Negative Polarity as Scope Marking

Chris Barker, *New York University, Linguistics*

April 9, 2017

DRAFT of April 9, 2017. Thanks to Dylan Bumford, Daniel Büring, Simon Charlow, Chris Collins, Anastasia Giannakidou, Kiyomi Kusumoto, Alejandro Pérez Carballo, Daniel Rothschild, Philippe Schlenker, Anna Szabolcsi, Zoltán Szabó, Chris Tancredi, and my excellent referees.

Abstract. What is the communicative value of negative polarity? That is, why do so many languages maintain a stock of special indefinites (weak Negative Polarity Items) that occur only in a proper subset of the contexts in which ordinary indefinites can appear? Previous answers include: marking the validity of downward inferences; marking the invalidity of veridical inferences; or triggering strengthening implications. My starting point for exploring a new answer is the fact that an NPI must always take narrow scope with respect to its licensing context. In contrast, ordinary indefinites are notorious for taking wide scope. So whatever else NPIs may do, they at least serve as an utterly reliable signal that an indefinite is taking narrow scope. As also proposed in recent work of Kusumoto and Tancredi, I will show that NPIs are only licensed in contexts in which the wide scope construal of an indefinite fails to entail the narrow scope. In other words, weak NPIs occur only in contexts in which taking narrow scope matters for interpretation. Thus one part of the explanation for the ubiquity and robust stability of negative polarity is that it signals scope relations.

1 Why negative polarity?

Negative polarity items (NPIs) occur only in a restricted set of linguistic contexts.

(1) a. Ann didn't see [anyone].
b. *Ann saw [anyone].

(2) a. Ann doubts Bill [ever] left.
b. *Ann hopes Bill [ever] left.

For instance, *anyone* is grammatical in the presence of negation as in (1a), but it is ungrammatical in (1b) when negation is missing. Likewise, in (2b), something about the negative flavor of *doubt* provides *ever* with what it needs for the result to be grammatical.

It is a widely held belief (that I endorse) that a major part of what characterizes the difference between good contexts and bad contexts for NPIs is a semantic property having some connection to negation (though see, e.g., Szabolcsi 2004 and Collins and Postal 2014 for approaches that emphasize syntactic factors). Figuring out exactly what is the right property is not so easy, as we'll discuss more below. But even assuming that we can identify the relevant condition, there is a distinct question that begs to be asked: why in the world should NPIs care about appearing in that kind of context?

One reason this is puzzling is because the meaning of an NPI is arguably compatible with any kind of context. That is, in the interpretations of the grammatical sentences above, the NPIs appear to contribute a simple existential quantification. So (1a) can be paraphrased as 'it is not the case that there exists a person x such that Ann saw x'. The same contribution to meaning would be perfectly coherent in the ungrammatical sentence in (1b), in which case it would mean 'there is some person x such that Ann saw x' (i.e., Ann saw someone).

So the question now becomes: why would there be expressions that can only occur in some of the contexts in which they could make sense?

It is important to emphasize that negative polarity is not a marginal or unstable or rare phenomenon. There is nothing tentative or negotiable about the ungrammaticality of most sentences containing an inappropriate NPI. Furthermore, there are dozens of NPIs in English, and English is not unusual in this regard. And although negative polarity items do enter or leave the language over time, the overall level of negative polarity sensitivity in languages is robust and stable. Clearly, negative polarity is deeply woven into the fabric of natural language.

Now, we should not expect to find a functional explanation for every feature of a natural language. After all, grammaticization—turning transparently well-motivated behavior into opaque purely formal constraints—is what languages do. But negative polarity is so ubiquitous and so robust, we should prepare ourselves to discover that negative polarity provides some important functionality. That is, we should at

least ask: What use is polarity? In particular, what communicative purposes might it serve?

There are a number of answers to this question in the literature, several of which will be discussed below. One familiar answer is that NPIs signal that downward inferences are safe. Another answer is that NPIs signal the presence of a non-veridical context. Yet another popular answer is that NPIs require contexts in which they help the speaker make logically stronger or more informative statements.

I will propose a new scope-based theory of negative polarity licensing: that NPIs are licensed just when the wide-scope reading of an indefinite fails to entail the narrow scope reading. I'll argue that this theory has good empirical coverage, and that it does better in certain respects than some other theories of NPI licensing.

However, my goal is not to argue that the proposal here is the one exclusive best theory of NPI licensing. It seems likely to me that there may be several distinct overlapping factors that condition NPI distribution. Just as genes can express proteins that play roles in many different physiological systems, perhaps negative polarity can play several different communicative roles. Perhaps NPIs signal narrow scope, and also signal the safety of downward inferences, and also the presence of a non-veridical context, and they also allow a speaker to make a stronger statement. Perhaps it is the conjunctive communicative utility of all of these overlapping factors that explains the high value that languages place on negative polarity.

2 NPIs disambiguate scope

This paper explores a hypothesis that as far as I know has never yet been defended: that whatever else they do, NPIs at least serve to disambiguate scope relations.

To illustrate the main idea, consider that plain indefinites are notoriously able to take wide scope (see Szabolcsi 2010, Charlow 2014, Barker 2015 for surveys and discussion).

(3) a. If [a relative of mine] dies, I'll inherit a fortune. [Rienhart 1997]
b. Wide, *a* $>$ *if*: There is a relative x such that if x dies, I inherit.
c. Narrow, *if* $>$ *a*: If even one relative dies, I inherit.

The sentence in (3a) is ambiguous, depending on whether the indefinite *a relative of mine* takes wide scope over the conditional, as in the paraphrase in (3b), or narrow scope, as in (3c).

(4) If [any relative of mine] dies, I'll inherit a fortune.

But if the negative polarity item *any* is used instead of a plain indefinite, as in (4), the only available reading is one on which the negative polarity item takes narrow scope relative to the conditional.

This leads to the following descriptive observation:

(5) The use of an NPI rather than a plain indefinite reliably signals that it must take narrow scope with respect to some other element in the sentence.

That NPIs take narrow scope is not a new discovery by any means. In fact, most theories of NPI distribution (perhaps all) require that the NPI take narrow scope with respect to some element in its licensing context. After all, the title of Ladusaw's pathbreaking 1979 dissertation is "Negative Polarity as inherent scope relations". However, in previous theories, the narrow scope requirement is always in addition to some other independent condition on the licensing context, such as downward entailment, non-veridicality, strengthening, etc. The locus of the explanatory power in those theories is always meant to reside in the independent licensing condition, and narrow scope is just a side requirement establishing a suitable compositional relationship between an NPI and its licensing element. The goal of the present paper is to see how far we can get by viewing scope disambiguation as a central explanatory element, rather than as an ancillary condition.

3 A scope-based explanation for NPI licensing

On the assumption that NPIs signal narrow scope, I will suggest that one way to arrive at a licensing condition for NPIs will be to answer the following question: when is it useful to disambiguate the scope relations of an indefinite?

(6) [Someone] left.

In (6), because there are no scope-taking elements apart from the indefinite, there is no wide scope interpretation distinct from the narrow scope interpretation. As a result, there is no utility in marking a preference for a narrow scope interpretation.

(7) *[Anyone] left.

And in this situation, as (7) shows, the use of an NPI instead of an ordinary indefinite is not grammatical.

(8) A woman read [a book].

In (8), there are two logically distinct scope interpretations. However, the wide scope interpretation and the narrow scope interpretation entail each other. They have the same truth conditions, so there is no communicative advantage to signaling that one reading should be preferred over the other. And in this case, NPIs are not licensed:

(9) *A woman read [any books].

Generalizing the reasoning suggested by these examples, it makes sense from a functional point of view that NPIs do not occur in contexts in which the two relevant scope interpretations entail each other.

What if only one of the scope readings entails the other? In the rich-relative example (3a) from above, the narrow scope interpretation entails the wide scope interpretation, at least in any situation in which the speaker has any relatives at all. And as we have seen, NPIs are fine in the antecedent of a conditional. This suggests that when the narrow scope reading entails the wide scope reading, NPIs are fine.

NPIs are not ok when the wide scope interpretation entails the narrow:

(10) a. Every woman read [a book].
b. *Every woman read [any book].

The wide scope interpretation (a single book that is read by every woman) entails the narrow scope interpretation. And in this case, the NPI is not licensed.

Borrowing an idea from the strengthening approaches discussed below, we can understand what is going on here in terms of informativity. If the wide scope interpretation entails the narrow scope interpretation, as in (8) and (10a), then the narrow scope interpretation is less specific, that is, less informative. So from the point of view of informativity, at least in these simple examples, signaling narrow scope will only be useful when the wide scope reading does not entail the narrow.

Let this reasoning motivate the following hypothesis:

(11) Definition (informal version): NPIs are SCOPE LICENSED in a context only if the wide scope interpretation of the NPI does not entail the narrow scope interpretation.

