

Structure and Ontology in Nonlocal Readings of Adjectives

Marcin Morzycki
Michigan State University

Abstract In certain uses, adjectives appear to make the semantic contribution normally associated with adverbs. These readings are often thought to be a peripheral phenomenon, restricted to one corner of the grammar and just a handful of lexical items. I'll argue that it's actually considerably more general than is often recognized, and that it admits two fundamentally different modes of explanation: in terms of the syntactic machinery that undergirds these structures and in terms of the ontology of the objects manipulated by its semantics. Both modes of explanation have been suggested for some of the puzzles in this domain, and I'll argue both are necessary. With respect to adjectives including *average* and *occasional*, the key insight is that their lexical semantics is fundamentally about kinds. But to arrive at a more general theory of adverbial readings, it is also necessary to further articulate the compositional semantics. In this spirit, I'll argue that these adjectives actually have the semantic type of quantificational determiners like *every*. If this way of thinking about adverbial readings is on the right track, it instantiates a means by which these two distinct modes of explanation—and the distinct aspects of cognition they may ultimately be associated with—both play a crucial role in bringing about the apparently aberrant behavior of this class of adjectives.

Keywords adjectives, nonlocal readings, average, occasional, kinds, natural language metaphysics

I've presented various portions and versions of this work to various people and audiences. Thanks to Adam Gobeski, Ai Kubota, Ai Taniguchi, Anna Szabolsci, Anne-Michelle Tessier, Barbara Partee, Berit Gehrke, Bernhard Schwarz, Cara Feldscher, Chris Barker, Curt Anderson, Daniel Gutzmann, Friederike Moltmann, Gabriel Roisenberg Rodrigues, Galit Sassoon, Haley Farkas, Hannah Forsythe, Henry Davis, Hotze Rullmann, Irene Heim, Josh Herrin, Kay Ann Schlang, Kyle Rawlins, Lisa Matthewson, Lucas Champolion, Manfred Krifka, Norbert Hornstein, Omer Preminger, Paul Pietroski, Rose-Marie Déchaine, Sebastian Löbner, Stephanie Solt, Taehoon Hendrik Kim, Yan Cong, and audiences at the University of Maryland, Sinn und Bedeutung in Tübingen, Boston University, New York University, the University of British Columbia, and of course the Cognitive Structures conference. Berit and Sebastian kindly provided written comments on an earlier draft of this paper, and I'm especially grateful to Sebastian for involving me in the intellectual community around Düsseldorf.

1 Introduction

It is, of course, not news that the way language organizes the world may tell us something about how the mind does so. Nor is it news that that perhaps the best window into how language organizes the world is how language works: what words mean and how grammars manipulate those meanings. This is the project that Emmon Bach memorably dubbed ‘natural language metaphysics’ (Bach 1986, 1989) or ‘natural language ontology’. Importantly, it’s a project that’s worthwhile even if—perhaps especially if—it should fail to coincide with metaphysics proper, because our theory of natural language metaphysics is a repository of linguistic analysis. If we’re doing it right, its structure explains the structure of language. The structure of the world is another matter entirely, best left to others.

There is an important trade-off in this domain, however, that I’d like to use to frame this paper. Structures in natural language ontology can serve to explain linguistic phenomena, and when they do, they may lighten the explanatory burden on other components of linguistic theory, including the syntax and semantics. Conversely, introducing complexity in the syntax or semantics can make possible a simpler ontology.

It may help to sketch an example of what I have in mind. It’s entirely independent of the one I’ll focus on primarily in this paper. It concerns polar antonyms of adjectives, such as *tall* and *short*. No one would defend the view that they are unrelated, of course, so only the question is where to install a theory of their difference. One possibility is ontological. Height is measured in abstract representations of measurement, *degrees*, which include things like ‘6 feet’. The set of degrees that measure height (as opposed to e.g. weight) tell us the dimension along which a given measurement exists, but they don’t actually tell us whether we’re measuring how tall someone is or how short. They tell us that the dimension is spatial extent, but not whether the scale is tallness or shortness. To know that, we must know at least one more thing: the ordering imposed on those degrees. ‘6 feet’ is a greater degree of tallness than ‘5 feet’, but a lesser degree of shortness. On such a theory, advocated in Kennedy (1997, 2001) and elsewhere, the key to the relation between *tall* and *short* is that they measure on scales that impose opposite orderings on degrees along the same dimension. Their denotations therefore need not reflect any direct connection between the adjectives beyond specifying which scale they use, because the connection is between the scales, not between adjectives themselves.

The alternative is to suppose that the relation between *tall* and *short* is a matter of grammar, not (primarily) ontology, and that they use precisely the same scale after all. One might suppose, with Heim (2006, 2008) and Büring (2007), that *short* involves a special kind of negation, present in the syntactic tree but not normally pronounced as a separate morpheme. *Short*, on this view, is actually a way of pronouncing ‘little tall’ or ‘untall’. There are a variety of arguments to be made for this more complicated syntax, and with it in place, the ontology needn’t provide an independent analysis of the connection between the two antonyms because the richer syntax already does.

It’s not the case, of course, that any analysis of any arbitrary phenomenon can be said to be primarily grammatical or ontological. In the context of this volume,

it's especially worth noting that an approach that involves decomposition into features might occupy an intermediate position with respect to this distinction: the decomposition is in some respects like decomposing *short* into 'little tall', but of course the decomposition needn't be implemented directly in the syntax in this way, and there are interesting discussions to be had about the relationship between decomposing word meanings and decomposing the underlying concepts themselves. The former seems still a robustly grammatical enterprise; the latter considerably more an ontological one.

All that said, at some point there's a danger of putting more weight on this distinction than it can bear. Its purpose here is chiefly just to situate another empirical puzzle for which a balance has to be struck between grammatical and ontological explanation: adjectives like *average*. The first thing to notice is the curious effect they often have on the referent of the nominal in which they occur:

(1) The average American has 2.3 children.

This sentence, Carlson & Pelletier (2002) point out, is doubly mysterious. What sort of entity is 'the average American'? Certainly, on its most natural reading, it doesn't refer to some particular American who is especially typical of Americans. Second, what sort of entity is '2.3 children'? If *the average American* referred to a particular American—say, one named Steve—it would suggest, alarmingly, that Steve has only a fraction of one of his children. That's not what the sentence means, at least ordinarily. Nor, indeed, is it possible to straightforwardly disentangle the strangeness of the first nominal from the strangeness of the second. Even if we avoid the reading under which (1) involves direct reference to Steve, it still fails to communicate that it is typical for Americans to have fractional children.

On its face, it would seem that to avoid such morally outlandish outcomes, we must embrace a metaphysically outlandish one. We must accept that there are such things as 'average Americans' in the model underlying the semantics, and indeed perhaps in some extended sense such things as '2.3 children'. I don't think we should dismiss this possibility too readily. For one thing, as Bach would remind us, our judgment in these matters must be guided by language, not a priori notions about what sorts of objects populate the actual world. That's the difference between natural language metaphysics and metaphysics proper. Indeed, this metaphysical direction is precisely the one in which Carlson & Pelletier head. For this reason, Hornstein (1984) was ultimately mistaken in saying that 'no one wishes to claim that there are objects that are average men in any meaningful sense'. Yet, he argued, nominals like *the average American* act no different from more referentially pedestrian ones. He concluded that this was an argument against the enterprise of formal semantics itself.

My aim here will be more modest. Kennedy & Stanley (2009) observed that sentences such as (1) can be analyzed as a special case of a more general phenomenon: readings of adjectives in which the adjective is interpreted as though it were an adverbial. This requires a more complex syntax, but that more complex syntax is a low price to pay for the metaphysical benefit. It frees us from having to posit any spookily abstract and therefore implausible entities in the ontology. I'll argue, building on Morzycki (2016b), these adverbial readings are

in fact part of a considerably more general pattern of readings available to a far wider range of adjectives than generally recognized. I'll argue that these readings actually fall into three classes, and that this leads us to an analysis distinct in important respects from Kennedy & Stanley but that, as they argued, places the explanatory burden on the syntax and compositional semantics rather than the ontology.

In section 1, following largely the argument in Morzycki (2016b), I'll present the case that what I'll call nonlocal readings of adjectives (following Schwarz 2006 et seq.) are far more general than is typically recognized, and that they fall into three distinct classes. In section 2, I'll review some ways these problems have been approached in the past, highlighting the interplay between grammatical and ontological explanation. In section 3, I'll propose a strategy for approaching these facts that I hope may eventually scale up to the larger empirical picture and that has components of both kinds of explanation. In particular, I'll combine elements of syntactic assumptions that have widely been made with a new ingredient in the compositional semantics: the idea that adjectives with external readings have determiner-like meanings, and as a consequence have the complex grammar associated with determiners. I'll sketch this idea in general terms for *average* in particular, relating it to Gehrke & McNally (2010, 2015)'s crucial insight that adjectives like *occasional* involve reference to kinds. Finally, in section 4, I'll very briefly return to the larger issues with which we began: the analytical balance between structure in the syntax and semantics and structure in the ontology.

2 Nonlocal readings of adjectives

2.1 On '*occasional*'

Let's begin with the classic example of a nonlocal reading of an adjective, which is *occasional* (Bolinger 1967, Stump 1981, Larson 1999, Zimmermann 2003, Schäfer 2007, Gehrke & McNally 2010, 2015, DeVries 2010). It's the best-studied such case, and this will serve as a useful background against which to consider *average*. The standard sentence is (2):

- (2) An occasional sailor strolled by.
 - a. INTERNAL: 'Someone who sails occasionally strolled by.'
 - b. EXTERNAL: 'Occasionally, a sailor strolled by.'

