds where John is not in New York.

While the dynamic approach is insightful in appealing to local contexts, the standard implementation (2) is problematic. It predicts that the local context of ψ in $\ulcorner\varphi \text{ or } \psi\urcorner$ is invariably $c[\neg\varphi]$. However, in sentences of the form $\ulcorner\triangle\varphi \text{ or } \psi\urcorner$ where $\triangle$ is a modal, the local context of the right disjunct often entails the negated prejacent $\neg\varphi$ rather than $\neg\triangle\varphi$. We'll call this problem the *prejacent problem*.

[1] We'll focus throughout on a satisfaction-based theory of presupposition of the kind developed by Stalnaker (1974, 1978), Karttunen (1973, 1974), and Heim (1983). The information is "temporary" in the sense that the interlocutors are not committed to the truth of it.

[2] Throughout, we assume that updates are intersective: $c[p] = c \cap \{w : [\![p]\!](w) = 1\}$ (or, $c[p] = c \cap [\![p]\!]$) for an atomic p. We also assume, following Heim (1983) and others, that negation is characterized by $c[\neg\varphi] = c \backslash c[\varphi]$. Conceptually, note that a classical evaluation function is still needed in our model, so a reasonable way to understand the picture is that the evaluation function outputs *semantic interpretations* and the context update potentials model *pragmatic interpretations* (Stalnaker, 1975).

To illustrate the basic problem, consider data concerning presupposition filtering. If we add an epistemic *must* to the left disjunct of (1), the standard theory predicts that the right disjunct's local context gets weakened and only entails $\neg\Box\varphi$, which amounts to $\Diamond\neg\varphi$. Hence, the nontrivial presupposition (5b) should hold globally, contrary to fact. In (5), no global presupposition is felt, just like in the case of (1).

(5) John must've never smoked, or he stopped smoking.

a. $\not\rightsquigarrow$ John previously smoked.
b. $\not\rightsquigarrow$ If John might not have previously smoked, he previously smoked.
c. $\rightsquigarrow$ If John previously smoked, he previously smoked.

As we show, this isn't a problem unique to presupposition filtering; it resurfaces in all the prominent constructions used to motivate the local contexts posited in (2) – for $\ulcorner \triangle\varphi$ or $\psi \urcorner$ where $\triangle$ is a modal, the local context of ψ sometimes only entails $\neg\varphi$. Moreover, the pattern is not confined to epistemic modals. For convenience, we will sometimes refer to the sentence whose negation updates the local context of a disjunct as "the restrictor" of that disjunct (and the proposition it expresses "the restrictor proposition"). So, the question becomes why and when φ (rather than $\triangle\varphi$) can act as ψ's restrictor. We propose a view where disjunctions have alternative-sensitive local contexts that are regulated by the QUD; the restrictor of a disjunct is a particular relevant alternative of the other disjunct.

Here's the plan for the paper. We first review the empirical data that motivate the standard dynamic theory of disjunction, and present the data it fails to capture in §2. In §3, we give a new account of local contexts within disjunctions that predicts the intuitive interpretations of the problematic data, and then show that given these interpretations, we easily account for the inferences licensed by our data, and in particular, the failures of the *or*-to-*if* inference pattern, that the standard theory fails to capture; we also provide a generalization of the relevant inferences from disjunctions to conditionals.

2 The local context of modalized disjunctions

In this section, we argue that the local context of the right disjunct is not always the global context updated by the negation of the left disjunct. In particular, when one disjunct is modalized, the restrictor of the other disjunct is often its prejacent rather than the entire modalized expression, contrary to the standard theory. That is, a modal disjunct often updates the other disjunct's local context with the negation of its prejacent. We call this general phenomenon *prejacent restriction.*

> **Prejacent Restriction.** When one disjunct is modalized, the restrictor of the other disjunct may be the former's prejacent.

In this section, we present three empirical arguments for prejacent restriction based on presupposition filtering, pronominal anaphora, and failures of the *or*-to-*if* inference.

2.1 Argument #1: presupposition filtering

The standard theory of presupposition satisfaction for disjunctions predicts that, given a disjunction $\ulcorner A \text{ or } B_p \urcorner$ where B presupposes p, the disjunction globally presupposes $\neg A \rightarrow p$, but not p simpliciter. And this prediction is delivered by postulating that B_p has a local context entailing $\neg A$, i.e., $c[\neg A]$. This explains how (6) isn't felt to globally presuppose anything; its presupposition (6b) is entailed by a default context. Similarly, (7) is said to only presuppose the trivial (7b).

(6) There is no King of France, or the King of France is bald.

a. $\not\rightsquigarrow$ There is a unique King of France.

b. $\rightsquigarrow$ If there is a King of France, there is a unique King of France.

(7) John never smoked, or he stopped smoking.

a. $\not\rightsquigarrow$ John previously smoked.

b. $\rightsquigarrow$ If John smoked previously, he smoked previously.

It follows from the standard theory that, for $\ulcorner A \text{ or } B_p \urcorner$, when $\neg A$ doesn't entail p, the presupposition p won't be locally satisfied (unless the global context already entails it). However, this prediction is not borne out in the case of (8). The standard theory predicts that (8b) should be globally presupposed and that the presupposition cannot be locally satisfied. However, it's clear that (8b) is not presupposed. In terms of global presupposition, (8) is on a par with (6), suggesting that the local context of the right disjunct doesn't merely entail the *might*-claim, but rather entails that there is a King of France.[3]

(8) There must be no King of France, or the King of France is bald.

a. $\not\rightsquigarrow$ There is a unique King of France.
b. $\not\rightsquigarrow$ If there might be a King of France, there is a unique King of France.

Similarly, it is predicted that (9) globally presupposes (9b), which again is not borne out. This suggests that the local context of the right disjunct entails that John previously smoked, not just that he might have smoked before.

(9) John must have never smoked, or he stopped smoking.

a. $\not\rightsquigarrow$ John previously smoked.
b. $\not\rightsquigarrow$ If John might have smoked previously, he smoked previously.

Moreover, this is not confined to cases where left disjuncts contain epistemic modals. When the left disjunct contains a root modal, the local context of the right disjunct can also be updated by the negation of the modal prejacent.[4] This is evidenced by (10) not presupposing (10a) but only (10b) intuitively.

(10) We can't build a stadium here, or the stadium will make the traffic worse.

a. $\not\rightsquigarrow$ If we can build a stadium here, there will be a unique stadium.
b. $\rightsquigarrow$ If we build a stadium here, there will be a unique stadium.

Therefore, a modal disjunct can sometimes have its prejacent be the restrictor of the other disjunct, so that the prejacent's negation updates the other disjunct's local context.

Finally, we should flag that the left disjunct's presupposition can also be locally satisfied if the negation of the right disjunct entails the presupposition, as evidenced by (11a) (Schlenker, 2009; Kalomoiros and Schwarz, 2021).

[3] Klinedinst and Rothschild (2012) make a similar style of argument about what they call non-truth-tabular disjunctions, where the second disjunct (typically) involves a counterfactual. Our data extend Klinedinst and Rothschild's observations.

[4] Note that the prejacent of the left modal is future-oriented, and so the natural restriction would have to be paraphrased as *There will be a stadium* because the futurate without *will* is odd; we will assume that it's the prejacent that's the restrictor (broadly following Kaufmann, 2005, on future reference in conditional antecedents) but our account is compatible with the restrictor being the *will*-alternative of the *can*-sentence.

(11) a. John stopped smoking, or he never smoked.

b. $\not\rightsquigarrow$ John previously smoked.

This observation is sometimes called symmetric filtering, in the sense that a left/right disjunct's presupposition can be satisfied by the right kind of right/left disjunct, regardless of the order in which they appear. Notably, dynamic semantics has difficulty capturing symmetric filtering without postulating multiplying entries of disjunctions;[5] modifying the entry in (2) along the following lines:

(12) $c[\varphi \text{ or } \psi] := c[\neg\psi][\varphi] \cup c[\neg\varphi][\psi]$

to allow the left disjunct to always have a local context updated with the negation of the right disjunct leads to undesirable result: in sentences like (6) where the right disjunct presupposes that there is a King of France, this entry predicts that it is globally presupposed.[6] We will show in §3.3.1 that our proposal of a modified dynamic entry is flexible enough to capture symmetric filtering without postulating ambiguity.

2.2 Argument #2: Partee disjunctions

In so-called "Partee disjunctions," also known as "bathroom sentences," like (13), the pronoun in the right disjunct is licensed because it's interpreted in a local context updated by the negation of the indefinite in the left disjunct, i.e., that there is a bathroom in the house:

(13) Either there is no^x bathroom in this house, or it$_x$'s in a funny place.

Likewise, it seems clear that the anaphora in the right disjunct of (14) is licensed, even though the standard theory predicts that its local context, again, only entails that there's an epistemically possible bathroom.

