

Neg-Raising through Question-Raising?

Wolfgang Schwarz

Abstract. I suggest that neg-raising may arise through a general mechanism by which negative utterances are enriched to convey positive information. Concretely, I explore whether the inference from $\neg P(\phi)$ to $P(\neg\phi)$ can be explained by assuming that $\neg P(\phi)$ addresses a question like $\{ P(\phi), P(\neg\phi) \}$, which presupposes $\neg P(\phi) \rightarrow P(\neg\phi)$. I present evidence in favour of this conjecture, but also identify some challenges that an account of neg-raising based on this model would need to address.

1 Introduction

Negation often has a positive side-effect. (1a), for example, denies that Jack sang *Hey Jude*, but also conveys that he sang something else.

- (1) a. Jack didn't sing [*Hey Jude*]_F.
- b. $\neg$ Jack sang something other than *Hey Jude*.

This phenomenon is well-known (see, e.g., [Horn 1989: sec.1.2], [Kratzer 1989: sec.6]), but I don't think it has a canonical name. I'll call it "positive contrast". I want to explore whether the mechanism behind positive contrast also explains the phenomenon of neg-raising, exemplified in (2).

- (2) a. Bill doesn't think that Sue is smart.
- b. $\neg$ Bill thinks that Sue is not smart.

In both cases, a statement of the form $\neg P(\phi)$ conveys not only that $P(\phi)$ is false, but also that some alternative $P(\phi')$ to $P(\phi)$ is true, even though this is not predicted by the standard semantics of P . In both cases, the inference is robust and hard to cancel, yet it appears to be pragmatic, insofar as we don't judge (1a) to be false if Jack didn't sing anything at all, nor is (2a) false if Bill doesn't have an opinion about Sue's intellectual abilities. In both cases, the effect bears signs of a presupposition, but doesn't project the way presuppositions normally do.

There is no agreement on how positive contrast comes about. The effect appears to be a kind of dual to exhaustification: while exhaustification strengthens a proposition ϕ by negating certain alternatives ξ_i to ϕ , producing $\phi \wedge \bigwedge_i \neg \xi_i$, positive contrast strengthens $\neg\phi$ by the disjunction of certain alternatives ξ_i to ϕ , producing $\neg\phi \wedge \neg \bigwedge_i \neg \xi_i$. One might suggest that both effects are explained by the presence of a grammatical operator *Exh*. Concretely, suppose *Exh*(ϕ) triggers the presup-

position $\phi \leftrightarrow \bigwedge_i \neg \xi_i$.¹ Since presuppositions project out of negation, both $Exh(\phi)$ and $\neg Exh(\phi)$ would strengthen the asserted content by this presupposition, yielding the desired inferences. The neg-raising inference in (2) might then be derived in the same way. We'd assume that (2a) is interpreted as $\neg Exh(\text{Bill thinks Sue is smart})$, which would presuppose that a relevant alternative to *Bill thinks Sue is smart* is true. If all such alternatives imply that Bill has an unfavourable view of Sue's intellectual abilities, we'd get the inference in (2).²

In this paper, I want to explore an approach that follows a similar line of reasoning, but doesn't require a dedicated grammatical operator. Positive contrast effects like (1) are a staple in theories of focus and discourse structure in the tradition of [Rooth 1992], [Carlson 1983], [Roberts 1996], and [Beaver and Clark 2008]. Here, the interpretation of a declarative utterance involves the reconstruction of a question that the utterance is taken to address. For (1a), a plausible question is what Jack sang. The question presupposes that he sang something. Combining the question's presupposition with the literal meaning of (1a), one can infer that he sang something other than *Hey Jude*. (See, e.g., [Rooth 1996: 293-295], [Beaver and Clark 2008: 45-49], [Simons et al. 2010: 318].)

The hypothesis I want to explore is whether this kind of derivation might work for (2), in the obvious way: (2a) elaborates on what Bill thinks about Sue's intellectual abilities; the question presupposes that Bill has an opinion on this matter; since (2a) denies that the opinion is favourable, one can infer that it is unfavourable.

The proposed derivation would belong to the family of pragmatic/semantic accounts of neg-raising (e.g., [Bartsch 1973], [Halpern 1976], [Gajewski 2005, 2007], [Romoli 2013], [Homer 2015], [Zejlstra 2018]) that contrasts with syntactic accounts in which the inference is explained by covert syntactic movement (e.g., [Fillmore 1963], [Collins and Postal 2014]). In spirit, the derivation is closest to the proposals of Bartsch and Gajewski, who also suggest that the inference arises from a presupposition that one of the alternatives to the negated proposition is true. What it would add are some new ideas about the nature and origin of this presupposition.

Let me spell out the assumptions that go into the proposed derivation. It will be useful to have a semantic conception of questions. What (1a) addresses isn't a syntactic object, but an *issue*. I'll represent issues as sets of propositions that would resolve them. The issue addressed by (1a) might be represented as (3), where D is a set of relevant songs.

$$(3) \quad \{ \text{Jack sang } x \mid x \in D \}$$

The elements of (3) don't cover all of logical space. Worlds where Jack didn't sing anything are excluded. This is the sense in which the issue presupposes that Jack sang something.

To derive the inference in (1), we need to assume that the question's presupposition is inherited by (1a). The idea is that a speaker who addresses a question (in an ordinary, direct way) thereby signals that it has a true answer. Given that the speaker of (1a) addresses (3), one may then infer that some

¹ See [Bassi et al. 2021] and [Del Pinal et al. 2024] for reasons to think that exhaustification is presuppositional rather than assertive. [Bassi et al. 2021] argue that the presupposition is $\phi \rightarrow \bigwedge_i \neg \xi_i$, but their data could equally be accounted for by my biconditional version. Note that if the ξ_i are mutually exclusive then $\phi \leftrightarrow \bigwedge_i \neg \xi_i$ is equivalent to the claim that exactly one of the alternatives is true. We would get the "soft presupposition" posited in [Abusch 2005] and [Gajewski 2005, 2007].

² See [Staniszewski 2022], [Jeretic 2022], [Guerrini and Wehbe 2025], for related, but not identical, proposals.

element of (3) is true. Combined with the literal meaning of (1a), it follows that Jack sang something other than *Hey Jude*.

How do we know that the speaker addresses (3)? In rare cases, the question may have been explicitly raised in the discourse. But even if a question has been explicitly raised, speakers may choose to address it indirectly, by (directly) addressing a different question. (See, e.g., [Roberts 1996].) So the directly addressed question is usually not given by the overt conversational context. Nor do I assume that it is readily introspectible by speakers or hearers. It is a theoretical posit in our model of discourse. One constraint that narrows down the addressed question is that a felicitous utterance must be congruent with the question it addresses. The precise nature of this congruence requirement is a matter of debate, but focus marking is agreed to play a key role. The intonational focus on *Hey Jude* in (1a) indicates that the addressed question has the form (3), where the subject (*Jack*) and the verb (*sang*) are held fixed and the object is varied. With focus on *sang*, we would instead take (1a) to elaborate on whether Jack sang or, say, recited *Hey Jude*.