It has what's called an internal and an external reading. The internal reading is interesting in a number of respects, but from our current perspective, it's the external reading that is most immediately relevant. On this reading, the adjective makes a semantic contribution that is, to all appearances, completely divorced from the nominal in which it finds itself. The sailors that strolled by are sailors simpliciter. There is no question about the frequency of their sailing. But the situation is more puzzling still. On the external reading, the sentence means more or less the same thing as (3), where the definite determiner replaces the indefinite:

- (3) The occasional sailor strolled by.

Yet the meaning is essentially the same (but see Gehrke & McNally 2015 for detailed discussion). Indeed, some adjectives of this class (*odd* and *rare*) have the external reading *only* with *the*.¹ Setting apart a subtle change of flavor, the external reading also occurs with *your* and in the bare plural:

- (4) a. Your occasional sailor strolled by.
b. Occasional sailors strolled by.

So there are three mysteries so far: an ambiguity, unexpectedly wide scope, and unexpected interpretations of the determiner.

There are more still. Another is that, on the external reading, the adjective must occupy the leftmost position in the structure of the nominal:

- (5) The angry occasional sailor strolled by.
a. INTERNAL: ‘Someone angry who sails occasionally strolled by.’
b. #EXTERNAL: ‘Occasionally, an angry sailor strolled by.’

Indeed, the range of determiners with which *occasional* is possible on the external reading is relatively limited:

- (6) $\left. \begin{array}{l} \text{Every} \\ \text{Some} \\ \text{Several} \\ \text{Many} \\ \text{Most} \end{array} \right\} \text{occasional sailor(s) strolled by.}$
a. INTERNAL: ‘*D* person/people who sail(s) occasionally strolled by.’
b. #EXTERNAL: ‘Occasionally, *D* sailor(s) strolled by.’

Yet another idiosyncrasy of the external reading is that it renders the adjective unable to coordinate with ordinary adjectives:

- (7) The occasional and angry sailor strolled by.
a. INTERNAL: ‘Someone angry who sails occasionally strolled by.’
b. #EXTERNAL: ‘Occasionally, an angry sailor strolled by.’

Another still: on this reading, the adjective becomes incompatible with degree words such as *very* or the comparative:²

- (8) The very occasional sailor strolled by.
a. INTERNAL: ‘Someone who sails very occasionally strolled by.’
b. #EXTERNAL: ‘Very occasionally, a sailor strolled by.’

¹Berit Gehrke (p.c.) points out that this fact doesn’t follow from what will be proposed here—but then, I don’t really have an analysis to offer here of the *occasional* class more generally. That said, this fact is precisely what one might expect if, as Larson (1999) has argued, syntactic incorporation into a determiner gives rise to some lexical idiosyncrasy here. See section 5 for more.

²For some speakers, even the internal reading is missing. Others can get an external reading marginally with *very*, but not with *more*.

- (9) The more occasional sailor strolled by.
 a. INTERNAL: ‘Someone who sails more occasionally strolled by.’
 b. #EXTERNAL: ‘More occasionally, a sailor strolled by.’

2.2 Returning to ‘average’

Having noted the crucial features of *occasional*, let’s return to *average* with them in mind. First, there was ambiguity. As Carlson & Pelletier (2002), Kennedy & Stanley (2009) among others noted, there is an ambiguity with *average* too:

- (10) An average American has 2 children.
 a. INTERNAL: ‘An American, who is typical, has 2 children.’
 b. EXTERNAL: ‘On average, an American has 2 children.’

For the internal reading to be available without counterpragmatically ghastly background assumptions, we must change our earlier sentence to *2 children*. On this reading, the claim is that there is an American somewhere that is typical and that he has two children. There is another reading of *average* that also occurs in (11) (Sebastian Löbner, p.c.), which is also internal, or in any case fails to be external:

- (11) He’s so average.

The external reading is the one with which we are now familiar from *occasional*. It’s worth noting that it paraphrases naturally with an adverbial, *on average*, which is analogous to how *occasional* morphed into *occasionally*.

Here we encounter a set of properties that elegantly mirror those of *occasional*. There are unexpected interpretations of the determiner. Switching to the definite determiner leaves us, on the external reading, with apparently the same interpretation, and *your* is not much different:

- (12) $\left\{ \begin{array}{l} \text{The} \\ \text{Your} \end{array} \right\}$ average American has 2 children.
 a. INTERNAL: ‘ $\left\{ \begin{array}{l} \text{The} \\ \text{Your} \end{array} \right\}$ American that’s a typical one has 2 children.’
 b. EXTERNAL: ‘On average, an American has 2 children.’

Again, on the external reading, other determiners don’t seem to work:

- (13) # $\left\{ \begin{array}{l} \text{Every} \\ \text{Most} \\ \text{Some} \\ \text{Several} \\ \text{Two} \end{array} \right\}$ average American(s) has/have 2.3 children.

And again, on the external reading *average* has to be leftmost among the adjectives in its nominal:

- (14) a. An average irritable American has 2.3 children.
b. #?An irritable average American has 2.3 children.

It is also unable to coordinate with other adjectives on the external reading:

- (15) #An irritable and average American has 2.3 children.

It is incompatible with degree modifiers on this reading:

- (16) #A very average American has 2.3 children.

So, once again, the same mysterious patterns manifested themselves as with *occasional*. At a minimum, this supports the connection between the two that Kennedy & Stanley (2009) posited—perhaps indeed more robustly than they intended. But the pattern is more widespread still.

2.3 *Wrong*

Before considering the bigger picture, it will be necessary to lay out a few more examples of the general phenomenon. A version of the now-familiar pattern emerges once again with *wrong* (Haik 1985, Schmitt 2000, Schwarz 2006, to appear). It too has an internal/external ambiguity, though perceiving it is slightly trickier. Suppose Floyd is a spy who is required to provide his interlocutor with false information and deprive her of true information. If he succeeds in this, (17) is true on the internal reading, on which the information provided was incorrect:

- (17) Floyd gave the wrong answer.
a. INTERNAL: 'Floyd gave an answer that was incorrect.'
b. EXTERNAL: 'Floyd gave an answer that it was wrong of him to give.'

On the external reading, (17) is false, because Floyd answered as he is supposed to. On the other hand, if Floyd slips up at some point and accidentally answers a question truthfully, the situation is flipped: (17) is still true, but only on the external reading: he provided information that he isn't supposed to provide, namely, true information. Something similar happens in (18):

- (18) Floyd killed the wrong person.
a. INTERNAL: 'Floyd killed a person that was wrong (perhaps prone to error or wrong in general).'

- b. EXTERNAL: 'Floyd killed a person that it was wrong of him to kill.'

Again, the internal reading in (18) is more easily discerned with some context. Consider a dystopian game show in which participants are executed for answering a quiz question incorrectly. Floyd is the executioner. If he killed the contestant that answered incorrectly, (18a) is true only on the internal reading. ('Clyde was wrong, so I killed him,' he might explain.) If Floyd accidentally killed a contestant that provided the correct answer, (18b) would be true only on the external reading.

There is again an odd fact about the interpretation of the determiner: *the* is interpreted as an indefinite. In (17), there need not have been only one wrong answer, and in (18), there need not have been only one person who must not be killed. The picture is slightly different, though. *Your* is impossible here except on its usual possessive reading, irrelevant here:

- (19) a. ?Floyd gave your wrong answer.
 b. ?Floyd killed your wrong person.

Strangely, it's not just that the definite determiner is interpreted as an indefinite, but it's the principal way to say this. The indefinite would be unusual on the external reading:

- (20) a. Floyd gave a wrong answer.
 b. Floyd killed a wrong person.

It's not actually fully clear what reading these receive. For me, an external reading is possible, but only when there is a desire to communicate that there are multiple answers that shouldn't be given and people that shouldn't be killed.

Apart from that quirk, again we encounter restrictions on the choice of determiner on the external reading:

- (21) #Floyd opened $\left. \begin{array}{l} \text{every} \\ \text{most} \\ \text{some} \\ \text{several} \\ \text{two} \end{array} \right\}$ wrong envelope(s).

As before, inherently quantificational determiners fail.

The requirement that the nonlocal adjective be structurally higher than other adjectives again emerges:

- (22) a. Floyd opened the wrong brown envelope.
 b. #Floyd opened the brown wrong envelope.

So does the ban on coordination:

- (23) #Floyd opened the wrong and brown envelope.

And so does the ban on degree modification:

- (24) #Floyd opened the very wrong envelope.

So a rather large class of adjectives that includes *wrong*, *average*, *typical*, *occasional* and a number of its synonyms seems to manifest quite a number of common properties.

2.4 ‘Whole’ and ‘entire’

The parallels continue with *whole* and *entire*, though there will be an important twist. As before, there is an ambiguity (Moltmann 1997, 2005, Morzycki 2002), which I’ll assume is a special case of the internal/external ambiguity:

- (25) A whole ship was submerged.
 a. INTERNAL: ‘A complete, structurally intact ship was submerged.’
 b. EXTERNAL: ‘A ship was wholly submerged.’
- (26) The whole apple is terrible.
 a. INTERNAL: ‘The complete, structurally intact apple, the one with no bites taken out of it, is terrible.’
 b. EXTERNAL: ‘All parts of the apple are terrible.’

The internal reading is actually the unusual one in these cases, and may take a moment to perceive. It’s what could be expressed more or less unambiguously with *complete*—indeed, I suspect that it’s precisely the existence of this unambiguous alternative that accounts (on broadly Gricean grounds) for the unnaturalness of the internal reading.