(14) There must not be a^x bathroom, or it$_x$'s in a funny place.

Examples like (14) are problematic for standard dynamic theories of epistemic modals, such as Veltman's (1996) Update Semantics (US) and its many derivatives, which give epistemic *might*

[5] Postulating an ambiguity means that besides (2) that does the left-to-right filtering, we also have $c[\varphi \text{ or}_2\ \psi] := c[\neg\psi][\varphi] \cup c[\psi]$, which does right-to-left filtering.

[6] This is because $c[\neg\psi] = c \backslash c[\psi]$ (fn. 2), where ψ is the right disjunct, which means ψ needs to update the local context c. By satisfaction theory, ψ's presupposition needs to be satisfied by its local context for an update to go through, meaning that c needs to entail ψ's presupposition.

and *must* a test semantics:[7]

(15) a. $c[\Diamond\varphi] := \{w \in c \mid c[\varphi] \neq \emptyset\}$
b. $c[\Box\varphi] := \{w \in c \mid c[\varphi] = c\}$

This means that epistemic *might* and *must* are *externally static* – an indefinite in their scope cannot introduce a discourse referent (dref) capable of licensing anaphora outside their scope. For the most part, it seems like epistemic modals behave as if they're externally static.[8]

(16) a. Bill might have a^x car. #It$_x$'s black.
b. John must have a^x guitar. #It$_x$'s a Fender Stratocaster. (Kibble, 1994)

To a certain extent, it's unsurprising that the same problem reappears in the context of anaphora, given recent work connecting anaphora to presupposition projection (Mandelkern, 2022; Chatain, 2025).

The problem is not isolated to epistemic modals. It's also easy to construct "practical" Partee disjunctions like (17). Here the indefinite licenses a subsequent anaphoric pronoun, even though the pronoun's predicted local context entails only that there are deontically ideal worlds in which there is a King of France:[9]

(17) There can be no^x King of France, or he$_x$ would abuse his power.

Similarly, in (18), although the local context as predicted by the standard theory only entails that in some deontically ideal worlds, the addressee's child has a credit card, the anaphora in the right disjunct is licensed.

(18) You can't let your kid have a^x credit card, or he would use it$_x$ for silly things.

Like epistemic modals, root modals appear to be externally static:

(19) a. Bill may order a^x pizza. #It$_x$'s pepperoni.
b. Bill has to order a^x pizza. #It$_x$'s pepperoni.

So, we have a puzzle: in disjunctions where the first disjunct is modalized, the local context

[7] Veltman's (1996) original system does not include an entry for epistemic *must* – the entry in (15) is simply the dual of the entry for *might* given in Veltman's paper. See Willer (2015) for further discussion of how epistemic *must* behaves in US. Moreover, strictly speaking, the external staticity pertaining to drefs applies to a test semantics where an information state c is a set of world-assignment pairs.

[8] However, see Elliott (2022) for several other examples that prove difficult for externally dynamic epistemic modals.

[9] Intuitively, in (17) *no* (usually taken as $\neg > \exists$) splits scope around *can*, and the left disjunct is equivalent to *there can't be a^x King of France.*

updated with the entire modal is seemingly incapable of licensing anaphora, yet such sentences *do* license anaphora. This suggests that, in these examples, the second disjunct's local context is updated, not with the entire modal expression, but just the modal's prejacent; this way, the local context entails the existence of the relevant antecedent. (While we won't derive details of how anaphora is licensed in these examples, we aim to derive the intuitively correct local context for the anaphora, so that standard treatments of anaphora licensing can be plugged in.)

2.3 Argument #3: *or*-to-*if* and modal domain restriction

Stalnaker (1975) observes that we naturally infer (20b) from sentences like (20a).

(20) a. John is in New York or Boston.
b. $\rightsquigarrow$ If John is not in New York, he is in Boston.

This inference has been dubbed the *or*-to-*if* inference:

(21) ***Or*-to-*if*.** $\varphi \vee \psi \rightsquigarrow \neg\varphi \rightarrow \psi$

However, given a variably-strict semantics for conditionals, *or*-to-*if* isn't classically valid. Indeed, if *or*-to-*if* were classically valid, then indicative conditionals would be equivalent to material conditionals – a thesis that most contemporary theorists reject.[10] Dynamic semantics offers a promising explanation for the intuitive validity of *or*-to-*if*: while the inference is not classically valid, i.e., truth-preserving, it is valid in a weaker sense given the standard entry of disjunction and a natural entry for the indicative.

In particular, the *or*-to-*if* inference is acceptance-preserving where a context c accepts a proposition φ just in case

(22) A context c *accepts* φ, written $c \vDash \varphi$, iff $c[\varphi] = c$. (Veltman, 1996)

We can explain the felt validity of *or*-to-*if* inferences by defining logical consequence in terms of acceptance-preservation rather than truth-preservation:

(23) $\varphi_1, \ldots, \varphi_k \vDash \psi$ iff for every s, $s[\varphi_1] \cdots [\varphi_k] \vDash \psi$, i.e., $s[\varphi_1] \cdots [\varphi_k] = s[\varphi_1] \cdots [\varphi_k][\psi]$.

In other words, an argument is valid if an information state that accepts all its premises also accepts its conclusion – i.e., the conclusion doesn't add any information to the common ground

[10] Stalnaker (1975) dubbed this argument for the material analysis the "direct argument."

not contained in the premises.[11]

It is crucial that the right disjunct in (20a) has a local context that entails that John isn't in New York. When coupled with a standard dynamic theory of indicative conditionals, like Gillies's (2004) suggested entry below, it immediately follows that any information state that accepts a disjunction accepts a corresponding conditional.

(24) $c[\varphi \rightarrow \psi] := \{w \in c : c[\varphi] = c[\varphi][\psi]\}$

Here's a simple argument that *or*-to-*if* is acceptance-preserving given the entry in (2). First, an information state c accepts a disjunction $\ulcorner \varphi$ or $\psi \urcorner$ only if the $\neg\varphi$-worlds in c are all ψ-worlds and the $\neg\psi$-worlds are φ-worlds. However, given the entry in (24), c accepts $\ulcorner \neg\varphi \rightarrow \psi \urcorner$ if and only if the $\neg\varphi$-worlds in c are ψ-worlds; hence, any c that accepts $\ulcorner \varphi$ or $\psi \urcorner$ also accepts $\ulcorner \neg\varphi \rightarrow \psi \urcorner$. Hence, while *or*-to-*if* isn't a truth-preserving argument, it *is* acceptance-preserving.[12]

The notion of local context also gives us an elegant explanation of modal domain restriction in disjunctions; intuitively, the disjunction in (25a) implies (25b) where the modal's domain is restricted to worlds where John is not in New York.

(25) a. John is in New York, or he must be in Boston.
b. $\rightsquigarrow$ If John is not in New York, he must be in Boston.

A promising explanation is simply that epistemic modals' (and some root modals') domains are restricted by their local context (Yalcin, 2007; Klinedinst and Rothschild, 2012; Mandelkern, 2019, 2024). For concreteness, consider the Veltman-style test semantics for epistemic *must* introduced in the previous section:

(26) $c[\Box\varphi] := \{w \in c : c[\varphi] = c\}$

Given either the entry in (2), the *must*-disjunct has a local context $c[\neg\textbf{in\_nyc}(j)]$. Given the Veltman-style semantics, epistemic *must* in (25a) tests whether all the worlds in c where John isn't in New York are worlds where he is in Boston.

[11]This notion of entailment is a direct decedent of Stalnaker's (1975) more pragmatic definition of "reasonable inference," see Santorio (2022) and Égré et al. (2025) for recent discussion, and corresponds to validity$_2$ in Veltman (1996). There are several other plausible acceptance-based notions of entailment one can also consider here – e.g., Veltman's validity$_3$ checks whether the information state which accepts each premise individually accepts the conclusion.

[12]Plausibly, (20a) also implies that if John is not in Boston, he's in New York. But this inference is indeed captured by (2) assuming that being in NY and being in Boston are incompatible. It's perhaps worth noting that the implausible symmetric entry (12) for disjunction would make the inferences to both conditionals acceptance-preserving.

However, here too we find evidence of the pattern we observed with presupposition filtering and pronominal anaphora: the local context of a disjunct isn't always the result of updating the global context with the negation of the other disjunct. As Klinedinst and Rothschild (2012) and Meyer (2015) observe, while the left disjunct of (27) contains an epistemic *must*, the disjunction intuitively licenses an inference to (27b) rather than an inference to (27a).[13]

(27) John must be in New York, or he must be in Boston.

  a. $\not\rightsquigarrow$ If John might not be in New York, he must be in Boston.
  b. $\rightsquigarrow$ If John is not in New York, he must be in Boston.

