To appear in Aitha, Akshay, Kutay Serova and Madeline Snigaroff (eds.). 2023. Proceedings of CLS 59. Chicago: Chicago Linguistic Society

Zeroing in on Exclusively Exclusive Content

Christopher Kennedy
University of Chicago

"If the force of the exclusive proposition is to exclude everything other than what is named in or by the subject-term from 'sharing in the predicate,' that is no reason for reading in an implication that something named by the subject-term does 'share in the predicate;' and we certainly cannot exclude from our logic predicables that are not true of anything." – Geach 1962, p. 208

1 Introduction

Bylinina & Nouwen (2018) (BN) discuss two puzzles involving the numeral *zero* and plural count nouns in sentences like (1), one empirical and one theoretical.

(1) Zero signals have been detected by our instruments.

On the empirical side, *zero* is similar to the negative indefinite quantifier *no* in creating downward entailing contexts in both its restriction and its scope:

(2) a. Zero signals have been detected by our instruments. →
Zero (weak) signals have been detected (this week) by our (most sensitive) instruments.

b. No signals have been detected by our instruments. →
No (weak) signals have been detected (this week) by our (most sensitive) instruments.

But as pointed out by Zeijlstra (2007), *zero* is unlike *no* in failing to license negative polarity items (NPIs):[1]

---

*Big thanks to Brian Buccola, Lucas Fagan, Nick Fleisher, Itamar Francez, Larry Horn, Julian Kennedy, Rick Nouwen, Paolo Santorio, Mandy Simons, Clemens Steiner-Mayr, and three anonymous reviewers for valuable feedback on the ideas advanced in this paper. I'm sorry I couldn't address all of the issues they raised, and none of them should be held responsible for the kooky ideas that remain. Thanks also to Akshay Aitha, Inî Mendoza, Kutay Serova, and Madeline Snigaroff, the organizers of CLS 59, for the opportunity to present this work and for putting together a great conference.

[1]Chen (2018) claims that examples like (5a) are acceptable, and concludes that *zero* licenses NPIs in its restriction but not in its scope. I have not encountered speakers who find a difference in acceptability between NPIs in the restriction vs. scope of *zero* (i.e. between (3a)-(4a) on the one hand, and (5a) on the other), but several audience members at CLS 59 informed me that they find all of (3a)-(5a) acceptable; that is, for such speakers, *zero* licenses NPIs across the board. (Others reported judgments identical to those described here.) For the purpose of this paper, I will assume that such speakers can accommodate a meaning for *zero* that is equivalent to that of *no* (see (10a) below), though the nature and scope of this variation might be worth exploring in future work. See also footnote 7 for another case of potential variation.

(3) a. # Zero signals have been detected by any of our our instruments.
b. No signals have been detected by any of our our instruments.

(4) a. # Zero signals have ever been detected by our instruments.
b. No signals have ever been detected by any of our our instruments.

(5) a. # Zero signals at all have been detected by our instruments.
b. No signals at all have been detected by our instruments.

On the theoretical side, *zero* presents a challenge for contemporary views of the semantics of numerals and plurals. The standard view of plurals since at least Link 1983 is that their domains have the structure of a join semi-lattice, containing the set of all objects that can be constructed by the join operation out of the atomic objects in the domain of the corresponding singular noun. Such denotations lack minimal elements, i.e. they bottom out in atoms, and not in an empty object "$\perp$" that is a proper part of all other objects.

As for numerals, a large body of work has converged on the idea that numerals saturate a predicate in the nominal projection whose meaning can be characterized as an extensive measure function over the objects in the domain of a plural noun. In some analyses, this predicate is the denotation of the numeral itself (Landman, 2003, 2004; Rothstein, 2011; Ionin & Matushansky, 2006); in others, the measure function is introduced by some other expression in the nominal projection and takes a numeral as its argument (Cresswell, 1976; Krifka, 1989; Hackl, 2000).

Putting these two ideas together, the semantic content of a sentence like (6a) with arbitrary numeral *N* is as in (6b), where n is a number, x is an object in the domain of the plural noun, and # is a function from such objects to numbers, which gives back the number of atoms of which its argument is composed.

(6) a. *N* signals have been detected.
b. $\exists x[\#(x) = n \wedge \textbf{signals}(x) \wedge \textbf{detected}]$

Non-*zero* numerals such as *three* can then be analyzed in one of two ways. (See Kennedy 2013; Buccola & Spector 2016 and Bylinina & Nouwen 2020 for discussion of the various implementations of these general ideas, and comparisons between approaches.) The first is as expressions that directly fill in the value for n, deriving "lower-bounded" truth conditions like (7b); such meanings can then be enriched either pragmatically or through compositional exhaustification to derive upper bounds and "two-sided" truth conditions.

(7) $\exists x[\#(x) = 3 \wedge \textbf{signals}(x) \wedge \textbf{detected}(x)]$

The second approach treats numerals as generalized quantifiers over numbers (a special case of the semantic type of degrees), which take a property of numbers and give back true just in case the numbers that satisfy it meet certain conditions. This approach allows for the possibility of analyzing numerals as directly introducing upper bounds: in (8a), for example, *three* is true of a property of degrees just in case the maximum number that satisfies it is 3. The numeral takes scope and composes with the property of numbers derived by abstracting over the its base position, giving back the two-sided truth conditions in (8c).

(8) a. $[\![\text{three}]\!] = \lambda P.max(P) = 3$

b. $\lambda n \exists x[\#(x) = n \wedge \textbf{signals}(x) \wedge \textbf{detected}(x)]$

c. $max(\lambda n.\exists x[\#(x) = n \wedge \textbf{signals}(x) \wedge \textbf{detected}(x)]) = 3$

The trouble with *zero* is that neither of these approaches seems to work, and the problem for both accounts stems from the fact that, on a standard characterization of plural noun denotations as join semi-lattices, there is no object x in the domain of a plural noun like *signals* such that $\#(x) = 0$. In a number-denoting analysis of numerals the truth conditions that we derive for *zero* sentences, shown in (9a), are contradictory. And an analysis in which numerals are maximizing degree quantifiers, as in (9b), fails to derive coherent truth conditions at all.

(9) a. $\exists x[\#(x) = 0 \wedge \textbf{signals}(x) \wedge \textbf{detected}]$

b. $max(\lambda n \mid \exists x[\#(x) = n \wedge \textbf{signals}(x) \wedge \textbf{detected}(x)]) = 0$

The maximization function that the metalanguage expression max in (9b) stands for takes an ordered set as input and returns the unique element of that set that that is ordered above all other objects; as such, it does not return a value for the empty set. This approach to numerals is therefore committed to the position that their meanings are defined only when their scope terms determine non-empty sets. But since there is no object x in the domain of a plural noun such that $\#(x) = 0$, this means that the truth conditions of *zero* are incompatible with its definedness conditions. The prediction of such an analysis, then, would seem to be that use of *zero* should lead to semantic anomaly.