As before, there are restrictions on the determiner, but they take a different form. First, *a*, *the*, and *your* retain their usual meanings, and don’t become interchangeable. Second, strong quantifiers are still incompatible with the external reading, but weak ones are perfectly compatible with it (I will now indulge in the habit of marking sentences with a # when they are impossible on the external reading):³

- (27) a. $\left. \begin{array}{l} \# \text{Every} \\ \# \text{Most} \\ \text{Many} \\ \text{Several} \\ \text{Two} \end{array} \right\} \text{ whole ship(s) } \left\{ \begin{array}{l} \text{was} \\ \text{were} \end{array} \right\} \text{ submerged.}$
- b. $\left. \begin{array}{l} \# \text{Every} \\ \# \text{Most} \\ \text{Many} \\ \text{Several} \\ \text{Two} \end{array} \right\} \text{ whole apple(s) } \left\{ \begin{array}{l} \text{is} \\ \text{are} \end{array} \right\} \text{ terrible.}$

The other, now increasingly familiar restrictions reemerge in their customary form. The external reading is only possible when the nonlocal adjective occurs high:

- (28) a. A whole enormous ship was submerged.
 b. An enormous whole ship was submerged. *(internal only)*

It’s incompatible with coordination:

³Sebastian Löbner (p.c.) points out that one might explain the ill-formed examples in (27) because one nominal can’t express two different quantifications (Löbner 2000), which would accord with the grammaticality of adverbial *entirely* in e.g. *Every ship was entirely white*.

(29) A whole and enormous ship was submerged. *(internal only)*

And it's incompatible with degree modification:

(30) An entirely whole ship was submerged. *(internal only)*

2.5 Epistemic adjectives

Abusch & Rooth (1997) observed a proposition-modifying interpretation of what they called 'epistemic adjectives' that now won't come as a shock. These adjectives include *unknown*, *undisclosed*, *unspecified*, and *unexpected*. They can receive a wide-scope reading:

- (31) Solange is staying at an unknown hotel. (Abusch & Rooth 1997)
- a. INTERNAL: 'Solange is staying at a hotel no one has heard of.'
 - b. EXTERNAL: 'Solange is staying at a hotel and it is not known which hotel she is staying at.'

The external reading systematically supports concealed-question paraphrases. For many years in the early 2000s, (32) was a kind of running joke in American political discourse, and it's actually very hard to make sense of its internal reading:

(32) Dick Cheney is hiding at an undisclosed location.

The external reading is that Dick Cheney is hiding at a location and it has not been disclosed, for his safety, what location that is. On its internal reading, perhaps it would have to be the very fact that it is a location that is not disclosed.

At this stage, we will encounter the same empirical refrain, and the reader can presumably sing along. On the external reading, there are again restrictions on the determiner. Although *the* and *a* seem to behave normally, strong inherently quantificational determiners remain impossible:

(33) Solange stayed at $\left. \begin{array}{l} \# \text{every} \\ \# \text{most} \\ \text{some} \\ \text{several} \\ \text{two} \end{array} \right\} \text{unknown hotel(s)}.$

As for *whole*, weak determiners are compatible with external readings.

The restrictions on the structural position of the adjective in the DP remain the same. The external reading is, as we have come to expect, possible only when the adjective is high:⁴

⁴Sebastian Löhnert suggests a mode of explanation of this fact: the concealed question-style semantics reveals these nominals denote individual concepts, which is incompatible with the sort of run-of-the-mill extensional intersective adjectival modification attempted in (34) and (35). Perhaps that strategy could help with the quantificational facts in (33) as well.

- (34) a. Solange stayed at an unknown horrible hotel.
 b. Solange stayed at a horrible unknown hotel. *(internal only)*

The external reading is unavailable when the adjective occurs in a coordinate structure:

- (35) #Solange stayed at a horrible and unknown hotel. *(internal only)*

It's incompatible with degree modification:

- (36) #Solange stayed at a very unknown hotel. *(internal only)*

2.6 *Same and different*

Other adjectives fall under broadly the same rubric. Among the best-studied of these are *same* and *different* (Nunberg 1984, Heim 1985, Carlson 1987, Keenan 1992, Moltmann 1992, Beck 2000, Lasersohn 2000, Majewski 2002, Alrenga 2006, 2007a,b, Barker 2007, Brasoveanu 2011). The facts in this domain are complicated in ways that muddy the waters considerably, and the terminology is different and confusing, but for our purposes the important point is that there is an ambiguity.

The main terminological confound is that the internal reading involves an anaphoric dependency on preceding discourse. This is in an important sense 'external', but it is not external in the *relevant* sense of seeming to require the adjective to access to the semantic content of the clause outside the nominal itself. This is clearer when considering the readings:

- (37) Floyd and Clyde read the same book.
 a. INTERNAL (ANAPHORIC): 'Floyd and Clyde read a book that is the same as the one previously mentioned.'
 b. EXTERNAL: 'Floyd and Clyde read a book in common.'
- (38) Floyd and Clyde read a different book.
 a. INTERNAL (ANAPHORIC): 'Floyd and Clyde read a book that is the different from the one previously mentioned.'
 b. EXTERNAL: 'The book Floyd read was not the same book as the one Clyde read.'

I won't rehearse the full song-and-dance yet again, in part because it presents, in this instance, complications that go considerably beyond the scope of this paper. Suffice it to say that on the external reading, *same* and *different* impose restrictions on the determiner with which they combine:

- (39) *Floyd and Clyde read $\left. \begin{array}{l} \text{every} \\ \text{most} \\ \text{some} \\ \text{several} \\ \text{two} \end{array} \right\}$ same book(s).

On this reading *same* and *different* are subject to the now familiar structural position requirement:

- (40) a. Floyd and Clyde read the same good book.
- b. *Floyd and Clyde read the good same book.

2.7 *Modal superlatives: 'possible' and its kin*

There is another important class of nonlocal readings of adjectives, which I will mostly set aside. These involve *possible*, *conceivable*, and the like ('modal superlatives'; Bolinger 1967, Larson 2000, Schwarz 2005, Cinque 2010, Romero 2013, Leffel 2014):

- (41) They interviewed every possible candidate.
 - a. EXTERNAL: 'They interviewed every candidate that it was possible to interview.'
 - b. INTERNAL: 'They interviewed every person who was possibly a candidate.'

There are important distinctions between these cases and the ones we've examined so far, but for the moment I will note only the similarity: again, there is an ambiguity between an internal and external reading.

2.8 *Miscellaneous obscurities and novelties*

Without further discussion, I'll note a few examples of nonlocal readings that are either obscure or, to my knowledge, novel:

- (42) The inevitable counterexample arose.
'Inevitably, a counterexample arose.'
- (43) He spooned a moody forkful. (P.G. Wodehouse; Hall 1973)
'Moodily, he spooned a forkful.'
- (44) An unlikely chiropractor discovered the solution.
'A chiropractor discovered the solution and it was unlikely that that chiropractor (or a chiropractor?) would do so.'
- (45) Clyde asked a random linguist.
'Clyde asked a linguist randomly.'
- (46) Floyd received an unfortunate grade.
'Floyd received a grade such that it was unfortunate to receive it.'

One shouldn't read too much into these without careful examination, of course, but they collectively suggest that more external readings lurk just over our analytical horizon.

3 Three classes of nonlocal readings

This paper is not a linguistic curio cabinet. We've established, I hope, that there are patterns in this domain. That's not to say that there aren't genuine mysteries here. It's just that the phenomena at issue are mysterious *in parallel ways*. The next stage is to systematize the patterns more robustly so we can move toward an analysis.

There are, I will argue, three distinct classes of nonlocal adjectives. The first class I will set aside here. It includes the aforementioned 'modal superlatives' like *possible*. They differ from the others most strikingly in which determiners are involved in the external readings. In these cases, universal quantifiers license the external reading, not inhibit it:

$$(47) \text{ We interviewed } \left. \begin{array}{l} \text{every} \\ \# \text{the} \\ \# \text{a} \\ \# \text{no} \\ \# \text{three} \end{array} \right\} \text{ possible candidate.}$$

Superlatives and *only* also license it:

$$(48) \text{ We interviewed } \left\{ \begin{array}{l} \text{the only} \\ \text{the best} \end{array} \right\} \text{ possible candidate.}$$

Analyses for these cases can be framed around ellipsis, along the lines first proposed in Larson (2000), with structures like (49):

$$(49) \text{ We interviewed the best candidate possible } \text{for us to interview.}$$

There is a satisfying account built from standard assumptions about superlatives in Romero (2013).

It will be the other two classes that will be of interest here. These are what I'll call the weak-quantifier class, which includes *whole* and *unknown* and which permits external readings with weak quantifiers, and what I'll call the no-quantifier class, which includes *occasional* and *average* and permits external readings only with non-quantificational determiners. Of course, describing various particular determiners as 'non-quantificational' is already a bit tendentious—though for the moment, I mean this only descriptively, in the sense of Heim (1982), Kamp (1981), and DRT more generally—so more needs to be said for explicitness.

It goes beyond the scope of this paper to advocate a particular theory of how determiner quantification works in general. All we require is some general conceptual machinery to characterize particular classes. I'll refer to *every* and *most* DPs as strong and inherently quantificational; definite descriptions and other DPs that arguably directly refer as strong but not inherently quantificational; and all others as weak.

Setting the ellipsis class aside, all nonlocal readings observe a generalization:

- (50) **STRONG QUANTIFIER RESISTANCE GENERALIZATION**
 Strong, inherently quantificational determiners (*every, most*) are incompatible with nonlocal readings.

This has been observed for specific lexical-semantic families of adjectives, but the important point is that it seems to be true of all of them.

As we've seen, a few nonlocal adjectives—*occasional, average, and wrong*—are even more constrained in that they are incompatible with any determiner apart from (some combination of) *the, a*, bare plurals, and generic *your*. Stating it more officially:

- (51) **BROADER QUANTIFIER RESISTANCE GENERALIZATION**
 Some adjectives with nonlocal readings idiosyncratically resist all inherently quantificational determiners.