Examples like (27) appear to constitute a clear counterexample to the validity of *or*-to-*if*. If the *or*-to-*if* inference were valid *tout court*, then speakers should draw the inference in (27a), contrary to fact. These authors also observe that this phenomenon isn't unique to epistemic modals. For example, consider (28) where the left modal is teleological.[14] The *or*-to-*if* inference clearly fails here too: (28) should imply (28a) according to *or*-to-*if*, but instead it implies (28b).

(28) John has to write a paper, or he will get a C.

  a. $\not\rightsquigarrow$ If John doesn't have to write a paper, he will get a C.
  b. $\rightsquigarrow$ If John doesn't write a paper, he will get a C.

From these data, we draw two conclusions. First, the *or*-to-*if* inference may fail in doubly-modal disjunctions of the form $\ulcorner \triangle_1 \varphi$ or $\triangle_2 \psi \urcorner$. And the conditional inferences licensed by the relevant disjunctions often have the form $\neg\varphi \rightarrow \psi$. Second, in those disjunction, we often want the domain of the right disjunct's modal to be restricted by $\neg\varphi$ rather than $\neg \triangle \varphi$. This is another case of prejacent restriction.

[13]Notably, Klinedinst and Rothschild (2012) do observe that *or*-to-*if* inference cannot provide a satisfactory explanation to these cases with prejacent restriction, but they don't take this to motivate revisions to the pattern like we do. Cf. Dever (2013) who also discusses instances of the *or*-to-*if* inference involving modal disjuncts; however, Dever's primary interest was explaining minimal pairs like the following:

i a Susan must not know the answer, or she would raise her hand.
  b # Susan would raise her hand, or she must not know the answer.

[14]This example is adapted from Meyer (2015).

3 Proposal

The key observation advanced so far has been that, sometimes, the local context of ψ in $\ulcorner \triangle\varphi$ or $\psi\urcorner$ is updated by $\neg\varphi$, instead of $\neg\triangle\varphi$. Why would the prejacent of a modal be available to update the local context? Since the prejacent is an *alternative* to the modal, we propose that local contexts in disjunctions are alternative-sensitive.[15]

As a side note, besides φ and ψ, we will be using $A, B, C, \ldots$ as metavariables for sentences and $p, q, \ldots$ for (formal) alternatives and presuppositions, and A,B,C,... and $\overline{\mathrm{A}}$, $\overline{\mathrm{B}}$, $\overline{\mathrm{C}}$ for propositions and their complements.

3.1 Ingredient #1: formal alternatives

Here we can appeal to Katzir's (2007) algorithm for formal alternatives familiar from the literature on scalar implicatures (see also Fox and Katzir, 2011). A formal alternative of a syntactic object S is one S' such that $S' \preceq S$, glossed as "S' is at most as complex as S," defined as follows:

(29) $S' \preceq S$ if S' can be derived from S by successive deletions and/or replacements of subconstituents with lexical items from the same grammatical category.

This algorithm predicts that $\Box A$ has the following formal alternatives:

(30) $Alt(\Box A) = \{\Box A, \Diamond A, A\}$

We obtain $\Diamond A$ by replacing $\Box$ with its dual (and scale-mate) $\Diamond$, since both are modal auxiliaries, and we can obtain A by simply deleting $\Box$.

The idea that local contexts are alternative-sensitive is familiar. The standard satisfaction theory of presupposition that we are adopting runs into a well-known problem: hearers accommodate presuppositions stronger than those predicted by the theory. Geurts (1996, 1998) dubbed this the "proviso problem." Geurts observed that, the satisfaction theory predicts that the conditional (31) should presuppose the conditional in (31b), but speakers typically accommodate the stronger presupposition (31a).

[15] Cf. Meyer (2015) who appeals to alternatives for the data but denies the connection with local contexts.

(31) If Theo hates sonnets, his wife hates them too.

a. $\rightsquigarrow$ Theo has a wife.

b. $\rightsquigarrow$ If Theo hates sonnets, Theo has a wife.

Singh (2007, 2008) proposes a solution to the proviso problem predicated on the idea that hearers reason about a set of formal alternatives that constitute possible candidates for accommodation. For example, given a sentence $A \rightarrow B_p$ where the consequent presupposes p, a hearer entertains the possible accommodations $\{A \rightarrow p, p\}$ associated with the formal alternatives $\{A \rightarrow B_p, A, B_p\}$; one update is picked among alternative global context updates. The problem of prejacent restriction, in many ways, seems like a local analogue of the proviso problem, where some alternative of one disjunct gets selected to be the restrictor of the other disjunct. In other words, one update is selected among alternative local context updates. That said, instead of postulating a pragmatic repair mechanism like Singh's, we aim to give a semantics that involves a principled algorithm that gives clear predictions about what the restrictor proposition is, in a context.

3.2 Ingredient #2: relevance

Not just any formal alternative of one disjunct can serve as the other disjunct's restrictor; crucially, the disjunct has to be *relevant*. In particular, the restrictor must address the Question Under Discussion (QUD) (Roberts, 2012). To illustrate, consider two variations on the following context: John's grades are currently in the A-range, but will drop to a B+ if he fails to write his final paper for the class. First, consider a context for (32) as uttered among the students, in which the professor has hinted that she's dropping the paper requirement for everyone. The professor has made her decision, but has yet to announce it. The left disjunct intuitively addresses the salient question whether the requirement holds, leading to the restriction of the right modal by the entire modal claim.

(32) **QUD:** Is there a paper requirement? $\{\Box\mathbf{write}(j), \overline{\Box\mathbf{write}(j)}\}$

a. John has to write a paper, or he will get an A.

b. $\rightsquigarrow$ If John doesn't have to write a paper, he will get an A.

c. $\not\rightsquigarrow$ If John doesn't write a paper, he will get an A.

Now imagine the exact same context, but instead of the professor contemplating whether to scrap the final paper, John is considering blowing off the final paper to party with his friends. When the QUD concerns whether the prejacent of the modal claim in the left disjunct is true, restriction by the prejacent is observed.

(33) **QUD:** Will John write a paper? $\{\mathbf{write}(j), \overline{\mathbf{write}(j)}\}$

a. John has to write a paper, or he will get a C.

b. $\not\rightsquigarrow$ If John doesn't have to write a paper, he will get a C.

c. $\rightsquigarrow$ If John doesn't write a paper, he will get a C.

Hence, the availability of prejacent restriction is constrained by the QUD:

> **Relevance Constraint.** For a disjunction $\ulcorner A$ or $B\urcorner$ given a QUD Π which partitions c: the local context can be updated with $\neg p$ with $p \in Alt(A)$ only if p is relevant to Π, and the same for $\neg q$ with $q \in Alt(B)$.

For concreteness, we assume Groenendijk and Stokhof's (1984) partition semantics, treating questions as partitions of the context c. If the QUD stack contains multiple questions, we can take Π to be the least common refinement of all the partitions. We spell the notion of relevance as follows:

> **Relevance.** p is (strongly) relevant to a question Π iff p partially or completely answers Π and is not identical to c – i.e., $\exists X \subset \Pi : X \neq \varnothing \; \& \; \bigcup X = [\![p]\!]$ (cf. Lewis, 1988).[16]

Note that if we assume, following Yalcin (2007), that the ordinary semantic value for epistemic modal claims $[\![\Diamond\varphi]\!]^c$ tests whether c is compatible with the prejacent, then the proposition expressed will be c itself or $\varnothing$; hence, an epistemic modal claim won't be (strongly) relevant.[17]

[16]Again, we assume the model contains both the classical evaluation function and, essentially, an evaluation function that outputs a sentence's update potential. See fn.1.

[17]That is, the truth-conditions of epistemic modal claims can be stated as $[\![\Diamond\varphi]\!]^{c,w} = \exists w' \in c : [\![\varphi]\!]^{c,w'} = 1$; which has the advantage of corresponding to the update rule we postulate.

3.3 The entry

We're now ready to sketch a simple algorithm for determining a disjunct's restrictor, out of the other disjunct's relevant alternatives. This algorithm comprises two steps:

- **step 1:** given a disjunct φ, we find the set of candidate restrictors $C(\varphi)$
- **step 2:** given the candidate restrictors $C(\varphi)$, we take the conjunctive closure of $C(\varphi)$ and let that be the restrictor of the other disjunct

The set of candidate restrictors should minimally be a subset of the given disjunct's formal alternatives relevant to the QUD. However, not any relevant alternative will do. For any expression φ, φ is a formal alternative to $\neg\varphi$; however, in $\ulcorner\neg\neg\varphi$ or $\psi\urcorner$ with simplex φ, it's clear that ψ's local context is always the negation of $\neg\neg\varphi$ and never the negation of $\neg\varphi$. Absent this constraint, we predict incorrectly that there's a reading of (34) that globally presupposes a tautology, while in fact it's infelicitous for presupposing that if John never smoked, he smoked previously.