This is not to say that it is impossible to provide a meaning for *zero*, consistent with the rest of the above assumptions about plurals and numerals, which returns correct, contingent truth conditions. We could, for example, define max as something other than maximization, such that $max(P) = 0$ when P is empty. Or even easier, we could give *zero* a quantificational denotation along the lines of (10a), which would derive (10b) as the truth conditions for (1) (cf. Chen, 2018).

(10) a. $[\![\text{zero}]\!] = \lambda P.\forall n > 0 : \neg P(n)$

b. $\forall n > 0 : \neg\exists x[\#(x) = n \wedge \textbf{signals}(x) \wedge \textbf{detected}(x)]$

However, as pointed out by BN, either move — and indeed any analysis that resolves the problems outlined above by assigning *zero* a denotation such that its semantic contribution is to exclude propositions involving non-zero values for n — would leave it a complete mystery why *zero* fails to license NPIs. In any such analysis, the scope of *zero*, on any theory of NPI-licensing, should be one in which such expressions flourish. Based on this fact, BN conclude that *zero* cannot have such a meaning.

Instead, they argue, *zero* sentences must be analyzed as in (9a), and they propose to resolve the problem of contradiction by discarding the standard treatment of plural nouns. Instead of treating plural noun denotations as semi-lattices, BN argue, they should be characterized as complete lattices, with a bottom element $\bot$ which is such that $\#(\bot) = 0$. If this is correct, then (9a) is not a contradiction, but a tautology. $\bot$ is a part of any object x in the domain of a plural noun, so (9a) is true not only when no signals have been detected, but also when one signal has been detected, or two, or three, and so on.

The semantic content of a *zero* sentence, on this view, is not particularly useful. However, as BN point out, it can be made useful by exhaustification. For present purposes, let us assume with Fox & Spector (2018) that exhaustification involves insertion of an operator with the meaning in (11): $[\![exh]\!](p)$ entails p and excludes all of alternatives to p that p does not entail.

(11) $[\![exh]\!] = \lambda p \lambda w.p(w) \wedge \forall p' \in ALT(p) : p \not\models p' \rightarrow \neg p'(w)$

Let us further assume for simplicity that the relevant alternatives in sentences involving numerals are all and only the sentences with other numerals. With these assumptions in hand, exhaustification of (9a) delivers the truth conditions in (12).

(12) $\exists x[\#(x) = 0 \wedge \textbf{signals}(x) \wedge \textbf{detected}] \wedge$
$\forall n > 0 : \neg\exists x[\#(x) = n \wedge \textbf{signals}(x) \wedge \textbf{detected}]$

It should be clear that (12) is both contingent and provides correct truth conditions for (1): the prejacent is a tautology, so it idles, but the exclusive proposition is equivalent to (10b).[2]

And crucially, unlike an analysis in which the truth conditions are derived from a meaning for *zero* like (10a), this analysis provides the basis for explaining the difference between *zero* and *no* with respect to NPI licensing. Specifically, BN follow Gajewski 2011 in taking the licensing conditions for NPIs to be such that at least the condition in (13-i) must be satisfied, where an environment is non-trivially downward-entailing just in case it is downward-entailing and not also upward-entailing.[3] (Strong NPIs must further satisfy the condition in (13-ii).)

(13) Given a structure [$_\alpha$ exh [$_\beta$ ... [$_\gamma$ NPI] ...]]:
(i) the environment γ is non-trivially downward-entailing in β
(ii) the environment γ is non-trivially downward-entailing in α

In the case of *no*, both of these conditions are met: *no* creates a non-trivially downward-entailing context, and exhaustification does not change things. In the case of *zero*, however, things are different. Although condition (13-ii) is met in

[2]Haida & Trinh (2020) call this analysis into question, arguing that exhaustification cannot rescue sentences that are tautologous in virtue of the logical properties of the expressions that compose them, which is the case with *zero* sentences if BN are correct. The analysis that I will propose below also has this feature, and so is also subject to Haida and Trinh's challenge. On the other hand, Haida and Trinh do not provide an alternative, positive account of the semantics of *zero* or its NPI-licensing properties, and indeed it is difficult to see how such an account could be formulated in a way that doesn't render a *zero* sentence tautologous or, as I will argue below, contradictory, and in need of rescuing by exhaustification. Resolution of this tension should be a focus of future work. See Section 5 and footnote 10 for additional discussion.

[3]I use Gajewski's licensing conditions in this paper because this is what BN use, but as far as I can tell, other theories could work just as well, such as one based on non-veridicality (Giannakidou, 1998) or one based on scope-marking (Barker, 2018), provided they include a means of deriving licensing domains and are made sensitive to non-trivial — or what I will refer to below as *contingent* — entailment.

virtue of exhaustification, condition (13-i) is not: since the prejacent of exhaustification is a tautology, it is both downward *and* upward entailing.

If Bylinina and Nouwen's analysis is correct, *zero* sentences provide a compelling argument for two different, theoretically significant conclusions. First, plurals must be analyzed as having denotations with the structure of a complete lattice, rather than a semi-lattice. And second, *zero* — and by extension, all numerals — must be assigned a semantic analysis that derives lower-bounded truth conditions, since, as we saw above, a quantificational analysis involving upper-bounded truth conditions incorrectly predicts *zero* to license NPIs, regardless of which characterization of plurals we adopt.

The purpose of this paper is to argue that although Bylinina and Nouwen's basic idea about why *zero* fails to license NPIs is fundamentally correct, their proposals about plurals and the semantics of *zero* are not. Instead, I will argue that plural nouns have denotations with the structure of a join semi-lattice, as is standardly assumed, and that *zero* should in fact be assigned a quantificational denotation as a maximizing degree quantifier. The argument will proceed in three steps. In section 2, I will highlight a theoretical drawback of BN's treatment of plurals that the standard account does not run into, and I will show that a set of facts that BN present as an independent argument in favor of analyzing plurals in terms of full lattices can be accommodated on the standard account. In section 3 I will discuss a set of data involving the interaction of *zero* and *only* that is problematic for BN's analysis. And finally, in section 4, I will show that *zero* actually does license NPIs, but only when it composes with nouns which, I will suggest, we have independent reasons to analyze as having domains with minimal elements. In other words, I will argue that when nouns have the kinds of domains that BN propose for plural count nouns, *zero* is an NPI-licenser, and that the contrast between such nouns and plural count nouns indicates that *zero* has the semantics of a maximizing degree quantifier.