These generalizations are the crucial element in the taxonomy, so it may help to summarize things in a table:

	strong quant. (<i>every, most</i>)	weak (<i>three, many</i>)
<i>occasional</i>	X	X
<i>average</i>	X	X
<i>wrong</i>	X	X
<i>same</i>	X	X
<i>whole</i>	X	✓
<i>unknown</i>	X	✓
<i>inevitable</i>	X	✓
<i>unlikely</i>	X	✓
<i>different</i>	X	✓
<i>possible</i>	✓	X


Of course, the challenge now is to explain these generalizations. That's a tall order, inasmuch as it requires a synthesis of a vast array of adjectives and (collectively) a vast literature and set of analytical approaches. This won't happen in any single paper. Nevertheless, having framed the challenge in this way, we are in a better position to assess what an explanation might look like.

4 Some background

4.1 Incorporation

First, we must dispense with a straw man. One might imagine that external readings of adjectives are brought about simply by moving the adjective from its base position to an adverbial position, where it is interpreted as an adverb. The idea is a natural one, and I'll argue that in a certain sense it's not entirely wrong—but formulated in this crude way, it's unenlightening. Why should this movement happen? Why would an adjective have an adverb meaning? How does this help us understand the interaction of the adjective with the determiner?

More enlightening alternatives are available. There are many analyses on the market of individual instances of the larger problem of nonlocal readings, but they aren't straightforwardly generalizable to the full range of facts. There is one idea, though, that constitutes an excellent starting point. It's Larson (1999)'s proposal (further developed in Zimmermann 2000, 2003) that, in the *occasional* construction, the adjective moves from its base position to incorporate into the determiner in a process of 'complex quantifier formation':⁵



This movement creates a single quantificational determiner, *an+occasional*. It is then possible to provide this determiner with a denotation, listed in the lexicon just like that of any other. The advantage of that is that it's straightforward to capture various idiosyncrasies. If we need to stipulate that for *occasional* and *average*, the denotations of *the*, *a*, and *your* should be identical but for *wrong* they shouldn't be, we can reflect it directly. Indeed, we should *expect* such idiosyncrasies, inasmuch as the lexicon is, after all, a repository of the idiosyncratic.

What's less comfortable is that we have to stipulate not just that *an+occasional*, *the+occasional*, and *your+occasional* all have identical denotations, but also to make precisely the same stipulation independently for *a+sporadic*, *the+sporadic*, and *the+sporadic*—and indeed for other combinations of *a*, *the*, and *your* with adjectives of this class (though see Zimmermann 2003 for some inroads on this).⁶


Some analysis is necessary of why these readings fail to occur with determiners other than *a*, *the*, and *your*. On this approach, it would simply be to fail to stipulate any complex determiners that fail to have these as components. It would be essentially an accidental lexical gap, a mere accident of the development of language.

This approach provides helps in one way right off the bat. Quantificational determiners have access to the VP by perfectly ordinary means: Quantifier Raising

⁵I use 'incorporation' here following Larson and Zimmermann, in the generalized sense derived from Baker (1985) that is standard in the generative syntactic literature.

⁶This isn't uniformly a flaw. Certain combinations of frequency adjectives and determiners do seem to lack external readings for mysterious reasons. *The odd sailor strolled by* gets an external reading, but it's far more difficult to get it for *?An odd sailor strolled by*, as Gehrke & McNally (2015) observe. I'm not entirely sure what to make of these facts, but they don't strike me as sufficient reason to give up on the cause of trying to derive these generalizations from something deeper. In this specific case, the independently pragmatic naturalness of the internal reading may be relevant.

(May 1977, 1985, Heim & Kratzer 1998). A generalized quantifier—the type of expression a quantified nominal denotes—takes a VP as its argument. The basic architecture of a quantified sentence is as in (54):



(55) $\llbracket \text{every dog} \rrbracket = \lambda Q_{\langle e, t \rangle} . \forall x[\mathbf{dog}(x) \rightarrow Q(x)]$

The determiner *every* here has ‘access’ to the VP in the sense that its denotation asks for a predicate, Q in (55), that it can subsequently manipulate. The manipulation of VP meanings is the signature property of adverbials, of course, so on the incorporation view, what makes it seem like *occasional* has an ‘adverbial’ external reading is that it incorporates into a quantificational determiner and therefore has access to a VP meaning. Its access to clausal material external to the DP is a side-effect of the access the VP it has by actually being, in a deeper sense, a determiner.

If an adjective is part of a quantificational determiner meaning, it will gain access to the VP as a matter of course.

Thus this approach accounts for the adverbial scope of *occasional* and its kin, for the idiosyncratic interpretations of determiners in this construction, and (by stipulation) for restrictions on the determiner. It also accounts for the restriction on coordination: any adjective in a coordinate structure would be unable to move out of it without violating the Coordinate Structure Island. In general, movement from outside of one on conjunct in a coordinate structure is not possible:

- (56) a. Floyd ate rice and beans.
 b. *Beans₁, Floyd ate rice and t₁.

That’s precisely the sort of movement that, on this view, would be required in (57) to achieve the impossible external reading:

- (57) a. The occasional and angry sailor strolled by.
 b. #[The+occasional₁] [t₁ and angry] sailor strolled by.

The obligatory high position of the adjective is explained as well—any adjectives above it would block its path to the determiner.


The incompatibility of external readings with degree modification would also be expected, because only a bare adjective, and not a phrasal constituent, can do head-to-head movement, the kind required here. *Occasional* on its own is the head of an AP, but *very occasional* is not. This approach may even shed light on Zimmermann (2003)’s observation that external readings are often absent where Quantifier Raising is blocked. This analysis can be extended to *average*, *wrong*, perhaps *same*, and maybe others.

Nevertheless, one might have some qualms. The movement required would seem to violate the Head Movement Constraint (Travis 1984), which would normally prevent a head from moving outside of an adjoined phrase (the AP, in this case) as in (53).

More worrying, perhaps: why are *a*, *the*, and *your* alone the determiners that have been targeted for complex quantifier formation? Could it in principle have been any other combination? And why is it that the denotations of these complex determiner-adjective combinations aren't unpredictable? If they're specified in the lexicon, one might imagine virtually arbitrary variation, but the generalizations we would like to explain aren't arbitrary. Whatever the answers to these questions, more would have to be said to make weak-determiner-compatible adjectives such as *whole*, *unspecified*, and *different* fit in.

4.2 Structure vs ontology: The first step

Framing the current project as a trade-off between structure and ontology, at least with respect to *average*, is as I've said not novel. What I propose here is a variation on a theme from Kennedy & Stanley (2009). They observe the connection between *average* and *occasional*, and that this connection affords an analytical opportunity. For them, *average* incorporates into the determiner, just like *occasional* does for Larson (2000). The actual combinatorics required to achieve the necessary readings are complicated in ways that can be set aside, but they require a non-standard scope-taking mechanism that Barker (2007) dubbed Parasitic Scope, though appeals to it without the brand name can be found in Sauerland (1998) and earlier. The structure they propose is this:



The variable n here ranges over real numbers, or what number terms like 2.3 denote. The denotation is built up using the complex determiner *th'average* as in (59):

$$(59) \quad \llbracket th'average \rrbracket (\llbracket \lambda n \lambda x . has \ n \ children \rrbracket) (\llbracket American \rrbracket) (\llbracket 2.3 \rrbracket)$$

The denotation of *th'average* applies to three arguments. The first is a relation between numbers and individuals that have that number of children. The second is a property indicating what population is being averaged over, in this case, Americans. The final one is a real number indicating the computed average.

The details of implementation won't be crucial here, but they involve computing a mean on the basis of the maximal number of children each individual

has⁷, and $|P|$ should be interpreted as the number of individuals that have the property P :

$$(60) \quad \llbracket \textit{th'average} \rrbracket = \lambda P_{(e,t)} \lambda f_{(e,nt)} \lambda n . \frac{\sum_{P(x)} \mathbf{max}\{n|f(x)(n)\}}{|P|}$$

The most important point, for current purposes, is that on this view *average* DPs don't refer to anything metaphysically exotic because they don't refer to anything at all. Rather, they have an exotically high semantic type, which, coupled with incorporation from an adjective into a determiner and an unusual scope-taking operation, add up to a semantics that yields the right reading. For them, the right reading is strictly adverbial. It's the reading that can be paraphrased 'on average, Americans have 2.3 children'. It's worth noting, though, that this analysis has many of the same costs as the basic incorporation analysis, including having to stipulate the equivalence of *the*, *a*, and *your* on this reading.

4.3 *The kind analysis of 'occasional'*

The balance between the compositional semantics and the ontology is tilted in precisely the other direction in Gehrke & McNally (2010, 2015), building on Schäfer (2007). The distinctive property of *occasional* nominals, for them, is not in their grammar but rather in their referential properties—and it is therefore there that we should locate an analysis. So they seek a simpler syntax-semantics and a richer ontology.

It would require navigating quite a bit off my intended course to do justice to their proposal, but at its heart is an idea on which I will build: kind reference. The observation is that *the occasional sailor* involves reference to realizations of sailor-kinds. Very approximately, the truth conditions of the now-familiar sailor sentence can be rendered as in (61):

- (61) The occasional sailor strolled by.
 Approximately: 'Suitably-distributed realizations of strolling-by event kinds involved realizations of the sailor kind.'

The major advantage to this strategy is that it doesn't require the compositional backflips that the incorporation analysis—and especially the Kennedy & Stanley (2009) variant for *occasional*—requires. Indeed, because there is no movement at all, it doesn't violate the Head Movement Constraint. It also provides insight into why *a*, *the*, and *your* should be the determiners that uniquely have a special status in this construction. This is precisely the class of determiners that have a special status with respect to kinds and genericity:

- (62) a. The domestic dog is a good friend.
 b. A dog is a better friend than a cat.
 c. Your purple-breasted snicklearbler is a magnificent bird. (*dialectal*)

⁷The maximality operator is required because anyone with three children also has two.

To the extent that this approach is successful, it requires no special stipulations about the denotations of determiners. And because of that, it helps explain why determiner interpretations don't vary freely. No special stipulations are necessary to explain why *your+occasional* or the unattested **every+occasional* don't just happen by chance to mean something they don't actually mean.