(34) #It's not the case that John never smoked before, or he stopped smoking.

An intuitive reason behind the constraint is that, the restrictor proposition based on a disjunct should be a relevant proposition that is most directly supported by the disjunct; and $\neg\varphi$ clearly doesn't support the relevant alternative φ in any obvious way to cash the idea out.

To implement this constraint, we require that, roughly, whichever cells of the partition are incompatible with the restrictor are also incompatible with the original disjunct. For a disjunct A, we define the set of candidate restrictors $C(A)$ as the set $Alt_R(A)$ of relevant alternatives to A such that, for every cell P in the partition of c, Π, if P rejects (i.e., accepts the negation of) p, then P also rejects A.

(35) $C(A) := \{p \in Alt_R(A) \,|\, \forall \mathsf{P} \in \Pi : \mathsf{P} \vDash \neg p \Rightarrow \mathsf{P} \vDash \neg A\}$

Equivalently, (35) can be stated in terms of an acceptance-based notion of consistency. We say that a context c is *update-consistent* with φ just in case some non-empty subset of c accepts φ:

(36) c is *update-consistent* with φ iff $\exists c' \subseteq c : (c' \neq \emptyset) \wedge (c' = c'[\varphi])$

I.e., c is update-consistent with φ if c already accepts φ or can accept it after undergoing an update.

Hence, we can see that (35) says that every cell in Π that is update-consistent with the original disjunct A must also be update-consistent with a candidate restrictor $p \in Alt_R(A)$. Let UC_Π be the function that takes a sentence and outputs the set of cells in Π update-consistent with that sentence. And p passes the test for a candidate restrictor based on φ just in case:

$$\mathsf{UC}_\Pi(A) \subseteq \mathsf{UC}_\Pi(p)$$

Given the definition in (35), we avoid the over-generation problem discussed above. Suppose, for example, that the partition Π is $\{\mathsf{A}, \overline{\mathsf{A}}\}$, and the left disjunct is $\neg A$, then out of the two relevant alternatives only $\neg A$ is a candidate restrictor, i.e., $C(\neg A) = \{\neg A\}$. For every cell that is update-consistent with the disjunct $\neg A$ is update-consistent with the relevant alternative $\neg A$, while A is eliminated as a candidate restrictor because the disjunct $\neg A$ is update-consistent with the $\overline{\mathsf{A}}$-cell but A is not ($\mathsf{UC}_\Pi(\neg A) \nsubseteq \mathsf{UC}_\Pi(A)$).

A further instructive case involves disjunctions of the form $\ulcorner(A \text{ or } B) \text{ or } C\urcorner$. Suppose that the partition Π is $\{\mathsf{AB}, \mathsf{A}\overline{\mathsf{B}}, \overline{\mathsf{A}}\mathsf{B}, \overline{\mathsf{AB}}\}$, then the left disjunct has the following formal alternatives all of which are relevant:[18]

(37) $Alt_R(A \vee B) = \{A \vee B, A, B\}$

It's obvious that the original disjunct $A \vee B$ is a candidate restrictor, since $\mathsf{UC}_\Pi(A \vee B) = \mathsf{UC}_\Pi(A \vee B)$. But the two disjuncts are not candidate restrictors, since $\mathsf{UC}_\Pi(A \vee B) = \{\mathsf{AB}, \mathsf{A}\overline{\mathsf{B}}, \overline{\mathsf{A}}\mathsf{B}\}$ and $\mathsf{UC}_\Pi(A) = \{\mathsf{AB}, \mathsf{A}\overline{\mathsf{B}}\}$ while $\mathsf{UC}_\Pi(B) = \{\mathsf{AB}, \overline{\mathsf{A}}\mathsf{B}\}$, and so $\mathsf{UC}_\Pi(A \vee B) \nsubseteq \mathsf{UC}_\Pi(A)$ and $\mathsf{UC}_\Pi(A \vee B) \nsubseteq \mathsf{UC}_\Pi(B)$. Therefore, the candidate restrictor set is a singleton:

(38) $C(A \vee B) = \{A \vee B\}$

So, the constraint in (35) predicts that the local context for $\ulcorner(A \text{ or } B) \text{ or } C\urcorner$ is updated by $\neg(A \vee B)$, given the above partition. This prediction seems desirable; intuitively, the disjunction in (39) suggests that if neither Alice nor Bob had signed the form, it would've been approved.

(39) The form was signed by either Alice or Bob, or it would not have been approved.

a. $\rightsquigarrow$ If the form had been signed by neither Alice or Bob, it would not have been approved.

b. $\not\rightsquigarrow$ If the form had not been signed by Alice, it would not have been approved.

[18] This assumes that the interpretation function $[\![\cdot]\!]$ assigns disjunctions classical Boolean truth conditions (recall that relevance checks if a proposition is a partial or complete answer). How the Boolean semantics and the update potential is connected is known as the explanatory problem (Schlenker, 2009), which we cannot address.

This is precisely the prediction our algorithm yields.

However, disjunctions of the form $\ulcorner(A \wedge B)$ or $C\urcorner$ complicate matters. Suppose that the partition Π is $\{\mathsf{AB}, \mathsf{A\overline{B}}, \mathsf{\overline{A}B}, \mathsf{\overline{AB}}\}$. Then the set of the left disjunct's relevant alternatives is $\{A, B, A \wedge B\}$. But, it is easy to see that the candidate restrictor set $\mathcal{C}(A \wedge B)$ is also $\{A, B, A \wedge B\}$. For each of the relevant alternative p, $\mathsf{UC}_\Pi(A \wedge B) \subseteq \mathsf{UC}_\Pi(p)$. In some sense, we run into an overdetermination problem as for which relevant alternative leads to the restrictor proposition.

To solve this problem, we take the restrictor to be the conjunctive closure of the candidate set based on a disjunct A: $\bigwedge \mathcal{C}(A)$. So, the restrictor proposition given $A \wedge B$ in this case is simply $[\![A \wedge B]\!] = [\![A]\!] \cap [\![B]\!] \cap [\![A \wedge B]\!]$.

More generally, we propose that the restrictor proposition given a disjunct A, written $[\![\dagger_A]\!]$, to be the conjunctive closure of the candidate set, assuming it exists:[19]

(40) $$[\![\dagger_A]\!] = \begin{cases} \bigcap_{p \in \mathcal{C}(A)} [\![p]\!] & \text{if } \mathcal{C}(A) \neq \emptyset \\ \bot & \text{otherwise} \end{cases}$$

Given our two-step procedure for finding $\dagger$, we can state the following update rule for disjunction:

(41) $c[\varphi \text{ or } \psi] = c[\neg\dagger_\psi][\varphi] \cup c[\neg\dagger_\varphi][\psi] \cup c[\dagger_\varphi \wedge \dagger_\psi][\varphi \wedge \psi]$[20]

The third conjunctive update $c[\dagger_\varphi \wedge \dagger_\psi][\varphi \wedge \psi]$ is needed. If we left it out from the entry, we would predict that *or* performs an exclusive update: assuming both disjuncts are relevant, after updating with "Bill went to New York or Boston", the context would also accept that Bill didn't go to both; but this is not desired especially when the disjunction gets negated. Including the conjunctive update circumvents this problem.

Since we usually work with $\dagger_A$'s negation, it's useful to spell out the local context update:

(42) $$c[\neg\dagger_A] = c \backslash c[\dagger_A] = \begin{cases} c \cap \bigcup_{\varphi \in \mathcal{C}(A)} [\![\neg\varphi]\!] & \text{if } \mathcal{C}(A) \neq \emptyset \\ c \cap \top & \text{otherwise} \end{cases}$$

[19] For ease of presentation, we take $\dagger$ as a covert object language operator, but it need not be. And we will assume that it only evokes the standard intersective update rule (fn. 1).

[20] For our purpose here, whether a conjunction evokes the intersective update rule (fn. 1) or a dynamic update rule (as in Heim, 1983) does not matter. For concreteness, we can assume the former.

Note that our algorithm has the nice prediction that, when the disjunct itself is relevant, it will be the restrictor of the other disjunct.

We make a general prediction about disjuncts with epistemic modals. Since epistemic modal claims don't count as relevant by our definition of relevance, we predict that restrictors based on simple epistemic modal claims are their prejacents (when they're relevant), or $\bot$ (when the prejacents aren't relevant).

Let's briefly consider the predictions our algorithm makes about pronominal anaphora. Consider a context where (43b) is uttered to address the QUD in (43a).

(43) a. **QUD:** Is there a bathroom in the house? $\Pi = \{\mathsf{B}, \overline{\mathsf{B}}\}$

b. Either there's no^x bathroom in the house, or it$_x$'s in a funny place.