I'll say that a constituent that is held fixed by a question in this manner forms part of the question's "background". In (3), Jack and the singing are in the background; *Hey Jude* is in the foreground. More precisely, the background is given by what all elements of the question have in common. All elements of (3) entail that Jack sang; they differ in what he sang.

Now return to the NR inference in (2).

- (2) a. Bill doesn't think that Sue is smart.
- b. $\neg$ Bill thinks that Sue isn't smart.

As before, we assume that a hearer who encounters (2a) needs to reconstruct a question or issue that the utterance might address. There are many possibilities, including (4), (5), and (6).

- (4) { Bill thinks Sue is smart, $\neg$ (Bill thinks Sue is smart) }
- (5) { Bill thinks Sue is x | $x \in \{ \text{smart, not smart} \}$ }
- (6) { Bill x Sue is smart | $x \in \{ \text{thinks, is sure, suspects, doubts} \}$ }

(4) is the polar question whether or not Bill thinks Sue is smart. (5) asks what Bill thinks about Sue's intellectual abilities: whether he thinks she's smart or not smart; it holds fixed that Bill has an opinion. (6) asks which attitude Bill has towards the proposition that Sue is smart, holding fixed that he has some attitude or other towards this proposition.

To derive the NR inference, we'd have to assume that (2a) is interpreted as addressing a question like (5). The derivation proceeds as explained above: (5) presupposes that Bill thinks either that Sue is smart or that she isn't smart; since (2a) denies the first possibility one can infer that Bill thinks that Sue isn't smart. With (4) or (6), the NR inference could not be derived.

A challenge for the proposed account is therefore to explain why (2a) should be taken to address a question like (5) – more generally, why an utterance of $\neg P(\phi)$, with a neg-raising predicate P , should be taken to address a question like { $P(\phi)$, $P(\neg\phi)$ }.³

³ The reason for the "like" modifier will become clear in a moment.

We can factor this assumption into three components. First, we need to assume that a negated utterance $\neg\phi$ can address a “positive” issue $\{ \phi, \phi', \dots \}$ whose elements vary ϕ . I’ll have little to say about this, except that it is independently plausible, as shown by the felicity of the following dialogues (which don’t involve NR predicates):

- (7) a. What did Jack sing?
b. He didn’t sing Hey Jude.
- (8) a. What did Bill say about Sue?
b. He didn’t say that she is smart.

It would, in fact, be odd to construe (7b) and (8b) as addressing a corresponding negative question: *What didn’t Jack sing?*, or *What didn’t Bill say?*.⁴

The other two assumptions are more substantive. The first is that the neg-raising predicate P is held fixed across the issue alternatives. The second concerns P ’s complements. Even if the question has the form $\{ P(\phi), P(\phi'), \dots \}$, there is no guarantee that we can derive the NR inference. For example, we can’t derive the inference in (2) if (2a) addresses (9).

- (9) $\{ \text{Bill thinks Sue is } x \mid x \in \{ \text{smart, tall} \} \}$

In this case, we could derive that Bill thinks that Sue is tall, which doesn’t tell us what he thinks about whether she’s smart.

Above, I assumed that the issue addressed by a neg-raising $\neg P(\phi)$ is $\{ P(\phi), P(\neg\phi) \}$, with the complements forming a polar issue $\{ \phi, \neg\phi \}$. This does the job. But weaker conditions often suffice. Many neg-raisers appear to be upward monotone. It then suffices that the complements form an issue $\{ \phi, \phi', \dots \}$ in which all elements other than ϕ entail $\neg\phi$. By upward monotonicity, all elements of $\{ P(\phi), P(\phi'), \dots \}$ other than $P(\phi)$ then entail $P(\neg\phi)$; so the question’s presupposition together with the literal meaning of $\neg P(\phi)$ entails $P(\neg\phi)$. For example, we can still derive the inference in (2) if (2a) addresses (10):

- (10) $\{ \text{Bill thinks Sue is } x \mid x \in \{ \text{smart, of average intelligence, stupid} \} \}$

Conjoined with the disjunction of the elements of this issue, (2a) entails (2b).

In sum, the proposed model predicts that NR inferences arise if and only if $\neg P(\phi)$ is taken to address a question $\{ P(\phi), P(\phi'), \dots \}$ in which P is backgrounded and the complements $\{ \phi, \phi', \dots \}$ form an embedded issue with certain properties, depending on structural properties of P .⁵ Alternative accounts of neg-raising generally don’t make these predictions. In the next two sections, I’ll put them to the test.

4 Some simple (and early) theories of question-answer congruence, such as [Rooth 1992], don’t predict the congruence between the questions and answers in (7) and (8). My proposal therefore requires a more flexible theory. But again, this hardly seems a cost.

5 There is no good way of spelling out the second condition in full generality that wouldn’t subsume the first condition. I still find it helpful to separate the two conditions.

2 Backgrounding

There are different ways of theorizing about discourse structure, with different conceptions of how assertions relate to issues. Without committing to a particular theory, there is no simple way to establish what issue a given assertion might address. Still, there is evidence that neg-raising predicates tend to be in the background of the addressed issue.

Let's start with the neg-raiser *think*. It belongs to the class of weak assertives, along with *believe*, *suppose*, *seem*, and *say*. [Hooper and Thompson 1973], followed by [Simons 2007] and others, have argued that these predicates often don't contribute to the main content of an utterance. Relevant data include the fact that (11b) can felicitously answer (11a), that *No* in (12b) can target the complement in (12a), and that weak assertives can be used parenthetically, as in (13).

- (11) a. Is Sue smart?
b. Bill thinks that she is.
- (12) a. Bill thinks that Sue is smart.
b. No, she isn't.
- (13) Sue is smart, Bill thinks.

The idea that weak assertives “don't contribute to the main content” can be developed in two ways. On a strong reading, we might hold that weak assertives (often) don't figure in the addressed issue at all, but merely serve as a hedge or to mark a source of information. If this is applied to the negated use in (2a), we get a new hypothesis about the addressed issue, besides the ones considered in the previous section: perhaps the issue is simply whether Sue is smart.

- (14) { Sue is smart, $\neg$ (Sue is smart) }

[Hooper and Thompson 1973: 482f.] suggest that this might explain NR: since *Bill thinks* contributes nothing to the addressed issue, the negation in *Bill doesn't think that Sue is smart* is effectively a negation of *Sue is smart*.

It is not clear to me how this would work, though. What we want to predict is not an inference to *Sue isn't smart*, but to *Bill thinks that Sue isn't smart*. Moreover, while it is somewhat plausible that the assertives in (11), (12), and (13) serve in a quasi-evidential capacity, this is much less attractive for negated occurrences. Without already presupposing the NR effect, it is unclear how *Bill doesn't think that Sue is smart* could be an answer to *Is Sue smart?* in which the reference to Bill's attitude (or lack thereof!) functions as a hedge or to mark a source of information. As Hooper and Thompson themselves point out, negated weak assertives also can't be used parenthetically:

- (15) # Sue is smart, Bill doesn't think.