Taken together, these observations will eventually lead to a semantic analysis of sentences in which *zero* composes with a plural count noun as *contradictions* rather than tautologies. In Section 5, I will resolve this apparent paradox by proposing, in line with the quote from Peter Geach above, that natural language includes a mechanism for mapping sentences to exhaustifications that consist exclusively of the exclusive proposition, not the prejacent. Geach himself intended his remarks as a proposal about the semantic contribution of *only*, and a similar analysis of *only* is proposed by McCawley (1981, pp. 226-227). Horn (1996) argues that such an analysis is untenable as an analysis of *only*, but I will argue that it is not untenable — it is in fact necessary — to detach entailment of the prejacent from sentence-level exhaustification, leaving only the exclusive proposition behind, in order to account for the behavior of *zero*. More specifically, I will suggest that, in fact, the exclusive proposition is all there is to the semantic content of exhaustification. In "normal" cases of exhaustification, the exclusive proposition is not-at issue, as recently proposed by Bassi *et al.* (2021), and entailment of the prejacent is just a by-product of composition. But in special cases, including *zero* sentences, the exclusive proposition is the exclusive at-issue content of exhaustification, and the content of the prejacent is effectively erased. This analysis will support an account of *zero*'s failure to license NPIs (when it does so fail) that has essentially the same structure as the one proposed by Bylinina & Nouwen (2018).

2 Triviality and contradiction

One theoretical drawback of an analysis of plural count nouns in terms of full lattices rather than semi-lattices is that it necessitates two kinds of existential quantification. If the domain of e.g. *signals* includes $\perp$, then a simple existential claim like (14a) turns out to be trivial, because $\perp$ satisfies (14b) even when no signals are detected; for the same reason, (15a) is a contradiction.

(14) a. Signals were detected.
b. $\exists x[\textbf{signals}(x) \wedge \textbf{detected}(x)]$

(15) a. Signals weren't detected.
b. $\neg\exists x[\textbf{signals}(x) \wedge \textbf{detected}(x)]$

To avoid this problem, we must assume two kinds of existential quantification: "$\exists$", which is satisfied by $\perp$ and used for *zero* (and potentially other numerals as well), and "$\exists_{>\perp}$", which is is not satisfied by $\perp$ and is used everywhere else. There is of course a good reason for using $\exists_{>\perp}$ in non-*zero* sentences: failure to do so would result in trivial or contradictory truth conditions. But the fact remains that adding $\perp$ to the domain of plural count nouns in order to accommodate *zero* forces us to give up on a general analysis of the existential import of plural indefinites. We might wonder whether the theoretical cost of such move is worth its empirical benefit.[4]

BN are well aware of this feature of their analysis, and point to a contrast discussed by Landman (2011) as independent evidence that the domains of plural count nouns include $\perp$. Landman argues that definite plurals pattern with universal quantifiers, and differently from definite singulars, in sometimes allowing for truth-conditionally trivial interpretations when their domains are empty. For illustration, consider a scenario involving a university administrator "Doctor A," a resource-challenged department chair "Professor B," and a visiting lecturer "Professor C," who just taught a graduate seminar that included no undergraduates. Professor B needs funding from Doctor A to hire Professor C for another year, and is willing to be misleading or underinformative if that will help her achieve her goals, but is unwilling to present herself as committed to something she knows to be false. In such a context, according to Landman, Professor B can respond to Doctor A's statement in (16A) with (16B) or (16B'), but not with (16B").

(16) A: If there is undergraduate support for Professor C, I'll fund their appointment for a second year.
B: Every undergraduate who took Professor C's seminar raved about it.
B': The undergraduates who took Professor C's seminar raved about it.
B": # The undergraduate who took Professor C's seminar raved about it.

Landman's explanation for the acceptability of (16B) is that universal quantifiers do not presuppose the existence of objects in their domains, but merely implicate it, thus B's utterance is tautologous and misleading, but neither strictly false nor undefined. He also shows that if plurals include $\perp$ in their domain, then Professor B's utterance in (16B') likewise turns out to be a tautology, and so can be

[4] If $\exists_{>\perp}$ is an alternative to $\exists$, we might also wonder why we don't have the option of using exhaustification to derive a meaning for (14a) that's equivalent to a *zero* sentence.

treated in the same way. In contrast, because singulars range over atoms, (16B") properly presupposes the existence of an undergraduate who took the seminar, and so Professor B cannot utter it without committing herself to something false.

There is, however, another possible explanation for the crucial contrast between (16B') and (16B") which maintains the standard treatment of plurals in terms of join semi-lattices: assume that definite plurals but not definite singulars come with a silent (universal) distributive operator when they compose with a distributive predicate. If this is the case, then (16B') is acceptable for the same reason that (16B) is: the quantifier ranges over an empty domain. If this alternative analysis is correct, then we expect definite plurals to pattern with definite singulars in similar examples involving collective predicates. The data in (17) suggest that this prediction is correct:

(17) A: If there is undergraduate support for Professor C, I'll fund her appointment for a second year.

B: # The undergraduates who took Professor C's seminar gathered in my office to rave about it.

B': # The undergraduates who took Professor C's seminar surrounded me in the lounge and raved about it.

My takeaway from these facts is that independent evidence for $\bot$ remains thin, and so the broader theoretical questions for BN's analysis articulated above remain.

3 *Only zero*

Elliott (2019) observes that modification of *zero* with *only* in sentences like (1) results in unacceptability; the same is true for exclusives like *just* and *solely*:

(18) # Only/just/solely zero signals have been detected by our instruments.

As Elliott points out, this fact is a complete mystery on the BN analysis of *zero* and plurals. On a standard semantics for exclusives, an utterance of *only* p presupposes p and asserts that no stronger alternative to p is true. On BN's analysis, the prejacent in (18) is a tautology, so the presupposition is always satisfied, and the exclusive component is equivalent to exhaustification. (18) should therefore be fine, and in fact equivalent to (1), given that its presuppositional component cannot fail to be satisfied. One might appeal to some kind of blocking principle to explain (18), and say that an overt exclusive is bad because the same interpretation can be derived by silent exhaustification. But this kind of account would also predict overt exclusives to be bad with other numerals (assuming they give rise to lower-bounded truth conditions, as BN do), which would clearly be wrong.

In contrast, the unacceptability of (18) follows straightforwardly on a standard analysis of plurals. As we saw in Section 1, such an analysis leads to contradictory truth conditions for the prejacent in (18), so the presupposition of the exclusive can never be satisfied. The unacceptability of (18) also follows, for similar reasons, if *only* does not presuppose the prejacent but rather entails it, a view with a history dating back to the Thirteenth Century, as documented by Horn (1996).[5]

Elliott also observes that it is not the case that there is a blanket ban on modification of *zero* by *only*, and uses the contrast between (19a) and (19b) to suggest that *only zero* is infelicitous just in the special case that *zero* picks out a scalar endpoint.

(19) a. The water here has only ever risen to zero_F degrees.
b. # The water here has only ever risen by zero_F degrees.