The main shortcoming of this approach, from the current perspective, is that it's not clear how to make it scale up. On its own, it seems convincing that kind-reference is going to be a crucial ingredient in the analysis of external readings. But it's not clear to me how to make it the principal ingredient in a fully general theory.

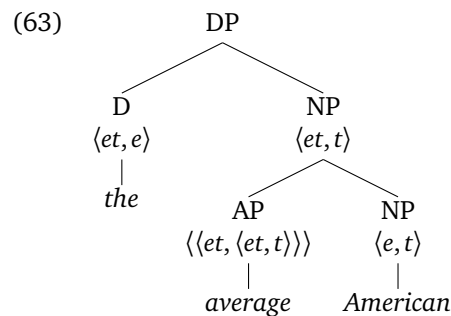
5 The modular strategy

5.1 Determiner-like adjectives


The aim of this paper is not to present a general theory of nonlocal readings, but taking a confident step in that direction requires a theory of how they arise that is modular: that is, one that relies on multiple interacting parts to arrive at an explanation. Such a theory makes it possible to activate or deactivate certain of these components to explain variation among subclasses of adjectives and—most directly at issue here—to explain the biggest split among nonlocal readings, the one between adjectives that give rise to Broader Quantifier Resistance and those that don't. (This sets aside, of course, the *possible* ellipsis class.)

One satisfying aspect of the incorporation analysis sketched above is that it reflects that nonlocal adjectives aren't prototypically adjective-like, even on a purely descriptive level. They don't pass standard diagnostics for adjectives, such as the ability to occur in comparatives, with degree modifiers, or in the complement position of *seem*. They don't conjoin with adjectives. Nor do they occur in the same positions as adjectives generally; rather, they are obligatorily high.

This might suggest incorporation or another form of syntactic differentiation, but all these properties also follow from simply assuming that nonlocal adjectives have an unusual semantic type. In the spirit of the incorporation approach, I'll assume these adjectives have precisely the same type of denotation as quantificational determiners, namely type $\langle et, \langle et, t \rangle \rangle$. Switching back to *average American*, the picture would be as in (63):



This has as a consequence that the node above the adjective, the NP *average American*, would denote a generalized quantifier. Following standard assumptions (see Heim & Kratzer 1998 for a review), it would therefore have to quantifier-raise and adjoin to the clause to avoid a type clash. I'll leave aside what happens higher in the clause for the moment to focus on the DP. The trace this movement leaves behind would standardly denote an individual. To make these LFs slightly easier to read later in the paper, I'll write it as a variable rather than a trace:



But this is hardly any help at all. It just gives rise to a different type clash: the NP would now denote an individual, but *the* is of type $\langle et, e \rangle$ and expects a property.


There is a natural solution. It's to adopt the standard BE type shift (Partee 1987), which shifts an individual to the property of being that individual:

(65) $\llbracket \text{BE} \rrbracket = \lambda x \lambda y [x = y]$

Applied to Floyd, for example, this shift would yield the property of being Floyd:

(66) $\llbracket \text{BE} \rrbracket (\mathbf{Floyd}) = \lambda y [\mathbf{Floyd} = y]$

Partee used it for copular constructions, and it has subsequently proven useful elsewhere. In this case, this resolves the type clash by providing *the* with the property-denoting argument it seeks in (64):



(68) $\llbracket \text{BE} \rrbracket (\llbracket x_1 \rrbracket) = \lambda y [x_1 = y]$

But as it turns out, at the next node up, this shift will achieve for us something more.

5.2 Determiners that work

One of the things we would like to explain is why *the*, *a*, and *your* seem to work robustly with a number of nonlocal adjectives, and why distinctions in their

interpretations seem to be neutralized in the presence of frequency adjectives and *average/typical*. That result follows from the type shift alone. There is one and only one individual that has the property of being Floyd, and it is Floyd. For this reason, *the person who is Floyd* and *Floyd* mean the same thing. So too, here *the* would combine with the property the shifted trace denotes to yield the unique individual that is identical to the one the unshifted trace denotes:

- (69) a. $\llbracket \textit{the} \rrbracket = \lambda P_{\langle e,t \rangle} . \iota y [P(y)]$
 b. $\llbracket \textit{the} \rrbracket (\llbracket \textit{BE } x_1 \rrbracket) = \iota y [x_1 = y] = x_1$

This is precisely the same individual as the one denoted by the trace alone. The effect is as though *the* were absent entirely, as though the nonlocal adjective and its NP sister had occurred in subject position on their own.

The semantically-bleached variant of *your* that occurs in e.g. *your average American* mostly amounts to a variant of *the* with a slight whiff of genericity about it, which would leave us in more or less the same place (see Gehrke & McNally 2010, 2015 for more).

As for *a*, the right result follows from a simple equivalence. To say that there's a person *x* such that *x* is wearing a hat and *x* is Floyd is just to say that Floyd is wear a hat. The same equivalence manifests itself in (70). The standard denotation of the indefinite article in (70a) when combined with the shifted trace denotation, as in (70b), yields an expression that asks for a predicate *Q* and says that some individual identical to x_1 satisfies *Q*:

- (70) a. $\llbracket a \rrbracket = \lambda P_{\langle e,t \rangle} \lambda Q_{\langle e,t \rangle} . \exists x [P(x) \wedge Q(x)]$
 b. $\llbracket a \rrbracket (\llbracket \textit{BE } x_1 \rrbracket) = \lambda Q_{\langle e,t \rangle} . \exists x [x_1 = x \wedge Q(x)]$

To say that there is an individual identical to x_1 of which the predicate *Q* holds is simply to say that *Q* holds of x_1 :

- (71) $\exists x [x_1 = x \wedge Q(x)] \Leftrightarrow Q(x_1)$

The result, again, is truth-conditionally identical to what would have happened had the determiner been absent entirely.

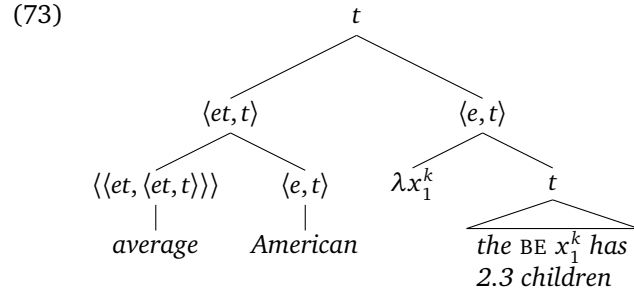
To articulate this a little bit further, let's adopt the toy denotation for *average* in (72a). This applies to the denotation of the modified NP, and predicates the VP meaning of the kind that corresponds to the NP meaning, using Chierchia (1998)'s \cap property-to-kind type shift:⁸

- (72) a. $\llbracket \textit{average} \rrbracket = \lambda P_{\langle e,t \rangle} \lambda Q_{\langle e,t \rangle} . Q(\cap P)$
 b. $\llbracket \textit{average American} \rrbracket = \lambda Q_{\langle e,t \rangle} . Q(\cap \textbf{American})$

This probably isn't adequate on its own as a theory of *average*, and much of Kennedy & Stanley (2009) may have to be layered on top of it. A few more

⁸Given this denotation, I could have equivalently dispensed with the λQ in the denotation of *average* and had *average American* denote a kind directly. This is possible here only because I am using a considerably simplified denotation, though.

words on this follow in 6.1 below. But it suffices to sketch the compositional machinery. Thus the updated tree would look like this (I've ornamented the tree with a superscript k to reflect that the trace of *average American* denotes a kind):



The result of the computation is just what we need:

- (74)
- $\llbracket \text{the BE } x_1^k \rrbracket = x_1^k$
 - $\llbracket \text{the BE } x_1^k \text{ has 2.3 children} \rrbracket = \text{has-2.3-children}(x_1^k)$
 - $\llbracket \text{average American} \rrbracket = \lambda Q_{\langle e, t \rangle} . Q(\cap \text{American})$
 - $\llbracket \text{average American} \rrbracket (\llbracket \lambda x_1^k \text{ the } [BE x_1^k] \text{ has 2.3 children} \rrbracket)$
 $= \text{has-2.3-children}(\cap \text{American})$

So the upshot is a semantics that requires that Americans generally have 2.3 children.

The crucial component to notice here is not the semantics of *average*, though, so much as the way the combination of the type shift, compositional assumptions, and kind-reference have achieved the effect of ensuring that precisely the determiners that systematically license external readings across yield the right result.

5.3 Determiners that don't work

What of determiners that *don't* work? Again, the nature of the movement and resulting type shift helps the situation—or rather, undermines it in the right way.

Strong determiners like *every* and *most* presuppose that their domain has more than one member. If there is only one person in the corner, for example, (75) gives rise to failure of presupposition:

- (75) Every person in the corner left.

I've spelled it out explicitly in the denotation of *every* in (76) (the colon indicates the presupposition; $|P|$, as before, indicates the cardinality of individuals that satisfy P)

- (76) $\llbracket \text{every} \rrbracket = \lambda P_{\langle e, t \rangle} : |P| > 1 . \lambda Q_{\langle e, t \rangle} . \forall x [P(x) \rightarrow Q(x)]$

In (77), *every* combines with the property $\llbracket \text{BE } x_1^k \rrbracket$:

- (77) a. #Every average American has 2.3 children.
 b. [average American] λx_1^k [every [BE x_1^k]]] has 2.3 children

(78) $\llbracket \text{BE } x_1^k \rrbracket = \lambda y [x_1^k = y]$

But (78) is a singleton property—there is only one individual that is identical to x_1^k . It therefore violates the presupposition *every* imposes on its first argument.

This presupposition is not a peculiarity of *every*, but rather a property of strong quantificational determiners in general. Thus *most* would work similarly. Because movement below the DP level, in the framework proposed here, systematically gives rise to such singleton properties, it systematically precludes combining with strong quantifiers.