I prefer a less radical take on the idea that weak assertives “don't contribute to the main content”. Granted, the main issue addressed by *S thinks that ϕ* often concerns ϕ : it has the form { ϕ , $\neg\phi$ } or { ϕ , ϕ' , ... }. But one may address this issue by raising and answering a secondary issue: what *S* thinks about the primary issue – which element of the primary issue *S* thinks is true. The secondary

issue then takes the form $\{ S \text{ thinks } x \mid x \in \{ \phi, \neg\phi \} \}$ or $\{ S \text{ thinks } x \mid x \in \{ \phi, \phi', \dots \} \}$, with *S thinks* in the background.

This is especially attractive for negated weak assertives, for which it's unclear how they could stand outside the addressed issue. In *Bill doesn't think that Sue is smart*, Bill's opinion is "at issue" insofar as it is relevant to the question being addressed. But it isn't the main issue. The main issue is (arguably) whether Sue is smart; the speaker addresses this issue by elaborating on what Bill thinks about it. *Bill doesn't think* is therefore in the background. Intuitively, the issue is what Bill thinks about Sue's intellectual abilities, not whether it is Bill who does the thinking, or whether we're dealing with a thinking or a doubting.⁶

If NR effects depend on backgrounding the relevant predicate, they should disappear if the predicate is in the foreground. The "MegaNegRaising" dataset from [An and White 2020] provides evidence that this is the case. For example, it shows that the NR effect is absent in (16) and weak in (17).

(16) Bill didn't think to get groceries.

(17) When I first met her, I didn't think Sue was smart.

Plausibly, these are contexts where *think* is in the foreground. In (16), the main issue is whether a certain act of remembering ("thinking") occurred: the issue isn't what Bill remembered, but whether he remembered a particular thing.⁷ More subtly, (17) encourages a biographical reading that also brings the speaker's thoughts into the foreground.

The MegaNegRaising dataset also shows that the NR effect is weaker for the passive *is thought to be* than for the active *thinks*. This, too, suggests that backgrounding is a relevant factor: passivization makes the matrix predicate more prominent and thereby more likely to be in the foreground. (The fact that NR effects are a matter of degree is noteworthy in itself.)

The relevance of backgrounding would also explain the focus-sensitivity of NR effects (discussed, e.g., in [Gajewski 2005] and [Romoli 2013: 223-226]). If a certain kind of intonational focus is put on a neg-raising predicate, the effect tends to disappear:

(18) a. Bill doesn't [think]_F that Sue is smart.

b. ↗ Bill thinks that Sue isn't smart.

This is what we would expect if the NR effect requires backgrounding the predicate. By the congruence requirement on questions and answers, focus on *think* indicates that this constituent is not held fixed in the addressed question: it is not in the background. Intuitively, (18a) addresses an issue like (6), not (5).

A related data point concerns defeasibility. [Bartsch 1973] already pointed out that NR inferences can sometimes be cancelled (see also [Homer 2015]):

⁶ It actually isn't important that *Bill* is in the background. In response to *Does everyone think that Sue is smart?*, *[Bill]_F doesn't think that Sue is smart* still has an NR reading. Here the addressed question is plausibly $\{ x \text{ thinks that Sue is } y \mid x \in D, y \in \{ \text{smart, not smart} \} \}$. In principle, we could even allow the predicate to be foregrounded, as long as it has only one salient and contrary alternative.

⁷ Admittedly, the predicate *think* in (16) may be different from the ordinary predicate *think*.

-
- (19) a. Bill isn't sure whether Brutus and Caesar lived at the same time. So, naturally, Bill doesn't think Brutus killed Caesar.
 b. ↗ Bill thinks Brutus didn't kill Caesar.

The first sentence in (19a) draws attention to Bill's incomplete knowledge. Together with the parenthetical *naturally*, this seems to encourage a non-standard reading of *Bill doesn't think...* on which it addresses what kind of attitude Bill has towards the proposition that Brutus killed Caesar.⁸

I've focussed on *think*, but similar points could be made for other neg-raisers like *want* or *should*. When we interpret *S doesn't want ϕ* , we normally take it to address a question like *What does S want?*, which presupposes that *S* wants either ϕ or some non- ϕ alternative. This interpretation is weakened if intonational focus is put on *want*, or if the conversational context makes clear that the speaker elaborates on their bouletic attitude towards a fixed proposition, as when a candidate at a job interview says *I don't want to make a lot of money*. (The example is from [Homer 2015].)

It can also be instructive to compare NR predicates with semantic relatives. The MegaNegRaising dataset shows that *want* is more likely to trigger an NR inference than *desire*. Let's look at an example.

- (20) a. Bill doesn't want to leave.
 b. Bill doesn't desire to leave.

Arguably, the more unusual and formal choice of *desire* changes the focus. While (20a) elaborates on what Bill wants to do, the question addressed by (20b) is more likely what attitude Bill has towards leaving. Foregrounding the attitude blocks the NR inference.

We can see the same contrast for *believe* and some of its variants:

- (21) a. Bill doesn't believe that Sue is smart.
 b. Bill doesn't hold that Sue is smart.
 c. Bill doesn't maintain that Sue is smart.
 d. Bill doesn't harbour the belief that Sue is smart.

The NR effect is much stronger for (21a) than for (21b)–(21d).⁹ Semantically, there is little difference between the relevant predicates. But *want* and *believe* are inconspicuous and common, in a way that *desire*, *hold*, *maintain*, and *harbour the belief* are not. We might expect that if someone uses an unusual expression, they are more likely to address a question in which that expression (or its meaning) figures in the foreground. (Indeed, it seems to me that the NR reading for (20b) becomes more accessible if the sentence is pronounced in a “posh” accent, in which case *desire* is no longer an unusual choice of word.)

The proposed model might therefore explain why neg-raisers are almost always simple and com-

⁸ I haven't described how, exactly, (19a) manages to foreground *think*. It's noteworthy that NR effects generally can't be cancelled merely by drawing attention to “middle” possibilities in which neither $P(\phi)$ nor $P(\neg\phi)$ is true (see, e.g., [Gajewski 2005] and [Romoli 2013: 300]). The hypothesis I'm exploring does not predict that cancellation should be that easy. This topic would deserve a more careful study.

⁹ The claim for *hold* and *maintain* is supported by the MegaNegRaising dataset. *Harbour the belief* is not included in the test; I here rely on my personal judgment.

mon;¹⁰ complex and unusual predicates are almost never NR, even if they closely resemble NR predicates in meaning.