The # in (19b) is Elliott's judgment; my own sense is that the status of this example is different from that of (18): (18) is semantically anomalous, but (19b) is a cheeky, Manner-violating way of saying that the water temperature has stayed the same. I base this claim on the fact that it is possible to find reasonably acceptable naturally occurring examples of *only zero* in a set of contexts which have in common the fact that *zero* can be construed as picking out a scalar minimum. These examples are not perfect, to be sure, and are marked compared to their counterparts without *only* in the same way that (19b) is. But the fact that they share the property of involving predicates that introduce scalar minima is, I believe, significant.

The first context is one that is semantically parallel to Elliott's (19b), and involves the differential measure phrase in a comparative. Here *zero* names the minimal degree of divergence from the standard of comparison, i.e. that value that represents no divergence from the standard at all. That *zero* has such a use is shown by the examples in (20):

(20) a. Guess what the difference in labor cost is to install crappy heat cable versus quality heat cable? Zero. It costs zero dollars more to install long-lasting, efficient ice dam heat tape than it does the cheap stuff.
b. [The Southern Border Wall] is now ZERO inches longer than the Border Wall was that existed before Trump's Inauguration.
c. In June 2024, Donald Trump will be 78 years old, or zero years older than Joe Biden is now.

And sure enough, modification by *only* is possible here:

(21) a. Infinitely better functionality and if you wait for a sale only 3/4 the price, or worst case scenario only $0 more.
b. For longer lasting polish, you may upgrade to gel polish for only $0 more.

Only is also acceptable in examples involving predicates that name scalar concepts with minimum values, such as cost, percentage, chance, probability and difference:

(22) a. Universidad Teologica del Caribe students pay only $0 to live on campus.

[5]It does not follow, however, on an analysis of *only* like the one suggested by Geach (1962) and McCawley (1981), in which *only* neither presupposes nor entails the prejacent. This fact will be crucial to the analysis of bare *only* sentences that I will present in Section 5.

b. Most of the scoring for the Golden Eagles comes from their starting lineup, as only zero percent of Oral Roberts' points come from bench players

c. ... due to novel drugs invented to combat leukemic cells and rescue the small patients which 10-20 years ago had only zero chance to survive.

d. The above [model] would be fully ranked for any $K \geqslant 3$, except at a few particular parametric values that could occur with only zero probability, due to the noises in the data model

e. It doesn't bog, or cut out. There is just ZERO difference between 1/2 or 3/4 throttle and full throttle.

Taken together, these examples suggest that *only zero* is, in fact, acceptable when it composes with an expression that uses a scale that has a minimum value: zero degree of comparative difference, zero chance, probability, difference, etc. This actually makes sense, since such scales easily allow for the possibility of (contingent) satisfaction of both the presupposition and the entailments of the exclusive. But this result also reinforces the problem that examples like (18) present for BN's analysis: if their proposals about the semantics of plural count nouns and *zero* are correct, there should be no difference in acceptability between (18) and the sentences in (21) and (22).

A final relevant set of facts comes from a contrast between what I will refer to as "extensive" mass nouns like *salt, acid, sugar* and *fat* and "intensive" mass nouns like *courage, interest, effort* and *talent*. Previous work has demonstrated that there is a distinction to be drawn between these nouns based on several systematic differences in grammatical behavior, including *wh*-determiner choice in Romance (van de Velde, 1996), acceptability with negative quantifiers in Romance (Tovena, 2001), interaction with degree modifiers and comparative morphemes in Wolof (Baglini, 2015), and acceptability with modifiers that pick out qualitative gradability in noun denotations (Morzycki, 2012; Francez & Koontz-Garboden, 2017). Several accounts of these differences have been proposed in the literature (including in the works just cited), which agree broadly on the idea that intensive mass nouns support an ordering of the objects in their extensions based on an intensive measure function which does not track mereology. Extensive mass nouns and plurals, on the other hand, support orderings based on an an extensive measure function which *does* track mereology.

What is of interest here is that these two classes of nouns differ with respect to exclusive modification of *zero*: *only zero* and *just zero* are acceptable when they compose with intensive mass nouns, but not when they compose with extensive mass nouns, as shown by the contrast between the examples in (23) and (24).[6]

(23) a. If the compensation scheme is a fixed salary, the employee puts in only zero effort and the solution is not efficient.

[6] *Zero* is otherwise acceptable with extensive mass nouns: (22a-c) were all constructed by adding *only* to naturally occurring sentences. In fact, van der Bent (2016) conducts a corpus analysis which shows that *zero* is unique among the numerals in occurring with all classes of nouns, which is presumably due to the fact that its meaning can be defined without presupposing individuation.

b. I've read the site a lot, but to have just ZERO interest when a woman is basically saying, "Look, I want to hook up. Your place or mine?"

c. Lee had only zero talent if you ignore the big shiny asterisk denoting that the zero talent was specifically in *traditional ninja practices.*

(24) a. # The most important characteristic of Tuscan bread is that it contains only ZERO salt.

b. # It's safest to drive when you have only zero alcohol in your body

c. # Using innovative technology, Montblanc is able to ensure that there is only zero oxygen inside the case.

According to Link (1983), mass nouns are like plural count nouns in having denotations that lack minimal elements. If this is correct, then the facts in (24) are unsurprising given what we have seen so far: there are no portions of salt, alchohol or water that extensively measure zero, so the prejacents of the exclusive modifiers are contradictory. But even if we assume with Francez & Koontz-Garboden 2017 that intensive mass noun denotations also have the structure of a join semi-lattice, there is a way to build on ideas in the literature cited above to make sense of the acceptability of the examples in (23). The crucial difference between intensive and extensive measures, as noted above, is that the former are not constrained to track mereology: the same extensive portion of effort, interest, etc. can correspond to different intensive degrees of effort, interest, etc. under different circumstances. There is therefore no obstacle to supposing that particular portions of effort, interest, etc. can have zero intensity, i.e. that nouns like *interest, effort*, etc. include "zero intensity" objects in their domains, rendering the prejacents in (23) contingent.

4 Zero *does* license NPIs

The previous sections challenged the analysis of plural count nouns as full lattices; in this section, I will consider data which calls into question BN's argument that *zero* cannot be analyzed as an upper-bounding degree quantifier, the central thesis of which is that such an analysis either fails to return a meaning in the first place, or else returns one that makes the wrong predictions about NPI-licensing.

As a starting point, recall that the scope term for a degree quantifier version of *zero* is degrees of the form in (25), which is true of a number/degree n if there is an object in the domain of the noun that measures n according to the noun's associated measure function $\#$.

(25) $\lambda n \exists x[...\textbf{noun}(x) \wedge \#(x) = n...]$

Now consider the hypothesis that *zero* denotes the maximizing degree quantifier in (26), in line with Kennedy's (2015) treatment of non-*zero* numerals (see also Buccola & Spector 2016; Fagan 2023).

