A unified approach to Neg-Raising

Zahra Mirrazi
Georg-August-University Göttingen, Seminar for English Philology,
Göttingen,
37073, Germany
zahrasadat.mirrazirenani@uni-goettingen.de

Hedde Zeijlstra
Georg-August-University Göttingen, Seminar for English Philology,
Göttingen,
37073, Germany
hzeijls@uni-goettingen.de

Abstract

This paper addresses a longstanding paradox about Neg-raising. While Neg-raising predicates appear to form a natural semantic class – suggesting a semantic or pragmatic basis for the phenomenon – they also give rise to grammatical effects, such as the licensing of strong/strict NPIs and Horn clauses. Although the semantic generality of Neg-raising predicates has been widely acknowledged, existing accounts– including semantic-pragmatic ones – treat the Neg-raising property as lexically specified, without explaining why certain predicates systematically exhibit it across languages while others never do.

The goal of this paper is to derive the semantic factors that distinguish Neg-raising predicates from non-Neg-raising predicates, while also accounting for the grammatical properties associated with Neg-raising. We propose that Neg-raising readings are available only with predicates that meet two conditions: (i) they are universal modals whose negated form is strictly equivalent to an existential counterpart scoping over negation, and (ii) this existential modal can be strengthened into a universal interpretation via an exhaustification operator, following Bar-Lev and Fox (2017).

We show that these conditions jointly explain both the grammatical effects

and the semantic generalizations associated with Neg-raising predicates. In doing so, our account uniquely explains both the semantic-pragmatic character of Neg-raising and its grammatical effects, unlike prior approaches which address only one or neither.

Keywords: Neg-raising; modals; Strict duality; exhaustification; NPIs; Horn clauses

1 Neg-raising: desiderata

Neg-Raising (NR) concerns the phenomenon, illustrated in (1)-(2), whereby certain negated predicates (e.g. *think*, *believe*, *expect*) can give rise to readings where the negation in the matrix clause seems to be interpreted in the embedded clause. For instance, (1a) has a reading in (1b), alongside the literal interpretation where indeed the speaker does not entertain the thought that the addressee is right. The same, *mutatis mutandis*, holds for (2a) and (2b).

(1) a. I don't think you're right.
b. I think you're not right.

(2) a. I don't believe you're right.
b. I believe you're not right.

Predicates that can give rise to such readings are standardly referred to as Neg-Raising predicates (NRPs), and the readings invoked by NR predicates, where a matrix negation appears to take narrow scope within the embedded clause, are called NR readings.

By contrast, other predicates, dubbed non-NR predicates, do not give rise to such inferences. The negation of predicates like *predict* or *claim* lack readings where negation is interpreted in the embedded clause, as shown in (3)-(4) below. Neither (3a) nor (4a) leads to the inferences given (3b) and (4b):

(3) a. She doesn't predict John is ill.
b. She predict John isn't ill.

(4) a. She doesn't claim John is ill.
b. She claims John isn't ill.

The underlying source of the NR phenomenon has been long debated. Three families of approaches accounting for NR effects have been proposed in the literature: a syntactic approach ((Fillmore, 1963; Horn, 1978; Collins and Postal, 2014)), where negation moves from the embedded clause into the matrix clause, and two semantic-pragmatic approaches, one in terms of excluded middle inferences (Bartsch, 1973; Gajewski, 2005,

2007; Homer, 2015; Romoli, 2012, 2013; Zeijlstra, 2018), and one in terms of exhaustification of domain alternatives (Staniszweski, 2021; Jeretič, 2021; Mirrazi and Zeijlstra, 2021, to appear)). Despite differences in the proposed features underlying NR, the current approaches share the assumption that the NR property is lexically encoded in certain predicates – though the set of NRPs does not necessarily correspond to a semantically natural class. We will briefly introduce them in turn.

Even though nowadays the standard approach to NR is often thought to be the semantic-pragmatic in nature, Collins and Postal (2014) have revived the position that NR involves syntactic movement of the negation from a lower clause into a higher clause, a proposal tracing back to Fillmore (1963), and also adopted in Horn (1971, 1972). Ignoring *do*-support effects, the syntactic structure of (1a) would then be as in (5), and the reading (1b) would follow from interpreting the negation in its base position–where <NEG> indicates a lower copy/trace of negation.

(5) I NEG think that you're <NEG> right.

For Collins and Postal (2014), NR is derived as a result of their assumption that only base occurrences of negation are interpreted. That is, the lowest copy of negation is semantically interpreted, whereas the highest copy of NEG is phonologically realized (in this case as *n't*). Under this approach, the NR property is characterized as a lexical feature of certain predicates that licenses the movement of negation from their embedded clause.

The first semantic pragmatic approach (Bartsch (1973); Horn (1989); Gajewski (2005); Romoli (2013); Homer (2015), and Zeijlstra (2018), among others) takes NR readings to be the result of an excluded middle inference that is a special lexical property of NRPs. This approach has two versions:

(6) a. *The presuppositional approach* (Gajewski, 2005, 2007): NR predicates come with an excluded middle *presupposition.*

b. *The implicature approach* (Romoli, 2012, 2013): NR predicates have excluded middle *alternatives*

The presuppositional approach to NR (Gajewski, 2005, 2007) takes NR predicates to come with an excluded middle presupposition. That is, the speaker is presupposed to be opinionated about the truth or falsity of the embedded proposition (P). The NR reading is then a logical consequence of this presupposition and the literal meaning of the sentence, as shown in (7).

(7) not [NRP [P]]
Assertion: $\neg$ NRP (P)
Presupposition: NRP (P) $\vee$ NRP $\neg$(P) (**?**; p.14).
$\therefore$ NRP $\neg$(P)

In light of problems with the projection behavior of the excluded middle presupposition, however, Romoli (2012, 2013) suggests that NRPs take the excluded middle as a lexical alternative. A NRP like *think*, then, has ($\mathbf{think}_x$p $\vee$ $\mathbf{think}_x\neg$p) as its lexical alternative, as shown in (8).[1]

(8) *Alt*(**think** p(x))={ $\mathbf{think}_x$ p, $\mathbf{think}_x$p $\vee$ $\mathbf{think}_x\neg$p }={$\Box_x$p, $\Box_x$p$\vee\Box_x\neg$p }

The set of alternatives of the negated NRP in (9) is given in (10a). Exhaustification of these alternatives (10b) will result in the strengthened NR reading (10c).

(9) a. John doesn't think that it is raining.
b. $\neg\mathbf{think}_j$ p

(10) a. *Alt*($\neg\mathbf{think}_j$ p)={ $\neg\mathbf{think}_j$ p, $\neg$($\mathbf{think}_j$ p $\vee$ $\mathbf{think}_j\neg$p) }
b. $[\![\text{EXH}]\!](\neg\mathbf{think}_j$ p)= $\neg\mathbf{think}_j$ p $\wedge$ $\neg\neg$($\mathbf{think}_j$ p $\vee$ $\mathbf{think}_j\neg$p) = $\neg\mathbf{think}_j$ p $\wedge$ ($\mathbf{think}_j$ p $\vee$ $\mathbf{think}_j\neg$p)
c. $\mathbf{think}_j\neg$p

Inspired by recent implicature approaches to the Free Choice inference (Bar-Lev and Fox, 2017), and Homogeneity (Magri, 2014; Bar-Lev and Fox, 2020), there have been new attempts to derive the NR reading using the machinery of exhaustification, without assuming a lexical alternative corresponding to the excluded middle (Mirrazi and Zeijlstra, 2021; Jeretič, 2022). While accounts within the second semantic-pragmatic approach differ in their implementations, they share the core idea that the group of NRPs lack a dual with a corresponding possibility meaning. The absence of a scalemate is then taken to facilitate a strengthening process via the application of an exhaustivity operator EXH on the subdomain alternatives of the sentence (S) containing a NRP. The result of this strengthening process, given in (11), is then taken to be equivalent to a strengthened NR reading.

(11) EXH [NEG [x *think* p]]
a. $S = \neg B\{w_1, w_2\}$ (notation for $\neg\forall w \in \{w_1, w_2\}.p(w)$)
b. $Alt(S) = \{\neg B\{w_1, w_2\}p, \neg B\{w_1\}p, \neg B\{w_2\}p\}$
c. $\text{EXH}[Alt(S)][S] \equiv \neg B\{w_1, w_2\}p \wedge \neg B\{w_1\}p \wedge \neg B\{w_2\}p \equiv B\{w_1, w_2\}\neg p$
(Jeretič, 2022)

Adjudicating among these approaches has proven to be difficult. A major challenge in arriving at a satisfactory theory of NR is to account for two opposite characteristics of NRPs. On the one hand, particular grammatical effects associated with NR, particularly

[1]Note that such alternatives violate the structural constraint that alternatives must not be more complex than the assertion Katzir (2007). However, Romoli (2012, 2013) assumes that certain items are associated with a set of lexical alternatives, which then grow to become alternatives of more complex expressions containing them.

(strict) NPI licensing and Horn-clauses, suggest that NR is a syntactic phenomenon. On the other hand, there are semantic generalities governing the membership in the group of NRPs that hint at the semantic nature of NR. First, NRPs across languages seem to form an (almost) natural semantic class (Horn, 1989). Second, there are certain predicates that are never NRPs in any language. However, pursuing a non-syntactic account of NRP is challenged by the existence of semantically unmotivated lexical exceptions to these generalizations (next to the grammatical reflexes of NR). Consequently, semantic-pragmatic approaches to NR as of yet have had to bite the bullet and follow the syntactic tradition of stipulating a lexical feature on certain predicates that is responsible for their NR property. Moreover, in light of difficulties in reconciling semantic and syntactic aspects of NR, many have abandoned the pursuit of a unified approach and posited that distinct properties of NR may have different sources (Zeijlstra, 2018; Crowley, 2019; Horn, 2020).

In the remainder of this section, we provide an overview of these challenges, the resolution of which constitutes the desiderata for a successful unified theory of NR.

1.1 Semantic generality and variation in NRPs

An observation made by Horn (1978), which has remained valid, is that NRPs generally fall within the following semantic classes:

(12) Classes of Neg-Raisers (Horn, 1978)

a. [OPINION] *think, believe, expect, suppose, imagine, reckon*

b. [PERCEPTION] *seem, appear, look like, sound like, feel like*

c. [PROBABILITY] *be probable, be likely, figure to*

d. [INTENTION/VOLITION] *want, intend, choose, plan*

e. [JUDGMENT/OBLIGATION] *be supposed, ought, should, be desirable, advise*

Moreover, certain predicates are never found among the NRPs of any language. For instance, factives are universally excluded from the class of neg-raisers (Kiparsky and Kiparsky, 1968). Similarly, existential modals such as *able, possible*, and *allow* are consistently absent from this class (Horn, 1989). While the current approaches to NR each have a way of excluding factive neg-raisers, Horn (1989) remains the only attempt to capture these semantic tendencies among neg-raisers. He proposes that whether or not a given predicate is a NRP is determined by its place in a positive scale whereby the predicates below the midpoint are distinguished by their tolerance property, as defined below.

(13) a. An operator P is tolerant if the conjunction $[P(p) \wedge P(\neg p)]$ is logically consistent

b. An operator P is intolerant if the conjunction [P(p) $\wedge$ P($\neg$p)] is logically inconsistent.