A similar piece of evidence comes from Horn’s observation that neg-raising predicates are typically “mid-scalar”. More precisely, they tend to designate a location between the midpoint and the maximum of a relevant scale. *Think, likely, should, want, and usually* are NR; the correspondingly stronger *know, be certain, have to, demand* and *always* are not. (See [Horn 1975, 1978].) How come?

Let’s compare *believe* and *be certain*.

- (22) a. Bill doesn’t believe that Sue is smart.
b. Bill isn’t certain that Sue is smart.

It’s natural to construe (22a) as addressing what Bill believes about Sue’s intellectual abilities: whether he believes that she’s smart or that she’s not smart. It is much less natural to read (22b) as addressing whether Bill is certain that Sue is smart or that she’s not smart. This issue would presuppose something extraordinary: that Bill is certain one way or the other. Intuitively, what’s addressed by (22b) is the strength of Bill’s belief: whether he is certain, somewhat confident, doubtful, and so on. The attitude is in the foreground. And so the NR effect is blocked.

Again, I’m not trying to give a full explanation of when NR inferences arise and when they don’t. My claim is merely that the distribution data seem to support the conjecture that NR effects depend on backgrounding the relevant predicate. If the conjecture is true, it is not surprising that NR effects often arise for simple, common, weak, and unfocused expressions, and not for complex, unusual, strong, and focused alternatives, even though the exact contribution of these factors still needs to be worked out.

One might object that a purely pragmatic account can’t explain why, say, *hoffen* has a neg-raising use in German, but its English equivalent *hope* does not, since there’s no significant difference in complexity, strength, or frequency of use.¹¹ Clearly an element of conventionalization plays a role.

But perhaps we can allow for this. Suppose the proposed account of NR is correct. The interpretation of $\neg P(\phi)$ then requires determining whether *P* is backgrounded or foregrounded in the addressed question. Focus marking, strength, complexity, and frequency of use may all influence this decision, but they may not yield a decisive verdict. If you learn German, you might still be unsure whether you should take *hoffen* to be backgrounded or foregrounded when interpreting utterances like (23).

- (23) Ich hoffe nicht, dass Bill bestraft wird.
[I don’t hope that Bill will be punished.]

The utterance itself can now provide a clue. A reading of (23) in which *hoffen* is backgrounded will generally make more sense in the conversational context. Without the NR strengthening, (23) makes

¹⁰ See, for example, the lists in [Horn 1989: 323,327] or in the MegaNegRaising dataset [An and White 2020]. [Zeijlstra 2018: 430] mentions *to be of the opinion* as a neg-raiser. Zeijlstra also mentions, however, that *to be of the opinion* differs from other NR predicates in an important respect: it doesn’t license strong NPIs. So maybe it’s a different kind of beast.

¹¹ According to SUBTLEX-DE and SUBTLEX-UK, *hop(e|es|ing)* is 2-3 times more common in English than *hoff(en|e|t)* in German, but the results for English include uses of *hope* as a noun.

an oddly weak claim. That’s why *don’t hope*, which only has this reading, is virtually unused in English, whereas *hoffe nicht* is fairly common in German.¹² The observed use of $\neg P(\phi)$ is therefore evidence that P can go in the background, which perpetuates this use. We might get a conventional regularity, in the sense of [Lewis 1969], without a change to either syntax or semantics.

3 The Embedded Issue

In the previous section, I’ve investigated whether NR effects depend on backgrounding the relevant predicate. I haven’t found conclusive proof, but the conjecture appears to be supported by the data. I now want to turn to the second condition for the proposed derivation of NR effects. I mentioned at the end of section 1 that the inference from $\neg P(\phi)$ to $P(\neg\phi)$ can’t be derived merely by assuming that $\neg P(\phi)$ is taken to address an issue of the form $\{ P(\phi), P(\phi'), \dots \}$ with P in the background. A further condition is needed. What it is depends on structural properties of P .

Let’s start with a few examples.

- (24) a. Bill doesn’t think that Sue is from [Italy]_F.
 b. $\neg\sim$ Bill thinks that Sue is not from Italy.

The intonational focus on *Italy* in (24a) suggests that the addressed question varies this constituent:¹³ the embedded issue $\{ \phi, \phi', \dots \}$ is unlikely to be the polar issue whether Sue is from Italy. More plausibly, (24a) addresses an issue like (25), where D is a contextually salient set of countries – say, $\{ \text{Italy, France, Spain} \}$.

- (25) $\{ \text{Bill thinks Sue is from } x \mid x \in D \}$

We can still derive the NR inference in (24). Instead of the “excluded-middle presupposition” $P(\phi) \vee P(\neg\phi)$, we now get the presupposition $\text{Think}(\text{from Italy}) \vee \text{Think}(\text{from France}) \vee \text{Think}(\text{from Spain})$. The literal meaning of (24a) rules out the first disjunct, so one can infer $\text{Think}(\text{from France}) \vee \text{Think}(\text{from Spain})$. Given the upward monotonicity of *think*, both disjuncts, and therefore their disjunction, entail $\text{Think}(\text{not from Italy})$.

In general, if P is upward monotone, it suffices that $\neg P(\phi)$ addresses an issue $\{ P(\phi), P(\phi'), \dots \}$ in which the complements $\{ \phi, \phi', \dots \}$ are mutually exclusive. But we don’t need strict, semantic upward monotonicity. A weaker, contextual form of monotonicity is enough. Even if *think* isn’t strictly upward monotone – cue the problem of logical omniscience –, one may reasonably suspect that if Bill thinks that Sue is from France then he thinks that she isn’t from Italy.

Now for a trickier case.

- (26) a. I don’t want cream in my coffee.
 b. $\neg\sim$ I want to not have cream in my coffee.

We might guess that (26a) addresses an issue of the form (27), where D is a relevant set of additives – say, $\{ \text{cream, milk, sugar} \}$.

¹² According to SUBTLEX-DE and SUBTLEX-UK, *hoffe nicht* is about 100 times more common in German than *don’t hope* in English.

¹³ See, e.g., the “Focus Principle” in [Beaver and Clark 2008: 37].

(27) { I want x in my coffee | $x \in D$ }

But then we couldn't explain the inference to (26b). (27) would presuppose that the speaker wants something or other in their coffee. Combined with the literal meaning of (26a), one could infer that they want something other than cream. But that doesn't get us to (26b). After all, the speaker might want sugar and be indifferent about cream.

So (27) isn't the right question for (26a). This is independently plausible: normally, (26a) would not convey that there's something else the speaker wants in their coffee.

One problem with (27) is that it is a "mention-some" question: its elements are not mutually exclusive. Ordinary declarative utterances like (26a) plausibly address a "mention all" question. Intuitively, (26a) answers a question like *What do you want in your coffee?*. This isn't a request to give an example of a desired additive. It's a request to specify one's preferred combination of additives.¹⁴

Importantly, the preferred combination may have no additives at all. The answer might be: *I want nothing in my coffee*. This is not merely the negation of *I want something in my coffee*. If the speaker prefers black coffee but is indifferent about sugar, both are false. *I want nothing in my coffee* rules out indifference towards sugar. So it doesn't just state what the speaker positively wants, as one might have thought.