(26) $[\![\text{zero}]\!] = \lambda P.max(P) = 0$

As we saw in Section 1, if plural count noun denotations are semi-lattices, and so include no x such that $\#(x) = 0$, then composition of *zero* with such a noun should lead to semantic anomaly, since its truth conditions are incompatible with its definedness conditions, which require its scope term to pick out a non-empty set. On the other hand, if plural noun denotations are full lattices, and *do* contain an x such that $\#(x) = 0$, then *zero*'s definedness conditions are compatible with its truth conditions, but we would wrongly predict it to be an NPI-licenser. Alternatively, if we assign *zero* a quantificational denotation like (27) (Chen, 2018), we can deliver the correct truth conditions regardless of how we analyze plural noun denotations, but we once again incorrectly predict *zero* to be an NPI licenser.

(27) $[\![\text{zero}]\!] = \lambda P.\forall n > 0 : \neg P(n)$

I think that BN's argument against an analysis of *zero* like (27) is solid: since this approach makes incorrect predictions about NPI-licensing no matter what we say about noun denotations, it can't be right. But I don't think the argument against the semantic analysis of *zero* as in (26) goes through, for two reasons.

First, strictly speaking, the definedness conditions on (26) don't predict *zero* sentences to be semantically anomalous across the board, they predict semantic anomaly only when *zero* composes with a noun whose denotation that lacks minimal objects, i.e., objects that measure 0 relative to the noun's associated measure function. On a standard analysis of plurals and mass nouns based on a join semi-lattice, nouns associated with extensive measure functions are guaranteed to lack such objects. But nouns associated with *intensive* measure functions are not. This means first that such nouns should be perfectly compatible with the maximizing denotation of *zero* in (26). It *also* means that *zero* should license NPIs with such nouns. This prediction is correct (see also Chen, 2018):

(28) *any*

a. I have zero sympathy for any voter of the red wall that voted for this lying charlatan and suffers because of his shite policies.

b. I have zero patience for any person who believes that they are "making a difference" by generalizing and spreading hate about a specific group of people.

c. I had zero familiarity with any characters in the main cast.

d. I have zero respect for any parent fighting against masks in schools.

(29) *ever*

a. The bones, which seem to have zero possibility of ever knowing life again, reveal how Israel feels about their discipline and exile.

b. Many widows and widowers have zero interest in ever engaging in a romantic relationship with another person.

c. I'm thinking of your highness/And crying long upon the loss/I've found/And on the plus and minus/It's a zero chance of ever/Turning this around. (Soundgarden 'Zero Chance', *Down on the Upside*)

(30) *at all*

a. Today I'll show you how to make your Kraft Mac and Cheese Better and DELICIOUS with almost ZERO effort at all!

b. I suppose it can sound a bit shocking to hear someone say with absolute certainty — with zero doubt at all — that there is nothing scary in the world and limitations don't exist.

c. I'm the biggest Rangers fan so trust me when I say they have zero chance at all as long as Quinn is still there coach.

Second, it is actually not the case that analyzing *zero* as in (26) automatically leads to semantic anomaly in sentences involving nouns whose denotations lack minimal elements. Kennedy (2015) shows that quantificational denotations of the sort in (26) can be lowered to number-denoting denotations through the application of Partee's (1987) **BE** and **iota** operations. The lowered denotation of *zero* is the number 0, which is semantically compatible with plural nouns, as we have seen, but derives "useless" truth conditions: tautological truth conditions for noun denotations that contain minimal elements, and contradictory truth conditions for noun denotations that do not. Bylinina & Nouwen (2018) show how to explain *zero*'s failure to license NPIs in the first case; here I will argue that the same account can be extended to the second case, with one modification.

Recall that BN's account goes as follows. When number-denoting *zero* composes with a noun whose denotation includes minimal elements, it is trivially downward entailing over its scope: for scope terms ϕ, ψ, *zero* ϕ entails *zero* ψ, no matter the logical relation between ϕ and ψ, because *zero* ϕ and *zero* ψ are both guaranteed to be true. Since (weak) NPIs are licensed only by expressions that generate non-trivially downward entailing contexts within a domain that excludes exhaustification, they are not licensed in the scope of *zero*.

My account is basically the same. When number-denoting *zero* composes with a noun whose denotation does *not* include minimal elements, it is *vacuously* downward-entailing over its scope: for scope terms ϕ, ψ, *zero* ϕ entails *zero* ψ, no matter the logical relation between ϕ and ψ, because *zero* ϕ and *zero* ψ are both guaranteed to be *false*. To capture the NPI facts, all we need to do is minimally revise the licensing conditions on NPIs as in (31), where an environment is contingently downward-entailing just in case it is both non-trivially downward-entailing and non-vacuously downward-entailing. (I will formalize contingent entailment in Section 5.)

(31) Given a structure [$_\alpha$ exh [$_\beta$ ... [$_\gamma$ NPI] ...]]:

(i) the environment γ is contingently downward-entailing in β

(ii) the environment γ is contingently downward-entailing in α

(31) handles all the cases that (13) did (involving Strawson-downward-entailment, definite descriptions, and so forth), and also accounts for *zero*'s pattern of NPI licensing. When it composes withan expression whose domain contains minimal elements, such as an intensive mass noun, it maintains its quantificational meaning and creates a contingently downward-entailing context, licensing NPIs. When it composes with an expression whose domain does not contain minimal elements, such as a plural count noun or extensive mass noun, it shifts to a number-denoting denotation, which derives contradictory truth conditions and creates a non-contingently

(vacuously) downward-entailing context, and so fails to license NPIs.[7]

We are now left with one very large open question: why aren't sentences in which *zero* composes with a plural count noun or concrete mass noun heard as contradictions? The next section provides an answer.

5 Exclusively exclusive content

I have argued that the distribution of exclusives with *zero* and the differential acceptability of NPIs with *zero* in examples involving intensive mass nouns vs. plural count nouns and extensive mass nouns are best explained by adopting a semantic analysis of *zero* as a maximizing degree quantifier, and a semantic analysis of plural count nouns (and extensive mass nouns) as having domains with the structure of a join semi-lattice. If this is correct, then *zero* must lower to a number-denoting meaning when it composes with a plural, deriving the truth conditions in (32b) for an example like (32a).

(32) a. Zero signals have been detected.
b. $\exists x[\#(x) = 0 \wedge \textbf{signals}(x) \wedge \textbf{detected}(x)]$

As we saw above, (32b) is a contradiction, because there is no object x in the domain of **signals** such that $\#(x) = 0$. Using the revised licensing conditions on NPIs in (31) that are stated in terms of contingent downward-entailingess, I was able to use this fact to explain why *zero* does not license NPIs, but now I have to say why examples like (32a), in actual use, have contingent truth conditions — why (32a) is intuitively true when no signals have been detected, and false otherwise.