He argues that tolerant predicates, like *possible*, are never neg-raisers. The predicates on the high end of the scale, like strong necessity modals, which imply the truth of the embedded proposition (veridical propositions), are also excluded from the class of NRPs. This leaves mid-scalar predicates, i.e. intolerant predicates that do not imply the truth of their embedded proposition, the only candidates to be NRPs. Given these robust tendencies, Horn (1989) concludes that NR, rather than being an idiosyncratic property of certain predicates, is an instance of "a fundamental grammatical, semantic and pragmatic phenomenon manifested across distinct, but systematically related, classes of predicates in generally and typologically diverse languages" (Horn, 1989, 309).

Despite the robustness of mid-scalar generalization, a considerable deal of cross-linguistic variation among NRPs is observed. There are some semantically similar lexical items, both within and across languages, that are different in their NR property. For instance, while *want* allows for NR readings, *wish* is not a NRP. In the same vein, it has been observed that cognates of NRPs in one language are not necessarily NRPs even in closely related languages: Dutch *hopen* and German *hoffen* are classified as NRPs, whereas their English translation *hope* does not allow for NR readings.

In sum, the semantic generality of NRPs strongly suggests that NR is a semantic/pragmatic phenomenon, but the idiosyncrasies of NRPs has pushed even the proponents of semantic/pragmatic approaches to NR to take NR to be a feature randomly encoded on some lexical items along the lines of the syntactic/lexical approach to NR. While the hope of the semantic/pragmatic approaches has been that a fine-grained semantics of NRPs will explain the apparent lexically-idiosyncratic exceptions (Jacobson, 2018; Horn, 2020), no attempt has been made to resolve this paradoxical situation.

1.2 Grammatical footprints of Neg-raising

Although the semantic generality of NRPs favors a semantic-pragmatic view of NR, such approaches face challenges from two grammatical phenomena consistently observed in NR contexts. These include the licensing of strong and strict NPIs under extraclausal negated NRPs, as well as the existence of so-called *Horn clauses*.

Under the syntactic approach, on the other hand, there is no explanation for the semantic commonalities among NRPs, as NR results from a lexically-encoded feature on certain predicates that allows for the movement of negation from their embedded clause. The only generalization that can be partially explained under the syntactic approach is why factive predicates are not Neg-raisers, by assuming that factive clauses are islands for movement. Nonetheless, the syntactic approach has been presented as the only approach that can deal with these grammatical reflexes of NR. Both strict NPIs and Negative

Inversion, as said, appear to require a local negation, and under the syntactic approach that requirement is trivially fulfilled. Since existing semantic pragmatic approaches to these grammatical reflexes of NR, have not been successful in accounting for these reflexes, the existence of such reflexes has been taken to be one the strongest arguments in favor of the syntactic approach to NR (Collins and Postal (2014)).

In the remainder of this section, we review the challenges that the licensing of strict NPIs and Horn clauses under NRPs pose for semantic-pragmatic approaches to Neg-raising.

1.2.1 Strong/strict NPIs

There is a class of NPIs, such as *in ages*, *breath a word*, *jackshit* or punctual *until* that deviate from NPIs like *any* or *ever* in that they are both strong and strict NPIs, while NPIs like *any* and *ever* are weak and non-strict.

Strong NPIs differ from weak NPIs in that these can only be licensed in anti-additive[2] contexts and not just in any downward entailing context. This is depicted by the contrast in the licensing environments of the strong NPI punctual *until* and the weak NPI *any* in (14) and (15).

(14) a. She didn't move in *until June*.
b. Nobody moved in *until June*.
c. *Few students moved in *until June*.
d. *At most three students moved in *until June*.

(15) a. She didn't eat *any cookies*.
b. Nobody ate *any cookies*.
c. Few students ate *any cookies*.
d. At most three students ate *any cookies*.

Collins and Postal (2014) show that these NPIs are not only strong NPIs but also strict NPIs. Strict NPIs differ from their non-strict counterparts in that they cannot be licensed when separated from their licenser by a clause or island boundary. Being in the scope of an anti-additive operator is thus not enough to license such NPIs, as shown in the examples below.

(16) a. *Stanley doesn't claim that Carolyn will *breathe a word* about it.
b. *Calvin didn't say that Mona moved in *until June*.

(17) a. *I didn't find a person who moved in *until June*.

[2]An anti-additive contexts is the scope of an anti-additive function, defined below.

(1) A function f is Anti-Additive iff for any a, b in the domain of f , $f(a) \wedge f(b) \subseteq f(a \vee b)$

b. *Nobody found a person who who moved in *until June*.

c. *I didn't find a person who had seen Mary *in years*.

d. *Nobody found a person who had seen Mary *in years*.

(Based on Collins and Postal (2014))

Non-strict NPIs, like *any* and *ever* are not subject to such restrictions:

(18) a. I didn't claim that she knows *any physics*.

b. I haven't said that she has *ever* been there.

(19) a. I didn't find a person who knows *any* physics.

b. Nobody found a person who knows *any* physics.

c. I didn't find a person who had *ever* been to France.

d. Nobody found a person who had *ever* been to France.

(Collins and Postal, 2014)

Strikingly, a negated NRP can license embedded strong/strict NPIs, as shown below.

(20) a. Stanley doesn't believe that Carolyn will *breathe a word* about it.

b. Calvin didn't think that Mona moved in *until June*.

This is surprising, since although the embedded clauses in (20) are anti-additive environments, at least under their strengthened readings, we have seen that anti-additivity is not enough to license strict NPIs when they are separated from the negation by clause-boundaries. Thus, these sentences are still expected to be ungrammatical.

Moreover, Collins and Postal (2014) show that licensing of strict NPIs under negated NRPs is subject to island constraints.

(21) a. *I don't believe the rumor that Tom has found the solution *in years*.

b. *I don't think Tom has found the solution *in years* and is a reliable chap.

(after Collins and Postal (2014))

Such locality constraints on the licensing of strict NPIs pose a challenge for semantic-pragmatic accounts of NR. By contrast, they follow naturally under a syntactic approach in which, at some point in the derivation, negation is generated within the embedded clause. This would create a local anti-additive context necessary for licensing strong and strict NPIs.

1.2.2 Horn-clauses

Horn-clauses are instances in which subject-auxiliary inversion is triggered not by a fronted negative phrase, as in standard negative inversion, but rather by a fronted NPI, which itself is licensed by a negated NRP. Examples illustrating this pattern are provided in (22)–(23).

(22) I don't think that *ever before* have the media played such a major role in a kidnapping.

(23) I don't suppose that *under any circumstances* would she help me.

Subject–auxiliary inversion under negation (also known as Negative Inversion), illustrated in (77), typically applies in a strictly local manner: the specifier of the raised auxiliary must contain a negative or downward-entailing element.

(24) a. Nowhere does he mention my book.
b. *Somewhere does he mention my book. (Büring, 2004)

Collins and Postal (2014) show that, much like the licensing of strict NPIs by a negated NRP, Horn clauses are subject to island constraints, as shown in (25).

(25) a. I don't think that *under any circumstances* would he agree.
b. * What I don't think is that *under any circumstances* would he agree.
(Collins and Postal, 2014)

(26) a. I don't believe that *under any circumstances* would he agree.
b. *I don't believe the rumor that *under any circumstances* would he agree.

Therefore, Collins and Postal (2014) take the existence of Horn clauses as strong evidence for a syntactic approach to NR. Only under such an approach could the negation in the main clause have occupied the specifier of CP at an earlier stage of the derivation, as shown in (27), where < ... > indicates a lower copy.

(27) I do NEG think [<NEG > anywhere] did he mention my book [<NEG anywhere>]

Collins and Postal (2014)'s proposal straightforwardly accounts for the contrast between (22)-(23) and (28)-(29): as *say* and *regret* are not NRPs, syntactic movement of negation is impossible in (28)-(29). Consequently, these predicates are correctly predicted to be unable to embed Horn clauses.

(28) *I don't say that *ever before* have the media played such a major role in a kidnapping.

(29) *I don't regret that *under any circumstances* would she help me.

As of date no alternative semantic/pragmatic accounts for Horn-clauses have been formulated. Also, existing semantic/pragmatic accounts have faced difficulties to explain the licensing of strong/strict NPIs under negated NRPs.

Reconciling the semantic and syntactic properties of Neg-raising has this proven challenging, leading many to abandon a unified approach to NR and propose that different

instances of NR may originate from different sources (Zeijlstra 2018; Jacobson 2018, 2020; Crowley 2019; Collins and Postal 2018a, among others).

While we agree that a successful theory of NR must account for licensing of embedded strong/strict NPIs and Horn-clauses, as well as the relative semantic uniformity among NRPs, we are hesitant to conclude that the attested observations necessitate a hybrid approach to NR. In this paper, we argue that a unified approach to NR that can account for both semantic and grammatical footprints of NR can be formulated.

2 Proposal

Our proposal shares the core insight of other exhaustification (EXH)-based approaches that the NR-inference arises because a weak asserted proposition is strengthened via the application of the EXH operator on its alternatives. In order to meet the desideratum that a successful theory of NR should able to derive the semantic factors that distinguish NRPs from non-NR ones while accounting for its grammatical properties, we propose a new implementation of the EXH-based approach, which consists of the following two components:

(30) a. *LF-equivalence under strict duality*: $\neg\forall \Leftrightarrow_{strict} \exists\neg$ (provided that $\forall$ and $\exists$ have the same presuppositions)

b. *Strengthening of subdomain alternatives* (Chierchia, 2013)

We agree with other EXH-based approaches that strengthening subdomain alternatives can successfully account for (i) the impact of contextual relevance in triggering the NR inference, which in turn explains why this inference is not always invoked by negated NRPs; and (ii) the parallel between homogeneity and NR. However, we argue that for the exhaustification mechanism to be an explanatory account of how NR inferences are generated, it has to be enriched with the assumption of LF-equivalence under strict duality. We argue that this novel component of our analysis is essential in explaining (i) why certain predicates cannot act as NRPs while others can; (ii) what parameters there are for cross-linguistic variation with respect to NR, and (iii) how the observed grammatical effects of NR can be accounted for. Moreover, incorporating this assumption allows our analysis to avoid some technical issues that challenge other EXH-based approaches.

2.1 LF-equivalence under strict duality

The first step in the strengthening mechanism is to identify the proposition and its set of alternatives that serve as input to EXH. Two hypotheses have been entertained about how EXH interacts with the asserted proposition. Under the first approach, EXH is considered as shorthand for general pragmatic reasoning involved in interpreting each utterance to

convey the most informative true proposition that is relevant to the question under the discussion (van Rooy, 2002). The second and perhaps the dominant approach takes EXH to be a covert lexical item that syntactically combines with a sentence ϕ (Fox (2007); Chierchia et al. (2012), among others). Whether pragmatic or grammatical in nature, EXH is a procedure that determines alternatives to ϕ. For EXH to be a cognitively viable procedure and descriptively adequate (e.g. to avoid *the symmetry problem*)[3], the set of alternatives to ϕ needs to be constrained. It is widely accepted that that two kinds of restrictions govern the calculation of alternatives: (i) alternatives to ϕ must be relevant in the context; and (ii) alternatives to ϕ cannot be more complex than ϕ (Katzir, 2007).