We can make sense of this if we assume that the complement of *want* specifies what is present in the most preferred configuration (compare [von Stechow 1999]). If the most preferred coffee configuration has nothing in the coffee, it follows that the speaker isn't indifferent about sugar. So *I want x in my coffee* seems to mean something like *My preferred coffee configuration is coffee plus x* ; Accordingly, the question addressed by (26a) might be (28):

(28) { My preferred coffee configuration is coffee plus x | $x \in \{ \text{cream, milk, sugar, nothing} \}$ }

With this, the NR inference to (26b) can be derived. The literal meaning of (26a) rules out that the preferred configuration is coffee + cream (or coffee + cream + sugar, etc., if we allow for multiple additives). The remaining configurations all involve no cream. So the speaker's preferred configuration is one without cream. It follows – with barely any monotonicity assumptions – that the speaker wants to not have cream in their coffee.

Let's look at a predicate that doesn't trigger an NR inference.

- (29) a. Bill didn't say that Sue is from Italy.
b. $\neg$ Bill said that Sue is not from Italy.

Say is a weak assertive, like the neg-raisers *think*, *believe*, and *seem*. We may therefore assume it tends to go into the background: (29a) addresses what Bill said about where Sue is from. Why don't we get an NR effect?

Two factors might be at play. One is that *say* is highly non-monotone. Suppose Bill said that Sue is from France. Did he (thereby) say that Sue is not from Italy? Arguably not. This blocks the derivation of the NR inference, assuming that (29) addresses a question like (30):

(30) { Bill said that Sue is from x | $x \in \{ \text{Italy, France, Spain, ...} \}$ }

¹⁴ Negative answers are generally poor responses to overt mention-some questions.

In general, the proposed model predicts that quotational predicates (*say*, *claim*, *assert*, *sing*, etc.) trigger NR inferences only if the embedded issue is polar, as in (31).

- (31) { Bill said that $x \mid x \in \{ \text{Sue is from Italy, Sue is not from Italy} \} \}$

And this carries a strong presupposition that isn't easily accommodated: of all the things Bill could have said, we are invited to presuppose that there are only two live possibilities. At the very least, a hearer should have doubts about whether (29a) addresses (31), rather than (30) or some other question with more options for what Bill might have said.

I'll return to uncertainty about the addressed question in a moment. First I want to discuss another possible explanation of what blocks the NR inference in (29). Intuitively, (29a) answers *What did Bill say about where Sue is from?*. As in the cream example, we might want to allow for *nothing* as a salient answer. This would block the derivation of the NR inference. If it's possible that Bill said nothing, (29a) can be true while (29b) is false.

But here we have to be careful. Allowing for a null answer would also block the derivation of the NR inference for *think*: if the question addressed by (24a) allows for the possibility that Bill has no opinion about where Sue is from, we can't derive (24b). And isn't *nothing* a felicitous question to *What does Bill think about where Sue is from?*

This issue connects to long-standing debate about focus marking and positive contrast, going back to [Jackendoff 1972: ch.6]. Note that a null answer would also block the derivation of the positive contrast effect in (1): if it's a live possibility that Jack sang nothing, (1a) could be true while (1b) is false. So the question-based account of positive contrast requires that null answers are somehow excluded from the addressed question. This is usually explained in terms of congruence (see, e.g., [Rooth 1996], [Beaver and Clark 2008]): by default, the cells of the addressed question are assumed to vary the denotation of the focused constituents; *Jack sang Let it Be* is a relevant alternative to *Jack sang [Hey Jude]_F*, but *Jack sang nothing* isn't. On this account, *nothing* is a felicitous answer to *What did Jack sing?*, but it isn't perfectly congruent: it rejects a "soft" presupposition carried by the question (compare [Abusch 2010]).¹⁵

In sum, there are independent reasons to think that ordinary, unmarked utterances of (1a) and (24a) address questions that don't include null answers among their cells.¹⁶ These reasons also seem to apply to (29a), suggesting that the hypothesized second factor for blocking the NR inference doesn't play a key role.¹⁷

So far, so good. But there is a more serious problem. The proposed model seems to predict too much. If (26a) addresses (28), we could infer not only that the speaker doesn't want cream, but also that they are not indifferent about sugar: every element of (28) specifies that the preferred coffee configuration contains sugar or that it doesn't. A similar problem arises for the Italy example (24). If (24a) addresses (25), (24a) should not just convey (24b), but also that there's a specific country other

¹⁵ One relevant piece of evidence: *I found out what Jack sang* strongly conveys that Jack sang something. Similarly, *I found out what Bill thinks about ϕ* strongly conveys that Bill has an opinion.

¹⁶ In the coffee/cream example, *nothing* isn't a null answer in the relevant sense. A null answer would be *I don't have a preference*.

¹⁷ The factor still might play some role. For example, if it is a highly salient possibility that Bill said nothing about where Sue is from, (29a) might be taken to address the polar question whether Bill said that Sue is from Italy.

than Italy that Bill thinks Sue is from.¹⁸

Let's see if we can avoid these unwelcome predictions. I'll start with the Italy case. Intuitively, it is clear what went wrong: Bill might think that Sue is from either France or Spain, or from Eastern Europe, or from Micronesia, without being able to pin down a specific country. That is, we should include unspecific alternatives to Italy in the addressed question.

But if we extend the domain D in (25) to unspecific locations, the elements of (25) may no longer be mutually exclusive. Suppose D contains both Italy and Southern Europe. Bill may well believe that Sue is from Italy and also that she is from Southern Europe. The two elements of (25) can be true at the same time. This failure of exclusivity would break the derivation of the NR inference. Suppose Bill thinks that Sue is from Southern Europe, without having a more specific belief about her place of origin. Then (24a) is literally true: it is not the case that Bill thinks that Sue is from Italy. But (24b) is false: Bill doesn't think that Sue is from outside Italy.

The problem is serious, and I'm not sure about the best solution. But I don't think it just affects my proposed account of neg-raising. Consider *Bill thinks Sue is from Italy*. What issue might this address? Intuitively, the issue is *Where does Bill think Sue is from?*. But how should this issue be modelled? It plausibly isn't $\{ \text{Bill thinks Sue is from } x \mid x \text{ is a country} \}$; we don't want to presuppose that Bill can pin down a specific country. But we also don't want to presuppose that he can't. So shouldn't we allow both *Italy* and *Southern Europe*? But then we get an issue with overlapping elements, yet we don't seem to deal with a "mention-some" question.

One might suggest that the issue is about the most specific location that Bill thinks Sue is from. But this doesn't seem true in general. In most contexts, *Bill thinks Sue is from Italy* doesn't convey that Bill has no more specific belief – for example, that Sue is from Tuscany. The main issue (indirectly) addressed by *Bill thinks Sue is from Italy* might be Sue's exact place of origin; but it might just as well be which country Sue is from, or whether she's from Italy or from somewhere else.