Recall that BN showed that the tautologous content assigned to *zero* sentences on their analysis — which is no more useful than contradictory content — can be saved by exhaustification: exhaustification of (32a) derives a proposition that is true just in case it is false that n signals have been detected, for all $n > 0$. I will make the same claim here: *zero* sentences come to have contingent truth conditions, of exactly this sort, through exhaustification. But in order to derive this result, I need to make two adjustments to the theory of exhaustification, one rather minor, the other less so.

The minor adjustment concerns the characterization of the exclusive proposition. Intuitively, we want exhaustification of a *zero* sentence to exclude all alternatives involving other numerals, just as in BN's analysis. The version of *exh* I adopted earlier in (11) follows Fox & Spector (2018) in excluding all propositions

[7]Let me quickly note that this analysis predicts no difference between the use of *zero* in distributive vs. collective predicates, either with respect to overall acceptability, or with respect to NPI licensing, since the semantics for *zero* in (26) will render sentences with the latter just as contradictory as those with the former. The contrasts between the *no* and *zero* variants of the following examples show that this prediction is correct (pace Chen 2018).

(i) a. No/#Zero friends had ever met in a final match until Jones played Smith.
b. No/#Zero nonbelievers have ever gathered in these hallowed halls.
c. No/#Zero trees had yet surrounded any of the houses.

that are not entailed by the prejacent. But as Nick Feisher (p.c.) pointed out to me, if *zero* sentences are contradictions, as I have claimed, then there are no alternatives that they do not *vacuously* entail, and so exhaustification won't exclude anything.

There is a (relatively) simple solution to this problem, however, which is to suppose that the entailment relation that exhaustification is sensitive to — just like the one I introduced above for NPI-licensing, or the one that Fleisher (2013) argues is involved in the calculation of Scope Economy — is a *contingent* one, as defined in (33) (Smiley, 1959, p. 240).

(33) p contingently entails q, $p \vDash q$, iff $p \rightarrow q$ is a tautology, and either:
  (i) neither $\neg p$ nor q are tautologies, or
  (ii) $p \rightarrow q$ is a substitution instance of some tautology $p' \rightarrow q'$, where neither $\neg p'$ nor q' are tautologies.

(33i) is the basic condition which ensures that, in general, entailment fails when either the premises are a contradiction or the conclusion is a tautology; (33ii) allows for the special case that a conclusion or tautology still counts as entailing itself.[8] Redefining entailment in this way makes intuitive sence since exhaustification is ultimately about excluding informationally stronger alternatives, and contradictions (and tautologies) are informationally inert. It will, moreover, have no effect on standard examples of exhaustification, in which the prejacent is contingent. But when applied to a *zero* sentence like (32a) as in (34a), the result will be exclusion of all non-*zero* alternatives, as in (34b).

(34) a. *exh* [zero signals have been detected]
  b. $\exists x[\#(x) = 0 \wedge \textbf{signals}(x) \wedge \textbf{detected}(x)] \wedge$
  $\forall m > 0 : \neg \exists x[\#(x) = m \wedge \textbf{signals}(x) \wedge \textbf{detected}(x)]$

This brings us to the second, more substantive adjustment to the theory of exhaustification. The standard view, exemplified by (34b), is that exhaustification returns the conjunction of the prejacent and the exclusive proposition. But in a *zero* sentence, the prejacent is a contradiction, and the conjunction of a contradiction with another proposition is still a contradiction. What we need to do is eliminate the content of the prejacent entirely, leaving behind only the content of the exclusive proposition. But how to achieve this result?

My answer is to propose that natural language includes a mechanism for mapping sentences to their corresponding exclusive content, exclusively, as argued (incorrectly, according to Horn (1996)) originally by Geach (1962), and later by McCawley (1981), in their analyses of *only*. The intuition is that among the sentences

[8]The reasons for adopting this kind of definition are nicely stated by Clark (1967) (who further elaborates Smiley's definition to account for other cases), who observes that "...if A entails B, then A & –B inconsistent. But the converse does not hold: (A & -A) & –B is inconsistent, but we do not want to allow that (i) A & –A entails B. For we see that the inconsistency belongs merely to A & –A; A & -A is not inconsistent *with* –B in the same way as A is inconsistent *with* –A. The extreme solution, which disallows entailments with non-contingent antecedents or consequents, would certainly debar (i), but it would also debar (ii) A & –A entails A & –A, because of the intrinsic inconsistency of the antecedent and of the consequent and in spite of the additional inconsistency *between* A & –A and –(A & –A)."

generated by natural language grammars are those that are useless in virtue of their own content, but are useful in virtue of their capacity to exclude alternatives, and that the grammar of exhaustification provides a mechanism for taking advantage of this. This idea is, in fact, precisely the same as the one that BN rely on to explain why *zero* sentences are not heard as tautologies, though for them it is enough to rely on the standard theory of exhaustification to achieve these results. My claim is that *zero* sentences reveal something more interesting: the grammar of exhaustification also provides the basis for making use of "predicables that are not true of anything," to use Geach's phrasing.

My implementation of this proposal builds on recent work by Bassi *et al.* (2021), who argue for a reanalysis of *exh*, in which the exclusive component is not part of the operator's at-issue content, but is rather not at-issue.[9] This analysis resolves a number of empirical challenges for the standard analysis of exhaustification, such as the disappearance of exclusive inferences in downward-entailing contexts and some complex patterns of projection, and they explain the fact that exclusive inferences often "feel" like at-issue content through appeal to systematic global accommodation.

Here I would like to propose a refinement of the Bassi et al. analysis: *the content of exhaustification is exclusively the exclusive proposition*. The default is for the exclusionary content to be not at-issue, as argued by Bassi et al., in which case entailment of the prejacent just reflects the lack of any contribution to, or modification of, the at-issue content of the prejacent by exhaustification. But under certain circumstances, the exclusionary content can be at-issue, in which case the content of the prejacent is effectively overwritten: it is plays a role in the calculation of the truth conditions of the exhaustified proposition, but it is not itself entailed by the exhaustified proposition. The proposal is spelled out in (35), where not-at-issue content is underlined, and at-issue content is $\boxed{\text{boxed}}$.

(35) $[\![exh]\!] =$

a. $\lambda p \lambda w : \underline{\forall p' \in ALT(p) : p \not\models p' \rightarrow \neg p'(w)}.\boxed{p(w)}$ DEFAULT

b. $\lambda p \lambda w.\boxed{\forall p' \in ALT(p) : p \not\models p' \rightarrow \neg p'(w)}$ SPECIAL CASE

(35a) is the Bassi et al. semantics for *exh*; (35b) is equivalent to the Geach/McCawley semantics for *only*, and is exactly what we need to derive both the correct truth conditions and the NPI-licensing properties of *zero* sentences involving plurals and extensive mass nouns.