The left disjunct has two relevant alternatives:

(44) $\{\neg\exists x[\mathbf{bathroom}(x)], \exists x[\mathbf{bathroom}(x)]\}$

However, we can rule out the negation-free alternative because, the original disjunct is update-consistent with the $\overline{\mathsf{B}}$-cell in the partition, but that alternative is not. So, the right disjunct's restrictor is $\neg\exists x[\mathbf{bathroom}(x)]$, and its local context entails the existence of a bathroom because $\neg\neg\exists x[\mathbf{bathroom}(x)] \Leftrightarrow \exists x[\mathbf{bathroom}(x)]$.[21] The exact same reasoning applies to (45), a case with prejacent restriction.

(45) There must be no^x bathroom in this house, or it$_x$'s in a funny place.

Here too the only candidate restrictor is $\neg\exists x[\mathbf{bathroom}(x)]$ because $\Box\neg\exists x[\mathbf{bathroom}(x)]$ is irrelevant given the QUD. Hence, we predict that (45) should license anaphora.

3.3.1 *Deriving asymmetric local contexts: presupposition satisfaction*

Let's start by highlighting one novel prediction of our account – it captures left-to-right and right-to-left symmetric filtering of presuppositions. First, the disjunction in (46) only presupposes (46a).

[21] To derive the licensing, this requires a dynamic system that validates double negation elimination, see e.g. Elliott and Sudo (2025). If we follow Heim (1982) and assume that *it* carries a familiarity presupposition, then the right disjunct would by itself be infelicitous, and not relevant, resulting in trivial restriction for the left disjunct.

(46) Either John never smoked, or John stopped smoking.

a. $\rightsquigarrow$ If John smoked previously, he smoked previously.
b. $\not\rightsquigarrow$ John smoked previously.

For warm-up, let's spell out in more detail how the standard entry in (2), repeated as (47), predicts this conditional presupposition.

(47) $c[\varphi \text{ or } \psi] := c[\varphi] \cup c[\neg\varphi][\psi]$

Presuppositions are understood as definedness conditions for updates. Putting things in the terms of our acceptance-based framework, for $c[A_p]$ to go through, c needs to accept A_p's presupposition p (Veltman, 1996). Since in an update with a disjunction, the right disjunct ψ is said to update the local context $c[\neg\varphi]$, $c[\neg\varphi]$ needs to accept ψ's presupposition. That is equivalent to the requirement that c needs to accept $\neg\varphi \rightarrow \psi$, given our informational entry for indicatives in (24). Hence (46a) is presupposed by (46).

Note that there's an intuitive asymmetry between the local contexts of the two disjuncts; the left disjunct here restricts the local context of the right disjunct but not the other way around. This asymmetry is hardwired into (47).

Given our system, this asymmetry arises when one disjunct is irrelevant to the QUD; as a consequence, the restrictor $\dagger$ it contributes to the other disjunct's local context will be $\bot$ – i.e., the empty proposition. For example, suppose that the disjunction in (46) addresses the polar question "Was John ever a smoker?" In such a context, (46)'s left disjunct is relevant to the QUD, while its right disjunct is irrelevant to the QUD; as a result, the restrictor based on the left disjunct $\dagger_{left}$ is the left disjunct itself while $\dagger_{right}$ is $\bot$. We thus predict the following update, agreeing with the standard theory and thus predicting the same presupposition (with the third update empty):

(48) $c[(46)] = c[\top][\neg\textbf{smoke}(j)] \cup c[\textbf{smoke}(j)][\textbf{stopped}(j)] \cup c[\neg\textbf{smoke}(j)][\textbf{stopped}(j)]$

We believe the asymmetry is justified conceptually. The question of whether John stopped smoking presupposes that he used to smoke and is only relevant against a context entailing that presupposition. Relative to a context that's uncertain about whether he smoked before, the right disjunct is not (yet) relevant.

Unlike the asymmetric entry (47), the exact same reasoning also explains how (49) carries the same filtered presupposition. For, holding fixed the QUD, our account predicts the same update potential for disjunctions no matter which disjunct comes first.

(49) Either John stopped smoking, or he never smoked.

a. $\rightsquigarrow$ If John smoked previously, he smoked previously.

Thus, we predict the same update potential as above (omitting the empty third update)

(50) $c[(49)] = c[\mathbf{smoke}(j)][\mathbf{stopped}(j)] \cup c[\top][\neg\mathbf{smoke}(j)]$

And a context accepting this disjunction, again, accepts the conditional "If John smoked previously, he stopped".

The same reasoning also applies to (9), given the same QUD, repeated as (51).

(51) John must have never smoked, or he stopped smoking.

The restrictor based on the right disjunct is, by the same reasoning as above, $\bot$. The restrictor based on the left disjunct is $\neg\mathbf{smoke}(j)$. The left disjunct itself $\Box\neg\mathbf{smoke}(j)$ cannot be relevant, and of the two relevant alternatives $\neg\mathbf{smoke}(j)$ and $\mathbf{smoke}(j)$, only the former counts as a candidate restrictor. This is intuitively correct, as we observe local satisfaction for the right disjunct, suggesting that its local context entails $\mathbf{smoke}(j)$.

Finally, consider an example involving prejacent restriction with root modals, repeated from (10).

(52) We can't build a stadium here, or the stadium will make the traffic worse.

The left disjunct's relevant alternatives will be $\neg\mathbf{build\_stadium}(u)$ and $\mathbf{build\_stadium}(u)$ when the disjunction addresses the QUD "Will the stadium be built here?". The original left disjunct "we can't build a stadium here," is only update-consistent with the negative answer, assuming it's common ground that prohibited stadium construction won't occur. Thus, a candidate restrictor needs to also be update-consistent with the negative answer. The negation-free alternative can't be a candidate restrictor, since it is only update-consistent with the positive answer "the stadium will be built here". In contrast, $\neg\mathbf{build\_stadium}(u)$ passes the test for candidate restrictors. Therefore, we correctly predict the right disjunct's local context:

(53) $c[\neg\neg\textbf{build\_stadium}(u)] = c[\textbf{build\_stadium}(u)]$

This predicts that (52) doesn't carry any substantive global presupposition. Since the right disjunct is not relevant, the local context for the left disjunct is the entire global context set c.

That said, we do suspect that when the order of disjuncts is reversed, which questions are relevant often changes. A natural hypothesis is that the left disjunct needs to be relevant by default, while the right disjunct doesn't have to be. When there is no overt question, the hearer will often accommodate a question whose answers correspond to the left disjunct.

3.3.2 Deriving symmetric local contexts: doubly-modal disjunctions

When it comes to disjunctions where no disjunct depends on the other one, we often have genuine symmetric local contexts. This is most clearly seen in the case of doubly modal disjunctions (DMDs). For example, (27), repeated as (54), intuitively licenses conditional inferences about both disjuncts, where the antecedent indicates the restriction of the modal domains.[22]

(54) John must be in New York, or he must be in Boston.

a. $\not\rightsquigarrow$ If John might not be in New York, he must be in Boston.
b. $\rightsquigarrow$ If John is not in New York, he must be in Boston.

In this subsection, we demonstrate how we derive symmetric prejacent restriction. In the next subsection, we show how to derive the inferences to the conditionals given prejacent restriction.

With natural assumptions, we derive the context update potential (58) for (57), with an empty conjunctive third update (John cannot be in both New York and Boston).

(55) $$c[(54)] = c[\neg\textbf{in\_boston}(j)][\Box\textbf{in\_nyc}(j)] \cup c[\neg\textbf{in\_nyc}(j)][\Box\textbf{in\_boston}(j)] \\ \cup c[\textbf{in\_boston}(j) \wedge \textbf{in\_nyc}(j)][\Box\textbf{in\_nyc}(j) \wedge \Box\textbf{in\_boston}(j)]$$

There are a few plausible QUDs which a speaker might use (57) to address; namely, either the polar question in (56a) together with (56b) or the *wh*-question in (56c).

[22] It's worth noting that Klinedinst and Rothschild (2012) suggest that in the following (their ex. (105))

i. John is probably in China, or he's in Singapore.

that John is probably in China is implied, while on our account only the conditional *if John's not in Singapore, he's probably in China* is implied. We think their observation is not obviously correct, because they would predict that (54) *John must be in Boston* is entailed, which would be false on the Yalcin-style semantics for *must* they assume. Therefore, we think the conditional inferences should be the baseline semantic inference, if the data form a unified class.