(12) Hypothesis: NPIs must be scope licensed.

There are many ways of implementing this generalization. For the sake of making concrete, testable predictions, I'll adopt a specific implementation, though I believe that other implementations may be worth exploring.

One design choice depends on whether licensing should be thought of as a relationship between an NPI and a licensing operator (e.g., sentence negation), or as a relationship between an NPI and a surrounding context (such as the rest of the clause containing the NPI). Homer 2008, 2012 discusses these two strategies, and argues on empirical grounds that the contextual approach is superior. Without recapitulating his arguments, I will adopt the contextual approach here. However, I'm not aware of any reason why the operator approach would be incompatible with the main hypothesis of the paper, though it would complicate the definitions.

A bit of notation: a CONTEXT $C[\ ]$ is a logical form expression with a hole in it somewhere. For any logical form expression ϵ, $C[\epsilon]$ is the expression created by replacing the hole in $C[\ ]$ with ϵ. For instance, if $C[\ ] = \lambda x.\textbf{saw}\,[\ ]\,x$, and $\epsilon = \textbf{ann}$, then $C[\epsilon] = C[\textbf{ann}] = \lambda x.\textbf{saw ann}\,x$.

Let $C[\ ]$ be a context whose hole is of generalized quantifier type, i.e., type $(e \to t) \to t$, and let x be a variable of type e that does not occur free in $C[\ ]$.

(13) Definition. An NPI is SCOPE LICENSED in context $C[\ ]$ just in case $\exists x.C[\lambda\kappa.\kappa x] \not\to C[\lambda\kappa\exists x.\kappa x]$.

That is, a position is scope licensed just in case a wide scope existential fails to entail a narrow scope existential.[1]

Note that although informativity considerations helped guide us towards which hypothesis to test, the hypothesis is stated entirely in terms of scope and entailment. That means that the predictions of the hypothesis can come apart from the predictions that would follow from exclusively from informativity considerations (see section 7). Note also that the scope licensing hypothesis is a framed as a necessary condition for NPI licensing, not necessary and sufficient. There certainly are additional requirements beyond scope licensing for NPIs to be licensed, as discussed below in section 5.

[1] Because the definition does not mention a restricting predicate, there is no de re / de dicto contrast to worry about. That is, the definition could have explicitly mentioned a restricting predicate, e.g., $\exists P\exists x.Px \wedge C[\lambda\kappa.\kappa x] \wedge \neg C[\lambda\kappa\exists x.Px \wedge \kappa x]$. Note that the official definition is strictly stronger than this alternative condition, since we can always choose $P = \lambda x.\textbf{true}$. The predicate-based definition is more permissive, and would be satisfied by any attitude predicate that is opaque with respect to the de re / de dicto contrast:

(14) *Ann thinks [any dancer] left.

If we choose $P = \textbf{dancer}$, the wide scope existential gives a de re reading on which there is some dancer x that Ann thinks left, without Ann necessarily being aware that x is a dancer. That does not entail that Ann has a narrow-scope de dicto thought about someone that she believes is a dancer. Because the official definition does not depend on a restricting predicate, it correctly predicts that (14) is ungrammatical, since the bracketed position is not scope licensed.

3.1 What class of NPIs are we talking about here?

There are many different types of NPIs in the world's languages, and many different types of NPIs in English. Although I hope that the ideas proposed here will be relevant for a broad range of NPI types, it seems prudent to start with a more manageable goal. Therefore I'll restrict my official aim to accounting for so-called weak NPIs in English, whose paradigmatic exemplars include *anyone*, *any N*, and *ever*.

There certainly are many types of NPIs whose licensing conditions are different from those for weak NPIs. To mention just one, strong NPIs are briefly discussed below in section 9.

To mention a second, quite important, class, free-choice *any* has a distribution that overlaps heavily—but does not coincide—with that of weak NPIs. There are reasons to believe that NPI *any* and free-choice *any* should have a unified analysis, as in, e.g., Kadmon and Landman 1993. For instance, there are many languages in which a single lexical item functions as both an NPI and as a free-choice item, just as *any* does in English. However, there are just as many languages in which NPIs and free choice meanings correspond to distinct lexical items.

Even in English, Horn 2000 argues that the two uses are distinct, creating ambiguity:

(15) If you sleep with (just) anyone, you should use a condom.

'If there is anyone that you sleep with' is the NPI interpretation; 'if you're in the habit of sleeping with just anyone' is the free-choice reading. One way to understand this ambiguity would be to suppose that the free-choice *any* scopes underneath a silent generic operator. If so, then scope licensing holds for free-choice *any* as well as for NPI *any*, since scoping the indefinite outside of the generic operator does not entail the narrow scope reading. The difference between NPI *any* and free-choice *any*, then, would be that free-choice *any* requires a licensing context that has an appropriately modal or generic meaning.

It's also important to note that there are polarity items that are not obviously indefinite (e.g., *need*, as in *You need not respond*). Exploring what the scope licensing approach has to say about these cases will have to wait for a different occasion.

3.2 Kusumoto and Tancredi's 2013 scope-based theory

Comparison of wide and narrow scope of an existential also plays a crucial role in the licensing theory of Kusumoto and Tancredi 2013. On their analysis, NPIs take

scope immediately beneath a licensing operator (such as, e.g., negation). The at-issue content contributed by the NPI is the usual existential quantifier, scoping just under the licensor. They assume that in addition, the NPI triggers a presupposition formed by moving the existential just outside the licensing operator. So *Ann didn't read anything* asserts that it is not the case that there is something Ann read, and presupposes that there is something Ann didn't read. In upward entailing contexts (e.g., **Ann read anything*), the presupposition entails the at-issue content. Assuming that it is incoherent to assert what is presupposed, this accounts for the ungrammaticality of the NPI in upward entailing contexts.

As Kusumoto and Tancredi note, this explanation does not generalize to cases involving embedded licensing contexts. Therefore they refine their official proposal (page 13) to say that a weak NPI presupposes that the context with the indefinite taking wide scope does not entail the context with the indefinite taking narrow scope. This refined presupposition essentially amounts to an implementation of the scope licensing hypothesis.

Although Kusumoto and Tancredi's descriptive generalization is essentially the same as my scope licensing, there are significant differences in technical implementation and explanatory goals. For instance, they do not address the question of why it would be useful for an indefinite to trigger a presupposition such as the one they attribute to NPIs, which is the analog of the central question of this paper. Despite these differences, I am encouraged that independent researchers arrived at the same descriptive generalization.

4 Other proposals for the utility of NPIs

4.1 Downward Entailment

Ladusaw 1979 proposes that weak NPIs can only occur in a downward entailing environment (Fauconnier 1975, Ladusaw 1979).

(16) Definition: a context $C[\,]$ is DOWNWARD ENTAILING iff for all P and Q such that $P \subseteq Q$, $C[Q] \rightarrow C[P]$.

That is, downward-entailing environments license inferences from supersets to subsets.

(17) Ann didn't eat [fish] last night.

(18) Ann didn't eat [cod] last night.

The fact that cod is a subtype of fish, along with the fact that (17) entails (18), shows that the bracketed position is downward entailing with respect to the rest of the sentence.

(19) Hypothesis: NPIs are licensed only in downward entailing contexts.

Since the context in (17) is downward entailing, this hypothesis correctly predicts that an NPI can be licensed in that environment.

(20) Ann didn't eat [any fish] last night.

In contrast, in upward entailing contexts, NPIs are not licensed.

(21) Ann ate [fish] last night.

(22) Ann ate [cod] last night.

(23) *Ann ate [any fish] last night.

Since (21) does not entail (22), the bracketed position is not in a downward entailing context, so the downward entailment theory correctly predicts that NPIs should not be good in (23).

The scope licensing hypothesis makes the same predictions: since the existence of some fish that Ann didn't eat is not enough to guarantee that she didn't eat any fish, the scope licensing hypothesis is consistent with the facts. And in example (21), there is no narrow scope reading distinct from the wide scope reading, so the scope licensing hypothesis correctly predicts that (23) is not grammatical.

On the explanatory side of the downward entailing theory, what use are NPIs? The usual answer (e.g., Dowty 1994) is that NPIs signal that downward inferences are safe. For instance, from the fact that Ann didn't eat any fish last night, we can automatically infer that Ann didn't eat any cod.

There is experimental evidence that does not favor this explanation. Szabolcsi et al. 2008 reasoned that if part of the utility of NPIs was to signal that downward inferences are safe, then the presence of an NPI should facilitate downward inferences. They discovered that the presence of an NPI slowed processing in general, and there was no facilitation of inferences from supersets to subsets.