Putting aside relevance for now, we must first define the intended notion of complexity. Following Buccola et al. (2021) and their neo-Katzirian framework of alternatives, we assume that alternatives are derived via a finite series of deletions, contractions, and replacements of constituents in the sentence ϕ with constituents of the same category taken from the *language of thought*, whose structure resembles that of classical Logical Form (LF), but contains a lexicon composed of abstract elements, which may or may not be lexicalized in the language. Complexity is then defined in terms of the 'primitiveness' of these lexical items and the complexity of the overall structure.

What we propose to add to this system is the conjecture that EXH is blind to the difference between ϕ and ϕ' where ϕ' is strictly equivalent to ϕ and is no more complex than ϕ. Consequently, any operations that apply to ϕ can also apply to ϕ'. Based on this, we assume that ϕ' can serve as the prejacent of EXH.

Strict equivalence is defined as follows:

(31) a. p is strict equivalent to q (p $\Leftrightarrow_{strict}$ q) iff p strictly entails q (p $\Rightarrow_{strict}$ q) and q strictly entails p (q $\Rightarrow_{strict}$ p).

b. p strictly entails q (p $\Rightarrow_{strict}$ q) iff in every world where p is true, q is true as well.

Consider a sentence containing a negated universal modal, ϕ : $\neg\forall w$: p(w). Given our conjecture, a dual sentence, ϕ' : $\exists w$: $\neg$p(w), can serve as a prejacent of an EXH applying to ϕ, if and only if ϕ' is strictly equivalent to ϕ and no more complex than ϕ. Under standard assumptions regarding the syntax and semantics of modal statements, it is evident that the two sentences do not differ in complexity, either in terms of overall structure or the primitiveness of their lexical items. Nonetheless, strict LF-equivalence does not hold for all modals.

To understand why, it is important to first note that in a trivalent logic system[4], which

[3]The symmetry problem refers to the challenge of providing a well-motivated theory for why a sentence S with a scalar implicature that $\neg A$ for some alternative A should not have $S \wedge \neg A$ as an alternative. If $S \wedge \neg A$ were an alternative, it would yield a scalar implicature that ($\neg(S \wedge \neg A)$, which would contradict what the S says and the scalar implicature $\neg A$ (Fox (2007); Katzir (2007); Breheny et al. (2018), among others).

[4]The possible truth-values are {1,0,# } and the presupposition failure is marked by #.

we adopt here, strict LF-equivalence requires the presupposition preservation. That is, two expressions are strictly LF-equivalent only if one is *defined* and true in a world, the other is *defined* and true. On this view, modals that carry or interact with presuppositions can disrupt strict duality, since the negated universal and its existential dual may differ in definedness across worlds. Factive *know* is a good example. Assume that $\Diamond_k$K is the existential dual knowledge operator of $\Box_k$. When the existential knowledge operator also carries the factivity presupposition that the embedded p is true, strict LF-equivalence does not hold. While the negated universal version of *know* presupposes that its prejacent is true, the existential dual that outscopes negation presupposes that the same prejacent is false:

(32) $\underline{\text{p(w)} = 1}.\neg\Box_k\ \underline{\text{p(w)}} \not\Leftrightarrow_{\textit{strict}} \underline{\neg\text{p(w)} = 1}.\Diamond_k\ \underline{\neg\text{p(w)}}$

Note that even when $\Diamond$K$_p$ doesn't carry any presupposition, strict LF-equivalence is still not valid. In a world where the factivity presupposition is not satisfied, $\neg\Box_k$ p(w) is #, but $\Diamond_k$ $\neg$p(w) is true:

(33) $\underline{\text{p(w)} = 1}.\neg\Box_k\ \underline{\text{p(w)}} \not\Leftrightarrow_{\text{STRICT}} \Diamond_k\ \neg\text{p(w)}$

By contrast, sentences with negated doxastic modals that lack a factivity presupposition, such as *think*, are strictly equivalent to a corresponding sentence containing their existential dual scoping over negation. By (31), we have $\neg\forall w\in W$:p(w) $\Leftrightarrow_{\textit{strict}}$ $\exists w\in W$:$\neg$p(w). Under our proposal, the strictly equivalent statement with the existential dual can also serve as the prejacent for EXH when applied to the original sentence. The question that naturally arises is whether applying EXH to strictly equivalent sentences always results in the same meaning or if it may yield a different meaning. In the next section, we argue not only that the latter is possible, but also that some readings can only be generated when a strictly equivalent statement is available to EXH. Specifically, we propose that NR readings are only generated when EXH applies to a sentence with an existential modal scoping over negation, provided this configuration is made available under strict equivalence.

2.2 Strengthening of subdomain alternatives

Before demonstrating the further importance of LF-equivalence under strict duality for a satisfactory theory of NR in terms of strengthening via exhaustification, we must now elaborate on the second component of our analysis, *strengthening of subdomain alternatives*.

Following Chierchia (2013, 2022); Bassi and Bar-Lev (2018), we take existential quantificational operators to trigger subdomain alternatives. Alternatives of the assertion are derived by substituting the domain variable with its non-empty subsets.

(34) $Alt(\exists_D) : \{\exists_{D'} : D' \subseteq D\}$

For concreteness, let us consider alternatives of the following sentence with possibility modals:

(35) $\exists w \in W$: $\neg$p(w)

Assuming that the domain of this modal consists of three worlds w_1, w_2 and w_3, the alternatives generated from replacing the domain variable with its non-empty subsets are given in (36).

(36) Alt ((35)) = $\{\exists w \in \{w_1, w_2, w_3\}$: $\neg$p(w), $\exists w \in \{w_1, w_2\}$: $\neg$p(w), $\exists w \in \{w_1, w_3\}$: $\neg$p(w), $\exists w \in \{w_2, w_3\}$: $\neg$p(w), $\exists w \in \{w_1\}$: $\neg$p(w), $\exists w \in \{w_2\}$: $\neg$p(w), $\exists w \in \{w_3\}$: $\neg$p(w)$\}$

In the process of strengthening, the main sentence and its alternatives will be input to an exhaustivity operator. We adopt the definition of the EXH operator by Bar-Lev and Fox (2017), according to which EXH takes a proposition (p), and a set of alternatives (C) as arguments, and returns the conjunction of all of the negated innocently excludable (IE) alternatives, and all of the asserted innocently includable (II) alternatives.

(37) Innocent Exclusion + Innocent Inclusion–based exhaustivity operator:
$[\![\text{EXH}]\!]^{IE+II}(C)(p)(w) \Leftrightarrow \forall q \in \text{IE}(p, C)[\neg q(w)] \wedge \forall \in \text{II}(p, C)[r(w)]$

The *IE* alternatives are all those that can be assigned false consistently with the prejacent. The *II* alternatives are those that can be assigned true consistently with the prejacent and the falsity of all *IE* alternatives.

(38) Given a sentence p and a set of alternatives C:

a. $IE(p, C) = \bigcap\{C' \subseteq C : C'$ is a maximal subset of C, s.t.
$\{\neg q : q \in C'\} \cup \{p\}$ is consistent $\}$

b. $II(p, C) = \bigcap\{$C" $\subseteq$ C : C" is a maximal subset of C, s.t.
$\{r : r \in$ C"$\} \cup \{p\} \cup \{\neg q : q \in \text{IE}(p, C)\}$ is consistent $\}$

None of the alternatives given in (36) can be consistently assigned false while asserting the prejacent, but all of them can be consistently assigned true. The IE and II alternatives among the set of alternatives given in (36) are given below.

(39) a. $IE((35), (36)) = \varnothing$

b. $II((35), (36)) = (36)$

Upon exhaustification of our toy example, we will have (40), which is equivalent to a statement with a necessity modal.

(40) $\text{EXH}^{IE+II}(Alt(\exists w{\in}\{w_1,\ w_2,\ w_3\}\colon\ \neg\text{p}(w))) = \exists w{\in}\{w_1,\ w_2,\ w_3\}\colon\ \neg\text{p}(w) \wedge\ \exists w{\in}\{w_1,\ w_2\}\colon\ \neg\text{p}(w)\ \wedge\ \exists w{\in}\{w_1,\ w_3\}\colon\ \neg\text{p}(w) \wedge\ \exists w{\in}\{w_2,\ w_3\}\colon\ \neg\text{p}(w)\ \wedge\ \exists w{\in}\{\ w_1\}\colon\ \neg\text{p}(w)\ \wedge\ \exists w{\in}\{\ w_2\}\colon\ \neg\text{p}(w) \wedge\ \exists w{\in}\{\ w_3\}\colon\ \neg\text{p}(w) = \boldsymbol{\forall w{\in}\{w_1,\ w_2,\ w_3\}\colon\ \neg\text{p}(w)}$

Note that the resulting interpretation—where a universal modal takes scope over a negated proposition—is exactly the neg-raising reading we aim to capture. Namely, we want to derive an interpretation of the form *think* $\neg p$ from an original sentence of the form $\neg$*think* p. We thus propose that NRPs are best characterized as universal modals that, when negated, are strictly equivalent to a possibility modal scoping over negation. Next, we entertain the hypothesis that NR readings arise when a statement with the meaning equivalent to a possibility modal scoping over negation is strengthened to one with a necessity modal scoping over negation. This allows us not only to maintain a standard universal modal semantics for NR propositional attitude verbs like *think* and *believe* (à la Hintikka and Hintikka 1969, and contra Staniszweski 2021; Staniszewski 2022), but also to derive the robust generalization that the class of NRPs excludes factive modals.

As demonstrated in the previous section, $\Diamond$K $\neg$p(w) is not strictly equivalent to $\neg\Box_k$p(w). Consequently, no equivalent statement with the existential modal $\Diamond$K $\neg$p(w) is available to serve as the prejacent of EXH. Since, under our analysis, the NR reading is the result of exhaustifying an existential modal statement to a universal one, such a strengthened reading cannot be derived for factives like *know*. This means that it is not NRPs that are special in allowing NR inferences, but rather that strictly non-NRPs are special in disallowing them. Universal modals with a presupposition that is incompatible with their existential dual may never yield NR readings, as no strictly equivalent existential statement is available for further strengthening. Since under our analysis, the factivity presupposition is what blocking the strict equivalence and in turn blocks NR, we predict that non-factive *know* can in principle be a neg-raisier. The example in (41), where the lawyer is obviously aware that overturning the result of an election is not constitutionally possible, supports this prediction, as it allows a NR reading:[5]

(41) Trump: I can overturn the result of the election.
Constitutional lawyer: I don't know that is constitutionally possible, sir.

In addition to excluding factives from the class of NRPs, our assumption about LF-equivalence under strict duality is crucial for generating well-formed alternatives that serve as inputs to EXH. Let us consider an obvious alternative to our proposal, explored by Jeretič (2021); Jeretič (2022), where the neg-raising reading is derived by applying

[5]Note though, that does not mean that every non-NRP must have particular presuppositions that render existential LF-equivalents impossible. E.g., strong modals like *must* or *need* do not seem to have such a presupposition. Hence, at this stage our approach may still overgeneralize. However, in Section 3, we account for the lack of NR readings of strong modals in terms of the obligatory pruning of their singleton alternatives.

EXH directly to the original statement involving a negated universal modal, bypassing the intermediate step of LF-equivalence under strict duality.