This points to another solution. We've assumed that every utterance addresses a unique issue, which the hearer is able to identify. This is unrealistic. An utterance like *Bill thinks Sue is from Italy* is congruent with many issues, even assuming that *Italy* is in the foreground and the rest in the background. The addressed issue might be any of (32), among many others.

- (32) a. $\{ \text{Bill thinks Sue is from } x \mid x \in \{ \text{Italy, France, Spain} \} \}$
b. $\{ \text{Bill thinks Sue is from } x \mid x \text{ is a country} \}$
c. $\{ \text{Bill thinks Sue is from } x \mid x \in \{ \text{Italy, outside Italy} \} \}$
d. $\{ \text{Bill thinks Sue is from } x \mid x \in \{ \text{Italy, France, Spain, Germany, Eastern Europe, Asia, ...} \} \}$

Instead of assuming that the utterance uniquely addresses one of these issues, we could assume that it addresses the whole family of issues – or a subfamily, depending on the conversational context. Each issue in the family has mutually exclusive elements, but the elements of different issues can overlap. Intuitively, a question like *Where is Sue from?* or *Where does Bill think Sue is from?* underspecifies the relevant alternatives. Are we asking whether Sue is from Italy or somewhere else? Are we asking

¹⁸ I can access such a reading if I read (24a) with rise-fall contour on *Italy*. But this inference is clearly less robust than the NR inference to (24b).

which country she is from, or which city, or region, or continent? When a speaker utters *Bill thinks Sue is from Italy*, they might not have settled on a specific issue. Even if they have, the hearer may not be able to determine that issue.

Suppose you hear an utterance, and all you can determine is that it addresses an issue in a certain family (or that it addresses a certain family of issues). The issues in the family might have different presuppositions – as is the case with the issues in (32). (32a) and (32b) presuppose that there’s a specific country that Bill thinks Sue is from; (32c) and (32d) do not. Plausibly, you shouldn’t assume that all presuppositions of all issues in the family are satisfied (which may even be impossible). Only if something is presupposed by every live issue in the family can you safely assume that it is the case.

With all this in place, we can return to the NR inference in (24). Like *Bill thinks Sue is from Italy*, *Bill doesn’t think Sue is from Italy* normally underdetermines the addressed issue. The issue might be any of (32), among others. The hearer thus can’t infer that there’s a specific country (other than Italy) that Bill thinks Sue is from. This inference would be licensed by some of the candidate issues, but not by others. What the hearer can safely infer is that Bill thinks that Sue is from somewhere other than Italy. This follows no matter which issue in the family is addressed.¹⁹

The problem with the cream example (26) might be solved in the same way. Above, I suggested that *I don’t want cream in my coffee* addresses (28).

(28) { My preferred coffee configuration is coffee plus x | $x \in \{ \text{cream, milk, sugar, nothing} \} \}$

This presupposes that the preferred configuration is uniform with respect to sugar. But the utterance doesn’t mention sugar. It’s not obvious that the speaker’s preference with respect to sugar is part of the addressed issue. Perhaps the issue is simply whether the speaker’s preferred configuration includes cream or not, as in (33).

(33) { My preferred coffee configuration has x | $x \in \{ \text{cream, no cream} \} \}$

Or perhaps the issue is about cream and milk, but not about sugar. And so on. If we can’t determine which of these issues is addressed, we can’t infer that the speaker has a preference about sugar. But we can infer that they prefer a configuration without cream.

4 Further Issues

Let’s recap. My goal is to explore whether theories of discourse structure might be able to explain NR inferences as arising through essentially the same mechanism as “positive contrast” inferences like (1): $\neg P(\phi)$ is taken to address a question like $\{ P(\phi), P(\neg\phi) \}$, which, together with the literal meaning of $\neg P(\phi)$, allows us to infer $P(\neg\phi)$. In the previous two sections, I’ve argued that this hypothesis is promising, in particular because it seems to explain the distribution of NR effects. I’ve also identified some challenges that such an account would need to meet. In this final section, I want to briefly review some other data that any theory of neg-raising should explain. In each case, I will

¹⁹ I don’t imagine that processing (24a) involves explicitly going through a list of questions the speaker might address and assessing which inferences are supported by all of them. There are obvious shortcuts. At the algorithmic level, processing $\neg P(\phi)$ might simply involve generating a varied sample of alternatives $P(\xi)$ to $P(\phi)$ and noticing that they all support the inference to $P(\neg\phi)$.

briefly indicate a direction in which the proposed model might account for the data. (I won't mention data that specifically cast doubt on syntactic movement accounts, as it is obvious that these don't challenge the proposed model.)

First, there's the puzzle of limited projection. If the NR inference from $\neg P(\phi)$ to $P(\neg\phi)$ arises from an excluded-middle presupposition $P(\phi) \vee P(\neg\phi)$, we might expect that the presupposition projects out of conditional antecedents and epistemic modals, as presuppositions normally do. But it doesn't (see, e.g., in [Gajewski 2007: 295], [Romoli 2013: 302f.]):

- (34) a. If Bill doesn't think that Sue is smart, he'll let us know.
 b. $\neg$ Bill has an opinion as to whether Sue is smart.

On the account I have outlined, this is not a surprise. The "excluded middle presupposition" is a purely pragmatic presupposition arising from the fact that $\neg P(\phi)$ is taken to address $\{ P(\phi), P(\neg\phi) \}$. The conditional *if* $\neg P(\phi)$ *then* Q plausibly addresses a different issue. It elaborates on what is the case if the antecedent is true – whether *if* $\neg P(\phi)$ *then* Q or *if* $\neg P(\phi)$ *then* $\neg Q$, perhaps. This question doesn't presuppose $P(\phi) \vee P(\neg\phi)$.

Next, the puzzle of partial cyclicity. When neg-raising predicates are stacked under (overt) negation, as in $\neg P_1(P_2(\phi))$, they sometimes have a "cyclic" reading $P_1(P_2(\neg\phi))$ in which the negation is interpreted as located below the lowest predicate. This happens, for example, with *think* and *want* in (35a), but not with the inverse order in (35b), as [Horn 1971] pointed out.

- (35) a. I don't think Sue wants to help.
 b. I don't want Sue to think I'm angry.

Can we make sense of this? On the model I have outlined, we would predict the cyclicity in (35a) if the addressed question is something like $\{ \text{Think}(\text{Want}(\phi)), \text{Think}(\text{Want}(\neg\phi)) \}$. We would predict the non-cyclicity in (35b) if the addressed question here is something like $\{ \text{Want}(\text{Think}(\phi)), \text{Want}(\neg\text{Think}(\phi)) \}$. But whence the difference?