What does the scope licensing hypothesis predict about processing? Since there is no guarantee of downward entailment, there is no prediction of facilitation for downward inferences, in agreement with Szabolcsi et al.'s findings regarding inferences. As for explaining the general slowdown in processing, such predictions depend on detailed assumptions about how scope is computed, which this paper need not (and does not) take a stand on.

Characterizing monotonicity inferences in natural language has blossomed into an enterprise with applications in logic and in computational linguistics; see Icard and Moss 2014 for a recent survey.

We will see in the following discussions a number of non downward entailing contexts in which NPIs are nevertheless licensed. This means that if it is part of the job of an NPI to signal that downward inferences are safe, then they are at best an unreliable signal.

4.2 Victories and challenges for downward entailment

One of the great victories of the downward-entailment theory is that it correctly predicts that some operators that don't overtly involve negation nevertheless license NPIs.

(24) Ann doubts Bill ate [any fish].

(25) Wide: There is some fish that Ann doubts Bill ate.

(26) Narrow: Ann doubts that there is any fish that Bill ate.

The verb *doubt* creates a downward-entailing environment (from (24) infer that Ann doubts Bill ate a herring). So the downward entailing theory correctly predicts that the NPI is licensed. The scope hypothesis makes the same prediction, of course, since the wide scope reading of indefinites embedded under *doubt* fails to entail the narrow-scope reading.

Another great victory for the downward-entailing theory is that it makes fine-grained predictions for the restriction and the nuclear scope of quantificational determiners. For instance, *every* is downward entailing with respect to its restriction (*Every student left* entails *Every tall student left*) but not with respect to its nuclear scope (*Every student left* does not entail *Every student left quickly*). And sure enough, NPIs are grammatical in the restriction of *every* (*Every student with any sense left*) but not in the nuclear scope (**Every student ever left*). The downward entailment theory makes similarly detailed predictions for other quantificational determiners, including *no* and *some*. Although I won't pause here to prove it, the scope licensing hypothesis makes all of the same predictions.

But the downward entailing theory does not make the right prediction for the restriction of *most*.

(27) Most students left. $\not\rightarrow$ Most tall students left.

(28) Most students with [any amount of money] left.

(29) Wide: There is some amount of money x such that most students with x left.

Most is not downward entailing in its first argument. This means that the downward entailing theory predicts that NPIs should be ungrammatical in the restriction of *most*, contrary to fact. On the scope licensing theory, however, the wide scope interpretation does not entail the narrow scope interpretation, so the scope-marking theory is consistent with these facts. Gajewski 2010a argues that *most* is downward entailing after all, though not with respect to the complete sentence. Whatever the virtues of that analysis, it doesn't lend support to the claim that NPIs signal the validity of downward inferences.

For a second example of a case in which the downward entailment theory makes inaccurate predictions, the antecedent of a conditional is generally not considered to be downward entailing (Heim 1984, von Fintel 1999): *If someone talks to me, I'll have a good time* does not entail *If someone mean talks to me, I'll have a good time*. The behavior of conditionals is consistent with the scope licensing theory, as we've already seen in (4).

These are only two problematic cases for downward entailment out of several. We'll consider some additional problem cases below; see Rothschild 2006 or Giannakidou 2015 for a list of challenges for the downward entailing hypothesis.

Chemla et al. 2011 offer a general response to these kinds of problematic cases that shores up the claim the NPIs are intended to signal the validity of downward inferences. They argue from behavioral evidence that the degree to which an NPI is acceptable correlates with the degree to which subjects *assume* that a context is downward entailing. For instance, they suggest that NPIs are acceptable in the restriction of *most* because subjects tend to (incorrectly) accept downward inferences in those contexts as valid. This makes sense from a functional point of view: if part of the usefulness of an NPI is to signal that downward inferences are safe, and if downward inferences are typically or usually safe in a context, then the use of an NPI has communicative value, even if it does not provide a strict logical guarantee.

In any case, the scope hypothesis handles many situations that are troublesome for downward entailment, including *most* and conditionals, as well as other cases discussed below, without any special moves.

4.3 Veridicality

Veridicality underwrites a prominent theory of negative polarity licensing (Giannakidou 1994, Zwarts 1998, Giannakidou 1998, 2002, 2011, 2015).

(30) Definition: a context $\mathrm{C}[\,]$ is VERIDICAL iff $\mathrm{C}[\mathrm{p}] \rightarrow \mathrm{p}$.

(We'll modify this definition shortly.) A context $C[\,]$ is veridical just in case inserting a proposition-denoting expression p produces a meaning that entails p. For instance, the context *Ann knows that [ ]* is veridical, since *Ann knows that Bill left* entails that Bill left.

(31) Hypothesis: an NPI is licensed only in non-veridical contexts.

A context is non-veridical, of course, just in case it fails to be veridical. Since the context *Ann doubts that [ ]* is non-veridical, this hypothesis correctly predicts that NPIs can occur in the complement of *doubt*:

(32) Ann doubts that Bill ever ate any fish.

Considering a wider range of attitude predicates motivates a refinement of veridicality (Giannakidou 2006, 2011):

(33) Definition: a context $C[\,]$ is VERIDICAL$_x$ iff $C[p] \rightarrow (\mathrm{Dox}(x) \subseteq p)$.

That is, a context $C[\,]$ is veridical relative to an individual x just in case inserting a proposition-denoting expression p produces a meaning that entails that x believes p.

(34) Ann believes that Bill (*ever) left.

For instance, the context *Ann believes that [ ]* is veridical relative to Ann, since it entails that Ann believes Bill left. Thus the revised definition correctly predicts that weak NPIs are ungrammatical in this context, since NPIs require non-veridicality.

In order to extend the veridicality hypothesis to quantificational determiners, Giannakidou 2015:section 3.1 suggests that the restriction of a quantificational determiner can be thought of as non-veridical if a sentence containing that determiner does not allow inference to the existence of an entity satisfying the restriction. For instance, the NPI *any* is licensed in *Every student who ate any fish left* because the truth of the sentence does not guarantee that there is any such student. The idea is that existence in the referential/nominal domain is the counterpart of veridicality in the propositional/clausal domain. See Bernardi 2002 and Appendix 2 of Gajewski 2010b for discussions of some of the formal details of extending (non-relativized) veridicality to encompass existence implications.

Although this extension of the core definition is natural enough, it does not make the right prediction for some quantificational determiners, including *most* (discussed above), *only*, and *exactly* n:

(35) Only Ann saw [anything].

(36) Exactly two robots saw [anything].

Only is a well-known non-downward-entailing NPI licensor, and the consensus in the literature is that *exactly* n can license NPIs (Linebarger 1987, Israel 2011:245, Rothschild 2006, inter alia; see section 5 for discussion of the factors that degrade acceptability for larger choices of n). Yet these contexts are veridical, and entail the existence of an entity corresponding to the NPI indefinite. They are counterexamples to the claim that non-veridicality (likewise, downward entailing) is a necessary condition for licensing NPIs.

On the veridicality hypothesis, the explanation for the functional utility of negative polarity is that tracking truth is useful. Instead of signaling that certain inferences are valid, as in the downward entailment theory, NPIs signal that certain other inferences are *in*valid, namely, veridical inferences.

But in fact, by paying special attention to the scope of the NPI indefinite (the topic of this paper), we can say a bit more.

(37) Ann doubts Bill met with two students.

This sentence is ambiguous, depending on whether the indefinite takes scope over the attitude verb or not. When it takes wide scope, the interpretation entails the existence of a particular pair of student such that Ann doubts that Bill met with them. When the indefinite is replaced with an NPI (*Ann doubts Bill met with any students*), the NPI indefinite is guaranteed to take narrow scope inside the non-veridical context.

So the presence of an NPI in a non-veridical environment not only (we're supposing) marks the surrounding context as non-veridical, it also signals that the indefinite remains in the scope of the non-veridical context, and therefore does not establish a discourse referent that persists outside of the veridical context. In other words, the NPI indicates that the indefinite itself is non-veridical.

4.4 Strengthening and gramaticized implicature

Kadmon and Landman's 1993 analysis of *any* gives a clear, compelling explanation for the function of NPIs: NPIs strengthen the claim expressed by the sentence in which they occur.

More specifically, they propose that NPIs quantify over a strictly wider set of individuals than their non-NPI counterparts, and stipulate that NPIs can only be used in a context in which this domain-widening leads to a stronger claim.

(38) I don't have ANY potatoes.

Assuming that *ANY potatoes* includes borderline potatoes, this negated sentence makes a stronger claim that it would if it contained an ordinary indefinite.

Krifka 1995 and others have criticized the strengthening account on both conceptual and empirical grounds. Conceptually, testing whether widening produces a stronger claim is a global, non-compositional computation.

Empirically, it simply isn't true that NPIs always strengthen the sentences they are part of. Horn 2006:32 offers

(39) Ann hasn't recovered [yet].