(42) *Alt* ($\neg\forall w \in\{w_1,\ w_2,\ w_3\}$: p($w$)) = $\{\neg\forall w\in\{w_1,\ w_2,\ w_3\}$: p($w$), $\neg\forall w\in\{w_1,\ w_2\}$: p($w$), $\neg\forall w\in\{w_1,\ w_3\}$: p($w$), $\neg\forall w\in\{w_2,\ w_3\}$: p($w$), $\neg\forall w\in\{w_1\}$: p(w), $\neg\forall w\in\{\ w_2\}$: p($w$), $\neg\forall w\in\{\ w_3\}$: p($w$)$\}$

The set of alternatives given in (42) includes propositions that negate a universal quantification over singleton domain: $\neg\forall w\in\{\ w_1\}$: p($w$), $\neg\forall w\in\{\ w_2\}$: p($w$), $\neg\forall w\in\{\ w_3\}$: p($w$). Recall that the process of exauhstification involves assigning truth values to a set of alternatives and determining whether their falsity (for IE alternatives) or truth (for II alternatives) is consistent with the truth of the main assertion. This makes it essential to be explicit about how we reason about universal quantification over singleton domains. As the infelicity of (43) shows, such sentences seem to be undefined.

(43) # Not all earth's moons have/lack signs of life.

Thus we must clarify what it means to negate or assert an undefined alternative. There are mainly two possible approaches to deal with such alternatives. Under the first approach, which is taken by Alxatib (2014) and Spector and Sudo (2017), EXH $\phi(w)$ is undefined if any of ϕ's alternatives is undefined. Under this approach, the undefined alternatives involving singleton propositions cannot take part in the exhaustification process. As innocent inclusion of these alternative is essential to have a strengthened NR reading, applying EXH directly to the original statement involving a negated universal modal does not yield the desired NR reading.

An alternative approach might allow for undefined alternatives to be excluded or included if certain conditions hold. For example, exclusion could be framed as the non-truth, rather than the falsity, of an alternative, enabling undefined alternatives to be excluded (See Spector and Sudo (2017) who entertain this option). Undefined alternatives may also be considered includable, but only if their presupposition can be accommodated. However, presupposition accommodation is subject to constraints: most notably, the accommodated content must not conflict with the asserted content. Given that subdomain alternatives are specifically characterized by the size of their domain, it remains unclear what kind of content, if accommodated, would make these innocently includable singelton alternative sentences assertable. We conclude, therefore, that the assumption of LF-equivalence is essential not only to explain why certain negated necessity modals are never neg-raisers, but also to derive a strengthened reading while avoiding generating undefined alternatives.

Finally, under our approach, the generated subdomain alternatives with an existential quantifier are equivalent to a set of disjunctive alternatives, as in (44). As none of these alternatives is equivalent to the conjunction in (45), this set of disjunctive alternatives is

not closed under conjunction. Thus, our approach groups NR together with homogeneity and Free Choice, attributing their analogous behavior to the common core of non-closure under disjunction (Fox, 2007; Bar-Lev and Fox, 2020; Bar-Lev, 2021).

(44) a. $\exists w \in \{w_1, w_2, w_3\}$: $\neg p(w) \Leftrightarrow \neg p(w_1) \vee \neg p(w_2) \vee \neg p(w_3)$
b. $\exists w \in \{w_1, w_2\}$: $\neg p(w) \Leftrightarrow \neg p(w_1) \vee \neg p(w_2)$
c. $\exists w \in \{w_1, w_3\}$: $\neg p(w) \Leftrightarrow \neg p(w_1) \vee \neg p(w_3)$
d. $\exists w \in \{w_2, w_3\}$: $\neg p(w) \Leftrightarrow \neg p(w_2) \vee \neg p(w_3)$
e. $\exists w \in \{w_1\}$: $\neg p(w) \Leftrightarrow \neg p(w_1)$
f. $\exists w \in \{w_2\}$: $\neg p(w) \Leftrightarrow \neg p(w_2)$
g. $\exists w \in \{w_3\}$: $\neg p(w) \Leftrightarrow \neg p(w_3)$

(45) $\neg p(w_1) \wedge \neg p(w_2) \wedge \neg p(w_3) \notin Alt$

2.3 Weak Readings

It has been observed that there are contexts in which negated neg-raising predicates do not give rise to NR readings (Bartsch, 1973; Gajewski, 2005, 2007; Homer, 2015). The examples in (46a) and (47a), where the NR readings (46b) and (47b) are absent, illustrate this fact:

(46) a. Unlike many people nowadays, my great-grandparents didn't want to spend all their spare time on the internet.
b. My great-grandparents wanted not to spend all their spare time on the internet.(Homer, 2015)

(47) *Context: At a job interview. . .*
a. I don't want to make a lot of money, you know.
b. I want not to make a lot of money.(Homer, 2015)

Romoli (2012, 2013) argue that this observation poses a challenge for both syntactic and presuppositional semantic approaches to NR. In contrast, under implicature-based accounts of neg-raising, the absence of a NR reading follows naturally from the pragmatic nature of the strengthening process, given the widely held view that implicature derivation is context-dependent. As we mentioned earlier, it is widely assumed that only alternatives that are considered to be *relevant* in the context enter implicature computation. Naturally, contextual factors like the question under discussion (QUD) play a crucial role in determining the relevant set of alternatives under consideration. While differing in implementation, Romoli (2012, 2013) and Jeretič (2022) both attribute the suspension of NR readings to the role of the QUD. According to Romoli (2012, 2013), non-NR readings arise when the set of *relevant* alternatives supplied to the EXH operator includes only

the assertion itself. Jeretič (2022), on the other hand, derives non-NR readings from the distribution of the EXH operator, which need not apply if the unstrengthened reading is a complete answer to the QUD.

Similar to Romoli (2012, 2013), and building on the account of non-maximal readings of definite plurals proposed by Bar-Lev and Fox (2020), we assume that all sentences are parsed by applying exhaustification and that weak readings of negated neg-raising predicates are the result of applying exhaustification over a subset of subdomain alternatives. The reading derived from the exhaustification, however, depends on which alternatives are input to the EXH operator, and which alternatives are to be ignored (i.e. are to be *pruned*). Crucially, a NR reading can only be derived when all subdomain alternatives, including singleton ones, are considered relevant for the process of exhaustification. In the remainder of this section, we show that applying EXH over any proper subset of subdomain alternatives that adheres to the constraints on pruning alternatives will result in a weak non-neg-raising reading.

The contextual nature of implicature calculation is incorporated to the theory of exhaustification via the the requirement that alternatives to a sentence must be relevant. Following Bar-Lev (2021, 2024), we adopt the perspective that, rather than evaluating the relevance of each alternative individually, we should consider different pruning choices and assess whether they result in relevant propositions when the EXH operator is applied.

Let us now consider potential pruning choices available to the set of subdomain alternatives of the following modal assertion $S = \neg\forall w \in \{w_1, w_2, w_3\} : p(w)$, where S is strictly equivalent to $S' = \exists w \in \{w_1, w_2, w_3\} : \neg p$. The first option is to prune all subdomain alternatives except for the assertion itself, as shown in 48a. The second option is to only prune singleton alternatives, as in (48b).

(48) a. $\exists w \in \{w_1, w_2, w_3\} : \neg p(w)$; ~~$\exists w \in \{w_1, w_2\} : \neg p(w)$~~; ~~$\exists w \in \{w_1, w_3\} : \neg p(w)$~~; ~~$\exists w \in \{w_2, w_3\} : \neg p(w)$~~; ~~$\exists w \in \{w_1\} : \neg p(w)$~~; ~~$\exists w \in \{w_2\} : \neg p(w)$~~ ; ~~$\exists w \in \{w_3\} : \neg p(w)$~~

b. $\exists w \in \{w_1, w_2, w_3\} : \neg p(w)$; $\exists w \in \{w_1, w_2\} : \neg p(w)$; $\exists w \in \{w_1, w_3\} : \neg p(w)$; $\exists w \in \{w_2, w_3\} : \neg p(w)$; ~~$\exists w \in \{w_1\} : \neg p(w)$~~; ~~$\exists w \in \{w_2\} : \neg p(w)$~~ ; ~~$\exists w \in \{w_3\} : \neg p(w)$~~

As the innocent inclusion of singleton subdomain alternatives is crucial to derive a strengthened NR reading, applying EXH to either set of alternatives in (48a) or (48b) only a weak non-NR reading is generated, as can be confirmed in (49).

(49) a. EXH^{IE+II}(48a) = $\exists w{\in}\{w_1, w_2, w_3\}$: $\neg$p(w)

b. EXH^{IE+II}(48b) = $\exists w{\in}\{w_1, w_2, w_3\}$: $\neg$p(w)$\wedge$ $\exists w{\in}\{w_1, w_2\}$: $\neg$p(w) $\wedge$ $\exists w{\in}\{w_1, w_3\}$: $\neg$p(w)$\wedge$ $\exists w{\in}\{w_2, w_3\}$: $\neg$p(w)

Note that exhaustifying the alternative that remains after the first pruning choice does not yield a meaning that is distinct from the plain assertion without the application of EXH. In other words, applying EXH at this point does not strengthen the meaning of

the sentence. This outcome violates widely accepted constraints on the distribution of EXH, which require that its application must add to the information contributed by the plain assertion– that is, it must lead to strengthening (Chierchia et al. (2012); Chierchia (2013); Magri (2011), among others). Romoli (2012) formalizes this constraint as follows:

(50) *Minimize Weakness:*
Do not insert EXH in a sentence S if it leads to an equivalent or weaker meaning than S itself.

The resulting meaning in (49b), in contrast—which conveys that p is false in most worlds in the quantification domain—is weaker than the NR-reading, which commits to the falsity of p in all worlds. Nonetheless, it is still stronger than an assertion that merely requires the existence of one world in which p is false.

There are many more pruning choices we could have considered but they are ruled out by the constraints on pruning mechanism. For instance, we could imagine to prune all subdomain alternatives other than $\exists w \in \{w_1\} : \neg p(w)$, where the resulting meaning after the exhaustification would be that $\exists w \in \{w_1,\ w_2,\ w_3\}$: $\neg$ p(w) $\wedge$ $\neg \exists\ w \in \{w_1\}$: $\neg$p(w). Such pruning choices, however, are ruled out by the principle that pruning can only weaken the meaning (Crnič et al., 2015; Bar-Lev, 2018, 2021).

So far, we have derived two distinct readings for a statement with a negated NRP, depending on which alternatives are input to EXH. It seems, however, that such a statement can have an even weaker readings. For instance, it is unlikely that the most desire worlds of the speaker of (47) are ones in which she does not make a lot of money. This observation leads us to reconsider the type of alternatives we have been assuming. Up to this point, we have assumed that the alternatives introduced by NRPs are subdomain alternatives. This raises the question of what happens if EXH applies to lexical alternatives in assertions involving NRPs. First, note that NRPs appear to lack a lexically overt existential counterpart (Jeretič, 2021; Jeretič, 2022). But what about under a conceptual view of alternatives, such as the one proposed by Buccola et al. (2021), which we have adopted in this paper? Given an assertion $\neg\Box p$, EXH can take its strictly equivalent LF $\Diamond\neg p$ as input. Substituting the scalar item $\Diamond$ with a stronger scale-mate $\Box$, we will have $\Box\neg p$ as a stronger alternative to the assertion. Thus, exhaustification can result in a stronger reading $\Diamond\neg p \wedge \neg\Box\neg p$ as the strengthened reading. Crucially, while this reading is distinct from the NR reading, it is still stronger than assertion without the application of EXH, and still distinct from the reading derived from the application of EXH over the set of subdomain alternatives remained after pruning singelton alternatives. Luckily, nothing in our account blocks the generation of this reading. Let us see why.