We have thus derived one of the two generalizations articulated earlier: the Strong Quantifier Resistance Generalization. All external readings observed it, so if this mechanism is crucial to deriving external readings, this explains it. Weak quantificational determiners do not have this presupposition, so they don't in general block external readings.

But what of the Broader Quantifier Resistance Generalization, the one only some adjectives observed? Some adjectives—like our test cases, *average* and *occasional*—do block the external reading with weak quantifiers too. But despite the absence of the fatal presupposition, these fail in another respect. The denotation of *three* is as in (79), a property of individuals that have a cardinality of 3:

(79) $\llbracket \text{three} \rrbracket = \lambda x [|x| = 3]$

When this combines with the shifted trace, it will combine intersectively with its denotation to yield (80):

(80) $\llbracket \text{three BE } x_1^k \rrbracket = \lambda y [x_1^k = y \wedge |y| = 3]$

This is a property satisfied by a plurality with three elements that is identical to the kind x_1^k . That means, naturally, that the kind x_1^k has to be a plurality of three elements. But kinds aren't pluralities, and they don't have cardinalities. This is pretty straightforward metaphysically, but again, linguistic evidence makes it clear. As Chierchia (1998) demonstrates especially convincingly, across languages kinds are essentially a kind of mass term. *Cheese*, for example, denotes a kind in English, and **three cheese* is ungrammatical.

So in this case, the problem that rules out weak quantifiers has to do with kinds, and it will be only nonlocal adjectives that leave behind kind-denoting traces that will be subject to this additional restriction. *Occasional* is also incompatible with weak quantifiers, and, as Gehrke & McNally (2010, 2015) demonstrate, its semantics also relies crucially on kinds. Nonlocal adjective with no kind overtones such as *whole* or *wrong* or *unspecified* should therefore avoid running afoul of this difficulty and be compatible with weak quantifiers even on their external readings. And indeed they are. More on both of these points follows in the subsequent two sections.

5.4 A word about ‘occasional’

Occasional and its kin aren’t the focus here, but a brief word about how they might work in this framework is appropriate. The approach to which I’m most sympathetic would be to simply combine the insights of two competing classes of approaches. Kinds must occupy a central place, for the reasons discussed above. But quantification can play a central role too. In particular, there is no reason not to adopt the Zimmermann (2000)’s quantifier OCCASIONAL, which quantifies jointly over the individuals and events, though here it will be crucial that it be kinds and events (with *s* as the type of events):

$$(81) \quad \llbracket \text{occasional} \rrbracket = \lambda P_{(e,t)} \lambda Q_{(e,st)} . \text{OCCASIONAL } x^k, e : \cap P(x)[Q(x^k)(e)]$$

This denotation would trigger movement to a position just below where the event argument is closed, and yield sentence denotations like (82):

$$(82) \quad \text{a. The occasional sailor strolled by.}$$

$$\text{b. } \llbracket \text{occasional sailor } \lambda x_1^k \text{ the BE } x_1^k \text{ strolled by} \rrbracket$$

$$= \text{OCCASIONAL } x^k, e : \cap \text{sailor}(x)[\text{strolled-by}(x^k)(e)]$$

This seems a reasonable happy medium between the two approaches.

5.5 The weak quantifier class

There remains to discuss the class of external readings that *are* compatible with weak quantifiers. For those, though, in one sense there is little to be said. What ensured incompatibility with weak quantifiers above was the role of kinds. Adjectives whose semantics makes no special reference to kinds don’t give rise to the problem of computing the cardinality of a kind.

To illustrate, the denotation of *unknown* could be characterized as in (83), where I’ve used $?x\phi$ to abbreviate the embedded question ‘which *x* is such that ϕ ?’.⁹

$$(83) \quad \text{a. } \llbracket \text{unknown} \rrbracket = \lambda P_{(e,t)} \lambda Q_{(e,t)} . \exists x[P(x) \wedge Q(x) \wedge \neg \text{known}(?y[Q(y)])]$$

$$\text{b. } \llbracket \text{unknown hotel} \rrbracket$$


$$= \lambda Q_{(e,t)} . \exists x[\text{hotel}(x) \wedge Q(x) \wedge \neg \text{known}(?y[Q(y)])]$$

What *unknown hotel* does is a little complicated. First, it requires that there exist a hotel that satisfies the property formed by raising the whole quantified NP. Second, it requires that it not be known which individuals satisfy this property.

It will help to see how this works in action. The tree for (84a) arrived at by raising would be as in (84b):

⁹One may freely substitute one’s favorite theory of indirect questions here, so far as I can see, though what I have in mind is that $\ulcorner ?y[Q(y)] \urcorner$ should be taken to be the set of propositions formed by varying the value of *y*, i.e., an abbreviation for $\ulcorner \{p : \exists x[p = Q(y)] \urcorner$.

- (84) a. Solange stayed at three unknown hotels.
b.



This assumes a null existential determiner in the head of the nominal, and that, standardly, it undergoes quantifier raising. The denotation of (84) would be as in (85):

- (85) a. $\llbracket \text{BE } x_1 \rrbracket = \lambda y[x_1 = y]$
 b. $\llbracket \text{three BE } x_1 \rrbracket = \lambda y[x_1 = y \wedge |x_1| = 3]$
 c. $\llbracket \exists \text{ three BE } x_1 \rrbracket = \lambda g_{\langle e, t \rangle} . \exists y[x_1 = y \wedge |x_1| = 3 \wedge g(y)]$
 d. $\llbracket \lambda x_1 [\exists \text{ three } x_1] \lambda x_2 \text{ Solange stayed at } x_2 \rrbracket$
 $= \lambda x_1 . \exists y[x_1 = y \wedge |x_1| = 3 \wedge \text{stay-at}(y)(\text{Solange})]$
 $= \lambda x_1 . [|y| = 3 \wedge \text{stay-at}(y)(\text{Solange})]$

This is a property that holds of any three-membered plural individual such that Solange stayed at its members.¹⁰

What *unknown hotels* adds to this is that this plurality is required to consist of hotels, and that it not be known which hotels precisely these are. The computation for the full sentence is in (86):

- (86) $\llbracket \text{unknown hotels } \lambda x_1 [\exists \text{ three } x_1] \lambda x_2 \text{ Solange stayed at } x_2 \rrbracket$
 $= \exists x \left[\text{hotel}(x) \wedge |x| = 3 \wedge \text{stay-at}(x)(\text{Solange}) \wedge \right.$
 $\left. \neg \text{known}(?y[\text{stay-at}(y)(\text{Solange})]) \right]$

The result, correctly, is that there must be three hotels at which Solange stayed, and it must not be known which hotels these are.

The crucial element in all this, though, is that there is nothing about *unknown* that prevents cardinalities from being computed, and so nothing that resists, in this instance, *three*, and more broadly any of its kin.

5.6 Summary

The result, then, is that there is no need for incorporation. The external scope facts follow from quantifier raising. The interpretation of determiners is standard. Restrictions on determiners follow from independent considerations. The general resistance of nonlocal adjectives to strong quantifiers follows from the

¹⁰I set aside questions of distributivity and collectivity here.

compositional circumstances of their movement, which invoke a type shift with which they are incompatible. The resistance of certain nonlocal adjectives to weak quantifiers follows from independent facts about the lexical semantics of the adjective—specifically, having a kind-based semantics. Other restrictions, like the lack of coordination with ordinary adjectives and absence of degree modifiers, follow from the quantifier type of these expressions.

This means it was not necessary to stipulate which determiners support incorporation and which don't, or what interpretations result for every combination. Nor was it necessary to stipulate why *the*, *a*, and *your* wind up making the same semantic contribution, or to do so repeatedly for each frequency adjective. It also wasn't necessary to stipulate anything about the interaction of quantificational force with external readings. This is possible in part precisely because what I have offered here is only a sketch. The devil, as always, is in the details. But I hope this illustrates an analytical approach to these facts that might hope to scale up to the broader analytical picture I sought to draw.

6 Taking stock

6.1 *Could things be so simple?*

One issue remains strikingly unresolved. I've characterized the denotation for *average* I've provided above as a toy denotation. I've said, perhaps a bit defensively, that things couldn't possibly be so simple. Surely, it couldn't suffice to say that *the average American* means, essentially, the same thing as the kind-denoting nominal *Americans*, and (87a) and (87b) mean more or less the same thing:

- (87) a. The average American has 2.3 children.
b. Americans have 2.3 children.

But the truth is, I think the simple toy version of the facts may be onto a deeper grammatical intuition than the more complicated one.

To be sure, we have the option of layering on components of the Kennedy & Stanley approach here, introducing elements of their machinery on top of the bits I propose to achieve their desired adverbial reading. There is a danger of redundancy, though. And the more one does that, the farther one gets to the connection to kind-reference, for which Gehrke & McNally provide ample evidence.

The defense of the naïve theory proceeds in several steps. The first is empirical. Suppose we adopt a theory that involves computing a mean. On such a view, (88a) and (88b) would both be predicted to be true, and, therefore, quite probably (88c) too:

- (88) a. The average human has one ovary.
b. The average human has one testicle.
c. The average human has one ovary and one testicle.

Yet they are all false, or in any case false outside of exceptionally odd statistical contexts. Any theory that revolves primarily around calculation of means would

fail to predict this. But in a theory that relies on kind reference, it's expected. On such a view, it's the *2.3 children* case that's puzzling.

That, I think, is precisely where we should *want* to be puzzled—that is the case that we should treat as exotic rather than as the core example. Most languages through most of human history had no reason to refer to fractions. Moreover, the semantics of fractions is independently puzzling. They are problematic completely independent of their role in *average* sentences. It makes sense, then, that the theory of *average* shouldn't be itself founded on this independent mystery.