(40) Bill doesn't read [much].

In fact, Krika claims (p. 215) that weak NPIs only strengthen when they participate in contrastive stress. In addition, he offers cases in which widening is not semantically possible.

(41) This sequence doesn't have [any prime numbers] in it.

There are no borderline prime numbers for widening to include. In contrast, the scope licensing hypothesis is better off here, since marking narrow scope does useful disambiguating work.

Despite these criticisms, Krifka is inspired by Kadmon and Landman's approach, and goes on to propose that weak NPIs have scalar meanings from which their behavior follows given general norms of assertion.

Here's how it work. Semantically, Krifka assumes that the relevant class of NPIs make a contribution that includes identifying a property and a set of alternative properties. For instance, the NPI *anyone* introduces the property of being a person, along with the set of alternatives consisting of all properties that correspond to subsets of people.

Krifka's composition rules are sensitive to the presence of alternatives in such a way that a sentence containing an NPI such as **Ann saw [anyone]* will denote a triple consisting of a background context, a focused value, and a set of alternatives:

$$\langle \lambda P \lambda w \exists x.Px \wedge \textbf{saw}(x)(\textbf{ann})(w),\ \textbf{person},\ \{P|P \subset \textbf{person}\}\rangle$$

Second, Krifka assumes that when a declarative sentence denotation has this kind of tripartite structure, it triggers a special pragmatic rule called SCALAR.ASSERT. SCALAR.ASSERT takes a triple such as the one just discussed and produces a

strengthened proposition consisting of the information contributed by applying the background to the focused value, conjoined with the negation of all propositions formed by applying the background continuation to each alternative that gives rise to a stronger proposition:

$$\textsc{scalar.assert}(\langle \kappa, \mathrm{P}, \mathcal{A} \rangle) = \lambda w.\kappa \mathrm{P} w \wedge \neg \exists \mathrm{Q} \in \mathcal{A} | \kappa \mathrm{Q} w \wedge (\kappa \mathrm{Q} \subset \kappa \mathrm{P}) \quad (42)$$

Here, $\kappa \mathrm{Q} \subset \kappa \mathrm{P}$ just in case the set of worlds delivered by applying the background continuation κ to Q is a proper subset of the set of worlds delivered by applying κ to P, that is, just in case Q inserted in the relevant context entails P in that context.

This is just a grammaticization of ordinary Gricean scalar implicature. Assuming cooperativity and sufficient information on the part of the speaker, assertion of a scalar can implicate the negation of all stronger scalar claims: if what you say is that Ann ate two hot dogs, you might mean to convey that Ann did not eat three hot dogs, since if she had, and you knew it, you would have said so.

Continuing the example of **Ann saw [anyone]*, SCALAR.ASSERT applied to the triple above yields a strengthened assertion that entails the conjunction of the following propositions: Ann saw a person, but she didn't see a tall person, and she didn't see a short person, she didn't see a happy person, she didn't see an unhappy person, and so on. So although she saw someone, there isn't anyone specific that she saw. These updates taken altogether are inconsistent, and so the updated context is predicted to be empty, which means that the conversation fails. The problem with asserting an NPI in an upward-entailing context, then, is that what is implicated contradicts what is said.

(43) Ann didn't see [anyone].

In contrast, when the context is downward entailing, there is no subproperty $\mathrm{Q} \subset$ **person** that leads to a stronger proposition. For instance, if Ann didn't see a tall person, that does not entail that she didn't see anyone. As a result, SCALAR.ASSERT does not add any information beyond the bare assertion that Ann didn't see anyone, and the (trivially) strengthened proposition is coherent and informative.

One of the appealing aspects of this proposal is that it says how the behavior of NPIs can follow from general conversational principles. On the other hand, it requires building neo-Gricean principles into the competence grammar in a way that blurs the line between what is said and what is meant in a controversial way. For instance, as discussed in section 6, checking for NPI licensing in embedded contexts requires assuming that SCALAR.ASSERT can apply to embedded expressions.

A second, more abstract, line of thought has been pursued in some detail by Chierchia 2006, 2013. Chierchia's approach is based on Krifka's, but with some

important differences. Chierchia adopts the idea that NPIs contribute a property whose alternatives are subproperties. He also adopts the idea that it is the communicative pressure to make the strongest claim possible that accounts for the affinity of NPIs for downward entailing contexts. However, instead of the SCALAR.ASSERT operator, there is a (silent) exhaustivity operator whose meaning is roughly equivalent to *only*. Like SCALAR.ASSERT, the presence of an exhaustivity operator is triggered by the presence of grammatically-encoded alternatives, but it is not necessary to think of an exhaustified expression as a speech act.

What, then, is the communicative function of NPIs on the grammaticized implicature accounts of Krifka and Chierchia?

Chierchia 2013:27,36,82 suggests that NPIs have a "proclivity" to evoke wider domains. But NPIs are not required to widen domains (2006:559, 2013:section 1.2, e.g., 37); they only *potentially* widen. When they do widen—especially when they undergo contrastive emphasis—the interaction between the alternatives they supply and the exhaustivity operator guarantees that they only occur in contexts in which widening results in strengthening. The "natural communicative function," then, of NPIs is that they offer "the possibility of adding emphasis" (Chierchia 2013:40).

If providing a locus for widening were the only function for NPIs, and if there is no widening in non-contrastive cases (Chierchia 2013:28), why would a speaker choose an NPI instead of an ordinary indefinite, given that the truth conditions are "perfectly interchangeable"?

The answer suggested here is that even when they do not strengthen, NPIs continue to signal narrow scope.

I want to emphasize that these two functional explanations are fully compatible: NPIs can both mark narrow scope, and sometimes strengthen. The more overlapping uses NPIs perform, the better explanation we have for why they are so ubiquitous and so stable across time.

5 Scope licensing and monotonicity

Most theories of NPI licensing predict that NPIs should only be licensed in downward-entailing contexts. In contrast, scope licensing, as we have already seen, can be satisfied in many non-monotonic contexts as well. It is reasonable to wonder whether scope licensing might be too permissive. This section defends scope licensing as a viable part of a comprehensive licensing strategy.

Here is how the argument will play out: there are many contexts in which NPIs

are licensed that are not downward entailing. Replacing downward entailing with Strawson downward entailing (as I'll explain) captures some of those exceptions, but not all. In addition, NPIs are never licensed in upward entailing contexts, even when those contexts are also downward entailing, so downward entailing alone is neither necessary nor even sufficient as a licensing condition.

Scope licensing covers the exceptions to the downward entailing hypothesis, and correctly rules out all upward entailing contexts, and so remains a good candidate for a necessary condition on weak NPIs. However, there are non-monotonic contexts that satisfy scope licensing but that don't license NPIs, so scope licensing must be supplemented by additional conditions (just as in other theories).

5.1 Strawson entailment

We've seen that *only* is not downward entailing.

(44) Only Ann read a book.

(45) Only Ann read a long book.

The truth of (44) does not guarantee the truth of (45), because (44) can be true even if the only book that Ann read was a short book. However, von Fintel notes that situations in which this downward-entailing inference does not go through are all situations in which the presuppositions of the second sentence are not satisfied. That is, (45) presupposes that Ann read a long book. If we restrict attention to situations in which the presuppositions of both sentences are satisfied, the entailment goes through: every situation in which Ann read a long book is a situation in which (44) entails (45).

Superlatives provide a second class of examples of an NPI-licensing context that is not downward entailing but that can be Strawson downward entailing.

(46) Ann is the smartest girl in my school.

(47) Ann is the smartest girl in my class.

(46) presupposes that Ann is in my school, and (47) presupposes that Ann is in my class. In any situation that satisfies both of these presuppositions, (46) entails (47), so this context is Strawson downward entailing. And sure enough, weak NPIs are good in this context (witness *Ann is the smartest girl I ever met*).

The discovery of Strawson downward entailing was a breakthrough: it just feels right. For the record, it feels right to me.

However, Strawson downward entailment is not a good candidate for a necessary condition on NPI licensing. Continuing the superlatives example, there are cases in which NPIs are licensed but that are not Strawson downward entailing.

(48) I bought the cheapest laptop with [any reviews].

(49) I bought the cheapest laptop with [any positive reviews].

This context is not downward entailing, since (48) does not entail (49). But it is not Strawson downward entailing either. It is quite possible for the unique cheapest laptop with any reviews to be a different model than the unique cheapest laptop with any positive reviews, so it is possible for (48) to be true at the same time that (49) to be false. As von Fintel 1999 comments, the Strawson downward entailingness of superlatives only "shines through" in predicative position.

Superlatives do satisfy scope licensing, by the way. For instance, the existence of some thing x such that I bought the cheapest laptop that had x does not entail the corresponding narrow scope reading.