There are three logical possibilities representing the speaker's belief states: speaker believes p ($\Box p$), speaker believes not-p ($\Box\neg p$), or speaker has not formed a belief and holds both p and not-p to be possible ($\Diamond p \wedge \Diamond\neg p$). While the plain assertion $\neg\Box p$ is

compatible with both second and third possibilities, $\Diamond p \wedge \Diamond \neg p$ uniquely describes the third logical possibility. Thus, exhaustification is not vacuous as the resulting meaning is still stronger than the assertion itself.

Note, however, that unless the exhaustification system is revised to allow for more complex lexical alternatives, like what is proposed by Romoli (2012, 2013), the NR reading can only be derived via exhaustification of sub-domain alternatives. When the question under discussion (QUD) is such that only beliefs held by the speaker matter, this other derived non-NR reading is not more informative than the assertion without the insertion of EXH. Note that both the assertion without EXH (i.e. $\neg \Box p$) and the output of this EXH (i.e. $\Diamond \neg p \wedge \neg \Box \neg p$) describe the third logical possibility (i.e. $\Diamond p \wedge \Diamond \neg p$) which will not be relevant to the QUD. In such contexts, the LF with an EXH that takes sub-domain alternatives as input is always preferred, as it yields the stronger reading (i.e. $\Box \neg p$). Thus, under our approach, both exhaustification can apply to either lexical and sub-domain alternatives. Among the possible outputs, however, the interpretation that provides the most informative answer to the QUD in a given context is always the one selected.

Having outlined the general schema of our proposal, we now turn to the desiderata for any theory of NR discussed in the first section, and examine how they are accounted for under our approach. The first desideratum concerns the observed semantic generality and variation among NRPs (Section 3), and the second addresses the grammatical reflexes of NR (Section 4).

3 Semantic generality and variation in NRPs

As we have already discussed, the semantic generality of NRPs and non-NRPs favors a semantic/pragmatic approaches to NR. Since the seminal work of Horn (1978), however, there have been no significant attempts to derive the semantic generalizations that govern membership in the class of NRPs. Facing the difficulty of accounting for apparent exceptions to these otherwise robust generalizations, semantic-pragmatic approaches, just like the syntactic approaches, have modeled a predicate's NR property as the result of the presence of a lexically-encoded feature on the NRP. A unique feature of our approach is that not only does it account for the robust mid-scalar generalization but also provides a new perspective on the idiosyncratic lexical exceptions of potential neg-raisers.

First, under our framework, only predicates whose negation is strictly equivalent to an existential modal with a negated prejacent can yield a strengthened NR reading through exhaustification. This excludes both existential modals and factive universal modals—respectively corresponding to low-points and the high endpoint, in Horn's scale of negation—from the group of NRPs. Moreover, our account also derives another well-attested generalization in the literature: the availability of NR inferences depends on

the eventivity or stativity of the predicate, with only stative predicates permitting Neg-raising readings. All Neg-raisers are stative predicates, and in cases where an eventive counterpart of a Neg-raiser exists, e.g., *think*, it does not give rise to a NR interpretations (Bervoets, 2014; Xiang, 2013; Ozyıldız, 2021; Bondarenko, 2022; Jeretic and Ozyıldız, 2022). Under our approach, eventive predicates are banned from being a NRP, as the negated form of these predicates involves a negation scoping over an existential quantifier over event variables (see also Jeretic and Ozyıldız (2022)). Consequently, they cannot be strictly equivalent to an existential scoping over a negated proposition. Note that since speech report predicates like *say* and *claim* are eventive (Moulton, 2009), they are correctly predicted to disallow NR inferences. However, as Moulton (2009) observes, there are cases where *say*, in combination with a modal, functions not as a speech report but as a belief report. As we will see in the next section, it is precisely in such contexts that *say* exhibits some grammatical properties characteristic of NRPs, such as licensing strong NPIs and Horn clauses.[6]

To fully derive Horn's mid-scalar generalization, however, we still need to explain why other non-factive high scalar items, such as modals like *must* or *have to*, are never NRPs.

First, note that strong universal modals can have strictly equivalent existential duals, as illustrated by the equivalence between the two statements below.[7]

(51) a. You don't have to leave.

b. You are allowed to not leave.

As such, the explanation for why factive *know* is not neg-raiser that we provided before, does not apply here. Instead, we propose that in the process of exhaustifying their corresponding strict equivalent LF – where an existential modal takes a negated prejacent – the singleton sub-domain alternatives of these modals are systematically and obligatorily pruned from the set of alternatives considered by EXH. Given that the inclusion of alternatives involving singleton sets, as we saw before, is necessary to get a strong NR reading, this obligatorily pruning of singleton alternatives rules out the derivation of a NR reading for these modals.

The question is what explains the obligatory pruning of the singleton sub-domain alternatives of strong modals. We observe that these modals are strong modals in the sense of Silk (2016, 2018, 2022), according to whom the defining characteristics of strong modals (both $\Box\phi$ and $\Diamond\phi$) is that they predicate the necessity or possibility of the prejacent ϕ of the actual world. Statements with weak modals, on the other hand, are accepted

[6]There are also cases where non-eventive *say* can be used in as a belief report without a modal. For example, in a sentence like *Nothing says that this solution is the only valid one*, the verb *say* does not refer to any speech event. It seems to us that in this usage, *say* may function as a potential NRP, as the sentence conveys *this solution is not the only valid one*.

[7]Our account does not rely on the availability of a lexical dual. However, for some strong modals like *have to*, the equivalence also holds at the lexical level.

without needing to be verified in the actual world. Strong modals thus presuppose that the actual world is part of their domain, but the quantificational domain of weak modals does not need to include actual world. We formulate this idea via the necessary inclusion of the actual world in the quantificational domain of strong modals (see also Mirrazi and Zeijlstra (to appear)).

Since the actual world is unique, its inclusion in the quantification domain of modals renders a non-homogeneous set of worlds, with one world differing from the rest by virtue of being the actual world. When a proposition containing a strong modal is input to an EXH operator, innocent inclusion of its singleton sub-domain alternatives would involve non-trivial task of distinguishing between the actual world and the rest of worlds in the domain. As Kratzer (2012) notes, singleton propositions that are uniquely predicated of the actual world are not cognitively viable, as they assume an omniscient speaker capable of distinguishing the actual world from all its duplicates. We propose that due to this cognitive infeasibility, singleton sub-domain alternatives of strong modals are systematically pruned. Since assertion of these alternatives are essential for deriving a strengthened reading, exhaustification of strong modals does not result in a NR reading.

Given the requirement for the availability of a strictly equivalent LF corresponding to an existential modal with a negated prejacent, and the obligatory pruning of sub-domain alternatives of strong modals, our approach limits the class of NRPs to weak necessity modals. Thus Horn's mid-scalar generalization is effectively derived under approach.

Finally, our approach opens up a new perspective to understand the cross-linguistic variation underlying NRPs.

Under our approach, NR inferences are not determined at the lexical level; rather, it is the semantic content of the entire proposition, as represented at LF, that determines whether a strict LF with an existential modal scoping over negation is available for further strengthening. As such, subtle aspects of predicate meanings including their presuppositions they carry and the specifications of the possible worlds in their quantification domain, are crucial in determining whether a predicate is a NRP. The reason is that these aspects of the meaning may disrupt strict duality or necessitate pruning of singleton alternatives, which in turn blocks the strengthening process.

While cross-linguistic research is necessary to map the range of variation in the class of neg-raisers across languages and to provide systematic explanations for each case, we end this section by noting that our approach sheds light on the long-standing puzzle of why Dutch *hopen* and German *hoffen* are classified as NRPs, whereas English *hope* is generally assumed to not permit NR.

Consider the contrast between Dutch and English in the examples below.

(52) If you hope to be a millionaire, you are irrational.

(53) ?? Als je miljonair hoopt te worden, ben je niet rationeel
if you millionaire hope to become, you are not rational
roughly ??'If you desire to be a millionaire, you are irrational.'

While the conditional in English is felicitous and true – since the content of the desire predicate *hope* must be consistent with the context set – the infelicity of the Dutch counterpart reveals that *hopen* can describe a desire without requiring it to necessarily be attainable in the actual world, as having the desire of being a millionaire is not irrational. Under our approach, this additional component of English *hope* invokes the inclusion of the actual world in its domain, restricting *hope* to desires that are attainable in the actual world. This, in turn, necessitates the pruning of singleton sub-domain alternatives, preventing the predicate from yielding a neg-raising reading through exhaustification. Hence, it cannot invoke NR reading. As there is no such requirement for Dutch *hopen*, a NR reading is available with it.

This exemplifies how subtle cross-linguistic differences in the semantic content of predicates can result in variation in their ability to license Neg-raising inferences. A predicate that qualifies as a NRP in one language may fail to do so in another due to fine-grained distinctions in lexical meaning. A comprehensive investigation of such cross-linguistic variation among NRPs is left to future work.

In the next section, we turn to grammatical properties of NRPs that *prima facie* appear to challenge a pragmatic-semantic approach to Neg-raising. However, we show that, upon closer examination, these cases in fact are more in line with our semantic pragmatic account than with syntactic approaches to Neg-raising.

4 Grammatical footprints of Neg-raising: Strict NPIs and Horn clauses

We have seen that our approach successfully accounts for the generalizations observed in the semantic properties of NRPs – a result that is, perhaps, to be expected from a semantic-pragmatic approach. As discussed in the first section, however, the real challenge for semantic-pragmatic approaches lies in explaining the grammatical properties traditionally associated with NR – namely, the licensing of strong NPIs and Horn clauses in NR environments – which are more readily captured under syntactic accounts.

In this section, we argue that our approach not only accounts for these grammatical reflexes of Neg-raising but in fact offers a more principled and empirically adequate explanation than syntactic theories.

4.1 Strict NPIs

As mentioned in Section 1.2, with examples in (20), NPIs like punctual *until* or *breath a word* are fine under negated NRPs, but not under (most) negated non-NRPs. The relevant examples are repeated below.

(54) a. Stanley doesn't believe that Carolyn will *breathe a word* about it.
b. Calvin didn't think that Mona moved in *until June*.

(55) a. *Stanley doesn't claim that Carolyn will *breathe a word* about it.
b. *Calvin didn't say that Mona moved in *until June*.

The licensing of such NPIs under negated NRPs is surprising, as we have seen earlier that anti-additivity alone is not enough to license strong strict NPIs, and that their licensing is subject to locality constraints. Existing semantic-pragmatic approaches have problems accounting for the locality property of strict NPIs. Specifically, for NPIs that require a local licenser, it remains unexplained why, in NR environments such as (54), negation can license a strict NPI that is not within the same local domain.

To this date, no semantic-pragmatic approach has offered a satisfactory solution to this locality problem. By contrast, the syntactic approach to NR has no problem with that. Under such an approach, the negation would start out in the embedded clause, license the strict NPI there, and then raise into the matrix clause.

At the same time, such a syntactic approach faces a number of difficulties of its own. Some of these problems have been addressed already in the literature[8]; others, like some of the issues discussed below, are novel.