I suspect that the general nature of weak assertives, mentioned in section 2, plays a role. Informally, the main issue addressed by (35a) isn't what the speaker thinks, but what Sue wants. The speaker's beliefs are a secondary issue. That is, the speaker elaborates on $\{ \text{Want}(\phi), \text{Want}(\neg\phi) \}$ by means of elaborating on what they think about this issue. By the same reasoning as in section 2, this means that the addressed question is $\{ \text{Think}(\text{Want}(\phi)), \text{Think}(\text{Want}(\neg\phi)) \}$. In (35b), by contrast, the main issue is what the speaker wants (specifically, what they want Sue to think). The speaker's attitudes don't form a secondary issue.

This line of explanation is supported by the distribution of cyclicity across NR-predicates (see, e.g., [Staniszewski 2017]). The clearest examples of cyclicity involve weak assertives like *think*, *seem* or *suppose* as the higher predicate P_1 . It's also notable that the cyclic reading seems strongest if the higher predicate combines with a first-person subject. Compare (35a) with (36).

- (36) Bill doesn't think Sue wants to help.

A cyclic reading of (36) is still possible, but it seems less robust.²⁰

²⁰ [Gajewski 2005: 53ff] explains the difference between (35a) and (35b) in terms of the different projection of presuppo-

Next, the problem of quantified subjects. Can we explain the NR effect in cases like (37) and (38)?²¹

- (37) a. No juror thinks Bill is innocent.
 b. $\neg$ Every juror thinks that Bill is guilty.
- (38) a. Not every juror thinks that Bill is innocent.
 b. $\neg$ Some juror thinks that Bill is guilty.

In response, one might suggest that the matrix clause in quantified sentences can address a local question. [Beaver and Clark 2008], for example, pursue this idea. But a global treatment also seems possible. Arguably, (37a) and (38a) address what each juror thinks about whether Bill is guilty:

- (39) $\{ \bigwedge_{x \in J} x \text{ thinks Bill is } f(x) \mid f \in \{ \text{guilty}, \text{innocent} \}^J \}$

D^J is the set of functions from J to D . Each element of (39) specifies for each juror whether they think that Bill is guilty or innocent. The question presupposes that each juror has an opinion. Since (37a) rules out all elements of (39) in which some juror thinks that Bill is innocent, one can infer that every juror thinks that Bill is guilty: this is implied by all remaining elements of (39). (38) can be handled in the same way, assuming it also addresses a question like (39).

Finally, there's the puzzle of NPI licensing. [Lakoff 1969] observed that strong NPIs like *until tomorrow* are licensed in the complement of NR predicates, but not in otherwise similar non-NR predicates, as illustrated in (40).

- (40) a. Bill doesn't think that Sue will leave until tomorrow.
 b. # Bill isn't certain that Sue will leave until tomorrow.

This observation was initially taken to support a syntactic account of neg-raising, on the assumption that NPIs are licensed only under local (clausemate) negation. Closer scrutiny, however, has revealed that the licensing conditions for NPIs are more liberal. Arguably, what is required is local negation not at the level of syntax, but at the level of interpreted meaning (see, e.g., [von Stechow 1999], [Chierchia 2013], [Gajewski 2011], [Romoli 2013], [Zeijlstra 2018], [Jacobson 2020]). Concretely, [Gajewski 2011] argues that strong NPIs are licensed in environments that are anti-additively downward-entailing once presuppositions are factored in; given that $\neg P(\phi)$ lexically presupposes $P(\phi) \vee P(\neg\phi)$, the enriched meaning $P(\neg\phi)$ creates an anti-additively downward-entailing environment for ϕ , which would explain the felicity of *until tomorrow* in (40a). On the proposal I've

sitions under *think* and *want*. Suppose, with Gajewski, that *Want*(ϕ) triggers the presupposition $\text{Want}(\phi) \vee \text{Want}(\neg\phi)$. Presuppositions under *Think* project into the doxastic alternatives. (35a) therefore implies that the speaker's doxastic alternatives satisfy either *Want*(*Help*) or *Want*($\neg$ *Help*). Since the former is ruled out by the utterance, one can infer that she thinks Sue wants to not help. Presuppositions under *Want*, however, project not into the bouletic alternatives, but into the doxastic alternatives. Given that *Think* presupposes $\text{Think}(\phi) \vee \text{Think}(\neg\phi)$, it follows that (35b) only implies that Sue has an opinion about whether the speaker is angry in all doxastic alternatives. But does (35b) really convey a belief that Sue has an opinion about whether the speaker is angry? I don't think so. Also, this explanation of the non-cyclicity in (35b) draws on a specific fact about *want*. It's not clear how it carries over, for example, to cases with *should* as the higher predicate.

21 This is a prominent question in the debate on conditional excluded middle; see e.g. [Williams 2010].

outlined, $\neg P(\phi)$ is also enriched to $P(\neg\phi)$ through a presupposition, but the presupposition isn't lexical; it is incurred through the question the utterance is taken to address. Whether the resulting variant of Gajewski's account is tenable remains to be determined, but it is at least noteworthy that the defeasibility and focus-dependency of NR-effects correlates with the felicity of NPIs, as shown in (41) and (42).

- (41) ? Bill doesn't [think]_F that Sue will leave until tomorrow.
- (42) Bill isn't aware that Sue has to leave. ? So, naturally, he doesn't think that she will leave until tomorrow.

A lot more should be said about this topic, and about many of the other topics that I've touched on in this paper. I have not presented a worked-out model of neg-raising. My goal is more generic. I want to make it plausible that neg-raising is not an isolated grammatical quirk, but a natural consequence of broader linguistic mechanisms that also give rise to positive contrast effects. To substantiate this conjecture, I've looked at one popular approach to positive contrast effects, based on theories of discourse structure. I've argued that this approach appears to provide an attractive account of NR effects, without requiring any new mechanisms or lexical assumptions. The appearance may of course be deceiving. Further work is needed to determine whether the details can be worked out.

One last caveat. Throughout, I've assumed that the neg-raising behaviour of *think*, *believe*, *want*, *should*, etc. has a uniform explanation. This may not be the case. A closer look at the data might reveal that different cases call for different explanations. If so, the approach I've presented might be part of the story, but not the whole story.

There are other phenomena, not usually discussed as cases of neg-raising, on which the approach might shed light. There is, for example, the puzzle that a negated conditional $\neg(A > C)$ seems equivalent to the corresponding conditional with a negated consequent, $A > \neg C$, as illustrated by (43), where I exploit the cyclical NR effect of *not think* to make the wide-scope negation more natural.

- (43) a. I don't think the coin would have landed heads if it had been flipped.
b. $\neg$ I believe that the coin would have not landed heads if it had been flipped.

This is puzzling because many seemingly competent speakers also judge that $A > C$ and $A > \neg C$ are both false if there is a nontrivial objective chance of C given A .²² These conflicting data points could be reconciled if we assume that $\neg(A > C)$ is normally taken to address the question what would (or will) be the case if A is the case: something like $\{A > C, A > \neg C\}$. This would license the inference from $\neg(A > C)$ to $A > \neg C$, even if the inference isn't semantically valid.