So there are contexts that license NPIs but that are not Strawson downward entailing. Other non-SDE NPI licensing contexts include the restriction of *most* and the nuclear scope of *exactly* n.

It is possible to generalize Strawson entailment. The more general idea (Horm 2016) is that NPIs only care about what is at issue, and ignore backgrounded implications of all sorts (where presuppositions are just one kind of backgrounded implication).

This generalization provides insight into the contrast between *only* and *exactly* n. Linebarger 1980, 1987 claims that *exactly* n licenses NPIs, but only for small choices of n.

(50) Exactly 4 people have [ever] read my dissertation.

(51) Exactly 4000 people have [?ever] read my dissertation.

This contrast has become a widely-accepted part of the lore. However, the acceptability of large n improves if the context supports the inference that the number, though large, is smaller than might be expected.

(52) It's somewhat surprising to learn that in the fifteen years we've been keeping precise records, exactly 712 of our economics majors have [ever] taken [any courses at the business school].

As Horn 2016:300 puts it, "Downward monotonicity, and hence scalar structure, is crucial, but is calculated only on the asserted / at-issue component of meaning, with inert material disregarded." Thus in (50), the point of the utterance is

to highlight how few people read the dissertation. From the point of view of truth conditions, *only* n and *exactly* n both mean *at least* n *and no more than* n. The reason that *only* always licenses NPIs is that it is part of its conventional meaning to emphasize the 'no more than' part of its implications, since that is the at-issue part of its meaning. *Exactly* n is neutral between its upward and downward entailing components, and so only licenses NPIs in contexts in which the news value of the utterance comes from the downward entailing part.

Two comments are in order: first, unlike downward entailment and the scope licensing condition, which depend only on entailments, both Strawson entailment and the more general at-issue condition depend on entailments in a specific situation, and so are essentially sensitive to discourse assumptions. Second, like many pragmatic conditions, the general statement of when an NPI use will be fully aligned with what is at issue is hard to make precise, though Crnič 2014, adapting Heim 1984, proposes a specific contextual condition based on alternatives to the NPI (one course, two courses, three courses, etc.).

I will assume that we need to supplement scope licensing with some version of a discourse-sensitive at-issue scalarity restriction, just as for other theories of NPI licensing (see, e.g., remarks in Chierchia 2013:section 4.3).

5.2 Upward entailing as an anti-licensor?

In addition to a scalarity condition, fans of downward entailment and of Strawson downward entailment must add a stipulation that NPIs are prohibited in upward entailing contexts. The reason is that it is possible for a context to be simultaneously downward entailing and upward entailing, as noted by Homer 2008, and Gajewski and Hsieh 2016 and works cited there. (See also unpublished slides of von Fintel circa 2008.)

The discussion centers around definite determiners, including *the* and *both*. The facts suggest that NPIs are licensed in (Strawson) downward entailing environments as long as the contexts are not also upward entailing. The discussions are complicated by presupposition, genericity, plurality, and other factors, and there isn't room to discuss these case thoroughly here.

But there are simpler examples that can serve to make the point. As far as I know, this hasn't been noted in the literature yet, but there are at least two general ways a context can be both downward and upward entailing. The first is if the context is true no matter what value is inserted into the hole position.

(53) *Zero or more students read [anything]. [Krifka 1995]

Krifka 1995:217 notes that such contexts do not license NPIs, though he does not remark that they are downward entailing.

Likewise, if a context is false no matter how the hole is filled, the context is both upward and downward entailing.

(54) *At most zero students and at least three students read [anything].

As (54) shows, these contexts also do not license NPIs. Incidentally, the always-false examples are counterexamples to the non-veridicality hypothesis as well.

For theories that predict that (Strawson) downward entailing contexts can license NPIs, it appears that it is necessary to stipulate that NPIs are prohibited in upward entailing contexts. Alternatively, in view of the many counterexamples to (Strawson) downward entailing, another strategy would be to abandon (Strawson) downward entailing as a licensing condition, and replace it with an anti-upward entailing condition. Authors who have suggested this strategy include Progovac 1992, 1994:135; Neale 1999, Nishiguchi 2003, Rothschild 2006, and Collins and Postal 2014:72.

To what extent is scope licensing, either as implemented by Kusumoto and Tancredi 2013, or as implemented here, essentially equivalent to an anti-upward entailment strategy? After all, upward entailing contexts can never satisfy the scope licensing condition. To see why, assume $\exists x.C[\lambda\kappa.\kappa x]$. It follows immediately from the definition of upward entailing that $C[\lambda\kappa\exists x.\kappa x]$, since for any x, $\lambda\kappa.\kappa x \subseteq \lambda\kappa\exists x.\kappa x$.[2]

But in any case, we should strongly prefer scope licensing over anti-upward entailment as our generalization, since scope licensing has an explanatory story, based on whether narrow scope matters for interpretation.

Let's take stock. Downward entailing is not a necessary requirement for NPI licensing, as shown by conditionals, superlatives, *most*, *only*, and *exactly* n. Of

[2] It turns out that scope licensing is not equivalent to non-upward entailing.

$$(\lambda\mathcal{P}.\mathcal{P}(\lambda z.z = \textbf{ann}) \wedge \mathcal{P}(\lambda z.z = \textbf{bill}) \wedge \neg\exists z.\mathcal{P}(\lambda x.\textbf{spoke-to}\ x\ z))\ [\,]$$

This context is not upward entailing (nor is it downward entailing), yet it does not satisfy the scope licensing condition. It is not clear that there are natural language expressions whose meaning corresponds to this kind of context. However, Steedman (e.g., 2012:49) claims that *Every boy admires, and every girl detests, some saxophonist* can be interpreted in a way in which the indefinite gets independently evaluated, once for each of the conjuncts of the right node raising construction. So we might try the following context:

Ann is identical to, and Bill is identical to, and no one spoke to, [ ].

this list of counterexamples, Strawson downward entailment covers *only* and some superlatives, and arguably conditionals, but does not cover the rest, so it is also not a necessary condition. Nor is non-veridicality (superlatives, *most*, *only*, *exactly* n). Furthermore, any licensing theory based on downward entailment or Strawson downward entailment must stipulate an independent prohibition against upward entailing contexts.

To be sure, there are cases in which scope licensing is satisfied, yet NPIs are not licensed, so scope licensing must be supplemented with additional constraints, just as in other semantic theories of NPI licensing. For instance, in addition to scalar contrasts such as the one discussed above for *exactly* n, there are also syntactic licensing constraints that are sensitive to linear order (see Barker and Shan chapter 8 for a discussion and an explanation).

Nevertheless, as far as I know, scope licensing holds for every example in which a weak NPI is licensed, so it remains a candidate for a necessary condition.

6 Double licensing

What happens when NPIs occur in the scope of more than one potential licensor? Schmerling 1971 shows that adding a second licensor can degrade acceptability:

(55) No one did anything to help.

(56) Ann didn't do anything to help.

(57) *No one didn't do anything to help.

Homer 2012 argues that double licensing fails when the two licensors are too close to each other in some syntactically defined sense. Without exploring his theory of licensing domains, it suffices for present purposes to note that there are many examples in which double licensing is acceptable.

(58) No one claimed that Ann didn't do anything to help.

(59) If he doesn't know anything about logic, I'll teach him. [Hoeksema 1986]

(60) She rarely doesn't eat anything for lunch. [Dowty 1994]

(61) It is not possible that John didn't understand anything. [Homer 2012]

(62) It is not the case that Ann didn't see [anyone].

(63) Ann doubts that no one said [anything].

(64) Ann never claimed that Bill didn't kill [anyone].

Each of these examples contains two elements, each of which are sufficient to license the negative polarity item.

Double licensing is compatible with the scope licensing hypothesis. Scope licensing requires the NPI to occur in a context in which the wide scope interpretation fails to entail the narrow scope interpretation. That condition is met in each of the examples here by choosing an embedded context (in most of the examples, the clause following *that*).

We can be sure that the embedded context is the operative one for licensing, because the sentences are not ambiguous: in each case, the NPI takes narrow scope with respect to the lower licensing context.

This means that the disambiguation function of the NPI is in no way diminished in the double-licensing construction: it remains useful to signal that the indefinite takes narrowest scope, and not intermediate or wide scope.

From the point of view of the downward entailing theory, the existence of double licensing is puzzling. After all, for most double-licensing situations (e.g., (62)), the NPI is in an *upward* entailing context with respect to the entire sentence. We can suppose that the NPI is licensed with respect to the embedded clause, just as we did for the scope licensing account. But if we do, we lose the functional motivation for restricting NPIs to downward entailing contexts. In particular, it is not valid to infer from (62) that it is not the case that Ann didn't see anyone tall. So if the job of an NPI is to signal that downward inferences are safe, this is a clear case in which the presence of the NPI gives an untrustworthy signal.