One such problem for Collins and Postal (2014) is that the set of predicates that, when negated, can license strict NPIs is not restricted to NRPs. Horn (2014) points out that non-factive *know, be aware*, and some other predicates, which he dubs *Cloud of Unknowing* (CoU) predicates, license strict NPIs as well (next to the ones exemplified before, also strong NPIs like *all that* or *any ... at all*), as shown below (examples taken from Collins and Postal (2018b)):

(56) a. I don't know that it was *all that* easy even back then.
b. Bruschetta is quite delicious as well, but I can't say I've indulged *in ages*.

However, in (56), there is no semantic reflection of negation in the embedded clause, i.e. the examples in (56) lack a NR reading. But if negation started out in the embedded position a NR reading should be available, contrary to fact.

To resolve this, Collins and Postal (2018b) propose that when a negation raises in a main clause with a CoU predicate, the underlying structure of (56a) contains two addi-

[8]See, for instance, Jacobson (2018, 2020), pace Crowley (2019) or Sato (2023).

tional negations in the embedded clause, one of which must raise into the matrix clause, followed by phonological deletion of the these two lower negations. We show this in (57):

(57) [I do NEG$_1$ know ~~NEG$_2$~~ [that it was <NEG$_2$ > ~~NEG$_3$~~ all that easy back then]]

It must be said that such a mechanism, where a downward entailing element may license the phonological deletion of two lower negations (of which one should move from the embedded into the main clause), is central to Collins and Postal (2014)'s work. For instance, they argue that it is also what underlies the NR reading behind examples like (58), which they assign the LF in (59):

(58) Nobody supposes that nuclear war is winnable

(59) Nobody NEG$_1$ supposes ~~NEG$_2$~~ [that < NEG$_2$ > nuclear war ~~NEG$_3$~~ is winnable]]

The reason why they need to allude to such a mechanism is that examples like (58) would otherwise be a major problem for the syntactic approach, the reason being that the strengthened reading of (58) is that *everybody* thinks that nuclear war is not winnable, not that *somebody* thinks that nuclear war is not winnable. Given that *nobody* is a negated existential quantifier / negative indefinite[9], this reading would never come about if the negation was base-generated and interpreted in the embedded clause and moved to the matrix clause. Then the NR reading would only be that that somebody thinks that nuclear war is not winnable.

However, the motivation for adopting such a more complex treatment appears to rest primarily on the assumption that no simpler alternative is currently available. In other words, the plausibility of (57) as an analysis depends on the view that the semantic-pragmatic approach cannot, on its own, fully account for the relevant data.

Apart from that, note that (57) is not the only available parse for (56a). (60), which does yield a NR reading, is in principle an available parse as well, given that in (56) NEG-movement must be allowed (otherwise NEG$_2$ could never move from the embedded into the main clause, which for Collins and Postal (2014) is a necessary condition for the phonological deletion).

(60) [I do NEG know [that it was <NEG> ever all that easy back then]]

Collins and Postal (2018b) rule out the availability of (60) by stipulating a condition that states that if a negation raises into a clause containing a negated COU predicate, this predicate must be under the scope of a distinct negation (Collins and Postal, 2018b). That is the case in (57), but not in (60). Apart from the fact that such a condition is purely stipulative, this condition also turns out to be empirically flawed. Note that (61) under that condition should be fine with a NR reading, as here the predicate is indeed outscoped by a distinct negation (*nobody*), contrary to fact:

[9]See Iatridou and Sichel (2011) for a battery of arguments why negated existential quantifiers / negative indefinites should be decomposed in a negation and an existential/indefinite.

(61) Nobody doesn't know that it was all that easy back then.

Hence, the existence of strict NPIs with CoU predicates remains problematic for the syntactic approach to NR. Thus, we conclude that the syntactic approach still suffers from some serious problems when accounting for the licensing of strong/strict NPIs under negated NRPs.

Instead, our proposed mechanism in terms of strict equivalence can actually explain the attested observations. We propose that when determining whether the conditions for the licensing of strict NPIs are satisfied, the grammar takes into account strictly equivalent LFs (see also Linebarger (1987))– provided that they are no complex than the original assertion .That is, the locality requirement is satisfied if the sentence under consideration has a strictly equivalent LF where the strict NPI is in the same local syntactic domain as its licenser and not separated from it by a relevant intervener, such as an island. As we characterized NRPs as (weak) necessity modals that when negated are strictly equivalent to a possibility modal scoping over a negated prejacent, it follows that NPIs under negated NRPs will be in the same local syntactic domain as the negation.

(62) $\neg\Box[_p$ ....NPI...$] \equiv \Diamond[\neg\ [_p$ ....NPI...$]]$

For instance, the conditions on licensing *breathe a word* are straightforwardly met when in (54a), represented at (63a), as the strictly equivalent LF in (63b) is available:

(63) a. Stanley $\neg\Box$ [Carolyn will breathe a word about it then].

b. Stanley $\Diamond[\neg$ Carolyn will breathe a word about it then].

Moreover, it can be independently shown that negation taking immediate scope over a clause can license strict NPIs within that clause, as evidenced by (64). This provides empirical support for the claim that such configurations satisfy the locality conditions required for NPI-licensing.

(64) a. She apologized—not that it helped at all.

b. I called—not that he gave a damn.

Note that under strict duality, a negated universal modal is equivalent to an existential modal scoping over negation, where the negation takes scope above the embedded proposition. In the absence of any intervening element, a strong, strict NPI inside the embedded proposition will be in the local scope of this negation. However, if an island intervenes between this negation and the strong strict NPI, we expect ungrammaticality as the locality requirement is not satisfied.

(65) $*\neg\Box$ [*intervener* $[_p$ ....NPI...$]] \equiv \Diamond[\neg[$ *intervener* $[_p$ ....NPI...$]]]$

This is precisely the pattern Collins and Postal (2014) observed. That is, strict NPIs in syntactic islands cannot be licensed by a clause-external negated NRP, as shown by the ungrammaticality of the examples in (66).

(66) a. *I don't believe the rumor that Tom has found the solution *in years*.

b. *I don't think Tom has found the solution *in years* and is a reliable chap.

Note that in the strict dual LFs of (66a) and (66b) respectively represented in (67a) and (67b), *in years* is still separated from the negation by island boundaries.

(67) a. I ◇[¬ [**the rumor** [that Tom has found the solution *in years*]]].

b. I ◇ [¬ [[Tom has found the solution *in years*] **and** [is a reliable chap]]].

In addition to deriving island constraints on the licensing of strict, strong NPIs, our account also rules out sentences with other elements like universal quantifiers that intervene between the NPI and its licensor (Linebarger, 1987; Guerzoni, 2006; Sailer, 2007; Chierchia, 2013).

(68) a. *Calvin doesn't think that everyone will *breathe a word* about it.

b. *Calvin ◇[¬ [**everyone** [will *breathe a word* about it]]].

Without further complicating their system, Collins and Postal (2014) cannot explain such intervention effects as in their system negation is originated as adjoining to the NPI, and is thus in local relationship to it.

It is important to note that our account only rules out cases in which a strict, strong NPI takes scope under a universal quantifier. A sentence containing the temporal adverbial *until June*, which takes scope above the subject, is judged to be grammatical.

(69) a. Calvin don't think that everyone tenant will move in *until June*.

b. Calvin ◇[¬ *until June* [**every tenant** will move in]]].

Under our approach, strict duality is a necessary but not sufficient condition for deriving NR inference. The strictly equivalent LF needs to be pragmatically enriched to yield NR inference. In contrast, the licensing of strong, strict NPIs requires only the existence of some LF – either the surface LF or a strictly equivalent one – in which the NPI is in the scope of a local negation. Therefore, we predict the set of licensers for strict, strong NPIs to be a superset of NRPs. Thus, the observation that in addition to negated NRPs, CoU predicates can also license strict, strong NPIs, as observed in (56), far from being a problem for our account, is in fact predicted by our account. We characterize CoU predicates as predicates that when negated, can have a strictly equivalent LF where negation scopes below–providing the configuration required for licensing strong, strict NPIs, irrespective of whether or not there is further pragmatic enrichment to result in NR inference arises.

As was discussed earlier, the presence of factive presuppositions is the reason why strict equivalence does not hold for factive *know*, which is therefore correctly predicted not to license strong, strict NPIs. By contrast, we also correctly predict that non-factive

uses of *know* – which do have a strictly equivalent LF with negation scoping below an existential dual – can license strong, strict NPIs. Moreover, as noted in the previous section, strict equivalence also accounts for why an eventive predicate like *say* is correctly excluded from the class of NRPs. Negating an existential quantifier over events does not have the required strictly equivalent LF. However, as observed by Moulton (2009), *say* also has a non-eventive use in which it functions more like a belief report than a speech report. While a detailed investigation of the semantic contribution of such uses of *say* is beyond the scope of this paper, we note that our account makes a clear prediction: in the absence of eventivity–i.e., when strict equivalence is not disrupted–and assuming Moulton's intuition that these non-eventive uses of *say* semantically resemble belief reports, it is not surprising that they would license strict, strong NPIs.

Before turning to Horn clauses, we would like to consider how other semantic-pragmatic approaches account for the licensing of strict, strong NPIs in NR environments. As mentioned earlier, these approaches fail to explain the locality constraints observed with such NPIs. Here, we further show that they also face difficulties accounting for why CoU predicates can license strict, strong NPIs.

Gajewski (2007) argues that his semantic pragmatic approach in terms of excluded middle presuppositions can account for the licensing of strong NPIs, as it could be maintained that what matters is whether the overall meaning contribution is anti-additive and not just the assertion. Since the enriched reading of a NR construction has the prejacent in an anti-additive environment, such NPIs would be licensed indeed. After all, the overall meaning contribution of, for instance, (54b) is that Calvin believes that Mona wouldn't move in until June, so *until* would be in the direct scope of a negation. By contrast, since (55b) lacks this overall meaning contribution, this solution wouldn't apply to this case, thus neatly explaining the contrast.

In light of problems with anti-additivity as a necessary property of the licenser, however, Gajewski (2011) and Chierchia (2013) propose that weak and strong NPIs are distinguished in terms of sensitivity to non-truth-conditional meanings. The idea is that, while for the licensing of weak NPIs, only the assertion needs to be downward entailing, the evaluation of downward entailment for the licensing of strong NPIs also takes presuppositions and implicatures into account (Gajewski (2011)). That a sentence like *Few students moved in until June* is bad under such approaches is not due to *until June* ending up in a non-anti-additive environment but because it infers that *some student(s) moved in until June*, and *some student(s)* is not downward entailing.

Under such more recent approaches to strong/weak NPI-hood, the negated NRPs (be they pragmatically enriched or not) and non-NRPs in (54)-(55) would both end up rendering the strong NPIs in a downward entailing context, and thus allow strong NRP to be licensed under both negated NRPs and negated non-NRPs. Hence, semantic pragmatic approaches to NR would overgenerate, unless it can be shown that non-NRPs trigger in-

ferences that make the overall meaning contribution no longer downward entailing. This is in fact what Romoli (2013) suggests. He proposes that while the NR inference of NRPs leaves the downward entailment of the environment intact, non-NRPs come with an existential inference that disrupt it, thus disallowing strong NPIs in their complements. For instance, consider the contrast between (70a) and (70b).