Another potential application is the well-known homogeneity of definite plurals, illustrated by (44).

- (44) a. I don't think the pigs are in the barn.
b. $\neg$ I believe that the pigs are not in the barn.

Assuming that *the pigs are in the barn* is true iff all (or almost all) the pigs are in the barn, the inference is semantically invalid. Why are we tempted to draw it? A possible explanation is that

22 E.g., [Lewis 1981: 331], [Joyce 1999: 172f.], [Bennett 2003: p.191], [Hájek 2009: p.217].

$\neg(\textit{The } x : Fx)Gx$ is taken to address a question like $\{ (\textit{The } x : Fx)Gx, (\textit{The } x : Fx)\neg Gx \}$, which presupposes that the F are uniformly G or non- G . Intuitively, (44a) addresses *where are the pigs?*, or *where do you think the pigs are?*, with *the* in the background.

References

- Dorit Abusch [2005]: “Triggering from Alternative Sets and Projection of Pragmatic Presuppositions”
- [2010]: “Presupposition Triggering from Alternatives”. *Journal of Semantics*, 27(1): 37–80
- Hannah Y. An and Aaron S. White [2020]: “The lexical and grammatical sources of neg-raising inferences”. *Society for Computation in Linguistics*, 3(1)
- Renate Bartsch [1973]: ““Negative Transportation” Gibt Es Nicht”. *Linguistische Berichte*, 27(7): 1–7
- Itai Bassi, Guillermo Del Pinal and Uli Sauerland [2021]: “Presuppositional Exhaustification”. *Semantics and Pragmatics*, 14: 11:1–48
- David I. Beaver and Brady Z. Clark [2008]: *Sense and Sensitivity: How Focus Determines Meaning*. John Wiley & Sons
- Jonathan Bennett [2003]: *A Philosophical Guide to Conditionals*. New York: Oxford University Press
- Lauri Carlson [1983]: *Dialogue Games: An Approach to Discourse Analysis*. Dordrecht, Netherlands: Reidel
- Gennaro Chierchia [2013]: *Logic in Grammar: Polarity, Free Choice, and Intervention*. Oxford University Press
- Chris Collins and Paul M. Postal [2014]: *Classical NEG Raising: An Essay on the Syntax of Negation*, vol 67. MIT press
- Guillermo Del Pinal, Itai Bassi and Uli Sauerland [2024]: “Free Choice and Presuppositional Exhaustification”. *Semantics and Pragmatics*, 17(3)
- Charles J. Fillmore [1963]: “The Position of Embedding Transformations in a Grammar”. *WORD*, 19(2): 208–231
- Jon Robert Gajewski [2005]: *Neg-Raising: Polarity and Presupposition*. PhD Thesis, Massachusetts Institute of Technology
- [2007]: “Neg-Raising and Polarity”. *Linguistics and Philosophy*, 30(3): 289–328
- [2011]: “Licensing Strong NPIs”. *Natural Language Semantics*, 19(2): 109–148

-
- Janek Guerrini and Jad Wehbe [2025]: “Homogeneity as Presuppositional Exhaustification”. *Journal of Semantics*: ffa007
- Alan Hájek [2009]: “Fifteen Arguments against Hypothetical Frequentism”. *Erkenntnis*, 70(2): 211–235
- Richard N. Halpern [1976]: “The Bivalence of Neg Raising Predicates”. *Studies in the Linguistic Sciences*, 6: 69–81
- Vincent Homer [2015]: “Neg-Raising and Positive Polarity: The View from Modals”. *Semantics and Pragmatics*, 8: 4:1–88
- Joan B. Hooper and Sandra A. Thompson [1973]: “On the Applicability of Root Transformations”. *Linguistic Inquiry*, 4(4): 465–497
- Laurence R. Horn [1971]: “Negative Transportation: Unsafe at Any Speed?” In *Proceedings of Chicago Linguistics Society*, vol 7. Chicago: CLS, 120–133
- [1975]: “Neg-Raising Predicates: Toward an Explanation”. *CLS*, 11: 94
- [1978]: “Remarks on Neg-Raising”. In P. Cole (Ed.) *Pragmatics*, chapter Pragmatics. Brill, 129–220
- [1989]: *A Natural History of Negation*. Chicago: University Of Chicago Press
- Ray Jackendoff [1972]: *Semantic Interpretation in Generative Grammar*. Cambridge, MA: MIT Press
- Pauline Jacobson [2020]: “Neg Raising and Ellipsis (and Related Issues) Revisited”. *Natural Language Semantics*, 28(2): 111–140
- Paloma Jeretic [2022]: “Neg-Raising with Belief Predicates as a Scaleless Implicature”. *Semantics and Linguistic Theory*: 891–914
- James Joyce [1999]: *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press
- Angelika Kratzer [1989]: “An Investigation of the Lumps of Thought”. *Linguistics and Philosophy*, 12(5): 607–653
- Robin Lakoff [1969]: “A Syntactic Argument for Negative Transportation”. In *Proceedings from the Annual Meeting of the Chicago Linguistic Society*, vol 5. Chicago Linguistic Society, 140–147
- David Lewis [1969]: *Convention: A Philosophical Study*. Cambridge (Mass.): Harvard University Press
- [1981]: “Causal Decision Theory”. *Australasian Journal of Philosophy*, 59: 5–30

-
- Craige Roberts [1996]: “Information Structure in Discourse: Toward a Unified Theory of Formal Pragmatics”. *Ohio State University Working Papers in Linguistics*, 49: 91–136
- Jacopo Romoli [2013]: “A Scalar Implicature-Based Approach to Neg-Raising”. *Linguistics and Philosophy*, 36(4): 291–353
- Mats Rooth [1992]: “A Theory of Focus Interpretation”. *Natural language semantics*, 1(1): 75–116
- [1996]: “Focus”. In Shalom Lappin (Ed.) *The Handbook of Contemporary Semantic Theory*, Oxford: Blackwell
- Mandy Simons [2007]: “Observations on Embedding Verbs, Evidentiality, and Presupposition”. *Lingua*, 117(6): 1034–1056
- Mandy Simons, Judith Tonhauser, David Beaver and Craige Roberts [2010]: “What Projects and Why”. *Semantics and Linguistic Theory*: 309–327
- Frank Staniszewski [2017]: “Partial Cyclicity and Restrictions on Neg-raising”. In *Proceedings of the 34th West Coast Conference on Formal Linguistics*, 490–497
- [2022]: *Modality and Time in Logical Context*. PhD Thesis, Massachusetts Institute of Technology
- K. von Stechow [1999]: “NPI Licensing, Strawson Entailment, and Context Dependency”. *Journal of Semantics*, 16(2): 97–148
- J. Robert G. Williams [2010]: “Defending Conditional Excluded Middle”. *Noûs*, 44(4): 650–668
- Hedde Zeijlstra [2018]: “Does Neg-Raising Involve Neg-Raising?” *Topoi*, 37(3): 417–433