Similar remarks apply to veridicality-based theories.

As Krifka 1995 points out, making an implicature approach work in general requires assuming that his SCALAR.ASSERT operator can apply to properly embedded expressions, i.e., not at the utterance level. As he puts it (page 245), "we must develop a framework in which illocutionary operators are part of the semantic recursion", i.e., that embedded constituents can count as illocutionary acts. If we can accept this radical assumption, double licensing can be explained as cases of embedded scalar assertion.

Likewise, on Chierchia's exhaustivity approach, double licensing works out fine, as long as the silent exhaustivity operator applies to the embedded clause rather than to the clause as a whole. The embedded clause will be strengthened, but this means that the larger utterance will be weakened. From the point of view of the sentence as a whole, the presence of the NPI signals a weaker statement, not a stronger statement. Whatever the virtues of this sort of approach, it attenuates the connection between the licensing behavior and the supposed functional motivation for NPIs. Why is it useful to strengthen the meaning of a subexpression, especially

if that weakens the strength of the utterance as a whole?

In contrast, the utility of marking narrow scope continues to be just as relevant and useful in double-licensing conditions as it is in single-licensing conditions.

7 Unscopable indefinites

What happens when an indefinite is unable to take wide scope for reasons that are independent of polarity sensitivity?

The scope licensing hypothesis was motivated by the expectation that NPIs should prefer environments in which marking narrow scope is informative. If NPI licensing were perfectly aligned with informativity, we might expect NPIs to be forbidden in environments in which indefinites are unable to take wide scope anyway. In this section and the next, we will see that NPIs can sometimes be ok in such contexts.

For a first case to consider, if an indefinite contains a bound pronoun, it is not logically coherent to give that indefinite wide scope over the binder of that pronoun.

(65) No one$_i$ brought [any of her$_i$ friends] to the wedding.

(66) *Everyone$_i$ brought [any of her$_i$ friends] to the wedding.

In (65), if the pronoun *her* is bound by the downward-entailing quantifier *no one* as indicated, it is not possible for the bracketed indefinite to take wide scope over the quantifier.

However, the scope licensing condition depends only on the semantic nature of the context, and is not sensitive to the internal structure of the NPI. Because of this, it correctly rules in (65), while still ruling out (66).

For a second case, indefinites in predicative positions are widely believed to have fixed scope.

(67) There is [someone] at the door. [pivot of existential *there*]

(68) I have [a sister]. [object of possessive *have*]

The assumption is that indefinites in the bracketed positions are never able to take wide scope.

(69) If there is [a person] at the door, I'll be surprised.

The prediction is that (69) is unambiguous, and only has an interpretation on which the indefinite takes narrow scope with respect to the conditional (compare with

(3a)). This is a safe assumption to make for present purposes, since if it turns out that predicate positions don't impose this limitation, they don't pose any potential problem for scope licensing.

If indefinites in predicate position can't take wide scope, then marking narrow scope does not provide any disambiguation. Nevertheless, if a suitable licensing environment is present, NPIs are good in these positions:

(70) There aren't [any students] in the lounge.

(71) If Ann has [any children], I'll be surprised.

Once again, this is perfectly consistent with the scope licensing condition. Scope licensing only requires that the wide scope reading fails to entail the narrow scope reading; it does not require that the sentence in question be able to express the wide scope reading. In particular, since the existence of a student who isn't in the lounge does not entail (70), this example is consistent with scope licensing.[3]

What these two cases show is that NPIs can be licensed even there is no possibility of the indefinite taking wide scope, that is, even in some cases in which marking narrow scope is not informative.

It is important to emphasize three things. First, NPIs in predicative position still take narrow scope, so the correlation of NPIs with narrow scope remains without exception. Second, neither of these cases constitute counterexamples to the scope licensing condition.

Third, if the indefinites in these positions can't take wide scope, that means that in these specific situations, marking narrow scope is not informative. The scope licensing hypothesis is in good company: as near as I can see, all explanations for the functional utility of negative polarity must face some situations in which NPIs are licensed, but the motivating function is obscure or missing. For instance, in double-licensing contexts (see section 6), neither downward inferences nor non-veridical inferences are valid, yet NPIs are licensed. Likewise, there are many situations in which NPIs are licensed but in which widening and strengthening does not occur, as discussed by Krifka 1995 and by Chierchia 2016 chapter 1. But we shouldn't be terribly surprised by this state of affairs—having constraints that apply

[3] A technical detail: there are many theoretical approaches to predicative constructions (see McNally 2016 for a survey). If you believe (along with, e.g., Landman 2004) that the indefinite in a predicative position denotes a property directly, without ever denoting a generalized quantifier, then it is necessary to adjust the scope licensing condition to deal in properties. Here is a natural extention: an NPI in a context $C[\,]$ in which the hole has type $\mathtt{e} \to \mathtt{t}$ will be scope licensed just in case $(\exists x.C[\textsc{ident}(x)]) \not\to C[\textsc{be}(\exists x \lambda\kappa.\kappa x)]$, where IDENT and BE are the familiar typeshifters from Partee 1987 (simplifications are possible).

in situations that extend beyond the basic motivating environment is the normal situation whenever a functional constraint becomes a grammatical constraint.

In the next section, we'll see a third kind of situation in which indefinites might not have a wide scope interpretation.

8 Questions

It is well known that questions license weak NPIs.

(72) Did Ann [ever] eat [any tofu]?

(73) Who [ever] ate [any tofu] unless they had to?

What does scope licensing have to say about questions?

8.1 Embedded interrogatives

The situation with embedded questions is straightforward.

(74) Ann knows if Bill read [a book].

(75) Wide: there is a book x such that Ann knows if Bill read x.

(76) Narrow: Ann knows if there is a book that Bill read.

(77) Ann knows if Bill read [any books].

The wide scope reading does not entail the narrow scope reading: knowing that Bill did not read one particular book can be enough to verify the wide scope reading, but it is certainly not enough to know whether Bill read any books at all. So the scope licensing hypothesis makes good predictions with respect to embedded questions.

8.2 Unembedded interrogatives

The situation for unembedded questions is more complicated. For starters, it is not self-evident what ought to count as entailment when questions are involved. For present purposes, I'll use the well-known notion of question entailment due to Groenendijk and Stokhof 1984. (This notion of question entailment has a similar treatment in more recent work, such as Inquisitive Semantics.) Their approach is based on treating the meaning of a question as a propositional concept.

(78) $[\![$*Which people left?*$]\!] =$
$\lambda ij.\{x|\textbf{person}_i\ x \wedge \textbf{left}_i\ x\} = \{x|\textbf{person}_j\ x \wedge \textbf{left}_j\ x\}$

(79) $[\![$*Which tall people left?*$]\!] =$
$\lambda ij.\{x|\textbf{tall}_i\ x \wedge \textbf{person}_i\ x \wedge \textbf{left}_i\ x\} = \{x|\textbf{tall}_j\ x \wedge \textbf{person}_j\ x \wedge \textbf{left}_j\ x\}$

When evaluated at any specific index w, the extension of a question meaning will be a proposition, construed as a the set of worlds in which the same people left as in w. Since any worlds that agree completely on which people left necessarily agree on which tall people left, $[\![$*(78)*$]\!]^w \subseteq [\![$*(79)*$]\!]^w$ for all w, so it is natural to say that (78) entails (79). One way of putting it is that any complete answer to the question of which people left will necessarily constitute a complete answer to the question of which tall people left.

Since we have a context that guarantees entailment from a set to a subset, this feels remarkably like downward entailment. Could the fact that questions are downward entailing in this sense be related to why they robustly license NPIs? Unfortunately, as soon as the wh-phrases in (78) and (79) are replaced with indefinites, the entailment relation disappears.

(80) $[\![$*Did a person leave?*$]\!] =$
$\lambda ij.(\exists x.\textbf{person}_i\ x \wedge \textbf{left}_i\ x) = (\exists x.\textbf{person}_j\ x \wedge \textbf{left}_j\ x)$

(81) $[\![$*Did a tall person leave?*$]\!] =$
$\lambda ij.(\exists x.\textbf{tall}_i\ x \wedge \textbf{person}_i\ x \wedge \textbf{left}_i\ x) = (\exists x.\textbf{tall}_j\ x \wedge \textbf{person}_j\ x \wedge \textbf{left}_j\ x)$

The set of worlds that agree with w on whether a person left is not a subset of the set of worlds that agree with w on whether a tall person left. In particular, if w is a world in which the only person who left was tall, a world in which only short people left can easily be a member of the first set (it agrees with w on whether someone left) but not the second (it disagrees with w on whether someone tall left). So the downward entailment condition is not satisfied, and we are left without an explanation for why the NPI in *Did anyone leave?* is grammatical.