(70) a. John doesn't think that Mary will arrive until tomorrow.
b. *John isn't certain that Mary will arrive until tomorrow.

Romoli (2013) proposes that strong NPIs are not licensed under *certain* because such non-NRPs come with an existential inference which disrupts the downward entailment of the enriched meaning. (71a) infers (71b)

(71) a. John isn't certain that Mary was here.
b. it's possible for John that Mary was here.

NRPs, on the other hand, lack such an existential inference. The NR inference (72b), of (72a) for instance, does not disrupt downward entailment.

(72) a. John doesn't think that Mary was here.
b. it's impossible for John that Mary was here.

Romoli (2012, 2013) thus proposes that there are two kinds of attitude predicates, which are universal quantifies:NRPs which come with an excluded middle alternative and non-NRPs which have their existential counterpart as alternatives. Negated NRPs create a downward entailing environment in both their assertive and enriched meanings, allowing for the licensing of strict, strong NPIs. In contrast, negated non-NRPs imply that their existential counterpart is true, thereby disrupting downward entailment and preventing the licensing of such NPIs. As a result of this binary distinction, the licensing of strict, strong NPIs is exclusively limited to NRPs. Therefore, this proposal cannot account for why CoU predicates can license strict, strong NPIs. Since CoU predicates are not NRPs, they have their existential counterparts as alternatives in Romoli's system. As a result, their negation is expected to trigger existential inferences that disrupt DEness. Yet, these predicates still license strong NPIs, which is problematic for this system.

We conclude that, apart from the locality facts discussed above, our approach also fares better than existing approaches when it comes to the licensing of strong NPIs under negated NRPs, CoU predicates.

4.2 Horn clauses

Along similar lines, we argue that the existence and distribution of Horn clauses are naturally accounted for under our system. In fact, we argue here as well that our proposal

is not just an alternative to syntactic accounts of Horn-clauses (most notably, Collins and Postal (2014)), but actually fares better in capturing the distribution of Horn-clauses in CoU predicates without appealing to complex stipulations.

This is because the analysis proposed by Collins and Postal (2014) encounters the same challenges we identified earlier in the context of licensing strong, strict NPIs. As with strong NPIs, Horn clauses can be licensed under negated CoU predicates, as in (73). However, Collins and Postal (2014)'s solution – in terms of the deletion of two lower negation copies – mirrors their treatment of strong NPIs and thus inherits the same problems discussed in the previous section.

(73) I don't know that ever before had all three boys napped simultaneously.

While we believe that the proposal by Collins and Postal (2014) falls short of providing a complete explanation for the observed distribution of Horn clauses, as we maintain that the licensing of Horn clauses and strict, strong NPIs under NRPs and under CoU predicates reflects the same underlying phenomenon and should receive a unified explanation. At the same time, following their observation that Horn clauses are subject to locality constraints – a property we have also observed in the case of strict, strong NPIs – we take this as further motivation to reconsider existing semantic-pragmatic approaches to NR.

As a starting point, we follow Collins and Postal (2014) in taking Horn clauses to involve negative inversion. Second, drawing on insights from Büring (2004) and Collins and Postal (2014) – we assume that the necessary conditions for the availability of the negative version parallels that of licensing strict, strong NPIs: in addition to being sensitive to the monotonicity of the environment, negative inversion is also constrained by locality. The unique aspect of our proposal is that, just as with strong, strict NPIs, the grammar considers strictly equivalent LFs when determining whether the conditions for negative inversion are met. Under our approach, both NRPs and CoU predicates – such as non-factive *know* – are, when negated, strictly equivalent to an LF in which negation scopes below the existential counterpart of the predicate. This enables our account to successfully explain the licensing of Horn clauses under CoU predicates without relying on more complex or less motivated mechanisms like those proposed by Collins and Postal (2014).

Let us begin with the first characteristics of negative inversion. According to Büring (2004), when a negative phrase is preposed to a position that scopes over existential closure of the eventuality[10] variable ($\exists e$) – that is, when it functions as sentential negation – subject–auxiliary inversion is obligatory (Acquaviva, 1994; Penka, 2011). This is illustrated in examples (74)–(76), taken from Büring (2004).

(74) a. Nowhere is this so noticeable as in the South of France.

[10]The term *eventuality* is meant to refer to any event-like entity; this includes events, states, situations, etc.

b. *Nowhere this is so noticeable as in the South of France.

(75) a. Not until the next morning did she realize how serious it was.

b. *Not until the next morning she realized how serious it was.

(76) a. Seldom has so much been owed by so many to so few.

b. *Seldom so much has been owed by so many to so few.

There are, however, cases where negative inversion appears to be optional. In the well-known example in (77), adapted from Klima (1964), both the inverted and non-inverted orders are acceptable:

(77) a. With no job is Kim happy.

b. With no job, Kim is happy.

Building on the observation that these two variants yield distinct interpretations, it has been proposed that they differ in underlying structure (Liberman, 1974; Büring, 2004). In the inverted version (77a), the preposed phrase *with no job* takes wide scope over $\exists e$, yielding sentential negation in the sense of Klima (1964). By contrast, in the non-inverted version (77b), *with no job* takes narrow scope below $\exists e$. This structural distinction is reflected in their meanings: (77a) conveys that there is no eventuality such that Kim is happy with a job, while (77b) asserts that there is an eventuality where Kim is happy, and in that eventuality, she has no job.

Another important characteristic of negative inversion is that its licensors extend well beyond simple negation. Much like the licensors of NPIs, the key unifying feature of negative inversion is that the fronted phrases create a downward entailing environment. Thus, we conclude with Büring (2004) and Collins and Postal (2014) that whatever licenses an NPI, when fronted, also licenses negative inversion, as shown in (78).[11]

(78) a. Only two of them did he find useful.

b. Only two of them *lifted a finger* to help.

Combined with our proposal that the necessary conditions for negative inversion can be satisfied at a strictly equivalent LF no more complex than the original sentence, we formulate the conditions on negative inversion as follows:

(79) *The Negative Inversion Condition*
If the structure $[[_{\text{X}} <...>] \exists e]$ is available at the LF of a sentence, or at a strictly equivalent LF no more complex than the original LF, such that:

a. X contains a fronted phrase and X is downward entailing;

[11] *Only* is not DE *strictu sensu* but rather Strawson DE, cf. Von Fintel (1999).

b. $\exists e$ is the existential closure over eventuality described by the main proposition and is in the local scope of X

c. there is no relevant intervener between X and $\exists e$;

Then, negative inversion is licensed.

Given (79), we observe that both NRPs and CoU predicates have a strict LF equivalent as represented in (80) where fronted existential NPIs in Horn clauses preserve the downward-entailingness and thereby create the appropriate environment for negative inversion:

(80) $\neg\Box[$ <NPI> $\exists e\ ] \equiv \Diamond[\neg$ <NPI> $\exists e]$

(81) a. I don't believe/ know that *under any circumstances* would he agree.

b. I $\Diamond$ [$\neg$ <*under any circumstances*> would he agree.]]

Note that an NPI, which denotes an existential quantifier, does not by itself create a downward environment; its ability to satisfy the conditions for negative inversion crucially depends on strict equivalence to provide a local negation. Furthermore, since negative inversion is subject to locality constraints, any intervening element between the negation and the existential closure of the eventuality described by the embedded proposition is expected to disrupt the licensing conditions and block negative inversion. Thus, we predict that negative inversion will be ungrammatical in the presence of an island or when the fronted element itself acts as an intervener – such as universal quantifiers or presuppositional elements that disrupts DEness.

(82) a. *I don't believe the rumor that *under any circumstances* would he agree.

b. *I $\Diamond$ [$\neg$ [**the rumor** [that <*under any circumstances*> would he agree.]]

(83) a. I didn't believe *all gorillas/*every gorilla/*many gorillas/*a lot of gorillas could they teach to speak Mohawk. (Collins and Postal, 2014)

b. *I $\Diamond$ [$\neg$ [<**all gorillas/every gorilla/many gorillas/a lot of gorillas**> could they teach to speak Mohawk.]]

Hence our approach, without requiring any additional stipulations specific to Horn clauses, beyond our independently motivated assumptions about strict equivalence and negative inversion, correctly predicts the existence and distribution of Horn clauses under NR and CoU predicates, unlike the syntactic approach to NR.

Before concluding this section, we would like to address an example presented by Horn (2014). Similar to other cases involving the licensing of Horn clauses under CoU predicates, which challenge the syntactic account proposed by Collins and Postal (2014), this example shows that a Horn clause can be licensed by*accept* – a predicate that is not an NRP.

(84) I didn't accept that any of those problems had she ever really solved.

We have noted that strict equivalence is not calculated lexically; rather, it is the overall semantic content of the proposition, as represented at LF, that determines whether a strict LF with an existential modal scoping over negation is available. This allows us to understand why *accept* despite not being a NRP can nonetheless license strict strong NPIs. Again, a full semantic analysis of *accept* and other CoU predicates is beyond the scope of this paper, but we would like to we offer a few speculative remarks as a starting point for future work. We take *accept* as a doxastic necessity modal that also encode a change of mental state, with a semantic roughly as given below.

(85) $[\![\text{accept}]\!] = \lambda p.\lambda x.\lambda e.\lambda t.\exists e[\text{Accept}(e,x) \wedge \tau(e) = t \wedge \mathit{Result}\ (e) = \forall w' \in \mathit{Dox}(x)(w) : p(w')\ \text{at}\ t \wedge \forall t' < t.\ \underline{\neg \forall w'' \in \mathit{Dox}(x)(w) : p(w'')}\ \text{at}\ t']]]$

The negation of *accept* conveys that the agent's belief state has not undergone any change with respect to the relevant proposition: the content of belief at time t remains identical to what it was prior to t– underlined in the formula above. As such, the belief content of the agent both prior and at t is strictly equivalent to $\exists w'' \in \mathit{Dox}(x)(w) : \neg p(w'')$.[12]

5 Conclusion

We began this paper by identifying two key desiderata that any successful theory of NR should meet. First, following the seminal work of Horn (1978, 1989), we noted that membership in the class of NRPs is not arbitrary. Rather, there is a robust cross-linguistic generalization: NRPs tend to be mid-scalar predicates. At the same time, there is notable variation among semantically similar lexical items in whether they allow NR. A satisfactory theory must therefore not only capture the general pattern but also account for the finer-grained distinctions observed both within and across languages.

Second, NR is associated with grammatical footprints that seem at adds with otherwise semantic pragmatic nature of NR. A successful theory must therefore do more than identify the semantic factors that distinguish NR predicates from non-NR ones; it must also provide a principled account of the grammatical properties associated with NR. Ideally, such a theory should unify these semantic and syntactic aspects, showing how they arise from a common underlying mechanism.

We have proposed that NR readings are available only with predicates that meet two conditions: (i) they are universal modals whose negated form is strictly equivalent to an existential counterpart scoping over negation, and (ii) this existential modal can be

[12] Abstracting away from the lexical item, we can observe that the belief component of *accept* can undergo strengthening in a manner similar to that observed with NRPs. For instance, (84) conveys that the speaker believes she did not really resolve any problems.

strengthened into a universal interpretation via an exhaustification operator, following Bar-Lev and Fox (2017). On this view, non-NRPs fail to meet one or both of these conditions.