That said, nothing in the general conception of external readings proposed here does not rest above all on any particular assumptions about kinds. Perhaps other notions could do the necessary work without putting us on thin ice with respect to sentences containing fractions. Indeed, I consider one such possibility in section 6.2 below. The only crucial role kinds play here is to rule out computing cardinalities, which in turn is crucial to distinguishing the weak-quantifier class from the no-quantifier class. That's not nothing, but there may be other means of accomplishing this. Nevertheless, it's worth recognizing that there are several converging lines of evidence that point to kinds or some form of genericity in these sentences: initial intuitions about what *average* sentences mean, the judgments in (88), the role kinds play in distinguishing classes of external readings, and its place in correctly predicting which determiners have which readings. One might be still able to explain *2.3 children* by simply adopting, Kennedy & Stanley, an extraordinarily high type, but it seems right that special stipulations should be required there and not elsewhere.

It's worth pointing out, though, that one could also follow in the spirit of Carlson & Pelletier (2002) and appeal to fictive entities in place of some form of kind. This analytical avenue may actually be more available on this approach. Kennedy & Stanley argue against the fictive entity approach in part on the grounds that it doesn't explain the limited inventory of determiners possible with *average*. Those facts, however, can be explained independently here. But again, if the relevant notion of fictive entities can emerge with an appropriate kind flavor, that seems preferable on independent grounds to the alternative.

None of that directly addresses what the semantics of *2.3 children* should be. My suspicion is that an ultimately satisfying answer requires not just an theory of nonlocal readings of adjectives, but a better theory of mathematical language, and in particular of what I elsewhere call 'semantic viruses' (Morzycki 2017), in the spirit of Sobin (1997) syntactic viruses (see also Lasnik & Sobin 2000, Schütze 1999). I argue there that some expressions associated with educated, often highly self-conscious language may use special semantic mechanisms not otherwise available in the semantics. Making this distinction may help us distinguish which operations what grammatical phenomena truly are exotic and may call for some brute-force high-type complexity, and where we should seek simplicity, even occasionally in the face of apparent counterevidence.

6.2 *Kinds and concepts*

Sebastian Löbner (p.c.) suggests that a number of the restrictions on external readings of *average* and *occasional* may involve characterizing more precisely the

concept types they give rise to. *Average American* on the relevant reading isn't a sortal concept—one that supports counting and is neither uniquely referential nor relational. That accounts for its incompatibility with strong quantifiers (*#every average American*), and perhaps for its incompatibility with stacked or conjoined adjectives (*#an average (and) irritable American*).

This mode of explanation in some respects has the same shape as a kind-based one, or indeed as one organized around fictive entities. They all seek to derive the properties of the expression from the ontological status of the extension of the nominal. Both kinds and the relevant non-sortal concepts are uncountable. It doesn't seem too far-fetched to claim that fictive entities might not be countable either, though that's less obvious. Insomniacs are sometimes advised to count sheep in order to fall asleep, yet under normal circumstances the livestock in one's bedroom are entirely fictive. Likewise, the resistance to quantification that I earlier attributed to a failure of presupposition could be attributed to countability as well. As I expressed it in (76), the presupposition involved determining the cardinality individuals that satisfy the property expressed by the nominal argument. In this implementation, that is not undefined. Even though this property has in its domain kinds, it denotes a property that holds of precisely one kind. Therefore it is countable. This follows from how the movement and type shift interact. One might imagine, though, an alternative analysis where the inherent countability of the noun is crucial. In order for the analysis of adjective stacking and conjunction to go through, however, one really would have to have the NP *average American* denote this concept kind quite low in the tree, before any type shifts have taken place. On this view, then, the crucial difference could be viewed as being in how high in the structure of the nominal kinds are invoked. But there are good reasons to think properties of kinds are to be found deep in the nominal extended projection, very near the noun (Zamparelli 1995 among others). So this fact too may be insufficient to distinguish these two approaches on a deep level, setting aside particular analytical choices I've made here.

The adjective order facts, however, might be of use. Most evidence for a layer in the nominal projection that is concerned with kinds rather than objects suggests that it is the lower of the two. So-called Bolinger contrasts (Bolinger 1967; see Morzycki 2016a and Leffel 2014 for extensive discussion) such as the one in (89) to show that adjectives lower in the nominal ascribe inherent or individual-level properties, and higher ones ascribe contingent or stage-level properties:

- (89) a. the invisible visible stars
b. #the visible invisible stars

On its only possible reading, (89a) refers to stars that are visible in principle but invisible at the moment, perhaps by clouds. But (89b) is contradictory, because it refers to stars that are invisible in principle but visible at the moment. A broadly similar fact, in the spirit of Larson (1998, 2000):

- (90) an ugly beautiful dancer
a. 'an ugly person who dances beautifully'
b. *'a beautiful person who dances in an ugly way'

- (91) a beautiful ugly dancer
 a. *‘an ugly person who dances beautifully’
 b. ‘a beautiful person who dances in an ugly way’

Larson marshals such facts to argue for a generic quantifier in the nominal projection. But be it about kinds or not, the domain of genericity in the nominal is low. Yet as we’ve seen, adjectives associated with external readings are exclusively high. A reminder:

- (92) An ugly occasional sailor strolled by.
 a. ‘An ugly person who sails occasionally strolled by.’ (internal)
 b. *‘Occasionally, an ugly sailor strolled by.’ (external)
- (93) An occasional ugly sailor strolled by.
 a. *‘An ugly person who sails occasionally strolled by.’ (internal)
 b. ‘Occasionally, an ugly sailor strolled by.’ (external)

If kinds or non-sortal properties were at issue lower in the nominal, this effect would be expected to be either reversed or absent entirely.

One appeal of such an approach, in either of these incarnations, would be that the quantificational facts and the facts about conjunction and stacking could be brought under the same rubric. As it stands, the latter derive from the quantificational type of the NP. A major disadvantage is that they wouldn’t readily extend to the rather large class of adjectives compatible with weak quantifiers. Nor, in the absence of a scope-taking mechanism, would they permit the adjective access to the VP denotation. Yet this access is precisely what seems to be required for e.g. epistemic adjectives such as *unknown*, as shown in 5.5.

7 Final remarks

To close, a few words about the commonly expressed intuition that nonlocal readings are a grammatical oddity. These adjectives are indeed odd, but in a precise and interesting sense. They are odd in the way that platypuses and lungfish are odd: they are—perhaps metaphorically, or perhaps more than metaphorically—transitional forms in an evolutionary progression, unusual because they combine features of two distinct categories that we normally regard as mutually exclusive. Over succeeding generations of speakers, certain adjectives may emerge from the swampy depths of the inner NP to which they are usually confined, and tentatively make their way onto the dry land of the determiner domain. They can’t be expected to make this leap in a single stride, so we can observe them in the midst of their evolutionary journey and thereby discover more about both their evolutionary origin and their destination. Like platypuses and lungfish, they are important and analytically revealing not despite their strangeness, but because of it.

Substantively, the proposal was that nonlocal adjectives have quantificational determiner denotations, trigger raising of the NP in which they occur, stranding the determiner, and sometimes require properties of kinds as their arguments.

This isn't a general theory of all nonlocal readings, naturally. That would be far too ambitious for any single paper. But it has the shape of a general theory, and my hope is that further research will be able to fill in the gaps in a similar spirit.

From the broader cognitive perspective, though, one of the larger lessons is the balance between the explanatory burden on the ontology and on the structural machinery. For *average*, for example, one might have gone in the direction of recognizing 'average Americans' as actual, if very abstract, objects in the model, 'fictive persons'. For *occasional*, I followed Gehrke & McNally (2010, 2015) in placing a great deal of explanatory weight on the notion of kinds, if perhaps not quite so much weight as they have.

On the other hand, structural component played a crucial role. For *average*, one could go so far as Kennedy & Stanley (2009) do, and invoke quite high-powered syntactic and semantic machinery to twist the tree into the shape we require. For *occasional*, Larson (1999), Zimmermann (2000) and others provide a path that also requires quite a bit of syntactic machinery.

It is misguided, I think, to ask where we wind up in each respect: how much compositional structure do we need, how much metaphysics, and what the right balance is. Rather, we should recognize that there may be some explanatory trade-offs, but that inevitably, we will need a bit of both modes of explanation—and it is up to language to tell us how much we need of either.

References

- Abusch, Dorit & Mats Rooth. 1997. 'Epistemic NP modifiers'. In A. Lawson (ed.), *Proceedings of Semantics and Linguistic Theory (SALT) 7*. CLC Publications, Ithaca, NY.
- Alrenga, Peter. 2006. 'Scalar (non-)identity and similarity'. In Donald Baumer, David Montero, & Michael Scanlon (eds.), *Proceedings of the West Coast Conference on Formal Linguistics (WCCFL) 25*. Cascadilla Press, Somerville, MA.
- Alrenga, Peter. 2007a. 'Comparisons of similarity and difference'. In *Proceedings of the 2005 Workshop on the Formal Analysis of Adjectives (JET Adjectifs)*. University of Paris 7 Jussieu.
- Alrenga, Peter. 2007b. 'Tokens, types, and identity'. In *Proceedings of the North East Linguistic Society (NELS) 38*. GLSA Publications, Amherst.
- Bach, Emmon. 1986. 'The algebra of events'. *Linguistics and Philosophy* **9**(1), 5–16.
- Bach, Emmon. 1989. *Informal Lectures on Formal Semantics*. State University of New York Press, Stony Brook.
- Baker, Mark C. 1985. *Incorporation: A Theory of Grammatical Function Changing*. Doctoral dissertation, MIT.
- Barker, Chris. 2007. 'Parasitic scope'. *Linguistics and Philosophy* **30**(4), 407–444.
- Beck, Sigrid. 2000. 'The semantics of 'different': Comparison operator and relational adjective'. *Linguistics and Philosophy* **23**(2), 101–139.
- Bolinger, Dwight L. 1967. 'Adjectives in English: Attribution and predication'. *Lingua* **18**, 1–34.