(56) a. Is John in New York? $\{\mathsf{N}, \overline{\mathsf{N}}\}$

b. Is John in Boston? $\{\mathsf{B}, \overline{\mathsf{B}}\}$

c. Where is John? $\{\mathsf{N}\overline{\mathsf{B}}, \overline{\mathsf{N}}\mathsf{B}\}$

Either way, the question(s) will always partition the context set into worlds where John is in New York and worlds where he is in Boston. This means that, for both disjuncts, the only formal alternative relevant to the QUD is the prejacent of the modal in that disjunct. Hence, the algorithm predicts that the restrictor based on the right disjunct $\dagger_{right}$ and the restrictor based on the left disjunct $\dagger_{left}$ are as follows:

$$[\![\dagger_{left}]\!] = [\![\mathbf{in\_nyc}(j)]\!]$$

$$[\![\dagger_{right}]\!] = [\![\mathbf{in\_boston}(j)]\!]$$

If we plug these values into our semantics, we arrive at the desired update potential in (58).

The exact same reasoning predicts that (57) has the update potential (58).

(57) John might be in New York, or he must be in Boston.

(58) $c[(57)] = c[\neg\mathbf{in\_boston}(j)][\Diamond\mathbf{in\_nyc}(j)] \cup c[\neg\mathbf{in\_nyc}(j)][\Box\mathbf{in\_boston}(j)]$
$\cup c[\mathbf{in\_boston}(j) \wedge \mathbf{in\_nyc}(j)][\Diamond\mathbf{in\_nyc}(j) \wedge \Box\mathbf{in\_boston}(j)]$

Our account also derives desirable predictions concerning DMDs with logically complex disjuncts. Consider the following example adapted from an example due to Cariani (2017):[23]

(59) Jones might be eating lunch, or he might be teaching, or he must be in his office.

a. $\rightsquigarrow$ If Jones is neither eating lunch nor teaching, he must be in his office.

b. $\not\rightsquigarrow$ If Jones isn't eating lunch, he must be in his office.

c. $\not\rightsquigarrow$ If Jones isn't teaching, he must be in his office.

The intuitive *or*-to-*if*-style inference (59a) suggests that the restrictor for the rightmost disjunct is the disjunction of the prejacents of the two *might*-claims. And our account derives the desired meaning of (59) and similar examples given the following QUD:

[23] Cariani presents examples like (59) as a counterexample to Klinedinst and Rothschild's (2012) proposal that, in DMDs, the right disjunct's information state can be restricted by any sub-clause of the left disjunct. Meyer's (2015) account of DMDs which allows the right disjunct's modal domain to be restricted by any formal alternative to the left disjunct's information state might predict the intuitive reading, but seems to overgenerate other readings. See our earlier work for further arguments against extant analyses of DMDs (XXXX, XXXX).

(60) Where is John? $\Pi = \{E\overline{TO}), \overline{E}T\overline{O}, \overline{ET}O\}$

The left disjunct of (59) has the following relevant alternatives:

(61) $\{\mathbf{is\_eating}(j) \vee \mathbf{is\_teaching}(j), \mathbf{is\_eating}(j), \mathbf{is\_teaching}(j)\}$

Only the disjunctive one is a candidate restrictor for the right disjunct, and thus the restrictor proposition is simply $[\![\mathbf{is\_eating}(j) \vee \mathbf{is\_teaching}(j)]\!]$. This leads to the context update potential:

(62) $$c[(59)] = c[\neg\mathbf{in\_office}(j)][\Diamond\mathbf{is\_eating}(j) \vee \Diamond\mathbf{is\_teaching}(j)]$$
$$\cup c[\neg\mathbf{is\_eating}(j) \wedge \neg\mathbf{is\_teaching}(j)][\Box\mathbf{in\_office}(j)]$$
$$\cup c[\mathbf{in\_office}(j) \wedge \neg(\neg\mathbf{is\_eating}(j) \wedge \neg\mathbf{is\_teaching}(j))][(\Diamond\mathbf{is\_eating}(j) \vee$$
$$\Diamond\mathbf{is\_teaching}(j)) \wedge \Box\mathbf{in\_office}(j)]$$

where the restrictor of the right disjunct is derived by de Morgan's laws, applied to the restrictor proposition. Again, the conjunctive update is empty.

Let's consider our predictions about data involving root modals, starting with (32), repeated as (63c). As we suggested, in the context given in (32), the QUD crucially includes the polar question Q_1 in (63a). And we assume that the right disjunct addresses Q_2 below.[24] Hence, the partition generated will be $\{\Box W \wedge A, \neg\Box W \wedge A, \Box W \wedge \neg A, \neg\Box W \wedge \neg A\}$ (W: **write**(j), A: **get-A**(j)).

(63) a. $\mathbf{Q}_1$**:** Is there a requirement to write a paper?
b. $\mathbf{Q}_2$: Will John get an A?
c. John has to write a paper, or he will get an A.

Consider the set of candidate restrictors based on the left disjunct. The only relevant alternative of $\Box\mathbf{write}(j)$ is itself, which meets the constraint for candidate restrictors. And so $\Box\mathbf{write}(j)$ is the restrictor for the right disjunct. Similarly, the restrictor based on the right disjunct is $\mathbf{get\text{-}A}(j)$. Thus, we predict that (63c) leads to the following update:

(64) $$c[(63c)] = c[\neg\mathbf{get\text{-}A}(j)][\Box\mathbf{write}(j)] \cup c[\neg\Box\mathbf{write}(j)][\textsc{woll}\ \mathbf{get\text{-}A}(j)]$$
$$\cup c[\Box\mathbf{write}(j) \wedge \mathbf{get\text{-}A}(j)][\Box\mathbf{write}(j) \wedge \textsc{woll}\ \mathbf{get\text{-}A}(j)]$$

A different set of questions seem to be more salient for (65c), where in the given context there is no uncertainty about what is required; instead, the uncertainty is about whether John will do what's required and write a paper. Hence the partition Π will be $\{WC, \overline{W}C, W\overline{C}, \overline{WC}\}$.

[24] As a simplification, we assume that the answers don't involve WOLL, but this giving up this assumption doesn't change the result as long as we assume WOLL is an alternative to teleological $\Box$. Moreover, this assumption is more or less harmless as WOLL doesn't seem to involve genuine modal quantification (Cariani and Santorio, 2018).

(65) a. $\mathbf{Q}_1$**:** Will John write a paper?

b. $\mathbf{Q}_2$: Will John get an C?

c. John has to write a paper, or he will get an C.

Consider the set of candidate restrictors based on the left disjunct. The only relevant alternative of the left disjunct is $\mathbf{write}(j)$. And there are two cells (WC and W$\overline{\text{C}}$) that are update-consistent with $\mathbf{write}(j)$. In a default context, both cells are naturally update-consistent with $\Box\mathbf{write}(j)$, i.e., that John has to (teleologically) write a paper. Hence $\mathsf{UC}_\Pi(\Box\mathbf{write}(j)) \subseteq \mathsf{UC}_\Pi(\mathbf{write}(j))$ and $\mathbf{write}(j)$ will be the candidate restrictor for the right disjunct. And $\mathbf{get\text{-}C}(j)$ will be the restrictor for the left disjunct. Therefore, we predict the following context update potential for (65c).

(66) $c[(65c)] = c[\neg\mathbf{get\text{-}C}(j)][\Box\mathbf{write}(j)] \cup c[\neg\mathbf{write}(j)][\textsc{woll}\ \mathbf{get\text{-}C}(j)]$

$$\cup c[\mathbf{get\text{-}C}(j) \land \mathbf{write}(j)][\Box\mathbf{write}(j) \land \textsc{woll}\ \mathbf{get\text{-}C}(j)]$$

Finally, consider example (67) from Simons (2005), which, in addition to a reading expressing uncertainty about Mary's obligation, has a reading that implies (67a) and (67b).

(67) Mary must sing, or she must dance.

a. $\rightsquigarrow$If Mary doesn't dance, she must sing.

b. $\rightsquigarrow$If Mary doesn't sing, she must dance.

These inferences arise when (67) addresses the following QUD:

(68) What will Mary do? $\Pi = \{\text{SD}, \text{S}\overline{\text{D}}, \overline{\text{S}}\text{D}, \overline{\text{SD}}\}$

Given this QUD, the only relevant alternative to the left disjunct $\Box\mathbf{sing}(m)$ is its prejacent $\mathbf{sing}(m)$, which is update-consistent with the SD and S$\overline{\text{D}}$ cells. If at least one of the cells is update-consistent with the original disjunct $\Box\mathbf{sing}(m)$, $\Box\mathbf{sing}(m)$ will the only candidate restrictor for the right disjunct. And the condition is plausibly met in any natural context, since Mary's singing is typically compatible with her being required to sing. Likewise, when similar conditions are met, $\mathbf{dance}(m)$ will be the left disjunct's restrictor. Hence, given the QUD described above, we predict that (67) has the following update potential in typical contexts:

(69) $c[(67)] = c[\neg\mathbf{dance}(m)][\Box\mathbf{sing}(m)] \cup c[\neg\mathbf{sing}(m)][\Box\mathbf{dance}(m)]$

$$\cup[\mathbf{sing}(m) \land \mathbf{dance}(m)][\Box\mathbf{sing}(m) \land \Box\mathbf{dance}(m)]$$

Hence, the current account predicts that prejacent restriction is possible in (67).