My proposal does not come without questions, however, the answers to which will determine its ultimate viability. The first concerns the details of the exhaustification operator's exclusion criteria. I have so far been assuming a simple characterization of exclusion based on (contingent) non-entailment, but there is a growing body of work which suggests that things are more complex than this (see e.g. Sauerland 2004; Fox 2007; Bar-Lev & Fox 2020; Bassi *et al.* 2021 and others; see Groenendijk & Stokhof 1984 for related points about *only*). In particular, this work

[9] Bassi *et al.* (2021) specifically treat the exclusive proposition as a presupposition, but allow for the possibility that it is some other kind of not at-issue content. For the purposes of this paper, this distinction does not matter.

argues that exclusion should be defined in a way that is sensitive to the consistency of (sets of) negated alternatives with the truth of the prejacent. Such a definition will, for obvious reasons, not exclude anything if the prejacent is a contradiction.

That said, the intuition driving these analyses is that exhaustification shouldn't lead to contradiction or arbitrary inclusion/exclusion of alternatives *given the truth of the prejacent*. This makes perfect sense if exhaustification entails the prejacent, as in (35a) but it arguably makes less sense if it does not, as in (35b). The challenge, then, for maintaining the hypothesis manifest in (35) — that the only difference between "regular" and "exclusively exclusive" exhaustification is whether the exclusionary content is not at-issue vs. at-issue — is to find a way of stating the exclusion criteria that is sensitive to this distinction.

A second, even more pressing question concerns the distribution of exclusively exclusive interprations of exh: why is (35a) the default, and under what conditions does (35b) emerge? It is clear that the exclusively exclusive interpretation is not freely available. If it were, there would be a use of *"Some of the signals were detected,"* for example, which would be consistent with the possibility that none of the signals were detected. Similarly, we would lose access to a set of analytical frameworks that rely on the interaction of exhaustification and grammatical triviality/contradiction to explain certain patterns of distribution, such as Chierchia's (2013) account of NPIs and the analyses of exceptives in Gajewski (2013); Hirsch (2016) and Crnič (2018). In a nutshell, these analyses argue that certain expressions — including *any* and *except* — trigger obligatory exhaustification, and the problem with examples like (36a-b) (compared to variants with negation and a universal quantifier, respectively) is that the prejacent and the exclusive proposition are contradictory, independent of arbitrary choices of content words — a property that Gajewski (2009) calls *L-triviality* — which leads to unacceptability.

(36) a. # Kim detected any signals.
b. # Someone except Kim was the first to detect these signals.

Clearly, the contradictions that these analyses rely on would not arise if the interpretation of exhaustification in (35b) were freely available, and so their explanatory power would disappear. Whether these kinds of analyses are on the right track is of course a matter of debate, but I also take it that it *should be* a matter of debate, and that it would be a strike against the proposal I have advanced here — or at least a legitimate basis for skepticism — if it were to wipe out an entire line of otherwise productive research.

One response to the distributional question is a brute force one: stipulate (via selection or some other mechanism) that the exclusively exclusive interpretation of *exh* is only available with *zero*. The theoretical cost of such a move would then need to be measured against the theoretical costs I tallied for BN's analysis in Section 2. But there is, however, an important difference between examples like (36a-b) and *zero* sentences, which may provide the basis for a more explanatory response: the prejacents in the former cases are contingent, and the prejacents in the latter are contradictory. Given that the "regular" interpretation of *exh* in (35a) returns the content of the prejacent as its exclusive at-issue contribution, it is natural to suppose that such a use also presupposes the contingency of the prejacent. At the same time, is equally natural to suppose that the exclusively exclusive use in (35b) does not — and indeed must not, if my analysis is correct — come with such a contingency presupposition, since it doesn't entail the prejacent. If these suppositions are right, they open the door to answering the distributional question by appeal to something like the principle of Maximize Presupposition (Heim, 1991; Sauerland, 2008): whenever the prejacent is contingent, exhaustification must be understood as in (35a), even if doing so leads in turn to (grammatical) contradiction. We would then derive the desired result that (35b) may be used only in the special case of "predicables that are not true of anything" — or, indeed, predicables that are true of everything.[10]

6 Conclusion

I have argued that the variable NPI licensing properties of *zero* are best explained by a semantic analysis of *zero* as a maximizing degree quantifier and a traditional semantic analysis of plural count nouns (and extensive mass nouns) in which their denotations have the structure of join semi-lattices, and lack minimal elements. In this analysis, simple sentences like *"zero signals were detected"* are contradictions in virtue of their semantic content, and they come to have contingent truth conditions because the grammar of exhaustification allows for the option of returning exclusively exclusive at-issue content.

It must be admitted that this is a rather big conclusion to draw from the behavior of a rather small word, and so it is appropriate to ask: is this conclusion justified? Morzycki (2017) suggests that *zero* is a "semantic virus" — a kind of add-on to the basic vocabulary and meanings of English, with unusual properties that sets it outside the core vocabulary and meanings of numerals, plurals, measure terms and so forth. But as we have learned all too well during the first part of the third decade of the Twenty-First Century, viruses require a certain kind of compatibility with their hosts, and are able to thrive by taking advantage of existing structures and operations. If my interpretation of the facts involving *zero* are correct, then we have some initial reasons to think that Geach and McCawley were on the right track in proposing that natural language includes mechanisms that allow for the expression of exclusively exclusive content. Whether this hypothesis turns out to be correct will, of course, depend a lot on whether we can find more than zero additional instances of its application.

[10] The prejacent of *exh* in a *zero* sentence is plausibly *grammatically* contradictory, in the sense described above (L-triviality), if the extensive measure functions associated with plurals and extensive mass nouns are introduced by functional structure. This will likely turn out to be important, because it also does not appear to be the case that (35b) can be used to "rescue" sentences that are contradictory in virtue of arbitrary lexical content, as pointed out to me by Clemens Steiner-Mayr. If it could, then a sentence like *"The grass is colorless and green"* could be used to communicate that the grass is some color other than green, assuming alternatives {*"The grass is colorless," "The grass is green"*}. That said, the idea that the exclusively exclusive meaning of *exh* emerges only in composition with an L-trivial prejacent is in direct conflict with Haida and Trinh's (2020) claim that exhaustification can't rescue a sentence from L-triviality, as I observed in footnote 2.

References

Baglini, R. 2015. *Stative Predication and Semantic Ontology: A Cross-linguistic Study*. University of Chicago dissertation.

Bar-Lev, M., & D. Fox. 2020. Free choice, simplification, and innocent inclusion. *Natural Language Semantics* 28. 175–223.

Barker, C. 2018. Negative polarity as scope marking. *Linguistics and Philosophy* 41. 483–510.

Bassi, I., G. Del Pinal, & U. Sauerland. 2021. Presuppositional exhaustification. *Semantics and Pragmatics* 14.

Buccola, B., & B. Spector. 2016. Modified numerals and maximality. *Linguistics and Philosophy* 39. 151–199.

Bylinina, L., & R. Nouwen. 2018. On "zero" and semantic plurality. *Glossa* 3. 1–23.