We have argued that the first condition is necessary not only to account for the grammatical properties of NR – such as the locality constraints on the licensing of strong or strict NPIs and Horn clauses – but also to explain why these properties are shared by both CoU predicates and NRPs. The locality constrain follow from the existence of a strict LF equivalent in which negation appears below the modal – an option available to both CoU predicates and NRPs alike – thereby enabling the local licensing of strong, strict NPIs and Horn clauses. This is something that purely semantic pragmatic approaches based solely on exhaustification cannot adequately explain.

Our account also captures the semantic commonalities that distinguish NRPs from non-NRPs, while also offering a way to understand cross-linguistic variation among NRPs. Under our approach, subtle aspects of predicate meaning, including the presuppositions they carry, are crucial in determining whether a strictly equivalent LF with negation scoping low is available for further strengthening. We have shown that the absence of such a strictly equivalent LF explains why factives and eventives are never NRPs. Moreover, the specification of possible worlds in the quantificational domain of modals plays an important role in whether a NR reading can be derived through exhaustification. This accounts for why strong universal modals that necessarily quantify over the actual world are excluded from the class of NRPs. As closely related lexical items may still differ in their presuppositions or in their quantificational domain requirements, which explains why some are NRPs while others are not.

Our proposal forms a novel contribution in that it is the only approach to NR so far that can deal with both the attested semantic pragmatic nature of NR and NRPs and the fact that it also triggers certain grammatical reflexes. Other approaches can only deal with one of these properties, or at closer scrutiny, neither of the two.

References

Paolo Acquaviva. The representation of operator-variable dependencies in sentential negation. *Studia Linguistica*, 48(2):91–132, 1994.

Sam Alxatib. Free choice disjunctions under only. In *Proceedings of NELS*, volume 44, pages 15–28, 2014.

Moshe Bar-Lev. Free choice, homogeneity, and innocent inclusion. *Hebrew University of Jerusalem dissertation*, 2018.

Moshe E Bar-Lev. An implicature account of homogeneity and non-maximality. *Linguistics and Philosophy*, 44:1045–1097, 2021.

Moshe E Bar-Lev. Obligatory implicatures and the relevance of contradictions. *Journal of Semantics*, 41(2):149–173, 2024.

Moshe E Bar-Lev and Danny Fox. Universal free choice and innocent inclusion. In *Semantics and Linguistic Theory*, volume 27, pages 95–115, 2017.

Moshe E Bar-Lev and Danny Fox. Free choice, simplification, and innocent inclusion. *Natural Language Semantics*, 28(3):175–223, 2020.

Renate Bartsch. "Negative transportation" gibt es nicht. *Linguistische Berichte*, 27(7), 1973.

Itai Bassi and Moshe Bar-Lev. A unified existential semantics for bare conditionals. In *Proceedings of sinn und bedeutung*, volume 21, pages 125–142, 2018.

Melanie Jane Bervoets. *On opining: modal verbs, dispositions, free choice, and negation*. University of California, Los Angeles, 2014.

Tatiana Igorevna Bondarenko. *Anatomy of an Attitude*. PhD thesis, Massachusetts Institute of Technology, 2022.

Richard Breheny, Nathan Klinedinst, Jacopo Romoli, and Yasutada Sudo. The symmetry problem: current theories and prospects. *Natural Language Semantics*, 26:85–110, 2018.

Brian Buccola, Manuel Križ, and Emmanuel Chemla. Conceptual alternatives: Competition in language and beyond. *Linguistics & Philosophy*, 2021.

Daniel Büring. Negative inversion. In *Proccedings of NELS*, 2004.

Gennaro Chierchia. *Logic in grammar: Polarity, free choice, and intervention*. OUP Oxford, 2013.

Gennaro Chierchia. 'people are fed up; don't mess with them.'non-quantificational arguments and polarity reversals. *Journal of Semantics*, 39(3):475–521, 2022.

Gennaro Chierchia, Danny Fox, and Benjamin Spector. The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. *Semantics: An international handbook of natural language meaning*, 3:2297–2332, 2012.

Chris Collins and Paul M Postal. *Classical NEG raising: An essay on the syntax of negation*, volume 67. MIT press, 2014.

Chris Collins and Paul M. Postal. Disentangling two distinct notions of NEG raising. *Semantics and Pragmatics*, 11(5), 2018a.

Chris Collins and Paul M. Postal. Dispelling the cloud of unknowing: More on the syntactic nature of neg raising. In Laurence R. Horn and Kenneth Turner, editors, *Pragmatics, truth and underspecification: Towards an atlas of meaning*, Current Issues in Semantics and Pragmatics (CRiSPI), pages 54–81. Brill, Leiden, 2018b.

Luka Crnič, Emmanuel Chemla, and Danny Fox. Scalar implicatures of embedded disjunction. *Natural Language Semantics*, 23(4):271–305, 2015.

Paul Crowley. Neg-raising and neg movement. *Natural Language Semantics*, 27(1):1–17, 2019.

Charles Fillmore. The position of embedding transformations in grammar. *Word*, 19: 208–231, 1963. doi: https://doi.org/10.1080/00437956.1963.11659796.

Danny Fox. Free choice and the theory of scalar implicatures. In *Presupposition and implicature in compositional semantics*, pages 71–120. Springer, 2007.

Jon Gajewski. Licensing strong npis. *Natural Language Semantics*, 19(2):109–148, 2011.

Jon Robert Gajewski. *Neg-raising: Polarity and presupposition*. PhD thesis, Massachusetts Institute of Technology, 2005.

Jon Robert Gajewski. Neg-raising and polarity. *Linguistics and Philosophy*, 30(3):289–328, 2007.

Elena Guerzoni. Intervention effects on npis and feature movement: Towards a unified account of intervention. *Natural Language Semantics*, 14:359–398, 2006.

Jaakko Hintikka and Jaakko Hintikka. *Semantics for propositional attitudes*. Springer, 1969.

Vincent Homer. Neg-raising and positive polarity: The view from modals. *Semantics and Pragmatics*, 8(4):1–88, 2015. doi: https://doi.org/10.3765/sp.8.4.

Laurence Horn. Remarks on neg-raising. *Syntax and Semantics*, 9:129–220, 1978.

Laurence R. Horn. *A Natural History of Negation*. University of Chicago Press, Chicago, IL, 1989.

Laurence R. Horn. The cloud of unknowing. In Jack Hoeksema and Dicky Gilbers, editors, *Black book*, pages 178–196. University of Groningen, Groningen, 2014.

Laurence R Horn. Neg-raising. In *The Oxford Handbook of Negation*. 2020.

Laurence Robert Horn. Negative transportation: Unsafe at any speed. In *Chicago linguistic society*, volume 7, pages 120–133, 1971.

Laurence Robert Horn. *On the semantic properties of logical operators in English*. PhD thesis, University of California, Los Angeles, 1972.

Sabine Iatridou and Ivy Sichel. Negative dps, a-movement, and scope diminishment. *Linguistic Inquiry*, 42(4):595–629, 2011.

Pauline Jacobson. Some people think there is neg raising, and some don't: Neg raising meets ellipsis. *Linguistic Inquiry*, 49(3):559–576, 2018.

Pauline Jacobson. Neg raising and ellipsis (and related issues) revisited. *Natural Language Semantics*, 28(2):111–140, 2020.

Paloma Jeretič. Scaleless implicatures, blocked by actuality entailments. In *Proceedings of Sinn und Bedeutung*, volume 25, pages 448–465, 2021.

Paloma Jeretic and Deniz Ozyıldız. Why neg-raising requires stativity. In *23 rd Amsterdam Colloquium*, page 134, 2022.

Paloma Jeretič. Neg-raising with belief predicates as a scaleless implicature. In *Semantics and Linguistic Theory*, pages 891–914, 2022.

Roni Katzir. Structurally-defined alternatives. *Linguistics and philosophy*, 30:669–690, 2007.

Paul Kiparsky and Carol Kiparsky. *Fact*. Linguistics Club, Indiana University, 1968.

Edward Klima. Negation in english. *The structure of language/Prentice Hall*, 1964.

Angelika Kratzer. *Modals and conditionals: New and revised perspectives*, volume 36. Oxford University Press, 2012.

Mark Liberman. On conditioning the rule of subj–aux inversion. In *Proceedings of NELS*, volume 5, pages 77–91, 1974.

Marcia C Linebarger. Negative polarity and grammatical representation. *Linguistics and philosophy*, pages 325–387, 1987.

Giorgio Magri. Another argument for embedded scalar implicatures based on oddness in downward entailing environments. *Semantics and Pragmatics*, 4(6):1–51, November 2011. doi: 10.3765/sp.4.6.

Giorgio Magri. An account for the homogeneity effect triggered by plural definites and conjunction based on double strengthening. In *Pragmatics, semantics and the case of scalar implicatures*, pages 99–145. Springer, 2014.

Zahra Mirrazi and Hedde Zeijlstra. Neg-raising without lexical excluded middle inferences: Resolving Horn clauses and other problems. In *Proceedings of SICOGG 23*, 2021.

Zahra Mirrazi and Hedde Zeijlstra. Missing words and missing worlds. In *Proceedings of Sinn und Bedeutung 27*, to appear.

Keir Moulton. *Natural selection and the syntax of clausal complementation*. PhD thesis, University of Massachusetts Amherst, 2009.

Deniz Ozyıldız. *The event structure of attitudes*. PhD thesis, University of Massachusetts Amherst, 2021.

Doris Penka. *Negative indefinites*. Number 32. Oxford University Press, 2011.

Jacopo Romoli. *Soft but strong. Neg-raising, soft triggers, and exhaustification*. PhD thesis, 2012.

Jacopo Romoli. A scalar implicature-based approach to neg-raising. *Linguistics and Philosophy*, 36(4):291–353, 2013.

Manfred Sailer. Npi licensing, intervention and discourse representation structures in hpsg. 2007.

Yosuke Sato. When vp-ellipsis and sluicing conspire against syntactic neg-raising. *Linguistic Inquiry*, pages 1–12, 2023.

Alex Silk. Update semantics for weak necessity modals. In *DEON*, pages 237–255, 2016.

Alex Silk. Weak and strong necessity modals. *Meaning, decision, and norms: Themes from the work of Allan Gibbard*, 2018.

Alex Silk. Weak and strong necessity modals: On linguistic means of expressing" a primitive concept ought". 2022.

Benjamin Spector and Yasutada Sudo. Presupposed ignorance and exhaustification: how scalar implicatures and presuppositions interact. *Linguistics and Philosophy*, 40(5): 473–517, 2017.

Frank Staniszewski. *Modality and Time in Logical Context*. PhD thesis, Massachusetts Institute of Technology, 2022.

Frank Staniszweski. A variable force analysis of positive polarity neg-raising modals. In *Proceedings of Sinn und Bedeutung*, volume 25, pages 805–822, 2021.

Robert van Rooy. Relevance only. In *Proceedings of EDILOG*. Citeseer, 2002.

Kai Von Fintel. Npi licensing, strawson entailment, and context dependency. *Journal of semantics*, 16(2):97–148, 1999.

Yimei Xiang. Neg-raising and aspect: Evidence from mandarin. *University of Pennsylvania Working Papers in Linguistics*, 19(1), 2013.

Hedde Zeijlstra. Does neg-raising involve neg-raising? *Topoi*, 37(3):417–433, 2018.