If simple downward-entailment doesn't work, what about scope licensing? In order to apply the scope licensing criterion, we would have to compare a wide-scope reading with the narrow-scope reading. Groenendijk and Stokhof 1984 offer a general method for allowing a variety of generalized quantifiers, including indefinites, to take scope over a question.

(82) Did someone leave?

(83) {Did Ann leave?, Did Bill leave?, Did Carl leave?}

On their method, the wide-scope reading of (82) denotes a set of ordinary question meanings, perhaps as in (83). Groenendijk and Stokhof call such interpretations

'choice readings'. On the choice reading, (82) conveys something like 'for some particular person x, did x leave?'. On the choice reading, one way to give a complete answer to (82) would be to resolve any one of the constituent questions, in which case *Bill didn't leave* could be a complete answer.

The formal status of choice readings is simpler in Inquisitive Semantics (e.g., Ciardelli et al. 2016). Instead of involving sets of ordinary question meanings, choice readings have the same type as any (potentially inquisitive) proposition, where a proposition is a set of downward-closed sets of worlds. Without pausing to define the logic underlying Inquisitive Semantics, the usual narrow scope reading of the indefinite corresponds to the formula $?\exists\mathrm{x}.\mathrm{L}\mathrm{x}$, and the wide-scope interpretation, i.e., the choice reading, corresponds to $\exists\mathrm{x}.?\mathrm{L}\mathrm{x}$ (see, e.g., Ciardelli et al. 2009:61). Just as with the earlier theory, the issue raised by the choice reading can be resolved by stating whether any of the individuals quantified over by the existential has property L or not.

In any case, the choice reading does not entail the narrow scope reading. That is, declaring that Bill didn't leave doesn't settle the issue raised by the narrow-scope question *Did anyone leave*?. Thus the scope licensing condition correctly predicts that questions should license NPIs.

Of course, it is far from clear that natural language questions can express choice readings. In particular, Krifka 2001 argues that if they do exist, they are difficult to access at best. Guerzoni and Sharvit 2014 and Szabolcsi 2016 make good progress addressing the cross linguistic possibility of disjoining questions, concentrating mostly on embedded interrogatives, but I take it that the debate continues.

However, it is not necessary for questions to be able to express a choice reading in order for us to decide that the scope licensing condition has been satisfied. As we saw above in section 7, the scope licensing conditions makes good predictions even in situations in which the indefinite in question is not able to take wide scope. There is no semantic reason why a question could not give rise to a wide scope reading, as we know from the behavior of embedded questions. Apparently, choice readings just aren't appropriate speech acts when using an unembedded interrogative. For the purposes of licensing NPIs, it is enough to know what the wide scope interpretation would be, and to show that it does not entail the narrow scope interpretation.

So questions, both unembedded and embedded, are consistent with the scope licensing hypothesis.

9 What about strong NPIs?

Strong NPIs include punctual *until*, *in weeks*, and additive *either*, and are supposed to be licensed only in anti-additive contexts (Zwarts 1998, van der Wouden 1997:32).

(84) Definition. A context $C[\ ]$ is ANTI-ADDITIVE iff for all properties P and Q, $C[P \vee Q] \leftrightarrow C[P] \wedge C[Q]$.

Expressed in scope talk, an anti-additive context guarantees that a narrow scope disjunction is equivalent to a wide scope conjunction.

For instance, *No one [ ]* is anti-additive, since *No one [left or sang]* is equivalent to *No one [left] and no one [sang]*. In contrast, despite the fact that the context *At most one person [ ]* is downward entailing and does license weak NPIs (e.g., *At most one person ever said anything kind*), it is not anti-additive, since *At most one person [left or sang]* is not equivalent to *At most one person [left] and at most one person [sang]*. (Only the second sentence can be true in a situation in which the set of people who either left or sang is two.)

The anti-additive hypothesis correctly predicts the following licensing facts:

(85) No one has responded in weeks.

(86) No one responded until Thursday.

(87) No one left. No one sang, either.

(88) *At most one person has responded in weeks.

(89) *At most one person responded until Thursday.

(90) At most one person left. *At most one person sang, either.

For theories that guarantee that NPI licensing contexts are downward entailing, the existence of strong NPI seems natural. In fact, a context is anti-additive only if it is also downward entailing, since left-to-right half of the equivalence (i.e., $C[P \vee Q] \rightarrow C[P] \wedge C[Q]$) is the same thing as being downward entailing, since $P \subseteq (P \vee Q)$ for all properties P and Q, and for any $P \subseteq Q$, $P \subseteq (P \vee Q)$. So the licensing requirement for strong NPIs is just the licensing requirement for weak NPIs strengthened from a material conditional to a biconditional.

Intriguingly from the point of view of the scope-based explanation for negative polarity advocated here, the definition of anti-additivity is already expressed in the form of a scope-taking generalization. Here is a candidate for a condition that approximates anti-additivity, and that might come close to capturing the distribution of strong NPIs:

(91) $\forall x.C[\lambda\kappa.\kappa x] \leftrightarrow C[\lambda\kappa\exists x.\kappa x]$

This condition holds for contexts for which a narrow scope indefinite is equivalent to a wide scope universal. This condition just exchanges disjunction and conjunction in the usual definition for their quantificational correlates, namely, existential quantification and universal quantification.

We should not take it as settled that anti-additivity is the right licensing condition for strong NPIs. Chapter 10 of Collins and Postal argues at length that anti-additivity is not the right criterion. Likewise, Gajewski 2010b has a different account as well. Just as many NPIs stake out the endpoint of a scale, and therefore strengthen the statements they occur in, strong-NPI contexts are *contexts* that are at the end of a scale, and therefore make the resulting statement stronger than it would otherwise be. For instance, *no one [ ]* is the endpoint of the scale of the number of people involved, and is anti-additive, in contrast with *at most one person [ ]*, which is not quite at the endpoint of the same scale, and is not anti-additive. This approach to understanding strong NPIs is also compatible with with the scope licensing hypothesis.

10 Conclusions

NPIs always take narrow scope with respect to some element in their licensing context. This means that no matter what other communicative purposes NPIs may serve, at the very least they provide an utterly reliable signal about the scope interpretation of the sentences they occur in.

Furthermore, I have argued that weak NPIs only occur in contexts in which a wide scope reading does not entail the corresponding narrow scope reading. Functionally, the explanation for this restriction is that it would be misleading to mark narrow scope if the wide scope interpretation were at least as informative.

This scope licensing hypothesis is a remarkably good candidate for a necessary condition on NPI licensing. It covers cases in which (Strawson) downward entailing theories, veridicality, and scalar implicature theories make the wrong predictions.

Scope licensing is certainly not both necessary and sufficient, however. In particular, some non-monotonic contexts require in addition that the NPI stake out a highly informative position on some scale provided by the discourse situation.

The scope licensing condition does not require that the context in question be able to be used to express a wide scope interpretation. NPIs containing bound pronouns, NPIs in predicative position, and NPIs in unembedded questions all show that NPIs can be used in positions where an ordinary indefinite could not

take wide scope. For the purposes of licensing NPIs, however, it is enough to know what the wide scope reading would be, and to show that it does not entail the narrow scope reading.

The scope licensing hypothesis is compatible with a range of independent theories of the communicative value of NPIs: NPIs may also signal the safety of downward inferences, or the presence of non-veridicality, or the potential for widening and strengthening. The more useful NPIs turn out to be, the better explanation we have for why they are so ubiquitous and so robust.

Whatever various communicative purposes negative polarity serve, NPIs certainly mark narrow scope, and they only occur in contexts in which narrow scope matters for interpretation.

11 References

Barker, Chris. 2015. Scope. In Shalom Lappin and Chris Fox (eds). *The Handbook of Contemporary Semantics*, 2d edition. Wiley-Blackwell. 47–87.

Barker, Chris and Chung-chieh Shan. 2014. *Continuations and Natural Language*. Oxford University Press.

Bernardi, Raffaella. 2002. *Reasoning with polarity in categorial type logic*. PhD, University of Utrecht.

Charlow, Simon. 2014. *On the semantics of exceptional scope*. PhD Dissertation, NYU.

Chemla, Emmanuel, Vincent Homer, and Daniel Rothschild. 2011. Modularity and intuitions in formal semantics: the case of polarity items *Linguistics and Philosophy* **34**:537–570.

Chierchia, Gennaro. 2006. Implicatures of Domain Widening and the "Logicality" of Language. *Linguistic Inquiry* **37.4**: 535–590.

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

Ciardelli, Ivano, Irma Cornelisse, Jeroen Groenendijk, Floris Roelofsen. 2009. Computing compliance. In X. He, J. Horty, E. Pacuit (eds). *International Workshop on Logic, Rationality and Interaction*. 55–65.