3.3.3 *Or-to-if failures and Generalized Free Choice*

Now that we've shown our predictions of when prejacent restriction arises, let's move on to the failure of the *or*-to-if inference. The failure of the *or*-to-if inference is closely connected to a phenomenon Meyer (2015) dubs "Generalized Free Choice." Meyer observes that, in a doubly-modal disjunction, speakers often draw the following conjunctive inference:[25]

Generalized Free Choice. $\triangle_1\varphi \vee \triangle_2\psi \rightsquigarrow (\neg\psi \to \triangle_1\varphi) \wedge (\neg\varphi \to \triangle_2\psi)$

What's particularly striking about Generalized Free Choice is that the inferred conditionals often have the negation of one modal disjunct's prejacent as its antecedent, and the other disjunct as its consequent; they involve prejacent restriction. For example:

(70) John must be in New York, or he must be in Boston.

a. $\rightsquigarrow$ If John is not in New York, he must be in Boston.
b. $\rightsquigarrow$ If John is not in Boston, he must be in New York.

In other words, Generalized Free Choice inferences are licensed only when *or*-to-*if* fails. Meyer derives Generalized Free Choice as an implicature via scalar strengthening.[26] This may be supported by the observation that Generalized Free Choice inferences appear conjunctive in matrix cases, but appear disjunctive in downward-entailing environments.[27]

(71) John doubts that he has to write a paper, or he will get a C.

a. $\not\rightsquigarrow$ John believes that [he either doesn't have to write a paper to not get a C or if he doesn't write a paper, he won't get a C]
b. $\rightsquigarrow$ John believes that [he both doesn't have to write a paper to not get a C and if he doesn't write a paper, he won't get a C]

However, we think the dynamic framework and the notion of acceptance offer a more parsimonious story: when the doubly modal disjunctions in question are accepted, the conditionals are

[25]Note that the two disjuncts may involve modals of different flavor and/or force. This is significant because some existing accounts of modal disjunction for sentences like (67) postulate that $\Box A \vee \Box B$ can be interpreted as $\Box(A \vee B)$ (similarly for $\Diamond$) (Simons, 2005; Ciardelli and Guerrini, 2026). But that type of account crucially only applies when the two disjuncts contain the same modal and doesn't apply to many instances Generalized Free Choice. That said, we leave it open whether that account is correct for sentences like (67).

[26]It's also worth noting that Klinedinst and Rothschild (2012) derive the inferences to the conditionals via the hypothesis that *or* has a conjunctive entry.

[27]Here we make the standard assumption that *doubt* $\approx$ *believe not* in (71).

also accepted. Within this story, the force of the disjunction in downward-entailing environments doesn't need a special explanation: disjunctions only ever have disjunctive force.

In general, Generalized Free Choice can be derived from prejacent restriction and the following fact about acceptance, given our entry:

(72) **Proposition 1** (*Cautious Or-to-If*) For any c, if c accepts $\ulcorner\varphi$ or $\psi\urcorner$, and $c[\neg\dagger_\psi]\cap c[\neg\dagger_\varphi][\psi] = \varnothing$ and $c[\neg\dagger_\psi][\varphi] \cap c[\neg\dagger_\varphi] = \varnothing$, then c accepts both $\neg\dagger_\psi \to \varphi$ and $\neg\dagger_\varphi \to \psi$.

In other words, if c accepts $\ulcorner\varphi$ or $\psi\urcorner$, and thus $c = c[\neg\dagger_\psi][\varphi]\cup c[\neg\dagger_\varphi][\psi]\cup c[\dagger_\varphi\wedge\dagger_\psi][\varphi\wedge\psi]$, then c accepts both $\neg\dagger_\psi \to \varphi$ and $\neg\dagger_\varphi \to \psi$ if $c[\neg\dagger_\psi] \cap c[\neg\dagger_\varphi][\psi] = \varnothing$ and $c[\neg\dagger_\psi][\varphi] \cap c[\neg\dagger_\varphi] = \varnothing$.

There is a stronger result that holds for disjunctions involving epistemic modals, given the update semantics we assume. Our account predicts that epistemic modals always lead to prejacent restriction (or trivial restriction), hence we have:

(73) **Proposition 2** (*Or-to-if with epistemic modals*)

For any c, if c accepts $\ulcorner\triangle_1 A$ or $\triangle_2 B\urcorner$ with $\triangle_1, \triangle_2 \in \{\Diamond_{epi}, \Box_{epi}\}$ via $c = c[\neg B][\triangle_1 A]\cup c[\neg A][\triangle_2 B]\cup c[A\wedge B][\triangle_1 A \wedge \triangle_2 B]$, then c accepts both $\neg B \to \triangle_1 A$ and $\neg A \to \triangle_2 B$.

This is because it's guaranteed that $c[\neg B][\triangle_1 A] \cap c[\neg A] = \varnothing$ and $c[\neg B] \cap c[\neg A][\triangle_2 B] = \varnothing$ given update semantics. This immediately explains the inferences (70) licenses.

(74) John must be in New York, or he must be in Boston.

a. $c[(74)] = c[\neg\textbf{in\_boston}(j)][\Box\textbf{in\_nyc}(j)] \cup c[\neg\textbf{in\_nyc}(j)][\Box\textbf{in\_boston}(j)]$

b. $\rightsquigarrow$ If John is not in New York, he must be in Boston.

c. $\rightsquigarrow$ If John is not in Boston, he must be in New York.

By similar reasoning, it's easy to verify that the following two conditional inferences are accepted whenever (59) (repeated as (75)) is, because of the context update potential we derived (omitting the empty conjunctive update).

(75) Jones might be eating lunch, or he might be teaching, or he must be in his office.

a. $c[(75)] = c[\neg\textbf{in\_office}(j)][\Diamond\textbf{is\_eating}(j) \vee \Diamond\textbf{is\_teaching}(j)]\cup$ $c[\neg\textbf{is\_eating}(j) \wedge \neg\textbf{is\_teaching}(j)][\Box\textbf{in\_office}(j)]$

b. $\rightsquigarrow$ If Jones is neither eating lunch nor teaching, he must be in his office.

c. $\rightsquigarrow$ If Jones is not in his office, he might be eating lunch or he might be teaching.

The inferences we can draw with root modals are sometimes weaker, depending on whether a veridical modal is involved.[28] Some *or*-to-*if*-style inferences with prejacent restriction involving root modals will be acceptance-preserving, while others will be licensed only in certain contexts. For example, we predicted that (63c), repeated as (76a), has the update potential in (76b) given the questions in (63a) about whether John has to write a paper and whether he will get an A.

(76) a. John has to write a paper, or he will get an A.

b. $c[(76a)] = c[\neg\textbf{get-A}(j)][\Box\textbf{write}(j)] \cup c[\neg\Box\textbf{write}(j)][\textsc{woll}\ \textbf{get-A}(j)]$

It's easy to check that, assuming $\textsc{woll}$ is veridical, $\neg\textbf{get-A}(j) \rightarrow \Box\textbf{write}(j)$ and $\neg\Box\textbf{write}(j) \rightarrow \textsc{woll}\ \textbf{get-A}(j)$ are accepted by any context that accepts (76a). This is just an instance of *Cautious Or-to-If.*

Next, consider (65c), repeated as (77a), addressing the questions about whether John will write a paper and if he will get a C. It has the update potential (77b).

(77) a. John has to write a paper, or he will get a C.

b. $c[(77a)] = c[\neg\textbf{get-C}(j)][\Box\textbf{write}(j)] \cup c[\neg\textbf{write}(j)][\textsc{woll}\ \textbf{get-C}(j)]$
$\cup c[\textbf{get-C}(j) \wedge \textbf{write}(j)][\Box\textbf{write}(j) \wedge \textsc{woll}\ \textbf{get-C}(j)]$

Again, assuming veridicality of $\textsc{woll}$, accepting the disjunction guarantees accepting $\neg\textbf{get-C}(j) \rightarrow \Box\textbf{write}(j)$ and $\neg\textbf{write}(j) \rightarrow \textsc{woll}\ \textbf{get-C}(j)$.

Finally, consider Simons's (2005) example in (78), where the conditional inferences depend on the context. There's a natural reading of (78) where it implies both (78a) and (78b).

(78) Mary must sing, or she must dance.

a. $\rightsquigarrow$If Mary doesn't dance, she must sing.

b. $\rightsquigarrow$If Mary doesn't sing, she must dance.