- Brasoveanu, Adrian. 2011. 'Sentence-internal *different* as quantifier-internal anaphora'. *Linguistics and Philosophy* **34**, 93–168.
- Büring, Daniel. 2007. 'More or less'. In *Proceedings of the Chicago Linguistics Society (CLS)* 43.
- Carlson, Greg & Francis Jeffrey Pelletier. 2002. 'The average american has 2.3 children'. *Journal of Semantics* **19**(1), 73–104.
- Carlson, Greg N. 1987. 'Same and different: Some consequences for syntax and semantics'. *Linguistics and Philosophy* **10**(4), 531–565.
- Chierchia, Gennaro. 1998. 'Reference to kinds across languages'. *Natural Language Semantics* **6**(4), 339–405.
- Cinque, Guglielmo. 2010. *The syntax of adjectives: a comparative study*, *Linguistic Inquiry Monographs*, vol. 57. MIT Press, Cambridge, MA.
- DeVries, Karl. 2010. 'Frequency adjectives'. BA thesis, Michigan State University.
- Gehrke, Berit & Louise McNally. 2010. 'Frequency adjectives and assertions about event types'. In Ed Cormany, Satoshi Ito, & David Lutz (eds.), *Proceedings of Semantics and Linguistic Theory (SALT)* 19, pp. 180–197. eLanguage, Ithaca, NY.
- Gehrke, Berit & Louise McNally. 2015. 'Distributional modification: The case of frequency adjectives'. *Language* **91**(4), 837–870.
- Haïk, Isabelle. 1985. *The syntax of operators*. Doctoral dissertation, MIT.
- Hall, Jr., Robert A. 1973. 'The transferred epithet in P. G. Wodehouse'. *Linguistic Inquiry* **4**(1), 92–94.
- Heim, Irene. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Doctoral dissertation, University of Massachusetts Amherst.
- Heim, Irene. 1985. 'Notes on comparatives and related matters'. Ms., University of Texas, Austin.
- Heim, Irene. 2006. 'Little'. In M. Gibson & J. Howell (eds.), *Proceedings of Semantics and Linguistic Theory (SALT)* 16. CLC Publications, Ithaca, NY.
- Heim, Irene. 2008. 'Decomposing antonyms?' In Atle Grøn (ed.), *Proceedings of Sinn und Bedeutung* 12, pp. 212–225. ILOS, Oslo.
- Heim, Irene & Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Blackwell Publishing, Oxford.
- Hornstein, Norbert. 1984. *Logic as Grammar*. MIT Press, Cambridge, MA.
- Kamp, Hans. 1981. 'A theory of truth and semantic representation'. In Jeroen Groenendijk, Theo Janssen, & Martin Stokhof (eds.), *Formal Methods in the Study of Language, Part 1*, vol. 135, pp. 277–322. Mathematical Centre Tracts, Amsterdam. Reprinted in Jeroen Groenendijk, Theo Janssen and Martin Stokhof (eds), 1984, *Truth, Interpretation, and Information; Selected Papers from the Third Amsterdam Colloquium*, Foris, Dordrecht, pp. 1–41.
- Keenan, Edward L. 1992. 'Beyond the Frege boundary'. *Linguistics and Philosophy* **15**(2), 199–221.
- Kennedy, Christopher. 1997. *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. Doctoral dissertation, UC Santa Cruz. Published in 1999 by Garland, New York.
- Kennedy, Christopher. 2001. 'Polar opposition and the ontology of "degrees"'. *Linguistics and Philosophy* **24**(1), 33–70.
- Kennedy, Christopher & Jason Stanley. 2009. 'On "average"'. *Mind* **118**(471), 583–646.

- Larson, Richard. 1998. 'Events and modification in nominals'. In D. Strolovitch & A. Lawson (eds.), *Proceedings of Semantics and Linguistic Theory (SALT) 12*. CLC Publications, Ithaca, NY.
- Larson, Richard. 1999. 'Semantics of adjectival modification'. Lecture notes, LOT Winter School, Amsterdam.
- Larson, Richard K. 2000. 'ACD in AP?' In Roger Billerey & Brook Danielle Lillehaugen (eds.), *Proceedings of the West Coast Conference on Formal Linguistics (WCCFL) 19*. Cascadilla Press, Somerville, MA.
- Lasersohn, Peter. 2000. 'Same, models and representation'. In B. Jackson & T. Matthews (eds.), *Proceedings of Semantics and Linguistic Theory (SALT) 10*, pp. 83–97. CLC Publications, Ithaca, NY.
- Lasnik, Howard & Nicholas Sobin. 2000. 'The who/whom puzzle: On the preservation of an archaic feature'. *Natural Language and Linguistic Theory* **18**(2), 343–371.
- Leffel, Tim. 2014. *The semantics of modification: Adjectives, nouns, and order*. Doctoral dissertation, New York University.
- Löbner, Sebastian. 2000. 'Polarity in natural language: Predication, quantification and negation in particular and characterizing sentences'. *Linguistics and Philosophy* **23**(3), 213–308.
- Majewski, Helen. 2002. 'The same and different: the same or different?' Ms., University of Massachusetts Amherst.
- May, Robert. 1977. *The Grammar of Quantification*. Doctoral dissertation, MIT, Cambridge, MA.
- May, Robert. 1985. *Logical form*. MIT Press.
- Moltmann, Friederike. 1992. 'Reciprocals and same/different: Towards a semantic analysis'. *Linguistics and Philosophy* **15**(4), 411–462.
- Moltmann, Friederike. 1997. 'Intensional verbs and quantifiers'. *Natural Language Semantics* **5**(1), 1–52.
- Moltmann, Friederike. 2005. 'Part structures in situations: The semantics of Individual and Whole'. *Linguistics and Philosophy* **28**(5), 599–641.
- Morzycki, Marcin. 2002. 'Wholes and their covers'. In Brendan Jackson (ed.), *Proceedings of Semantics and Linguistic Theory (SALT) 12*. CLC Publications, Ithaca, NY.
- Morzycki, Marcin. 2016a. *Modification*. Key Topics in Semantics and Pragmatics. Cambridge University Press, Cambridge.
- Morzycki, Marcin. 2016b. 'Toward a general theory of nonlocal readings of adjectives'. In Nadine Bade, Polina Berezovskaya, & Anthea Schöller (eds.), *Proceedings of Sinn und Bedeutung*.
- Morzycki, Marcin. 2017. 'Some viruses in the semantics'. In Nicholas LaCara, Keir Moulton, & Anne-Michelle Tessier (eds.), *A Schrift to Fest Kyle Johnson*. University of Massachusetts Linguistics Open Access Publications.
- Nunberg, Geoffrey. 1984. 'Individuation in context'. In M. Cobler et al. (ed.), *Proceedings of the West Coast Conference on Formal Linguistics (WCCFL) 2*, pp. 203–217. Stanford Linguistics Association, Stanford.
- Partee, Barbara. 1987. 'Noun phrase interpretation and type-shifting principles'. In Jeroen Groenendijk, Dick de Jongh, & Martin Stokhof (eds.), *Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers*, pp. 115–143. Foris, Dordrecht. Reprinted in Partee & Portner (2002).

- Partee, Barbara H. & Paul Portner (eds.). 2002. *Formal Semantics: The Essential Readings*. Blackwell Publishing, Oxford.
- Romero, Maribel. 2013. 'Modal superlatives: A compositional analysis'. *Natural Language Semantics* 21, 79–110.
- Sauerland, Uli. 1998. *The Meaning of Chains*. Doctoral dissertation, MIT, Cambridge, MA.
- Schäfer, Roland. 2007. 'On frequency adjectives'. In E. Puig-Waldmüller (ed.), *Proceedings of Sinn und Bedeutung 11*. Universitat Pompeu Fabra, Barcelona.
- Schmitt, Cristina. 2000. 'Some consequences of the complement analysis for relative clauses, demonstratives and wrong adjectives'. In Artemis Alexiadou, Paul Law, André Meinunger, & Chris Wilder (eds.), *The syntax of relative clauses*, pp. 309–348. John Benjamins Publishing, Amsterdam.
- Schütze, Carson. 1999. 'English expletive constructions are not infected'. *Linguistic Inquiry* 30(3), 467–484.
- Schwarz, Bernhard. 2005. 'Modal superlatives'. In Effi Georgala & Jonathan Howell (eds.), *Proceedings of Semantics and Linguistic Theory (SALT) 15*. CLC Publications, Ithaca, NY.
- Schwarz, Bernhard. 2006. 'Attributive wrong'. In David Montero Donald Baumer & Michael Scanlon (eds.), *Proceedings of the West Coast Conference on Formal Linguistics (WCCFL) 25*, pp. 362–370. Cascadilla Press, Somerville, MA.
- Schwarz, Bernhard. To appear. 'The wrong number: non-local adjectival modification'. In Lisa Matthewson, Cécile Meier, Hotze Rullmann, & Thomas Ede Zimmermann (eds.), *The Blackwell Companion to Semantics*. Blackwell Publishers.
- Sobin, Nicholas. 1997. 'Agreement, default rules, and grammatical viruses'. *Linguistic Inquiry* 28(2), 318–343.
- Stump, Gregory T. 1981. 'The interpretation of frequency adjectives'. *Linguistics and Philosophy* 4(2), 221–257.
- Travis, Lisa. 1984. *Parameters and Effects of Word Order Variation*. Doctoral dissertation, MIT.
- Zamparelli, Roberto. 1995. *Layers in the Determiner Phrase*. Doctoral dissertation, University of Rochester. (As revised in 2000.)
- Zimmermann, Malte. 2000. 'Pluractional quantifiers: The Occasional construction in English and German'. In Brendan Jackson & Tanya Matthews (eds.), *Proceedings of Semantics and Linguistic Theory (SALT) 10*. CLC Publications, Ithaca, NY.
- Zimmermann, Malte. 2003. 'Pluractionality and complex quantifier formation'. *Natural Language Semantics* 11(3), 249–287.