Ciardelli, Ivano, Jeroen Groenendijk, Floris Roelofsen. 2016. *Introduction to Inquisitive Semantics*. Manuscript.

Collins, Christopher and Paul Postal. 2014. *Classical NEG Raising: an essay on the syntax of negation*. MIT Press.

Crnič, Luka. 2011. *Getting* even. MIT PhD dissertation .

Crnič, Luka. 2014. Against a Dogma on NPI licensing. In L. Crnič and U. Sauerland (eds). *The Art and Craft of Semantics: A Festschrift for Irene Heim*, Volume 1: 117–145.

Collins, Chris, and Paul M. Postal. 2014. *Classical NEG raising: an essay on the syntax of negation*. MIT Press.

Dowty, David. 1994. The role of negative polarity and concord marking in natural language reasoning. *Semantics and Linguistic Theory IV*: 114–144.

Fauconnier, Giles. 1975. Polarity and the scale principle. Chicago Linguistics Society 11. 188–199.

von Fintel, Kai. 1999. NPI licensing, Strawson entailment, and context dependency. *Journal of Semantics* 16.2: 97–148.

Gajewski, Jon. 2010a. Superlatives, NPIs and most. *Journal of Semantics* **27.1**: 125–137.

Gajewski, Jon. 2010b. Licensing strong NPIs. *Natural Language Semantics* **19.2**:109–148.

Gajewski, Jon and I-ta Chris Hsieh. Comments on negative polarity items in definite descriptions. In: Luka Crnic and Uli Sauerland (eds). *The Art and Craft of Semantics: A Festschrift for Irene Heim*, vol. 1, MITWPL 70, pp. 181–198.

Giannakidou, A. 1994. The Semantic licensing of NPIs and the Modern Greek subjunctive. In Ale de Boer, Helen de Hoop, Henriette de Swart (eds). *Language and Cognition 4, Yearbook of the Research Group for Theoretical and Experimental Linguistics*. University of Groningen: 55–68.

Giannakidou, Anastasia. 1998. *Polarity Sensitivity as (Non)veridical Dependency*. John Benjamins, Amsterdam-Philadelphia.

Giannakidou, A. 2002. Licensing and sensitivity in polarity items: from downward entailment to non-veridicality. In Maria Andronis, Anne Pycha and Keiko Yoshimura (eds). *CLS 38: Papers from the 38th Annual Meeting of the Chicago Linguistic Society, Parasession on Polarity and Negation*.

Giannakidou, Anastasia. 2011. Negative and positive polarity items. In C. Maienborn, K. von Heusinger, and P. Portner (eds) *Handbook of Semantics*. Berlin/New York: de Gruyter. 1660–1712.

Giannakidou, Anastasia. 2015. Polarity. Oxford Encyclopedia.

Groenendijk, Jeroen and Martin Stokhof. 1984. *Studies on the Semantics of Questions and the Pragmatics of Answers*. PhD dissertation, University of Amsterdam.

Guerzoni, Elena, and Yael Sharvit. 2007. A question of strength: on NPIs in interrogative clauses. *Linguistics and Philosophy*. **30**: 361–391.

Guerzoni, Elena, and Yael Sharvit. 2014. Whether or not anything but not whether anything or not. In Luka Crnič and Uli Sauerland (eds). *The Art and Craft of Semantics: A Festschrift for Irene Heim*. Vol. 1, MITWPL 70: 199–224.

Heim, Irene. 1984. A Note on Negative Polarity and Downward Entailingness. In NELS 14, 98–107. GLSA, UMass, Amherst.

Heim, Irene and Angelika Krazer. 1998. *Semantics in Generative Grammar*. Blackwell.

Hoeksema, Jack. 1986. Monotonicity phenomena in natural language. *Linguistic Analysis* **16**: 25–40.

Horn, Larry. 2000. *Any* and *(-)ever*: Free choice and free relatives. In Adam Wyner (ed), *Proceedings of the 15th Annual Conference of the Israeli Association for Theoretical Linguistics*: 71-111.

Homer, Vincent. 2012. Domains of Polarity Items. Unpublished manuscript.

Homer, Vincent. 2008. Presuppositions Can Be Disruptors Too: A Case against Strawson-Entailment. In Natasha Abner and Jason Bishop (eds). *Proceedings of the 27th West Coast Conference on Formal Linguistics*. Somerville, MA: Cascadilla Proceedings Project. 220–228.

Horn, Larry. 2006. The border wars: A neo-Gricean perspective. In K. von Heusinger and K. Turner et al. (eds). *Where semantics meets pragmatics*. 21–48. Elsevier.

Icard, Thomas, and Lawrence Moss. 2014. Recent progress on monotonicity. In *LiLT (Linguistic Issues in Language Technology) 9*.

Israel, Michael. 2011. *The Grammar of Polarity: Pragmatics, Sensitivity, and the Logic of Scales*. Cambridge University Press.

Kadmon, Nirit and Fred Landman. 1993. Any. *Linguistics and Philosophy* 16.4: 354–422.

Krifka, Manfred. 1995. The semantics and pragmatics of polarity items. *Linguistic Analysis*. **25**: 209–257.

Krifka, Manfred. 2001. Quantifying into question acts. *Natural Language Semantics* **9.1**: 1–40.

Kusumoto, Kiyomi, and Christopher Tancredi. 2013. Weak NPIs as Double Scope Quantifiers. Reports of the Keio Institute of Cultural and Linguistic Studies 44.

Ladusaw, William A. 1979. *Polarity Sensitivity as Inherent Scope Relations*. PhD thesis, University of Texas at Austin. Published 1980 by Garland.

Landman, Fred. 2004. *Explorations in Semantics: Indefinites and the Type of Sets*. Wiley-Blackwell.

Linebarger, Marcia. 1980. Polarity *any* as an existential quantifier. In Kreiman and

Ojeda (eds). *Proceedings of the Sixteenth Regional Meeting of the Chicago Linguistic Society*.

Linebarger, Marcia. 1987. Negative polarity and grammatical representation. *Linguistics and Philosophy* **10.3**: 325–387.

McNally, Louise. 2016. Existential. Mark Aronoff (ed), *Oxford Bibliographies in Linguistics*. New York: Oxford University Press. Entry Launch Date: 28/4/2016. DOI: 10.1093/OBO/9780199772810-0070.

Nicolae, Andreea C. 2015. Questions with NPIs. *Natural Language Semantics* **23.1**: 21–76.

Nishiguchi, Sumiyo. 2003. Non-monotonic negativity. In *Proceedings of the 17th Pacific Asia Conference on Language, Information, and Computation (PACLIC 17)*. Singapore: COLIPS. 204–215.

Neale, Stephen. 1999. Persistence, Polarity, and Plurality. In U. Egli and K. von Heusinger (eds). *Reference and Anaphoric Relations*, Dordrecht: Kluwer. 147–153.

Partee, Barbara. 1987. Noun phrase interpretation and type shifting principles. In J. Groenendijk, D. de Jongh and M. Stokhof (eds). *Studies in Discourse Representation Theory and the theory of generalized quantifiers*. 115–141.

Progovac, Ljiljana. 1992. Negative Polarity: a Semantico-Syntactic Approach. Lingua 86, 271- 299.

Progovac, Ljiljana. 1994. *Negative polarity and positive polarity: a binding approach.* Cambridge University Press.

Reinhart, Tanya. 1997. Quantifier Scope: how labor is divided between QR and choice functions. *Linguistics and Philosophy* **20**:335–397.

Rothschild, Daniel. 2006. Non-monotonic NPI-licensing, definite descriptions, and grammaticalized implicatures. In Proceedings of SALT 16.

Schmerling, Sue. 1971. A note on negative polarity. *Papers in Linguistics* **4.1**: 200–206.

Schwarzschild, Roger. 2002. Singleton Indefinites. *Journal of Semantics* 19: 289–314.

Steeman, Mark. 2012. *Taking scope*. MIT Press.

Szabolcsi, Anna. 2004. Positive polarity — negative polarity. *Natural Language and Linguistic Theory* **22.2**:409–452.

Szabolcsi, Anna. 2010. *Quantification*. Cambridge University Press.

Szabolcsi, Anna. 2016. Direct vs. indirect disjunction of wh-complements, as diagnosed by subordinating complementizers. Manuscript.

Szabolcsi, Anna, L. Bott, and Brian McElree. 2008. The Effect of Negative Polarity Items on Inference Verification. *Journal of Semantics*. **25.4**:411–450.

van der Wouden, Ton. 1997. *Negative Contexts: Collocation, Polarity and Multiple Negation*. Routledge.

Zwarts, Frans. 1998. Three types of polarity. In F. Hamm and E. W. Hinrichs (eds). *Plurality and Quantification*. Kluwer: 177–238.