We've shown that (78) can have the following meaning if it addresses the question "What will Mary do?":

(79) $c[(78)] = c[\neg\textbf{dance}(m)][\Box\textbf{sing}(m)] \cup c[\neg\textbf{sing}(m)][\Box\textbf{dance}(m)]$
$\cup[\textbf{sing}(m) \wedge \textbf{dance}(m)][\Box\textbf{sing}(m) \wedge \Box\textbf{dance}(m)]$

[28]We assume that root modals receive a standard Kratzer-style (1981) semantics for root modals and treat their associated updates as intersective. There are several options about linking them with local contexts: one could extend test semantics to root modals along the lines of Goldstein (2019) or supplement a static semantics for root modals with principles linking local contexts and modal domains (Mandelkern, 2019, 2024).

Here's one way to derive the conditional inferences. We can assume that an update of a local context with deontic *must* is information-sensitive. It is sensitive to the local context in that, an update with it is defined only if the modal domain from every world in the context entails the context, i.e., it contains the information encoded in the local context (cf. Kolodny and MacFarlane, 2010; Klinedinst and Rothschild, 2012; Mandelkern, 2019, 2024). Given *Cautious Or-to-If*, both conditional inferences (78a) and (78b) are licensed by a context that accepts the disjunction, if the modal updates are all defined. The definedness condition requires that $\forall w' \in c[\neg\mathbf{dance}(m)] : R(w') \subseteq c[\neg\mathbf{dance}(m)] \subseteq [\![\neg\mathbf{dance}(m)]\!]$ and $\forall w' \in c[\neg\mathbf{sing}(m)] : R(w') \subseteq c[\neg\mathbf{sing}(m)] \subseteq [\![\neg\mathbf{sing}(m)]\!]$. Thus, $c[\neg\mathbf{dance}(m)] \cap [\neg\mathbf{sing}(m)][\Box\mathbf{dance}(m)] = \emptyset$ and $c[\neg\mathbf{dance}(m)][\Box\mathbf{sing}(m)] \cap [\neg\mathbf{sing}(m)] = \emptyset$.[29] Hence, the inferences in (78) will follow from *Cautious Or-to-If.*

4 Conclusion

In this paper, we have argued that standard dynamic theories of disjunctions are mistaken. While the notion of local context is useful for presupposition filtering, anaphora licensing and *or*-to-*if* and modal domain restriction, that the local context of a disjunct always entails the negation of the other disjunct (i.e., that a disjunct always has the other disjunct as its restrictor) is not correct. We give a new theory of local contexts in disjunction where the restrictor of a disjunct is a particular relevant alternative of the other disjunct. This new theory accounts for the data the standard theory fails to capture. In particular, we showed that while *or*-to-*if* is not, in general, acceptance-preserving, we can have a generalization that predicts more generally when and what conditional inferences are licensed by a disjunction, which we called *Cautious Or-to-If.* For reasons of space, we hinted that our account predicts the right local contexts for Partee's bathroom disjunctions, without working through the details. Some adaptation of the current account to a more complex system designed to deal with Partee disjunctions would be needed for that purpose, which we

[29] We are assuming that the two deontic modals involve the same accessibility relation R which has consistent outputs given worlds in c. And suppose all the modal updates are defined. If there was an arbitrary world $w \in c[\neg\mathbf{dance}(m)]$ and $w \in c[\neg\mathbf{sing}(m)][\Box\mathbf{dance}(m)]$, we'd have $R(w) \subseteq [\![\neg\mathbf{dance}(m)]\!]$ and $R(w) \subseteq [\![\neg\mathbf{sing}(m) \cap \mathbf{dance}(m)]\!]$, and $R(w)$ will have to be $\emptyset$, which contradicts the consistency assumption. Thus $c[\neg\mathbf{dance}(m)] \cap [\neg\mathbf{sing}(m)][\Box\mathbf{dance}(m)] = \emptyset$. Similarly, $c[\neg\mathbf{dance}(m)][\Box\mathbf{sing}(m)] \cap [\neg\mathbf{sing}(m)] = \emptyset$.

leave for future work.

References

Beaver, D. I. (2001). *Presupposition and assertion in dynamic semantics*, Volume 29. CSLI publications Stanford.

Cariani, F. (2017). Choice points for a modal theory of disjunction. *Topoi 36*(1), 171–181.

Cariani, F. and P. Santorio (2018). Will done better: Selection semantics, future credence, and indeterminacy. *Mind 127*(505), 129–165.

Chatain, K. (2025). Reducing pronoun accessibility to presupposition satisfaction. *Journal of Semantics*, ffaf006.

Ciardelli, I. and J. Guerrini (2026). Against wide-scope free choice. *Semantics and Pragmatics 19*, 4–EA.

Dever, J. (2013). The revenge of the semantics-pragmatics distinction. *Philosophical Perspectives 27*, 104–144.

Égré, P., L. Rossi, and J. Sprenger (2025). Certain and uncertain inference with indicative conditionals. *Australasian Journal of Philosophy 103*(3), 569–596.

Elliott, P. D. (2022). Partee conjunctions. Unpublished manuscript, Heinrich-Heine-Universität Düsseldorf, https://ling.auf.net/lingbuzz/006857.

Elliott, P. D. and Y. Sudo (2025). Free choice with anaphora. *Semantics and Pragmatics 18*, 2–1.

Fox, D. and R. Katzir (2011). On the characterization of alternatives. *Natural Language Semantics 19*(1), 87–107.

Geurts, B. (1996). Local satisfaction guaranteed: A presupposition theory and its problems. *Linguistics and Philosophy*, 259–294.

Geurts, B. (1998). Presuppositions and anaphors in attitude contexts. *Linguistics and Philosophy 21*(6), 545–601.

Gillies, A. S. (2004). Epistemic conditionals and conditional epistemics. *Noûs 38*(4), 585–616.

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

Groenendijk, J. and M. Stokhof (1984). *Studies on the Semantics of Questions and the Pragmatics of Answers.* Ph. D. thesis, University of Amsterdam.

Heim, I. (1983). On the projection problem for presuppositions. In *Formal semantics–the essential readings*, pp. 249–260. John Wiley & Sons.

Heim, I. R. (1982). *The semantics of definite and indefinite noun phrases.* Ph. D. thesis, University of Massachusetts Amherst.

Kalomoiros, A. and F. Schwarz (2021). Presupposition projection from disjunction is symmetric. *Proceedings of the Linguistic Society of America 6*(1), 556–571.

Karttunen, L. (1973). Presuppositions of compound sentences. *Linguistic Inquiry 4*(2), 169–193.

Karttunen, L. (1974). Presupposition and linguistic context. *Theoretical Linguistics 1*, 181–194.

Katzir, R. (2007). Structurally-defined alternatives. *Linguistics and Philosophy 30*, 669–690.

Kaufmann, S. (2005). Conditional truth and future reference. *Journal of semantics 22*(3), 231–280.

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 (2012). Connectives without truth tables. *Natural Language Semantics 20*, 137–175.

Kolodny, N. and J. MacFarlane (2010). Ifs and oughts. *The Journal of philosophy 107*(3), 115–143.

Kratzer, A. (1981). The notional category of modality. *Formal semantics: The essential readings*, 289–323.

Lewis, D. (1988). Relevant implication. *Theoria 54*(3), 161–174.

Mandelkern, M. (2019). Bounded modality. *Philosophical Review 128*(1), 1–61.

Mandelkern, M. (2022). Witnesses. *Linguistics and Philosophy 45*(5), 1091–1117.

Mandelkern, M. (2024). *Bounded meaning: The dynamics of interpretation.* Oxford University

Press.

Meyer, M.-C. (2015). Generalized free choice and missing alternatives. *Journal of Semantics 33*(4), 703–754.

Roberts, C. (2012). Information structure: Towards an integrated formal theory of pragmatics. *Semantics and pragmatics 5*, 6–1.

Santorio, P. (2022). Path semantics for indicative conditionals. *Mind 131*(521), 59–98.

Schlenker, P. (2009). Local contexts. *Semantics & Pragmatics 2*(3), 1–78.

Simons, M. (2005). Dividing things up: The semantics of or and the modal/or interaction. *Natural Language Semantics 13*(3), 271–316.

Singh, R. (2007). Formal alternatives as a solution to the proviso problem. In *Semantics and Linguistic Theory (SALT)*, pp. 264–281.

Singh, R. (2008). *Modularity and Locality in Interpretation.* Ph. D. thesis, Massachusetts Institute of Technology.

Stalnaker, R. (1974). Pragmatic presuppositions. In R. Stalnaker (Ed.), *Context and Content*, pp. 47–62. Oxford University Press.

Stalnaker, R. (1975). Indicative conditionals. *Philosophia 5*(3), 269–286.

Stalnaker, R. (1978). Assertion. In *Pragmatics*, pp. 315–332. Brill.

Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic 25*(3), 221–261.

Willer, M. (2015). An update on epistemic modals. *Journal of Philosophical Logic 44*(6), 835–849.

Yalcin, S. (2007). Epistemic modals. *Mind 116*(464), 983–1026.