Bylinina, L., & R. Nouwen. 2020. Numeral semantics. *Language and Linguistics Compass* .

Chen, S. Y. 2018. *Zero* degrees: numerosity, intensification, and negative polarity. In *Proceedings of CLS 54*, 53–67, Chicago, IL. Chicago Linguistics Society.

Chierchia, G. 2013. *Logic in Grammar*. Oxford, UK: Oxford University Press.

Clark, M. 1967. The general notion of entailment. *The Philosophical Quarterly* 17. 231–245.

Cresswell, M. J. 1976. The semantics of degree. In *Montague Grammar*, ed. by B. Partee, 261–292. New York: Academic Press.

Crnič, L. 2018. A note on connected exceptives and approximatives. *Journal of Semantics* 35. 741–756.

Elliott, P. 2019. #only zero. *Snippets* 35. 1–2.

Fagan, L. 2023. Numerals are doubly bounded: Evidence from exclusives and polarity. In *Proceedings of NELS 53*, ed. by S.-L. Lam & S. Ozaki, Address. GLSA, University of Massachusetts.

Fleisher, N. 2013. Comparative quantifiers and negation: Implications for scope economy. *Journal of Semantics* 32. 139–171.

Fox, D. 2007. Free choice disjunction and the theory of scalar implicatures. In *Presupposition and Implicature in Compositional Semantics*, ed. by U. Sauerland & P. Stateva, 71–120. Palgrave MacMillan.

Fox, D., & B. Spector. 2018. Economy and embedded exhaustification. *Natural Language Semantics* 26. 1–50.

Francez, I., & A. Koontz-Garboden. 2017. *Semantics and Morphosyntactic Variation: Qualities and the Grammar of Property Concepts*. Oxford, UK: Oxford University Press.

Gajewski, J., 2009. On analyticity in natural language. Unpublished ms., University of Connecticut.

Gajewski, J. 2011. Licensing strong NPIs. *Natural Language Semantics* 19. 109–148.

Gajewski, J. 2013. An analogy between a connected exceptive phrase and polarity items. In *Beyond* any *and* ever, volume 262, 183–212, Berlin. Walter de Gruyter.

Geach, P. T. 1962. *Reference and Generality*. Ithaca, NY: Cornell University Press.

Giannakidou, A. 1998. *Polarity Sensitivity as (Non)Veridical Dependency*. Amsterdam and Philadelphia: John Benjamins.

Groenendijk, J., & M. Stokhof. 1984. *Studies on the Semantics of Questions and the Pragmatics of Answers*. Amsterdam: University of Amsterdam dissertation.

Hackl, M. 2000. *Comparative Quantifiers*. Massachusetts Institute of Technology dissertation.

Haida, A., & T. Trinh. 2020. Zero and triviality. *Glossa* 5. 1–14.

Heim, I. 1991. Articles and definiteness. In *Semantics. An International Handbook of Contemporary Research*, ed. by A. von Stechow & D. Wunderlich, 487–535. Berlin: Walter de Gruyter & Co.

Hirsch, A. 2016. An unexceptional semantics for expressions of exception. In *Proceedings of the 39th Annual Penn Linguistics Conference*, volume 221. University of Pennsylvania Working Papers in Linguistics.

Horn, L. 1996. Exclusive company: *Only* and the dynamics of vertical inference. *Journal of Semantics* 13. 1–40.

Ionin, T., & O. Matushansky. 2006. The composition of complex cardinals. *Journal of Semantics* 23. 315–360.

Kennedy, C. 2013. A scalar semantics for scalar readings of number words. In *From Grammar to Meaning: The Spontaneous Logicality of Language*, ed. by I. Caponigro & C. Cecchetto, 172–200. Cambridge University Press.

Kennedy, C. 2015. A "de-Fregean" semantics (and neo-Gricean pragmatics) for modified and unmodified numerals. *Semantics and Pragmatics* 8. 1–44.

Krifka, M. 1989. Nominal reference, temporal constitution and quantification in event semantics. In *Semantics and Contextual Expression*, ed. by R. Bartsch, J. van Benthem, & P. van Emde Boas, 75–115. Stanford, CA: CSLI Publications.

Landman, F. 2003. Predicate argument mismatches and the adjectival theory of indefinites. In *From NP to DP*, ed. by M. Coene & Y. d'Hulst, volume 1, 211–238. John Benjamins.

Landman, F. 2004. *Indefinites and the Type of Sets*. Oxford University Press.

Landman, F. 2011. Boolean pragmatics. In *This Is Not a Festschrift*, ed. by J. van der Does & C. Dulith Novaes. Amsterdam, Netherlands: Universiteit van Amsterdam.

Link, G. 1983. The logical analysis of plurals and mass terms: A lattice theoretical approach. In *Meaning, Use, and the Interpretation of Language*, ed. by R. Bäuerle, C. Schwarze, & A. von Stechow, 302–323. Berlin: de Gruyter.

McCawley, J. D. 1981. *Everything that Linguists have Always Wanted to Know About Logic* *but were ashamed to ask*. Chicago: University of Chicago Press.

Morzycki, M. 2012. The several faces of adnominal degree modification. In *Proceedings of WCCFL 29*, ed. by J. Choi, E. A. Hogue, J. Punske, D. Tat, J. Schertz, & A. Trueman, 187–195. Somerville, MA: Cascadilla Press.

Morzycki, M. 2017. Some viruses in the semantics. In *A Schrift to Fest Kyle Johnson*, 281–291. Linguistics Open Access Publications.

Partee, B. 1987. Noun phrase interpretation and type shifting principles. In *Studies in Discourse Representation and the Theory of Generalized Quantifiers*, ed. by G. Groenendijk, D. de Jongh, & M. Stokhof, 115–143, Dordrecht, Reidel, Netherlands. Foris.

Rothstein, S. 2011. Counting, measuring and the semantics of classifiers. *The Baltic International Yearbook of Cognition, Logic and Communication* 6.

Sauerland, U. 2004. Scalar implicatures in complex sentences. *Linguistics and Philosophy* 27. 267–391.

Sauerland, U. 2008. Implicated presuppositions. In *The Discourse Potential of Underspecified Structures*, ed. by A. Steube, 581–600. De Gruyter.

Smiley, T. J. 1959. Entailment and deducibility. *Proceedings of the Aristotelian Society* 233–254.

Tovena, L. 2001. Between mass and count. In *Proceedings of WCCFL 20*, ed. by K. Megerdoomian & L. Bar-el, 565–578, Somerville, MA. Cascadilla Press.

van de Velde, D. 1996. *Le Spectre Nominal*. Paris: Peeters.

van der Bent, R., 2016. The meaning and use of zero: An exploration of the semantic interpretation of the number word *zero*. BA thesis.

Zeijlstra, H. 2007. Zero licensers. *Snippets* 16. 21–22.